servicenow

Brazil IT Operations Management

Last updated: 10/08/2026

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Some examples and graphics depicted herein are provided for illustration only. No real association or connection to ServiceNow products or services is intended or should be inferred.

ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Please read the ServiceNow Website Terms of Use at www.servicenow.com/terms-of-use.html

Company Headquarters
2225 Lawson Lane
Santa Clara, CA 95054
United States
(408) 501-8550

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

servicenow

Table of Contents

IT Operations Management 5

Install ITOM 8

AI in ITOM 10

Install the ServiceNow Otto for IT Operations Management (ITOM) application 13

Use the ITOM MCP Server Console 13

Applications installed with ServiceNow Otto for ITOM 27

Now Assist for ITOM 27

ITOM/OT SU Licensing and subscriptions 28

Explore 29

Configure 38

Use 39

Reference 45

ITOM Cloud Accelerate 51

Cloud Account Management 52

Cloud Services Catalog 141

Cloud Migration Assessment 206

Cloud Provisioning and Governance 206

Cloud Configuration Governance 649

Cloud Action Library 723

ITOM Cloud Accelerate reference 753

ITOM Optimization 757

ITOM Optimization reference 758

ITOM Visibility 759

Install ITOM Visibility 762

AI in ITOM Visibility 763

Discovery 768

Firewall Audits and Reporting 1486

Certificate Inventory and Management 1512

Service Mapping 1609

Discovery patterns used by ITOM Visibility 2074

Cryptographic Asset Compliance 2944

AI Agent Topology Mapping 2957

ITOM Content Service 2997

Tag Governance 3005

Cloud Discovery Workspace 3029

Cloud License Estimator 3039

ITOM Infra Services Workspace 3043

Learning Enhanced Automation Platform (LEAP) 3050

ITOM Visibility reference 3112

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM AIOps....3476
Install ITOM AIOps apps using Setup Hub....3476
Explore....3478
AI in ITOM AIOps....3501
Configure....3505
AIOps Dashboards in SOW for ITOM....3551
Event Management....3572
Health Log Analytics....4253
Integrations Launchpad in Service Operations Workspace for ITOM....4626
Service Level Objective Management....4653
Service Observability....4678
Service Reliability Management....4765
Synthetic monitoring....4843
ITOM AIOps reference....4870
Agent Client Collector....4878
Exploring Agent Client Collector....4879
Configuring Agent Client Collector....4898
Agent Client Collector Monitoring....5118
Agent Client Collector Log Analytics....5123
View live CI data with ACC....5139
AI Control Tower....5141
Agent Client Collector reference....5145
Metric Intelligence....5497
Exploring Metric Intelligence....5498
Configuring Metric Intelligence....5520
Optimizing Metric Intelligence....5526
Metric Intelligence reference....5547
Access to cloud environments for ITOM products....5548
Install and configure MID Servers to access cloud environments....5548
Set up a cloud service account....5553

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IT Operations Management

Get better visibility into your infrastructure and services, prevent service outages, and expand your organization's operational agility with ServiceNow® IT Operations Management.

See more and do more with IT Operations Management

View and download the full infocard for a highlight of IT Operations Management features.

	Get visibility into your infrastructure and services Take charge of your infrastructure and services from one source. Improve service quality, strengthen change processes, reduce risk, optimize your infrastructure spend, and minimize software compliance issues.
	Avoid and solve service issues effectively Free up your IT staff from time-consuming, error-prone tasks and drive down service outages with AIOps. Accelerate issue resolution with relevant and contextual change, problem, and incident data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get visibility into your infrastructure and services

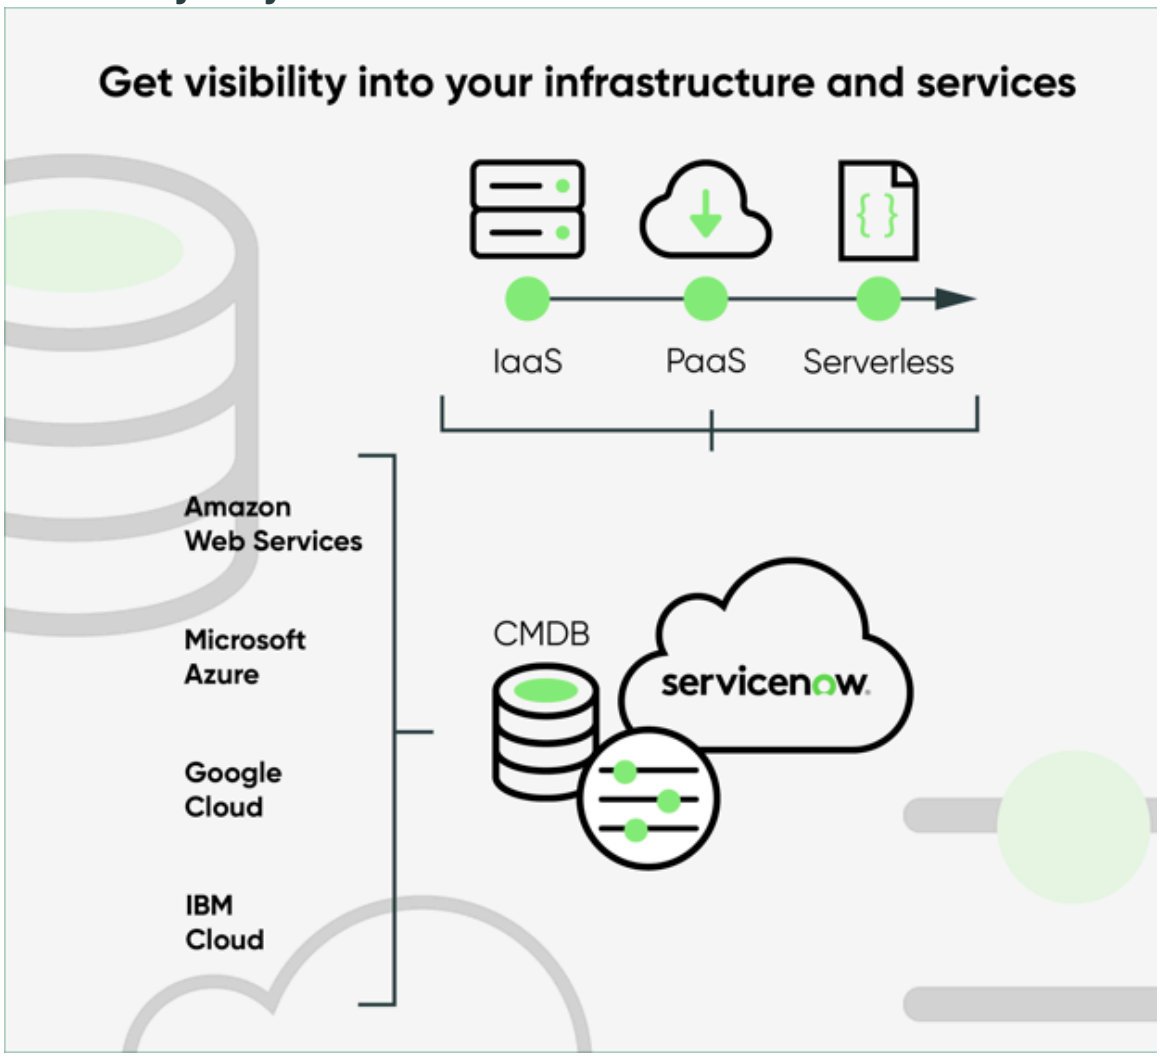

Instantly see the service impact of your infrastructure issues and changes, simplify root-cause analysis, and reduce mean time to repair (MTTR). End-to-end discovery and service mapping gives you an accurate, up-to-date CMDB view of your IT infrastructure and services. The resulting complete and reliable record is used by other ServiceNow applications, such as ITOM AIOps, ITOM Optimization, and Software Asset Management. In this way, you can improve service quality, strengthen change processes, reduce risk, optimize infrastructure spend, and minimize software compliance issues. For more information, see ITOM Visibility.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Avoid and solve service issues effectively

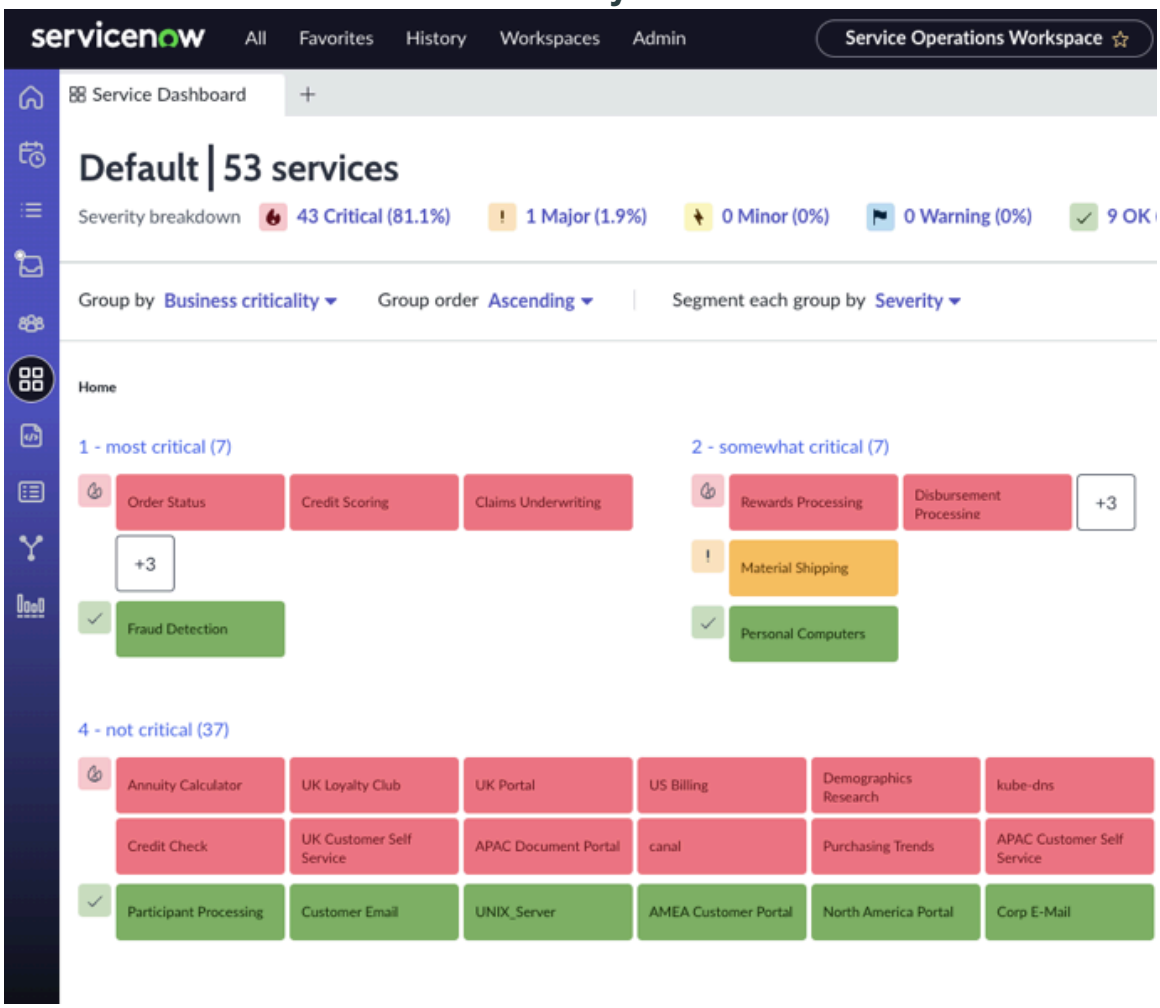

ITOM AIOps uses AI to turn events into actionable alerts. Pinpoint service issues and rapidly identify and remediate the root cause. Unlike legacy event management systems, which are static and rule based, ITOM AIOps applies machine learning and advanced analytics to correlate events, adapting automatically to rapidly evolving virtualized and cloud environments. For more information, see ITOM Health.

Accelerate cloud strategy and reduce costs

Accelerate your cloud strategy with ITOM Cloud Accelerate. Improve your organization's cloud utilization, drive down costs, automate requests, and aid in approvals and governance of the infrastructure that supports native cloud environments.

IT Operations Management products licensing

The ServiceNow AI Platform® uses a licensing method where your organization is billed for using the following IT Operations Management products: ServiceNow® ITOM Visibility, ServiceNow® ITOM Discovery, ServiceNow® ITOM AIOps, ServiceNow® Health Log Analytics, ServiceNow® ITOM Optimization, ServiceNow® ITOM Cloud Accelerate, ServiceNow® Service Reliability Management, ServiceNow®SLO Management, ServiceNow® Service Observability, and ServiceNow® synthetic monitoring.

The ServiceNow Product Documentation doesn't provide information on prices, packaging, or other details determined by your organization customer contract. For general information about licensing and subscriptions, see ITOM/OT SU Licensing and subscriptions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The ServiceNow Store regularly releases updates to applications created by ServiceNow. To access the latest features and updates for ITOM applications, deploy the latest upgrades from the ServiceNow Store. Each application may be on a different release schedule and is not specifically tied to a family release.

Get started

- Work with an implementation specialist to achieve your desired business outcomes. To learn more, visit the Customer Success Center.
- Choose the training that's best for you at the ServiceNow University center.

Learn

- What is AIOps?
- What is a configuration management database (CMDB)?
- What is IT operations (ITOps)?
- What is a network operations center (NOC)?
- What is a common service data model (CSDM)?
- What is cloud computing?
- What is cloud orchestration?
- What is IT operations management (ITOM)?
- What is network discovery?
- What is observability?

Applications and features

- ITOM licensing and subscriptions
- ITOM Cloud Accelerate
- ITOM Visibility
- Visibility: Discovery
- Visibility: Service Mapping
- Health: Event Management
- Health: Health Log Analytics
- Health: Agent Client Collector
- ITOM Optimization
- Optimization: Cloud Provisioning and Governance
- Service Reliability Management
- SLO Management
- Service Observability
- Synthetic monitoring

Install ITOM using ServiceNow Otto for Setup

ServiceNow Otto for Setup provides a guided setup experience to install and configure IT Operations Management.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow Otto for Setup for ITOM overview

You access ServiceNow Otto for Setup from the Admin Home page. The Manage your products section organizes your licensed ServiceNow products into dedicated spaces for focused setup and configuration. Select **IT Operations Management** to access tailored administrative tools and overviews.

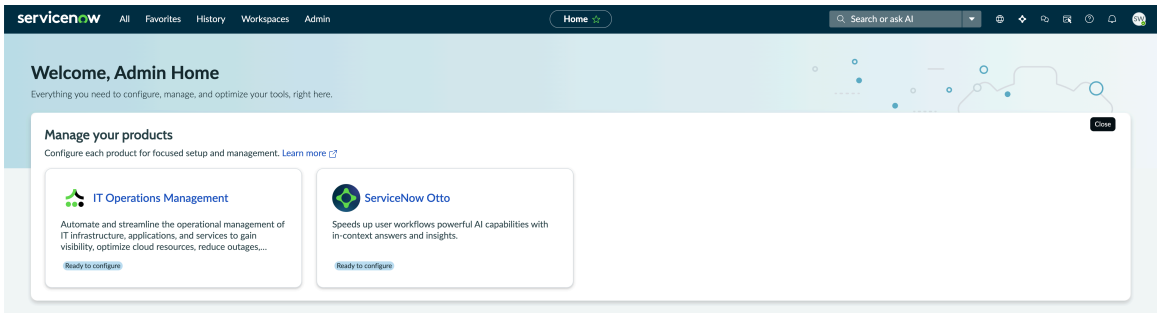

The Welcome! Let's get started section walks first-time administrators through the initial ITOM setup process. Select **Start setup** to install all required applications and plug-ins needed to begin managing your IT infrastructure.

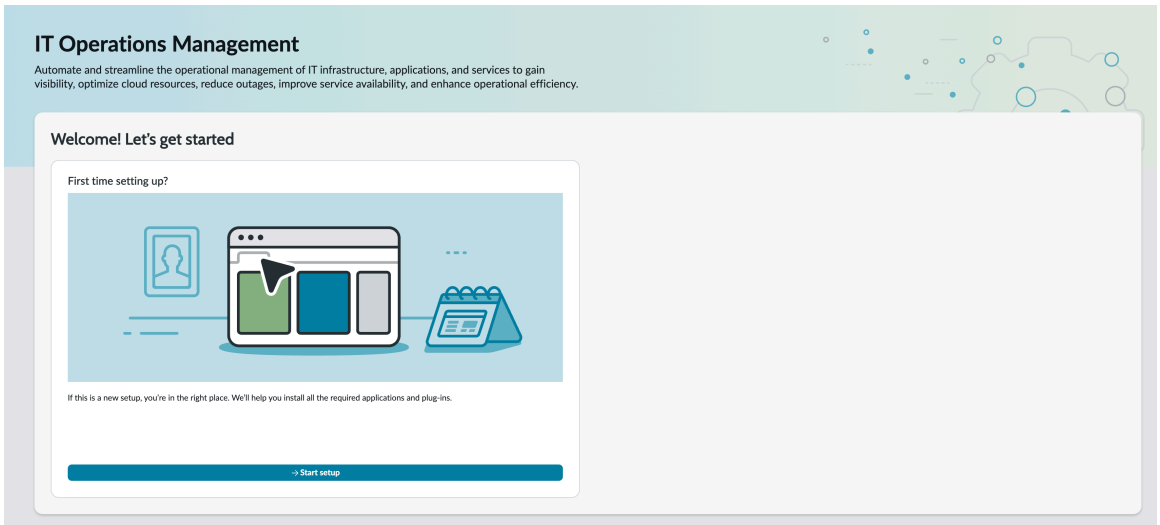

Instance information provides visibility into your ITOM setup progress, including installed applications, available entitlements, and ServiceNow Otto skill deployment status. This helps administrators track implementation milestones and verify that all required components are properly configured.

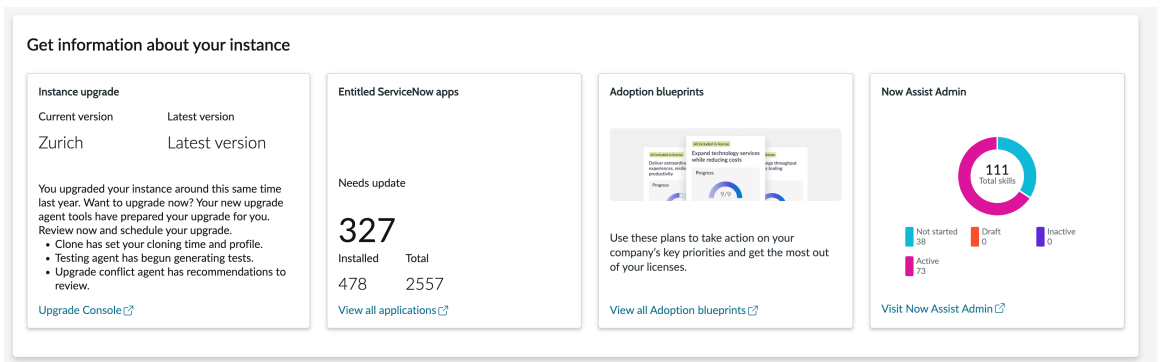

The Resources section helps you build foundational knowledge of ITOM. Watch an introductory video to get a quick overview of ITOM capabilities before diving into configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

General guidelines: Access strategies and guidelines for deploying, configuring, and optimizing ITOM to get the most out of your implementation.

Documentation: Find comprehensive guides, technical references, and general guidelines to support your ITOM setup and ongoing administration.

Community: Connect with other ServiceNow users and experts to share insights, ask questions, and get advice on ITOM.

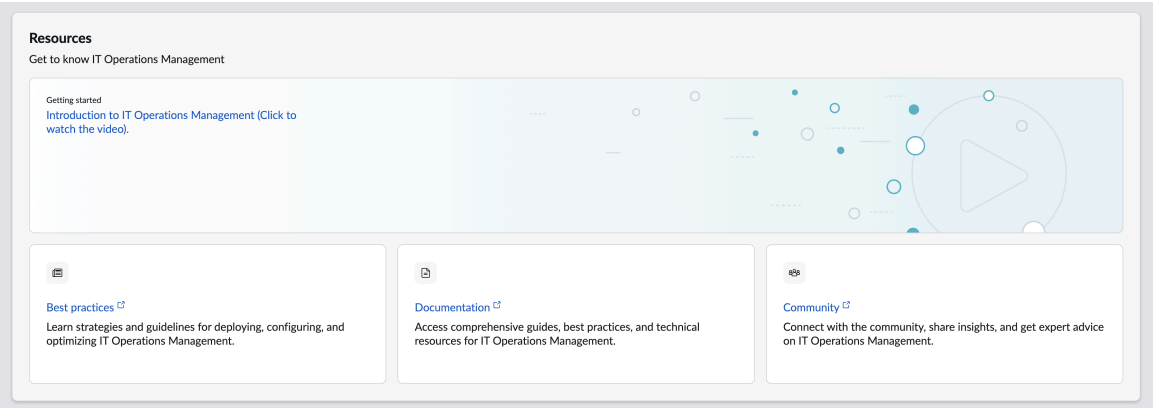

Install and configure ITOM applications

Once you use ServiceNow Otto for Setup to install ITOM, you can also use it to install ITOM Visibility and AIOps. Then, use the ITOM Configuration Console to configure Discovery and Event Management, two essential ITOM applications.

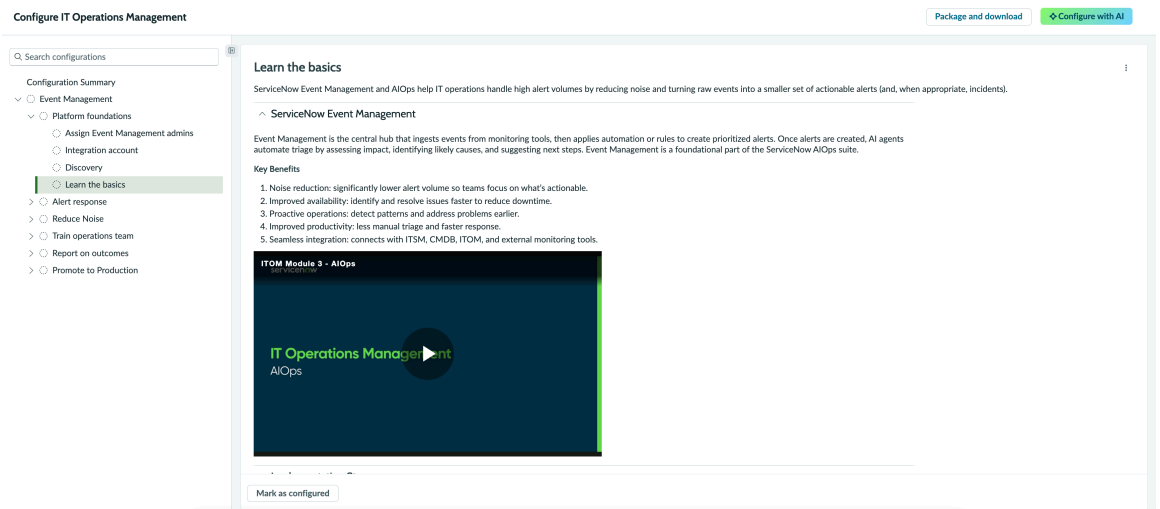

AI in IT Operations Management

AI features across IT Operations Management help operators triage alerts, investigate incidents, automate service mapping, diagnose infrastructure, and generate resolution artifacts using generative AI and agentic workflows.

The ServiceNow AI Platform now brings you a new AI experience with three licensing tiers available:

- Foundation: AI basics to deliver insights
- Advanced: AI to boost productivity across relevant use cases
- Prime: Act autonomously with all AI assets, and create your own

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see ServiceNow product tiers.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

For information about AI assets that are available on the ServiceNow AI Platform, see the following topics:

- Generative AI skills
- Agentic workflows

ITOM Visibility

For a summary of all AI features across ITOM Visibility, see AI in ITOM Visibility.

For AI features by product, see the following topics:

- AI in Discovery
- AI in Service Mapping
- AI in Firewall Audits and Reporting
- AI in Certificate Inventory and Management
- AI in LEAP

ITOM AIOps

For a summary of all AI features across ITOM AIOps, see AI in ITOM AIOps.

For AI features by product, see the following topics:

- AI in Express List
- AI in Service Observability
- AI in SLO Management
- AI in Agent Client Collector Monitoring

Important: Some generative AI skills, AI agents, and agentic workflows are turned on by default. For more information, see AI agents, skills, and agentic workflows on by default.

AI model providers

You can use Now LLM Service, Azure OpenAI, Google Gemini or Anthropic Claude on AWS as the AI model provider for all generative AI skills and AI agents. Use the Configuration Controls in AI Control Tower to define which options are available, then set the skill-level preferences in the AI Admin Hub console. For more information, see Large language models on the ServiceNow AI Platform.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

- Not all model providers are available for customers with in-country SKUs, and some AI products/features are currently unavailable for in-country customers. For more information, see the KB1584492 article in the Now Support Knowledge Base. Be sure to check for model provider availability updates in future releases.
- Some AI products/features are currently unavailable for customers in the FedRAMP, NSC DOD IL5, or Australia IRAP-Protected data centers, self-hosted customers, or in other restricted environments. For more information, see the KB0743854 article in the Now Support Knowledge Base. Be sure to check for availability updates in future releases.
- Some AI products/features are currently available only for customers in some regions. Be sure to check for availability updates in future releases.
- Some AI products and skills are not available in Regulated Markets. For more information, see KB2593939: Regulated Markets AI Products/Skills Not Available. Be sure to check for availability updates in future releases.

Installed agents

There might be AI agents installed on your instance that are not used in agentic workflows. To learn how to see all agents that are available to you, see Find AI agents.

Security

Enable security settings to run AI agents and agentic workflows using Access Control Lists (ACLs) and user identities. You can configure and manage the ACLs in AI Agent Studio. See Implement access control in AI agents for more information.

Dependencies

For a list of plugin dependencies, see Applications installed with ServiceNow Otto for ITOM.

Domain separation and generative AI

By default, all skills exist in the global domain. When you use AI in a domain-separated environment, users are only able to access data in their domain. For example, if a user uses the summarization skill, AI only uses material that exists in the user's domain when generating that summary. Additionally, there is no co-mingling of data for domain-separated instances when using generative AI skills. The data resides only on the instance, and the shared services used for generative AI do not persist any requests (prompts) and responses. For more information, see Domain separation in the AI Admin Hub console. (Note that global domain is not the same as global scope. For more information, see Exploring Next Experience pickers.)

Data processing

This application requires data to be transferred from ServiceNow customers' individual instances to a centralized ServiceNow environment, which may be located in a different data center region from the one where your instance is, and potentially to a third-party cloud provider, such as Microsoft Azure. This data is handled per ServiceNow's internal policies and procedures, including our policies available through our CORE Compliance Portal.

Data collection

ServiceNow collects and uses the inputs, outputs, and edits to outputs of this application to develop and improve ServiceNow technologies including ServiceNow models and AI products. Customers can opt out of future data collection at any time, as described in the AI Opt-Out page.

AI limitations

This application uses artificial intelligence (AI) and machine learning, which are rapidly evolving fields of study that generate predictions based on patterns in data. As a result, this application may not always produce accurate, complete, or appropriate information. Furthermore, there is no guarantee that this application has been fully trained or tested for your use case. To mitigate these issues, it is your responsibility to test and evaluate your use of this application for

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

accuracy, harm, and appropriateness for your use case, employ human oversight of output, and refrain from relying solely on AI-generated outputs for decision-making purposes. This is especially important if you choose to deploy this application in areas with consequential impacts such as healthcare, finance, legal, employment, security, or infrastructure. You agree to abide by ServiceNow's AI Acceptable Use Policy, which may be updated by ServiceNow.

Troubleshoot and get help

- ServiceNow Community on AI and Intelligence
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Install the ServiceNow Otto for IT Operations Management (ITOM) application

Install the ServiceNow Otto for IT Operations Management (ITOM) application.

Before you begin

Install the AIOps Experience [sn_sow_aiops] application from the ServiceNow Store. See Install Service Operations Workspace for ITOM application for more information.

Role required: admin

About this task

Install ServiceNow Otto for IT Operations Management (ITOM) plugin to enable generative AI and agentic AI on your instance.

Procedure

1. Go to the ServiceNow Store website.
2. In the search bar, enter `ServiceNow Otto for ITOM` and then select **Search**.
3. Select the ServiceNow Otto for ITOM widget and then select **Get** to make this application available for installation on your instance(s).
4. Install the ServiceNow Otto for ITOM application.
   a. Navigate to **All** > **System Definition** > **Plugins**.
   b. Search for ServiceNow Otto for IT Operations Management (ITOM).
   c. Select **Install**.

Related topics

View an alert analysis by ServiceNow Otto in Express List

Using the ITOM MCP Server Console to perform ITOM tasks

Connect an AI-enabled MCP Client application to your ServiceNow environment using the ITOM MCP Server Console. Use the MCP Client to investigate alerts, review configuration item (CI) reliability, assess incident impact, and create service level objectives (SLOs).

ITOM MCP Server Console overview

The ITOM MCP Server Console bridges your ServiceNow ITOM environment to any AI-enabled MCP Client, such as Moveworks or AWS Claude. When agents use an MCP Client to ask questions about alerts or CIs, the ITOM MCP Server Console fetches data from the ServiceNow ITOM environment. It performs actions on behalf of the agent through natural language conversation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The ITOM MCP Server Console supports these core capabilities for alert investigation and CI reliability:

- Retrieves a list of alerts based on parameters such as time frame, state, assignment, or severity.
- Retrieves alert details based on the alert number.
- Investigates, analyzes, and assesses potential impact of an alert.
- Performs actions on alerts: assign, close, reopen, acknowledge, and add work notes or comments. Supports single-alert and bulk operations.
- Reviews the reliability status of a CI, including related incidents, alerts, and SLOs.
- Reviews the reliability topology of a CI across upstream and downstream dependencies.
- Creates a standard availability SLO for a CI that doesn't already have one.

The ITOM MCP Server Console includes AIOps, SRM, and LEAP tools. For details on LEAP tools, see LEAP MCP Server.

ITOM MCP Server users

Users	Description
IT administrators	Activate the ITOM MCP Server Console.
IT operators	Use the ITOM MCP Server Console to investigate alerts, review CI reliability, assess incident impact on reliability, and create SLOs from an MCP Client application without switching between windows. Ask questions through the MCP Client application and request actions directly in the chat. Perform alert actions such as assign, close, reopen, and acknowledge on single alerts or alert groups.

How the ITOM MCP Server Console works

The ITOM MCP Server Console operates as a secure intermediary between any MCP Client application and the ServiceNow instance:

1. An agent opens their MCP Client application and asks a question about an alert.
2. The MCP Client application sends the question to the ITOM MCP Server Console using the Model Context Protocol (MCP).
3. The ITOM MCP Server Console authenticates the user against ServiceNow role-based access control.
4. The ITOM MCP Server Console retrieves the requested data or performs the requested action in the ServiceNow instance.
5. The ITOM MCP Server Console returns the result to the AI client application, which displays it to the agent.

Activate the ITOM MCP Server Console

Activate the ITOM MCP Server Console to enable AI-driven alert management capabilities on your ServiceNow instance through Model Context Protocol (MCP) connectivity.

Before you begin

Verify that the following plugins are installed on your instance:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- ServiceNow Otto for IT Operations Management (ITOM) (sn_itom_gen_ai)
- Event Management (com.glideapp.itom.snac)
- ITOM MCP Server Console (sn_itom_mcp_server)
- Model Context Protocol Server (sn_mcp_server)

To use the Alert Hypothesizer described in Investigate alerts using an MCP Client, you must also install Health Log Analytics plugin.

To use the CI reliability and SLO tools described in Review CI reliability with an MCP Client, you must also have the following plugins installed:

- Service Reliability Management (sn_sow_srm)
- Service Level Objective Management (sn_sow_slo)
- AI agents for SLO (sn_ai_agents_slo)

Refer to Connecting to an MCP server from an MCP client for more information about setting up the OAuth and connecting to the MCP Server Console.

Role required: sn_mcp_server.admin

Note: To set up your own OAuth connection, you need either an oauth_admin, or admin role to configure your OAuth client entry.

Procedure

1. Activate the ITOM MCP Server Console.

   a. Navigate to **All** > **MCP Server Console**.

   b. From the Configuration tab, select **Servers**.

   c. Select the **ITOM MCP Server**.

   Note: You must change the application scope to **ITOM MCP Server**

   The **MCP Server Console** screen opens with all fields populated.

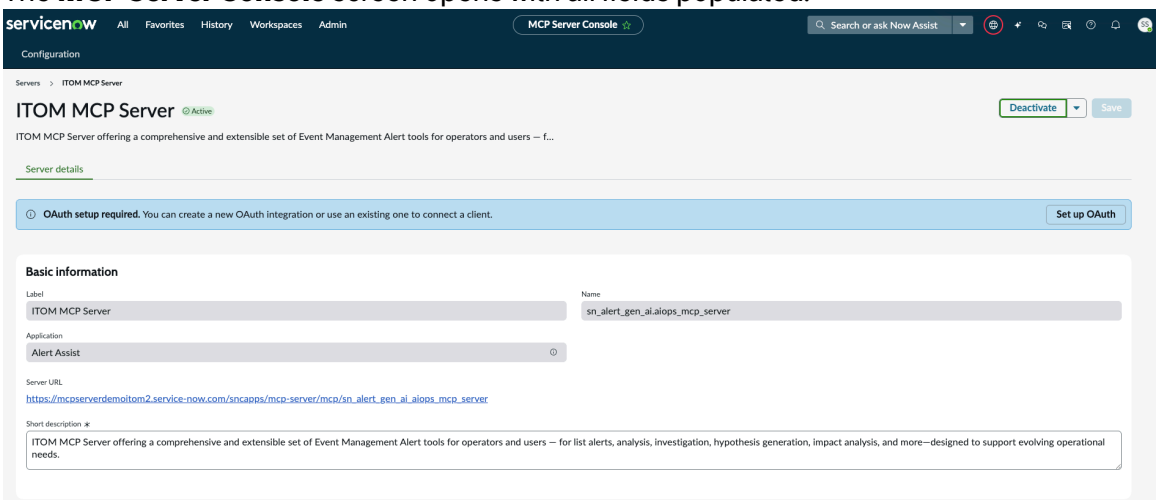

   d. From the Deactivate drop-down menu, select **Activate**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Set up OAuth to securely authenticate the ITOM MCP Server Console with your ServiceNow instance.
   Refer to the instructions in Connecting to an MCP server from an MCP client to setup the OAuth and connect to the MCP Server Console.

   Note: To set up your own OAuth connection, you need either an oauth_admin, or admin role to configure your OAuth client entry.

Related topics

- MCP Server Console
- Configuring MCP Server Console
- Connecting to an MCP server from an MCP client
- Install Model Context Protocol Client

Alert investigation with an AI client

Investigate alerts using the ITOM MCP Server Console with an MCP Client application.

The ITOMMCP Server Console enables you to analyze and investigate alerts using an MCP Client application, such as Moveworks or AWS Claude, connected to your ServiceNow instance. You can ask the AI client application questions about alerts through natural language prompts. MCP tools retrieve data from your ServiceNow instance and perform actions on your behalf.

The following tools are available to investigate alerts.

ITOMMCP Server Console alert investigation tools

Tools	Description
List Alert Record	Provides a list of alert records using parameters provided in the search query. **Note:** Alert record lists are limited to the first 50 alerts.
Alert Investigation	Checks for existing AI-generated investigation results for the alert and returns them if available. If no results exist, triggers the alert investigation pipeline in the background and returns a processing status until results are ready. This tool consolidates three separate tools: Alert Investigation, Alert Analysis, and Alert Impact Analysis.

Investigate alerts using an MCP Client

Use the ITOM MCP Server Console to investigate alerts through an MCP Client application, such as Moveworks or AWS Claude.

Before you begin

For more information about the Alert Investigation tool, see.

Role required: evt_mgmt_admin, evt_mgmt_operator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open your MCP Client application, such as AWS Claude, that is connected to your ServiceNow instance using the ITOM MCP Server Console.
   Your system administrator configures this integration during setup.
2. Ask the AI tool a question about an incident.

   Use natural language to ask questions. Here are common types of prompts:

   - Retrieve a list of alerts for today: "List alerts created today."
   - List only critical alerts: "List only the critical alerts created today."
   - Analyze and investigate the alerts: "Analyze and investigate the alerts."

     Note: Information returned through alert analysis may include several lines of text per alert.

   - Analyze a specific alert: "Analyze Alert 00100032."
3. Review the MCP Client application response.

   The MCP Client application retrieves data from your ServiceNow instance and presents it in the chat. The response includes only data you have permission to access.

Reviewing CI reliability with an MCP Client

Use an AI-enabled MCP Client that is connected to your ServiceNow environment through the ITOM MCP Server Console to review configuration item (CI) reliability and topology. In addition, the MCP Client can assess incident impact on reliability and create service level objectives (SLOs).

The following MCP tools retrieve information from your ServiceNow instance and, in some cases, perform actions.

Tool	Description
CMDB CI Reliability Status	Review the current reliability status of a CI, including active SLOs, related incidents and alerts, and affected dependent CIs.
CMDB CI Reliability Topology	Review the reliability topology of CI across upstream and downstream dependencies.
CMDB SLO for CMDB CI	Create a standard availability SLO for a CI that doesn't already have one.
Incident linked CI reliability	Review the reliability impact of an incident on a linked CI, including related alerts and impacted SLOs.

Review CI reliability with an MCP Client

Use an MCP Client with the ITOM MCP Server Console to review configuration item (CI) reliability and topology, assess incident impact on reliability, and create service level objectives (SLOs).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Before you begin
Verify that the ITOM MCP Server Console is active and the required plugins are installed on your instance. For more information, see Activate the ITOM MCP Server Console.

Role required: sn_sow_slo.slo_operator

Procedure
1. Open your MCP Client application.
2. In the MCP Client, ask a question or request an action for a CI.

   The following are common types of prompts:

   - Review the reliability impact of an incident: "What is the reliability impact of incident INC0010217?"
   - Review CI reliability status: "Show the reliability status of Customer Portal Service."
   - Review CI reliability topology: "Show the reliability topology of Customer Portal Service. I want to understand whether other CIs are affected."
   - Create an SLO for a CI: "Create an SLO for Auth Service."
3. Review the MCP Client response.

   The MCP Client retrieves data from your ServiceNow instance and presents it in the chat. The response includes only data that you can access.

What to do next
If you created an SLO, you can learn more about SLOs, including SLO types and compliance periods, in Reliability metrics in SLO Management. To edit or deactivate an SLO, see Edit a reliability metric.

Alert actions with an MCP Client

Manage alerts through an AI agent or MCP Client. The Alert Actions tool lets the agent or client act on alerts directly.

Alert Actions tool overview

MCP tools retrieve data from your ServiceNow instance and perform alert actions.

The Alert Actions tool enables you to assign, close, reopen, acknowledge, and annotate alerts through natural language prompts without leaving your MCP Client. The tool complements capabilities for retrieving and investigating alerts.

Important: The Alert Actions tool acts on alerts through an AI agent or MCP Client. You're interacting with AI. AI-generated responses may be inaccurate or incomplete. Review the confirmation and the linked alert after each action to verify the result.

Single vs. bulk alert operations

The Event Management Alerts table represents both individual alerts and alert groups. The Alert Actions tool automatically detects the scope:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Single-alert operations act on one alert at a time. Typical for direct investigation and targeted response.

  Users with the evt_mgmt_operator role can perform single-alert actions. They can assign to self, close after investigation, acknowledge receipt, and add diagnostic notes.

- Bulk operations act on multiple related alerts (alert group) in a single request. Efficient for coordinated team response to widespread issues.

  Users with the evt_mgmt_admin role can perform bulk operations on alert groups. They can reassign multiple related alerts to the network team, close a batch of resolved alerts, and acknowledge a group of new alerts.

Both modes support all actions described in the Supported actions section.

Supported actions

The Alert Actions tool supports the following operations on individual alerts and alert groups:

Alert Actions tool supported actions

Action	Description
Assign	Assign an alert to a user or group. Useful for routing work to the right team.
Close	Mark an alert as closed. Use when the underlying issue is resolved.
Reopen	Reopen a previously closed alert. Required if a closed alert recurs.
Acknowledge	Acknowledge receipt of an alert. Helps prevent duplicate escalations and signals active monitoring.
Add work notes	Add internal notes to an alert to record investigation details for other agents and operators.
Add comments	Add comments visible to other agents and operators working on the alert.

Use cases

The Alert Actions tool is useful in the following scenarios:

- An operator investigates an alert in an MCP Client, identifies the root cause, and immediately closes it with a note—all in the chat.
- A manager receives a notification about critical alerts, lists them in the MCP Client, assigns the entire batch to the on-call engineer, and tracks confirmation in real time.
- An operator acknowledges time-sensitive alerts to help prevent duplicate escalations, without navigating to the ServiceNow Event Management UI.

Workflow

The Alert Actions tool operates as part of the ITOM MCP Server Console workflow:

1. An operator or manager opens their MCP Client application and describes the action they need.

   For example: "Assign this alert to the network team".

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. The MCP Client sends the request to the ITOM MCP Server Console using the Model Context Protocol (MCP).
3. The ITOM MCP Server Console authenticates the user against ServiceNow role-based access control (RBAC).
4. The ITOM MCP Server Console performs the requested action on the alert record.
5. The ITOM MCP Server Console returns a confirmation to the MCP Client, including a deep link to the alert for verification in the UI if needed.

Perform actions on single alerts from an MCP Client

Use the Alert Actions tool in an MCP Client to perform single-alert operations. For example, to assign, close, reopen, acknowledge, or annotate one alert at a time.

Before you begin

Verify that:

- The ITOM MCP Server Console is active on your ServiceNow instance.
- Your MCP Client application is connected to the ITOM MCP Server Console with valid OAuth credentials.

Important: The Alert Actions tool acts on alerts through an AI agent or MCP Client. You're interacting with AI. AI-generated responses may be inaccurate or incomplete. Review the confirmation and the linked alert after each action to verify the result.

For more information about the Alert Actions tool, see Alert actions with an MCP Client.

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task

Individual alerts require targeted action based on investigation. For example, to investigate and close an alert:

1. In the MCP Client application (such as AWS Claude), ask: "Analyze alert 00100032."

   The alert details and analysis are retrieved from ITOM.
2. Review the analysis and determine the root cause.
3. Ask: "Close alert 00100032 with this note: Root cause identified as disk space depletion on DB server. Applied automated cleanup job."

   You receive confirmation: "Alert 00100032 closed. Work note added."
4. Select the provided link to verify in the ServiceNow UI.

Procedure

1. Open your MCP Client application.
2. Ask the MCP Client to perform an action on an alert.
   Use natural language to describe the action. For example:

   - "Assign alert 00100032 to the network team."
   - "Close alert 00100032. Root cause was a transient DNS issue."
   - "Reopen alert 00100032. The issue is recurring."

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- "Acknowledge alert 00100032."
- "Add a note to alert 00100032: Investigated at 10:15 UTC. CPU load returned to normal."

3. Review the response from your MCP Client.
   The response includes:
   - Confirmation of the action performed.
   - The updated alert state or field values.
   - A deep link to the alert in ServiceNow for verification.
4. **Optional:** Verify the action in the ServiceNow UI.
   Select the deep link provided in the response to open the alert in ServiceNow and confirm the changes.

Result
The Alert Actions tool performs your requested action on the alert. The alert state, assignment, or notes are updated in ServiceNow immediately.

What to do next
After performing an action:

- If you assigned the alert, notify the assigned user or team through your preferred communication channel.
- If you closed the alert, monitor for recurring symptoms or related alerts.
- If you reopened the alert, continue investigation or escalate if needed.

Perform bulk alert actions from an MCP Client

Use the Alert Actions tool in an MCP Client to assign, close, reopen, or acknowledge multiple related alerts as a group in a single request.

Before you begin
Verify that:

- The ITOM MCP Server Console is active on your ServiceNow instance.
- Your MCP Client application is connected to the ITOM MCP Server Console with valid OAuth credentials.

Important: The Alert Actions tool acts on alerts through an AI agent or MCP Client. You're interacting with AI. AI-generated responses may be inaccurate or incomplete. Review the confirmation and the linked alert after each action to verify the result.

For more information about the Alert Actions tool, see Alert actions with an MCP Client.

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task
Alerts often occur in groups related to the same underlying issue or infrastructure failure. The Alert Actions tool performs bulk operations that coordinate a response across multiple related alerts. Bulk operations apply the same action to every alert in the group, which keeps handling consistent across the group.

For example, to respond to a widespread outage:

1. In the MCP Client application (such as AWS Claude), ask: "List all critical alerts created in the last hour."

   12 related alerts are returned. These alerts span database, network, and storage systems.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Determine the root cause affects all systems.

   Ask: "Assign all 12 of these alerts to the infrastructure team. Add note: Outage traced to power supply failure in DC1. Maintenance started at 14:30 UTC."

   You receive: "All 12 alerts assigned to infrastructure team. Work note added to each alert."

3. Verify by clicking the link to the alert group.
4. Later, after the issue is resolved, ask: "Close all 12 of these alerts. Add note: Power supply replaced. Full system recovery confirmed."

   You receive: "All 12 alerts closed."

Procedure

1. Open your MCP Client application.
2. Ask the MCP Client to list related alerts.
3. Identify the alert group or set of related alerts you want to manage.
   For example: "List all critical alerts related to the database service created in the last 2 hours."
4. Request a bulk action on the alert group.
   Use natural language to describe the bulk action. For example:
   - "Assign all these alerts to the database team."
   - "Close all database-related alerts created today. Add note: Issue resolved by maintenance window."
   - "Reopen all alerts related to the load balancer. New symptoms detected."
   - "Acknowledge the 5 network alerts listed above."
   - "Add this note to all storage alerts: Escalated to infrastructure team for investigation."
5. Review the response from your MCP Client.
   The response includes:
   - Count of alerts affected by the bulk action.
   - Confirmation that the action succeeded.
   - Deep links to the alert group or individual alerts for verification.
6. **Optional:** Verify the bulk action in the ServiceNow UI.
   Select the provided link to open the alert group and confirm all changes were applied correctly.

Result

The Alert Actions tool performs the bulk action on all selected alerts. Each alert is updated with the new state, assignment, or notes in ServiceNow immediately.

What to do next

After performing a bulk action:

- If you assigned a group of alerts, notify the assigned team about the scope and nature of the alerts.
- If you closed a group, monitor for recurrence or new related alerts that might indicate incomplete resolution.
- If you made bulk assignments, follow up to confirm the team is investigating.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert remediation with an MCP Client

Get remediation suggestions for an alert from an AI agent or MCP Client, matching what's available in Service Operations Workspace.

Alert remediation tool overview

MCP tools retrieve data from your ServiceNow instance and perform alert actions.

The Alert remediation tool lets you request a remediation suggestion for an alert through natural-language prompts in your MCP Client, without switching to the ServiceNow UI. The tool routes your request to the existing remediation agent that already runs in Service Operations Workspace, so you get the same suggestion regardless of channel.

Important:

The Alert remediation tool suggests actions through an AI agent or MCP Client. You're interacting with AI. AI-generated responses may be inaccurate or incomplete. Review the suggestion and the linked alert before taking action.

Supported capabilities

The Alert remediation tool supports the following capabilities:

- Accepts a request for a specified alert and returns a remediation suggestion based on existing workflows, knowledge base articles, and related context.
- Returns the remediation recommendation directly in the response and prompts you to execute or dismiss it.
- Lets you execute the suggested remediation directly through the MCP Client.
- If a remediation fails to execute, the tool offers a fallback action, such as creating an incident, and requests confirmation before running it.

Considerations

Consider the following when using the Alert remediation tool:

- Initial availability may cover only a subset of alert types or remediation actions.
- Remediation suggestions are scoped to the alerts you're authorized to view.

Use cases

The Alert remediation tool is useful in the following scenarios:

- An operator asks their MCP Client, such as AWS Claude, "What should I do about this alert?" and receives a remediation suggestion grounded in existing workflows and knowledge base articles.
- An operator working from an external AI agent reviews a suggested remediation and confirms before executing it.

Workflow

The Alert remediation tool operates as part of the ITOM MCP Server Console workflow:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. An operator asks the external agent what to do about a specified alert.
2. The MCP Client sends the request to the ITOM MCP Server Console using the Model Context Protocol (MCP).
3. The ITOM MCP Server Console authenticates the request against ServiceNow role-based access control.
4. The request is routed to the existing remediation agent, which evaluates the alert against workflows, knowledge base articles, and related context.
5. The ITOM MCP Server Console returns the remediation recommendation to the MCP Client, which prompts the operator to execute or dismiss it.

Related topics

Use the ITOM MCP Server Console

Alert actions through an MCP Client

Remediate alerts with an MCP Client

Get and execute an alert remediation workflow or suggestion through an MCP Client.

Before you begin

Verify that:

- The ITOM MCP Server Console is active on your ServiceNow instance.
- Your MCP Client application is connected to the ITOM MCP Server Console with valid OAuth credentials.
- The alert you want to remediate has been investigated, either by the AI Specialist or through the Alert Investigation tool.

  If the alert has not yet been investigated, initiate an investigation using the Alert Investigation tool. For more information, see Investigate alerts using an MCP Client.

⚠ **Warning:** The Alert remediation tool provides suggestions through an AI agent or MCP Client. You're interacting with AI. AI-generated responses may be inaccurate or incomplete.

For more information about the Alert remediation tool, see Alert remediation with an MCP Client.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Open your MCP Client application.
2. Ask the MCP Client how to handle a specified alert, using natural language.
   Either the recommended remediation workflow for the alert is returned, or if none is available, a remediation suggestion is returned.
3. Confirm whether to execute or dismiss the recommendation.
   If you confirm, the remediation is executed.
4. If the remediation fails to execute, confirm whether to proceed with the fallback action.
   For example, creating an incident.

MID-less log ingestion for HLA from an MCP Client

Set up MID-less log ingestion integrations for Health Log Analytics (HLA) directly from an AI-enabled MCP Client, without opening the ServiceNow UI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using an MCP Client to set up a MID-less integration

If the HLA app is installed on your ServiceNow instance, the ITOM MCP Server Console lets you set up HLA integrations for MID-less log data ingestion. You can set up integrations from any AI-enabled MCP Client. You can use an MCP Client, such as AWS Claude, to create the integrations entirely through natural-language prompts, without navigating to the Integrations Launchpad.

The HLA ingest tools associated with the ITOM MCP Server Console enable you to perform the following tasks from the MCP Client:

- Find existing MID-less ingest integrations by name.
- Create an ingest integration or rename an existing one. The MCP Client returns the integration ID, HTTP endpoint, gRPC endpoint, and access token for a newly created integration.
- Turn an integration on or off. An integration must be deactivated before it can be renamed.
- Generate a new access token for an integration. The previous token remains valid until its original expiration.

Workflow

MID-less ingestion integration setup is part of the ITOM MCP Server Console workflow:

1. An administrator opens their MCP Client application and describes the integration they need.

   For example: "Create an OTel ingestion integration named otel-prod."

   Note: By default, integrations are created in an inactive state. However, the MCP tools are capable of creating an active integration. For example, tell the MCP Client: "Create an active OTel ingestion integration named otel-prod."

2. The MCP Client sends the request to the ITOM MCP Server Console using the Model Context Protocol (MCP).
3. The ITOM MCP Server Console authenticates the user against ServiceNow role-based access control.
4. The ITOM MCP Server Console creates or updates the ingest integration record in HLA.
5. The ITOM MCP Server Console returns the integration details, including endpoint and token information, to the MCP Client.

Related topics

Set up MID-less log ingestion using an MCP Client

Use the ITOM MCP Server Console to perform ITOM tasks

MID-less log streaming for HLA

Set up MID-less log ingestion using an MCP Client

Create, configure, and activate MID-less log ingestion integrations for Health Log Analytics from an MCP Client.

Before you begin

Verify that:

- The ITOM Gateway is set up and configured on your ServiceNow instance. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.
- The ITOM MCP Server Console is active on your instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Your MCP Client application is connected to the ITOM MCP Server Console with valid OAuth credentials.
- The Health Log Analytics plugin is installed on your instance.

For more information about setting up MID-less integrations from an MCP Client, see MID-less log ingestion for HLA from an MCP Client.

Role required: evt_mgmt_admin or sys_admin (for JWT provider and token configuration)

Procedure

1. Open your MCP Client application.
2. In the MCP Client application, ask to create a MID-less integration.
   Use natural language. For example: "Create an OTel integration named otel-prod."
   If no integration with that name exists, the integration is created in a draft, inactive state. The following information is returned:
   - Integration ID
   - HTTP endpoint
   - gRPC endpoint
   - Access token
3. Review the returned integration details.
4. Ask the MCP Client to activate the integration.
5. Request the setup information for your integration type.
   The following table lists what to ask for, based on integration type.

Integration type	What to do
OTel	Ask your MCP Client to generate an OTel collector `exporters` configuration block, then add it to your `collector.yaml` file.
Other MID-less integration types	Ask your MCP Client to help you complete the setup information specific to that integration type outside ServiceNow. For example, for Amazon Data Firehose, ask it to help you set up a Firehose stream in your AWS account.

6. **Optional:** Ask the MCP Client to rename the integration.

   Note: The integration must be inactive to be renamed. The integration is automatically deactivated, renamed, and reactivated.

Result

The integration is active and ready to receive log data at the returned endpoint. On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

What to do next

Verify that log data is arriving in HLA.

To rotate credentials, ask the MCP Client to refresh the integration's token. The previous token remains valid until its original expiration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Applications installed with ServiceNow Otto for ITOM

Table that lists applications installed with the ServiceNow Otto for IT Operations Management (ITOM) application. When you update your application, any newly required application dependencies are installed.

Applications installed with ServiceNow Otto for ITOM (sn_itom_gen_ai)

Application name	Description
ServiceNow Otto for Platform (sn_genai_platform)	Enables AI skills and features across various BU applications.
AI agents for AIOps (sn_aiops_ai_agents)	Enables IT operations with AI Agents that can help autonomously triage alerts, assess business and technical impact, investigate root causes, and recommend or execute remediation steps. These agents operate within coordinated agentic workflows.
AI Agents for Service Mapping (sn_sm_gen_ai)	Provides autonomous AI Specialists for Service Mapping. It automates the creation of ML-based service maps from Application Service Candidates and the connection of Business Applications to discovered Application Services in the CMDB.
AI agents for Observability (sn_obs_aia)	Helps IT operators assess business and application service impact, and formulate probable cause theories. Helps prioritize investigations by analyzing data from ServiceNow and collaborating with third-party AI agents from leading APM and observability vendors.
LEAP (sn_itom_leap)	LEAP (Learning-Enhanced Automation Playbook) leverages AI-driven insights to mine historical incident data, dynamically prioritize tasks, and generate actionable resolution playbooks.

Related topics

- Plugins or applications installed with ITOM Cloud Accelerate
- Plugins or applications installed with ITOM Visibility
- Plugins or applications installed with ITOM AIOps
- Plugins activated with ITOM Optimization

Now Assist for IT Operations Management (ITOM)

Set up, manage, and optimize your generative-AI solutions on the ServiceNow AI Platform from a single workspace.

Get started with AI

Starting with the Australia release, Now Assist is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

The ServiceNow AI Platform now brings you a new AI experience with three licensing tiers available:

- Foundation: AI basics to deliver insights
- Advanced: AI to boost productivity across relevant use cases
- Prime: Act autonomously with all AI assets, and create your own

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see ServiceNow product tiers.

For information about AI assets that are available on the ServiceNow AI Platform, see the following topics:

- Generative AI skills
- Agentic workflows

Important:

- Not all model providers are available for customers with in-country SKUs, and some AI products/features are currently unavailable for in-country customers. For more information, see the KB1584492 article in the Now Support Knowledge Base. Be sure to check for model provider availability updates in future releases.
- Some AI products/features are currently unavailable for customers in the FedRAMP, NSC DOD IL5, or Australia IRAP-Protected data centers, self-hosted customers, or in other restricted environments. For more information, see the KB0743854 article in the Now Support Knowledge Base. Be sure to check for availability updates in future releases.
- Some AI products/features are currently available only for customers in some regions. Be sure to check for availability updates in future releases.
- Some AI products and skills are not available in Regulated Markets. For more information, see KB2593939: Regulated Markets AI Products/Skills Not Available. Be sure to check for availability updates in future releases.

You can use Now LLM Service, Azure OpenAI, Google Gemini or Anthropic Claude on AWS as the AI model provider for all generative AI skills and AI agents. Use the Configuration Controls in AI Control Tower to define which options are available, then set the skill-level preferences in the AI Admin Hub console. For more information, see Large language models on the ServiceNow AI Platform.

AI limitations

This application uses artificial intelligence (AI) and machine learning, which are rapidly evolving fields of study that generate predictions based on patterns in data. As a result, this application may not always produce accurate, complete, or appropriate information. Furthermore, there is no guarantee that this application has been fully trained or tested for your use case. To mitigate these issues, it is your responsibility to test and evaluate your use of this application for accuracy, harm, and appropriateness for your use case, employ human oversight of output, and refrain from relying solely on AI-generated outputs for decision-making purposes. This is especially important if you choose to deploy this application in areas with consequential impacts such as healthcare, finance, legal, employment, security, or infrastructure. You agree to abide by ServiceNow's AI Acceptable Use Policy, which may be updated by ServiceNow.

Data processing

This application requires data to be transferred from ServiceNow customers' individual instances to a centralized ServiceNow environment, which may be located in a different data center region from the one where your instance is, and potentially to a third-party cloud provider, such as Microsoft Azure. This data is handled per ServiceNow's internal policies and procedures, including our policies available through our CORE Compliance Portal.

Data collection

ServiceNow collects and uses the inputs, outputs, and edits to outputs of this application to develop and improve ServiceNow technologies including ServiceNow models and AI products. Customers can opt out of future data collection at any time, as described in the AI Opt-Out page.

ITOM/OT SU Licensing and subscriptions

ServiceNow® IT Operations Management (ITOM) licensing calculates and displays the usage of ITOM subscriptions based on subscription units.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Unified statistics for all ITOM products

Use the ITOM licensing module to view subscription-related information for all ITOM products: ServiceNow® ITOM Visibility, ServiceNow® ITOM Discovery, ServiceNow® ITOM AIOps, ServiceNow® Health Log Analytics, ServiceNow® ITOM Optimization, ServiceNow® ITOM Cloud Accelerate, ServiceNow® Digital End-User Experience (DEX), and ServiceNow® Service Observability.

The ServiceNow Product Documentation doesn't provide information on prices, packaging, or other details determined by your organization customer contract.

Consumption level

Easily see how many subscriptions for IT Operations Management applications your organization purchased and consumed.

Explore	Install	References
Explore ITOM/ OT SU Licensing and subscriptions	Learn how to install ITOM/OT SU Licensing	Learn about schedule jobs, tables, subscription form and dashboard

To know about Operational Technology Management (OTM) licensing and subscriptions, see Subscriptions form for the OTM products.

Exploring ITOM/OT SU Licensing

Learn more about ServiceNow® ITOM/OT SU Licensing, the different ITOM subscriptions it includes, and its benefit.

ITOM/OT SU Licensing overview

The ITOM/OT SU Licensing application calculates and displays how many subscriptions for IT Operations Management (ITOM) applications your organization purchases and consumes. Subscriptions for the following ITOM applications are available: ServiceNow® ITOM Visibility, ServiceNow® ITOM Discovery, ServiceNow® ITOM AIOps, ServiceNow® Health Log Analytics, ServiceNow® ITOM Optimization, ServiceNow® ITOM Cloud Accelerate, ServiceNow® Digital End-User Experience (DEX), and ServiceNow® Service Observability. You can purchase subscriptions for each ITOM application separately (a la carte) or together (in a bundle).

ITOM/OT SU Licensing benefits

Benefit	Feature	Users
Viewing the usage of ITOM applications to make informed decisions about your ITOM subscriptions. You can use it to confirm your expectation of their license consumption, and	Review resource usage against allocated subscription units	sn_itom_license.reader

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefit	Feature	Users
if the measurement is higher than planned, either alter your usage of ITOM or increase your subscription.		
Analyzing resource usage and status in comparison to your purchased subscriptions is facilitated by dashboard reports. These reports visually present both the daily usage count and the average consumption of subscription units over a 90-day period and let you see which resources are being counted for each ITOM application.	ITOM/OT SU Licensing dashboard	admin
Comparing your current CI consumption against the predicted usage under the new 2026 Packaging SKU, so you can plan your subscription requirements before migration to 2026 packaging.	Predict CI consumption for 2026 Packaging SKU	admin

Subscriptions for IT Operations Management

The ServiceNow AI Platform® uses a licensing method where your organization is billed for using the following IT Operations Management products: ServiceNow® ITOM Visibility, ServiceNow® ITOM Discovery, ServiceNow® ITOM AIOps, ServiceNow® Health Log Analytics, ServiceNow® ITOM Optimization, ServiceNow® ITOM Cloud Accelerate, ServiceNow® Service Reliability Management, ServiceNow®SLO Management, ServiceNow® Service Observability, and ServiceNow® synthetic monitoring.

For licensing purposes, licensable resources are assigned to CI categories. License usage is measured using subscription units, which may cover a different number of resources, depending on the resource CI category. The customer contract for your organization determines the CI categories from the available options, and also the ratios of CI allocation for subscription units in each CI category. Information from the customer contract is synchronized with the ServiceNow platform. The latest version of ServiceNow® ITOM/OT SU Licensing supports the following CI categories:

- Containers
- PaaS
- FaaS
- Servers
- Unresolved Monitored Object
- Unclassed category
- End User Devices
- IoT
- Customer Prem Devices

You have the option to buy subscriptions for individual ITOM products separately (a la carte) or as part of a bundle with multiple products. You can also purchase both—a bundle subscription that covers most required applications and an a la carte subscription for applications that consume more subscription units. For more information, see ITOM/OT SU Licensing subscription types.

ServiceNow® charges for the following IT Operations Management products:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM Visibility

The ServiceNow® ITOM Visibility product consists of ServiceNow® Discovery, Firewall Audits and Reporting, Certificate Inventory and Management, Service Mapping, Cryptographic Asset Compliance, AI Agent Topology Mapping, ITOM Content Service, Tag Governance, Cloud Discovery Workspace, and Cloud License Estimator applications. Discovery and Service Mapping give you a unified, connected view of your entire IT network and the services that it supports. The Service Graph Connectors and CMDB 360 applications provide data to ITOM Visibility.

In addition, Service Graph Connectors populate the CMDB by bringing in data from third-party monitoring and management systems. CMDB 360 enables you to track how the CMDB contains various discovery sources at the CI attribute level. Use Certificate Inventory and Management to avoid service outages due to expired TLS certificates. Firewall Audits and Reporting helps you to discover and manage your firewall security policies to heighten the security awareness for firewall audits.

For licensing purposes, ITOM Visibility provides the CI count it discovers, both during horizontal and top-down discovery. It also reports CIs that the Service Graph Connectors populate the CMDB with.

You can also use limited Tag Governance functionality under the ITOM Visibility subscription. You can identify on-premises resources that are inconsistent and don't comply with the tag policies of your organization.

ITOM Discovery

ServiceNow® ITOM Discovery is an entry-level package alternative where customers can build a healthy data platform. This package measures licensed usage of the Discovery feature and Service Graph Connectors. Discovery is a "subset" of ITOM Visibility that is also based on subscription units. ITOM Discovery includes the following basic features: ServiceNow® Discovery, Service Graph Connectors, and CMDB 360. ServiceNow® Discovery finds applications and devices on your network, and then updates the CMDB. Service Graph Connectors supplement the population of the CMDB, while CMDB 360 retains the complete history of discovery sources and proposed values.

ITOM Discovery provides the count of CIs found by ServiceNow® Discovery and brought in by Service Graph Connectors.

ITOM AIOps

The ServiceNow® ITOM AIOps product includes the ServiceNow® Event Management and ServiceNow® Metric Intelligence applications, which help you track and maintain the health of services in your organization.

ITOM AIOps reports the number of CIs, where it captures alerts, events, and metric data. ITOM AIOps reports the CI count for all CI categories.

Health Log Analytics

Health Log Analytics CI count is based on log data received from licensable resource types. The consumption is reflected in ITOM AIOps as well as in Health Log Analytics. For that reason, you need the ITOM AIOps subscription in addition to the Health Log Analytics subscription.

ITOM Cloud Accelerate

The ITOM Cloud Accelerate product consists of the following workflow applications: Tag Governance, Cloud Migration Assessment, Cloud Configuration Governance,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

and Cloud Action Library. In Tag Governance, you can update CI tags in the cloud environment, in addition to updating them in the CMDB.

For licensing purposes, the ITOM Cloud Accelerate license provides the CI count for CIs managed and monitored by Cloud Migration Assessment, Cloud Configuration Governance, or Cloud Action Library. If different ITOM Cloud Accelerate features report the same CIs as their resources, the CI count is manipulated to remove the duplication.

ITOM Optimization

The ServiceNow® ITOM Optimization application provides automation for the cloud workflows used to manage the cloud resources throughout their life cycle. It enables certified and enterprise-compliant cloud deployment, cost visibility, and other cloud management processes.

Service Observability

The ITOM Service Observability feature creates actionable dashboards and observability tabs that centralize data from existing third-party monitoring tools. It allows you to track the performance of business-critical services, enables faster root cause analysis, and reduces mean time to resolution (MTTR). Application and platform telemetry insights from multiple observability tools are combined into a centralized workflow.

Related topics

Subscription Management

Data collection and aggregation for licensing process

ITOM/OT SU Licensing application counts CIs for ITOM applications and uses a daily average count for the last 90 days to produce license statistics for purchased subscription units.

ServiceNow charges for various ITOM applications, including ITOM Visibility, ITOM Discovery, ITOM AIOps, Health Log Analytics, ITOM Cloud Accelerate, and ITOM Optimization. For specific information about the products and features covered by ITOM subscriptions, see Subscriptions for IT Operations Management.

The process of collecting and aggregating information for licensing purposes involves the following steps:

1. The ITOM licensing mechanism counts the Configuration Items (CIs) managed by each ITOM product daily, categorizing them into licensable CI categories. For more details, see KB0748149: ITOM Subscription Unit license calculation logic.
2. If the same CIs are managed by different features within the same ITOM products, the CI count is adjusted to eliminate the duplication. For example, if both Discovery and Service Mapping discover the same CIs, these CIs appear only once in the CI count. Container CIs are excluded from SU count if they have a relationship to a container image.
3. The licensing module aggregates CI counts from ITOM applications to determine the average daily CI count for the last 90 days.
4. The licensing module then compares the daily average CI counts for the last 90 days from ITOM applications with the licensing information in the customer contract to generate license statistics.

As a result, you can view the statistics on how your organization uses purchased subscription units.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM application license flow

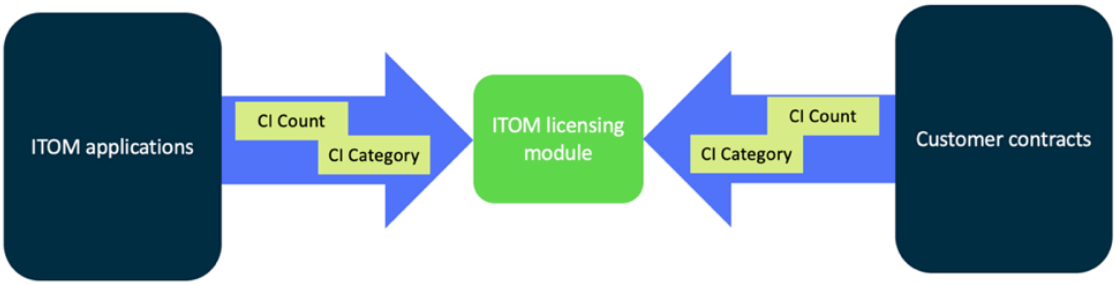

ITOM/OT SU Licensing subscription types

Purchase licensing subscriptions for ITOM products individually (a la carte), bundle several products together, or opt for both individual and bundled subscriptions.

You have the option to buy subscriptions for individual ITOM products separately (a la carte) or as part of a bundle with multiple products. You can also purchase both—a bundle subscription that covers most required applications and an a la carte subscription for applications that consume more subscription units.

If your organization uses more subscriptions than estimated, you have the option to buy additional subscription units as needed. These units can be purchased for either a bundle or individual products (a la carte). However, if your organization uses ITOM applications in a bundle at various levels, you can choose to buy both — a bundle subscription covering most required applications and an a la carte subscription for those that require more subscription units.

Note: ITOM bundles also provide entitlement to Integration Hub to select spokes and protocols. ITOM Cloud Accelerate requires the additional Integration Hub Enterprise bundle entitlement.

ITOM bundles also provide entitlement to Workflow Data Fabric to select spokes and protocols. ITOM Cloud Accelerate requires the additional Workflow Data Fabric Enterprise bundle entitlement for uses of Cloud Action Library from anything other than the ITOM Cloud Accelerate applications. Cloud Account Management requires the Cloud Governance Suite SKU.

ITOM/OT SU Licensing Bundle subscription

You can purchase subscriptions for several ITOM products together as part of a bundle.

The ITOM/OT SU Licensing application calculates and displays subscription consumption for bundle subscriptions. When your organization purchases subscriptions in a bundle, you get the same number of subscription units for all ITOM applications covered by the bundle. You cannot use bundle subscriptions for other applications, even if these other applications are part of the same bundle. The consumption level of the bundle is defined by the highest number of consumptions by one of the products in the bundle.

When you purchase subscriptions in a bundle, you receive an equal number of subscription units for all ITOM applications covered by that bundle. For instance, if you buy a bundle of 500 units covering ITOM Visibility and ITOM AIOps, your organization receives 500 units for each application.

Note: You can only purchase one bundle at a time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subscriptions window displaying subscriptions purchased by bundle

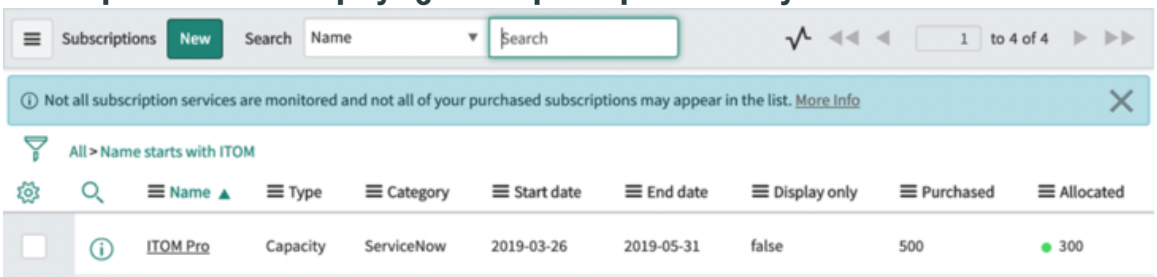

Bundle subscriptions are specific to certain applications. You cannot utilize bundle subscriptions for other applications, even if those applications are part of the same bundle. For instance, if you purchase a bundle of 500 units covering ITOM Visibility and ITOM AIOps, you receive 500 units for each application. If you use all 500 ITOM Visibility units, the remaining ITOM AIOps units cannot be used for ITOM Visibility.

The consumption level of the bundle is determined by the highest usage among the products included in the bundle.

Automatic adjustment of bundle to actual consumption level

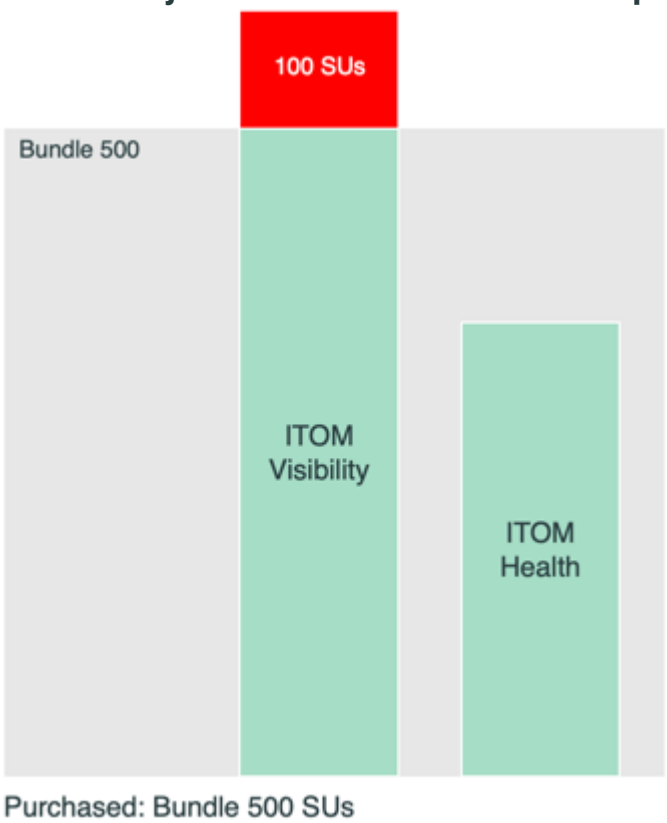

The Subscriptions window displays the actual number of consumed subscriptions under **Allocated**. The red color code indicates that the bundle is overdrawn.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subscriptions window showing overdraft on subscriptions purchased by bundle

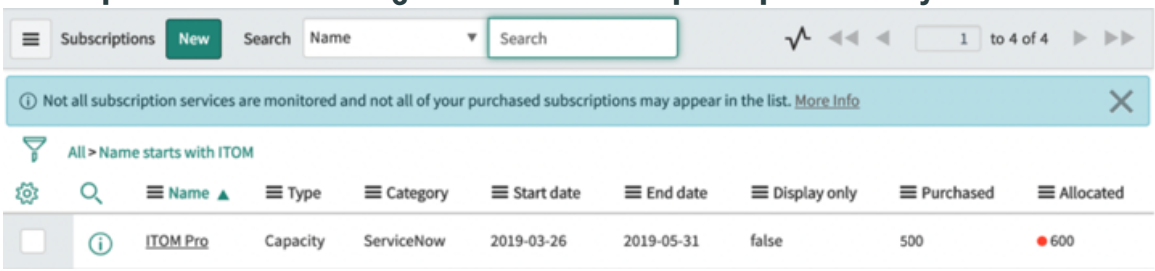

ITOM/OT SU Licensing Individual (ala carte) subscription

You can purchase subscriptions for individual ITOM applications (ala carte).

The ITOM/OT SU Licensing application calculates and displays subscription consumption for ITOM applications purchased individually (a la carte). If your organization exceeds the number of subscription units, the Subscriptions window shows the relevant a la carte subscription as overdrawn. You cannot use unconsumed subscriptions for overdrawn subscriptions of another application.

The Subscriptions window displays information about the purchased and allocated subscriptions for ITOM applications.

Subscriptions window displaying subscriptions purchased a la carte

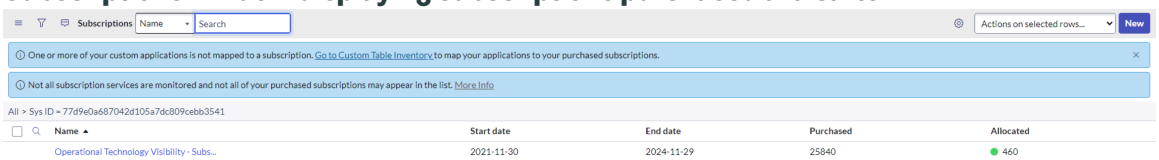

If your organization surpasses the number of purchased subscriptions for an ITOM application, you cannot utilize the remaining subscriptions for another application.

Calculation of the a la carte subscriptions

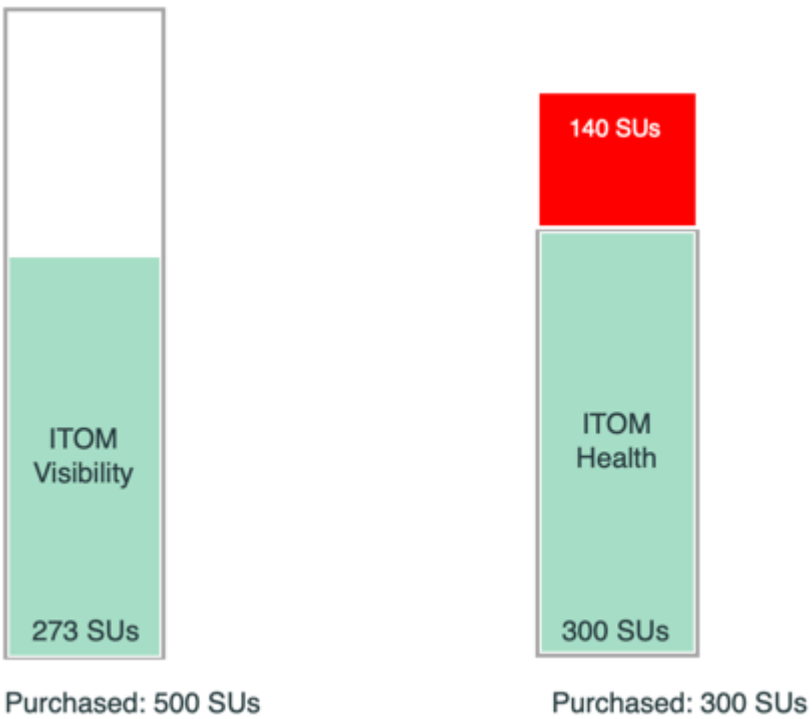

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If your organization exceeds the total number of subscription units purchased, the Subscriptions window indicates the relevant a la carte subscription as overdrawn. In the following illustration, ITOM AIOps consumed more subscriptions than the amount purchased for it.

A la carte subscriptions in an overdraft

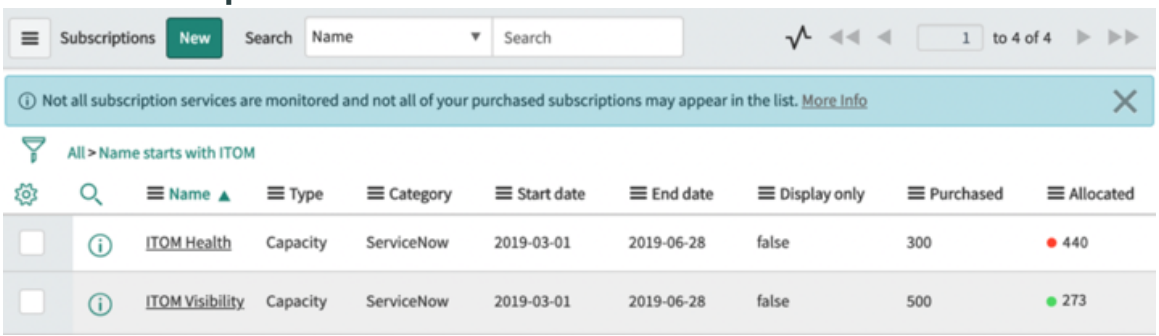

Name	Type	Category	Start date	End date	Display only	Purchased	Allocated
ITOM Health	Capacity	ServiceNow	2019-03-01	2019-06-28	false	300	440
ITOM Visibility	Capacity	ServiceNow	2019-03-01	2019-06-28	false	500	273

ITOM/OT SU Licensing Bundle and Individual (ala carte) subscription

You can purchase subscriptions for individual ITOM applications (a la carte) or as a bundle covering multiple ITOM applications.

The ITOM/OT SU Licensing application calculates and displays the subscription usage for ITOM products. It displays the ITOM products individually or as part of a bundle, based on your subscription. If your organization has purchased ITOM subscriptions both as a bundle and individually (a la carte), the licensing module deducts from the bundle first. It then deducts from the individually purchased subscriptions.

If your organization surpasses the allocated number of subscription units, the Subscriptions window indicates the corresponding subscriptions as overdrawn. If the combined number of subscriptions from both bundle and a la carte options exceeds the allocation, the licensing module identifies the surplus. The surplus is taken from the a la carte subscription. If the total subscriptions for an application covered solely by the bundle surpass the allocation, the Subscriptions window displays the bundle as overdrawn. The individual application is not listed separately.

Subscription deduction process: Bundle usage and a la carte scenario

You can purchase ITOM subscriptions as a bundle and also acquire a la carte subscriptions for an application included in the bundle. The licensing module deducts from the bundle first. It then deducts from the a la carte subscriptions.

In the following example, an organization utilized all subscriptions allocated by the ITOM Pro bundle for ITOM AIOps. Subsequently, the licensing module deducted from the a la carte subscriptions for ITOM AIOps.

Subscription consumption by bundles

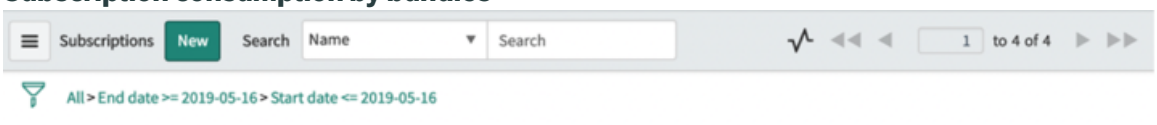

Name	Type	Category	Start date	End date	Display only	Purchased	Allocated
ITOM Health	Capacity	ServiceNow	2019-03-01	2019-06-28	false	250	0
ITOM Visibility	Capacity	ServiceNow	2019-03-01	2019-06-28	false	250	0
ITOM Pro	Capacity	ServiceNow	2019-03-26	2019-05-31	false	500	250

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subscription display: Application covered by bundle and a la carte subscriptions

If your organization surpasses the total number of subscriptions in both the bundle and a la carte subscriptions, the licensing module regards it as overdrawn on the a la carte subscription. However, if your organization exceeds the subscriptions specifically for an application covered solely by the bundle, it is considered overdrawn on the bundle.

Licensing module calculates bundle and a la carte subscriptions for the same ITOM applications

Subscription calculation illustration: Bundle and a la carte allocation

The diagram illustrates how the licensing module calculates subscriptions through bundle and a la carte methods for specific applications covered by the bundle.

Overdrafts on the bundle and a la carte subscriptions

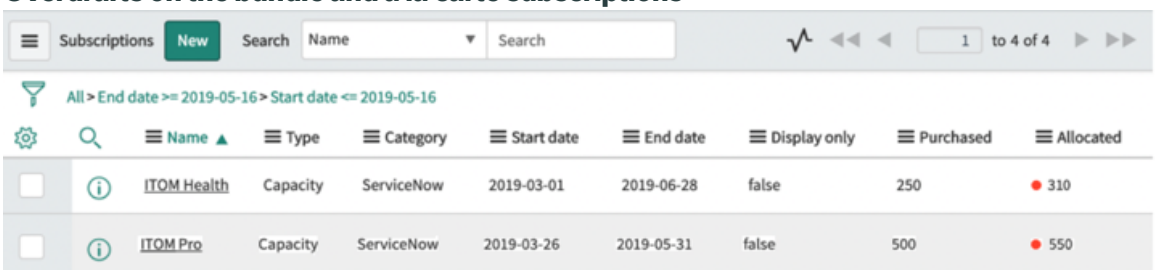

In this scenario, the Subscriptions window indicates the a la carte subscriptions for this application as overdrawn. If your organization exceeds the total subscriptions for an application covered solely by the bundle, the Subscriptions window displays the bundle as overdrawn.

In the provided figure, ITOM AIOps consumed 860 subscriptions, surpassing the subscriptions provided by the bundle (500) and a la carte (250). Consequently, the Subscriptions window indicates the ITOM AIOps a la carte subscription as overdrawn.

For ITOM Visibility, which consumed 550 subscriptions, and with no a la carte subscriptions purchased, the ITOM Pro bundle covering ITOM Visibility appears as overdrawn.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subscriptions for Operational Technology Management

The ServiceNow platform uses Operational Technology Management (OTM) to manage licenses for the manufacturing industry. OTM contains the licenses that are included in IT Operations Management, as well as licenses unique to OTM.

Monitor OTM licenses in the same way as IT Operations Management licenses. The ServiceNow instance indicates OTM license types on the **ITOM Licensing Category MetaData** (**OTM License** > **Licenses by CI types**) page, where the **SKU Type** column contains a value of **otm**. By default, the license filter is set to **SKU Type contains otm**.

The same license options available for IT Operations Management are also available for OTM, under the **OTM License** menu item in a ServiceNow instance.

For details on subscriptions for OTM, see Subscriptions for IT Operations Management.

Configuring ITOM/OT SU Licensing

Gain the capability to efficiently manage and optimize ITOM subscriptions within the ServiceNow framework for streamlined operations.

Configuration overview

Begin by installing ServiceNow® ITOM/OT SU Licensing, setting the groundwork for a comprehensive understanding of managing ITOM subscriptions.

Update ServiceNow® ITOM/OT SU Licensing

Update the ServiceNow® ITOM/OT SU Licensing [com.snc.itom.license] application to ensure you use the latest licensing functionality. Updating the application installs related ServiceNow® Store applications and plugins if they are not already installed.

Before you begin

- Ensure you have an ITOM product installed on your ServiceNow instance for the ITOM/OT SU Licensing application to work. Installing the ITOM/OT SU Licensing application is not required, as the system automatically installs the application alongside any ITOM product.
- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the ITOM/OT SU Licensing application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

You receive notifications when updates for ITOM/OT SU Licensing application are available on ServiceNow Store.

The following items are installed with ITOM/OT SU Licensing:

- Scheduled jobs
- Tables

For more information, see Components installed with ITOM/OT SU Licensing.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the ServiceNow ITOM/OT SU Licensing application (sn_itom_licensing) using the filter criteria and search bar.

   You can search for the application by its name or ID.
3. In the **Updates** tab, select the ServiceNow ITOM/OT SU Licensing application.
4. Select **Proceed to update**.

   In the Install dialog that is displayed, any dependencies that are installed when the application is updated are listed.
5. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
6. Select the version you want to upgrade to and select **Install**.

Using ITOM/OT SU Licensing

Ensure optimal utilization of your organization's ITOM subscriptions by analyzing the resource usage of your ITOM products.

Using ITOM/OT SU Licensing overview

Explore how to view and track subscriptions for ITOM applications, analyzing resource usage against allocated subscription units to ensure optimal utilization. Verify the count of Configuration Items (CIs) used for ITOM subscriptions and identify which CI types specifically require ITOM subscription units. Finally, generate lists showcasing the current countable CIs for each ITOM application, allowing you to visualize their consumption of subscription units.

View subscription statistics for ITOM

View how many subscriptions for IT Operations Management applications your organization purchased and consumed.

Before you begin

Ensure that your organization has purchased ITOM subscriptions. You cannot view the information in the **ITOM License** module without subscriptions.

Ensure that you installed the latest available version of the ITOM/OT SU Licensing from ServiceNow Store.

Role required: sn_itom_license.reader

About this task

Review the configuration item (CI) allocation and allocation level to see how your organization uses IT Operations Management subscriptions and plan upcoming subscription needs.

Procedure

1. To view subscription information for IT Operations Management subscriptions purchased a la carte and in bundles, navigate to **ITOM License** > **License Summary**.
   Use subscription names to differentiate between subscriptions bought in bundles and a la carte.
2. Review the information on the form as covered in Subscriptions form for the ITOM products.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

KB0748149: ITOM Subscription Unit license calculation logic

Subscriptions for IT Operations Management

Review resource usage against allocated subscription units

Review statistics for the resources that can be managed by ITOM products. Compare this data to the average number of allocated subscription units.

Before you begin

Role required: sn_itom_license.reader

See ITOM detailed licensing data over time with the ITOM Licensing Dashboard. View daily CI counts or average of the last 90 daily counts. Domain-specific information and specific CI listings for each daily count enable you to monitor and analyze resource usage. You can generate a list of currently countable CIs for each of the ITOM applications as described in View CIs consuming ITOM subscription units.

Procedure

1. Navigate to **All** > **ITOM License** > **ITOM Licensing Dashboard**.
2. Select the relevant tab to display the dashboard for the ITOM product.
   For example, select **Visibility**.

ITOM Visibility dashboard

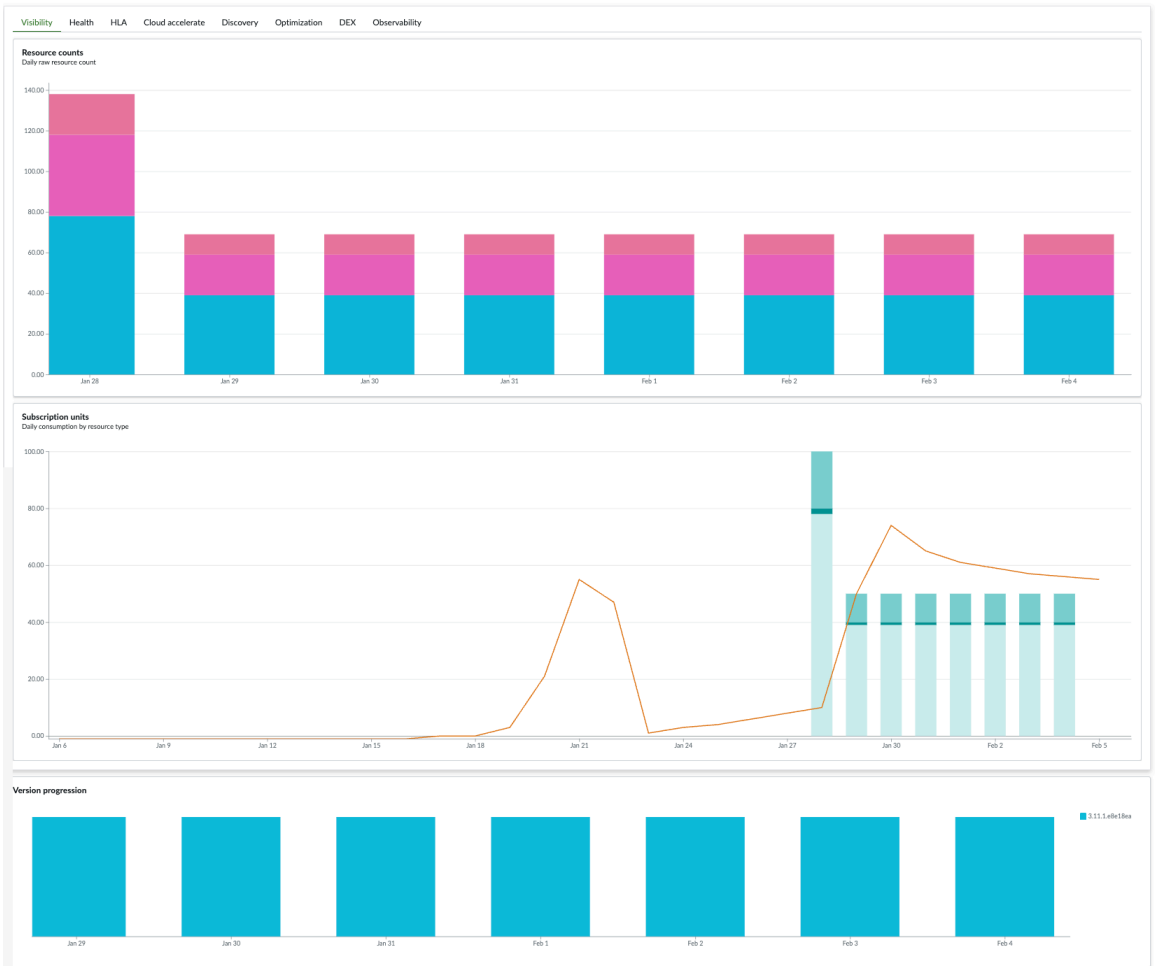

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Point to a bar to view the number of CIs in each category.
4. Select the bar to view a list of the counted CIs.
5. Review the dashboard described in ITOM/OT SU Licensing dashboard.

Check CI count used for ITOM subscriptions

View daily CI counts or the average of the last 90 daily counts.

Before you begin

Ensure that your organization has purchased ITOM subscriptions. You cannot view the information in the **ITOM License** module without subscriptions.

Role required: sn_itom_license.reader

About this task

ITOM applications, such as ITOM Visibility, ITOM Discovery, ITOM AIOps, Health Log Analytics, ITOM Cloud Accelerate, and ITOM Optimization provide information on the resources they handle. These resources, called configuration items (CIs), are what ITOM applications find, monitor, and store in the CMDB. The ITOM/OT SU Licensing module combines this CI data with the subscription information of your organization to produce statistics on the subscription usage and consumption of the ITOM applications. For more information about the licensing workflow process, see Data collection and aggregation for licensing process.

This information does not include any statistics on subscriptions purchased in bundles. For complete information on subscriptions, view subscription statistics for IT Operations Management.

Procedure

1. Navigate to **ITOM License** > **License Report**.

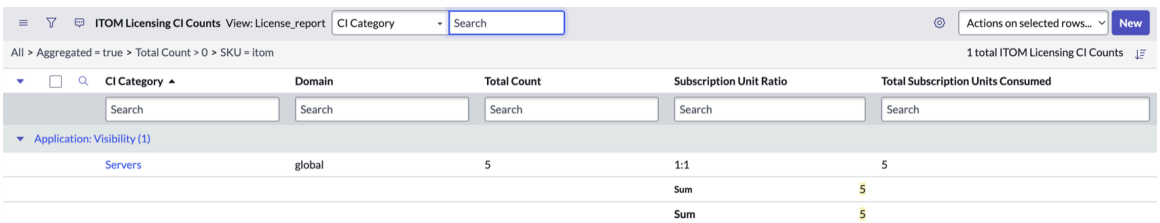

   To know more about the fields in the Licence Report window, see License Report form.

2. In the **Total Count** column, view the average daily CI counts for the last 90 days.
3. View the daily CI count in one of the following ways.
   - Set the **Aggregated** filter condition to **[Aggregated] [is] [false]** and select **Run**.
   - Navigate to **ITOM License** > **License Daily Usage Count**.

Check which CI types require an ITOM subscription

View the list of configuration item (CI) types that require IT Operations Management subscription units.

Before you begin

Ensure that your organization has purchased ITOM subscriptions. You cannot view the information in the **ITOM License** module without subscriptions.

Role required: sn_itom_license.reader

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

For licensing purposes, licensable resources are assigned to CI categories. License usage is measured using subscription units, which may cover a different number of resources, depending on the resource CI category. The customer contract for your organization determines the CI categories from the available options, and also the ratios of CI allocation for subscription units in each CI category. Information from the customer contract is synchronized with the ServiceNow platform. The latest version of ServiceNow® ITOM/OT SU Licensing supports the following CI categories:

- Containers
- PaaS
- FaaS
- Servers
- Unresolved Monitored Object
- Unclassed category
- End User Devices
- IoT
- Customer Prem Devices

The system stores information about CI categories and CI types that belong to them in the ITOM Licensing MetaData [itom_lu_metadata] table. This table includes information only about parent CI types. All CI types that are created by extending these parent CI types are included by default. For example, the Apache Web Server CI type is included in the CI count, because it is an extension of the Web Server CI type.

Containers are only counted toward SU consumption if they don't have a relationship to a container image CI.

Procedure

1. Check which CI category is licensable:

   a. Navigate to **All** > **ITOM License** > **License Summary**.

   b. Select the relevant ITOM subscription.
   For example, **ITOM Visibility**.

   c. Scroll down and select the **ITOM CI subscription unit ratio** tab.
   The tab lists licensable CI categories generated from the customer contract. These are the licensable CI categories.

   d. Make a note of the licensable CI categories provided for this ITOM subscription.

   e. **Optional:** Return to the **License Summary** list and check CI categories generated for other ITOM subscriptions.

   The customer contract for your organization determines the ratios of CI allocation for subscription units in each CI category, and this information is synchronized with the ServiceNow platform.

2. Navigate to **All** > **ITOM License** > **Licenses by CI Type**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: The ITOM Licensing MetaData [itom_lu_metadata] table contains all CI categories, including categories of CIs that do not require ITOM subscription units.

3. Review the list to identify licensable CI categories that appear on the **ITOM CI subscription unit ratio** tab of the **License Summary** list.
   The customer contract for your organization determines the ratios of CI allocation for subscription units in each CI category, and this information is synchronized with the ServiceNow platform.
4. Select the arrow adjacent to the licensable CI category to expand it and access the associated CI types.

   The CI types listed under CI categories require ITOM subscription units for access.

View CIs consuming ITOM subscription units

Generate a list of currently countable CIs for each of the ITOM applications: ITOM Visibility, Discovery, ITOM AIOps, Health Log Analytics, ITOM Cloud Accelerate, and ITOM Optimization.

Before you begin

Ensure that your organization has purchased ITOM subscriptions. You cannot view the information in the **ITOM License** module without subscriptions.

Role required: sn_itom_license.reader

About this task

The generated CI list is closely correlated to the latest daily CI count. At the same time, the number of CIs on the generated CI list may be slightly different from the latest daily CI count, if changes occurred since the last daily count.

Note: To estimate the predicted usage under 2026 Packaging SKUs, see Predict CI consumption for 2026 Packaging SKU.

Procedure

1. Navigate to **All** > **ITOM License** > **Report ITOM Licensable CIs**.
   The **Report ITOM Licensable CIs** page appears.
2. Select the application for which you want to view licensed CIs:
   - Discovery
   - Optimization
   - Visibility
   - HLA
   - Governance
   - Health
3. To create a report for the selected application, select **Populate licensable CIs**.
4. To accept the confirmation message and generate the report, select **Yes**.

   The new report replaces the data in the previously generated report. You can cancel the report by selecting the application and then selecting **Cancel Job**.
5. Wait for several minutes and then refresh the page.
   The **Status** of the application appears **Completed** when the report is complete.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Select the application and then select **Show licensable CIs**.
   The ITOM Licensable CIs page displays the list of CIs with a license.

Predict CI consumption for 2026 Packaging SKU

Forecast your licensing needs by comparing your current ITOM subscription unit usage against the predicted usage under 2026 Packaging SKUs.

Before you begin

Check your entitlements to determine whether you have ITOM subscriptions.

Ensure that you installed the latest available version of the ITOM/OT SU Licensing from ServiceNow Store.

Role required: admin

About this task

Before migrating to the 2026 Packaging SKUs, you can estimate post-migration consumption without requiring an active ITOM license. Predict SU consumption per category by applying SU ratios to the latest CI data in your environment. SU ratios represent the number of CIs in each category that correspond to one Subscription Unit, as set by your contract. Use this prediction to confirm your expectation of license consumption. If the measurement is higher than planned, either alter your usage of ITOM or increase your subscription.

Procedure

1. Navigate to **All** > **ITOM License** > **Report ITOM Licensable CIs**.
2. Select the check box for the application for which you want to view licensed CIs and select **Show Licensable CIs**.
   After the *`Usage count estimator`* job completes, results are added to the ITOM Licensable CIs [itom_lu_licensable_cis] table and are visible in the **Report ITOM Licensable CIs** view.

Prediction results example

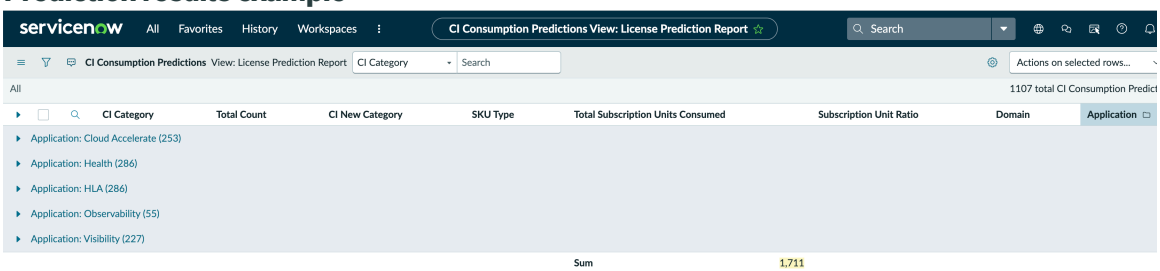

> **Note:** *`CI New Category`* equals *`true`* if this category is not covered by your current license. It will be introduced as a new licensable category under the 2026 Packaging SKU.

3. See only the prediction records and not production licensing records by setting the filter conditions to **Contains Estimate**.
4. View the latest daily SU consumption prediction results per CI category created by usage count estimator jobs that run automatically on a daily basis.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to the ITOM Licensable CIs [itom_lu_licensable_cis] table by typing the table name in the navigation filter.

b. See the latest daily prediction results per each CI category by setting the filter conditions to **Created-on-today**.

What to do next

Compare the predicted SU consumption data to the estimate you had before this forecast.

ITOM/OT SU Licensing Reference

Reference topics provide additional information about scheduled jobs, tables, subscriptions form and dashboard that you use to configure and administer ITOM/OT SU Licensing.

Components installed with ITOM/OT SU Licensing

Several types of components are installed with activation of the ITOM/OT SU Licensing plugin, including tables and scheduled jobs.

Scheduled jobs installed

Scheduled job	Description
ITOM Exclusion Tables Update Store	Updates the exclusion list.
ITOM Licensing Aggregator Store	Calculates the average of daily CI counts for the last 90 days.
ITOM Licensing Cloud Accelerate CI Listing Store	Creates the list of licensable CIs for ITOM Cloud Accelerate.
ITOM Cloud Accelerate Licensing Usage Count Store	Calculates the daily CI counts for ITOM Cloud Accelerate.
Event Management - Node Count Store	Calculates the node count used in Event Management.
ITOMHealthCIReporterWithOTOMCountITOMStore	Compiles the list of licensable CIs for ITOM.
ITOMHealthCIReporterWithOTOMCountOTOMStore	Compiles the list of licensable CIs for Operational Technology Management (OTM).
ITOM Health Licensing Usage Count Store	Calculates the daily CI count for ITOM AIOps.
ITOM Optimization Licensing Usage Count Store	Calculates the daily CI count for ITOM Optimization.
ITOM Licensing Optimization CI Listing Store	Compiles the list of licensable CIs for ITOM Optimization.
OTOM Licensing Visibility CI Listing Store	Compiles the list of licensable CIs for ITOM Visibility (OTM).
ITOM Licensing Discovery CI Listing Store	Compiles the list of licensable CIs for Discovery.
ITOM Licensing Visibility CI Listing Store	Compiles the list of licensable CIs for ITOM Visibility.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scheduled job	Description
ITOM Visibility Licensing Usage Count Store	Calculates the daily CI count for ITOM Visibility.
OTOM Licensing Discovery CI Listing Store	Compiles the list of licensable CIs for Discovery (OTM).
ITOM DEX Licensing Usage Count Store	Calculates the daily CI counts for ITOM DEX.
ITOM Licensing DEX CI Listing Store	Creates the list of licensable CIs for ITOM DEX.
ITOM Licensing Service Observability CI Listing Store	Creates the list of licensable CIs for ITOMService Observability.
ITOMService Observability Licensing Usage Count Store	Calculates the CI counts for ITOMService Observability.
ITOM Licensing Historical Table Update Store	Creates a table to store all the licensable CIs daily

Tables installed

For information on how ITOM/OT SU Licensing counts licensable CIs for different ITOM products, refer to ITOM Subscription Unit license calculation logic (KB0748149).

Table	Description
ITOM LU Discovery Source Mapping [itom_lu_discovery_source_mapping]	Contains the list of licensable discovery source for each category.
ITOM LU Governance App Mapping [itom_lu_governance_app_mapping]	List of records that contain the mapping of governance applications to their respective licensable CIs.
ITOM LU Governance CIs [itom_lu_governance_ci]	Contains the list of CIs counted under the Governance license.
ITOM License Exclusion Metadata [itom_license_exclusion_metadata]	Contains the list of exclusion rules applicable to different license.
License Exclusions [license_exclusion_list]	Contains the list of CIs that need to be excluded from the license count based on the exclusion rule.
Visibility LU Temporary [visibility_lu_temp]	Contains the list of CIs counted under the Discovery license.
ITOM Licensing Category MetaData [itom_lu_category_metadata]	Contains licensing metadata.
ITOM Licensing Discovery Sources	Contains the categories for all discovery sources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[itom_lu_discovery_sources]	
ITOM LU DEX CIs [sn_itom_licensing_itom_lu_dex_ci]	Contains the list of CIs counted under the DEX license.
Historical Licensable CIs [sn_itom_licensing_itom_lu_historical_cis]	Contains the list of Historical Licensable CIs.

Subscriptions form for the ITOM products

Description of the fields on the Subscriptions form for the ITOM products.

ITOM products subscription form fields

Field	Description
Name	Name of either subscription bundle or ITOM application, if your organization purchased subscriptions per application separately (a la carte).
Start date	Date the subscription was activated.
End date	Date till when the subscription is valid.
Purchased	Number of purchased subscriptions per bundle or application (a la carte).
Allocated	Number of consumed subscription units from a bundle or application (a la carte). For bundles, this field shows the highest number of consumed subscription units by applications that are part of the bundle. For example, if ITOM Visibility consumed 300 subscription units and ITOM AIOps consumed 200 subscription units, this field shows 300 for the bundle covering these applications. The color code indicates the percentage of the subscription units that your organization consumed. By default, the color code threshold is 90%. • Green: Your organization has used less than 10% of purchased subscriptions. • Yellow: Your organization has used more than 90%, but less than 100% of purchased subscriptions. • Red: Your organization has used 100% or more and exceeded the number of purchased subscriptions. Purchased subscriptions are overdrawn.

ITOM/OT SU Licensing dashboard

Use the ITOM Licensing dashboard to review the statistics of the resource consumption and status against your purchased subscriptions. This dashboard contains reports for each ITOM application. The reports offer visualization of the daily usage count and of the average consumption of subscription units for 90 days.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The ITOM Licensing dashboard is a part of ITOM Licensing application. See also OTM Licensing Dashboard at **All** > **OTM** > > **OTM License** > **OTM Licensing Dashboard**.

Role required: sn_itom_license.reader

Access the Assessment dashboard

Ensure that you have an sn_itom_license.reader role to access the ITOM Licensing Dashboard located at **All** > **ITOM License** > **ITOM Licensing Dashboard**.

ITOM Licensing dashboard

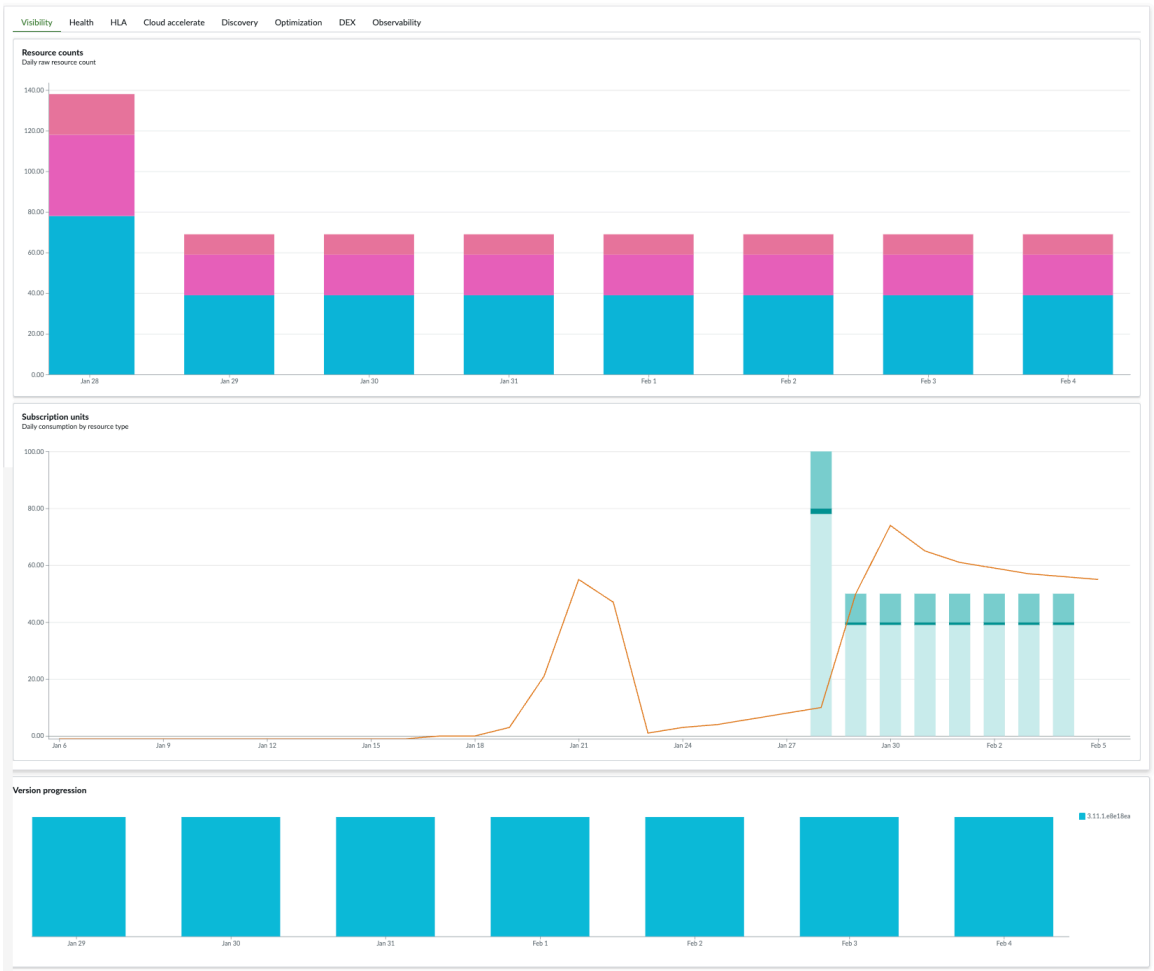

Use cases

As an example of how people in your organization would use this dashboard, see the following use case.

User	Dashboard use
sn_itom_license.reader	Validate the resource usage for different ITOM products. Report the cases where the organization exceeded the number of purchased subscription units for specific resources.

Note: ServiceNow applications refer to devices and applications that comprise a service instance as configuration items (CIs).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data visualization

The dashboard includes the following visualization:

Usage	Type	Source table	Description
• ITOM Visibility • ITOM AIOps • Health Log Analytics • ITOM Cloud Accelerate • ITOM Discovery • ITOM Optimization • ITOM Digital End-User Experience (DEX) • ITOM Service Observability	§	ITOM Licensing CI Counts [itom_lu_ci_counts] and UsageAnalytics Counts for Tables [usageanalytics_count]	The first area displays the daily raw resource counts that include all the detected resources before sorting them into specific CIs. Hover over the bars to see the table of resource types. Each type is shown with an absolute number of resources and a share that it takes among others. The second area displays bars that represent Subscription units counts for different licensable categories for the last 120 days per ITOM application. The dashboard also displays the brown line that represents the Usage Analytics counts table. This count is based on average consumption of subscription units for the last 90 days. SU ratios are one server, three PaaS resources, three containers, or one unresolved monitored object equal 1 SU. Hover over the vertical bar in the Subscription units area for the desired day to open the Available CIs list. The table shows each CI category included in the daily

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Usage	Type	Source table	Description
			CI count, along with the absolute number of resources and the share that it takes among others. The third area shows which version of the ITOM Licensing is used on the instance each day. It helps explain spikes in count every time the new version is installed.

License Report form

Description of the fields on the License Report form.

License Report form fields

Field	Description
CI Category	Grouping of Configuration Items (CIs) in each application based on shared characteristics or attributes. For the list of supported CI categories, including 2026 Packaging SKUs, see Subscriptions for Operational Technology Management
Domain	Domain to which the CI belongs.
Total Count	Average of 90 daily counts of CIs categorized by CI types, within each ITOM application. **Note:** The total count adheres to the above definition only if the **Aggregated** condition in the filter is set as **true**. When the **Aggregated** condition is set as **false**, the total count represents the daily count for a specific application and category on that day.
Subscription Unit Ratio	Number of CIs, in a specific category that need a subscription, that are counted as one Subscription Unit. SU ratio is determined by your customer contract.
Total Subscription Units Consumed	Total number of subscription units consumed by ITOM applications within a specific timeframe. The licensing module calculates this number by using the subscription units ratio with the count of CIs in each CI category.
Sum	Total consumption of subscription units by the application and all ITOM applications.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

License Report form fields (continued)

Field	Description
	Appears highlighted under Total Subscription Units Consumed.

Note: To estimate the predicted usage under 2026 Packaging SKUs, see Predict CI consumption for 2026 Packaging SKU

ITOM Cloud Accelerate

ITOM Cloud Accelerate workflows streamline cloud automation across the cloud adoption journey via self-service catalogs and controlled workflows. It expedites application migration by facilitating assessment, planning, and resource migration tracking.

Features of ITOM Cloud Accelerate

ITOM Cloud Accelerate enables organizations to validate their cloud assets against predefined standards and correct any deviations.

ITOM Cloud Accelerate

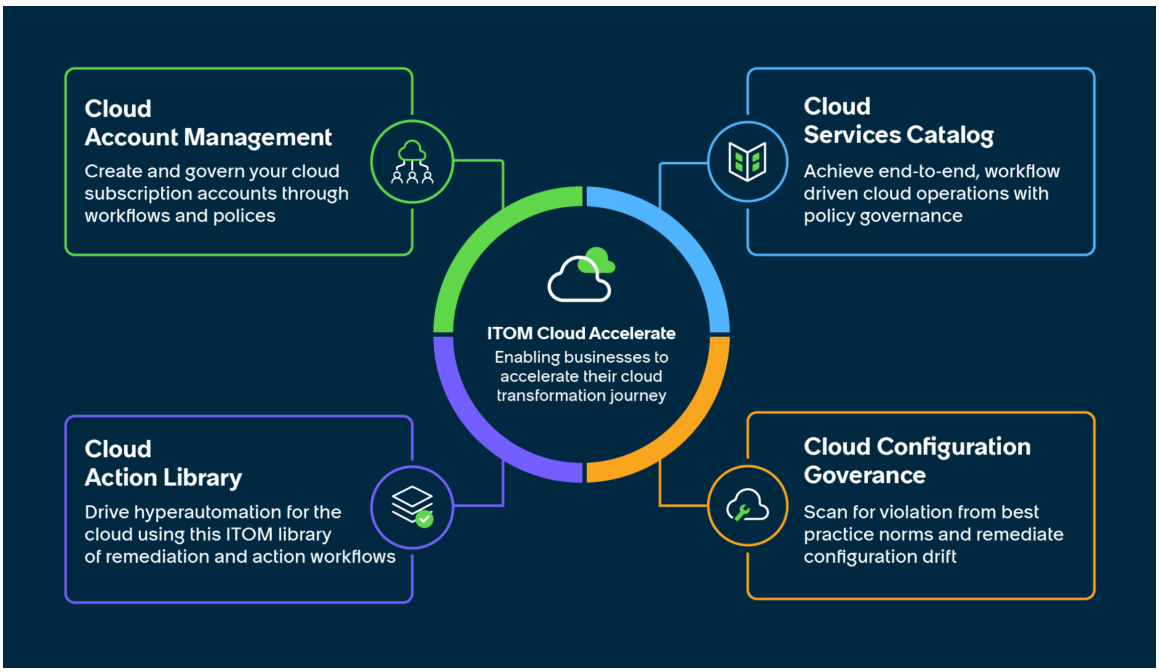

These automated processes empower businesses to quicken their journey towards cloud transformation by automating operations on Day 0, Day 1, and Day 2. It offers the following features:

Cloud Account Management

Use the ServiceNow® Cloud Account Management app for efficiently managing and provisioning of cloud accounts. The app allows you to either obtain manual approvals for account creation—ensuring that the process is closely monitored and verified by relevant stakeholders—or to automate the approval process based on predefined rules and conditions, saving time and reducing manual intervention.

Cloud Services Catalog

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the ServiceNow® Cloud Services Catalog app for simplified, assisted, and featured interface to access cloud resources. You can also publish cloud offerings to a catalog and manage the usage and life cycle of those resources.

Cloud Migration Assessment

Use the ServiceNow® Cloud Migration Assessment application to plan, organize, and track the process of relocating your enterprise IT resources and workloads to cloud platforms. Analyze the IT resources using dashboards that visualize and summarize information on resources by different categories.

Cloud Configuration Governance

Use the ServiceNow® Cloud Configuration Governance application to check the configuration settings of cloud resources in your organization against a set of policies to identify violations. After identifying the violation, use remediation workflows to mitigate them.

Cloud Action Library

Use the ready-to-use actions and subflows of the ServiceNow® Cloud Action Library application to interact with the cloud resources of the organization. ServiceNow® ITOM Cloud Accelerate features, such as Cloud Configuration Governance, use these actions to operate.

ITOM Cloud Accelerate licensing

The ServiceNow AI Platform® uses a licensing method where your organization is billed for using ITOM Cloud Accelerate. The ServiceNow Product Documentation doesn't provide information on prices, packaging, or other details determined by your organization customer contract. For general information about licensing and subscriptions, see ITOM/OT SU Licensing and subscriptions.

Deploying ITOM Cloud Accelerate

ITOM Cloud Accelerate features are available as separate apps on ServiceNow Store.

For licensing purposes, the ITOM Cloud Accelerate license provides the CI count for CIs managed and monitored by Cloud Migration Assessment, Cloud Configuration Governance, or Cloud Action Library. If different ITOM Cloud Accelerate features report the same CIs as their resources, the CI count is manipulated to remove the duplication.

Using guided setup to implement IT Operations Management applications

IT Operations Management Guided Setup provides a sequence of tasks that help you configure IT Operations Management applications on your ServiceNow instance. To open IT Operations Management guided setup, navigate to **Guided Setup** > **ITOM Guided Setup**. For more information about using the guided setup interface, see Using guided setup.

Cloud Account Management

Cloud Account Management simplifies account management by automating tasks like creation and provisioning. Predefined policies enable easy suspension, reactivation, and certification of accounts. Potentially reduce provisioning times from weeks to minutes by automating requests, approvals, and setups.

Cloud Account Management brings together ServiceNow cloud solutions in a unified experience as part of the Cloud Governance Suite.

The Cloud Account Management feature automates the management process for AWS accounts and Azure subscription creation. Key benefits include automating account management,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

streamlining approval workflows, reducing provisioning time, and enhancing security with automated certifications.

Cloud Account Management workflow

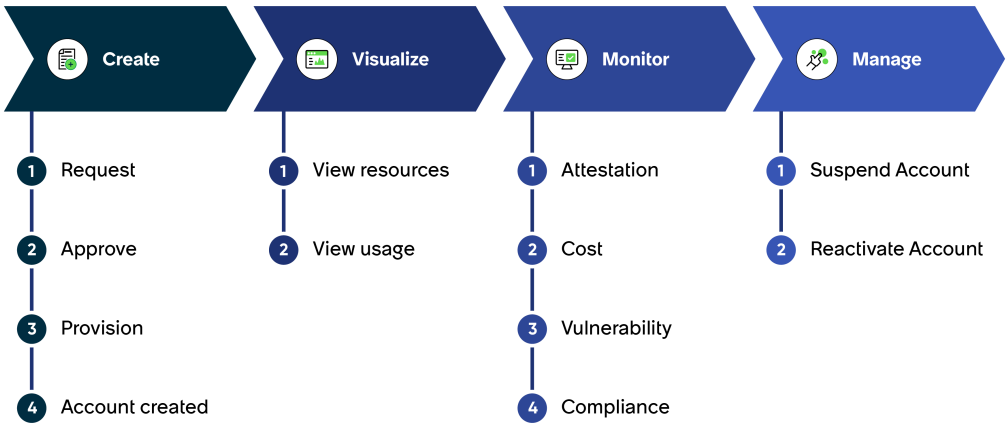

Get Started

Explore

Learn about Cloud Account Management concepts and features.

Configure

Configure environment and user access.

Use

Become familiar with workflow to create and manage cloud accounts.

Reference

Get details about components like fields, tables, and properties.

Exploring Cloud Account Management

The ServiceNow Cloud Account Management in Cloud Workspace application provides a framework to streamline the cloud account creation and management process.

Cloud Account Management overview

The automation capabilities of Cloud Account Management in Cloud Workspace simplify administrative tasks related to account management, such as creation, suspension, reactivation,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

and certification. Data certification confirms the integrity and security of user accounts within the organization's cloud environment.

As a Cloud Account Management admin, you can view the account statistics displayed on the **Home** tab, which include critical severity accounts, high severity accounts, suspended accounts, accounts due for certification, and accounts with undefined budgets.

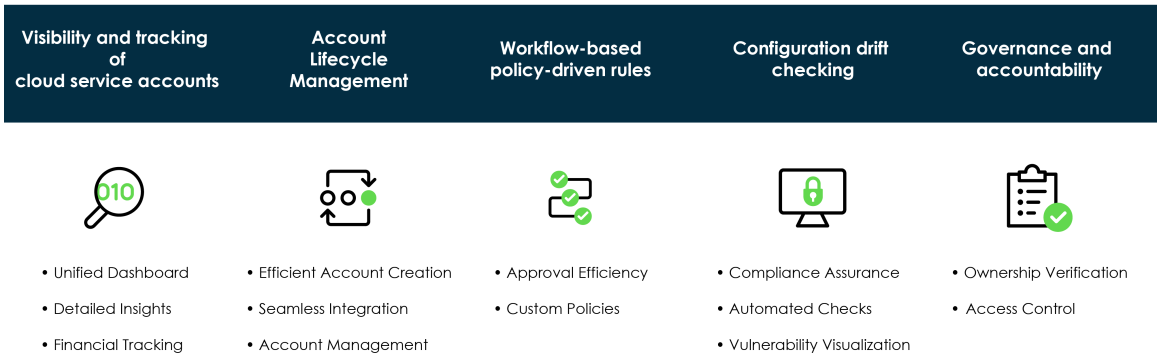

About Cloud Workspace entitlements

Cloud Account Management dynamically adjusts the features available to you based on your entitlements.

- Cloud Account Management entitlement becomes available with the ITOM Cloud Accelerate license.
- The Cloud Governance Suite (CGS) license is a prerequisite to have Cloud Workspace. The CGS license provides the following capabilities:
  - Access the home page
  - Access the asset explorer, which provides an overview of your cloud assets and asset details
- The Cloud Governance Suite with the Cloud Account Management entitlements (CGS + CAM) provides the following additional capabilities:
  - View an overview of your accounts and account details
  - Submit, view, or approve account requests
  - Certify an account
  - View the compliance dashboard
  - Configure Cloud Account Management or view configuration details
  - Create request policies to automate the complete account creation and suspension process

Compliance dashboard in Cloud Workspace

The compliance dashboard consolidates data from ITOM Visibility, Cloud Account Management to provide a unified view of cloud data and key metrics critical for security and compliance reporting.

For more details, see Viewing the compliance dashboard.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud account terminology

Cloud Account Management uses the term "cloud account," but other cloud providers use different terms for similar concepts. The terms are as follows:

- Microsoft Azure Cloud (Azure) refers to subscriptions
- Amazon AWS Cloud (AWS) refers to management and member accounts
- Google Cloud Platform (GCP) refers to projects

Cloud providers offer two main access mechanisms for performing operations:

- Console user access
- Programmatic user access

For Cloud Account Management, these operations are referred to as IAM service accounts. However, other cloud providers use different terminology:

- AWS refers to the IAM user
- Azure refers to service principals
- GCP refers to service accounts

Cloud Account Management user personas

Personas used in Cloud Account Management

User	Description
Requester	Initiates cloud account creation requests and requests for suspension or reactivation of their own accounts.
Approver	Reviews account requests and either approves or denies them.
Admin	Confirms that the Cloud Account Management configurations align with cloud configurations. Customizes the default data certification policy. Onboards accounts created outside the Cloud Account Management application.
Certifier	Approves an account as certified or failed. Acts as a verification entity by confirming the accuracy and integrity of the data.
Asset viewer	Can view all the configuration items (CIs) in Asset Explorer and access the compliance dashboard.
Account manager	Can view all account details and associated assets. Account managers have edit access to accounts with primary ownership and read-only access to those with secondary ownership.

For more information about Cloud Account Management groups and responsibilities, see Cloud Account Management ACL groups, roles, and responsibilities.

Cloud Account Management Benefits

The Cloud Account Management provides several benefits.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Account Management Benefits

Benefit	Feature	Users
Simplifies subscription account creation by defining standardized procedures and user roles and permissions to promote consistency and compliance with security policies.	Request a cloud account Cancel a Cloud account request	Requester
Enables suspending or reactivating accounts, and adding unmanaged accounts. Offers a visualization dashboard to manage accounts and request policies to automate account creation, approvals, and budget checks.	Suspend a Cloud account Reactivate a cloud account Add an unmanaged cloud account Creating configurations About data visualization in Cloud Account Management Review request policies	Admin
Streamlines performing verifications that a person or entity has legitimate ownership or control over an account for security, compliance, and regulatory purposes.	Certify an account	Certifier

Related topics

- About Amazon Web Services API permissions
- Components installed with Cloud Account Management
- About provision modes in Cloud Account Management

About data visualization in Cloud Account Management

The Cloud Account Management overview dashboard gives a clear view of all cloud accounts and their compliance status. Cloud Account Management pulls data from Cloud Configuration Governance to create a simple, visual overview of compliance. As an admin, you can take the necessary actions for the reported vulnerabilities.

Data visualization process

The workflow for creating a scan configuration for data visualization.

1. Select policy sets:
   - Choose a predefined policy set or customize policies to specify what to scan for account vulnerabilities.
   - Cloud Account Management provides default policies for AWS, Azure, and GCP, which you can adjust to meet specific compliance or security requirements.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Select accounts: Identify the accounts to scan by selecting either a management account (which includes all its member accounts) or individual member accounts.

3. Define a scan schedule or execute an on-demand scan: Set up a scan schedule or run an on-demand scan to evaluate vulnerabilities based on the selected policies for the specified accounts.

4. Daily scans:
   - Conduct daily scans to identify vulnerabilities and compliance issues in accounts.
   - These scans generate reports on violations, severity levels, and account aliases, providing a detailed view of account compliance.

5. Track costs:
   - Cloud Account Management integrates with the Cloud Cost Management Workspace to track budgets and usage of each account.
   - Confirm Cloud Cost Management Workspace version 8.1 or above is installed, and billing is configured.

To set up data visualization, see Set up scan configuration for data visualization.

About Amazon Web Services API permissions

Cloud Account Management interacts with Amazon Web Services to create and manage subscription accounts.

Note: You must establish an AWS service account for Cloud Account Management that is separate from the account for Cloud Discovery.

The following API permissions are required to start a new subscription account in AWS:

- budgets: CreateBudgetAction
- budgets: DescribeBudgetAction
- budgets: ModifyBudget
- budgets: ViewBudget
- organizations: AttachPolicy
- organizations: CreateAccount
- organizations: CloseAccount
- organizations: DescribeAccount
- organizations: DescribePolicy
- organizations: DescribeOrganization
- organizations: DescribeOrganizationalUnit
- organizations: DescribeCreateAccountStatus
- organizations: ListRoots
- organizations: ListAccounts
- organizations: ListTagsForResource

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- organizations: ListAWSServiceAccessForOrganization",
- organizations: ListAccounts
- organizations: ListParents
- organizations: ListOrganizationalUnitsForParent
- organizations: MoveAccount
- organizations: TagResource
- iam: GetAccountSummary
- sts: AssumeRole

Note:

For more details on API permissions, download the Cloud Discovery REST API permissions spreadsheet so you can research and grant the user permissions required for running the discovery process.

About provision modes in Cloud Account Management

Cloud Account Management supports both Terraform (Cloud and Enterprise) and cloud native interface provision modes.

Terraform support

Terraform templates are configuration files that specify the infrastructure resources needed for a specific application or environment using the Hashicorp Configuration Language (HCL). Cloud Account Management dynamically retrieves Terraform credentials and templates from the Git repository during the Terraform account creation process. Using the Terraform template, the ServiceNow instance interacts with the cloud organization's IAM (Identity and Access Management) to retrieve the user's identity.

Cloud native interface support

An AWS or Azure admin creates the IAM user with the specified permissions, and the Cloud Account Management admin enters the credentials of the IAM user into the ServiceNow instance. No additional steps are required for users to access the cloud native interface.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Account Management provision modes

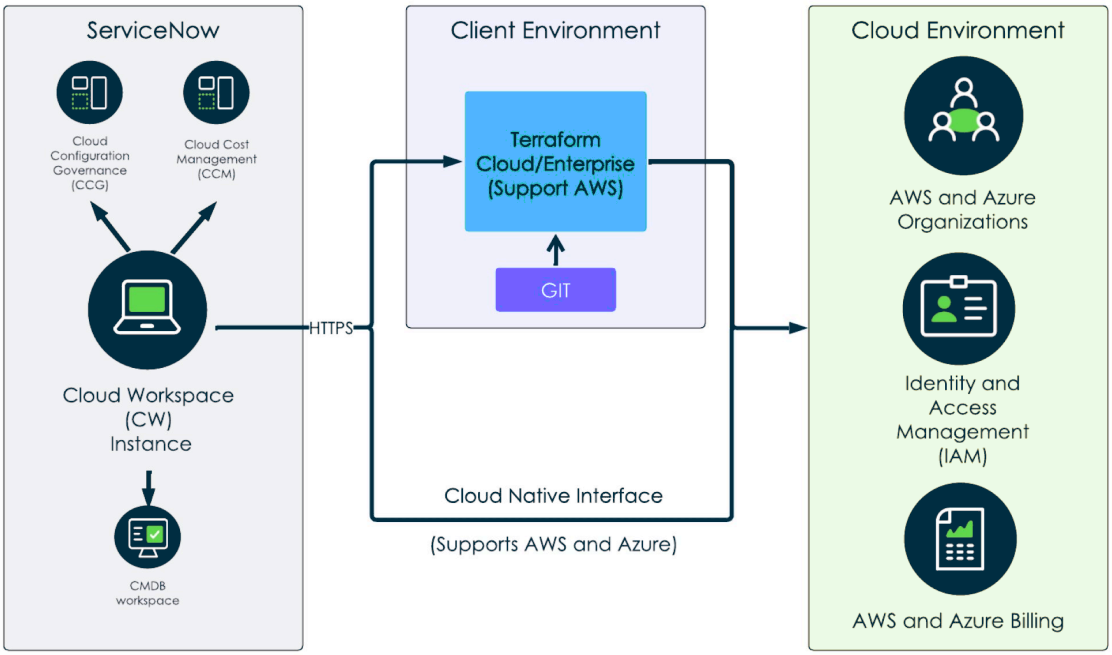

About policies in Cloud Account Management

Cloud Account Management provides policy-driven rules to automate account provisioning, approvals, and security. These policies improve governance, control costs, enhance security, and promote sustainability in cloud operations. Regular reviews promote consistency and efficiency in managing cloud infrastructure.

Cloud Account Management provides policies that you can use for the following purposes:

- Suspending or reactivating AWS accounts. For more details, see Set up suspension of an AWS account using service control policy.
- Auditing accounts and promoting compliance with security and governance standards. For more details, see Review default Cloud Account Management certification policy.
- Automating general approval, budget approval, and configuration assignments through the request policies listed in the following table.

Cloud Account Management request policies

Request policy	For more information
Subscription Account Creation General Approval	Automate General Approval policy process
Subscription Account Creation Budget Approval	Automate Budget Approval policy process
Subscription Account Creation Configuration Assignment	Automate Configuration Assignment policy process

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefits of Cloud Account Management request policies

- Improved governance: Promotes consistent policies across all cloud accounts.
- Cost savings: Helps prevent overspending by managing accounts efficiently.
- Enhanced security: Includes security controls as part of performance and availability.

You can automate the account approval process, budget approval process, and configuration assignment process. To help promote consistency with the cloud subscription account creation process, regularly review and update your policies. For more information, see Review request policies.

Configuring Cloud Account Management

Configuring the Cloud Account Management application involves a set of required tasks, integrations, and setups to be completed before initiating service account creation and provisioning.

The steps involved are different for AWS and Azure. These configurations are listed in the following tables.

Configuration steps for AWS

No.	Task	Task Owner	Impacted Feature	Input	Output	Purpose
1	Set up and verify root email in AWS	Email Admin	Create an account	None	Email Alias	A single AD email across multiple AWS accounts streamlines management, enhances security, and boosts collaboration.
2	Set up an IAM account for a ServiceNow user in AWS	AWS Admin	All features	None	AWS Access Key AWS Secret Key	A centralized IAM user in the management account efficiently manages multiple AWS accounts via a CloudFormation template.
3	Set up suspension of an AWS account using service control policy	AWS Admin	Account suspension	None	SCP Policy ID	Adding an account number to the AWS organization's Service Control Policy blocks the creation of new resources and helps prevent overspending, while existing resources remain unaffected.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration steps for AWS (continued)

No.	Task	Task Owner	Impacted Feature	Input	Output	Purpose
4	 Setting up Terraform and GitHub **Note:** This step isn't required if cloud native interface is used for provisioning the account.	Terraform Admin DevOps Admin	Create an account for AWS	AWS Access Key AWS Secret Key	Terraform API Key Token Terraform OAuth Token ID Terraform Org VCS Identifier Terraform URL	GitHub templates deployed to Terraform Cloud or Enterprise streamline account creation and promote consistent configurations.
5	 Create a cloud native interface account configuration **Note:** This step isn't required if Terraform is used for provisioning the account.	ServiceNow AI Platform Admin	Create an account for Cloud Native Interface	None	None	No additional procedure is needed for cloud native interface. You can proceed to Install Cloud Workspace.
6	 Populate Cloud Account Management dependent tables	ServiceNow AI Platform Admin	All features	Discovery Credentials	Cloud Organization AWS Org Unit Cloud Service Account	Cloud Discovery or Service Graph Connector for AWS to import cloud organization structures and create subscription accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration steps for AWS (continued)

No.	Task	Task Owner	Impacted Feature	Input	Output	Purpose
	1. Create a service account 2. Define and run discovery schedule					
7	 Provisioning modes for Cloud Account Management in Cloud Workspace • Set up Terraform API key in ServiceNow • Set up AWS API configuration information in ServiceNow	ServiceNow AI Platform Admin	Create an account	Email Alias Terraform Cloud API Key Token Terraform Cloud OAuth Token ID Terraform Cloud Org VCS Identifier	None	Streamline provisioning with flexible modes like Terraform and cloud native interface, centralizes management, enhances security, optimizes costs, and improves governance.
8	 Set up scan configuration for data visualization	ServiceNow AI Platform Admin	Visualization	None	None	Data visualization requires scanning all account violations based on the policy set.
9	 Add members to the group	ServiceNow AI Platform Admin	All features	None	None	Confirm that members are assigned to the correct group for them to perform the account request, approval, provisioning, and certification process. Assign members to the correct groups for proper permissions, helping

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration steps for AWS (continued)

No.	Task	Task Owner	Impacted Feature	Input	Output	Purpose
						prevent unauthorized access and promoting security.
10	Creating configurations	Cloud Workspace Admin	All features	None	None	Configure before creating, suspending, or scanning accounts.
11	Review default Cloud Account Management certification policy	Cloud Workspace Admin	All features	None	None	Certify all cloud service accounts either created or boarded after discovery can be certified. The admin can also customize the default policy or create a policy. The default policy helps to certify all the available cloud service accounts once every 90 days.
12	Review request policies	Cloud Workspace Admin	All features	None	None	Request Policies are rules that govern the behavior of request workflows by applying data checks and conditions. Regularly reviewing and updating these policies confirms that your cloud account request process remains consistent with the procedure for creating cloud subscription accounts.

Configuration steps for Azure

No.	Name	Task Owner	Impacted feature	Input	Output	Purpose
1	Configure account	Azure Admin	Account suspension	None	OAuth Client ID OAuth Client Secret	Configure the permission and assign the permission to a user who suspends or reactivates the account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration steps for Azure (continued)

No.	Name	Task Owner	Impacted feature	Input	Output	Purpose
	suspension in Azure				Tenant ID	
2	Set up Azure Service Principal credential	ServiceNow AI Platform Admin	All features	OAuth Client ID OAuth Client Secret Tenant ID	None	Provide the Azure credentials obtained from your Azure administrator. These credentials are used to create a suspension profile and enables you to suspend temporarily or terminate Azure accounts as needed.
3	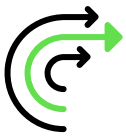 Create a cloud native interface account configuration **Note:** This step isn't required if Terraform is used for provisioning the account.	ServiceNow AI Platform Admin	Create an account for Cloud Native Interface	None	None	No additional procedure is needed for cloud native interface. You can proceed to Install Cloud Workspace.
4	Populate Cloud Account Management dependent tables	ServiceNow AI Platform Admin	All features	Discovery Credentials	Cloud Organization Azure Management Group Cloud Service Account	Cloud Discovery import cloud organization structures and create subscription accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration steps for Azure (continued)

No.	Name	Task Owner	Impacted feature	Input	Output	Purpose
	1. Create a service account 2. Define and run discovery schedule					
5	 Set up scan configuration for data visualization	ServiceNow AI Platform Admin	Visualization	None	None	Data visualization requires scanning all account violations based on the policy set.
6	 Add members to the group	ServiceNow AI Platform Admin	All features	None	None	Confirm that members are assigned to the correct group for them to perform the account request, approval, provisioning, and certification process. Assign members to the correct groups for proper permissions, helping prevent unauthorized access and promoting security.
7	 Creating configurations	Cloud Workspace Admin	All features	None	None	Configure before creating, suspending, or scanning accounts.
8	 Review default Cloud Account Management certification policy	Cloud Workspace Admin	All features	None	None	Certify all cloud service accounts either created or boarded after discovery can be certified. The admin can also customize the default policy or create a policy. The default policy helps to certify all the available cloud service accounts once every 90 days.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuring cloud providers

The topics in this section cover various cloud configurations required to set up and work with Cloud Account Management in Cloud Workspace.

Setting up AWS cloud

This section covers all the topics required to set up AWS cloud. The topics are arranged based on configuration priority.

Setting up a root email

The app creates accounts with a unique request ID and root email, promoting distinct emails for each account and simplifying management.

Shared root email account benefits:

- Reduces administrative burden: No need for individual email addresses for AWS accounts.
- Centralizes notifications: All notifications from the AWS account delivered to a single mailbox.
- Improves collaboration: Shared email access enables the DevOps team to view and collaborate on tasks.

Note: Make sure you note the following points when provisioning a new subscription account:

1. The app dynamically appends a request ID (for example, CWSAREQ0000002) to a predefined root email address (`aws-ccoe@mycompany.com`). This process creates a unique email address for each account in a structured format (`aws-ccoe+CWSAREQ0000002@mycompany.com`).
2. After the account is provisioned, AWS sends emails in this format that are received at a central email account `aws-ccoe@mycompany.com`.

Set up and verify root email in AWS

Create a unique email address to manage several accounts. Having multiple accounts and emails can be a significant challenge for an AWS admin. A single AD email for multiple AWS accounts simplifies management.

Before you begin

Role required: Email / AD admin

About this task

Procedure

1. Create an email address similar to `aws-ccoe@mycompany.com`.
2. Send a test email from your email to `aws-ccoe+test@mycompany.com`.
3. Confirm that the test email is received in the `aws-ccoe+test@mycompany.com` mailbox.

What to do next

Share the confirmed email address with the Provisioner or ServiceNow AI Platform admin for configuration purposes.

Once you receive an account creation confirmation email on `aws-ccoe@mycompany.com`, set the root credentials manually based on your company's security standards.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up an IAM account for a ServiceNow user in AWS

AWSCreate an AWS CloudFormation Identity Access Manager (IAM) user in the management account using the CloudFormation template (CFT). This allows you to manage multiple AWS accounts with a centralized IAM user.

Before you begin

- Role required: AWS admin
- The ServiceNow IAM user account requires specific AWS IAM permissions to perform basic operations. To learn more about the permissions, see About Amazon Web Services API permissions.

Procedure

1. Log in to the AWS Management console and enter `CloudFormation` in the search bar and then select it.
2. On the CloudFormation console, select **Create Stack**.

   A stack is a collection of AWS resources that you can manage as a single unit.
3. Copy the CloudFormation template (CFT) code into a file and save the file.

   The file extension should be `.yml`. For example, `CreateServiceNowUser.yml`.

```
AWSTemplateFormatVersion: '2010-09-09'
Description: This script is executed in Management Account
 where the  ServiceNow user is created.

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: User Credentials
        Parameters:
          - SNUserName
    ParameterLabels:
      SNUserName:
        default: User Name

Parameters:
  SNUserName:
    Type: String
    Description: User name for CW Service Account user
    MinLength: '6'
    MaxLength: '15'
    ConstraintDescription: The username must be between 6 and
 15 characters

Resources:
  SnowCWMemberAccountAccessGroup:
    Type: 'AWS::IAM::Group'
    Properties:
      GroupName: SnowCWMemberAccountAccessGroup

  SnowCWAccountAccessPolicy:
    Type: 'AWS::IAM::Policy'
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    Properties:
      PolicyName: SnowCWAccountAccessPolicy
      PolicyDocument:
        Statement:
          - Sid: ServiceNowCWUserAccess
            Effect: Allow
            Action:
              - 'organizations:ListRoots'
              - 'organizations:ListTagsForResource'
              - 'organizations:DescribeAccount'
              - 'organizations:CreateAccount'
              -
'organizations:ListAWSServiceAccessForOrganization'
              - 'organizations:TagResource'
              - 'organizations:ListAccounts'
              - 'organizations:ListRoots'
              - 'organizations:ListTagsForResource'
              - 'organizations:DescribeAccount'
              - 'organizations:CreateAccount'
              -
'organizations:ListAWSServiceAccessForOrganization'
              - 'organizations:TagResource'
              - 'organizations:ListAccounts'
              - 'organizations:CloseAccount'
              - 'organizations:DescribeOrganization'
              - 'organizations:DescribeOrganizationalUnit'
              - 'organizations:ListParents'
              -
'organizations:ListOrganizationalUnitsForParent'
              - 'organizations:MoveAccount'
              - 'organizations:DescribeCreateAccountStatus'
              - 'organizations:AttachPolicy'
              - 'organizations:DescribePolicy'
              - 'organizations:ListAccounts'
              - 'budgets:CreateBudgetAction'
              - 'budgets:ModifyBudget'
              - 'budgets:ViewBudget'
              - 'budgets:ListTagsForResource'
              - 'iam:GetAccountSummary'
              - 'iam:GetAccountPasswordPolicy'
              - 'budgets:DescribeBudgetAction'
              - 'iam:ListAccountAliases'
              - 'sts:AssumeRole'
              - 'iam:GetRole'
              - 'iam:ListAccountAliases'
              - 'iam:GetAccountPasswordPolicy'
              - 'organizations:DetachPolicy'
              - 'iam:GetAccountAuthorizationDetails'
            Resource: '*'
      Groups:
        - !Ref SnowCWMemberAccountAccessGroup

  CreateServicenowUser:
    Type: 'AWS::IAM::User'
    Properties:
      Path: /
      UserName: !Ref SNUserName
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
  AddSnowUserToSnowCWMemberAccountAccessGroup:
    Type: 'AWS::IAM::UserToGroupAddition'
    Properties:
      GroupName: SnowCWMemberAccountAccessGroup
      Users:
        - !Ref SNUserName
    DependsOn: CreateServicenowUser

Outputs:
  ServiceNowUserARN:
    Description: ARN of ServiceNow user
    Value: !GetAtt
      - CreateServicenowUser
      - Arn

  ServiceNowUser:
    Description: ServiceNow user
    Value: !Ref SNUserName
```

4. On the **Create Stack** page, select **Choose an existing template** > **Upload a template file** > **Choose file** to choose a template file from your local computer.
5. Select **Next** to continue and to validate the template.
6. On the **Specify stack details** page, enter a `CloudFormation stack name` in the **Stack name** field.

   The stack name is an identifier that helps you find a particular stack from a list of stacks. A stack name can contain only alphanumeric characters (case-sensitive) and hyphens. It must start with an alphabetic character and can't be longer than 128 characters.
7. In the **User Name** field, enter a Cloud Account Management service account user name.
8. Select **Next** to continue.
9. On the **Configure stack options** page, select the **I acknowledge that AWS CloudFormation might create IAM resources with custom names** check box and select **Next**.
10. On the **Review and create** page, select **Submit** to launch your stack.
    The IAM account gets created.
11. To add security credentials:

    a. Select **Resources**.

    b. Select the user name link.

    c. In the Access Keys section, select **Create access key**.

    d. On the **Access key best practices & alternatives** page, choose the **Third-party service** option.

       You plan to use this access key to enable access for a third-party application or service that monitors or manages your AWS resources.

    e. Select the **I understand the above recommendation and want to proceed to create an access key** check box and select **Next**.

    f. Select **Create access key**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

g. In the **Review access keys** page, download the access key file and select **Done**.

> **Note:** Share the access key and secret key with the Terraform admin and ServiceNow AI Platform admin.

What to do next

Set up suspension of an AWS account using service control policy

Set up an Identity Access Manager in AWS GovCloud

Create AWS CloudFormation Identity Access Manager (IAM) resources for three AWS account types in the management account using the CloudFormation templates (CFT) to integrate ServiceNow Cloud Account management (CAM) with AWS GovCloud.

Before you begin

Verify that the target user account has the required AWS IAM permissions. For more information, see About Amazon Web Services API permissions.

> **Note:** Always follow your organization's security policies when implementing IAM permissions.

Role required: AWS admin

About this task

AWS GovCloud accounts are always linked to commercial accounts, which means IAM resources must be created across three AWS account types:

- AWS Commercial Management Account: The standard AWS account that links to GovCloud.
- AWS GovCloud Root Management Account: The main account in the GovCloud environment.
- AWS GovCloud Organization Management Account: The account for managing organization resources in the GovCloud environment.

Procedure

1. Copy each CloudFormation template (CFT) code into different files with a `.yml` extension and save the files.
2. Log in to the AWS Management console using your root user credentials.
3. Follow the steps to execute each script in sequence.

   a. **AWS Commercial Management Account** CFT script

      i. Log in to the Commercial AWS console.

      ii. Access the CloudFormation console by entering `CloudFormation` in the search bar and selecting it.

      iii. On the CloudFormation console, select **Create Stack** > **Choose an existing template** > **Upload a template file** > **Choose file** to choose a template file from your local computer.

      A stack is a collection of AWS resources that you can manage as a single unit.

      iv. Use the saved Commercial CFT script and set the `SNUserName` parameter.

      v. Record the output parameters: `SNUserARN` and `SNUser`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
AWSTemplateFormatVersion: '2010-09-09'
Description: This script creates a ServiceNow user in the AWS
 Commercial Management Account for Cloud Account Management
 (CAM) integration.

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: User Credentials
        Parameters:
          - SNUserName
    ParameterLabels:
      SNUserName:
        default: User Name in the AWS Commercial environment

Parameters:
  SNUserName:
    Type: String
    Description: User name for ServiceNow Cloud Workspace
 application in AWS Commercial Environment
    MinLength: '6'
    MaxLength: '15'
    ConstraintDescription: The username must be between 6 and
 15 characters

Resources:
  # AWS Commercial Resources
  SnowCWMemberAccountAccessGroup:
    Type: 'AWS::IAM::Group'
    Properties:
      GroupName: SnowCWMemberAccountAccessGroup
    Description: IAM Group for ServiceNow Cloud Workspace
 users with permissions to manage AWS Organizations

  SnowCWAccountAccessPolicy:
    Type: 'AWS::IAM::Policy'
    Properties:
      PolicyName: SnowCWAccountAccessPolicy
      PolicyDocument:
        Version: '2012-10-17'
        Statement:
          # Create Account Workflow Permissions
          - Sid: ServiceNowCreateAccountAccess
            Effect: Allow
            Action:
              - 'budgets:CreateBudgetAction'
              - 'budgets:DescribeBudgetAction'
              - 'organizations:CreateGovCloudAccount'
              - 'organizations:DescribeCreateAccountStatus'
              - 'organizations:MoveAccount'
              - 'organizations:TagResource'
              - 'organizations:ListRoots'
            Resource: '*'
          # Account Management Workflow Permissions
          - Sid: ServiceNowAccountManagementAccess
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
            Effect: Allow
            Action:
              - 'budgets:ListTagsForResource'
              - 'budgets:ModifyBudget'
              - 'budgets:ViewBudget'
              - 'iam:GetAccountPasswordPolicy'
              - 'iam:GetAccountSummary'
              - 'iam:GetRole'
              - 'iam:ListAccountAliases'
              - 'organizations:AttachPolicy'
              - 'organizations:DescribePolicy'
              - 'organizations:DetachPolicy'
              - 'sts:AssumeRole'
            Resource: '*'
          # Discovery Workflow Permissions
          - Sid: ServiceNowDiscoveryAccess
            Effect: Allow
            Action:
              - 'organizations:DescribeAccount'
              - 'organizations:DescribeOrganization'
              - 'organizations:DescribeOrganizationalUnit'
              - 'organizations:ListAccounts'
              -
 'organizations:ListAWSServiceAccessForOrganization'
              -
 'organizations:ListOrganizationalUnitsForParent'
              - 'organizations:ListParents'
              - 'organizations:ListRoots'
              - 'organizations:ListTagsForResource'
            Resource: '*'
      Groups:
        - !Ref SnowCWMemberAccountAccessGroup
    Description: Policy that grants permissions required for
 ServiceNow CAM in the AWS Commercial environment

  CreateServicenowUser:
    Type: 'AWS::IAM::User'
    Properties:
      Path: /
      UserName: !Ref SNUserName
      Tags:
        - Key: Purpose
          Value: ServiceNowCAMIntegration
    Description: IAM user for ServiceNow CAM integration with
 AWS Commercial environment

  AddSnowUserToSnowCWMemberAccountAccessGroup:
    Type: 'AWS::IAM::UserToGroupAddition'
    Properties:
      GroupName: SnowCWMemberAccountAccessGroup
      Users:
        - !Ref SNUserName
    DependsOn: CreateServicenowUser
    Description: Adds the ServiceNow user to the access group

Outputs:
  SNUserARN:
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    Description: ARN of ServiceNow user in AWS Commercial
environment
    Value: !GetAtt CreateServicenowUser.Arn

  SNUser:
    Description: Username of the created ServiceNow user in
AWS Commercial environment
    Value: !Ref SNUserName
```

b. AWS GovCloud Root Management Account CFT script

i. Log in to the GovCloud AWS console.

ii. Access the CloudFormation console by entering `CloudFormation` in the search bar and selecting it.

iii. On the CloudFormation console, select **Create Stack** > **Choose an existing template** > **Upload a template file** > **Choose file** to choose a template file from your local computer.

iv. Use the saved GovCloud root CFT script and set the `SNGovUserName` parameter.

v. Record the output parameters: `SNGovUserARN` and `SNGovUser`.

Note:

For simplicity, this template uses a wildcard (*) in the Resource ARN for the SnowCWGovCloudUserPolicy, enabling the ServiceNow user to assume the role in any GovCloud account. While this works functionally, it grants more permissions than necessary.

For production environments:

- Add a parameter for the organization account ID: `OrgAccountId: Type: String`
- Update the Resource field:

```
Resource: !Sub

  'arn:aws-us-
gov:iam::${OrgAccountId}:role/ServiceNowGovCrossAccountR
ole'
```

- For working with multiple accounts, see CloudFormation Parameters and Using CommaDelimitedList
- Review AWS IAM Best Practices for more guidance on securing IAM policies

```
AWSTemplateFormatVersion: '2010-09-09'
Description: This script creates a ServiceNow user in the AWS
 GovCloud Management Account for Cloud Account Management
 (CAM) integration.

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: User Credentials
        Parameters:
          - SNGovUserName
    ParameterLabels:
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      SNGovUserName:
        default: AWS User Name in the GovCloud environment

Parameters:
  SNGovUserName:
    Type: String
    Description: User name for ServiceNow Cloud Workspace
 application in AWS GovCloud Environment
    MinLength: '6'
    MaxLength: '15'
    ConstraintDescription: The username must be between 6 and
 15 characters

Resources:
  # AWS GovCloud Resources
  SnowCWGovCloudAccessGroup:
    Type: 'AWS::IAM::Group'
    Properties:
      GroupName: SnowCWGovCloudAccessGroup
    Description: IAM Group for ServiceNow Cloud Workspace
 users in AWS GovCloud environment

  CreateGovCloudServicenowUser:
    Type: 'AWS::IAM::User'
    Properties:
      Path: /
      UserName: !Ref SNGovUserName
      Tags:
        - Key: Purpose
          Value: ServiceNowCAMIntegration
    Description: IAM user for ServiceNow CAM integration with
 AWS GovCloud environment

  SnowCWGovCloudUserPolicy:
    Type: 'AWS::IAM::Policy'
    Properties:
      PolicyName: SnowCWGovCloudUserPolicy
      PolicyDocument:
        Version: '2012-10-17'
        Statement:
          - Effect: Allow
            Action: 'sts:AssumeRole'
            Resource:
'arn:aws-us-gov:iam::*:role/ServiceNowGovCrossAccountRole'
          - Effect: Allow
            Action: 'sts:AssumeRole'
            Resource:
'arn:aws-us-gov:iam::*:role/OrganizationAccountAccessRole'
      Groups:
        - !Ref SnowCWGovCloudAccessGroup
    Description: Policy that grants basic permissions and
 ability to assume the cross-account role

  AddSnowUserToSnowCWGovCloudAccessGroup:
    Type: 'AWS::IAM::UserToGroupAddition'
    Properties:
      GroupName: SnowCWGovCloudAccessGroup
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      Users:
        - !Ref SNGovUserName
    DependsOn:
      - CreateGovCloudServicenowUser
      - SnowCWGovCloudAccessGroup
    Description: Adds the ServiceNow GovCloud user to the
 access group

Outputs:
  SNGovUserARN:
    Description: ARN of ServiceNow user in AWS GovCloud
 environment
    Value: !GetAtt CreateGovCloudServicenowUser.Arn

  SNGovUser:
    Description: Username of the created ServiceNow user in
 AWS GovCloud environment
    Value: !Ref SNGovUserName
```

c. **AWS GovCloud Organization Management Account** CFT script

i. Log in to the GovCloud Organization AWS console.

ii. Access the CloudFormation console by entering and selecting `CloudFormation` in the search bar.

iii. On the CloudFormation console, select **Create Stack** > **Choose an existing template** > **Upload a template file** > **Choose file** to choose a template file from your local computer.

iv. Use the saved GovCloud Organization CFT script and set the `SNGovUserARN` parameter to the output recorded from the **AWS GovCloud Root Management Account** CFT script.

v. Record the output parameter: `RoleARN`.

```
AWSTemplateFormatVersion: '2010-09-09'
Description: This template creates the
 ServiceNowGovCrossAccountRole in the AWS Organization
 Management Account for ServiceNow CAM integration.

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: Role Configuration
        Parameters:
          - SNGovUserARN
    ParameterLabels:
      SNGovUserARN:
        default: ServiceNow GovCloud User ARN

Parameters:
  SNGovUserARN:
    Type: String
    Description: ARN of the ServiceNow GovCloud user that
 will assume this role
    AllowedPattern:
 '^arn:aws(-[a-z0-9-]+)?:iam::[0-9]{12}:user/.*$'
    ConstraintDescription: Must be a valid AWS IAM user ARN
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Resources:
  ServiceNowGovCrossAccountRole:
    Type: 'AWS::IAM::Role'
    Properties:
      RoleName: ServiceNowGovCrossAccountRole
      Description: Role used by ServiceNow CAM for
 cross-account operations in AWS GovCloud
      AssumeRolePolicyDocument:
        Version: '2012-10-17'
        Statement:
          - Effect: Allow
            Principal:
              AWS: !Ref SNGovUserARN
            Action: 'sts:AssumeRole'
      Policies:
        - PolicyName: ServiceNowGovCloudPolicy
          PolicyDocument:
            Version: '2012-10-17'
            Statement:
              # Create Account Workflow Permissions
              - Sid: ServiceNowGovCreateAccountAccess
                Effect: Allow
                Action:
                  -
 'organizations:InviteAccountToOrganization'
                  - 'organizations:ListRoots'
                  - 'organizations:MoveAccount'
                  - 'organizations:TagResource'
                Resource: '*'
              # Account Management Workflow Permissions
              - Sid: ServiceNowGovAccountManagementAccess
                Effect: Allow
                Action:
                  - 'iam:GetAccountAuthorizationDetails'
                  - 'iam:GetAccountPasswordPolicy'
                  - 'iam:GetAccountSummary'
                  - 'iam:GetRole'
                  - 'iam:ListAccountAliases'
                  - 'organizations:AttachPolicy'
                  - 'organizations:DescribePolicy'
                  - 'organizations:DetachPolicy'
                Resource: '*'
              # Discovery Workflow Permissions
              - Sid: ServiceNowGovDiscoveryAccess
                Effect: Allow
                Action:
                  - 'ec2:DescribeRegions'
                  - 'organizations:DescribeAccount'
                  - 'organizations:DescribeOrganization'
                  -
 'organizations:DescribeOrganizationalUnit'
                  - 'organizations:ListAccounts'
                  -
 'organizations:ListAWSServiceAccessForOrganization'
                  -
 'organizations:ListOrganizationalUnitsForParent'
                  - 'organizations:ListParents'
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                  - 'organizations:ListRoots'
                  - 'organizations:ListTagsForResource'
                Resource: '*'

Outputs:
  RoleARN:
    Description: ARN of the ServiceNowGovCrossAccountRole
    Value: !GetAtt ServiceNowGovCrossAccountRole.Arn
```

4. Select **Next** to continue and validate the template.
5. On the **Specify stack details** page, enter a name for the template in the **Stack name** field.

   A stack name:

   - Can contain only case-sensitive alphanumeric characters and hyphens
   - Must start with an alphabetic character
   - Cannot be longer than 128 characters
6. In the **User Name** field, enter a Cloud Account Management service account user name.
7. Select **Next**.
8. On the **Configure stack options** page, select the **I acknowledge that AWS CloudFormation might create IAM resources with custom names** check box and select **Next**.
9. On the **Review and create** page, select **Submit** to launch your stack.
   The IAM account gets created.
10. To add security credentials for the two users created:

    a. Select **Resources**.

    b. Select the user name link.

    c. In the Access Keys section, select **Create access key**.

    d. On the **Access key best practices & alternatives** page, choose the **Third-party service** option.

    You can use this access key to enable access for a third-party application or service that monitors or manages your AWS resources.

    e. Select the **I understand the above recommendation and want to proceed to create an access key** check box and select **Next**.

    f. Select **Create access key**.

    g. In the **Review access keys** page, download the access key file and select **Done**.

    Note: Share the access key and secret key with the Terraform admin and ServiceNow AI Platform admin.

Configuring Service Control Policy in AWS

The AWS admin configures a Service Control Policy (SCP) and shares its ID with the ServiceNow AI Platform admin for Cloud Account Management setup. Cloud Account Management enforces the SCP via API calls to block resource creation in the account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The process involved in creating and using a Service Control Policy:

1. The AWS admin sets up a Service Control Policy (SCP) in the management account, using a list of resources from the CFT script. The admin can customize this list by adding more resource types.
2. The AWS admin shares the SCP ID with the ServiceNow admin, who uses it for registration during the Cloud Account Management configuration process.
3. The Cloud Account Management app submits a suspension request, triggering an API call to the AWS Organizations Attach Policy API. This API enforces the SCP to block users from creating resources in that account.

Copy the following content to a file and save the file as CloudFormation template (extension as *.cft):

```
AWSTemplateFormatVersion: 2010-09-09
Description: SCP policy for ServiceNow Cloud Workspace to
 restrict creation of new resources.
Resources:
  PolicyTestTemplate:
    Type: AWS::Organizations::Policy
    Properties:
      Type: SERVICE_CONTROL_POLICY
      Name: CAM_SCP_SuspendAccount_Policy
      Content:
        Version: 2012-10-17
        Statement:
          - Sid: CAMSCPSuspendAccountPolicy
            Effect: Deny
            Action:
              - 'ec2:RunInstances'
              - 'ec2:CreateVolume'
              - 'ec2:CreateSnapshot'
              - 'ec2:CreateImage'
              - 's3:CreateBucket'
              - 'iam:CreateUser'
              - 'iam:CreateRole'
              - 'iam:CreatePolicy'
              - 'dynamodb:CreateTable'
              - 'sqs:CreateQueue'
              - 'sns:CreateTopic'
              - 'lambda:CreateFunction'
              - 'ec2:CreateVpc'
              - 'ec2:CreateSubnet'
              - 'ec2:CreateInternetGateway'
              - 'ec2:CreateRoute'
              - 'rds:CreateDBInstance'
              - 'redshift:CreateCluster'
            Resource: '*'
```

Set up suspension of an AWS account using service control policy

Set up a restriction on cloud account creation. For example, when the account owner isn't in the organization or if there are budget constraints. Using Cloud Account Management, admin adds the account number to the AWS organization's service control policy, promoting existing accounts to continue to function normally while blocking new account creation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Verify you have saved the CloudFormation template. For more details, see Configuring Service Control Policy in AWS.

Role required: AWS admin

Procedure

1. Log in to the AWS Management console.
2. Enter `CloudFormation` in the search bar and then select it.
3. On the CloudFormation console, select **Create Stack**.
4. On the **Create Stack** page, select **Choose an existing template** > **Upload a template file** > **Choose file** to choose a template file from your local computer.
5. Select **Next** to continue and to validate the template.
6. On the **Specify stack details** page, enter a `CloudFormation stack name` in the **Stack name** field.

   The stack name is an identifier that helps you find a particular stack from a list of stacks. A stack name can contain only alphanumeric characters (case-sensitive) and hyphens. It must start with an alphabetic character and can't be longer than 128 characters.
7. Select **Next** and then select **Submit**.

Result

Once the CloudFormation Template (CFT) creates the policy, an Amazon Resource Name (ARN) is assigned. An ARN typically follows the format:

```
arn:aws:organizations::1234567890:policy/o-99t3h155el/service_control_policy/p-328wg3yb
```

To view the ARN, go to **AWS Organization** > **Policies** > **Service control policies** > **CAM_SCP_SupsendAccount_policy**.

The key element that you must provide to the ServiceNow AI Platform admin is the policy element, which in this example is: p-328wg3yb.

What to do next

Provisioning modes for Cloud Account Management in Cloud Workspace

Provisioning modes for Cloud Account Management in Cloud Workspace

Cloud Account Management in Cloud Workspace offers flexible provision modes such as Terraform and cloud native interfaces while centralizing management, enhancing security, optimizing costs, and improving governance.

Overview on provision modes

Provisioning modes for creating AWS accounts are:

- Terraform- Setting up Terraform and GitHub.
- Cloud native interface- There's no additional procedure required for cloud native interface. You can proceed to Install Cloud Workspace.

> **Note:** To know more on API permissions, see About Amazon Web Services API permissions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Account Management provision modes

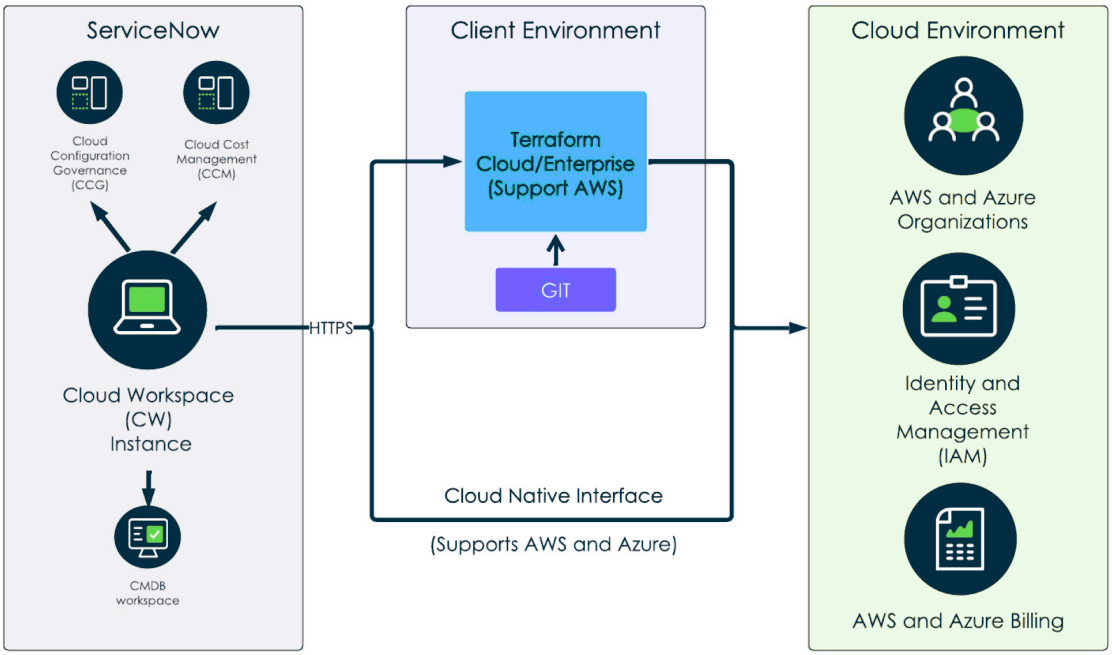

Setting up Terraform and GitHub

Simplify cloud account provisioning using Terraform. The workflow provides the basic steps involved in setting up Terraform and GitHub.

Publish Terraform templates

Publish the Terraform template defined by the application in GitHub to enable version control, collaboration, and centralized storage.

Before you begin

Role required: Terraform admin

Procedure

1. Create a GitHub repository dedicated to your Terraform templates.
2. Save the following codes in a separate file.
   - File name: awsaccount.tf

```
terraform {
  required_providers {
    aws = {
      source  = "hashicorp/aws"
      version = "~>5.0"
    }
  }
}

provider "aws" {
  region = var.region
}

resource "aws_organizations_account" "account" {
  email = var.root_email
  name = var.account_name
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
  tags = var.tags
  parent_id = var.parent_id
  close_on_deletion=var.close_delete
  create_govcloud=var.gov_cloud
}

resource "aws_budgets_budget" "cost" {
  name  = join("-", ["SN-CAM-Monthly-Budget",
 aws_organizations_account.account.id])
  count = var.monthly_budget > 0 ? 1 : 0
  budget_type  = "COST"
  limit_amount = var.monthly_budget
  limit_unit   = "USD"
  time_unit    = "MONTHLY"

  cost_filter {
    name = "LinkedAccount"
    values = [
      aws_organizations_account.account.id
    ]
  }

  notification {
    comparison_operator        = "GREATER_THAN"
    threshold                  = 100
    threshold_type             = "PERCENTAGE"
    notification_type          = "FORECASTED"
    subscriber_email_addresses = [var.root_email,
 var.notification_email]
  }

  notification {
    comparison_operator        = "GREATER_THAN"
    threshold                  = 85
    threshold_type             = "PERCENTAGE"
    notification_type          = "ACTUAL"
    subscriber_email_addresses = [var.root_email,
 var.notification_email]
  }

  notification {
    comparison_operator        = "GREATER_THAN"
    threshold                  = 100
    threshold_type             = "PERCENTAGE"
    notification_type          = "ACTUAL"
    subscriber_email_addresses = [var.root_email,
 var.notification_email]
  }
}
```

- File name: variables.tf

```
variable "region" {
  type = string
  default = "us-east-1"
}
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
variable "account_name" {
  type = string
  description = "(Required) Account Name"
}

variable "root_email" {
  type = string
  description = "(Required) Account Email"
}

variable tags{
  type = map(string)
  description = "(Required) Tags for the resource"
}

variable "close_delete" {
  type = bool
  description = "Close Account on deletion"
  default = true
}

variable "gov_cloud" {
  type = bool
  description = "Gov Cloud Account"
  default = false
}

variable "parent_id" {
  type = string
  description = "(Required) Account Parent Organizational
 Unit"
}

variable monthly_budget {
  type = number
  default = 0
}

variable notification_email {
  type = string
  description = "Additional email where Budget Notifications
 are to be sent"
}
```

3. Push the Terraform files to the remote repository.
4. Make a note of the location or URL once the template is stored in the GitHub.

This template is then referenced in Terraform Cloud and the ServiceNow instance to automate the creation of a subscription account.

Important:

Don’t modify the template as the variables are closely connected to the application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a Terraform organization for Cloud Account Management in Cloud Workspace

Integrate with Terraform Cloud/Terraform Enterprise to organize accounts, define workflows, enforce security, improve collaboration, and enhance scalability.

Before you begin

Role required: Terraform admin

Attributes required for this setup:

- AWS Access Key
- AWS Secret Key

Make sure you have access to Terraform Cloud or Terraform Enterprise to perform the following procedure.

Procedure

1. Log in and navigate to the Terraform Cloud (`https://app.terraform.io/session`) or Terraform Enterprise portal (`https://<tfe_server_url>/api/v1`).
2. Select the user name or organization name.
3. In the modal window, enter a unique org name and other details.
4. Select **Create Organization**.

   Note:

   - After a new organization is created, you'll receive confirmation within the platform. You can manage your organization settings, add members, and configure workspaces within your newly created organization.
   - To see to the official Terraform documentation:
     - Terraform Cloud: https://developer.hashicorp.com/terraform/cloud-docs/users-teams-organizations/organizations
     - Terraform Enterprise: https://developer.hashicorp.com/terraform/cloud-docs/users-teams-organizations/organizations

5. In Terraform org, navigate to **Settings** > **Variable set** > **Create variable set**.

   A variable set is a collection of variables shareable across multiple workspaces.

6. Provide a descriptive name like AWS credentials.
7. Select the appropriate scope (organization or a workspace).
8. Create two variable sets within the **Variables** section as follows:
   - Key: AWS_ACCESS_KEY_ID and Value: <AWS access key>
   - Key: AWS_SECRET_ACCESS_KEY and Value: <AWS secret key>

     Note: Select as **Sensitive** to restrict the visibility of credentials to unauthorized personnel with Terraform access.

   For more details on variable sets, see https://developer.hashicorp.com/terraform/tutorials/cloud-get-started/cloud-create-variable-set#.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integrate Terraform Cloud with GitHub

Integrate GitHub with Terraform to link workspaces to repositories, enabling version control, change tracking, and rollbacks.

Before you begin

Role required: Terraform admin or DevOps team member

Procedure

1. Go to the Terraform org created for this application.
2. Navigate to **Settings** > **Providers**.
3. Select **Add a VCS provider**.
4. From the GitHub drop-down list, choose a **GitHub.com (Custom)**.
5. Select the link **Register a new OAuth Application**.

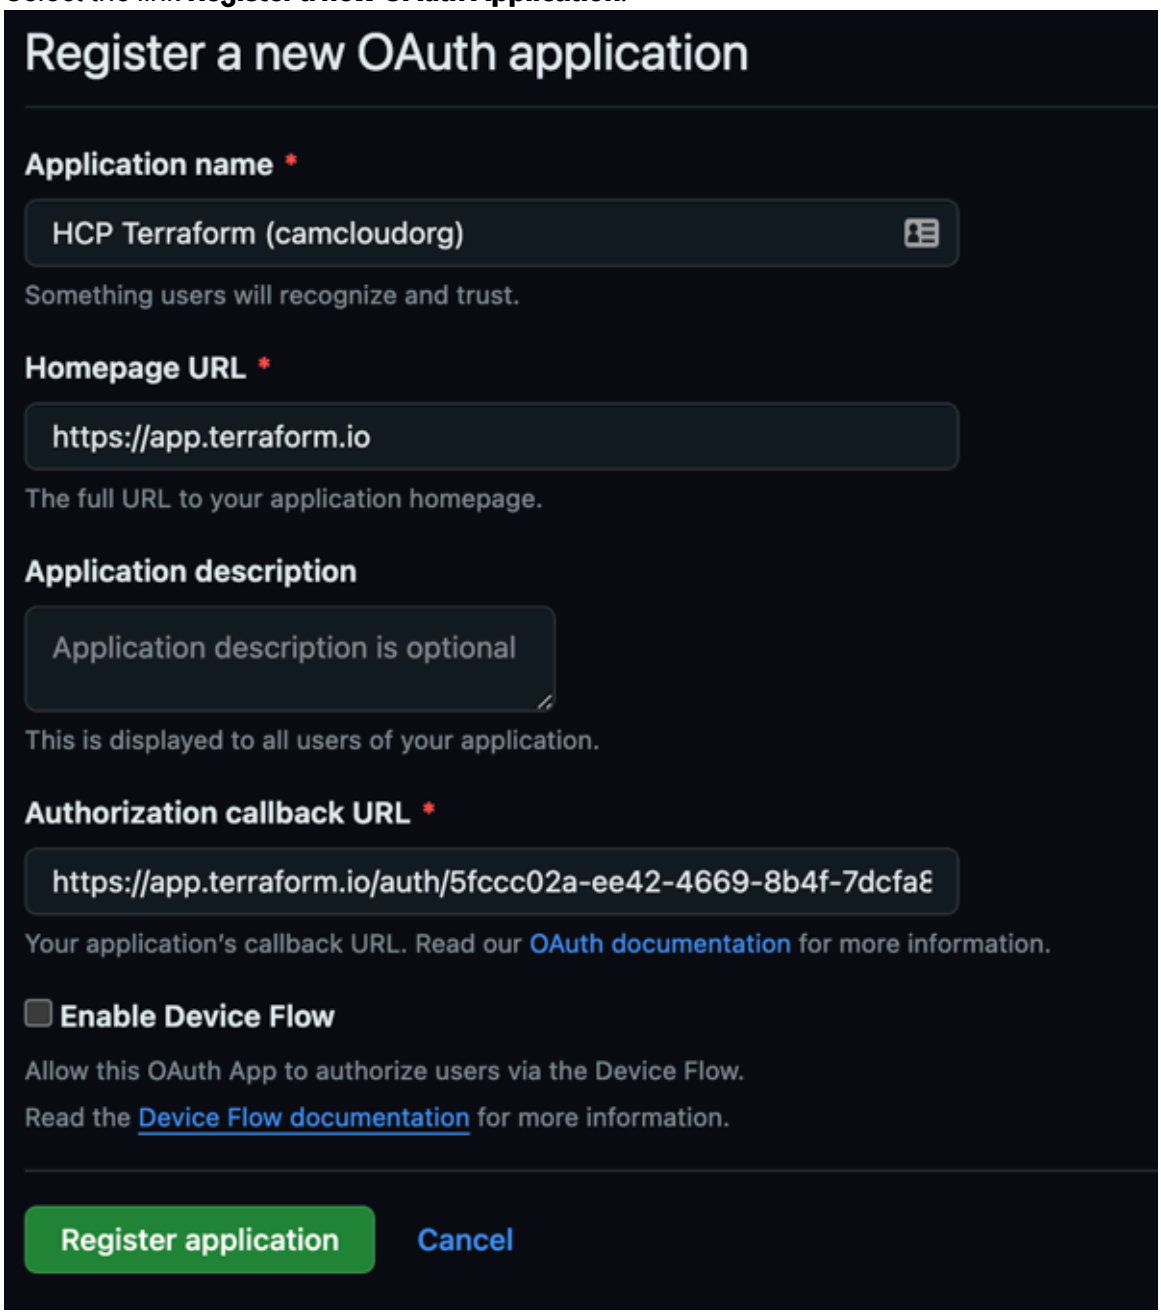

6. Select **Register application**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Copy the client ID.
8. Select **Generate a new client secret** and copy the generated Client Secret.
9. Go back to the Terraform console where the **Add a VCS provider** page is open and paste both Client ID and Client Secret.

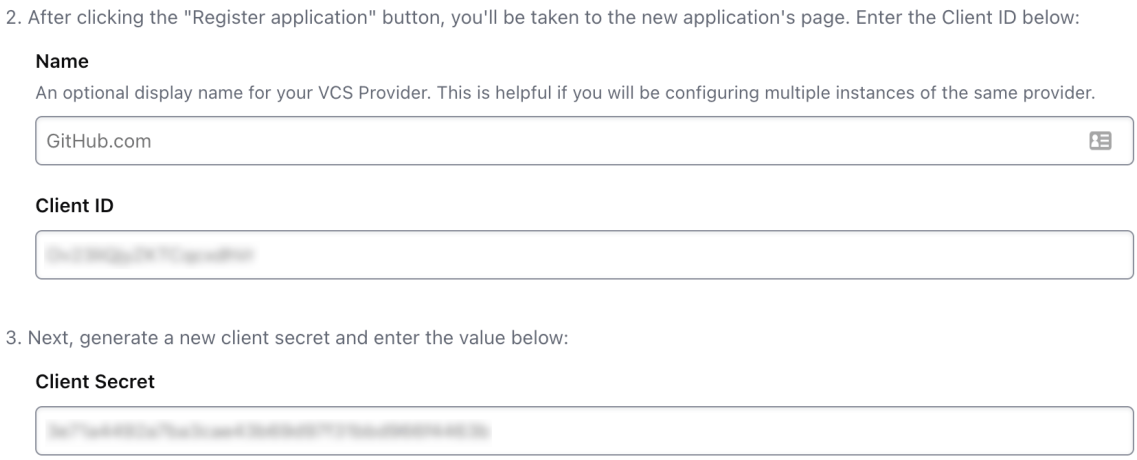

10. Select **Connect and continue**.
11. Select **Authorize**.

> **Note:** Follow the platform-specific instructions to authorize Terraform Cloud to access your GitHub repository.
>
> For more details, see the official documentation: https://developer.hashicorp.com/terraform/cloud-docs/vcs/Github-enterprise.

12. On the **Advance Settings** page, select **Skip and finish**.

What to do next

Make a note of the OAuth Token ID and share it with the ServiceNow admin.

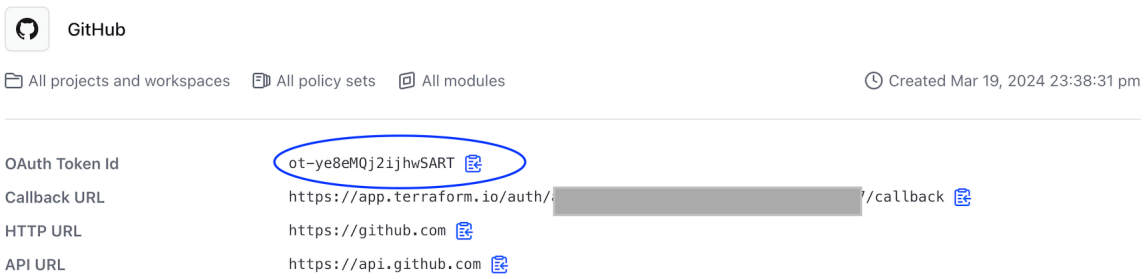

Create Terraform API token

Generating API tokens with limited permissions enhances security, enables fine-grained control, facilitates automation, and provides temporary access within your Terraform organization.

Before you begin

Role required: Terraform admin

Procedure

1. Log in to the Terraform console.
2. Select **Account Settings** from the user profile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Tokens** and then select **Create an API token**.
4. Enter **Description**.
5. Set the **Expiration** date.

   Note: The expiration date should align with the security standards of your company.

6. Select **Generate token**.
7. Share the following details with the ServiceNow admin to register in the ServiceNow instance:
   - VCS Identifier or the GitHub repository location, where you have stored the Terraform template, for example: `https://github.com/<your_repository>/cam-create-aws-account`.
   - Terraform Organization Name
   - Terraform OAuth Token ID
   - Terraform API Key Token
   - Terraform URL

What to do next

Setting up Cloud Account Management in Cloud Workspace

Add members to the group

Set up Terraform API key in ServiceNow.

Set up scan configuration for data visualization.

Setting up Azure cloud

This section covers all the topics required to set up Azure cloud. The topics are arranged in order of configuration priority.

Configure account suspension in Azure

Manage an Azure subscription using the permission and by assigning the role to a user. The role must have the permission to execute the APIs for suspending and reactivating an Azure account.

Before you begin

Role required: Azure admin

Procedure

1. Sign in to the Azure organization.
2. Search and select **Management groups**.
3. Select **Tenant Root Group**.
4. Select **Access Control** > **Roles**.
5. Select **Add** and then select **Add custom role**.
6. In the **Custom role name** box, specify a name for the custom role.
   The name must be unique for the Microsoft Entra directory.
7. Select **Next**.
8. On the **Permissions** tab, select **Next**.
9. On the **Assignable scopes** tab, you specify where your custom role is available for assignment, such as management group, subscriptions, or resource groups.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Select Add assignable scopes to open the Add assignable scopes pane.
11. On the **JSON** tab, paste the following code:

```
{
    "properties": {
        "roleName": "Policy Lock/Unlock Manager",
        "description": "Allows locking and unlocking Azure
 Policy assignments at management group level",
        "assignableScopes": [

 "/providers/Microsoft.Management/managementGroups/<ManagementG
roupId>”
        ],
        "permissions": [
            {
                "actions": [

 "Microsoft.Authorization/policyAssignments/write",

 "Microsoft.Authorization/policyAssignments/delete",

 "Microsoft.Authorization/policyAssignments/read"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
}
```

12. Select **Next** and then select **Create**.

What to do next

You must assign the permission to a role. For more information, see https://learn.microsoft.com/en-us/azure/role-based-access-control/role-assignments-portal.

Configure a lock/unlock policy. For more information, see Set up suspension of a subscription using Azure policy.

Set up Azure Service Principal credential

Provide the Azure credentials obtained from your Azure administrator. These credentials are used to create a suspension profile and enables you to suspend temporarily or terminate Azure accounts as needed.

Before you begin

Review the Permissions required for Azure Service Principal topic before creating a Azure Service Principal credential.

Role required: ServiceNow AI Platform admin

Procedure

1. Log in to the ServiceNow instance.
2. Navigate to **All** > **Connections & Credentials** > **Credentials**.
3. From the list, select **Azure Service Principal**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the **Connections & Credentials** form, provide the following information.

Field	Value
Name	Name of the service principal to register with the instance. For example, `Azure service principal credentials`.
Tenant ID	Directory ID value in the Azure configuration.
Client ID	Application ID value in the Azure configuration.
Secret key	Application key value in the Azure configuration.
Credential Alias	The credential alias that you want to tie to the OAuth 2.0 credential. By default, use the predefined alias name as: sn_itom_cam.CAM_Azure_Connection

5. In the **API Key** field, enter the key obtained from the admin.
6. To complete the process, select **Submit**.

Set up suspension of a subscription using Azure policy

Create the policy at the root level to lock or unlock an Azure account. As an Azure administrator, lock an Azure subscription, resource group, or resource to avoid accidental deletions and modifications.

Before you begin

Make sure you have already configured the permission. For more information, see Configure account suspension in Azure.

Role required: Azure admin

Procedure

1. Sign in to the Azure organization.
2. Search for and select **Users**.
3. Under the user name list, select a user.
4. From the navigation pane, select **Assigned roles**.
5. Under the Administrative roles, add the **Global Administrator** role by selecting **Add Assignments**.
6. Enter `Microsoft Entra ID` in the search box and select **Microsoft Entra ID** > **Properties**.
7. Under **Access Management for Azure resources**, set the toggle to **Yes**.
   Using this option, you can manage access to all Azure subscriptions and management groups in this tenant.
8. Search for and select **Policy**.
9. Select **+ Policy definition**.
10. Enter the following:
    a. Select the tenant root group using the ellipsis on **Definition location**.

       You can select either a management group or a subscription. If you select the tenant root group, all child subscriptions can also be managed.

    b. Enter the name of the policy definition. For example, LockAccount_policy.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. The description of what the policy definition is intended to do.

d. Under POLICY RULE, copy the following JSON code:

```
{
    "properties": {
        "displayName": "CAM_LockAccount_Policy",
        "policyType": "Custom",
        "mode": "All",
        "description": "Blocks the creation of resources and
configurations that fall under Azure Policy enforcement.",
        "parameters": {
            "allowedResourceTypes": {
                "type": "Array",
                "metadata": {
                    "displayName": "Allowed Resource Types",
                    "description": "List of resource types
that are allowed for creation. Any resource type not listed
here will be blocked.",
                    "strongType": "resourceTypes"
                },
                "defaultValue": []
            },
            "allowedLocations": {
                "type": "Array",
                "metadata": {
                    "displayName": "Allowed Locations",
                    "description": "List of allowed Azure
regions for resource creation."
                },
                "defaultValue": []
            },
            "effect": {
                "type": "String",
                "metadata": {
                    "displayName": "Effect",
                    "description": "The effect determines
what happens when the policy rule is evaluated to match"
                },
                "allowedValues": [
                    "Audit",
                    "Deny",
                    "Disabled"
                ],
                "defaultValue": "Deny"
            }
        },
        "policyRule": {
            "if": {
                "anyOf": [
                    {
                        "not": {
                            "field": "type",
                            "in":
"[parameters('allowedResourceTypes')]"
                        }
                    },
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                    {
                        "not": {
                            "field": "location",
                            "in":
 "[parameters('allowedLocations')]"
                        }
                    }
                ]
            },
            "then": {
                "effect": "[parameters('effect')]"
            }
        }
    }
}
```

e. Select **Save**.

Note: To verify if the policy is created, go to PolicyDefinitions. Change the filter Scope and Policy type to find the policy.

f. Select the policy name and copy the **Definition ID**.

Note: This Definition ID is required when you Create a suspend account configuration.

What to do next

Setting up Cloud Account Management in Cloud Workspace

Add members to the group

Set up Azure connection

Set up scan configuration for data visualization

Setting up Cloud Account Management in Cloud Workspace

This section covers all the general topics required to set up Cloud Account Management. The topics are arranged in order of configuration priority.

Install Cloud Workspace

You can install the Cloud Workspace application (sn_itom_cam) if you have the admin role. The application includes demo data and installs related ServiceNow® Store applications and plugins if they are not already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Cloud Account Management application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

Roles are installed with Cloud Workspace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Plugins installed in Cloud Account Management.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Cloud Workspace application (sn_itom_cam) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the available versions are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. **Optional:** If demo data is available and you want to install it, select the **Load demo data** check box.
   Demo data are the sample records that describe application features for common use cases. Load the demo data when you first install the application on a development or test instance.
6. Select **Install**.

Populate Cloud Account Management dependent tables

Provide the required information for cloud account creation and provisioning.

Before you begin

Role required: sn_itom_cam.cw_admin

Procedure

1. Add environment elements to help categorize cloud accounts based on multiple environments.

   a. Navigate to **All** and search for `sys_choice.LIST`.

   b. On the Choice page, select **Element** from the drop-down list and enter `environment` in the search field.

   c. If environment elements you want to use do not exist, create them.

      i. Select **New**.

      ii. On the form, fill in the required fields.

Choice New record form

Field	Value
Table	Choice [sys_choice] table
Element	`environment` (case sensitive)
Label	The name for the environment you want to appear on the cloud account request page, for example, Production.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Value	The same value as the label.

iii. Select **Submit**.

2. Populate the remaining tables.

   a. Navigate to **All** and search for the table name.

   b. Add the relevant information.

Table name	Description
Environment (sys_choice)	Contains translated text for options that appear in lists.
Cost Center (cmn_cost_center)	Contains commonly used reference between financial systems and IT. Cost center records represent business entities and have a related list of CI Cost Center Relationships that measure the cost center's consumption of business services.
Business Unit (business_unit)	Contains information on the business units within your organization, including functions like Finance, HR, or IT.
Department (cmn_department)	Contains the names of the departments in your organization.
System User (sys_user)	Contains user information for the ServiceNow instance.

Create a service account

Create a service account in CAM. A service account holds the credential and account information that you created in your provider account. Discovery uses the information to access your provider account and then obtain information on each logical datacenter that is associated with the account.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Service Accounts**.
2. Fill in the fields.
   For more information, see Cloud Service Account form.
3. Select **Submit**.

Define and run discovery schedule

The discovery schedule is the control point for running discoveries. The schedule controls when discovery runs, defines the MID Server to use, the type of discovery that should run, and the IP addresses to query.

Before you begin

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.

   The Cloud Admin Portal opens in a new browser tab.

2. On the Cloud Admin Portal, navigate to **Manage** > **Discovery Schedules**.
3. Select **Add New Schedule**.
4. Fill in the details (by selecting the provider).

   For more information on the field details, see Create a new discovery schedule form reference.

   If the system can't detect an appropriate MID Server, you're prompted to configure a MID Server.

5. Select **Next**.

   After the schedule runs, the system discovers resources from the org and org units linked to the AWS account or Azure tenant and management groups. The details are populated in the following tables:

   - cmdb_ci_cloud_org_list.do
   - For AWS: cmdb_ci_aws_org_unit_list.do
   - For Azure: cmdb_ci_azure_management_group

Create an application service for Cloud Account Management

Create an application service to standardize the organization, maintenance, and monitoring of services in your organization.

Before you begin

Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Configuration** > **Application Services**.
2. Select **New**.
3. In the **Basic Details** section, enter a `name`.
   This name appears in the account request form once you select the environment.
4. Select **Environment** from the drop-down list.
   The selected environment is mapped to the application service.
5. Select **Next**.
6. Select **Next** > **Done**.
   An application service is successfully created.

Review default Cloud Account Management certification policy

Review the default certification policy provided in Cloud Account Management to validate data against the policy set. This default policy certifies all available cloud service accounts every 90 days. Administrators can customize the default policy or create one as needed.

Before you begin

Role required: sn_itom_cam.cw_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Account Certification**.

   The policies dashboard opens in CMDB Workspace, displaying existing policies under the **Published Policies** section.

2. Select **CW Certification Policy** > **View Policy** to review policy details.
3. For creating or customizing a policy, see Create a CMDB Data Manager policy.

Note:

- To stop data validation or edit a policy, select **Deactivate Policy**.
- After the policy is executed, view tasks by selecting the **My Work** tab.
- Once the tasks are assigned to the Cloud Account Management Certifier group, it enables certifiers to review read-only fields and modify those fields that they're permitted to edit.
- For certifying an account, see Certify an account.

What to do next

Review request policies

Configure a custom name for Cloud Account Management

Configure a custom name for Cloud Account Management. You can optionally re-brand the Cloud Account Management with your company or any custom name. The name change is applicable only on the Cloud Account Management home page and navigation.

Before you begin

Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Now Experience Framework** > **Experiences**.
2. From the list of titles, search for the Cloud Account Management and double-click on the name to make the modification.
3. Enter a custom name and select the check mark.
4. Re-login to see the change.

Create a variable set for Request cloud account catalog

Customize the Request cloud account catalog by creating or editing variable sets. Build a variable set to group related variables into a single, reusable collection. Select only the necessary variables for specific cloud account requests to promote an organized and uniform process.

Before you begin

Role required: admin

Procedure

1. Navigate to **Workspaces** > **Cloud Workspace**.
2. Enter `sc_cat_item.LIST` in the navigation filter.
3. In the sc_cat_item.LIST table, search for and click the *`Request Cloud Account`* property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Click the **Variable Sets** tab and click **New**.
   The Variable Set screen appears.
5. Fill in the form details (as shown in the table).

Field	Description
Title	Unique name for the variable set
Internal Name	A (unique) variable set name for internal use. **Note:** If this field is empty, its value is auto-populated based on the **Title** field for all variable types except break, container split, and container end.
Layout	Layout display selections from: **1 Column Wide**, **2 Columns Wide alternating sides**, or **2 Columns Wide, one side, then the other**.
Order	Order number for the variable set.
Description	Description for the variable set.

6. Click **Submit**.
   The new variable set is created and appears in the **Variable Sets** tab. Open the new variable set and click **New** in the **Cloud Variables** tab to create variables for the set.

Edit a variable for Request cloud account catalog

Customize the cloud account request process by creating or editing variable to collect essential information during cloud account request.

Before you begin

Avoid removing the existing Request cloud account variables, as these variables are essential for creating a valid CAM request.

Role required: sn_itom_cam.cw_admin

About this task

Modify the Request cloud account form by creating or editing variables in the catalog. You can edit a variable to change its order, make the variable active or inactive, modify its permission and availability settings.

Procedure

1. Navigate to **Workspaces** > **Cloud Workspace**.
2. Enter `sc_cat_item.LIST` in the navigation filter.
3. In the sc_cat_item.LIST table, search for and click the *`Request Cloud Account`* property.
4. Click the **Variable Sets** subtab, and click the variable set.
5. Click the variable that you want to edit, or to create a variable, click **New**.
6. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Variable form

Field	Description
Type	Variable type to create. If you select **Lookup Select Box**, the **Type Specifications** subtab displays these fields: ○ **Use Pool Filter**: Select this check box to add a pool filter to the field. ○ **Pool**: Select a pool to offer the cloud service user a list of values from an existing table in the catalog order form. For example, you can select CloudAccountPool to get a list of all the cloud accounts. ○ **Pool Filter**: Select a filter to further narrow down the choices in the pool.
Variable Set	The variable set that the variable belongs to.
Mandatory	Option to make the variable mandatory as part of the ordering process.
Active	Option to make the variable available for use.
Order	Order that the variable is placed on the page for the catalog item. The variables are organized from top to bottom, from the least to greatest order value. For example, a variable with an order value of 1 is placed above other variables with higher-order values.

Questions

Field	Description
Question	Question to ask users ordering the catalog item. **Note:** If you select **Select Box** in the **Type** field, a **Question Choices** subtab appears at the end of the screen. You can enter a specific list of choices for the question.
Name	A name to identify the question. **Note:** If this field is empty, its value is auto-populated based on the Question field for all variable types except break, container split, and container end, HTML, container start, label, UI page, and duration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Tooltip	Tooltip text to display when users point to the variable. Enter a brief note to describe the purpose of the question.
Example text	Hint that's displayed in the question field before a user enters a value. Applicable for the following variables: ○ IP address ○ Email ○ URL ○ Single Line text ○ Multi Line text ○ Date ○ Date/Time

Annotation

Field	Description
Show Help	If selected, displays the help text for the variable.
Always Expanded	This option appears only when the **Show help** check box is selected. If selected, the **Help text** field value is expanded by default when the catalog item page loads.
Help tag	If the **Always Expanded** check box is cleared, click the value specified in this field to display the **Help text** field value.
Help text	Help information for a service catalog variable.
Instructions	Information that requires rich text formatting or adding images to support help information. **Note:** For HTML tables, use sizes that are within the width of the variable.

Type Specifications (The fields in this section vary for each variable type)

Field	Description
Variable Width	Width for the variable on the catalog item page, to specify what percentage of the screen size that it can span. This field appears for all variable types except for break, container end, container

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	start, container split, container layout, and label variables.
Validation (Regex)	Regular expression that validates the user input for the property.
Validation Error (Regex)	If the expression isn't valid, an error appears.

Default Value

Field	Description
Default Value	Default value for the variable.
Variable Attributes	Attributes for this variable.

Permissions

Field	Description
Read Roles	Roles that can view the variable before or after requesting the catalog item or record producer. Only a user with the roles specified in this field can view the variable.
Write Roles	Roles that can edit the variable in the variable editor after requesting the catalog item or record producer. If a user doesn't have the roles specified in this field, the variable is read-only in the variable editor.
Create Roles	Roles that can create values for the variable before requesting the catalog item or record producer. If a user doesn't have the specified role, the variable is read-only before requesting the catalog item or record producer.

Availability

Visible	If selected, the variable is visible in the item form before ordering and in the cart view. The variable is also visible in the Variable editor after ordering.
Read-only	If selected, the variable is read-only in the Cloud User Portal and cannot be edited.

7. Click **Update** once you have made all the changes to the variable.

Configure a custom catalog ID in Cloud Account Management account request

Configure the catalog ID used for Cloud Account Management request to map the custom service catalog and streamline account provisioning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Verify whether the catalog is a valid service catalog tested on various portals, such as the Employee Center.

Role required: sn_itom_cam.cw_admin

About this task
The Catalog ID uniquely identifies the Catalog used for Cloud Account Management. To map your custom cloud account request catalog, update the default sys_id value with the custom catalog sys_id value.

Procedure
1. Navigate to **Workspaces** > **Cloud Workspace**.
2. Enter `sys_properties.list` in the navigation filter.
3. In the System Properties [sys_properties] table, search for and click the *`sn_itom_cam.cam_catalog_id`* record.
4. In the **Value** field, enter the unique sys_id of the custom catalog.

   Note: You can add only one sys_id in the **Value** field.

5. Click **Update**.

Add members to the group

Ensure accurate assignment of members to their respective groups to streamline account requests, approvals, provisioning, and certification. Grant necessary permissions through proper group assignment to prevent unauthorized access and maintain security.

Before you begin
Role required: ServiceNow AI Platform admin

Procedure
1. Navigate to **All** > **Users and Groups** > **Groups**.
2. Search for `*ITOM` in the search field.
3. Assign a member to each ITOM group by selecting an entry.
   To know more about ITOM Cloud Account Management user roles, see Components installed with Cloud Account Management.
4. Under the **Roles** section, select **Edit**.
5. Search the user and move to the group members list.
6. Select **Save**.

Set up AWS API configuration information in ServiceNow

The credentials provided by your AWS administrator are used in this procedure to create a suspension profile, enabling you to temporarily suspend or terminate AWS accounts as needed.

Before you begin
The below procedure is only required if you're managing AWS accounts on Cloud Account Management.

Role required: ServiceNow AI Platform admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Log in to ServiceNow instance.
2. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
3. Select **New**.
4. In the **Name** field, enter the alias name.
   The format should be 'cloud org name-key' (for example, org1-tfkey).
5. From the **Type** drop-down list, select **Credential**.
6. Select **Submit**.
7. Select the aliases name from the list that you entered earlier.
8. Under the **Credentials** tab, select **New**.
9. From the list, select **AWS Credentials** as your credential type.
10. In the **Name** field, enter the alias name.
11. In the **API Key** and **Secret Access Key** fields, enter the keys obtained from the admin.
12. Select **Submit**.

Set up Terraform API key in ServiceNow

Add the Terraform API key to the ServiceNow instance to ensure consistent and error-free infrastructure provisioning.

Before you begin

The below procedure is only required if you're integrating Terraform.

Role required: ServiceNow AI Platform admin

Procedure

1. Log in to the ServiceNow instance.
2. Navigate to **All** > **Connections & Credentials** > **Credentials**.
3. From the list, select **API Key Credentials**.
4. In the **Name** field, enter the name.
5. In the **API Key** field, enter the key obtained from the Terraform admin.
6. To complete the process, select **Submit**.

Set up Azure connection

Add and configure an Azure connection with your Azure portal. Using the connection credentials, the Cloud Account Management application creates Azure subscriptions. This is a one-time configuration step.

Before you begin

The below procedure is only required if you're creating Azure subscriptions.

Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Flow Designer**.
2. Select **Integrations** > **Connections** > **Outbound**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Search and select **Azure Connection (sn_itom_cam)**.
4. Select **View Details**.
5. Select **Configure**.
6. In the **Configure Connection** dialog box, provide the following information.

Configure Connection

Field	Description
Connection name	Unique name for the Azure OAuth connection.
Connection URL	The default URL is automatically populated. If necessary, you can modify the URL.
OAuth Client ID	The client ID of the application registered in the Azure OAuth server.
OAuth Client Secret	The client secret of the application registered in the Azure OAuth server.
OAuth Token URL	The URL that specifies the endpoint to get the access token from the OAuth server. The default URL is automatically populated. If necessary, you can modify the URL.

7. Select **Configure and Get OAuth Token**.
8. **Optional:** To view the connection alias, select **View connection alias**.

Set up scan configuration for data visualization

Visualize all account violations by scanning them using the selected policy set. Display the severity of all violations on the admin dashboard to take appropriate actions.

Before you begin

Make sure you have referred to the default policies (Policy list for scanning cloud accounts). To create a custom policy set, see Create a policy set.

Role required: sn_itom_cam.cw_admin or sn_itom_ccg.scheduling_admin

Procedure

1. Navigate to **All** > **Cloud Configuration Governance** > **Scan Configurations**.
2. Create a scan configuration by selecting **New**.
3. Identify the configuration by entering `Name` and `Description`.
4. Under the **Configuration** tab, select a **Cloud Provider Type** from the lookup.
5. Select the **Unlock Policy Sets** icon ( ) and select a policy set from the list.
6. Select the **Lock Policy Sets** icon ( ).
7. Select Additional actions icon ( ) or right-click on the top bar and select **Save** to save the current settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

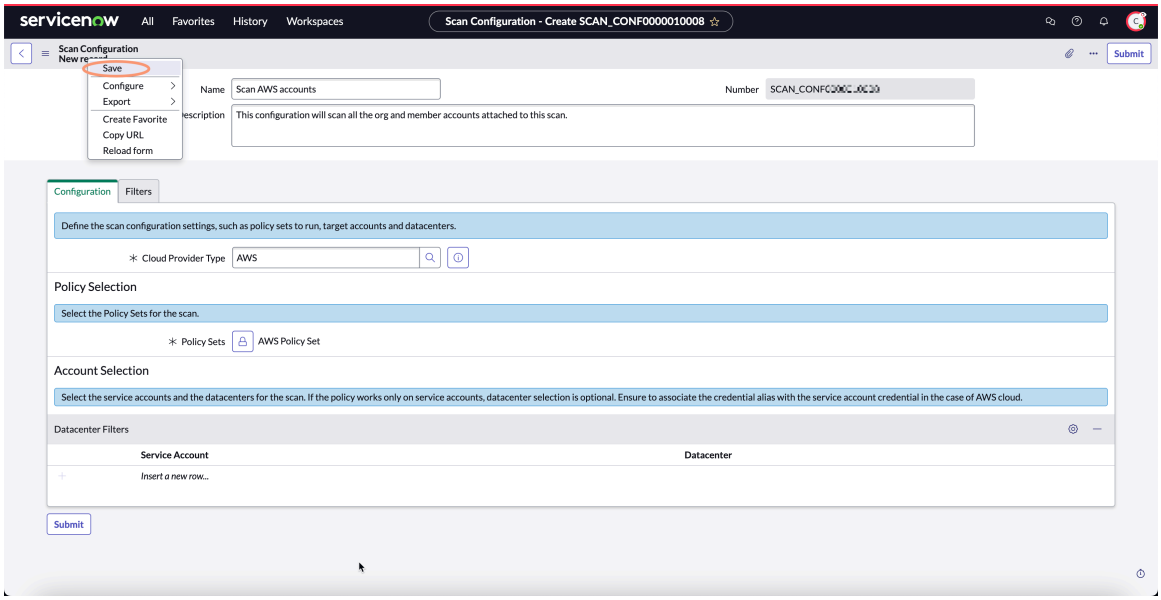

8. Under the **Service Account** column, select all the management accounts and member accounts from the lookup.

> **Note:** Selecting a management account scans all current and future member accounts automatically. Choosing individual member accounts requires manual selection for each scan.

9. Select **Setup Schedule** and set the schedule as necessary.
10. Select **Update**.
11. Synchronize the configuration with the Cloud Account Management app by selecting **Sync CAM Config**.

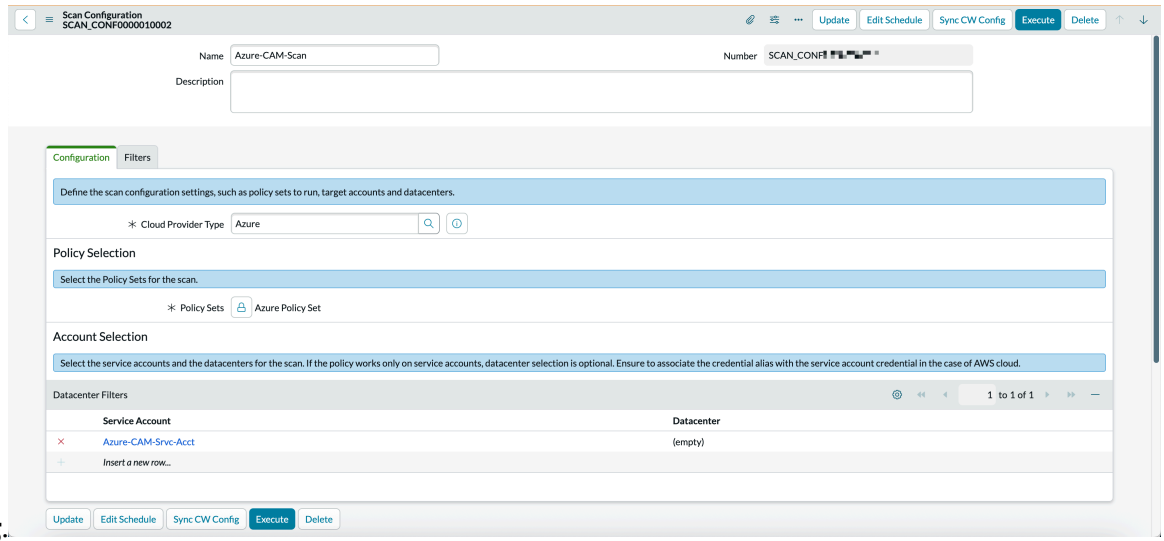

> **Note:**
>
> - To view the synchronized configuration, navigate to **All** > **Cloud Workspace** > **Configuration**.
> - If you have synchronized the configuration, you don't have to create a Scan Account configuration mentioned in Create a scan account configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

12. **Optional:** To run an on-demand scan, select **Execute**.

> **Note:** You can create unlimited configurations for each account. An account may have multiple scan configurations, with violations and assessments based on all of them.

Creating configurations

Setting up the Cloud Account Management app is a prerequisite task before proceeding to create, suspend, or scan a service account. By configuring Cloud Account Management, you can ensure that the application is ready to effectively execute account management tasks and deliver the desired results.

Create a Terraform account configuration

Define and set up configurations based on Terraform to automate cloud account provisioning, including required credentials and infrastructure templates.

Before you begin
Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Configuration**.
2. Select **New**.
3. Enter the details in the **New Configuration** form.

   To fill the form fields, see Terraform account configuration fields.
4. Select **Submit**.

What to do next
Review default Cloud Account Management certification policy

Create a cloud native interface account configuration

Configure cloud accounts using the native interface of supported cloud providers like AWS AWS or Azure to enable seamless provisioning and integration.

Before you begin
Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Configuration**.
2. Select **New**.
3. Enter the details in the **New Configuration** form.

   To fill the form fields, see Cloud native interface configuration account fields.
4. Select **Submit**.

What to do next
Review default Cloud Account Management certification policy

Create a suspend account configuration

Set up a suspension profile that allows temporary suspension or termination of cloud accounts to support cost control or policy enforcement.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_itom_cam.cw_admin

Procedure
1. Navigate to **All** > **Cloud Workspace** > **Configuration**.
2. Select **New**.
3. Enter the details in the **New Configuration** form.

   To fill the form fields, see Suspend account configuration fields.
4. Select **Submit**.

What to do next
Review default Cloud Account Management certification policy

Create a scan account configuration

Configure the Cloud Account Management app before scanning a service account to confirm it's ready for effective account management.

Before you begin
Role required: sn_itom_cam.cw_admin

Procedure
1. Navigate to **All** > **Cloud Workspace** > **Configuration**.
2. Select **New**.
3. In the **New Configuration** form, provide the following information.

Configuration type: Scan accounts

Field Name	Description
Cloud provider	Name of the cloud provider platform such as AWS, Azure, or GCP.
Name	Name for the configuration.
Service accounts	The service account and its member accounts. **Note:** Selecting a management account automatically scans all current and future member accounts. However, if you select individual member accounts, each must be chosen manually for scanning.
CCG scan configuration	Cloud Configuration Governance (CCG) scan configuration used for getting the violations for all accounts in the selected cloud organization.
Description	Information that can be recorded in work notes.

4. Select **Submit**.

What to do next
Review default Cloud Account Management certification policy

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Review request policies

Review and update request policies to confirm they align with your cloud account request process. These policies enforce data checks and conditions, promoting consistency in creating cloud subscription accounts.

Before you begin

Role required: sn_itom_cam.cw_admin

About this task

Certain steps in request processing often require manual intervention. These interventions include general approval, finance approval, and assigning configuration and AWS organizational unit. However, these steps can be automated when specific criteria are met. Cloud Account Management integrates with the PaCE plugin to define and manage these criteria as policies.

For each step, an entry is created in the *sn_itom_cam_automation_activity* table, and the associated policies are displayed under the **Request Policies** tab in the Cloud Workspace. Each policy is linked to an activity via the **Mappings** tab within the policy, although only one policy is supported according to automation activity.

Note: Currently, Azure account creation can't be automated.

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Request Policies**.
   The existing policies appear under the **All Policies** section and are further classified as active policies and inactive policies.
2. View or customize an existing policy by selecting the policy name.
3. Under the **Policy Builder** tab, review the policy conditions.
4. **Optional:** To modify an existing policy, select **Create copy**.
5. **Optional:** To add a condition, select **New condition set**.
6. **Optional:** To remove a condition, select the delete icon .

Note: For more information on managing policies, see Manage PaCE policies.

What to do next

These sections explain how to modify, test, and publish the automation policy.

- Automate General Approval policy process
- Automate Budget Approval policy process
- Automate Configuration Assignment policy process

Automate General Approval policy process

Automate the general approval or rejection of requests based on predefined conditions. Requests are automatically approved when the criteria are met. Otherwise, the system waits for manual approval. You can modify the default policy to include additional conditions as needed.

Before you begin

Role required: sn_itom_cam.cw_admin

About this task

This procedure demonstrates how to approve or reject an account using the general account approval policy. For example, if an account request includes a predefined approval policy, this process explains how the request is approved or rejected automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Request Policies**.
2. Select **Subscription Account Creation General Approval**.
3. Customize the policy by selecting **Create copy**.
4. Under the **Policy builder** tab, in the **If** condition block, make a note of the source fields against the highlighted values.

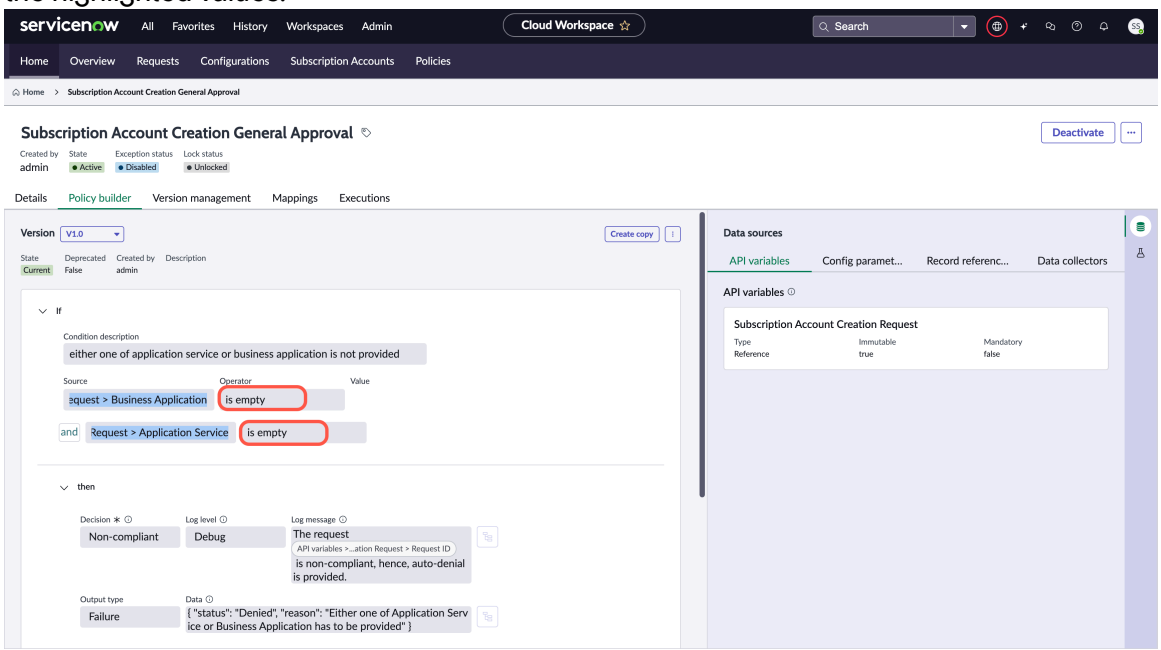

> **Note:** In this section, you're validating two conditions:
>
> **a.** Business Application = empty
>
> **b.** Application Service = empty

5. Select **Mappings**.
6. Verify if the Document assigned is **Automation Activity: General auto approval for subscription account creation**.

> **Note:**
>
> - If no document is linked, select **Add** and choose the **General auto approval for subscription account creation** document.
> - Map only one policy according to automation.
> - View or create documents in the `sn_itom_cam_automation_activity.list` table.

7. **Optional:** To remove auto-approval, go to the **Mappings** tab and delete the mapped document.

What to do next

To test the policy, see Test updates to a PaCE policy version.

To publish the policy, see Publish a PaCE policy version.

To add a new approval step in the flow designer, see Exploring Workflow Studio.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Automate Budget Approval policy process

Automate finance approvals for Cloud Account Management requests based on predefined conditions. Approve requests automatically when criteria are met. Otherwise, wait for manual approval. Modify the default policy to include additional conditions as necessary.

Before you begin

Role required: sn_itom_cam.cw_admin

About this task

This procedure demonstrates how to approve or reject a budget using the account budget approval policy. For example, if an account request includes a predefined approval policy, this process explains how the budget request is approved or rejected automatically.

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Request Policies**.
2. Select **Subscription Account Creation Budget Approval**.
3. Customize the policy by selecting **Create copy**.
4. Under the **Policy builder** tab, in the **If** condition block, make a note of the source fields against the highlighted values.

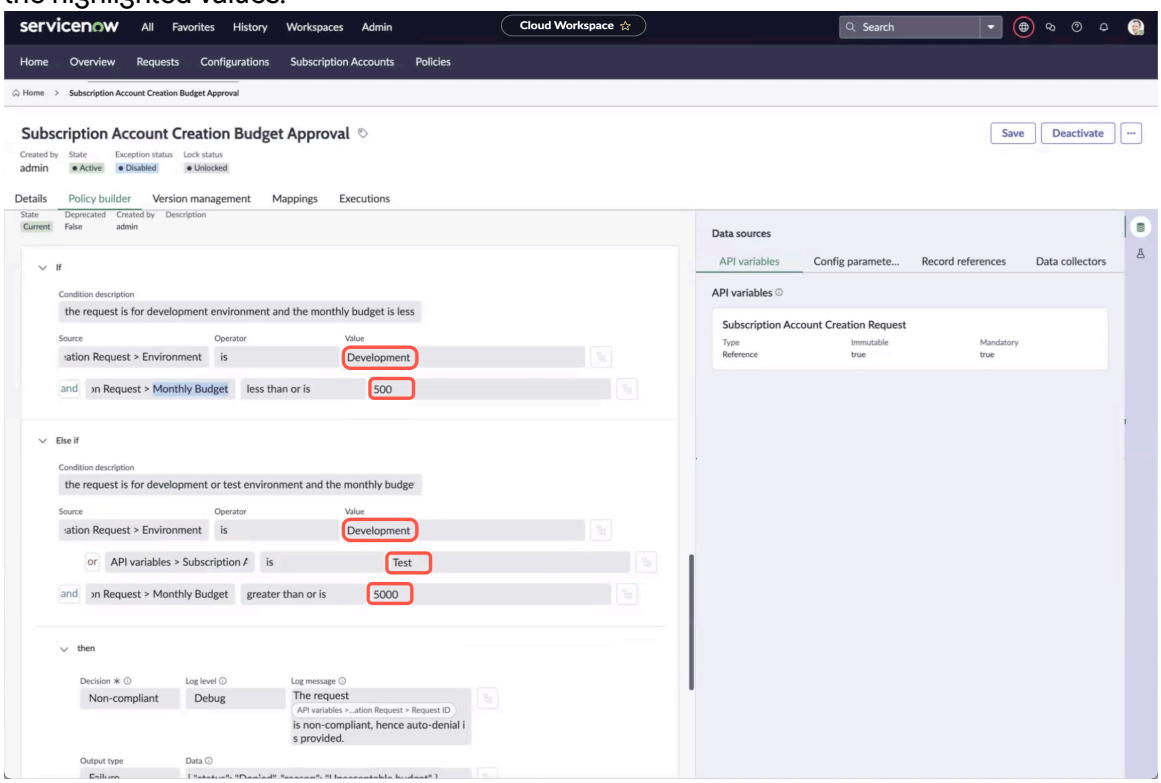

Note: In this section, you're validating two conditions:

a. Environment = Development

b. Monthly budget <= 500

a. Environment = Development or Test

b. Monthly budget = > 5000

5. Select **Mappings**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Verify if the Document is assigned to **Automation Activity: Finance auto approval for subscription account creation**.

   Note:

   - If no document is linked, select **Add** and choose the **Finance auto approval for subscription account creation** document.
   - Map only one policy according to automation.
   - View or create documents in the `sn_itom_cam_automation_activity.list` table.

7. **Optional:** To remove auto-approval, go to the **Mappings** tab and delete the mapped document.

What to do next

To test the policy, see Test updates to a PaCE policy version.

To publish the policy, see Publish a PaCE policy version.

To add a new approval step in the flow designer, see Exploring Workflow Studio.

Automate Configuration Assignment policy process

Automate assigning configurations and AWS organizational units in requests based on specified conditions. When conditions are met and approvals are complete, requests move forward without admin input. Update the default policy to include specific conditions and values for configuration and AWS organizational unit.

Before you begin

Role required: sn_itom_cam.cw_admin

About this task

This procedure demonstrates how to assign account configuration using the assignment policy. For example, when the account creation request is approved and awaits for the provisioning configuration information, this automation is executed.

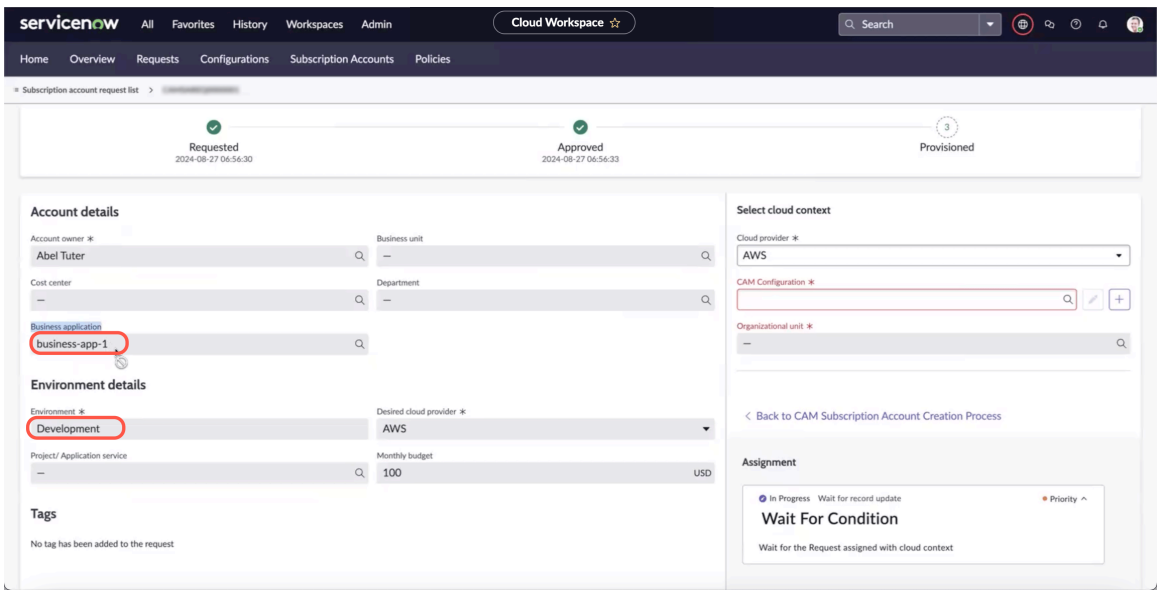

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Request Policies**.
2. Select **Subscription Account Creation Configuration Assignment**.
3. Customize the policy by selecting **Create copy**.
4. Under the **Policy builder** tab, in the **If** condition block, make a note of the values that you're validating.

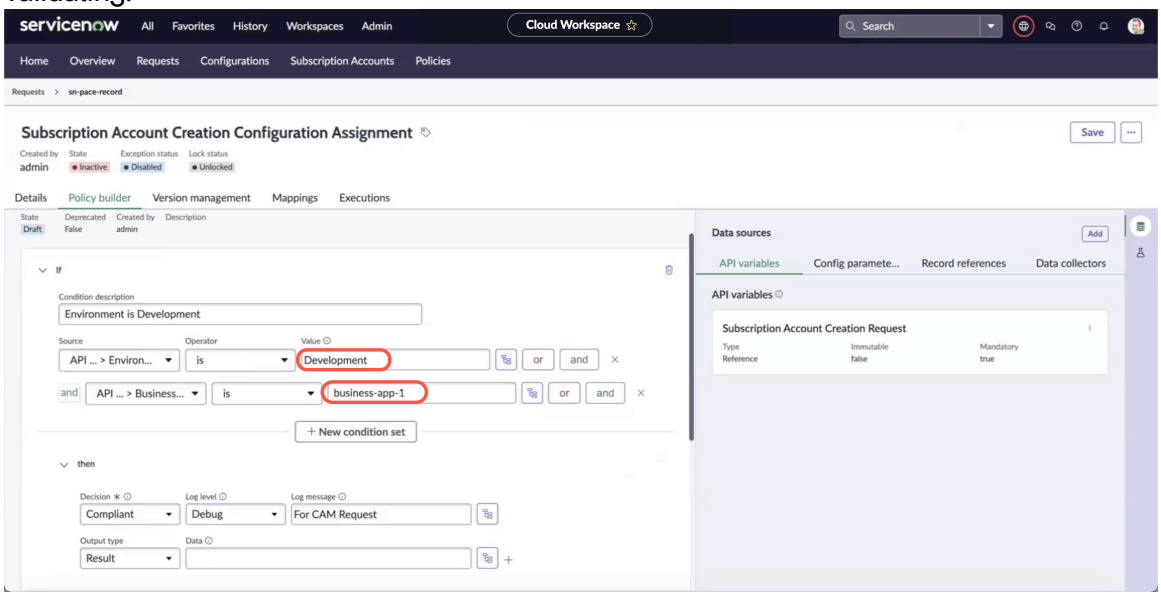

Note: In this section, you're validating:

   a. Environment = Development

   b. Business unit = business-app-1

5. To select the configuration and AWS org unit output, select **Data collectors** from the Data sources section.
6. Select **Add**.
7. Select **Get Config Assignment Results**.
8. In the **Add data collector** pop-up, enter `Label` and `Variable name`.
   These field values are only used to identify the data collector.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

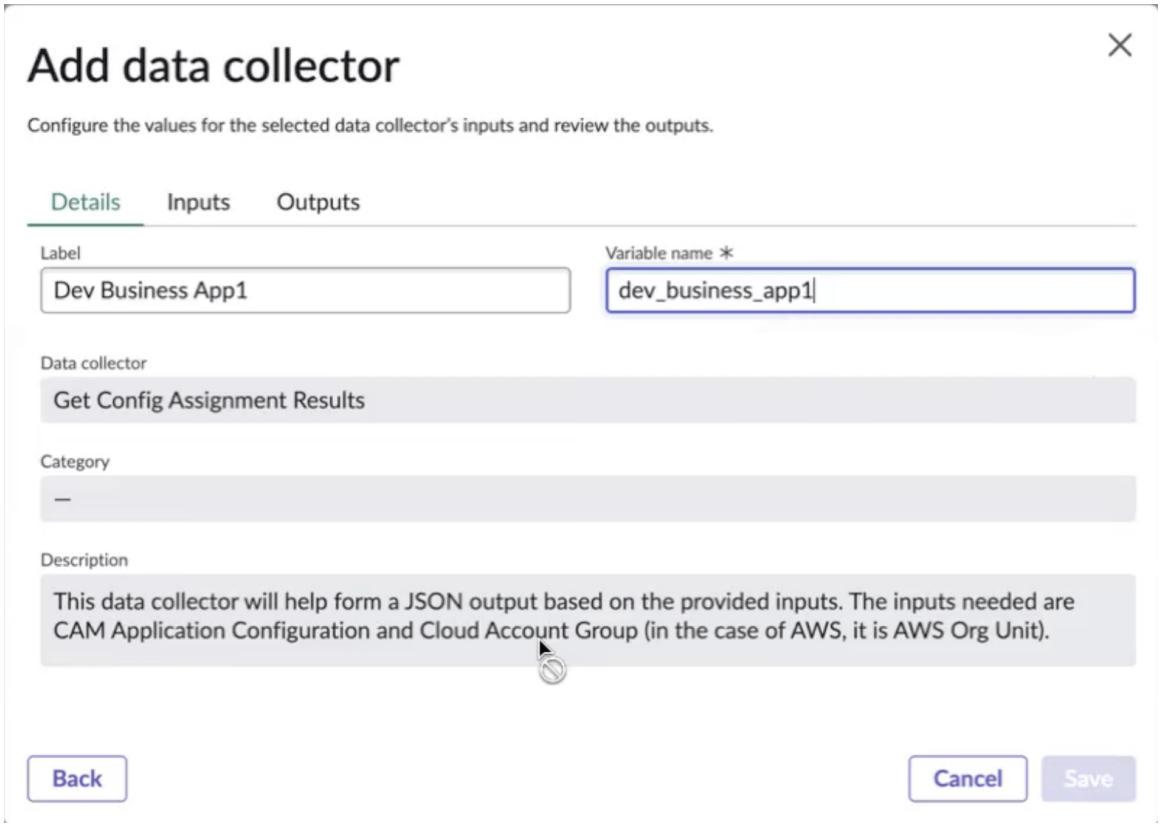

9. Under the **Inputs** tab, select **CW Application Configuration** and **CW Account Group** (same as the AWS org unit that an admin selects on the account request screen).

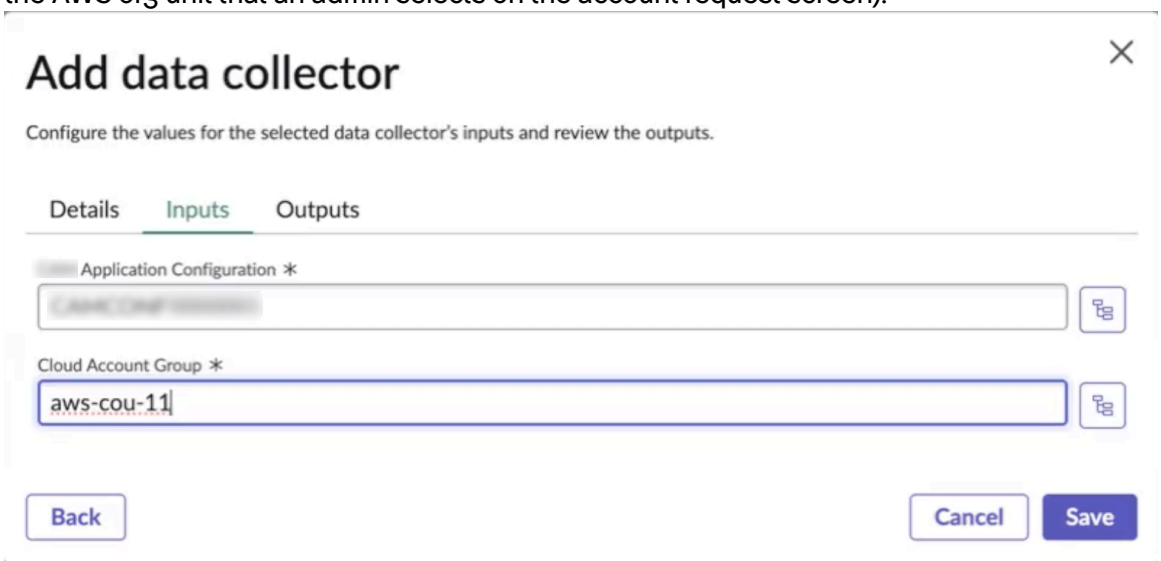

Note: The selected AWS org unit in the data collector should be part of the same Cloud Organization provided in the configuration.

10. Select **Save**.
    The new data collector entry is saved. Only when the **If** condition is met, the **then** condition get executed.
11. Select the output result as follows.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

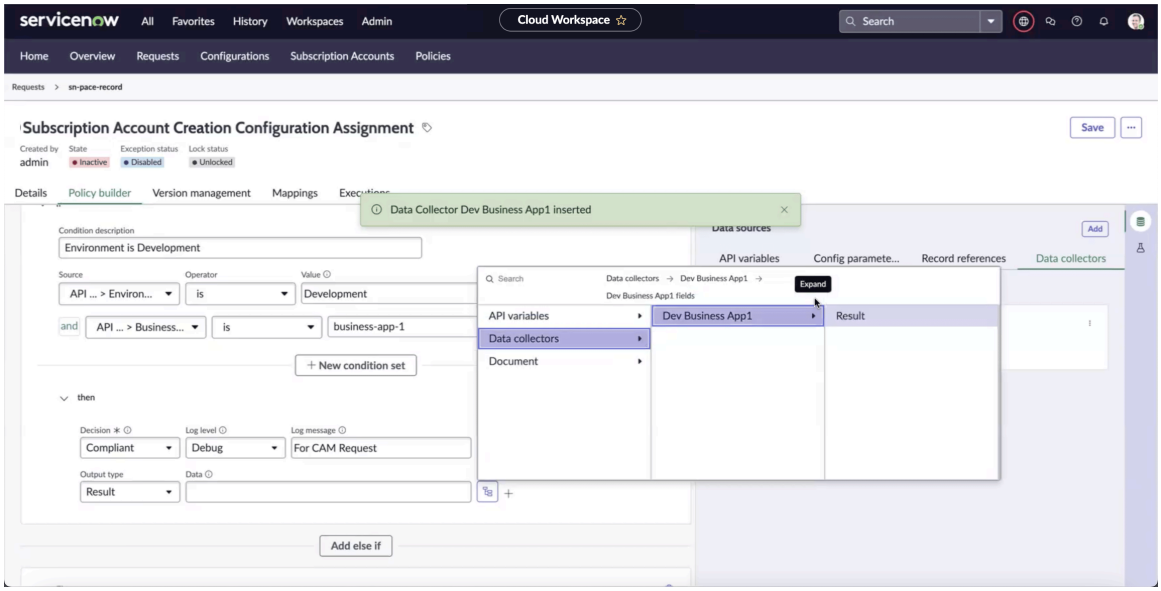

12. To save the policy, select **Save**.

Note:

- If necessary, you can enhance the Log message for the result. For some examples, you can see the default-published policies.
- Map only one policy according to automation.
- This configuration approval policy is inactive by default. You can deactivate to stop the automation.

What to do next

To test the policy, see Test updates to a PaCE policy version.

To publish the policy, see Publish a PaCE policy version.

To add a new approval step in the flow designer, see Exploring Workflow Studio.

Configure cloud asset collection table

Configure the Cloud Asset Collection Configs table to collect cloud assets or add comments to the CI entries.

Before you begin

Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Cloud Workspace**.
2. Remove the URL after *service-now.com/ and enter the table name `sn_itom_cam_asset_collection_config_list.do`.
   The table entries are as follows.

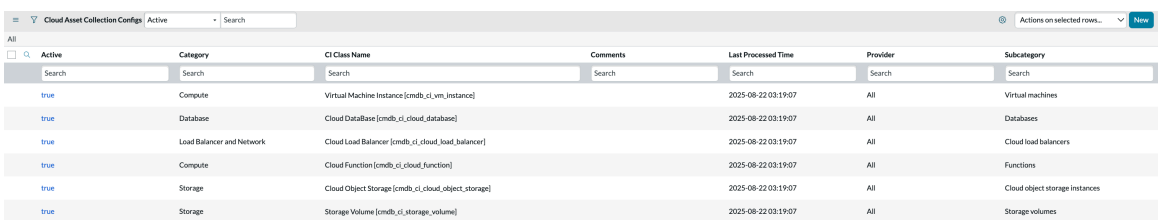

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Double-click an entry to update the Active or Comments column, then select the Save icon.

Note: You can stop collecting cloud assets by setting the Active column as false.

Using Cloud Account Management in Cloud Workspace

The Cloud Account Management app provides the capability to manage your cloud accounts.

Viewing Cloud Account Management dashboards

Cloud Account Management has two types of dashboards: one for requesters and another for admins. Depending on the logged- in user, an appropriate dashboard appears.

Viewing the compliance dashboard

The Cloud Workspace compliance dashboard provides insights that help security teams identify cloud resources without ownership information and flag vulnerable or at-risk resources. This information helps team members conduct security audits and confirm compliance with industry regulations across cloud environments.

As an admin, you can access the dashboard by navigating to **All** > **Cloud Workspace** > **Monitor and track**.

Compliance dashboard

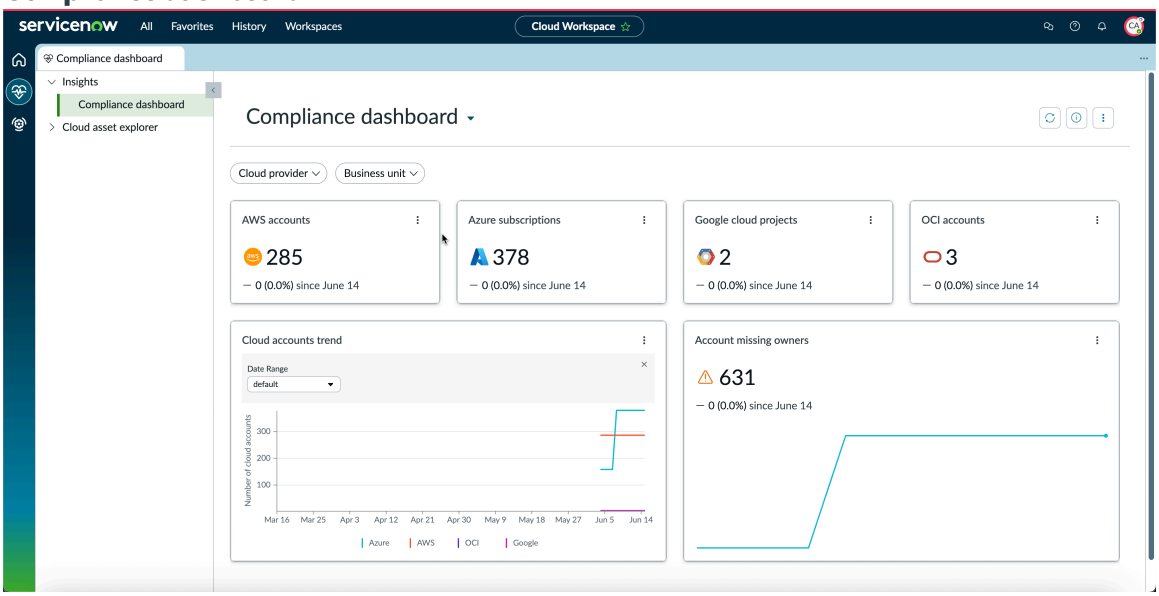

Note:

- The Compliance dashboard displays only the accounts discovered in your instance.
- The Compliance dashboard hides the provider card if no accounts are provisioned.

Viewing the admin home page

View account statuses and violations, actionable insights, and administrative controls on the admin home page.

As an admin or asset viewer, you can access the homepage by navigating to **All** > **Cloud Workspace**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Admin or asset viewer dashboard

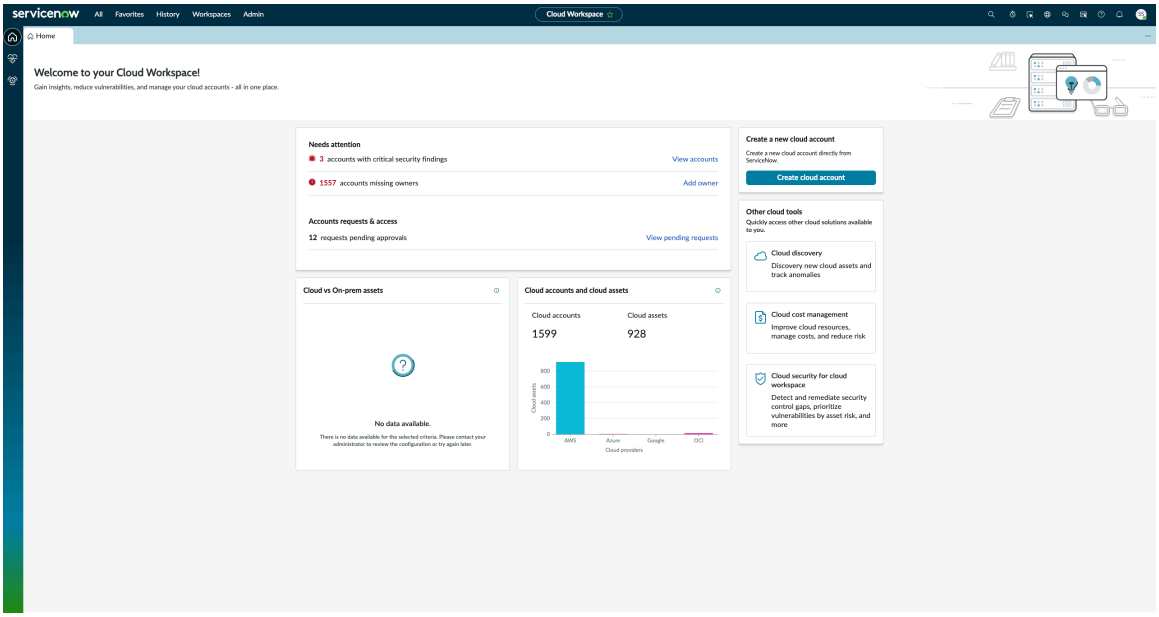

Note:

- The Needs attention widget appears only for admins with the Cloud Account Management entitlement.
- The Other cloud tools widget displays the applications installed on your instance. Cloud Discovery is a prerequisite for installing Cloud Account Management.

Viewing the requester home page

View and manage cloud account requests on requester page. The home page provides a view of requested, approved, denied, pending, and active accounts.

Requesters can access the requester home page by navigating to **All** > **Cloud Workspace**

Requester home page

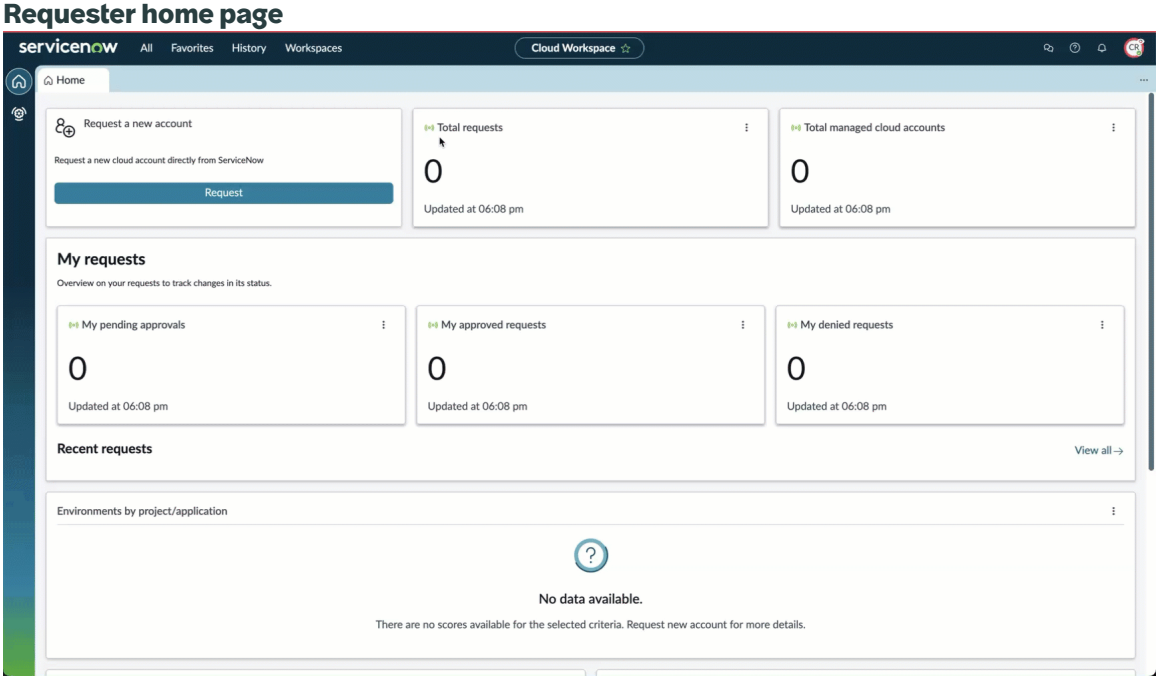

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Viewing the Cloud asset overview page

The Cloud asset overview page helps cloud governance and operations teams monitor, track, and act on compliance and asset details more efficiently. As an asset viewer, you can view all assets and drill down to see detailed information for each configuration item (CI).

can access the Cloud asset overview page by navigating to **All** > **Cloud Workspace** and then in the **Monitor and track** tab, select **Cloud asset explorer** > **Overview**.

Cloud asset overview page

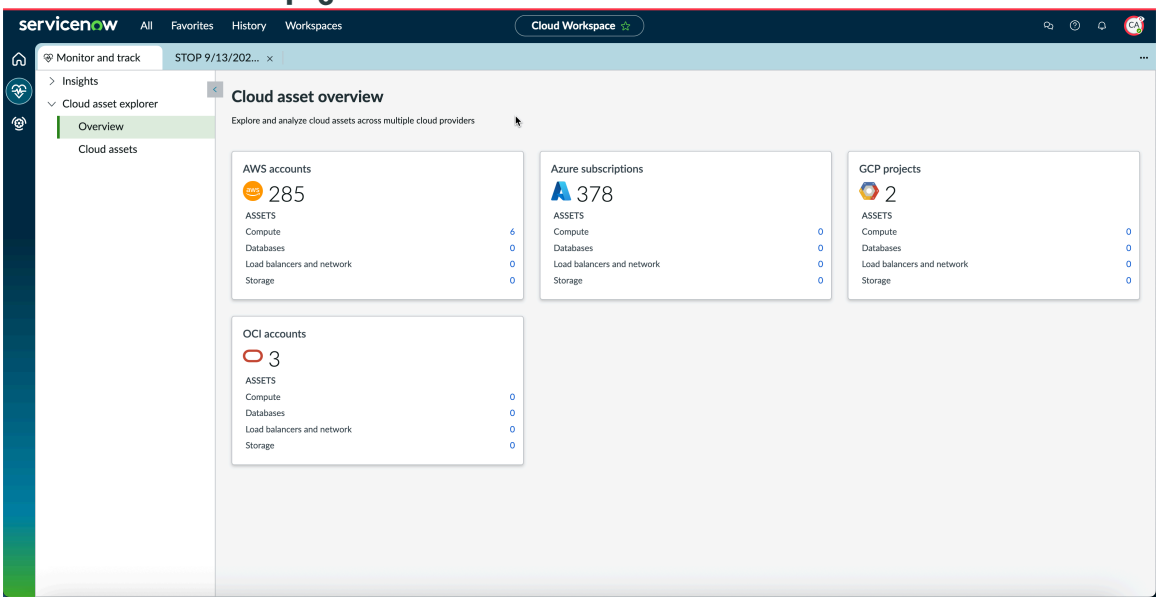

Note: You can enable or disable the Asset collection configuration to list the CIs. For more information, see Configure cloud asset collection table.

Viewing the requester home page

View and manage cloud account requests on requester page. The home page provides a view of requested, approved, denied, pending, and active accounts.

You can access the home page in either of two ways.

- Navigate to **All** > **Cloud Workspace**
- Navigate to **Workspaces** > **Cloud Workspace**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Requester home page

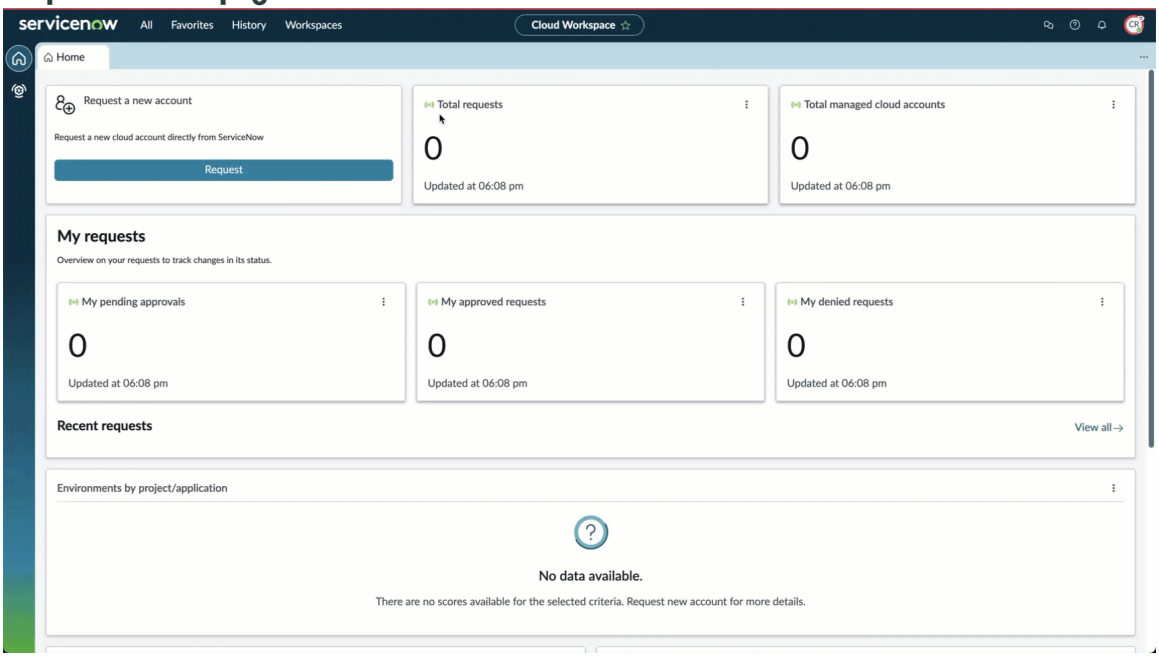

Viewing the home page

View account statuses and violations, actionable insights, and administrative controls on the home page.

Access the home page

As an admin or asset viewer, you can access the home page in either of two ways.

- Navigate to **All** > **Cloud Workspace**
- Navigate to **Workspaces** > **Cloud Workspace**

Admin dashboard

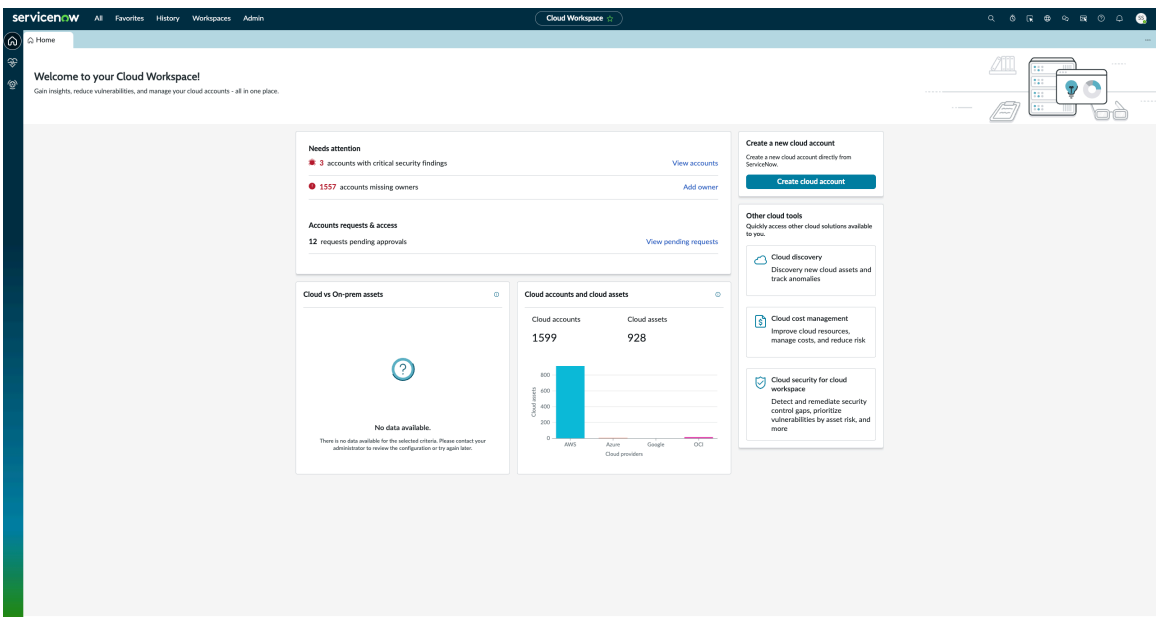

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- The Needs attention widget appears only for admins with the Cloud Account Management entitlement.
- The Other cloud tools widget displays the applications installed on your instance. Cloud Discovery is a prerequisite for installing Cloud Account Management.

Viewing the compliance dashboard

The Cloud Workspace compliance dashboard provides insights that help security teams identify cloud resources without ownership information and flag vulnerable or at-risk resources. This information helps team members conduct security audits and confirm compliance with industry regulations across cloud environments.

Access the compliance dashboard

As an admin, you can access the dashboard by navigating to **All** > **Cloud Workspace** > **Monitor and track**.

Compliance dashboard

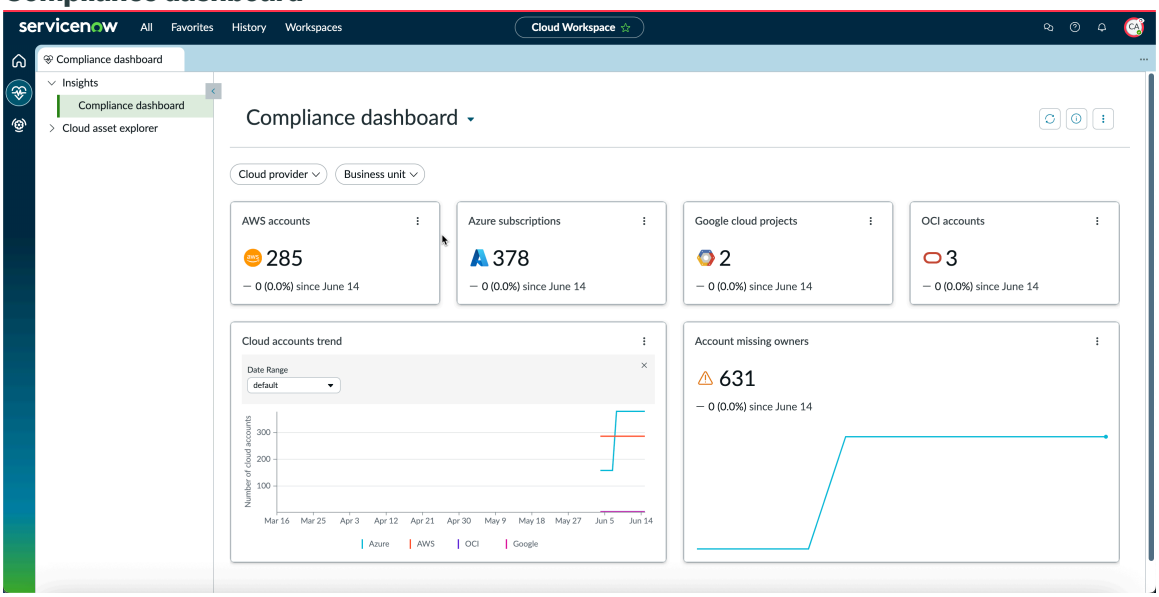

Note:

- The Compliance dashboard displays only the accounts discovered in your instance.
- The Compliance dashboard hides the provider card if no accounts are provisioned.

Viewing the cloud asset explorer

The dashboard helps cloud governance and operations teams monitor, track, and act on compliance and asset details more efficiently. As an asset viewer, you can view all assets and drill down to see detailed information for each configuration item (CI).

You can access the dashboard in either of two ways.

- Navigate to **All** > **Cloud Workspace** > **Cloud asset explorer**
- Navigate to **Workspaces** > **Cloud Workspace** > **Cloud asset explorer**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud asset explorer dashboard

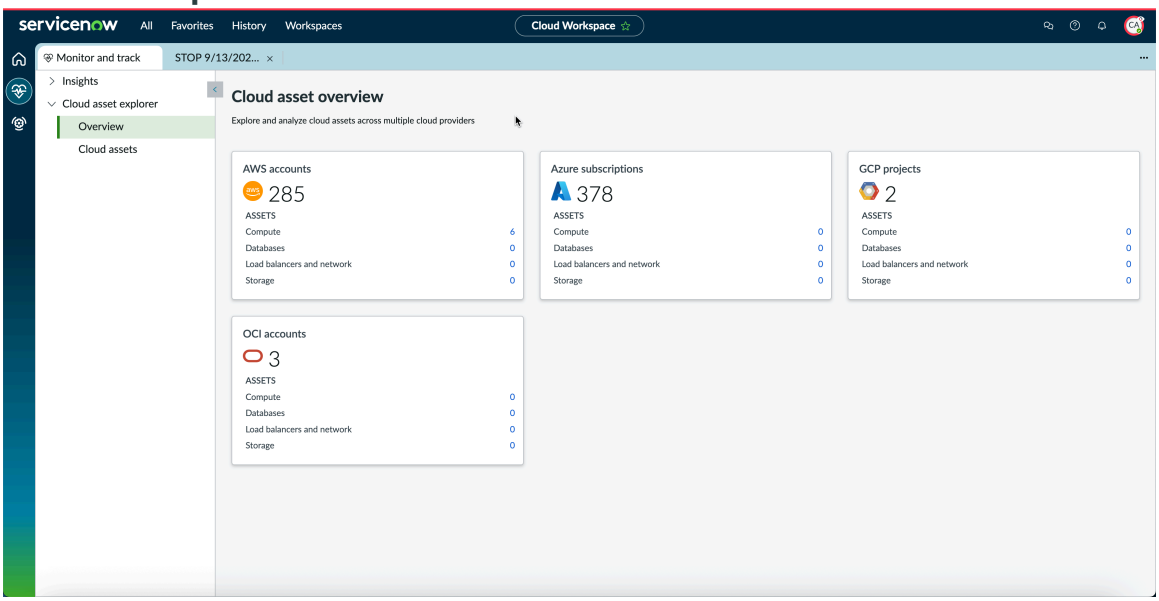

Note: You can enable or disable the Asset collection configuration to list the CIs. For more information, see Configure cloud asset collection table.

Managing the Cloud Account Management cloud asset details

Use the cloud asset form in the Cloud Account Management to view, edit, and manage cloud asset details such as attributes, health, and records.

The cloud asset form is a centralized, highly organized, and comprehensive set of sections with cloud asset details. Use this form to examine and edit cloud asset attributes, relationships, tags, services and offerings, and account health. You can also view data associated with the cloud asset, resources, activities, and related services.

The following sections describe the various sections of the cloud asset form.

Summary

The Summary section provides a consolidated view of key configuration item (CI) information for a Virtual Machine (VM) instance. This page helps you quickly access metadata, ownership details, operational status, discovery timestamps, and critical attributes required for monitoring and decision-making.

Attributes

The Attributes section displays all UI sections configured in the Explore CI view for the CI class. Select Attributes to access the available subsections. By default, the Discovery and Operational subsections align with the corresponding sections in the Explore CI view. The All subsection lists all CI attributes. Mandatory attributes are marked with an asterisk (*) and must have a value.

Tags

This section lists tags stored as key/value pairs in the Key Values [cmdb_key_value] table. These cloud tags are typically associated with CIs and resources from providers like AWS and Google Cloud Platform. Cloud tags are used in ITOM applications such as Tag Governance and are key to Tag-based discovery in Service Mapping. They are automatically populated by Discovery and Service Graph Connectors and can be grouped and managed as needed within your organization.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resources

This Resources section displays the resources linked to the CI and corresponds to the Related Lists defined in the Explore CI view.

Activity

This section shows related items such as open incidents and the activity stream for the CI on the following tabs:

Related Items

Pie charts that by default include Open incidents (Incidents table), Open changes (Change Request table), and Open problems (Problem table), broken down by priority or state.

CI history

History of the activities related to the CI, such as value changes in the CI attributes, listed in chronological order.

Relationships

This section shows the CI relationships, which include:

- Total relationships: Total number of relationships for this CI
- Duplicate relationships: Total number of duplicate relationship for this CI
- Stale relationships: Total number of stale relationship for this CI

Services and Offerings

This section shows cards with counts of the following services and offerings associated with the CI:

- Service Instance: Any CI related to the Service Instance [cmdb_ci_service_auto] table where Service Classification is Application Service.
- Business Service Offering: Any Service Instance, which has a relationship to Service_Offering where Service Classification is Business Service.
- Technology Management Offering: Any Service Instance that has a relationship to Service_Offering where Service Classification is Technical Service. Also includes CIs related to a Dynamic CI Group that is part of a Technical Service Offering.
- Technology Management Service: Parent of a Technology Management Offering.

Health

This section shows aggregations from the latest CMDB Health test results for the Completeness, Correctness, Compliance, and Relationship KPIs, and the aggregated state of attestation. Color codes are used to denote pass/fail scores.

Note: For more information on each section, see Manage CI details using CI Form in CMDB Workspace.

Viewing cloud assets

This page displays the Cloud assets view. It lists all compute assets across connected cloud providers such as AWS, Azure, GCP, and OCI. Details include location, cloud account, installation status, owner, discovered dates, ownership attestation, and cost center.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

As an admin, you can access the dashboard by navigating to **All** > **Cloud Workspace** > **Monitor and track**

Cloud assets

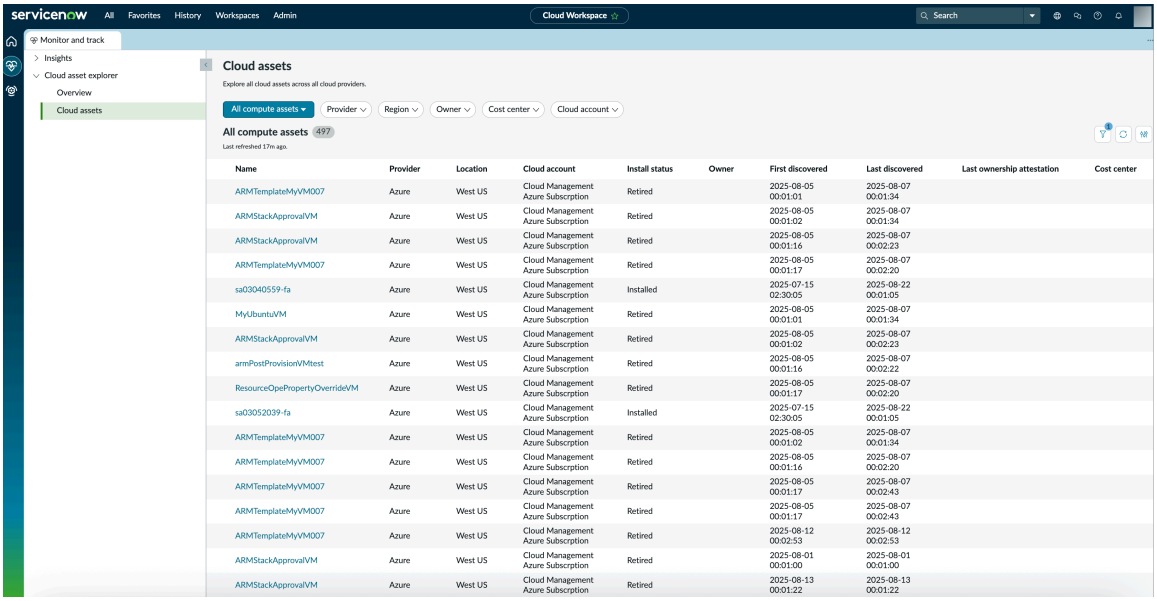

Click the respective name to view the details.

Note: Duplicate LDC entries no longer appear in the Region filter in Cloud Asset Explorer. After upgrading your instance to the latest version, the Region filter may be empty until the background job completes its initial run. Depending on data volume, this can take several hours. The job runs approximately every 30 minutes automatically.

View cloud asset details

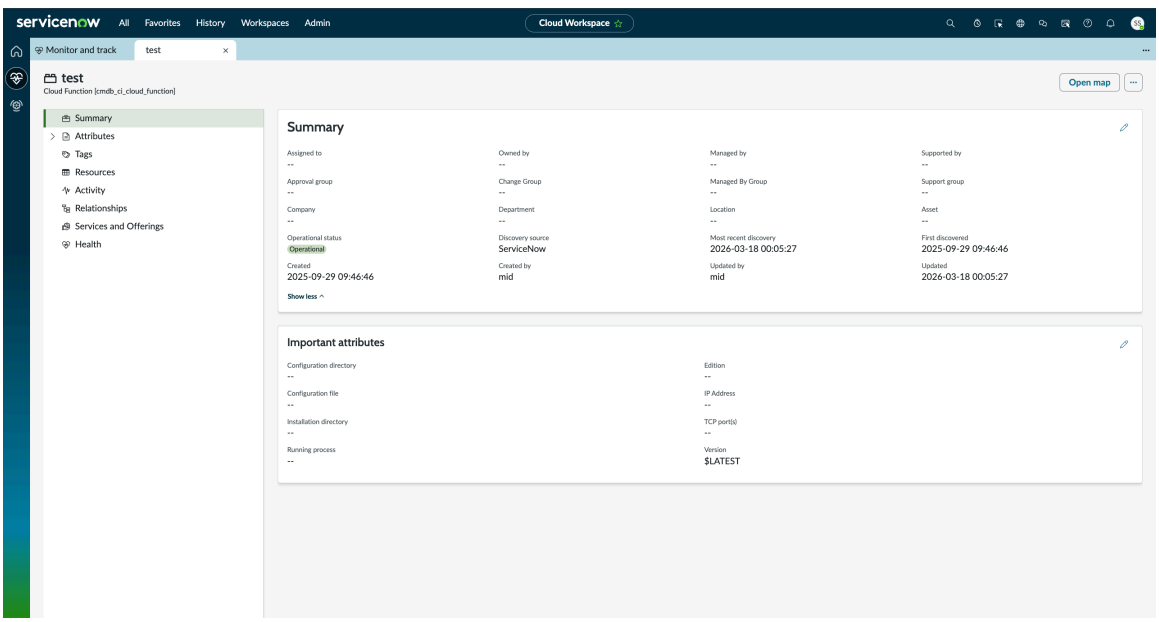

In Asset Explorer, retired configuration items (CIs) remain in the asset table as long as they exist in the CMDB, even after retirement. To manage this, users can configure a custom retention period (in seconds) that automatically removes retired CIs from the asset table after the defined duration. For example, setting the property to 86,400 seconds (1 day) ensures that a retired CI is removed from the asset table one day after retirement. The CI may continue to exist in the CMDB. This helps improve performance for high-volume environments by ensuring that only active CIs remain in the asset table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information about the various sections, see Managing the Cloud Account Management cloud asset details.

Viewing the cloud accounts page

The Cloud accounts page displays information used to access and manage the life cycle of multiple cloud subscription accounts.

Access the dashboard

As a Cloud Account Management (CAM) admin, you can access the Accounts page in either of two ways.

- Navigate to **All** > **Cloud Workspace**
-
- Navigate to **Workspaces** > **Cloud Workspace**

Cloud accounts page

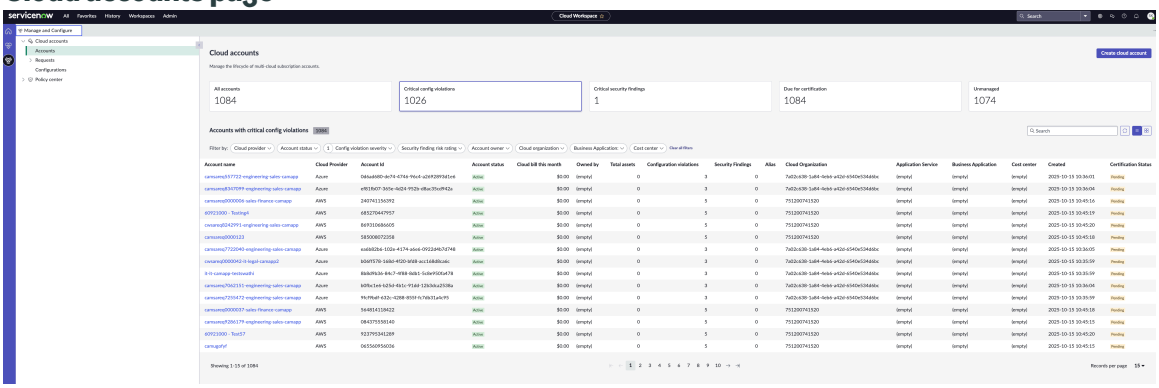

Note: The configuration violations data is populated only if the user is a CAM admin or account manager and the Cloud Configuration Governance content pack is installed.

The following metrics are displayed:

- **All accounts**: The total number of cloud accounts.
- **Critical config violations**: The number of cloud accounts with critical configuration violations.
- **Critical security findings**: The number of cloud accounts with critical security findings.
- **Due for certification**: The number of accounts due for certification.
- **Unmanaged**: The number of unmanaged accounts.

For more information, see Cloud account details.

Viewing cloud account records

The Cloud account records page is a comprehensive dashboard that provides information to help you manage a cloud account.

Before you begin

Role required: sn_itom_cam.cw_admin

Procedure

1. Access *`Cloud Workspace`*.
   As an admin, you can access the Cloud Accounts page using one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Navigate to **All** > *Cloud Workspace*
- Navigate to **Workspaces** > *Cloud Workspace*

2. Select **Manage and Configure**.
3. Select the account name to view the details of the corresponding account.
   The details provided for the account are listed in the following table.

Cloud account record dashboard tabs

Tabs	Description
Overview	This section provides the following critical information about the account: ○ **Total assets**: The total number of assets associated with the account. ○ **Last month spend**: Actual spend on your service accounts for the previous month. **Note:** This field appears only if the Cloud Cost Management workspace plugin is enabled. ○ **Critical configuration violations**: The number of critical configuration violations. ○ **Critical security findings**: The total number of critical security findings. ○ **Account details**: Critical information about the account including cloud provider, created date, alias, cost center, business unit, application service, department, cloud org, and cloud account group. ○ **People responsible for the account**: The individuals or teams responsible for managing the account such as the account owner, email group, and primary and secondary teams. ○ **Tags**: The discovered tag key-value pairs from the cloud provider. ○ **Certification**: The certification status, last certification date, last certified by, and certification due date of the account.
Assets	Assets discovered in the cloud account.
Configuration violations	Violations associated with the cloud account.
Security findings	Cloud account misconfigurations and critical host findings.
Costs (appears only if Cloud Cost Management Workspace plugin is enabled)	This section provides critical budgeting information for the cloud account: ○ **Last month spend**: Actual spend on your service accounts for the previous month. ○ **Monthly budget**: Valid budget of your cloud assets for that month. ○ **Monthly Variance**: Difference between the monthly budget and last month spends, indicating whether spending was over or under the approved budget.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tabs	Description
	○ **YTD annual spend**: The cumulative total expenditure from the start of the current year until date. ○ **Monthly spend chart**: A rolling 13-month visualization that plots actual spend against a constant budget line, enabling trend analysis over time.

Creating a cloud account

Creating a cloud account involves a requester submitting an account request, an approver approving or rejecting the request, and that admin provisioning the account.

Request a cloud account

Request a cloud account as a requester. A notification email is sent when the requester creates an account request.

Before you begin

The ServiceNow admin must have populated the data for the account owner, cost center, business unit, department, and business application.

Role required: sn_itom_cam.cw_requester

Procedure

1. Navigate to **All** > **Cloud Workspace**.
2. From the Home page, select **Request**.
3. On the Request new cloud account form, fill in the fields.

   For a description of the field values, see New cloud account request fields.

   Note: To create your custom catalog, see Configure a custom catalog ID in Cloud Account Management account request.

4. Select **Submit**.

Approve a cloud account request

Approve or deny a cloud account request as an approver. A notification email is received when the requester creates an account request. When a cloud account is requested, the table captures and retrieves the details of the request. Once the Cloud Account Management request is approved, it's assigned to a provisioner, such as CCOE, DevOps, or SRE.

Before you begin

Role required: sn_itom_cam.cw_approver

Procedure

1. Navigate to **All** > **Cloud Workspace** > **View accounts**.
2. Select **Requests** > **Pending approvals**.
3. Review the request and select **Approve** or **Deny**.
4. Confirm your decision.

Provision a cloud account

Provision a cloud account as an admin through Cloud Workspace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Confirm that the cloud context has been configured. For more information, see Creating configurations.

Discovery must have been performed on the cloud admin portal.

Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Requests**.
2. On the Account requests page, select the **Approved requests** tile.
3. Select the account to provision.
4. In the **Select cloud context** section, fill in the fields.

Field	Description
Cloud Provider	This field is automatically set to the cloud provider platform in the request.
CAM configuration	Configuration name of the cloud organization.
Organizational unit (appears for AWS)	Organizational unit (OU) are the list of units populated automatically after discovery is run.
Management group (appears for Azure)	Management groups help you manage multiple subscriptions by applying governance settings that automatically apply to all subscriptions within them. This field is automatically set to the management group found after discovery is performed.

5. If you are provisioning an Azure account, determine whether to use an existing billing account or create and add a new billing account.

Option	Action
Add an existing billing account	Select an existing account from the drop-down list.
Create and add a new single account	**a.** Select the Create new billing account icon (+). **b.** Select **Single account**. **c.** In the **Account name** field, enter the account name. **d.** In the **Display name** field, enter the display name. **e.** Select **Add**.
Create and add new bulk accounts	**a.** Select the Create new billing account icon (+). **b.** Select **Bulk upload**. **c.** Search for and select the Tenant name. **d.** Below the **Enter API response** field, select the link

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
	e. Sign in to your Azure portal. **f.** From the Azure portal, copy the sample JSON response that appears under the **EnrollmentAccountsListByBillingAccount** section. **g.** In the Cloud Account Management application, paste the sample JSON response in the **Enter API response** field. **h.** Select **Add**.

6. If you are provisioning an Azure account, determine whether to use an existing enrollment account or create and add a new enrollment account.

   An enrollment account is a management unit within a billing account for organizations with an Enterprise Agreement (EA). It organizes and controls Azure subscriptions and resources.

Option	Action
Enrollment account	Select an existing account.
Create and add a new single account	**a.** Select the Create new enrollment account icon (+). **b.** Select **Single account**. **c.** In the **Account name** field, enter the account name. **d.** In the **Display name** field, enter the display name. **e.** Search for and select the Tenant name. **f.** Select **Add**.
Create and add new bulk accounts	**a.** Select the Create new enrollment account icon (+). **b.** Select **Bulk upload**. **c.** Search for and select the Tenant name. **d.** Below the **Enter API response** field, select the link **e.** Sign in to your Azure portal. **f.** From the Azure portal, copy the sample JSON response that appears under the **EnrollmentAccountsListByBillingAccount** section. **g.** In the Cloud Account Management application, paste the sample JSON response in the **Enter API response** field. **h.** Select **Add**.

7. **Optional:** Provide a comment.
   a. In the **Work notes** section, enter your comments in the **Comments** text field.
   b. Select **Post Comment**.
8. At the top-right corner of the screen, select **Provision**.
9. In the dialog box, confirm the provisioning request by selecting **Yes**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

- Based on the defined context, Cloud Account Management creates a subscription account within the specified cloud organization and AWS organization unit or Azure management group
- When an Azure subscription is requested through Cloud Account Management, an alias is appended to the request name. The alias consists of an 8-character randomly generated suffix. For example: `CWSAREQ0000001-businessapp1-UEFwnlF8`.

Managing cloud accounts

Managing a cloud account enables you to accomplish a range of tasks such as canceling, suspending, or reactivating cloud accounts, and adding unmanaged accounts.

Account management activities

Activity	Description
Cancel a Cloud account request	Requesters can cancel their own account requests.
Suspend a Cloud account	Suspended accounts can still be used along with its resources, but new resources can't be created for it. Before requesting account suspension, the ServiceNow admin must have already set up a suspension configuration.
Reactivate a cloud account	Suspended accounts can be reactivated.
Add an unmanaged cloud account	You can import cloud accounts to manage them.

Cancel a Cloud account request

Cancel an account request that you have submitted.

Before you begin

Role required: sn_itom_cam.cw_requester

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Home**.
2. Select the **Requests** tab.
3. Select an account under the **Requested** pane.
4. Select **Cancel Request**.
5. Select **Yes** to confirm.

Suspend a Cloud account

Suspend an account when there's a budget constraint or the account owner isn't available in the organization. Only admins and account owners can see the suspended accounts.

Before you begin

An account requires a suspend profile to be suspended. For more information, see Set up suspension of an AWS account using service control policy.

Role required: sn_itom_cam.cw_admin or sn_itom_cam.cw_requestor

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Home**.
2. Select **Subscription Accounts**.
3. Under the **Managed Accounts** section, select **Active**.
4. Select an account check box and select **Suspend**.
5. In the confirmation dialog box, enter the reason for suspension and select **Yes**.
   - Once you have created a suspension request, the request is sent to the admin for approval.
   - Suspended accounts are listed under the **Managed Accounts** > **Suspended** section.

Reactivate a cloud account

Reactivate an account that was previously suspended if there is a need retrieve the account. After reactivating an account, the account status changes from suspended to active. Only Cloud Account Management managed accounts can be reactivated.

About this task

Only an admin or account owner can view the reactivated account.

Before you begin

A suspend profile must exist for an account to be reactivated. For more details, see Set up suspension of an AWS account using service control policy.

Role required: sn_itom_cam.cw_admin or sn_itom_cam.cw_requestor

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Home**.
2. Select **Subscription Accounts**.
3. In the **Managed Accounts** section, select **Suspended**.
4. Select an entry in the list of suspended accounts and select **Reactivate**.
5. Enter a reason for reactivation in the **Reason** field and select **Yes**.
   The reactivation request is sent for approval. Reactivated accounts are listed in the **Managed Accounts** section.

Add an unmanaged cloud account

Managing a cloud account involves tracking budgets, checking for configuration violations, and validating certificates. The procedure helps you to onboard an unmanaged cloud account.

Before you begin

Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Subscription accounts**.
2. In the **Unmanaged Accounts** section, select **All Accounts**.
3. Select an account and select **Manage Account**.

   Note: Currently, only AWS accounts are supported.

4. On the Request form, fill in Cost center, Department, Account owner, and Business Unit fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a description of the field values, see New cloud account request fields.

5. Select **Update Account Details & Manage**.
   A confirmation message appears saying your account is onboarded.

Update cloud account details

Keep account details up to date so you can identify and resolve policy violations effectively.

Before you begin

Role required: sn_itom_cam.cw_admin

Procedure

1. Navigate to **All** > **Cloud Workspace** > **View accounts**.
2. Under **All cloud accounts**, select an account.
   The **Overview** tab displays General details, Ownership details, and Certification details.
3. **Optional:** Update account information if necessary.

   a. Select the Edit icon ( ).

   b. Update the information.

Fields	Description
Cost center	Cost center name to associate with this account.
Business unit	Business unit name to associate with this account.
Application service	The project or application to associate with this account.
Department	Department name to associate with this account.

4. **Optional:** Update the account owner details if necessary.

   a. Select the Edit icon ( ).

   b. Select the account owner from the **Account Owner** from the drop-down list.

   c. Select the primary teams from the**Primary Teams** from the drop-down list.

   d. Select the secondary teams from the**Secondary Teams** from the drop-down list.

5. **Optional:** Propagate the name change to associated cloud assets or stop propagation.
   - Propagate the owner information to the cloud account and all associated CIs directly linked to the cloud account, overriding any existing owner information, by selecting the **Propagate Ownership to Cloud Assets** check-box.

     Associated CIs include items like S3 bucket or VM. CIs not directly linked to the cloud account include things like Azure Availability Zones.

   - To remove future propagation of the owner information, clear the **Propagate Ownership to Cloud Assets** check-box.

     The owner name is no longer assigned to newly associated CIs but is retained for the existing ones.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

If you have a large number of CIs, the propagation might take some time and changes might not be reflected immediately. You can check the progress by accessing the sn_itom_cam_subscription_account_batch_job_list table and checking the status of "owner_name" jobs listed in the **Parameters** column.

6. Select the Save icon ( ).

Certify an account

A certification task represents the work of verifying and certifying the data associated with a record. The certifier reviews tasks created after data validation against the published policy. Account certification helps resolve issues like outdated ownership, inactive user profiles, and improper permissions in cloud accounts.

Before you begin

Make sure that the admin has reviewed the Cloud Account Management default policy. Optionally, the admin can customize or add new policies based on the organization requirement. For more information, see Review default Cloud Account Management certification policy. The ITOM Cloud Account Management Certifier Group must contain at least one user.

Note: Cloud Account Management admin can add users to the Cloud Account Management Certifier Group by navigating to **All** > **Account Certification** > **ITOM CAM Certifier Group** > **Group Members**.

Role required: sn_itom_cam.cw_certifier

About this task

Procedure

1. Navigate to **All** > **Cloud Workspace** > **Account Certification Tasks**.

   The list under Certification Tasks shows pending certification tasks.

2. In the Certification tasks tab, select the certification task that you want to review.
3. Review the details of the certification task, then select **Review records to certify**.
4. Select the records that you want to conclude the review for in the Review not completed tab, and then select **Certify** or **Fail**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Account Management certification review

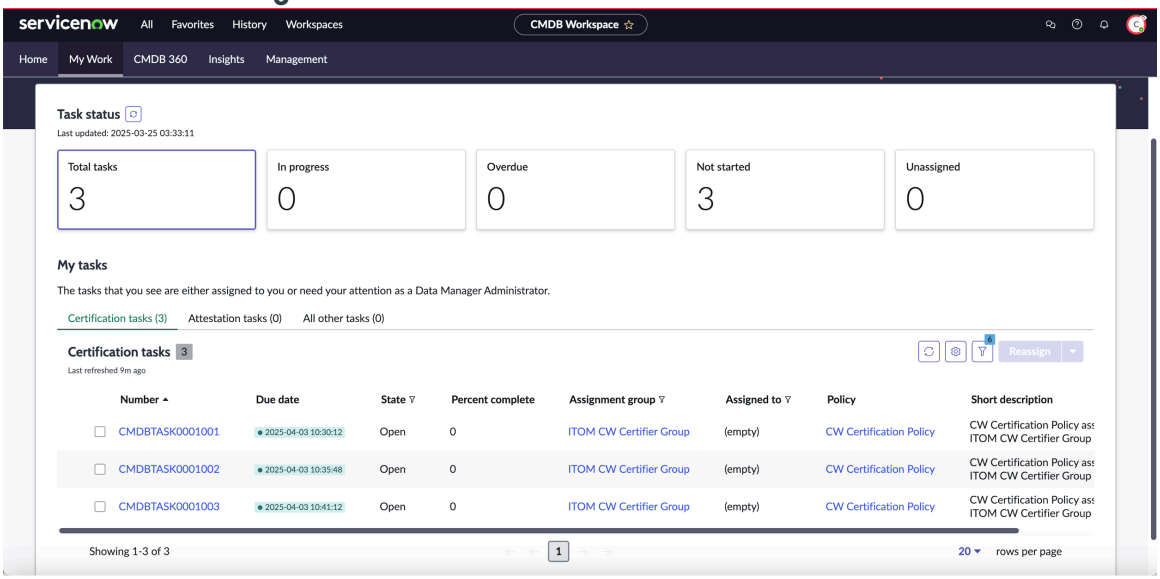

Note: If you select **Fail**, the task is moved to the registry and Cloud Account Management admin will take the next action. You can also reassign the ownership of a task.

5. Select the records that you want to review in details and then select **Review records**.
6. Select **Submit** when all the records have been reviewed.
   The **Submit** button isn't available until all the records in the task are reviewed and the task is complete.
7. **Optional:** Assign owners (if they're empty) by select and hold (or right-click) on the **Owned By** entry and selecting the name from the lookup.

Note:

The account certification state changes to pending if it meets the admin's policy conditions for recertification.

For more information on certifying a task, see Review certification tasks.

Cloud Account Management reference

Reference topics provide additional information for configuring and using the Cloud Account Management in Cloud Workspace application.

Components installed with Cloud Account Management

Several user roles are installed with the Cloud Account Management app.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles installed

Role	Description	Group that provides the role
Requester [sn_itom_cam.cw_requester]	Team lead or representative of the application. • Initiates cloud account creation requests. • Create requests for suspension or reactivation of their own accounts.	ITOM CAM Requester Group
Approver [sn_itom_cam.cw_approver]	Supervisor, manager, or finance approver. Reviews requests and either approves or denies them.	ITOM CAM Approver Group
Admin [sn_itom_cam.cw_admin]	Cloud Center of Excellence (CCOE) or Sites Reliability Engineer (SRE) team member. • Confirms that Cloud Account Management configurations align with the cloud and Terraform configurations. • Customizes the data certification policy. • Creates accounts that are outside the Cloud Account Management application.	ITOM CAM Admin Group
Certifier [sn_itom_cam.cw_certifier]	Marks an account as certified or failed.	ITOM CAM Certifier Group

To return to the main topic, see Exploring Cloud Account Management.

Terraform account configuration fields

Terraform account configuration form and related list field descriptions.

Configuration type: Create account

Field	Description
Provision mode	Name of the provision mode. Here it should be Terraform.
Cloud provider	The name of the cloud provider platform. Currently, only AWS is supported.
Cloud organization	The cloud organization where the subscription account is created. For more details, see Populate Cloud Account Management dependent tables.
Name	A name to identify the configuration.
Root email	Cloud Account Management uses an email address to generate the root email ID for newly created account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration type: Create account (continued)

Field	Description
	For example, if your root email ID is myawsaccount@example.com, the Cloud Account Management application appends a request ID to generate myawsaccount-CWSAREQ0000001@example.com. **Note:** To create a root email, see Set up and verify root email in AWS.
Terraform URL	A default URL appears. You can customize the URL by selecting the edit icon .
Terraform OAuth Token ID	The OAuth token ID provided when you configured the VCS provider in Terraform.
Terraform credentials	The Terraform credentials that you had configured. The entries are populated from the CMDB records.
Terraform organization	The Terraform organization name used as a workspace to create a subscription.
VCS identifier	GitHub (or Version Control System (VCS)) location of the Terraform template.
Description	Option to provide a description for the Terraform account.

To return to the procedure, see Create a Terraform account configuration.

Cloud native interface configuration account fields

Cloud native interface configuration account form and related list field descriptions.

Configuration type: Create account

Field	Description
Provision mode	Name of the provision mode. Here it should be Cloud native interface.
Cloud provider	Name of the cloud provider platform. Currently supports AWS and Azure platforms.
Cloud organization	The cloud organization where the subscription account is to be created.
Name	A name to identify the configuration.
Root email (AWS)	Cloud Account Management uses an email address to generate the root email ID for newly created account. For example, if your root email ID is myawsaccount@example.com, the Cloud Account Management application appends a request ID to generate myawsaccount-CWSAREQ0000001@example.com. **Note:** To create a root email, see Set up and verify root email in AWS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration type: Create account (continued)

Field	Description
Cloud credential alias	Alias created to access all the Cloud APIs. The entries are populated from the CMDB records. **Note:** Only Cloud Account Management admin can configure the cloud credential alias. To configure the cloud credential alias, see Set up AWS API configuration information in ServiceNow.
Description	Information that can be recorded in work notes.

To return to the procedure, see Create a cloud native interface account configuration.

New cloud account request fields

New cloud account request form and related list field descriptions.

Request a Cloud Account

Fields	Description
Environment	Name of the environment.
This is a GovCloud account (applicable only for AWS)	Check-box to request an AWS GovCloud account.
Desired cloud provider	Name of the cloud provider platform such as AWS or Azure.
Cloud account name suffix	Suffix a name for easier tracking.
Cloud account name	Auto-generated name derived from the cost center, business unit, business application, and suffix values. Special characters are automatically removed.
Monthly Budget Limit	Monthly budget for this account. **Note:** If the budget threshold is reached, an email alert is sent to the root user of the AWS organization account and owner of the newly created account. Even if the threshold is exceeded, there are no restrictions in accessing the account but you can't create resources for the suspended account.
Project/ application service	The project or application to associate with this account.
Cost center	Cost center to associate with this account.
Business Unit	Business unit to associate with this account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request a Cloud Account (continued)

Fields	Description
Department	Department to associate with this account.
Business application	Business application to associate with this account.
Tags	Key-value information to categories the account. For example, by region, country, organization, or a branch.
Account owner	Owner of this account.
Propagate ownership to cloud assets.	The account owner information is propagated to all assets for this account.
Primary teams	The teams that will manage the account and assets, acting as the responsible party when the owner is absent.
Secondary teams	The secondary team that will support the account.

Note: For AWS and AWS GovCloud accounts, the values entered in the **Cost center**, **Business Unit**, **Department**, **Project/Application service**, **Business application**, and **Environment** fields are synced to AWS as tags. AWS restricts tag values to letters, numbers, spaces, and the following special characters: `+ - = . _ : / @`. Any other characters are automatically replaced with spaces. If all characters in a field value are unsupported, the tag is skipped. A tooltip on the form indicates this behavior. These actions are recorded in the execution log.

To return to the procedure, see Request a cloud account.

Suspend account configuration fields

Suspend the account configuration form and related list field descriptions.

Configuration type: Suspend account

Field	Description
Cloud provider	Name of the cloud provider platform. The entries are retrieved from the CMDB records.
Cloud organization	The cloud organization that creates the subscription account.
Name	A name to identify the configuration.
Cloud credential alias	The cloud credential alias that you had configured. The entries are populated from the CMDB records.
Suspension SCP policy ID (applicable only for AWS)	The AWS suspension policy ID. This policy ID is used to suspend the AWS account. Service control policies limit the resources, services, and API operations. To configure AWS suspend policy, see Set up suspension of an AWS account using service control policy.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration type: Suspend account (continued)

Field	Description
Azure suspension policy ID (applicable only for Azure)	The Azure suspension policy ID. This policy ID is used to suspend the Azure account. Service control policies limit the resources, services, and API operations. To configure Azure suspend policy, see Set up suspension of a subscription using Azure policy.
Description	Information that can be recorded in work notes.

To return to the procedure, see Suspend a Cloud account.

Policy list for scanning cloud accounts

A list of default policies provided for scanning the cloud accounts.

Default policies for scan accounts

Policy Name	Description
Check AWS Discovery Schedule	Verifies whether an AWS account has a discovery schedule attached. Running discovery regularly helps facilitate the identification and management of potential security risks.
Check AWS Account Alias	Verifies an AWS account has a unique alias to improve account management, reduce errors, and promote clarity and traceability within your AWS infrastructure. **Note:** Make sure you have API permission for `iam:ListAccountAliases`.
Check AWS Account Owner Tag	Verifies whether an AWS account has a designated owner to enhance accountability, streamline incident response, and facilitate communication within your AWS environment.
Check AWS Custom Password Policy	Verifies whether a custom password policy is set for every AWS account. A robust password requirement for all IAM users significantly increases the difficulty for attackers to crack passwords through brute-force attacks or credential theft attempts, ultimately enhancing the overall security of your AWS infrastructure. **Note:** Make sure you have API permission for `iam:GetAccountPasswordPolicy`.
Check AWS Failed Certification	Verifies the AWS account certification status. Failed certifications indicate potential security vulnerabilities because compromised credentials might not be deactivated promptly and provide a window of opportunity for attackers to exploit these weaknesses.
Check AWS Pending Certification	Verifies whether an AWS account certification is in a pending state to enable the prompt resolution of pending certifications and avoid potential security vulnerabilities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Policy Name	Description
Check AWS Strong Password Policy	Verifies whether an AWS account adheres to a strong password policy to promote security. This policy mandates robust password complexity requirements, significantly bolstering your AWS environment's defense against unauthorized access. **Note:** Make sure you have API permission for `iam:GetAccountPasswordPolicy`.
Check Azure Discovery Schedule	Verifies whether Azure account has a discovery schedule attached. This policy helps secure and up-to-date resource landscape to facilitate the identification and management of potential security risks.
Check Azure Account Owner Tag	Verifies if Azure account has a designated owner tag to enhance accountability and facilitate communication within your Azure environment. This policy readily identifies the responsible party for each account, promoting a culture of ownership and streamlined incident response.
Check Azure Failed Certification	Verifies Azure account certification status for failure to promote strong access control by proactively monitoring for any service account with a failed certification status. Failed certifications indicate potential security vulnerabilities, as compromised credentials might not be deactivated promptly. This policy minimizes the window of opportunity for attackers to exploit these weaknesses.
Check Azure Pending Certification	Verifies if Azure service account certification is in a pending state. This monitoring enables prompt resolution of pending certification and avoids potential security vulnerabilities.
Check GCP Discovery Schedule	Verifies whether GCP account has a discovery schedule attached. This policy helps secure and up-to-date resource landscape to facilitate the identification and management of potential security risks.
Check GCP Account Owner Tag	Verifies if GCP account has a designated owner tag to enhance accountability and facilitate communication within your GCP environment. This policy readily identifies the responsible party for each account, promoting a culture of ownership and streamlined incident response.
Check GCP Failed Certification	Verifies GCP account certification status for failure to promote strong access control by proactively monitoring for any service account with a failed certification status. Failed certifications indicate potential security vulnerabilities, as compromised credentials might not be deactivated promptly. This policy minimizes the window of opportunity for attackers to exploit these weaknesses.
Check GCP Pending Certification	Verifies if GCP account certification is in a pending state. This monitoring enables prompt resolution of pending certification and avoids potential security vulnerabilities.

To return to the procedure, see Set up scan configuration for data visualization.

Cloud Account Management ACL groups, roles, and responsibilities

Access control lists (ACLs), groups, and roles in Cloud Account Management control how access permissions are organized and managed within a cloud environment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Groups and roles in Cloud Account Management

Cloud Account Management groups and responsibilities

Group	Role	Responsibilities
ITOM Cloud Account Management requester group	sn_itom_cam.cw_requester	• Initiates cloud account requests: Triggers the process for creating cloud accounts. • Justifies account creation: Provides clear and compelling reasons for why each cloud account is needed. • Gathers and submits account information: Collects and submits all the necessary details required for successful cloud account creation. • Manages account status: Request suspension or reactivation of cloud accounts as needed.
ITOM Cloud Account Management approver group	sn_itom_cam.cw_approver	• Evaluates and approves account requests: Review account creation requests, carefully assessing their validity and alignment with budget constraints. Following this evaluation, the administrator grants approval or denial of the request. • Verifies budgetary compliance: Confirms that the requested cloud account aligns with available budgetary resources before granting approval. • Communicates decisions with clear rationale: Communicates clearly the approval or denial decision to the requester by adding a comment to the request itself. This comment should provide a concise and informative explanation for the decision.
ITOM Cloud Account Management admin group	sn_itom_cam.cw_admin	• Verifies configuration alignment: Verifies that Cloud Account Management configurations are consistent and aligned with both cloud provider configurations and Terraform configurations. This process promotes secure access management practices across all environments.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Account Management groups and responsibilities (continued)

Group	Role	Responsibilities
		• Provisions cloud accounts: Configures the cloud context to provision new cloud accounts, streamlining the account creation process. • Troubleshoots provisioning issues: Diagnoses and resolves any problems encountered during cloud account provisioning, promoting successful account creation. • Customizes data certification policy: Tailors the Cloud Account Management default data certification policy to meet specific organizational requirements, promoting appropriate data handling procedures. • Manages all accounts: Oversees all cloud accounts, including CW-managed accounts and on-boarded accounts, guaranteeing centralized control and visibility.
ITOM Cloud Account Management certifier group	sn_itom_cam.cw_certifier	• Certifies discovered accounts: Reviews discovered accounts listed in the Subscription Accounts tab and designates them as certified if they meet compliance standards or failed if they don't. This process verifies proper vetting and ongoing monitoring of cloud accounts. • Manages verification tasks: Accesses and reviews verification tasks associated with cloud accounts. Following review, the administrator can certify tasks to indicate their successful completion.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How they work together

Cloud Account Management groups and responsibilities

Group	Role	Responsibilities
ITOM Cloud Account Management requester group	sn_itom_cam.cw_requester	• Initiates cloud account requests: Triggers the process for creating cloud accounts and provides the information for creating the account. • Manages account status: Requests suspension or reactivation of cloud accounts.
ITOM Cloud Account Management approver group	sn_itom_cam.cw_approver	• Evaluates and approves account requests. • Verifies budgetary compliance of requested cloud accounts and alignment with available budgetary resources.
ITOM Cloud Account Management certifier group	sn_itom_cam.cw_certifier	• Certifies discovered accounts. • Accesses and reviews verification tasks associated with cloud accounts.
ITOM Cloud Account Management admin group	sn_itom_cam.cw_admin	• Verifies that Cloud Account Management configurations are consistent and aligned with both cloud provider configurations and Terraform configurations. • Provisions cloud accounts. • Troubleshoots provisioning issues: Diagnoses and resolves any problems encountered during cloud account provisioning. • Customizes data certification policies to meet specific organizational requirements. • Oversees all cloud accounts, including CW-managed and on-boarded accounts, guaranteeing centralized control and visibility.

To return to the main topic, see Exploring Cloud Account Management.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Permissions required for Azure Service Principal

This table provides the permissions needed to create, close or cancel an Azure subscription, download billing details, and tag subscriptions.

Roles required for Azure Service Principal

Role	Action	Role definition ID
EnrollmentReader	Enrollment readers can view data at the enrollment, department, and account scopes. This data includes charges for all subscriptions under these scopes, even across tenants, and displays the AzurePrepayment (previously called monetary commitment) balance associated with the enrollment.	24f8edb6-1668-4659-b5e2-40bb5f3a7d7e
DepartmentReader	Download the usage details for the department you administer. You can also view the usage and charges associated with your department.	db609904-a47f-4794-9be8-9bd86fbffd8a
Microsoft.Billing/billingAccounts/read	Read the list of billing accounts.	
Microsoft.Management/managementGroups/subscriptions/write Microsoft.Management/managementGroups/write	Move subscription to the appropriate location once created	
Microsoft.Resources/tags/write	Add tags to the subscription.	
Microsoft.Billing/billingAccounts/billingSubscriptions/cancel/write	Close or cancel a subscription.	

Cloud Asset Collection table

Stores metadata of cloud assets collected from providers for IT operations.

A Cloud Asset Collection Table is a structured repository in IT Operations Management that stores and manages metadata of cloud resources collected from multiple cloud service providers. The table stores all discovered or synced cloud assets, such as virtual machines, storage buckets, and databases after collection from cloud environments.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Asset Collection Configuration fields

Field	Description
Active	Indicates whether this specific cloud asset collection configuration is enabled or turned off. • True – The configuration is active for the selected CI. It has been processed and populated in the flat table, and assets are collected according to the defined schedule. • False – The configuration is inactive, and no assets are collected.
Category	High-level classification of collected cloud assets, such as Compute, Storage, Networking, Database, or Security.
Subcategory	Granular classification under the main category. Examples include Virtual Machines, Containers, and Functions under Compute; Block Storage, Object Storage, and File Storage under Storage.
CI Class Name	CMDB class where collected assets are stored such as `cmdb_ci_vm_instance` for virtual machines.
Provider	Cloud service provider from which assets are collected.
Last Processed Time	Timestamp of the last successful Asset Collection run. Useful for: • Tracking data freshness • Troubleshooting delays or failures
Comments	Free-text field for notes or context about the configuration.

Cloud account details

The Cloud account details table provides detailed information about each account.

Cloud account details table

Entry	Description
Account name	The name of the cloud account.
Cloud Provider	The cloud service provider, such as AWS.
Account Id	The unique identifier for the account.
Account status	The current status of the account. The available statuses are: • Active • Suspended • Closed
Cloud bill this month	The cloud expenditure for the current month.
Owned by	The owner of the account.
Total assets	The total number of assets associated with the account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud account details table (continued)

Entry	Description
Configuration violations	The number of configuration violations.
Security Findings	The total number of security findings.
Alias	An alias for the account.
Cloud Organization	The cloud organization associated with the account.
Application Service	The application service associated with the account.
Business Application	The business application associated with the account.
Cost center	The cost center associated with the account.
Created	The date and time when the record was created in the ServiceNow® instance.
Certification Status	The certification status of the account. The possible statuses are: • Pending (yellow) • Failed (red) • Certified (blue)

Cloud Services Catalog

With the ServiceNow® Cloud Services Catalog application, you can use a simplified interface to access cloud resources. You can publish cloud offerings to a catalog and manage the usage and life cycle of those resources.

Important: Starting with the Brazil release, Cloud Services Catalog is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

Integrated with cloud providers, including Amazon Web Services, Microsoft Azure, and Google Cloud Platform, Cloud Services Catalog (CSC) enables you to use the Azure DevOps and CI-CD Pipeline Integration. You can integrate Ansible with Cloud Services Catalog as a Day 1 activity to deploy an Ansible job template via a ServiceNow catalog item. Use the multi-repo support with Terraform Connector IaC Discover capabilities, for improved hierarchical structure of multiple repositories.

Note: Cloud Services Catalog supports self-hosted environment. Make sure that the MID Server has the internet access. For more information, see the Deployment Guide for the self-hosted environment [KB0563731] article in the Now Support Knowledge Base.

AWS, Azure and Google Cloud Platform Out Of Box Catalogs are available with this release of Cloud Services Catalog. You can either order one of the catalog items as demonstrated in this video, or create one.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore 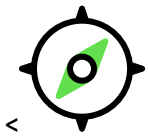 Explore and learn how the Cloud Services Catalog application provides the single place where you can automate cloud management tooling.	Configure Configure your Cloud Services Catalog and integrate it with other applications.	Use Cloud Services Catalog Use the Cloud Services Catalog application for searching and requesting cloud services, working with your stacks and resources, and so on.
Explore My Stacks Explore My Resources 	Order a Catalog Use the Cloud Services Catalog application to order from a variety of Out Of Box Catalog items, actions, and policies.	Reference Get more details about the permissions and workarounds.

Exploring Cloud Services Catalog

The Cloud Services Catalog application is the single place that you have to go to for automating cloud management tooling. It replaces and improves on the capabilities that were offered by Cloud Provisioning and Governance user and admin portals.

Cloud Services Catalog overview

The Cloud Services Catalog Request page on Employee Center enables you to request cloud management actions. For example, you can request provisioning (Day 1), patching and maintenance actions (Day 2), or life-cycle operations like starting, stopping, or deprovisioning.

Cloud Services Catalog benefits

The following table lists the benefits of Cloud Services Catalog.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefits of using Cloud Services Catalog

Benefit	Feature
Assistance for users	Guided Setup
Out Of Box content for Amazon AWS Cloud request and automation	AWS Out Of Box Catalogs
Out Of Box content for Microsoft AzureMicrosoft Azure Cloud Services request and automation	Azure Out Of Box Catalogs
Out Of Box content for Google Cloud Platform request	Out Of Box Catalogs using Cloud Services Catalog
Requests for process and pages for ordering Cloud Services	Cloud Services Request process
My Stacks management page for viewing and operating on stacks	Stacks Management
My Resources management page for viewing and operating on resources	Resources Management
Azure DevOps Continuous Integration Continuous Deployment (CICD) Pipeline Integration	Azure DevOps and CI-CD Pipeline Integration
Integration of Ansible with Cloud services	Integrate Ansible
Microsoft Azure and Amazon AWS Cloud setup	Cloud Setup
Multi-repo support with Terraform Connector app	Exploring Cloud Services Catalog Terraform Connector

What to do next

See Plugins or applications installed with ITOM Cloud Accelerate to find the list of plugins you activate with Cloud Services Catalog.

Configuring Cloud Services Catalog

You can configure the Cloud Services Catalog application so that so can begin to automate cloud management tooling.

Configuration overview

The configuration is in a sequential order where you create and set up the prerequisites first, and then integrate the credentials with your Cloud account profile. These setup steps are identified as Day 1 and Day 2 tasks.

The Cloud Services Catalog Day 1 configuration includes tasks like creating MID Servers, setting up your cloud providers, Amazon Web Services, or Microsoft Azure Cloud. You can also associate it with your cloud account, and schedule the datacenter Discovery.

An overview of the Day 1 setup process is outlined here:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Request the Cloud Services Catalog application, if needed.
2. Assign the appropriate roles to your cloud users.
3. Set up the MID Servers that handle the secure communications with the provider API endpoints.
4. On the provider portal, collect your account settings and the credentials that the Discovery process uses (through a MID Server) to access your provider accounts.
5. Securely associate the account settings and credentials with a service account in Cloud Services Catalog.
6. Set up a cloud account to represent your entire managed cloud infrastructure, and set up a service account that works with one of your provider accounts.
7. Specify the datacenters in the service account that should be included in the cloud account.

On Day 2, you can set up the additional cloud accounts and service accounts from the same or other providers.

1. Use the Configuration Management Database (CMDB) data to help you request and manage the cloud resources and manage your cloud infrastructure.
2. Manually run the Discovery process on each datacenter in the service account to populate the CMDB with the resource data for all datacenters.
3. Configure Discovery to run on a regular schedule to ensure that the data remains updated.

Using Guided Setup to implement Cloud Services Catalog

Guided Setup provides a sequence of tasks that help you configure CSC on your ServiceNow instance. To open Guided Setup for Cloud Services Catalog, navigate to **All** > **Cloud Services catalog- Guided Setup** on your instance.

For more information, see Guided Setup.

Install Cloud Services Catalog

You can install the Cloud Services Catalog application purchased from the ServiceNow® Store, to make it available on your instance.The application also includes demo data and installs related to ServiceNow applications and plugins, if you don't have them installed already.

Important: Starting with the Brazil release, Cloud Services Catalog Content Pack is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

Before you begin

- Cloud Services Catalog application from the ServiceNow® Store.
- Ensure that the application and all of its associated ServiceNow® Store applications have valid ServiceNow entitlements. See Get entitlement for a ServiceNow product or application.

Role required: admin, sn_appclient.app_client_company_installer (can only install applications that match the instance company), or sn_appclient.app_client_user

About this task

The following items are installed with Cloud Services Catalog:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Plugins
- Out Of Box Catalog
- Roles

Tip:

If you're an existing user of Cloud Services Catalog or CSC Innovation Labs:

- An upgrade from the previous version isn't currently supported. To upgrade an existing instance with Cloud Services Catalog or CSC apps, do the following actions:
- ○ Uninstall the previously installed CSC apps with table clean-ups (CSC content pack and CSC in order).
  - ○ Install new CSC apps (CSC and CSC content pack in the same order).
  - ○ You can run Guided Setup only once in an instance. Manually place new Terraform templates in the Terraform box, update the Out Of Box Terraform config provider accordingly, and run Discovery.

For more information, see the components links.

Procedure

1. Navigate to **All** > **System Application** > **All Available Applications** > **All**.
2. Search for the application by its name or ID.
   If you can't find the application, you might have to request it from the ServiceNow® Store. CSC is available as a separate subscription and requires the Cloud Provisioning and Governance plugin (com.snc.cloud.mgmt).
3. Select a version from the list and select **Install**.
   In the Install dialog box that is displayed, any dependencies that are installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow® Store to get any additional entitlements for dependencies.
5. Select **Install**.

Cloud Services Catalog setup for the Day 1 configuration task

You're ready to do the Day 1 configuration task of setting up the Cloud Services Catalog application for the first time. You can do the optional Day 2 setup and configure procedures as needed, in any order.

You must do the Cloud Services Catalog setup steps in order. On Day 1, you set up the cloud service for access, security, and billing purposes for Amazon Web Services and Microsoft Azure Cloud.

Before you can achieve these tasks, you must integrate Cloud Services Catalog with your Amazon Web Services or Microsoft Azure Cloud and progress to Day 2 configuration tasks.

When you finish all the Day 1 and Day 2 tasks, see the Cloud Services Catalog administration guide for more information on using the Cloud Services Catalog application in your organization.

Set up Amazon Web Services on Cloud Services Catalog

Integrate Cloud Services Catalog with your Amazon Web Services account as a part of the Day 1 configuration task of setting up the application for the first time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Collect your account settings and credentials from the provider portal. These settings and credentials help you to distinguish the resources in one account from another account and sign in or authenticate to AWS. Permission policies specify your approvals within the resources in the account.

Role required:

- Administrator for operation in AWS or AWS Management Console
- sn_cmp.cloud_admin role for operations in Cloud Services Catalog

Procedure

1. Assign roles to AWS users of Cloud Services Catalog and to individual users by their activities and responsibilities.
   For more information, see Assign roles to AWS users of Cloud Services Catalog.
2. Install and configure MID Servers for Cloud Discovery and Cloud Provisioning and Governance.

   To ensure secure and reliable communications, the Discovery process communicates with your cloud provider accounts and cloud resources through one or more MID Servers. For more information, see Install and configure MID Servers for Cloud Discovery and Cloud provisioning and governance.
3. Set up the MID Server on your network or in one of your cloud networks and configure it for Cloud Services Catalog to use a proxy server.
   Using a proxy server supports all cloud-based activities, such as running Discovery, billing downloads, provisioning virtual machines, and running life-cycle operations on virtual machines.
4. Configure one or more MID Servers for each datacenter for high performance and security. Data is encrypted to the MID Server and between the MID Server and the API endpoint.
5. Configure the MID Server even if you have already configured other MID Servers while you were setting up Cloud Services Catalog for another cloud provider.
6. Configure access to the AWS accounts by using permanent AWS credentials.

   - To securely access data on your provider account, the Discovery process must present the appropriate credentials.
   - To make the credentials available to Discovery and Cloud Services Catalog, you first create a user with programmatic access to the AWS Management Console.
   - You can securely store the credentials in a service account on the ServiceNow AI Platform®.
   - For more information, see Configure access to the AWS accounts using permanent AWS credentials.
7. Set up a cloud account and a service account for AWS.

   A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get the data on each resource in each specified datacenter. A cloud account is the logical representation in Cloud Services Catalog of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts, including the service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account. For more information, see Set up a cloud account and service account for AWS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Setting up Google Cloud Platform on Cloud Services Catalog

Integrate Cloud Services Catalog with your Google Cloud Platform account as a part of the Day 1 configuration task toward setting up CSC for the first time.

Before you begin

Collect your account settings and credentials from the provider portal. These settings and credentials help you to distinguish the resources in one account from another account and sign in or authenticate to Google Cloud. Permission policies specify your approvals within the resources in the account.

Role required:

- Administrator for operation in Google Cloud or Google Cloud Management Console
- sn_cmp.cloud_admin role for operations in CSC

To learn more about setting up Google Cloud with CSC Terraform Connector app for the first time, see Day 1 setup guide for Google Cloud through Cloud Services Catalog Terraform Connector.

Procedure

1. Assign roles to Google Cloud users of CSC and to individual users by their activities and responsibilities.

   For more information, see Assign roles to Google Cloud Platform users.

2. Install and configure MID Servers for Cloud Discovery and Cloud Provisioning and Governance.
   To ensure secure and reliable communications, the Discovery process communicates with your cloud provider accounts and cloud resources through one or more MID Servers. For more information, see Install and configure MID Servers to access cloud environments.
3. Set up the MID Server on your network or in one of your cloud networks and configure it for CSC to use a proxy server.
   Using a proxy server supports all cloud-based activities, such as running Discovery, billing downloads, provisioning virtual machines, and running life-cycle operations on virtual machines.
4. Configure one or more MID Servers for each datacenter for high performance and security. Data is encrypted to the MID Server and between the MID Server and the API endpoint.
5. Configure the MID Server even if you have already configured other MID Servers while you were setting up Cloud Services Catalog for another cloud provider
6. Create Service Account and the keys (credentials) in Google Cloud if you don not have it. You can securely store the credentials in a service account. You would need Service Account Key Admin (roles/iam.serviceAccountKeyAdmin) IAM role on the project, or the service account whose keys you want to manage. These permissions can be provided by administrator. To securely access data on your provider account, the Discovery process must present the appropriate credentials.

   For more information, see Set up Google Cloud Platform service accounts.

7. Set up a cloud account and a service account for Google Cloud Platform.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get the data on each resource in each specified datacenter.
- A cloud account is the logical representation in Cloud Services Catalog of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts, including the service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

For more information, see Set up a cloud account and service account for Google Cloud Platform.

Set up Microsoft Azure Cloud on Cloud Services Catalog

Integrate Cloud Services Catalog with your Microsoft Azure account as a part of the Day 1 configuration task toward setting up CSC for the first time.

Before you begin

Collect your account settings and credentials from the provider portal. These settings and credentials help you to distinguish the resources in one account from another account and sign in or authenticate to Microsoft Azure. Permission policies specify your approvals within the resources in the account.

Roles required:

- Microsoft Azure or Microsoft Azure AD (Active Directory) Administrator
- Application Administrator
- Application Developer
- Cloud Application Administrator and the Resource Policy Contributor role to create or modify the resource policies.
- sn_cmp.cloud_admin role for operations in Cloud Services Catalog

Procedure

1. Assign roles to Azure users of Cloud Provisioning and Governance.

   You assign Cloud Provisioning and Governance roles to the user groups and individual users based on the user activities and responsibilities. For more information, see Assign roles to Azure users of Cloud Provisioning and Governance.

2. Install and configure MID Servers for Cloud Discovery and Cloud Provisioning and Governance.

   To ensure secure and reliable communications, the Discovery process communicates with your cloud provider accounts and cloud resources through one or more MID Servers. For more information, see Install and configure MID Servers to access cloud environments.

3. Set up the MID Servers on your network or in one of your cloud networks.
   Data is encrypted to the MID Server and between the MID Server and the API endpoint.
4. Configure one or more MID Servers for each datacenter under management for high performance and security.
   Configure the MID Server even if you have already configured other MID Servers while setting up Cloud Services Catalog for another cloud provider.
5. Create a Microsoft Azure service principal.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To access the resources and billing data securely on your Microsoft Azure account, the Discovery process must present appropriate Microsoft Azure account credentials. You create a special programmatic account—a Microsoft Azure service principal—to generate the required credentials. For more information, see Create a Microsoft Azure service principal.

6. Store the Azure service principal credentials in the instance.

   To access the data on your provider account securely, the Discovery process must present the appropriate credentials. To make the credentials available to Discovery, you first create Microsoft Azure service principal credentials in the Azure Portal. You can then securely store the credentials in a service account in your instance. For more information, see Store the Azure service principal credentials in the instance.

7. Add an Azure service account. Add an Azure service account

   A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter. For more information, see Add an Azure service account.

8. Discover all datacenters in a service account.

   Discovery uses the information in the service account to identify all logical datacenters that are associated with the provider account.

9. Set up a cloud account for Azure.

   A cloud account is the logical representation in Cloud Services Catalog of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts, including the service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account. For more information, see Set up a cloud account for Azure.

10. Set the capacity limits on user requests for resources.

    Capacity limits place restrictions on the attributes of the cloud resources such as the number of virtual machines, virtual CPUs, or aggregate storage. You can set the limits on the resources separately for each logical datacenter in a cloud account.

Integrating Azure DevOps and the CI-CD pipeline

You can process the release request from the Cloud Services Catalog workflow with Azure DevOps for end-to-end automation of the development process, including provisioning the resources, and deploying and monitoring the applications.

The Azure DevOps and Continuous Integration-Continuous Deployment (CI-CD tool) pipeline integration is a release deployment process. You can use Cloud Services Catalog and the CI-CD tool to manage the deployed resources in Configuration Management Database (CMDB), directly through flows and actions.

Benefits of Azure DevOps pipelines with Cloud Services Catalog

By using the Azure DevOps pipelines with Cloud Services Catalog, your organization gets the following benefits:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Ability to work with CI-CD pipelines regardless of the underlying languages that are used in the pipelines (JavaScript, Java, Node.js, or Python).
- Support for all major version control systems such as GitHub, GitLab, Azure DevOps, Atlassian Bitbucket, or Team Foundation Version Control (TFVC).
- Support for Amazon Web Services and Microsoft Azure Cloud services.
- Ability to deploy the application anywhere (on-cloud or on-premises, containers, or virtual machines) or in any environment. For more information, see ADO and Release and Life-cycle Management.
- Ability to deploy Kubernetes, Platform as a Service (PaaS), and so on, without any development change needs.
- Automatic updates to CMDB through the required tags.
- Approvals and governance steps that are built into the workflow.

Key steps in integrating Azure DevOps and CI-CD pipeline

1. Use the release pipelines and Azure DevOps by Cloud Services Catalog.
2. Set up the Microsoft Azure DevOps console.
3. Run Discovery on Azure DevOps config provider.
4. Order a Microsoft Azure DevOps catalog item.

Release pipelines and Azure DevOps by Cloud Services Catalog

Azure DevOps has release pipelines, such as the Continuous Integration-Continuous Deployment (CI-CD) pipelines, that offer a build, test, and deploy approach. You can use these fully automated processes to rapidly deploy build requests to the production environment.

Cloud Services Catalog offers the Azure DevOps release pipeline as a CI-CD tool. You can build, test, and deploy applications for the entire development life cycle across different platforms and cloud providers.

Azure DevOps within Cloud Services Catalog offers additional features for managing work items, version control, project management, and reporting. Resource requests via Azure DevOps pipelines are processed by approval through the governance and conformation workflow.

Set up the Microsoft Azure DevOps console

Set up the Microsoft Azure DevOps console as the first step before you run Discovery and order a catalog item in the Cloud Services Catalog application.

Before you begin

1. Add a deploymentID variable to the Microsoft Azure DevOps pipelines in your organization. The Microsoft Azure DevOps integration discovers pipelines with a deploymentID variable only.
2. Tag all the provisioned resources with a deploymentID key in the pipeline so that you can discover the resources after provisioning.

The subscriptionID isn't required for Azure DevOps and Ansible templates.

Role required: none

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Name a pipeline variable deploymentID that you can use for pointed Discovery to update the resources in CMDB.
2. Create a Personal Access Token (PAT) credential that you can use in Azure DevOps Discovery.
3. In the **Update variable** dialog box, override the pipeline variables by selecting the **Let users override this value when running this pipeline** check box.

Update variable dialog box

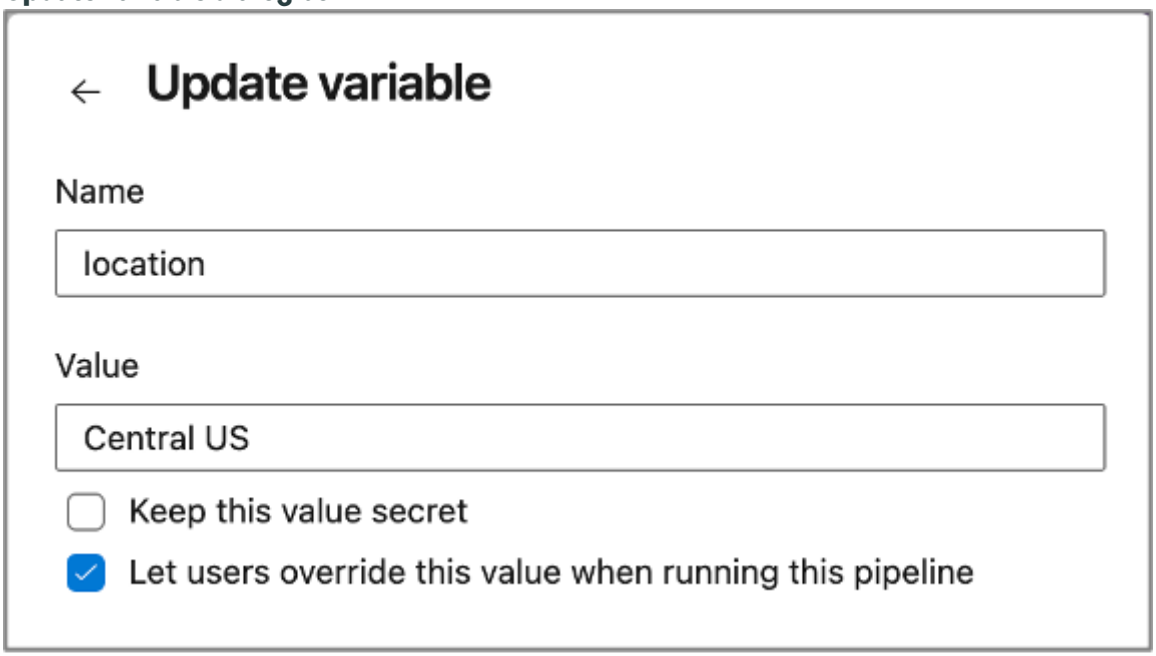

The pipeline works with the CSC DevOps Integration catalog item. The Microsoft Azure DevOps integration discovers pipelines with the deploymentID variable.

Run Discovery on Azure DevOps config provider

Add the Azure DevOps config provider and run Discovery to discover all projects, pipelines, and pipeline variables in an organization by using the Cloud Services Catalog application.

Before you begin

Role required: none

Procedure

1. Navigate to the **Admin** portal and select **ManageCredentials**.
2. Add the **API key** credential that you created with the Personal Access Token (PAT) credential. For more information on setting PAT privileges, see Azure DevOps permissions.
3. Prefix the Personal Access Token (PAT) with a colon and encode it to base64 format.

   a. Add a colon before your PAT value.

   Example
   Format: `:PAT`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Encode the string to base64 format.

Example
The result is an encoded string.

4. Prefix the encoded base64 string with `Basic` (include a space after Basic).

Example
Format: `Basic (encoded_base64_string)`

5. Add a required alias.
6. Navigate to **Manage** > **Config Management** and add a config provider.

Important:

Infrastructure as Code (IaC) Discovery is supported. You can discover the Azure DevOps pipeline Ansible job templates by setting up a scheduled discovery.

7. Discover all projects, pipelines, and pipeline variables in an organization by running Discovery.

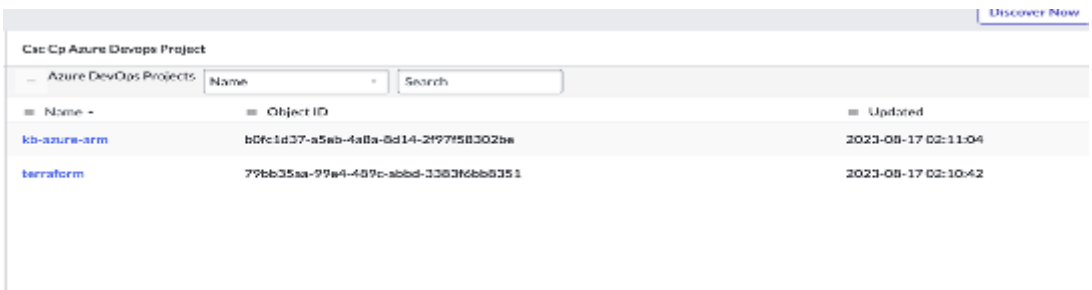

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Find the Config installables under the selected Azure DevOps project.

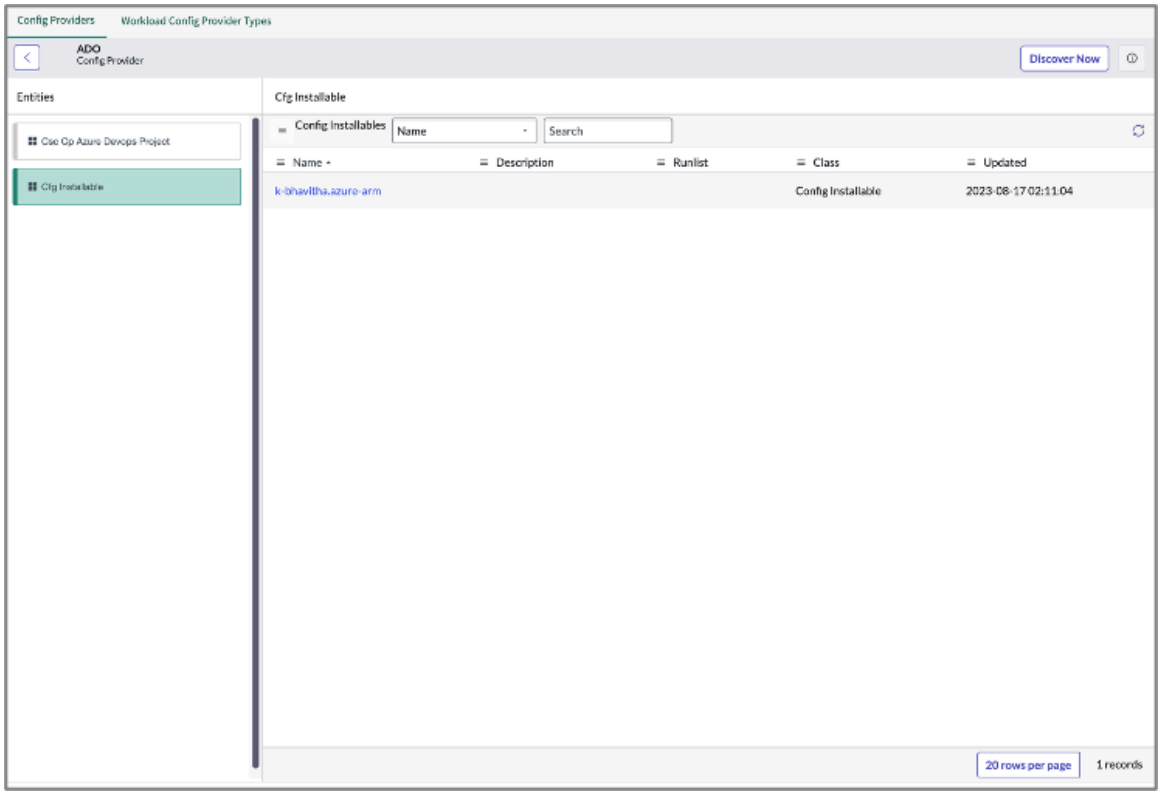

Result

You can now order an Azure DevOps catalog item from the Azure DevOps catalog order form on Employee Center.

Order a Microsoft Azure DevOps catalog item

Provision the Microsoft Azure DevOps catalog by using the Microsoft Azure DevOps catalog order form in the Cloud Services Catalog application.

Before you begin

To use Microsoft Azure DevOps for Cloud Services Catalog, assign Azure DevOps and Cloud Management capabilities to the MID Server. For more information, see Configure MID Server capabilities.

Set the deploymentID as variables in the AWS pipelines that you want to discover. The subscriptionID isn't required for Microsoft Azure pipelines.

Tag all the resources with the deploymentID key.

Role required: none

About this task

Follow this short video to learn how to order an Azure DevOps catalog item.

Procedure

1. In the **Update variable** dialog box, override the pipeline variables by selecting the **Let users override this value when running this pipeline** check box.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

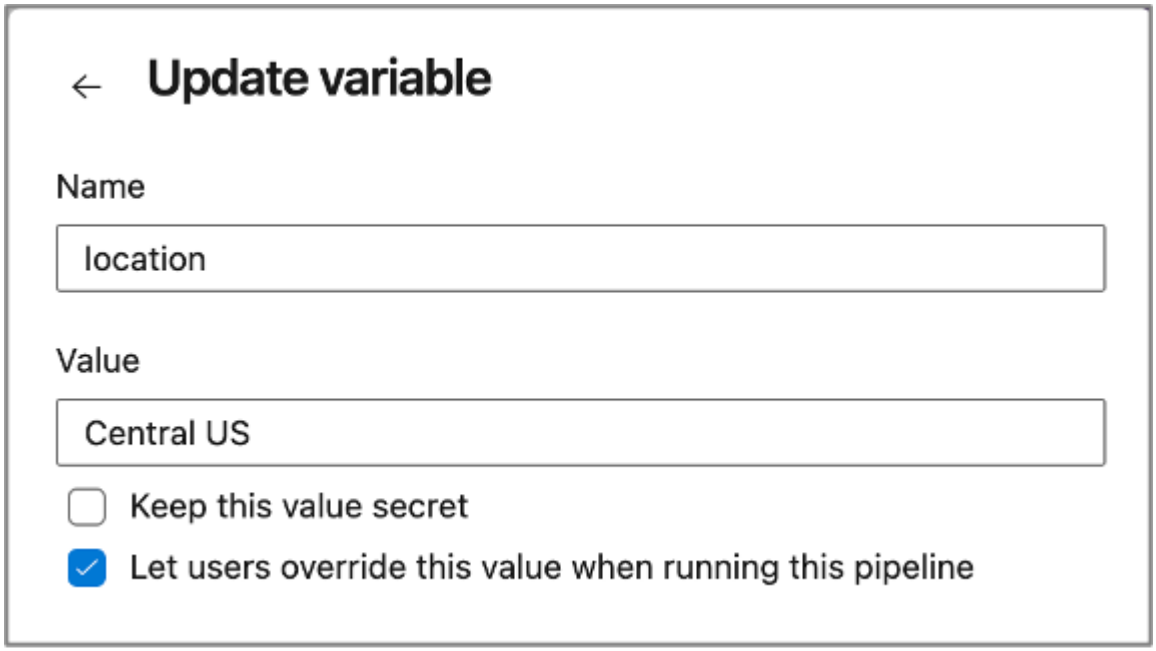

2. Navigate to **Employee Center** > **Cloud Services** > **Integrations** > **CSC DevOps Integration**.

   Use this Azure DevOps Catalog item order form to call an ADO pipeline.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

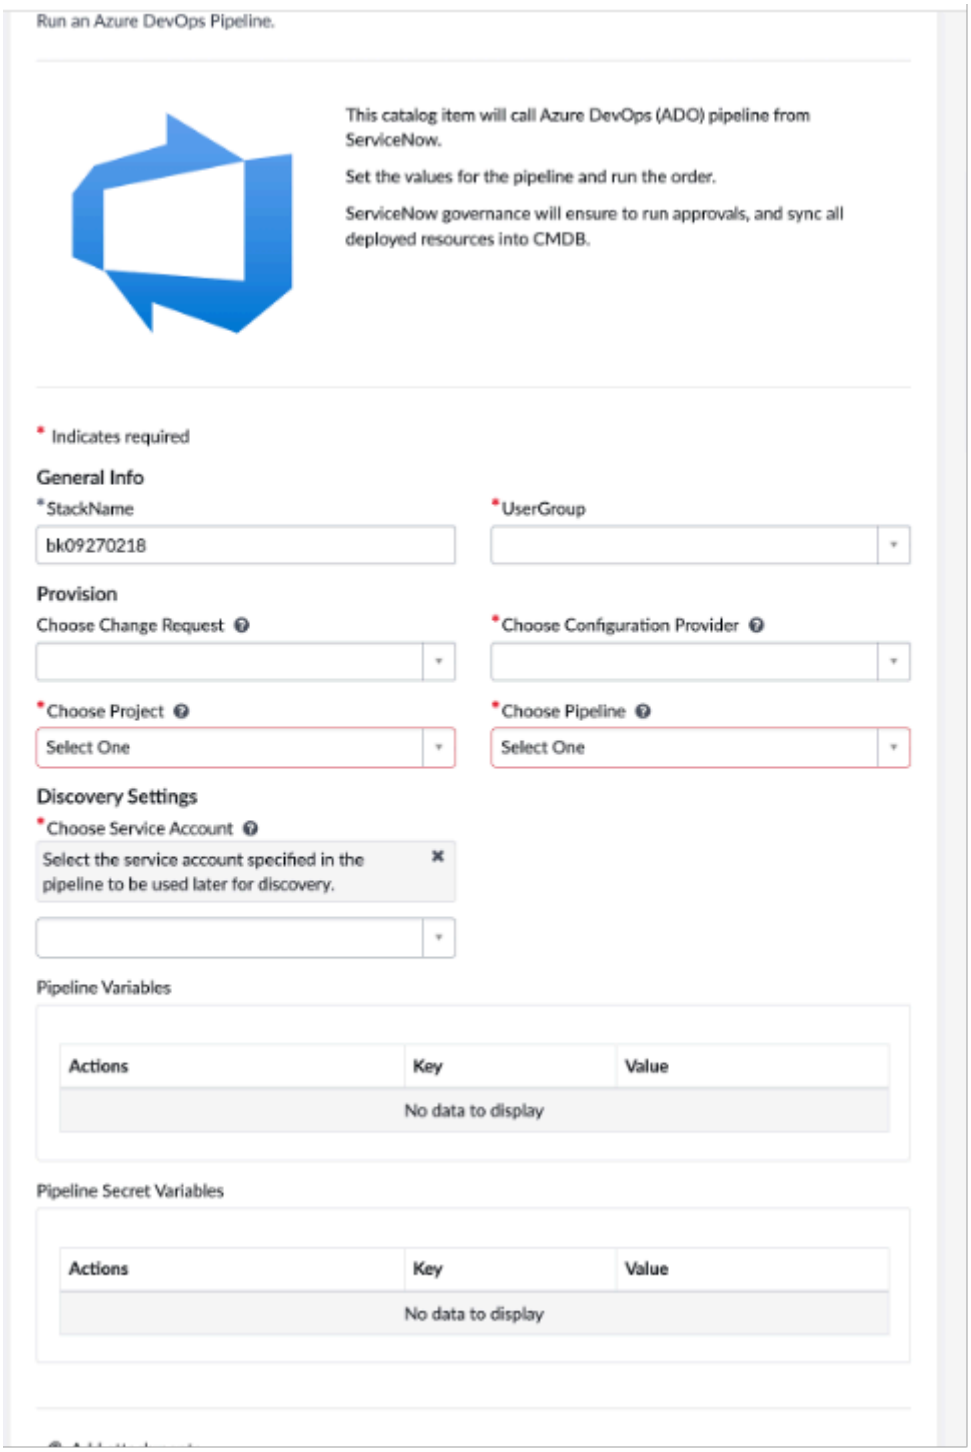

3. On the form, fill in the **UserGroup** and **Change Request** fields.
4. In the Choose Configuration Provider list, select the config provider that is associated with the Azure DevOps organization.
5. Select the project and associated pipeline in the organization.

   The pipeline variables are automatically filled in. You can modify the variables. For example, although the pipeline secret variables are masked, you can modify these values.
6. Select the **Service Account**, which is the same option that you selected as a service connection in the Azure DevOps pipeline.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For Amazon Web Services, a location list is displayed when you select **Service Account**.
- You can expose a pipeline variable location or region and use it in a pipeline for provisioning. This pipeline variable is overridden by the value that you chose in the location list to ensure that provisioning and Discovery are at the same location.
- A location is required for Discovery in Microsoft Azure, so no list is displayed when you select **Service Account**. These settings ensure the correct resources that are related to the stack are populated in CMDB. You can run Discovery based on those settings.

7. If the location is pre-filled for the AWS pipeline, select the same location in the catalog for Discovery.
8. Set up the identifying tags for the provisioned resources.
   The DeploymentID field is read-only and auto-generated for the order. The DeploymentID is assigned as a tag to all the provisioned resources if you added it as a tag in the Azure DevOps pipeline template of the stacks in the cloud. The DeploymentID is also assigned as a tag to the stack. If the pipeline has any of these three variables (Application, BusinessService or CostCenter), they're displayed in the General info in the Azure DevOps UI.
9. From the list, select the field values to be updated in the pipeline variables.

   These entries are considered for the current order only.

> **Note:** You can't modify these read-only variables from the pipeline variable section. These variables are added as tags to the stack in the cloud and are inserted in the CMDB key value table as well. You can do only the resource level actions on a provisioned stack because the stack level action is turned off. The previous stack is marked as obsolete if you run the pipeline again, although the resources remain in the cloud.

What to do next

Verify that a blueprint approval policy (CSC Content Approval Policy ADO Integration) is applied on the Microsoft Azure DevOps catalog. This policy mandates that an approval must be obtained from the Change management group before any provisioning can take place.

Integrating Ansible with Cloud Services Catalog

Integrate Ansible with Cloud Services Catalog as a Day 1 activity. With this activity, you can deploy an Ansible job template via a ServiceNow catalog item.

As a cloud_service_user, you can access the config provider and job templates and deploy a job template through a catalog item.

Day 0 support	Day 1 support	Day 2 support
1. Retrieve the resources from the cloud and update the CMDB 2. Verify that there are no errors. If errors are displayed, then the Discovery or job template deployment has failed	1. Order an Ansible job template from Service Catalog 2. Reference the discovered Microsoft Azure pipelines that are specific to Ansible 3. Select inputs that are needed for the release deployment pipeline 4. Submit a catalog	1. Fix the mapping of Out Of Box via Cloud Services Catalog content 2. Update documentation on limiting the job template to run on a specific host 3. Test and document how to support the job template deployment on unmanaged nodes 4. Start your Day 2 support for Ansible AWX

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Day 0 support	Day 1 support	Day 2 support
	5. Publish an Out Of Box catalog item as Cloud Services Catalog content 6. Select the job templates that you have access to and deploy them	
Store in CMDB as referrable objects.	1. Wait for the order approval 2. Track the progress of your template order 3. Update CMDB with the deployed resources	1. Ensure all the Day 2 support actions execute successfully 2. Switch between the blue (current application) and green (new application) environment

Set up Ansible as a Day 1 task

Set up the Ansible console as a Day 1 task before running Discovery and deploying job templates via catalog items.

Before you begin

All extra variables must have a variable called deploymentID. The value of this deploymentID is sent from the catalog request.

Select the **Prompt on launch** option next to the variables section in the template. This action allows you to update or override the default extra_vars with values from the catalog order form when launching the template.

Tag all resources with the deploymentID key to enable Discovery after provisioning.

Role required: admin

Procedure

1. Save the extra variables in a JSON format locally, to discover and update Configuration Management Database (CMDB).
2. Create Ansible tower credentials to be used in Ansible Discovery.
   For more information, see Configure Ansible Tower user name and password.
3. Allow the extra variables and job templates to be overridden.
   The default Ansible job template must be overridden for the pipeline to work with a Cloud Services Catalog Ansible automation catalog item.

Run Discovery on the Ansible config provider

Add the Ansible config provider and run Discovery to discover what's in the inventory, the host group, and the available job templates.

Before you begin

Do this procedure only if you have Ansible Tower credentials.

Role required: none

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Cloud Admin Portal** > **Manage** > **Credentials** > **New** > **Ansible Tower Credentials**.
2. On the form, fill in the fields as per the information received from the Ansible Tower.

Ansible Tower credentials

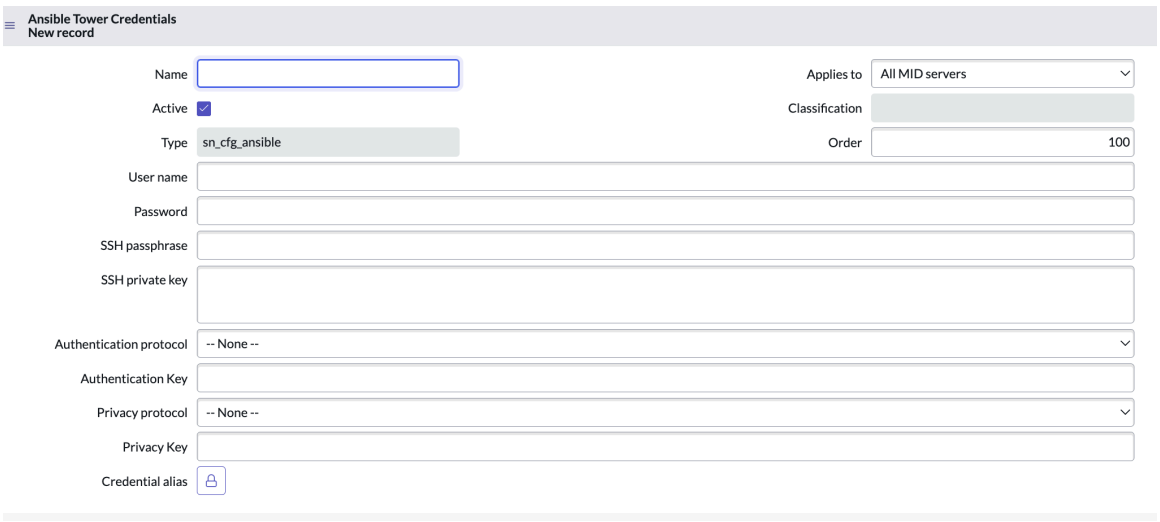

3. Select **Submit**.
4. Create a config provider by selecting **New**.

Create Config Provider form

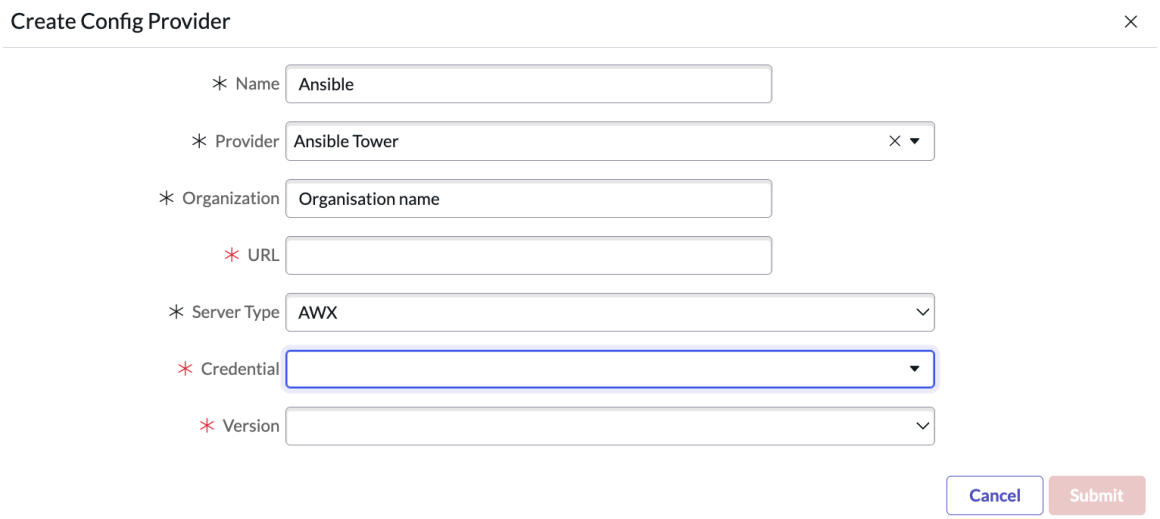

Important:

Infrastructure as Code (IaC) Discovery is supported. You can discover job templates, inventory, and groups.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Run the Discovery to discover all projects and extra variables in an organization.

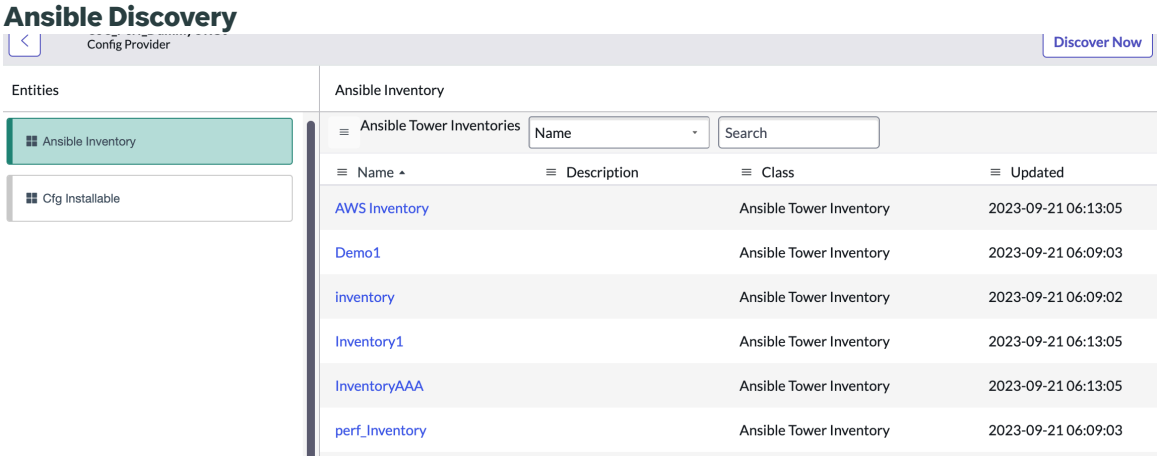

Result

You're all set to take the next steps to order an Ansible item.

Order an Ansible job template

Use Cloud Services Catalog to order an Ansible job template called from the catalog item.

Before you begin

Role required: cloud_service_user

About this task

You can use the Ansible Job template deployment form to call an Ansible job template and run it via a ServiceNow catalog item for deployment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Employee Center** > **Cloud Services** > **Integrations** > **CSC Ansible Automation**.

Ansible Job template deployment form

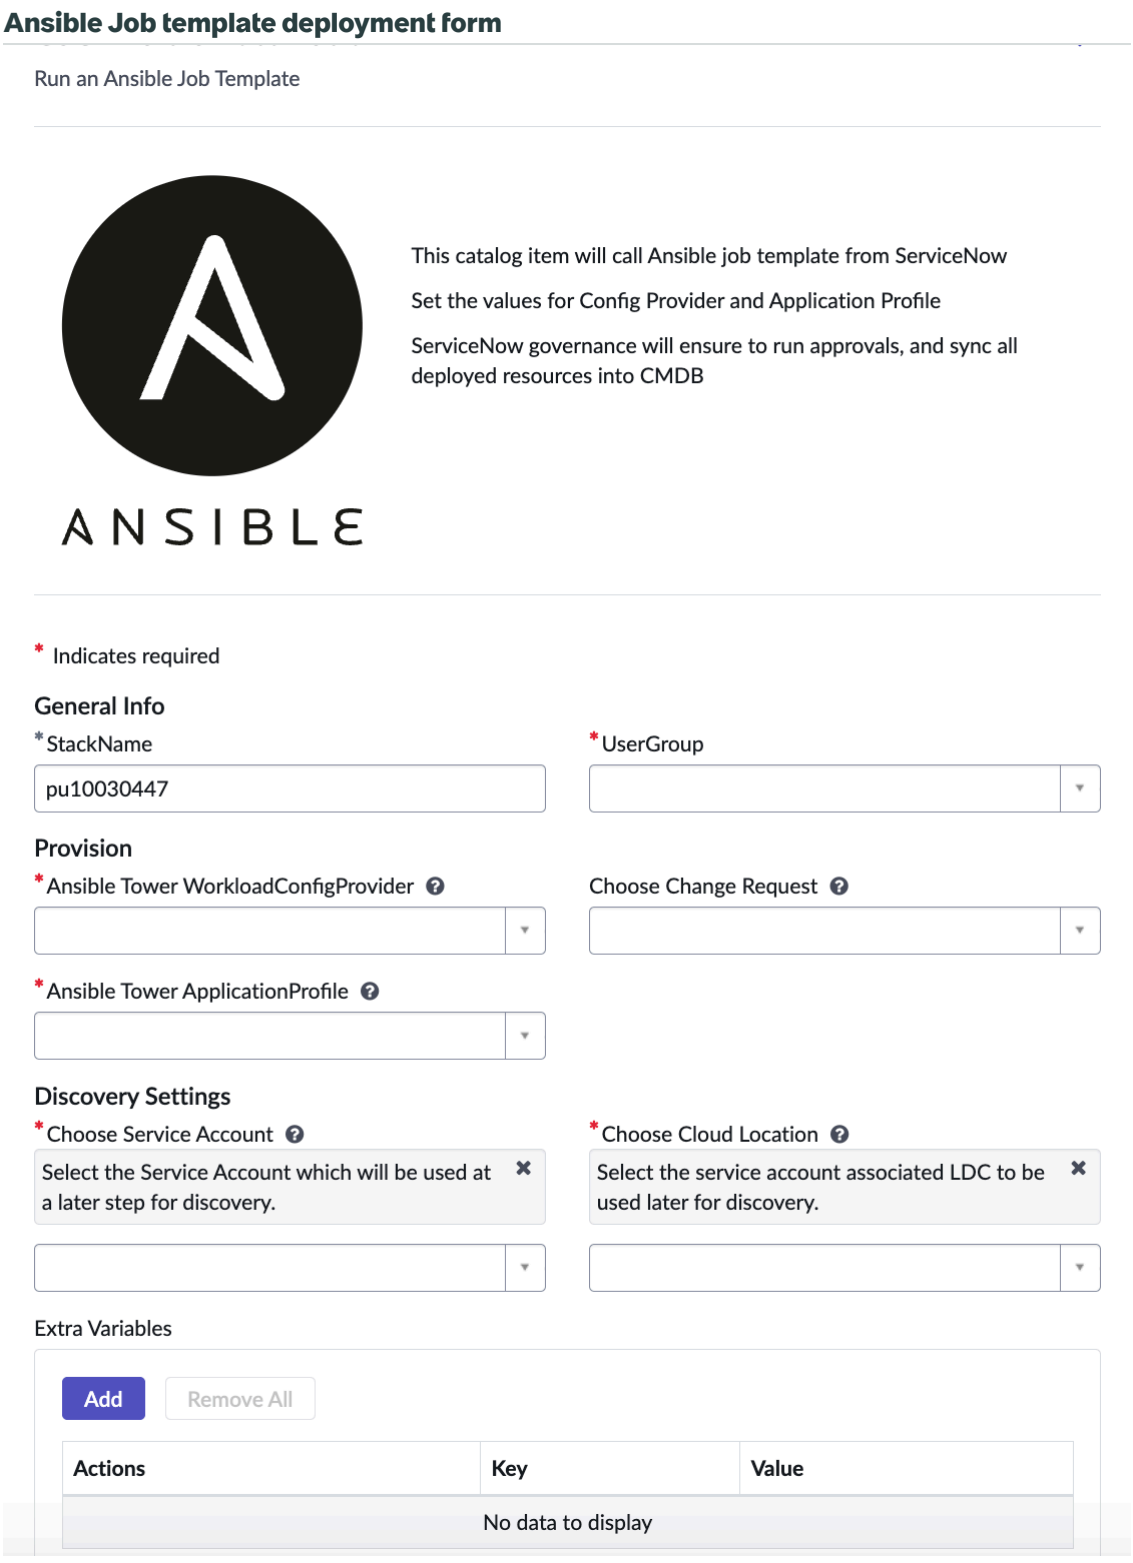

Ansible Template form fields

Fields	Description
UserGroup	User group that is assigned to the user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields	Description
Provision	Provisioning details. Select the Workload Config Provider and ApplicationProfile from the Ansible Tower.
Choose Service Account	Pipeline. For Amazon Web Services, the Location list is displayed when you select **Service Account**. The location is not required for Discovery in Microsoft Azure. **Tip:** You can expose an extra variable that was saved in a JSON format at the Ansible tower location or region and use it in the template for provisioning. This extra variable is overridden by the value that you selected in the Location list. By using this extra variable, you ensure that provisioning and Discovery are at the same location. These settings ensure that the correct resources that are related to the stack are populated in CMDB. You can run Discovery based on those settings.
Choose Cloud Location	Service account that is associated logical datacenter (LDC) for Discovery.
Extra Variables	Extra variables that you can modify even if the values are automatically filled in and masked.

2. Set up the identifying tags.
   The DeploymentID field is read-only and auto-generated for the order. The DeploymentID is assigned as a tag to all the provisioned resources if you added it as a tag in the Azure DevOps pipeline template of the stacks in the cloud. The DeploymentID is also assigned as a tag to the stack. If the pipeline has any of these three variables (Application, BusinessService or CostCenter), they are displayed in the General info in the Ansible Job Template form UI.
3. Select the field values from the list to be updated in the pipeline variables.

   These entries are considered for the current order only.

   Note: You can do only the resource level actions on a provisioned stack because the stack level action is turned off. The previous stack is marked as obsolete if you run the pipeline again, although the resources remain in the cloud.

What to do next

Verify that a blueprint approval policy (CSC Content Approval Policy Ansible) is applied. This policy mandates that approval must be obtained from the Change management group before any provisioning can take place.

Cloud Services Catalog administration guide

After you set up the Cloud Services Catalog application, you must set up some additional items so that your users can request and manage any cloud resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuring governance

You can use governance tools to restrict the provisioning of cloud resources, such as quotas and policies. For example, you can configure the approval workflow that is used after a user requests a cloud resource. You can also configure whether an email is sent after a cloud resource is stopped. You can configure which datacenter is automatically selected when a user selects a resource from a cloud provider. For more information about quotas, policies, and other governance types, see Cloud Governance.

Populating the Cloud Services Catalog

After you apply the governance quotas and policies, you must create Cloud Services Catalog items. You can create a catalog item in the following ways:

1. Import a template from Amazon Web Services, Microsoft Azure or Google Cloud Platform.
2. Use or modify an out-of-the-box catalog item.
3. Use a Terraform template.

For information about domain separation, multi-repo support with Terraform app, cloud accounts, the Cloud admin portal, or the Cloud API, see the following topics:

- Use multiple repositories structure with Terraform Connector app
- Domain separation and Cloud Services Catalog: Get additional information about domain separation.
- Cloud accounts: Learn about the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure.
- Cloud Admin Portal: Manage, design, govern, operate, and analyze all your cloud resources from this portal.
- Cloud API (CAPI): Integrate Cloud Provisioning and Governance with cloud providers by using REST APIs.

Cloud Services Catalog and Terraform Connector

Use the Cloud Services Catalog app with Terraform Connector support with all its features that help in provisioning and managing cloud resources across various public and private clouds.

Key features of Terraform Connector:

- Use Terraform as a configuration provider to provision resources in the following clouds:
  - Amazon Web Services (AWS)
  - Microsoft Azure
  - Google Cloud Provider
- Discover the Infrastructure as Code (IaC) templates to discover the IaC (Terraform) templates from the Terraform environment.
- Generate catalog items using the Terraform templates, utilizing them to provision resources across different cloud environments.
- Utilize the multiple repositories structure support with the Terraform app. See Use multiple repositories structure with Terraform Connector app to learn more.

For more information on Terraform Connector, configuring the advanced features, and how to use it, visit Exploring Cloud Services Catalog Terraform Connector.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Permissions management for Cloud Services Catalog

Permissions are your group-level access rights to features and actions in Cloud Services Catalog.

The permission policies in Cloud Provisioning and Governance User Portal control the actions that are allowed or limited for users based on the user group. See Permissions Management for permission types and ideas.

Currently, this does not apply for catalog provisioning in Cloud Services Catalog. For CSC, all users with roles permission can access, view and provision the catalogs, by default.

Policies and permission levels on services

Configure policies with the necessary level of permissions to provide access to the AWS, Azure, ADO, Ansible and Google Cloud services, respectively.

- Cloud Services Catalog AWS Out Of Box catalog permissions
  - Amazon Web Services EC2 Linux Out Of Box permissions
  - AWS EC2 Windows OOB permissions
  - AWS Dynamo DB permissions
  - AWS S3 Bucket
- Azure OOB catalog permissions
- Ansible permissions
- ADO permissions
- Cloud Services Catalog Google Cloud Platform Out Of Box catalog permissions
  - Google Cloud Platform (GCP) Big Query Out Of Box permissions
  - Google Cloud Platform Cloud SQL Out Of Box permissions
  - Google Cloud Platform Linux Out Of Box permissions
  - Google Cloud Platform Load Balancer Out Of Box permissions
  - CSC GCP Windows OOB permissions

Using Cloud Services Catalog

The Cloud Services Catalog application is a dynamic portal for all your day-to-day cloud activities. It enables you to manage stacks, work with requests, and access the Out Of Box catalog items and other popular topics.

You can explore Cloud Services by performing a search on the ServiceNow AI Platform® platform. You can also find related links, like **Browse Cloud Catalog** and **CSC-Guided Setup**.

Cloud Services is also designed to appear on the Employee Center toolbar as a new tab. Look for the options provided or navigate to **Browse Cloud Services** to experience Cloud Services Catalog.

The Cloud Services Catalog application displays the features in the following table.

Cloud Services features

Feature	Description
My Stacks	Lists all the stacks that are assigned to you. Select a catalog item to view the stacks that the item launches.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Services features (continued)

Feature	Description
	Select a stack to view the details such as all the resources, resource properties, activities, and life-cycle operations.
Cloud Services Topics	Provides the quick links to popular and relevant topics.
Browse Cloud Services	Provides the lookup feature that is powered by AI Search. It returns relevant Cloud Services, provides filters, and sorts results that help you to zoom in on the Cloud Service that you want. This feature is separate from the global search featuring on the Employee Center landing page. **1.** Use Filter by and Sort by to narrow your search. **2.** You can also use the Card or List view as needed. **3.** Select one of the sources to filter the search results. **4.** Use a topic filter to narrow the results.
Recent Stacks	Lists all the stacks with the recent activities that were performed on them.
My Resources	Observes and interacts with the discovered resources.

Requesting Cloud Services

Submit a request for one of the available service catalog item in cloud services. You can find the catalogs under **Browse Cloud Services**.

Before you begin

Role required: sn_cmp.cloud_service_user

Procedure

1. Navigate to **Browse Cloud Services** under **Cloud Services** on the Employee Central portal page in Cloud Services Catalog.
2. Find a suitable catalog item from options displayed.
3. Select **Request**.
   This opens the request form with information and attachments to be provided.
4. Enter the details in the form fields to order.
   When you submit a request for an item in the service catalog, the system typically provisions the stack immediately. Sometimes, your request must be approved first.

Exploring the My Stacks tab

Sort, view, or manage your stack of cloud resources and services as a part of the Day 2 activities by using the **My Stacks** tab on the Employee Center toolbar in Cloud Services Catalog.

The **My Stacks** tab has multiple options that let you categorize, sort, or view your stacks with a different managed status. The following example shows what the **My Stacks** tab looks like on Employee Center.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

My Stacks tab

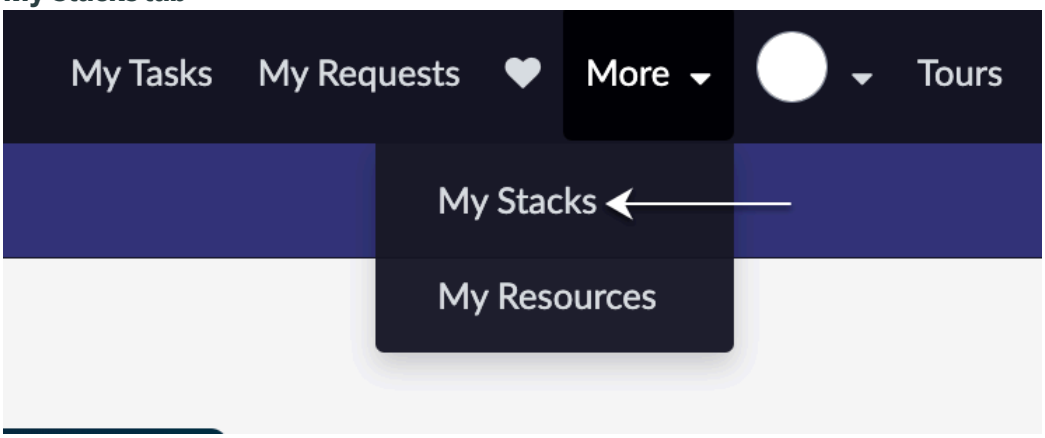

On the **My Stacks** tab, you can see the following statuses:

- Managed Stacks: All the stacks in all service categories.
- Stack by Catalog Items: Stacks that are listed by the catalog item.
- Stacks with Issues: Stacks that have errors due to incomplete or unapproved items.
- Stacks with follow-up actions: Stacks that have a pending action to be followed up.
- Discovered Stacks (unmanaged): Stacks that don't have any catalog that is associated with them.

Find the **My Stacks view** tab at the end of the My Stacks ribbon to access one of the following views:

- Stacks I own: Stacks that you created.
- Stacks owned by groups I am part of: Stacks that are created by you and by all the groups you're associated with.
- Stacks owned by anyone: All the stacks in the instance.

Manage My Stacks

Use the **My Stacks** tab to access all the stack properties such as the creation, update, and deletion details.

Before you begin

Role required: sn_cmp.cloud_service_user

Procedure

1. Access the Stacks page by going to the toolbar or homepage.

   - On the toolbar, select the **My Stacks** tab.
   - On the My Stacks homepage, select the **Recent Stacks** tab.

     By default, all stacks in all service categories, except the obsolete status, are listed.

2. In the navigation pane, use a filter category to display the matching stacks.
   The filters are displayed as editable conditions.

Result

The stacks are displayed with the properties that are in the following table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Stack properties

Property name	Description
Name	Name that you specified when requesting the stack.
Service Catalog Item	Service that was created when you ordered the catalog that resulted in this stack.
Stack Status	Status of the stack: • Initializing • Active (provisioned stacks) • Error • Processing (stacks being processed) • Terminating • Terminated (de-provisioned stacks) • Unmanaged (discovered stacks) • Resolving Error (stacks with errors) • Follow-up required • Obsolete
Owned by	User name of the requester.
Owner Group	Group name that the stack is assigned to.
Service Request Item	Record that was created in the Request Item table when you ordered the catalog that resulted in this stack.
Request	Request number that is mapped to the Request Item.
Created	Date and time that the stack was requested.
Updated	Date and time that the stack details were last updated.

Work with stacks

Start your Day 2 operations by starting, stopping, deprovisioning, or modifying your stacks.

Before you begin

Role required: sn_cmp.cloud_service_user.root_admin, admin, or sn_cmp.cloud_admin

About this task

When you request a life-cycle operation on a stack or resource, the system generates a change request. An approval policy specifies either that the change is auto-approved or that a user on the approver list must approve the change. For more information, see .

Procedure

1. In the navigation pane, in the **Name** field, select a stack and then select **Track Operations**.
2. View the request status and the status of the stacks by selecting **Status**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. View the properties and request parameters by selecting **Request Details**.
   You can also display the stack details in a pop-up window by selecting **View Full Details**.
4. Search and sort all the operations that you performed by selecting **Stack Operations**.

View stack dependency

Explore the details of the configuration items, services, and other tasks that are associated with the stacks and resources by using View Dependency.

Before you begin

Role required: admin

About this task

A stack can be dependent on multiple stacks, and multiple stacks can depend on a single stack. Stacks that comprise multiple configuration items (CIs) and services make up the stack dependencies in the Cloud Services Catalog application. The dependency map provides a simple and detailed graphical representation of these complex settings. You can navigate through the references to understand the event impact that one branch has on another.

Procedure

1. On the Managed Stacks page, select a stack to open the stack operations page.
2. On the View Full Details ribbon, select **View dependency**.
3. In **Search for CI**, enter a CI to view the dependency map.
4. On the form, fill in the fields.
   The following example shows the settings for viewing the stack dependencies. The settings include Save Custom Settings, Load Custom Settings, Predefined Filters, Dependency Type, and Max Levels.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Viewing stacks dependencies

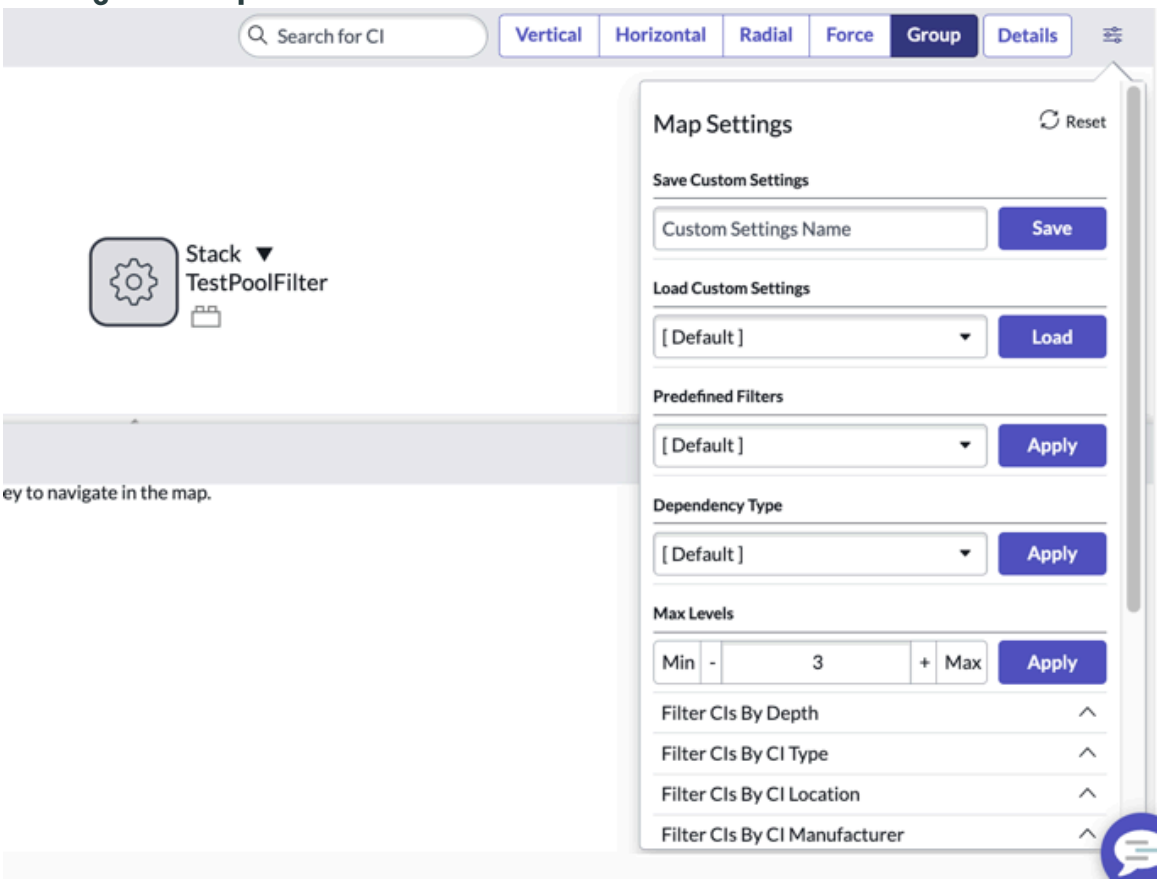

5. Undo all the CI filter selections and reset the map setting by selecting **Reset**.

Stack activities view

Display the cloud actions on your stacks and resources by using the **View Activities** tab.

You can do multiple actions in the stacks life cycle such as track operations, view activities, view dependency, and do actions. The following example shows the stack activities that you can perform and the following table describes the activities.

Stack activities

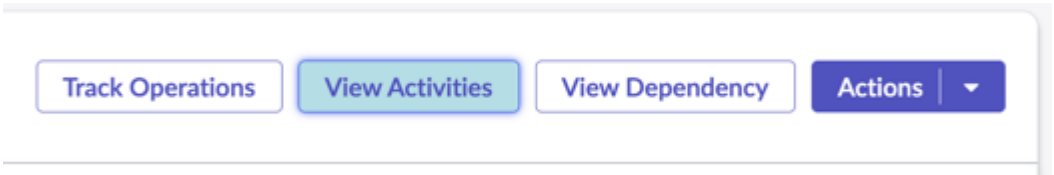

Stacks activities

Activity	Description
Change Requests	Change requests that have been submitted for stacks or resources that are assigned to you. The system generates change requests when you perform any life-cycle operations on a stack or resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Stacks activities (continued)

Activity	Description
Incidents	Incidents for stacks and their associated resources that are assigned to you.
Tasks	Catalog tasks that are created when a request for a stack fails to provision.
Scheduled Operations	List of scheduled operations that are associated with stacks and are displayed as these features: • Operation Time • Target • Operation • Status

Actions on stacks

Access a particular stack to perform a Day 2 or life-cycle operations such as stop, start, deprovision, ModifyLease, or ModifySchedule.

When you request a life-cycle operation on a stack or resource, the system generates a change request. An approval policy specifies that the change is auto-approved or that a user on the approver list must approve the change.

The stacks and resources undergo life-cycle operations from the time that they're started, provisioned, de-provisioned, and then stopped. At times, you must update or stop a particular resource's schedule for configuration issues. You can select a stack or resource to perform these operations. The actions include only the appropriate operations for the selected resource or stack. Not all operations are supported for all providers or for all service categories (resource types) by default.

The following example shows the actions that the Cloud Services Catalog application supports on stacks. The actions are stop, start, deprovision, ModifyLease, or ModifySchedule.

Actions on stacks

The following table lists the operations that the Cloud Services Catalog application supports as actions on stacks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Function
Stop	Stops the selected virtual machine (VM). The resource status changes to off. No setting is required.
Start	Starts the VM that you select. The resource status changes to on. No setting is required.
Deprovision	Terminates the selected stack and sends a notification to the stack owner. No setting required.
ModifyLease	Changes the lease end date.
ModifySchedule	Modifies the schedule profile and schedule time zone.

Redirection to My Stacks

You can use redirection to your **My Stacks** tab as an alternate but simplified way to locate a stack from the **My Requests** tab. For example, you can use redirection while you're looking up a stack that was recently provisioned.

You can view all your stacks with any relevant status, but the displayed list may be lengthy.

Redirection to stack details via the **My Requests** tab in Employee Center enables you to find the requested stack by the Requested Item (RITM). Use the **Cloud Request Status** tab on the Cloud Services Catalog application to display the request details and status of the stacks.

The following example shows how you can redirect to the **My Stacks** tab from the **My Requests** tab, instead of looking at a comprehensive list of stacks on the Stacks homepage.

Redirection to My Stacks

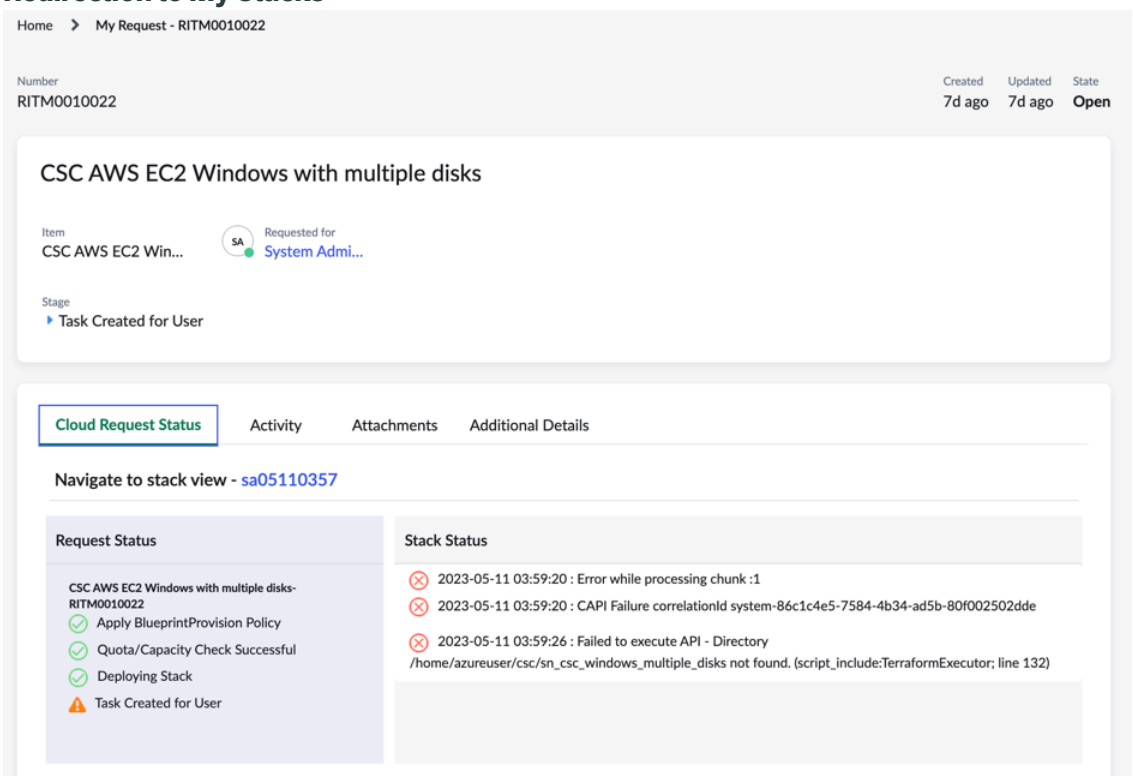

Exploring My Resources

Observe and interact with discovered resources by using the **My Resources** option under **More** tab on the Employee Center toolbar.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can designate the operators or owners for the identified resources so that they can handle these resources with standard Day 2 operations. By performing these actions through the ServiceNow AI Platform of the Cloud Services Catalog application, you can leverage the workflows for proper change management, audit trails, and overall governance.

The following example shows the **My Resources** tab.

My Resources tab

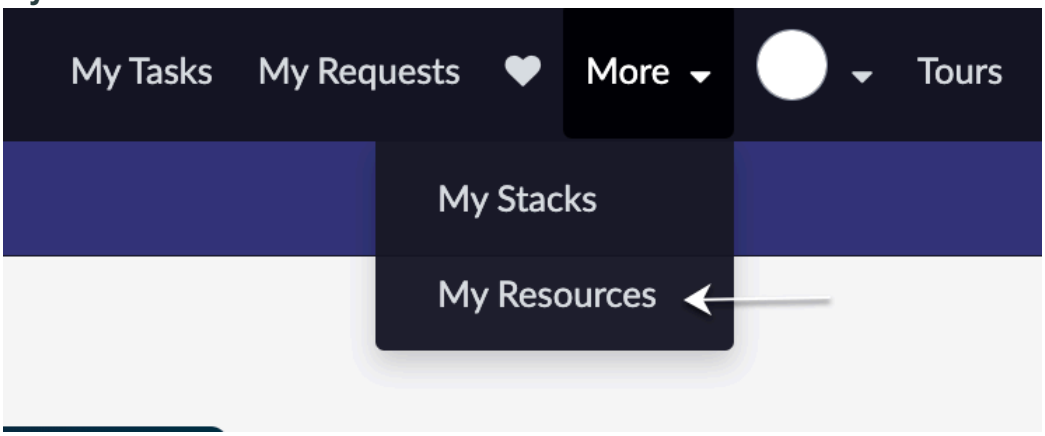

The following example shows a list of the resource categories and filters. You can use these to access the relevant resource and perform other functions, such as save and export.

Options on the All Resources tab

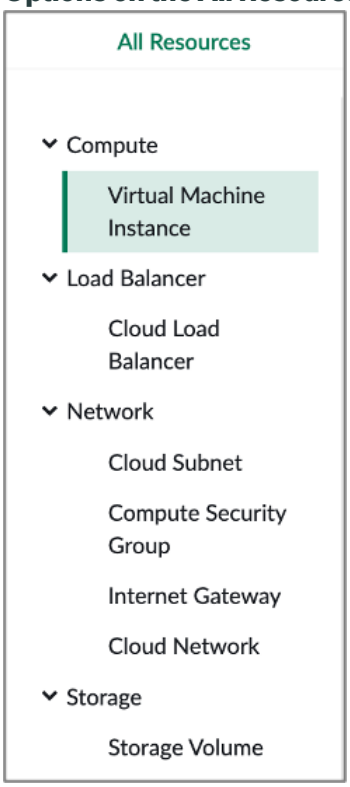

The cloud resources are categorised as above. You can access one of the categories to manage them.

Manage My Resources and resource filters

Access the resources that you own, resources that are owned by a group, or anyone's resources by using the resource filter. You can then find the information that you need by using sorting functions, conditions, categories, and keywords.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: admin

Procedure
1. In the navigation pane, select the resource category, apply a filter (pre-added or new), and select **Run** to display and export a set of resources.
2. Sort, view the details, and export as required.

Resource filter options

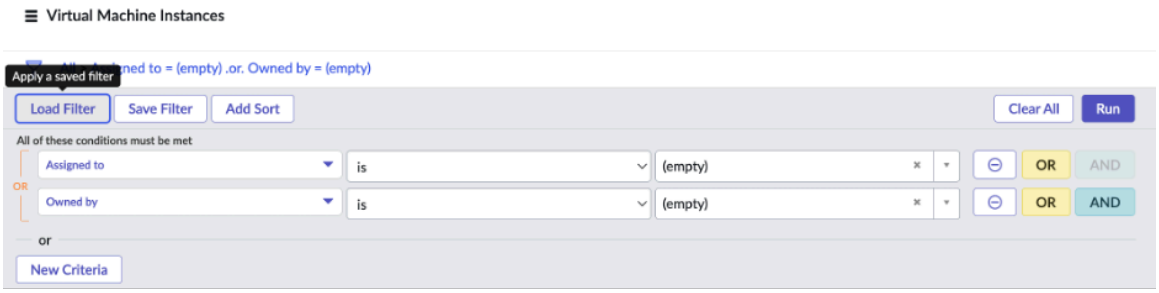

3. Apply an existing or a saved filter by selecting **Load Filter**.
4. In the Load Filter window, select **Manage Filter** to find the list of previously added filters.

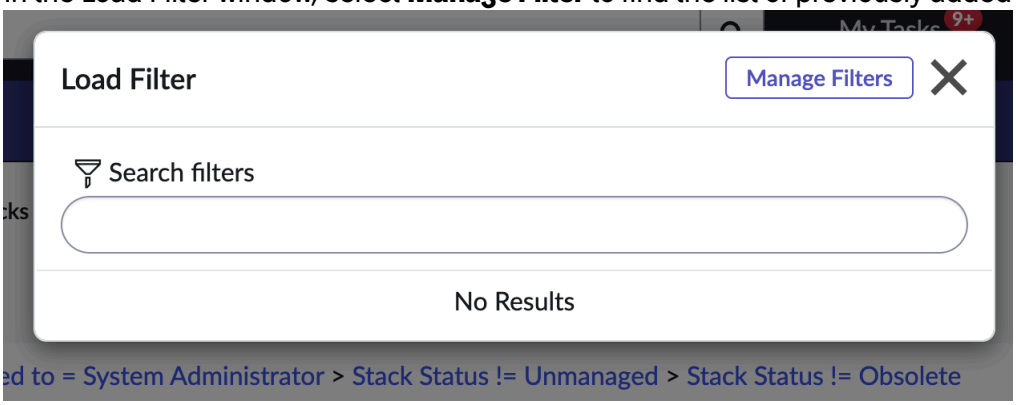

5. Select the menu option along **Filters** to view the various options to export the filtered resources.

6. You can also select a filter to display its details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

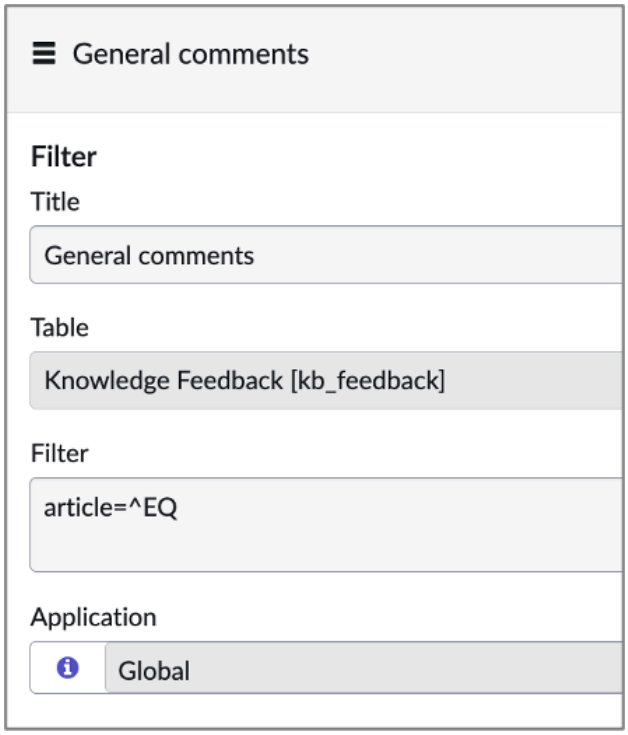

7. View more options and details about the filter that you chose by selecting the menu option next to General Comments.

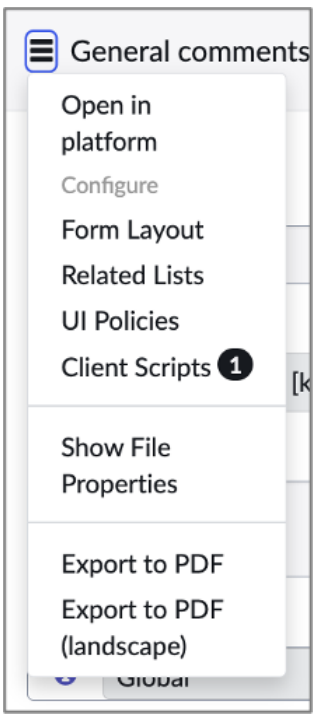

You can use the keyword search to find a filter directly by using keywords. You can choose export options such as PDF, Excel, or CSV.

8. Create and add one or more sets of conditions by navigating to **New Criteria** > **Create another set of conditions**.
9. Save the created filter by selecting **Save Filter**.
10. In the Name, Query, and Visibility fields, assign the name, query, and visibility.
You can also add sorting options by selecting **Add Sort**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Assign a sorting order from the list and then select **Save**.
12. From the View options, select one of the following options:
    - Resources I own: This option is selected by default
    - Resources owned by groups I am part of: This option displays all the groups that you are part of.
    - Resources owned by anyone: This option displays all the resources in the selected category, in Configuration Management Database (CMDB), and assigned.

Out Of Box Catalogs using Cloud Services Catalog

Cloud Services Catalog has Out Of Box Catalog items, actions and policies that you can render, as required.

Use Out Of Box catalog items for Azure, GCP and Amazon Web Services Cloud Service requests and automations. You can either order one of the catalog items, use Flow Designer based actions or provision policies.

Note: Use the catalog items as-is, and do not modify or update them. Making any modifications or customizations will change the Out Of Box behavior of the catalog and not yield the expected results. This will also impact their upgrades. You can create a copy of the templates and recreate your catalog items considering the Out Of Box content.

Order from these Cloud Services Catalog Out Of Box Catalog Items:

Amazon Web Services EC2 Linux Out Of Box Catalog items

Cloud Services Catalog Linux VM with agent client collector (ACC), up to 10 additional disks on EC2 or with security groups.

This is a fully functional sample catalog item that can be used to request a Linux instance of any instance type from Amazon EC2.

The video demonstrates how to order an Amazon Web Services EC2 Linux virtual machine.

Linux VM orders and features:

Order	Features
Order a Linux virtual machine on EC2, that allows installation of agent client collector. ACC is compulsory for production instances. VM sizes other than t2.nano, t2.micro require approval. Cloud configuration scan is run for compliance checking.	For instance types other than t2.nano and t2.micro, approval is needed from administrator, by policy. The catalog item has built in client rules to drive an ACC (Agent Client Collector) installation when the deployment environment is marked for production. To select the credential alias key for ACC deployment from mid, the appropriate ACC admin role permissions are required to be bestowed to the user. For prerequisites check “Deploy ACC-* in post provisioning” below. The work-flow automation runs a CCG scan on the deployed Linux instance and checks the policy rule with the configuration of the VM. If there is any violation, the stack will be set with 'Follow up required' state to notify that the stack deployment is not as per norms.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Order	Features
Order a Linux virtual machine on EC2 and specify up to 10 additional disk. For up to two disks, no approval is necessary. Seek approval for anything beyond that.	This catalog item form allows ordering up to 10 additional disks (volumes) attached to the VM of varying sizes and types. If more than two additional disks are added to the request, seek approval from administrator, by policy.
Order a secure Linux virtual machine with enhanced security by assigning a Network Security Group for highest level of security.	This catalog item form allows capability to assign one Network Security Group to the compute instance so that the network access to the VM is restricted. Security groups are listed for the selected VPC and subnets.

Note:

Tag (key-value) is assigned to all deployed resources. Example: VM, network, storage as present in the stack.

The key-value is updated only in the cmdb_key_value table, not in public cloud at this time. This will be fixed in future releases to update the tags in cloud.

Amazon Web Services EC2 Linux Out Of Box permissions

Amazon Web Services EC2 Linux OOB catalog items permissions

Catalog item	Services	Permissions
CSC AWS EC2 Linux VM	AWS EC2 provision	1. DescribeVolumes 2. DescribeImages 3. DescribeInstances 4. DescribeVolumes 5. DescribeInstanceCreditSpecifications 6. DescribeInstanceAttribute 7. DescribeInstances 8. DescribeTags 9. DescribeInstanceTypes 10. DescribeInstances 11. DescribeVpcs 12. RunInstances 13. DescribeImages 14. GetCallerIdentity
	AWS EC2 instance- stop	1. DescribeInstances 2. StopInstances

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Catalog item	Services	Permissions
	AWS EC2 instance- start	1. DescribeImages 2. DescribeInstances 3. StartInstances
	AWS EC2 instance-deprovision vm **Note:** This de-provisions storage volume as well.	1. DescribeAddresses 2. TerminateInstances
	AWS EC2 instance-deprovision stack **Note:** Terminates everything including nic.	1. GetCallerIdentity 2. DescribeInstances 3. DescribeImages
CSC AWS EC2 Linux multiple disks	AWS EC2 Linux multiple disks provision	1. DescribeVolumes 2. DescribeImages 3. DescribeInstances 4. AttachVolume 5. DescribeInstances 6. CreateVolume 7. DescribeInstanceCreditSpecifications 8. DescribeInstanceAttribute DescribeVpcs 9. DescribeInstanceTypes 10. DescribeTags RunInstances 11. GetCallerIdentity
	AWS EC2 Linux multiple disks stop	1. DescribeInstances 2. StopInstances
	AWS EC2 Linux multiple disks start	1. DescribeImages 2. DescribeInstances 3. StartInstances
	AWS EC2 Linux multiple disks stack deprovision-terminates all resources	1. DescribeInstances 2. TerminateInstances 3. ModifyInstanceAttribute

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Catalog item	Services	Permissions
		4. DescribeVolumes 5. DeleteVolume 6. GetCallerIdentity 7. DetachVolume 8. DescribeInstanceCreditSpecifications 9. DescribeVpcs 10. DescribeInstanceTypes 11. DescribeTags 12. DescribeImages
CSC AWS EC2 Linux with security groups	AWS EC2 Linux with security groups- provision	1. DescribeVolumes 2. DescribeImages 3. DescribeInstances 4. DescribeInstanceCreditSpecifications 5. DescribeInstanceAttribute 6. DescribeVolumes 7. DescribeTags 8. DescribeVpcs 9. DescribeInstanceTypes 10. RunInstances 11. GetCallerIdentity
	AWS EC2 Linux with security groups stop	1. DescribeInstances 2. StopInstances
	AWS EC2 Linux with security groups start	1. DescribeInstances 2. DescribeImages 3. StartInstances
	AWS EC2 Linux with security groups stack deprovision- terminate all resources	1. DescribeInstances 2. TerminateInstances 3. ModifyInstanceAttribute 4. GetCallerIdentity 5. DescribeInstanceCreditSpecifications 6. DescribeInstanceAttribute 7. DescribeVolumes 8. DescribeTags

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Catalog item	Services	Permissions
		9. DescribeInstanceTypes 10. DescribeVpcs 11. DescribeImages

Amazon Web Services EC2 Microsoft Windows Out Of Box Catalog items

Cloud Services Catalog Windows VM, up to 10 additional disks on EC2 or with security groups.

This is a fully functional sample catalog item that can be used to request a Windows instance of any instance type from Amazon EC2. These mandatory steps are required for provisioning:

1. In the Terraform box, add the private key pem file used for provisioning.
2. Copy the path of the PEM file from Terraform box and update that path in the default value of the Management keypath variable.

Order	Features
Order a Microsoft Windows virtual machine on EC2. Cloud configuration scan is run for compliance checking.	For instance types other than t2.nano and t2.micro, approval is needed from administrator, by policy. The work-flow automation runs a CCG scan on the deployed Linux instance and checks the policy rule with the configuration of the VM.
Order a Windows virtual machine on EC2 and specify up to 10 additional disks. For up to two disks, no approval is necessary, but for anything beyond that, approval will be required.	This catalog item form allows ordering up to 10 additional disks (volumes) attached to the VM of varying sizes and types. If more than two additional disks are added to the request, seek approval from administrator, by policy.
Order a secure Windows virtual machine on EC2, assign security groups. Select as many required for the highest level of security.	This catalog item form allows capability to assign one or more security groups (firewall-port rules) to the compute instance so that the network access to the VM is restricted. Security groups are listed for the selected VPC and subnets.

Note:

Tag (key-value) is assigned to all deployed resources. Example: VM, network, storage as present in the stack. The key-value is updated only in the cmdb_key_value table, not in public cloud at this time. This will be fixed in future releases to update the tags in cloud.

The provisioning work-flow will notify the windows admin password to the requester via email. It can take a minimum of four minutes to obtain the windows admin password from Amazon Web Services.

Amazon Web Services EC2 Windows Out Of Box permissions

Amazon Web Services EC2 Windows OOB catalog items permissions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Catalog item	Services	Permissions
CSC AWS EC2 Microsoft Windows VM	AWS EC2 provision	1. GetPasswordData 2. DescribeInstanceCreditSpecifications 3. DescribeInstanceAttribute 4. DescribeInstances 5. DescribeTags 6. DescribeInstanceTypes 7. DescribeVolumes 8. DescribeVpcs 9. RunInstances 10. GetCallerIdentity 11. DescribeImages
	AWS EC2 instance- stop	1. DescribeInstances 2. StopInstances
	AWS EC2 instance- start	1. DescribeImages 2. DescribeInstances 3. StartInstances
	AWS EC2 instance- deprovision	1. DescribeInstances 2. GetCallerIdentity 3. ModifyInstanceAttribute 4. TerminateInstances 5. GetPasswordData 6. DescribeInstanceAttribute 7. DescribeInstanceCreditSpecifications 8. DescribeImages 9. DescribeTags 10. DescribeVolumes 11. DescribeInstanceTypes 12. DescribeVpcs
CSC CSC AWS EC2 Windows multiple disks	AWS EC2 Windows multiple disks provision	1. DescribeVolumes 2. DescribeVolumes 3. AttachVolume 4. DescribeInstances

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Catalog item	Services	Permissions
		5. GetPasswordData 6. CreateVolume 7. DescribeInstanceAttribute 8. DescribeInstanceCreditSpecifications 9. DescribeTags 10. DescribeVpcs 11. DescribeInstanceTypes 12. RunInstances 13. DescribeImages 14. GetCallerIdentity
	Amazon Web Services EC2 Windows multiple disks stop	1. DescribeInstances 2. StopInstances
	AWS EC2 Windows multiple disks start	1. DescribeImages 2. DescribeInstances 3. StartInstances
	AWS EC2 Windows multiple disks deprovision	1. DescribeInstances 2. TerminateInstances 3. ModifyInstanceAttribute 4. DescribeVolumes 5. DeleteVolume 6. DetachVolume 7. GetCallerIdentity 8. DescribeVpcs 9. DescribeInstanceAttribute 10. DescribeInstanceTypes 11. DescribeTags 12. DescribeInstanceCreditSpecifications 13. GetPasswordData 14. DescribeImages
CSC AWS EC2 Windows with security groups	AWS EC2 Windows with security groups- provision	1. GetPasswordData 2. DescribeVolumes 3. DescribeInstanceAttribute 4. DescribeInstance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Catalog item	Services	Permissions
		5. DescribeInstanceTypes 6. DescribeVpcs 7. DescribeTags 8. DescribeInstanceCreditSpecifications 9. RunInstances 10. GetCallerIdentity 11. DescribeImages
	AWS EC2 Windows with security groups stop	1. DescribeInstances 2. StopInstances
	AWS EC2 Windows with security groups start	1. DescribeInstances 2. DescribeImages 3. StartInstances
	AWS EC2 Windows with security groups deprovision	1. DescribeInstances 2. TerminateInstances ModifyInstanceAttribute 3. GetCallerIdentity 4. DescribeInstanceAttribute 5. DescribeInstanceCreditSpecifications 6. DescribeVpcs 7. DescribeVolumes 8. DescribeTags 9. GetPasswordData 10. DescribeImages 11. DescribeInstanceTypes

Amazon Web Services DynamoDB Out Of Box Catalog Items

Cloud Services Catalog AWS DynamoDB table for highly scalable NoSQL database.

This is a fully functional sample catalog item that can be used to order a Dynamo DB table- a name-value database table on Amazon Web Services or AWS cloud. The standard inputs to be provided are the table name, and the partition key as well as sort key names. The details of the dynamo DB will appear in the stack details returned.

Order an AWS DynamoDB table, which is a key-value store, with primary keys and secondary indices. This is useful for developers who require a highly scalable and fully managed NoSQL database.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon Web Services DynamoDB Out Of Box permissions

Amazon Web Services DynamoDB OOB catalog items permissions. DynamoDB supports provision and deprovision only.

Catalog item	Services	Permissions
Cloud Services Catalog AWS DynamoDB	AWS DynamoDB provision	**1.** DescribeTimeToLive **2.** DescribeTable **3.** DescribeContinuousBackups **4.** ListTagsOfResource **5.** UpdateTimeToLive **6.** UpdateContinuousBackups **7.** GetCallerIdentity **8.** CreateTable **9.** GetCallerIdentity
	AWS DynamoDB deprovision	**1.** DescribeTable **2.** DeleteTable **3.** GetCallerIdentit **4.** ListTagsOfResource **5.** DescribeContinuousBackups **6.** DescribeTimeToLive **7.** DescribeTable

Amazon Web Services S3 Bucket Out Of Box Catalog items

Cloud Services Catalog AWS S3 Bucket with secure access and encryption.

This is a fully functional sample catalog item that can be used to order an object storage bucket from Amazon cloud's S3 service. Validation is done for the bucket name, to ensure that the s3 naming policy is adhered to. The underlying Terraform template is used to prevent public access to the bucket and setting up standard encryption on data.

Order an AWS S3 storage bucket that includes security measures such as the prevention of public access and server-side encryption that uses Amazon S3 managed keys (SSE-S3). These security features will ensure that data stored in the bucket is protected and secure.

Amazon Web Services S3 Bucket Out Of Box permissions

AWS S3 Bucket OOB catalog items permissions. S3 Bucket supports provision and deprovision only.

Catalog item	Services	Permissions
CSC AWS S3 Bucket	AWS S3 Bucket provision	**1.** GetBucketPolicy **2.** PutBucketPublicAccessBlock

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Catalog item	Services	Permissions
		3. GetBucketLifecycle 4. GetBucketReplication 5. GetBucketTagging 6. CreateBucket 7. GetBucketLogging 8. GetBucketPublicAccessBlock 9. GetBucketRequestPayment 10. GetBucketAcl 11. GetBucketWebsite 12. GetAccelerateConfiguration 13. GetBucketVersioning 14. GetBucketEncryption 15. GetBucketCors 16. GetBucketObjectLockConfiguration 17. PutBucketEncryption 18. GetCallerIdentity
	AWS S3 Bucket deprovision	1. DeleteBucketEncryption 2. DeleteBucketPublicAccessBlock 3. DeleteBucket 4. GetBucketObjectLockConfiguration 5. GetBucketLifecycle 6. GetBucketEncryption 7. GetBucketReplication DescribeTable

Google Cloud Platform (GCP) Big Query

Cloud Services Catalog with Google Cloud Platform BigQuery Serverless data warehouse.

BigQuery Serverless Data Warehouse is a fully functional sample catalog item that streamlines the process of ordering a data warehouse solution powered by BigQuery from Google Cloud Provider.

Order	Features
Order a serverless, cost-effective, multi-cloud data warehouse with BiqQuery for streaming data in real time with built-in machine learning, and share analytics with just a few clicks.	This fully functional sample catalog item streamlines ordering a serverless, cost-effective, multi-cloud data warehouse solution powered by BigQuery from the Google Cloud Provider.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Order	Features
	You can set up real-time data streaming and analytics sharing effortlessly with a few clicks. BigQuery's serverless architecture allows SQL-based data analysis and offers customization options for scalability.

Google Cloud Platform (GCP) Big Query Out Of Box permissions

GCP Big Query Out Of Box catalog items permissions

GCP permissions

Catalog item	Services	Permission
CSC GCP BigQuery	Provision	• bigquery.datasets.create • bigquery.tables.create • storage.objects.get
	Pointed discovery	• bigquery.tables.list • bigquery.datasets.get
	Deprovision	• bigquery.tables.delete • bigquery.datasets.delete • bigquery.routines.delete • bigquery.models.delete

Google Cloud Platform Cloud SQL

Cloud Services Catalog with GCP Cloud SQL Database.

Cloud SQL database is a sample catalog item you can use to order a Cloud SQL database solution.

Order	Features
Order a Cloud SQL database solution that integrates with Google Cloud services.	This is a fully functional sample catalog item that can be used to order a Cloud SQL database solution. It integrates with Google Cloud services to create and connect to your first scalable database in minutes. Upon completion, comprehensive details of the Cloud SQL configuration will be provided in the stack details. It automates essential functions such as backups, replication, patching, encryption, and storage capacity adjustments, ensuring that applications

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Order	Features
	benefit from enhanced reliability, scalability, and security measures.

Google Cloud Platform Cloud SQL Out Of Box permissions

Google Cloud Platform Cloud SQL Out Of Box catalog items permissions

GCP Permissions

Catalog item	Services	Permission
CSC GCP Cloud SQL	Provision	• cloudsql.instances.create • cloudsql.instances.get • cloudsql.users.delete
	Pointed discovery	• cloudsql.databases.list • cloudsql.instances.list • cloudsql.users.list • compute.regions.list • compute.zones.list
	Deprovision	cloudsql.instances.delete

Google Cloud Platform Linux Out Of Box Catalog items

Cloud Services Catalog Linux VM with Agent Client Collector (ACC), up to 10 additional disks or secure Linux deployment.

This is a fully functional sample catalog item that can be used to request a Linux instance of any instance type from Google Cloud Platform.

Linux VM orders and features:

Order	Features
CSC GCP Linux VM: Order a Linux Virtual machine that allows installation of Agent Client Collector. ACC is compulsory for production instances. VM sizes other than e2-micro and e2-small require approval.	This is a fully functional sample catalog item that can be used to request a Linux instance of any instance type from Google Cloud Provider. For instance types other than e2-micro and e2-small, approval is needed from administrator, by policy. The catalog item has built-in client rules to drive an ACC (Agent Client Collector) installation, when the deployment environment is marked for production. To select the credential alias key for ACC deployment from mid, the appropriate ACC admin role permissions must be provided to the user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Order	Features
CSC GCP Linux multiple disks: Order a Linux virtual machine and specify up to 10 additional disk. No approval is necessary for up to two disks, but for anything beyond that, approval will be required.	This catalog item form allows ordering up to 10 additional disks (volumes) attached to the VM of varying sizes and types. If more than two additional disks are added to the request, seek approval from administrator, by policy.
CSC GCP Linux Cloud Firewall: Order a secure Linux virtual machine with enhanced security. Control traffic to and from your VM instances based on your configuration.	This is a fully functional sample catalog item that can be used to request a Linux instance of any instance type from Google Cloud Provider. This catalog item form allows you to assign one or more security groups (firewall-port rules) to the compute instance. This restricts network access to the VM. Security groups are listed for the selected VPC and subnets.

Note:

Tag (key-value) is assigned to all deployed resources. Example: VM, network, storage as present in the stack.

The key-value is updated only in the cmdb_key_value table, not in public cloud at this time. This will be fixed in future releases to update the tags in cloud.

Google Cloud Platform Linux Out Of Box permissions

Google Cloud Platform Linux Out Of Box catalog items permissions

GCP permissions

Catalog item	Services	Permission
CSC GCP Linux VM	Provision	• compute.disks.create • compute.disks.get • compute.instances.create • compute.instances.get • compute.instances.setMetadata • compute.networks.get • compute.subnetworks.get • compute.subnetworks.use • compute.subnetworks.useExternalIp • compute.zones.get
	Pointed discovery	• compute.disks.list • compute.images.list • compute.instanceTemplates.list • compute.instances.list

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP permissions (continued)

Catalog item	Services	Permission
		• compute.machineTypes.list • compute.snapshots.list • compute.zones.list • compute.regions.list
	Day 2 VM: Virtual Server - StopNode	• compute.instances.stop • compute.regions.get • compute.machineTypes.get
	Virtual Server - StartNode	compute.instances.start
	Virtual Server - Resize	• compute.instances.stop • compute.machineTypes.get • compute.regions.get • compute.instances.setMachineType • compute.instances.start
	Virtual Server - Deprovison	• compute.instances.delete • compute.regions.get
	Stack Deprovision	• compute.instances.delete • compute.disks.delete **Note:** Day 2 disk permissions are also required. These are same as Linux mutiple disks.
CSC GCP Linux Multiple Disks	Provision	• compute.disks.create • compute.disks.get • compute.disks.use • compute.instances.create • compute.instances.get • compute.instances.setMetadata • compute.networks.get • compute.subnetworks.get • compute.subnetworks.use • compute.subnetworks.useExternalIp

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP permissions (continued)

Catalog item	Services	Permission
		• compute.zones.get • compute.zoneOperations.get
	Pointed discovery	• compute.disks.list • compute.images.list • compute.instanceTemplates.list • compute.instances.list • compute.machineTypes.list • compute.snapshots.list • compute.zones.list • compute.regions.list
	Day 2 Disk storage Storage -Detach	compute.instances.detachDisk
	Storage Deprovision	compute.disks.delete
	Storage Attach disk	compute.instances.attachDisk
	Stack Deprovision	• compute.instances.delete • compute.disks.delete **Note:** Day 2 VM permissions are also required. These are same as basic Linux VM.
CSC GCP Linux Secure Deployment	Provision	• compute.disks.create • compute.disks.get • compute.instances.create • compute.instances.get • compute.instances.setMetadata • compute.instances.setTags • compute.firewalls.create • compute.firewalls.get • compute.networks.get • compute.networks.updatePolicy • compute.subnetworks.get • compute.subnetworks.use

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP permissions (continued)

Catalog item	Services	Permission
		• compute.subnetworks.useExternalIp • compute.zones.get
	Pointed discovery	• compute.disks.list • compute.images.list • compute.instanceTemplates.list • compute.instances.list • compute.machineTypes.list • compute.snapshots.list • compute.zones.list • compute.regions.list • compute.networks.list • compute.firewalls.list
	Stack Deprovision	• compute.firewalls.delete • compute.networks.updatePolicy • compute.instances.delete • compute.disks.delete **Note:** Day 2 VM disk permissions are also required.

Google Cloud Platform Load Balancer

Cloud Services Catalog with Google Cloud Platform HTTP Load Balancer for high availability.

This is a sample catalog item designed to streamline the process of provisioning an HTTP load balancer from the Google Cloud Provider. Its primary objective is to mitigate performance issues and enhance user experience by efficiently managing incoming traffic.

Order	Features
Order a HTTP load balancer to distribute user traffic across multiple instances of your applications to reduce the risk of performance issues.	This fully functional catalog item creates a Classic Application Load Balancer - Global. Only external load balancers are available as global load balancers. Choose a global load balancer when you want the load balancer to be distributed globally. Each global load balancer is distributed across all regions. It is unaffected by zonal and regional outages.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Google Cloud Platform Load Balancer Out Of Box permissions

Google Cloud Platform HTTP Load Balancer Out Of Box permissions

GCP permissions

Catalog item	Services	Permission
CSC GCP HTTP Load Balancer	Provision	• compute.backendServices.create • compute.backendServices.get • compute.backendServices.use • compute.globalAddresses.create • compute.globalAddresses.get • compute.globalAddresses.use • compute.globalForwardingRules.create • compute.globalForwardingRules.get • compute.healthChecks.create • compute.healthChecks.get • compute.healthChecks.useReadOnly • compute.instanceGroups.create • compute.instanceGroups.get • compute.instanceGroups.update • compute.instanceGroups.use • compute.instances.use • compute.targetHttpProxies.create • compute.targetHttpProxies.get • compute.targetHttpProxies.use • compute.urlMaps.create • compute.urlMaps.get • compute.urlMaps.use • compute.networks.get
	Pointed discovery	• compute.backendServices.list • compute.forwardingRules.list • compute.healthChecks.list • compute.httpHealthChecks.list • compute.httpsHealthChecks.list • compute.instanceGroupManagers.list • compute.instanceGroups.list • compute.instances.list • compute.targetHttpProxies.list

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP permissions (continued)

Catalog item	Services	Permission
		• compute.targetHttpsProxies.list • compute.targetSslProxies.list • compute.targetTcpProxies.list • compute.urlMaps.list • compute.regions.list • compute.zones.list
	LB- Deprovision	• compute.globalForwardingRules.delete • compute.targetHttpProxies.delete • compute.urlMaps.delete
	Stack - Deprovision	• compute.globalAddresses.delete • compute.backendServices.delete • compute.healthChecks.delete • compute.instanceGroups.delete • compute.globalForwardingRules.delete • compute.targetHttpProxies.delete • compute.urlMaps.delete

Google Cloud Platform Windows Out Of Box Catalog items

Cloud Services Catalog Windows VM, up to 10 additional disks on Google Cloud Platform or with cloud firewall.

This is a fully functional sample catalog item that can be used to request a Windows instance of any type from Google Cloud Platform.

Linux VM orders and features:

Order	Features
CSC GCP Windows VM: Order a Windows Virtual machine that allows installation of Agent Client Collector. VM sizes other than e2-micro and e2-small require approval.	This is a fully functional sample catalog item that can be used to request a Windows instance of any instance type from Google Cloud Provider. For instance types other than e2-micro and e2-small, approval is needed from administrator, by policy. The Windows credentials are provided by the user in catalog order form. The credentials to meet complexity requirement based on windows version. The workflow automation runs a CCG scan on the deployed Windows instance, and checks the policy rule

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Order	Features
	with the configuration of the VM. If there is any violation, the stack will be set with 'Follow up needed' state to notify that the stack deployment is not as per norms.
CSC GCP Windows multiple disks: Order a Windows virtual machine and specify up to 10 additional disk. No approval is necessary for up to two disks, but for anything beyond that, approval will be required.	This is a fully functional sample catalog item that can be used to request a Windows instance of any instance type from Google Cloud Provider. This catalog item form allows ordering up to 10 additional disks (volumes) attached to the VM of varying sizes and types. If more than two additional disks are added to the request, approval is needed from administrator, by policy. The Windows credentials are provided by the user in catalog order form. The credentials to meet complexity requirement based on windows version.
CSC GCP Windows Cloud Firewall: Order a secure Windows virtual machine with enhanced security. Control traffic to and from your VM instances based on your configuration.	This is a fully functional sample catalog item that can be used to request a Windows instance of any instance type from Google Cloud Provider. You can assign one or more security groups (firewall-port rules) to the compute instance to restrict network access to the VM. Security groups are listed for the selected VPC and subnets. The Windows credentials are provided by the user in catalog order form. The credentials to meet complexity requirement based on windows version.

Note:

Tag (key-value) is assigned to all deployed resources. Example: VM, network, storage as present in the stack.

The key-value is updated only in the cmdb_key_value table, not in public cloud at this time. This will be fixed in future releases to update the tags in cloud.

CSC GCP Windows OOB permissions

Google Cloud Platform Windows Out Of Box permissions

GCP permissions

Catalog item	Services	Permission
CSC GCP Windows VM	Provision	• compute.disks.create • compute.disks.get • compute.instances.create • compute.instances.get • compute.instances.setMetadata • compute.networks.get • compute.subnetworks.get • compute.subnetworks.use

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP permissions (continued)

Catalog item	Services	Permission
		• compute.subnetworks.useExternalIp • compute.zones.get
	Pointed discovery	• compute.disks.list • compute.images.list • compute.instanceTemplates.list • compute.instances.list • compute.machineTypes.list • compute.snapshots.list • compute.zones.list • compute.regions.list
	Day 2 VM: Virtual Server - StopNode	• compute.instances.stop • compute.regions.get • compute.machineTypes.get
	Virtual Server - StartNode	compute.instances.start
	Virtual Server - Resize	• compute.instances.stop • compute.machineTypes.get • compute.regions.get • compute.instances.setMachineType • compute.instances.start
	Virtual Server - Deprovison	• compute.instances.delete • compute.regions.get
	Stack Deprovision	• compute.instances.delete • compute.disks.delete **Note:** Day 2 disk permissions are also required. These are same as Windows mutiple disks.
CSC GCP Windows Multiple Disks	Provision	• compute.disks.create • compute.disks.get • compute.disks.use

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP permissions (continued)

Catalog item	Services	Permission
		• compute.instances.create • compute.instances.get • compute.instances.setMetadata • compute.networks.get • compute.subnetworks.get • compute.subnetworks.use • compute.subnetworks.useExternalIp • compute.zones.get • compute.zoneOperations.get
	Pointed discovery	• compute.disks.list • compute.images.list • compute.instanceTemplates.list • compute.instances.list • compute.machineTypes.list • compute.snapshots.list • compute.zones.list • compute.regions.list
	Day 2 Disk storage Storage -Detach	compute.instances.detachDisk
	Storage Deprovision	compute.disks.delete
	Storage Attach disk	compute.instances.attachDisk
	Stack Deprovision	• compute.instances.delete • compute.disks.delete **Note:** Day 2 VM permissions are also required. These are same as basic Windows VM.
CSC GCP Windows Secure Deployment	Stack Deprovision	• compute.firewalls.delete • compute.networks.updatePolicy • compute.instances.delete • compute.disks.delete **Note:** Day 2 VM disk permissions are also required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Azure Linux Out Of Box Catalog items

Cloud Services Catalog Linux VM with agent client collector (ACC), up to 10 additional disks, security groups or with scalable web servers or with security groups.

This is a fully functional sample catalog item that can be used to request a Linux instance of any instance type from Microsoft Azure.

Order	Features
Order a Linux virtual machine on Microsoft Azure, that allows installation of agent client collector. ACC is compulsory for production instances. VM sizes other than Standard_A1_v2 and Standard_A2_v2 require approval. Cloud configuration scan is run for compliance checking.	For instance types other than Standard_A1_v2 and Standard_A2_v2, approval is needed from administrator by policy. The catalog item has built in client rules to drive a ACC (Agent Client Collector) installation, when the deployment environment is marked for production. To select the credential alias key for ACC deployment from mid, the appropriate ACC admin role permissions are required to be bestowed to the user. The work-flow automation runs a CCG scan on the deployed Linux instance, and checks the policy rule with the configuration of the VM. If there is any violation, the stack will be set with 'Follow up required' state to notify that the stack deployment is not as per norms.
Order a Linux virtual machine and specify up to 10 additional disks. For up to two disks, no approval is necessary. Seek approval for anything beyond that.	This catalog item form allows ordering up to 10 additional disks (volumes) attached to the VM of varying sizes and types. If more than two additional disks are added to the request, seek approval from administrator, by policy.
Order a secure Linux virtual machine, assign security groups. Select as many you require, for the highest level of security.	This catalog item form allows capability to assign one or more security groups (firewall-port rules) to the compute instance so that the network access to the VM is restricted. Security groups are listed for the selected VPC and subnets.
Order a Linux virtual machine scale set to deploy identical virtual machines, enabling automated scaling and ensuring high availability.	This catalog item provides the capability to deploy identical virtual machines using a Linux virtual machine scale set, which supports automated scaling for enhanced availability. Email URL is added as a post provision operation which emails the deployed application's end point URL to the requester.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

Tag (key-value) is assigned to all deployed resources. Example: VM, network, storage as present in the stack.

The key-value is updated only in the cmdb_key_value table, not in public cloud at this time. This will be fixed in future releases to update the tags in cloud.

Microsoft Azure Windows Out Of Box Catalog items

Cloud Services Catalog Windows VM, up to 10 additional disks on Microsoft Azure or with security groups.

This is a fully functional sample catalog item that can be used to request a Windows instance of any instance type from Microsoft Azure.

Order	Features
Order a Microsoft Windows VM that allows installation of agent client collector. VM sizes other than Standard_A1_v2 and Standard_A2_v2 require approval. Cloud configuration scan is run for compliance checking.	For instance types other than Standard_A1_v2 and Standard_A2_v2, approval is needed from administrator, by policy. The work-flow automation runs a CCG scan on the deployed Linux instance, and checks the policy rule with the configuration of the VM.
Order a Windows virtual machine on Microsoft Azure and specify up to 10 additional disks. For up to two disks, no approval is necessary, but for anything beyond that, approval will be required.	This catalog item form allows ordering up to 10 additional disks (volumes) attached to the VM of varying sizes and types. If more than two additional disks are added to the request, seek approval from administrator, by policy. The Windows admin password is obtained from Microsoft Azure at the time of creation.
Order a secure Windows virtual machine on Microsoft Azure, assign security groups. Select as many required for the highest level of security.	This catalog item form allows capability to assign one or more security groups (firewall-port rules) to the compute instance so that the network access to the VM is restricted. Security groups are listed for the selected VPC and subnets.

Note:

Tag (key-value) is assigned to all deployed resources. Example: VM, network, storage as present in the stack. The key-value is updated only in the cmdb_key_value table, not in public cloud, at this time. This will be fixed in future releases to update the tags in cloud.

Microsoft Azure Out Of Box permissions

Cloud Services Catalog Microsoft Azure Out Of Box catalog items permissions.

These catalog items are available out of box in CSC Content Pack applications.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Services	Resources
Microsoft.Compute	1. virtualMachines/* 2. disks/* 3. availabilitySets/* 4. images/* 5. virtualMachineScaleSets/* 6. locations/* 7. vmSize/* 8. register/action/* 9. unregister/action/* 10. capacityReservationGroups/* 11. cloudServices/* 12. diskAccesses/privateEndpointConnectionsApproval/* 13. diskEncryptionSets/* 14. galleries/* 15. hostGroups/* 16. proximityPlacementGroups/* 17. restorePointCollections/* 18. sshPublicKeys/* 19. sharedVMExtensions/* 20. sharedVMImages/* 21. skus/read snapshots/*
Microsoft.Network	1. register/action/* 2. unregister/action/* 3. checkTrafficManagerNameAvailability/action/* 4. internalNotify/action/* 5. getDnsResourceReference/action/* 6. checkFrontDoorNameAvailability/action/* 7. privateDnsZonesInternal/action/* 8. applicationGateways/* 9. expressRouteCircuits/* 10. loadBalancers/* 11. networkInterfaces/* 12. networkSecurityGroups/* 13. privateDnsZones/* 14. publicIPAddresses/*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Services	Resources
	15. locations/* 16. virtualNetworks/* 17. virtualNetworks/subnets 18. virtualNetworks/virtualNetworkPeerings/* 19. natGateways/* 20. localnetworkgateways/* 21. connections/* 22. routeTables/* 23. subnets/*
Microsoft.Storage	1. register/action/* 2. unregister/action/* 3. storageAccounts/* 4. blobs/*
Microsoft.Resources	1. tags/* 2. resourceGroups/* 3. subscriptions/* 4. resources/* 5. providers deployments/*
Microsoft.Sql	1. servers/* 2. servers/databases/* 3. locations/*/read
Microsoft.ResourceHealth	AvailabilityStatuses/*
Microsoft.ContainerService	1. managedClusters/* 2. fleets/*
Microsoft.Web	1. serverFarms/* 2. sites/*
Microsoft.ResourceGraph	resources/*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Out Of Box Catalogs using Cloud Services Catalog

Microsoft Azure Functions App Catalog items

Cloud Services Catalog Microsoft Azure Function App with pay-as-you-go pricing benefit.

This is a fully functional sample catalog item that enables users to request a Windows instance of any instance type from Microsoft Azure.

Order	Features
Order a function app to build and deploy small, event-driven functions without managing infrastructure. It will offer a server-less computing experience with pay-as-you-go pricing.	This catalog item deploys a NodeJS basic "hello world" application onto the provisioned Microsoft Azure function app. The versions of NodeJS supported are the same as the ones supported in Azure portal. The Windows admin password is obtained from Microsoft Azure at the time of creation.

Note:

Tag (key-value) is assigned to all deployed resources. Example: VM, network, storage as present in the stack. The key-value is updated only in the cmdb_key_value table, not in public cloud at this time. This will be fixed in future releases to update the tags in cloud.

Out Of Box Actions

The Out of the Box Actions comprise OOB deployment, tagging, retrieving and post-provisioning operations. All these actions are based on Flow Designer and implemented as Integration Hub Sub-flows.

Actions	Flow Designer Sub-flow Name	Description
Deploy ACC-* in post provisioning	CSC Content- Post Provisioning- Deploy ACC in Stack VMs	ACC (agent client collector) is deployed through a Flow Designer action. This is called as post-provision step for the Linux basic VM deployment catalog item. Pre-requisites: 1. ACC Listener must be configured in any MID Server. This will generate a mid web server endpoint and a MID Server API Key. 2. In the post provisioning operation parameters of the Basic Linux VM catalog, these two parameters – web server end point (typically in format wss://<IP_Address>:3389/ws/events) and MID Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Actions	Flow Designer Sub-flow Name	Description
		API Keys, need to be set. Out of the box, these parameters have empty values. **Note:** If the generated WS endpoint URL returns the private IP of the MID Server, then that private IP needs to be replaced with public IP of the MID Server VM for the internet-based connectivity. **3.** There should be connectivity over SSH port from the MID Server to the provisioned VM. If the agent fails to install on the provisioned VM, a task is created and the stack is set to a 'follow-up required' state. The stack automatically turns active after all follow-on tasks are closed.
CCG Scan in post provisioning	CSC Content- Create CCG Policy Scan on Stack Items	CCG Scan with the default OOB Policy Set in CCG is called with the simple windows and Linux deployment. When a CCG scan fails owing to violation, the stack is set to a follow-up needed state. The operator must ensure that the violation is remediated and close the violation task. Once all the follow-on tasks are closed, the stack automatically turns to active state and can be used by requester.
Tagging Action	CSC Content- Post Provisioning- Update Key Values in CMDB	All catalog items have a tagging action run as post-provision step. This action only updates the key-value entries in CMDB, not in cloud. The standard tags updated are StackName, Application, CostCenter and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Actions	Flow Designer Sub-flow Name	Description
		BusinessService. Fields with empty values will not be saved as key-value records.
Retrieve Windows Password	CSC Content- Post Provisioning- Retrieve Windows Password	All windows catalogs have a retrieve password action, which emails the Administrator, the password of the provisioned windows VM to the requester. The password is obtained from AWS after the deployment is complete using a ssh key ('pem') file to decrypt the password. Follow these pre-requisite steps to set up this pem file on the Terraform Linux server: **1.** Ensure that the SSH Private Key credential to the VM with Terraform CLI installed, is associated with the credential alias named 'TFO Server Credential Alias'. This alias is already created with installation of CSC Content pack. This alias is used to SSH into the terraform server to get the windows password from output variable. **2.** Ensure that in the Catalog's Provision variable set, set the 'keypath' variable default value to the management key path on the terraform VM. This key decrypts the password. **3.** Ensure that the requester has set up the email **Note:** For custom catalogs, this operation is available for use only, with Terraform Opensource- Linux based catalog items.
Email URL For Microsoft Azure Scalable Web Server Application	CSC Content- Email URL For Microsoft Azure Scalable Web Server Application	This action is added as a post provision operation in 'CSC Microsoft Azure Linux with

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Actions	Flow Designer Sub-flow Name	Description
		Scalable Web Server' catalog item. It emails the deployed application's end point URL to the requester. The URL is obtained from Microsoft Azure after the deployment is complete. Ensure the requester has set up the email.
Retrieve Microsoft Azure Function App Host name	CSC Content- Post Provisioning- Retrieve Microsoft Azure Function App Host name	This action is added as a post provision operation in 'CSC Microsoft Azure Function App catalog item which retrieves the deployed function app's host name to the requester.

Out Of Box Policies

The Cloud Services Catalog Out of the Box Policies comprise Approval, Naming, Lease and Quota Provisioning Policies.

Provisioning Policies	Description
Approval	ServiceNow Approval action called via Policy Trigger, on specific occasions. Example: When a requester orders VM of type other than t2.micro and t2.nano, or when more than two extra disks are requested for certain catalog items.
Naming Policy	A good approach for resources deployed as part of a single request is to tag the resources with a common name, such as the stack's name. If the stack name follows standard convention, for easy ownership identification, and for age, that is useful. The naming policy that is shipped out of the box is enabled for all cloud catalog items. When the policy is active, a stack name is suggested in a particular naming format for Stack Name field in the form. The naming format uses a combination of initial letters of the requesters name, followed by the current date and time in mm dd hh mm format. Example: if the requester name is John Smith, and the date/time is 10:30 am on 15-May, the stack name is generated as 'js05151030'.
Lease Policy	Lease Policy is used to trigger de-provisioning of unused/unwanted stacks to save costs on cloud. The lease policy will de-provision active stacks after the lease period expires. A week prior to this, reminders are sent to stack owners mentioning the stack is expiring.
Quota Policy	Quota Policy is used to decide if a request which is run once quota is exceeded, should go to approval or to be aborted. See Resource order controls for template-based catalog items

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CSC references

Reference topics provide additional information about the lists and forms that you use to configure, administer and use Cloud Services Catalog.

Domain separation and Cloud Services Catalog

Domain separation is supported in Cloud Services Catalog. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

Recommendations for domain separation in Cloud Services Catalog

Domain separation recommendations in Cloud Services Catalog:

CSC Guided setup

CSC Guided Setup should be executed solely by the system administrator within the global domain. It is important to note that the local domain is not the suited environment for this task.

Terraform Opensource in global domain

All the catalogs and config providers are installed in the global domain with the installation of Cloud Services Catalog plugin. The Out Of Box catalogs work only if you create the Terraform Opensource (TFO) configuration and credentials in the global domain. This also enables the Managed Service Provider (MSP) users to use single TFO server for multiple domains.

Service providers credentials

Service providers' (AWS, Microsoft Azure and Google Cloud Platform credentials should be domain specific. However, Azure DevOps) and Ansible config installables should be in global domain for provisioning.

MID Server in global domain

Set the MID server in the global domain as TFO and related config management is also set in the global context. This enables accurate global discovery.

MSP users can set global MID servers and also use domain specific MID for different service providers for Day 1.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Next steps

For more information on how domain separation works in Cloud Provisioning and Governance, visit Domain separation and Cloud Provisioning and Governance.

Related topics

Domain separation for service providers

Azure DevOps permissions

Azure DevOps discovery and provision permissions.

Permission type	Description	PAT Access
Cloud Services Catalog Azure DevOps Discovery	Discovers Projects and Pipelines in Azure DevOps Organisation & adds to CMDB.	# Build (Read)
Cloud Services Catalog Azure DevOps Provision	Provision via Azure DevOps pipeline triggered from Cloud Services Catalog DevOps Integration Catalog.	# Build (Read & Execute)
Cloud Services Catalog Azure DevOps tag:GetResources permission	Discover AWS resources provisioned through Azure DevOps pipeline.	Not applicable

Ansible permissions

Ansible Config Provider, Ansible Catalog and operations permissions.

Ansible Day 0, Day 1 and Day 2 permissions

Permission type	Description
Ansible Config Provider Discovery on Day 0	1. Configure Ansible credentials reference link: https://docs.ansible.com/ansible-tower/latest/html/userguide/credentials.html 2. Creating Ansible tower users: https://docs.ansible.com/ansible-tower/3.2.5/html/userguide/users.html 3. Roles/permissions can be granted to Ansible tower users per project/ inventory/ job template/credentials as required. For more details: https://docs.ansible.com/ansible-tower/latest/html/userguide/users.html#users-permissions 4. CAPI uses basic authentication token based mechanism which uses the Ansible tower user's credentials. OAuth tokens based on basic authentication credentials are generated. 5. Projects that contain the GitHub credentials of the playbook (this is where .yml files exist) is configured manually on Ansible. For more

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Ansible Day 0, Day 1 and Day 2 permissions (continued)

Permission type	Description
	information: https://docs.ansible.com/ansible-tower/latest/html/userguide/credentials.html#source-control 6. Uses Ansible credentials for Discovery of Inventory, projects, cfg installable– job templates (fetching playbooks and its extra variables) using CAPI.
Ansible Catalog provision/ de-provision on Day 1	Uses Ansible tower user credentials and service account credentials of Amazon Web Services/ Azure given in the Ansible job template.
Ansible tag:GetResources for discovery	Uses `tag:GetResources` permission to discover AWS resources provisioned through Ansible pipeline.
Ansible Day 2 operations	Uses Ansible tower user credentials, service account credentials of Amazon Web Services/ Azure and private ssh key pair or windows user name/password, to connect to the particular cloud resource provisioned.

Stages of Azure DevOps in Release life cycle management

The components and stages of an Azure DevOps pipeline are comprehensive and highly customizable. You can define multiple stages, tasks, and integrations based on your specific requirements and technology stack using Cloud Services Catalog.

Day 0 support	Day 1 support	Day 2 support
1. Scan the projects and pipelines 2. Gather pipeline variables from Azure DevOps	1. Order from Service Catalog 2. Reference from discovered Azure DevOps projects and pipelines 3. Select inputs needed for the release deployment pipeline 4. Submit catalog	1. SRE operations – patching, backup, etc, on production resources 2. Daily maintenance – restart, reboot, suspend, etc. 3. Tagging resources, Running hardening, etc.
Store in CMDB as referable objects	1. Wait for approval 2. Track the progress 3. Update CMDB with deployed resources	1. Ensure resilience in production. 2. Switch between blue and green environment

Workaround to known issues

The workaround steps to some of the known issues are compiled and described in the Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Workaround	Link
Cloud Services Catalog: Known Issues and Workarounds	https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB1362282

Related topics

https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB1362282

Cloud Migration Assessment

This application is no longer deployed, enhanced, or supported.

For details, see the Deprecation Process [KB0867184] article in the Now Support knowledge base.

Cloud Provisioning and Governance

The ServiceNow® Cloud Provisioning and Governance (CPG) application serves as a unified interface for accessing cloud resources, delivering cloud offerings to a catalog, and overseeing resource usage. This application is transformed as Cloud Services Catalog application, offering refined and streamlined management of usage and life cycle of cloud resources.

Important: Starting with the Brazil release, Cloud Provisioning and Governance is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

Cloud Provisioning and Governance is integrated with both private and public cloud providers, such as Amazon Web Services, Microsoft Azure, Google Cloud Platform and VMware.

Note:

Migration from workflow to subflows:

- If the default workflows aren't customized, the flows are migrated to subflows.
- If the default workflows are customized, the flows continue to use the existing workflow.

The Cloud Provisioning and Governance application is incompatible and cannot be used simultaneously with the legacy v1 Cloud Management application. The Cloud Provisioning and Governance (previously Cloud Management) application is compatible starting with the Jakarta release.

Workflow migration

As the automation experience evolves, legacy workflows are no longer supported starting with the Zurich release. You can continue using existing custom workflows and deprecated base system workflows only if already in place, but use the new base system flows for all future needs. For more information, see the following articles in the Now Support Knowledge Base:

- How to update Cloud Provision Service Request PAD or Cloud Resource Operation Request PAD [KB2265927]
- How to migrate Cloud Service Catalog to use flow in case any of the Out of Box Workflows has been modified [KB2262737]

All default predefined workflows have been redesigned using Workflow Studio.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Policy Rule Actions can now use Workflow Studio subflows instead of legacy workflows. If you have not customized the default workflows, the legacy workflows are migrated to subflows automatically. If you have customized the default workflows, they aren't migrated, and their use cases continue to work with the existing workflows.

The following workflows have been migrated to Workflow Studio flows or subflow:

- Cloud operation change request
- Cloud task processing
- Cloud approval workflow
- Post-process workflow
- Pre-process workflow
- Cloud resource operation request
- Blueprint request

The following workflows have been deprecated:

- Retrieve cloud billing data
- Retrieve cloud billing month
- Cloud operation step workflow launcher

Explore	Set up	Administer
• Brazil • Domain separation and Cloud Provisioning and Governance	• Day 1 setup guide for Amazon Web Services on Cloud Provisioning and Governance • Day 1 setup guide for Microsoft Azure Cloud on Cloud Provisioning and Governance • Day 1 setup guide for VMware on Cloud Provisioning and Governance	Cloud Provisioning and Governance administration guide

Use	Develop	Troubleshoot and get help
• Cloud Admin Portal • Cloud User Portal	• Developer training • Developer documentation	• Search the Known Error Portal for known error articles • Ask or answer questions in the ITOM community • Contact Customer Service and Support

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using guided setup to implement Cloud Provisioning and Governance

Cloud Provisioning and Governance guided setup provides a sequence of tasks that help you configure Cloud Provisioning and Governance on your instance. To open Cloud Provisioning and Governance guided setup, navigate to **Guided Setup** > **ITOM Guided Setup**.

For more information about using the guided setup interface, see Using guided setup.

Request the Cloud Provisioning and Governance application

The Cloud Provisioning and Governance application is available as a separate subscription and requires the Cloud Provisioning and Governance plugin (com.snc.cloud.mgmt).

Before you begin

Role required: admin

About this task

To purchase a subscription, contact your ServiceNow account manager. The account manager can arrange to have the plugin activated on your organization's production and subproduction instances, generally within a few days.

If you don't have an account manager, decide to delay activation after purchase, or want to evaluate the product on a subproduction instance without charge, follow these steps.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Related topics

- Day 1 setup guide for Amazon Web Services on Cloud Provisioning and Governance
- Day 1 setup guide for Microsoft Azure Cloud on Cloud Provisioning and Governance
- Day 1 setup guide for Google Cloud through Cloud Services Catalog Terraform Connector
- Day 1 setup guide for VMware on Cloud Provisioning and Governance

Product Name Change

Starting with the Paris release, the ServiceNow® Cloud Management application was called ServiceNow® Cloud Provisioning and Governance. The application is transformed into Cloud Services Catalog application providing improved and efficient management of the usage and life cycle of cloud resources.

The standalone subscription for Cloud Provisioning and Governance is discontinued. It is restructured and incorporated into Cloud Accelerate. Furthermore, it has been segmented into Cloud Migration Assessment, Cloud Account Management, Cloud Services Catalog, Cloud Configuration and Governance, and Cloud Action Library.

Product name references throughout the product documentation and the graphical user interface (GUI) are being updated. You might come across occasional references to their old names, based on the context. For some core components, the product name remains intentionally unchanged to maintain existing functionality.

Component name matrix

Component	Former name	New name
ITOM Optimization > Cloud Provisioning and Governance	Cloud Management application	Cloud Provisioning and Governance application
com.snc.cloud.mgmt	Cloud Management plugin	Cloud Provisioning and Governance plugin
com.snc.cloud.core	Cloud Management Core plugin	Cloud Provisioning and Governance Core plugin
Application Context (**Settings** > **Developer** > **Application** > **Cloud Provisioning and Governance**)	Cloud Management	Cloud Provisioning and Governance
com.sn_cmp_ibm	Cloud Management: IBM Cloud Connector	Cloud Provisioning and Governance: IBM Cloud Connector
MID Server Capability > Cloud Management	MID Server Capability > Cloud Management	Cloud Provisioning and Governance
Unchanged component names		
MID Server Application > Cloud Management	MID Server Application > Cloud Management	Unchanged

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Day 1 setup guide for Amazon Web Services on Cloud Provisioning and Governance

To set up Cloud Provisioning and Governance for the very first time, you perform the procedures in this "Day 1" setup guide. Be sure to perform the procedures in order. After you have performed Day 1 setup, you can perform optional Day 2 setup and configuration procedures as needed and in any order. Detailed instructions for each procedure follow this overview.

Request the Cloud Provisioning and Governance application

The Cloud Provisioning and Governance application is available as a separate subscription and requires the Cloud Provisioning and Governance plugin (com.snc.cloud.mgmt). See Request the Cloud Provisioning and Governance application.

Roles required to set up Cloud Provisioning and Governance

- Operations in the AWS Management Console require the administrator role.
- Operations in Cloud Provisioning and Governance require the sn_cmp.cloud_admin role.

About terms that Cloud Provisioning and Governance uses

Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings. In AWS, the region-specific containers for virtual resources are called regions. In Cloud Provisioning, regions are called datacenters or logical datacenters (LDCs). The term logical is used to reinforce the idea that Cloud Provisioning is provider-agnostic. All infrastructure or applications that are deployed using Cloud Provisioning are associated with a datacenter.

Quick overview of the setup process

Your setup process includes these tasks:

- If needed: Request the Cloud Provisioning and Governance application.
- Assign appropriate roles to cloud users.
- Set up the MID Servers that will handle secure communications with the provider API endpoints.
- On the provider portal, collect your account settings and the credentials that the Discovery process will use (through a MID Server) to programatically access your provider accounts. Securely associate the account settings and credentials with a service account in Cloud Provisioning and Governance.
- Set up a cloud account to represent your entire managed cloud infrastructure and set up a service account that works with one of your provider accounts. You specify which datacenters in the service account should be included in the cloud account. (Later, on "Day 2", you can set up additional cloud accounts and service accounts from the same or other providers.)
- Cloud Provisioning uses CMDB data to help users request and manage cloud resources and to help you manage your cloud infrastructure. To populate the CMDB with resource data for all datacenters, you manually run the Discovery process on each datacenter in the service account. Then, to ensure that the data continues to be updated, you configure Discovery to run on a regular schedule. Your cloud account might look like this:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Structure of a cloud account on Day 1

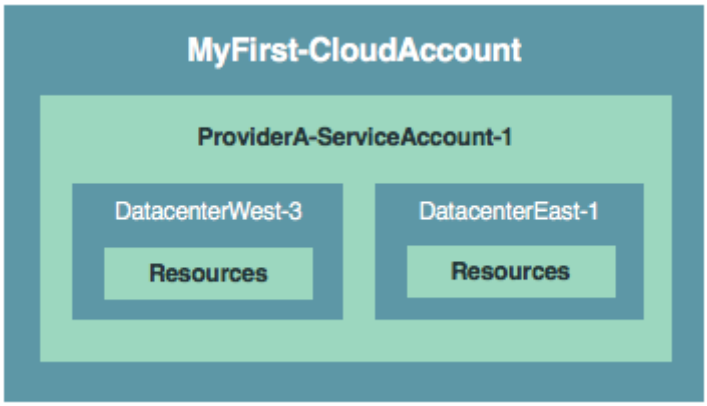

- Providers offer services that can auto-update the CMDB whenever a create/modify/terminate life-cycle change or configuration update occurs to a resource. You can configure the service to integrate directly with Cloud Provisioning and Governance.

What you will do to integrate Cloud Provisioning and Governance with your AWS accounts

Detailed instructions for each procedure follow this overview.

1. **Assign roles to AWS users of Cloud Provisioning and Governance**

   You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

2. **Install and configure MID Servers to access cloud environments**

   To ensure secure and reliable communications, the Discovery process communicates with your cloud provider accounts and cloud resources through one or more MID Servers. You can set up the MID Servers on your network or in one of your cloud networks. You can configure the MID Server for Cloud Provisioning and Governance to use a proxy server. Using a proxy server supports all cloud-based activities such as running Discovery, billing downloads, provisioning virtual machines, and running life-cycle operations on virtual machines.

   Note: Data is encrypted to the MID Server and between the MID Server and the API endpoint. To ensure high performance and security, you should configure one or more MID Server for each datacenter under management. Configure the MID Server even if you have already configured other MID Servers while setting up Cloud Provisioning and Governance for another cloud provider.

3. **Configure access to the AWS accounts using permanent AWS credentials**

   To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery and Cloud Provisioning and Governance, you first create a user with programmatic access in the AWS Management Console. You then securely store the credentials in a service account at ServiceNow AI Platform.

4. **Create AWS GovCloud credentials for Cloud Provisioning and Governance**

   Note: Skip this procedure if your organization does not use AWS GovCloud (US).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To securely access data on your provider account, the Discovery process must present appropriate credentials. An AWS GovCloud (US) region is an isolated AWS region that meets stringent US government security and compliance requirements to host sensitive workloads. Cloud Provisioning and Governance supports all AWS GovCloud (US) services.

5. Create a service account for AWS GovCloud

Note: Skip this procedure if your organization does not use AWS GovCloud (US).

If your organization uses AWS GovCloud (US) region, you must create a service account in the region where you provision the resources. These credentials that you create are used for Cloud Discovery, Cloud Provisioning and Governance, and Cloud Cost Management.

6. Set up a cloud account and service account for AWS

A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter. A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

8. Define the schedule for downloading AWS billing data

Define the scheduled job that regularly uses a MID Server to download billing data from the provider. Cloud Provisioning and Governance saves the data in a cost table and uses the information to generate reports.

Next steps

When you have finished all Day-1 and Day-2 procedures in this setup guide, see the Cloud Provisioning and Governance administration guide for information on using the Cloud Provisioning and Governance application in your organization.

Assign roles to AWS users of Cloud Provisioning and Governance

You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

Before you begin

Role required: user_admin or admin

Procedure

Assign the following roles to groups and users as appropriate:

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
Root administrator `[sn_cmp.cmp_root_admin]`	Highest level of application access for Cloud Provisioning and Governance.	All

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
Cloud administrator `[sn_cmp.cloud_admin]`	Configures the Cloud Provisioning and Governance application and sets up the cloud infrastructure.	Cloud infrastructure: Service accounts and cloud account Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • IBM Cloud Update • VMware Events
Governor `[sn_cmp.cloud_governor]`	Monitors overall cloud usage and enforces compliance of the organization's rules, quotas, and policies. Also manages tags and permissions to various objects.	Governance: • Policies • Pools • Quotas • Permissions
Service Designer `[sn.cmp.cloud_service_designer]`	Creates blueprints, ARM and CloudFormation templates, and catalog items.	Cloud Service Design: • Cloud templates • Blueprints • Blueprint catalog items • Resource blocks
Cloud user `[sn_cmp.cloud_service_user]`	Requests and manages stacks and resources.	Cloud User Portal. When you assign the role to a group, all members of the group share quota limitations and ownership of certain resources. You can access all the task records, assigned to you or otherwise, if you have the cloud service user role.
Cloud operator `[sn_cmp.cloud_operator]`	Monitors and troubleshoots the Cloud Provisioning and Governance application.	Dashboards and reports: • Cloud Operations Dashboard • Cloud Root Cause Analysis

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
		• Cloud Orchestration Trail • Cloud API Trail
Cloud Event Integration `[sn_cmp.cloud_event_integration]` **Note:** Not supported by IBM Cloud Connector.	Authorizes access to the instance for external services that auto-update the CMDB when cloud events occur. This role gives the access to the cloud event REST endpoint.	Cloud infrastructure: • Service accounts and cloud account • Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • VMware Events
Cloud Group administrator `[sn_cmp.cloud_group_admin]`	Grants admin access to any group that you belong to.	

Related topics

Assign a role to a group

Assign a role to a user

Set up a cloud account and service account for AWS

A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter. A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

Before you begin

Role required: none

- Operations in the AWS Management Console require the administrator role.
- Operations in Cloud Provisioning and Governance require the sn_cmp.cloud_admin role.

About this task

When you finish the Day-1 setup process, your cloud account might look like this:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Structure of a cloud account on Day 1

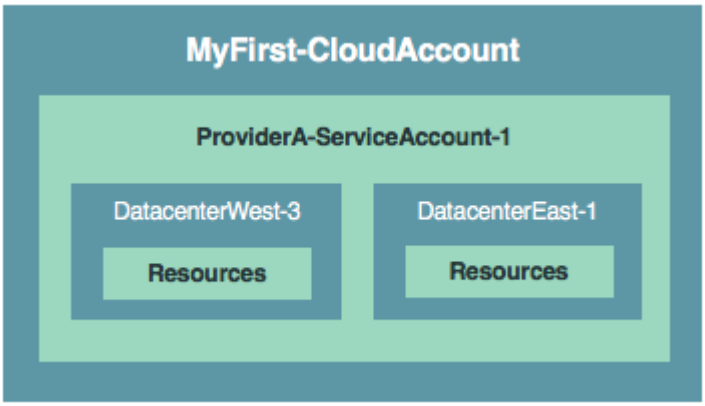

Procedure

1. On the Cloud Admin Portal, navigate to **Manage** > **Cloud Accounts**.
2. Select or create the cloud account:
   - Select an existing cloud account:
     a. On the Cloud Account page, click **Configure**.
     b. On the General Information page, click **Next** and then continue with the next step.
   - Create a cloud account:
     a. On the General Information page, click **New** and then enter a unique and meaningful **Name** and **Description** for the cloud account.
     b. Select the provider.
     c. Click **Next** and then continue with the next step.
3. On the Datacenters page, create a service account: Click the **+** next to the **Service Account** field, and then fill in the form for the service account.

Cloud Service Account form

Field	Description
Name	The unique and meaningful name for this service account.
Account ID	12-digit user account number. Expand the list under the AWS account name on the AWS Management Console to view the number. **Note:** In the **Account ID** field, remove the hyphen characters (-) from the number.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	IAM user account number 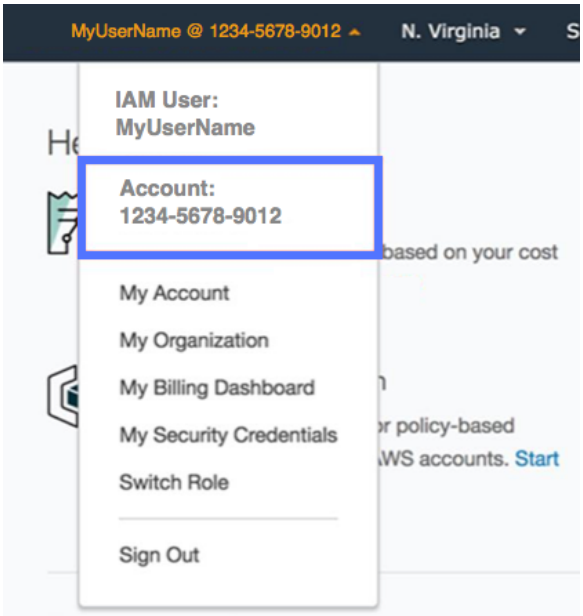
Discovery credentials	The credentials needed for ServiceNow applications to access this account. You may configure this field at a later stage, while configuring access to AWS accounts. ○ If you configured AWS credentials at ServiceNow AI Platform , select the magnifying glass icon, and then select the name of the relevant AWS credential. ○ To use other AWS accounts to access this account, leave the field blank. For example, you don't need to specify the AWS credentials for accounts assuming IAM roles or member accounts using their management account for access.
Datacenter URL	URL of the datacenter. This field is required for AWS China region and AWS GovCloud (US) accounts. For example: ○ AWS China region: `https://organizations.cn-northwest-1.amazonaws.com.cn` ○ AWS GovCloud (US): `https://ec2.us-gov-west-1.amazonaws.com`
Datacenter type	Type of the datacenter where the account is hosted. Select **AWS datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Parent account	Name of the management account (also referred to as the master account in some ServiceNow UI screens) that represents the organization in AWS that this member account belongs to. This field appears when you select **AWS datacenter**. Leave the field empty if this account is not part of an AWS organization.
Is master account	Management account flag. This field is also referred to as the master account flag in some ServiceNow UI screens. This field appears when you select **AWS datacenter** from the Datacenter Type list. Select the check box to associate the AWS service account with the management account. Select this check box only for accounts that you previously configured in the AWS Management Console as management accounts with some member accounts belonging to them. See the AWS documentation for information on AWS Organizations. **Note:** You must set up the correct permission in AWS or the Organization role for a standard credential. For more information, see Control AWS access and permissions using policies.
Accessor account	Name of the trusted account. Configure this field only for accounts that don't use permanent AWS credentials and rely on IAM roles for access.

4. Click **Submit** to create the service account.
5. Click **Discover Datacenters**.
6. When the datacenters appear, select one or more datacenters to include in the cloud account and then click **Save**.

   Note:
   - Select only those LDCs/Regions where your infrastructure resources are present. If you don't have any resources under these new data centers, exclude them while you run the full discovery.
   - To discover resources in certain regions or LDCs, raise an access request to the required endpoints with GCP Support. These regions include North America (US East, US West, and Canada Central), South America, Europe, Africa, Middle East, and Asia Pacific.

   The Cloud Account dashboard appears. The datacenters that you selected appear on the **Datacenters** tab.
7. Click a datacenter.
   The following lists appear:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Resources**: Cloud resources for the current datacenter. Run Discovery to populate the CMDB for the datacenter and populate the tab. See Discover all datacenters in a service account on-demand.
- **Capacity Limits**: Limits on virtual CPUs, virtual networks, storage volume size, and other services. See Set capacity limits on user requests for resources.
- **Discovery Log**: The process creates a log record for each action associated with a discovery status.

What to do next

- Configure additional permissions to provision Amazon AWS Cloud resources using Cloud Provisioning and Governance.
- After you complete all procedures in this initial "Day 1" setup, you can create additional cloud accounts and service accounts to organize and compartmentalize your cloud infrastructure.

Related topics

Set up an additional cloud account

Add an AWS service account to the cloud account

Configure a custom AWS member role

Customize the AWS roles that a MID Server can assume to receive temporary credentials for member accounts. You can configure additional parameters to improve security and customize the way that the member account's role is assumed when discovering cloud resources.

Before you begin

Role required: admin

About this task

Values in the Cloud Service Account > AWS Org Assume Role Params [cloud_service_account_aws_org_assume_role_params] table are passed as parameters to the AWS AssumeRole API.

Procedure

1. Navigate to **All** > **Cloud Provisioning and Governance** > **Organization Access Parameters** > **AWS Org Assume Role Parameters**.

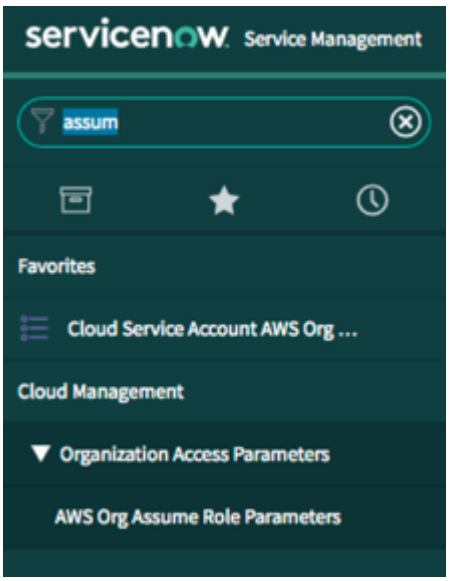

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Click **New** and then complete the form using these parameters:

Field	Description
Access role name [access_role_name]	Name of the AWS role in the member account which is used by the management account to acquire temporary credentials. Default: OrganizationAccountAccessRole
Role session name [role_session_name]	Name for the session using the temporary security credentials that may help in distinguishing use of a role by a principal or purpose. This session name is visible in the AWS Cloud Trail logs. See Cloud API Trail and the AWS documentation on AWS Cloud Trail for details. Default: master_account_id__<management account ID number> An example of this is: `master_account_id__321003876149`.
Credential TTL in seconds [credential_ttl_seconds]	Time in seconds for the temporary security credentials to live. Default: Calculated as follows: **a.** Retrieve the value in the *`mid.aws.sts.assume_role.credential_ttl_minutes`* MID Server property. **b.** Constrain this value to be between 15 and 720 minutes. If the setting in the property is less than 15 minutes, the system enters 15 minutes. If the setting is greater than 720 minutes, the system enters 720 minutes. **c.** Convert the resulting value into seconds.
External ID [external_id]	Unique identifier required by the trust policy of the role being assumed. Default: ServiceNow_MID_Server
Session policy [session_policy]	IAM policy in JSON format that further restricts the permissions of the temporary security credentials beyond the role configured policy. (JSON in AWS policy language.) Default: Blank
MFA [multifactor authentication]	Serial number of the Multi-Factor Authentication (MFA) device (hardware or virtual) used to authenticate the management account. Default: Blank
MFA token code [mfa_token_code]	Token code supplied by the MFA device (hardware or virtual) used to authenticate the management account. Default: Blank
Cloud service account [cloud_service_account]	Required. Service account to associate with the access parameters that you pass to the AWS AssumeRole API. Enter an account ID,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	either a management account or a member account, from the Service Accounts [cmdb_ci_cloud_service_account] table.

Note: For more details on how these parameters are used and what they mean, see the AWS documentation on the AWS Security Token Service API for the AssumeRole action.

Create AWS GovCloud credentials for Cloud Provisioning and Governance

Skip this procedure if your organization does not use AWS GovCloud (US). To securely access data on your provider account, the Discovery process must present appropriate credentials. An AWS GovCloud (US) region is an isolated AWS region that meets stringent US government security and compliance requirements to host sensitive workloads. Cloud Provisioning and Governance supports all AWS GovCloud (US) services.

Before you begin

Note:

- ServiceNow® ITOM products are not officially certified or supported against Germany and China regions.
- Skip this procedure if your organization does not use AWS GovCloud (US).

- You will need the AWS GovCloud (US) access key ID and the secret access key that you generated on the AWS Management Console.

Role required: sn_cmp.cloud_admin

About this task

Cloud Provisioning and Governance accesses GovCloud regions using a set of credentials for each region. To configure Cloud Provisioning and Governance to support a GovCloud region, you create one standard AWS account for each region (required for billing), obtain the credentials for the account, and then create a service account for the region. For more information on billing, see the "AWS GovCloud (US) Billing and Payment" page in the AWS documentation.

Procedure

1. On the AWS Management Console, enter `IAM` in the **AWS services** search box to open the Identity and Access Managements (IAM) service.
2. On the IAM Resources portal, click **Users**.
3. Create the Discovery user that has programmatic access to your AWS GovCloud (US) resource and billing data.
4. Select **Add user**.
   On the Details page, configure the user settings, and then select **Next**.

User name	Name for the programmatic user, for example, `servicenowcloud`.
Access type	Select **Programmatic access**.

5. On the Permissions page, attach the user to a policy by configuring the following settings and then click **Next**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set permissions for <username>	Select **Attach existing policies directly**.
Attach one or more policies …	Select the appropriate policy. **Note:** The **AdministratorAccess** policy has the most powerful permission level, including permission to provision cloud resources. The policy enables the same access that would be granted to the instance if you were not using IAM and used your AWS account Access Key ID and Secret Access Key. You might instead prefer to create a policy or combine multiple policies to grant the appropriate permission level. See Control AWS access and permissions using policies for details.

6. On the Review page, verify your selections and then click **Create user**.
7. On the Complete page, perform two steps:

   Note: Do not leave the page until you have completed both steps. The **Secret access key** value will not appear again. Paste the values that you generate in these steps into a Cloud Provisioning and Governance form.

   a. Click **Show** to display the Secret access key. Copy the value.
   b. Click **Download .csv** to save the CSV-format file that contains the user name, Access key ID, and the Secret access key value. You create the file as a backup in the case that you lose the values. Verify that the file was created and then store the file securely.
8. On your instance, navigate to **Orchestration** > **Credentials & Connections** > **Credentials**.
9. Click **New**, select **AWS Credentials**, enter a unique and meaningful **Name** (for example, `AWS GovCloud Creds 01`), and then fill in the form.

AWS Credentials form fields

Field	Input value
Name	Unique and descriptive name for the AWS credentials.
Active	Option to use the credential.
Access Key ID	The **Access key ID** that you generated on the AWS Management Console, such as: APIAIOSFODNN7EXAMPLE.
Secret access key	The **Secret access key** that you generated on the AWS Management Console, such as: wPalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Click **Submit**.
    The instance records the credentials. Next, you Create a service account for AWS GovCloud.

Create a service account for AWS GovCloud

If your organization uses AWS GovCloud (US) region, you must create a service account in the region where you provision the resources. These credentials that you create are used for Cloud Discovery, Cloud Provisioning and Governance, and Cloud Cost Management.

Before you begin

Note: Skip this procedure if your organization does not use AWS GovCloud (US).

- Role required: sn_cmp.cloud_admin
- AWS GovCloud (US) credentials for each GovCloud region.
- AWS GovCloud (US) account ID (from the AWS Management Console).

About this task

A service account holds the credential and account information that you created in your provider account. Discovery uses the information to access your provider account and then obtain information on each logical datacenter that is associated with the account.

Procedure

1. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal** > **Service Accounts**.
2. Click **New**, enter a unique and descriptive name for the account (for example, `AWS GovCloud SA 01`) and then fill in the Cloud Service Account form.

Field	Description
Account ID	Account ID to which this credential belongs.
Discovery credentials	Select the name of the credentials that you created in the Create AWS GovCloud credentials for Cloud Provisioning and Governance procedure. In the example, you used the name `AWS GovCloud Creds 01`.
Datacenter URL	URL of the datacenter. For example, https://ec2.us-gov-west-1.amazonaws.com
Datacenter type	Select the CMDB table that represents the type of datacenter. For example, the [`cmdb_ci_aws_datacenter`] table.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.

3. Click **Update** or **Submit**.
   The system creates the service account and displays the list of all discovered datacenters.

What to do next

Repeat the process to create additional service accounts as needed. Run Discovery for each datacenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Define the schedule for downloading AWS billing data

Define the scheduled job that regularly uses a MID Server to download billing data from the provider. Cloud Provisioning and Governance saves the data in a cost table and uses the information to generate reports.

Before you begin

Important:

Starting with the Vancouver release, the Billing dashboard is no longer available if you have downloaded and activated the ServiceNow Store Cloud Cost Management app. The following changes occur:

- You are redirected to the Cloud Cost Management home page by default.
- The View Dashboard widget in the Cloud User portal is replaced by the View Resources widget.
- The Current Month Spend widget and the Budget widget on the Cloud User portal do not show any data if Cloud Cost Management is activated on the instance.

If you have activated the Cloud Cost Management app, you can only navigate to the Billing Dashboard, if you are using Cloud Provisioning and Governance on a domain separated instance.

- Set up a payer account that gets billing information for all linked accounts.
- Create the Amazon S3 bucket that is the data storage area for the daily detailed billing records on the AWS account. The credential that is used to access the report must have permissions to the S3 bucket. Set the S3 bucket based on the Order location. The `AssumeRole` policies MUST include an S3 access policy. You might assume that the AWS account that you always use can create an S3 bucket in the region where the CFT is being provisioned. If you use STS-based authentication, however, the AWS account might not be able to create the bucket because authentication is tied to the `AssumeRole` policies. See the "Setting Up an Amazon S3 Bucket for AWS Cost and Usage Reports" page in the AWS documentation for details.
- Enable detailed billing reports on AWS. See the AWS documentation for details.

Note: The Detailed Billing Report feature is withdrawn by Amazon AWS Cloud for new customers. For more information, and for migration details see the AWS documentation on Detailed Billing Reports.

Role required: sn_cmp.cloud_governor

About this task

This procedure involves two tasks:

1. Fill in the AWS Billing Schedule form with information that you copy from the AWS portal. This action enables the instance to download AWS billing data.
2. Specify the schedule for downloading AWS billing information.

Procedure

1. In the Cloud Admin Portal, navigate to **Analyze** > **Billing**.
2. On the **Billing Schedules** tab, click **New**, enter a unique and meaningful **Name**, and then fill in the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Provider	Select **AWS**.
Service Account	The AWS service account that you created during initial setup. (The AWS account to run the billing report on.)
Bucket	Data storage area for the daily billing records on the AWS account. Follow this procedure to obtain the value: **a.** Log in to the AWS Management Console and select **My Billing Dashboard** in the menu. **b.** From the dashboard, select **Billing Preferences**. **c.** In the **Cost Management Preferences**section, expand **Detailed Billing Reports [Legacy]** and ensure that the **Turn on the legacy Detailed Billing Reports feature to receive ongoing reports of your AWS charges** check box is selected. **d.** Copy the value from the **Save the S3 bucket:** field into the **Bucket** field on the form. **Note:** By default, AWS selects the S3 bucket associated with the AWS account that you are currently signed into. A green tick indicates that the bucket is configured correctly.
Active	Select the check box to activate the download job.
Run	Frequency that the download job should run.
Time	Time of day that the download job should run.

3. Click **Save**.
   To run the schedule on demand, click **Execute Now**. On the **Retrieve Billing Data** window that appears, select a date range and then click **Download**.

What to do next
View cloud costs on the Billing dashboard.

Related topics
Cloud credentials

Configure the download size of AWS billing data

A single set of cloud billing data can be large. The MID Server, therefore, sends the data to the ECC queue in manageable chunks. You can optionally configure a system property to limit the size of each chunk to avoid performance issues caused by large data transfers.

Before you begin
Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open the System properties table by entering `sys_properties.list` in the application filter.
2. Filter the list to open the `sn_cmp.billing.page_size` property.
   The default value is 3 MB.
3. Modify the value as needed.
   The maximum is 5 MB.
4. Save the property.

Add an AWS service account to the cloud account

During Cloud Provisioning and Governance Day 1 setup, you added one service account to the cloud account. To compartmentalize your infrastructure or to include different datacenters, you can add another service account. A particular datacenter, however, cannot be selected in more than one service account in a cloud account.

Before you begin

Role required: sn_cmp.cloud_admin

About this task

A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter.

In this example, you added the service account named **ProviderB-ServiceAccount-1** and selected three datacenters to include in the cloud account:

A service account with three selected datacenters

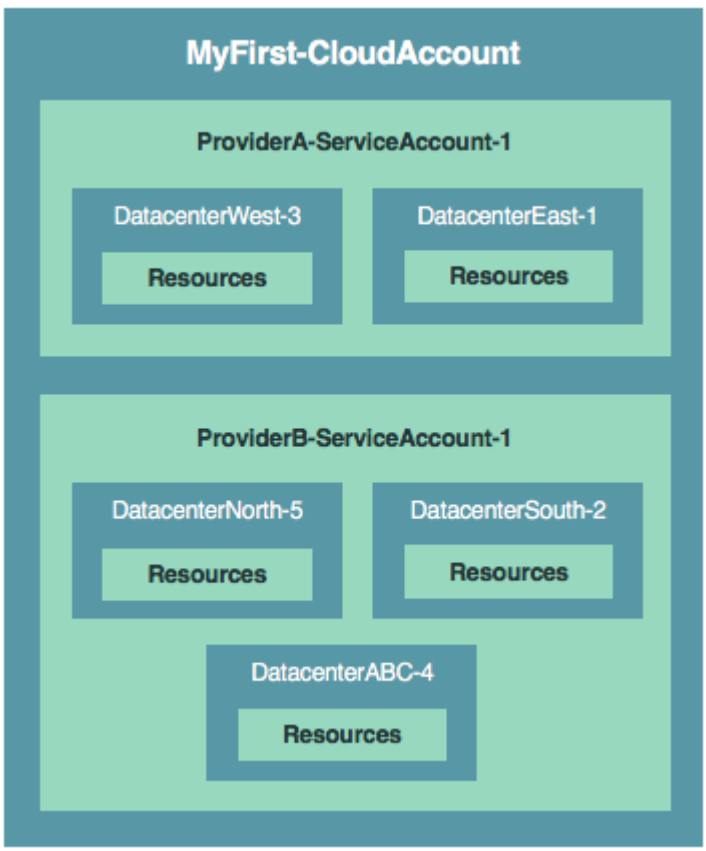

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: In a cloud account, you cannot select a particular datacenter in two different service accounts.

Procedure

1. Navigate to **Cloud Admin Portal** > **Service Accounts**.
2. Click **New**, enter a unique and meaningful **Name**, and then fill in the form.

Cloud Service Account form

Field	Description
Name	The unique and meaningful name for this service account.
Account ID	12-digit user account number. Expand the list under the AWS account name on the AWS Management Console to view the number. **Note:** In the **Account ID** field, remove the hyphen characters (-) from the number. **IAM user account number** 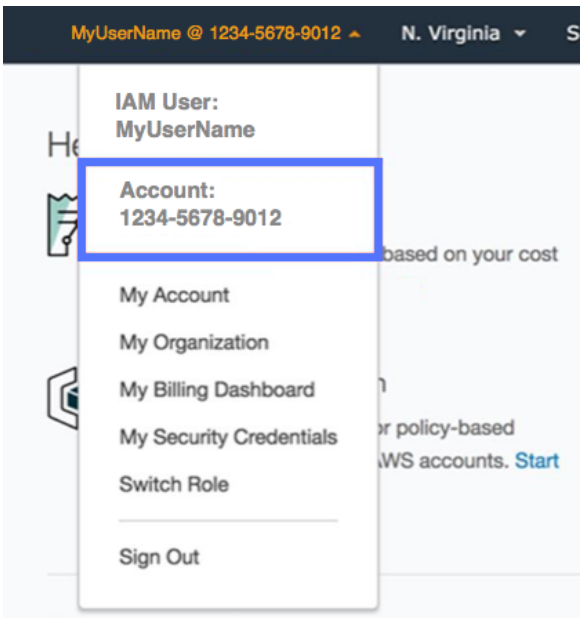
Discovery credentials	The credentials needed for ServiceNow applications to access this account. You may configure this field at a later stage, while configuring access to AWS accounts. ○ If you configured AWS credentials at ServiceNow AI Platform , select the magnifying glass icon, and then select the name of the relevant AWS credential. ○ To use other AWS accounts to access this account, leave the field blank. For example, you don't need to specify the AWS credentials for accounts assuming IAM roles or member accounts using their management account for access.
Datacenter URL	URL of the datacenter. This field is required for AWS China region and AWS GovCloud (US) accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	For example: ○ AWS China region: `https://organizations.cn-northwest-1.amazonaws.com.cn` ○ AWS GovCloud (US): `https://ec2.us-gov-west-1.amazonaws.com`
Datacenter type	Type of the datacenter where the account is hosted. Select **AWS datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.
Parent account	Name of the management account (also referred to as the master account in some ServiceNow UI screens) that represents the organization in AWS that this member account belongs to. This field appears when you select **AWS datacenter**. Leave the field empty if this account is not part of an AWS organization.
Is master account	Management account flag. This field is also referred to as the master account flag in some ServiceNow UI screens. This field appears when you select **AWS datacenter** from the Datacenter Type list. Select the check box to associate the AWS service account with the management account. Select this check box only for accounts that you previously configured in the AWS Management Console as management accounts with some member accounts belonging to them. See the AWS documentation for information on AWS Organizations. **Note:** You must set up the correct permission in AWS or the Organization role for a standard credential. For more information, see Control AWS access and permissions using policies.
Accessor account	Name of the trusted account. Configure this field only for accounts that don't use permanent AWS credentials and rely on IAM roles for access.

3. Click **Update** or **Submit**.
   The system creates the service account and displays the list of all discovered datacenters.
4. Repeat the process to add as many service accounts as needed.

Day 1 setup guide for Microsoft Azure Cloud on Cloud Provisioning and Governance

To set up Cloud Provisioning and Governance for the very first time, you perform the procedures in this "Day 1" setup guide. Be sure to perform the procedures in order. After you have performed Day 1 setup, you can perform optional Day 2 setup and configuration procedures as needed and in any order. Detailed instructions for each procedure follow this overview.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request the Cloud Provisioning and Governance application

The Cloud Provisioning and Governance application is available as a separate subscription and requires the Cloud Provisioning and Governance plugin (com.snc.cloud.mgmt). See Request the Cloud Provisioning and Governance application.

Roles required to set up Cloud Provisioning and Governance

- Operations on the Azure portal require one of the following roles:
  - Azure or Azure AD (Active Directory) Administrator
  - Application Administrator
  - Application Developer
  - Cloud Application Administrator

  and the Resource Policy Contributor role to create or modify resource policies.
- Operations in Cloud Provisioning and Governance require the sn_cmp.cloud_admin role.

About terms that Cloud Provisioning and Governance uses

Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings. In Azure, the region-specific containers for virtual resources are called regions. In Cloud Provisioning and Governance, regions are called datacenters or logical datacenters (LDCs). The term logical is used to reinforce the idea that Cloud Provisioning is provider-agnostic. All infrastructure or applications that are deployed using Cloud Provisioning are associated with a datacenter.

Quick overview of the setup process

Your setup process includes these tasks:

- If needed: Request the Cloud Provisioning and Governance application.
- Assign appropriate roles to cloud users.
- Set up the MID Servers that will handle secure communications with the provider API endpoints.
- On the provider portal, collect your account settings and the credentials that the Discovery process will use (through a MID Server) to programatically access your provider accounts. Securely associate the account settings and credentials with a service account in Cloud Provisioning and Governance.
- Set up a cloud account to represent your entire managed cloud infrastructure and set up a service account that works with one of your provider accounts. You specify which datacenters in the service account should be included in the cloud account. (Later, on "Day 2", you can set up additional cloud accounts and service accounts from the same or other providers.)
- Cloud Provisioning uses CMDB data to help users request and manage cloud resources and to help you manage your cloud infrastructure. To populate the CMDB with resource data for all datacenters, you manually run the Discovery process on each datacenter in the service account. Then, to ensure that the data continues to be updated, you configure Discovery to run on a regular schedule. Your cloud account might look like this:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Structure of a cloud account on Day 1

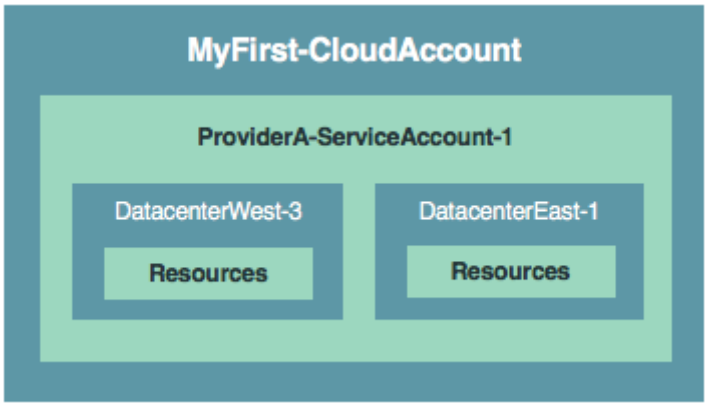

- Providers offer services that can auto-update the CMDB whenever a create/modify/terminate life-cycle change or configuration update occurs to a resource. You can configure the service to integrate directly with Cloud Provisioning and Governance.

What you will do to integrate Cloud Provisioning and Governance with your Microsoft Azure accounts

Detailed instructions for each procedure follow this overview.

1. Assign roles to Azure users of Cloud Provisioning and Governance

You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

2. Install and configure MID Servers to access cloud environments

To ensure secure and reliable communications, the Discovery process communicates with your cloud provider accounts and cloud resources through one or more MID Servers. You can set up the MID Servers on your network or in one of your cloud networks.

Note: Data is encrypted to the MID Server and between the MID Server and the API endpoint. To ensure high performance and security, you should configure one or more MID Server for each datacenter under management. Configure the MID Server even if you have already configured other MID Servers while setting up Cloud Provisioning and Governance for another cloud provider.

3. Create a Microsoft Azure service principal

To securely access resource and billing data on your Microsoft Azure account, the Discovery process must present appropriate Microsoft Azure account credentials. You create a special programmatic account — a Microsoft Azure service principal — to generate the required credentials.

4. Store the Azure service principal credentials in the instance

To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery, you first create Azure service principal credentials in the Azure Portal. You then securely store the credentials in a service account in your instance.

5. Add an Azure service account

A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

to access your provider account to get data on each resource in each specified datacenter.

6. Discover all datacenters in a service account on-demand

Discovery uses the information in the service account to identify all logical datacenters associated with the provider account.

7. Set up a cloud account for Azure

A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

8. Set capacity limits on user requests for resources

Capacity limits place restrictions on the attributes of cloud resources such as the number of virtual machines, virtual CPUs, or aggregate storage. You can set limits on resources separately for each logical datacenter in a cloud account.

10. Create Microsoft Azure credentials for billing download

Define the scheduled job that regularly uses a MID Server to download billing data from the provider. Cloud Provisioning and Governance saves the data in a cost table and uses the information to generate reports.

Next steps

When you have finished all Day-1 and Day-2 procedures in this setup guide, see the Cloud Provisioning and Governance administration guide for information on using the Cloud Provisioning and Governance application in your organization.

Set up Microsoft Azure Government Cloud on Cloud Provisioning and Governance

Set up Microsoft Azure Government Cloud on Cloud Provisioning and Governance for the first time. You can discover, provision, and manage Microsoft Azure Government Cloud resources using Cloud Provisioning and Governance, thereafter.

Before you begin

To set up Microsoft Azure Government cloud on Cloud Provisioning and Governance, you must:

- Set up a Microsoft Azure Service Principal on the Azure Government Portal and register the service account on the Cloud Provisioning and Governance application.
- Configure a MID Server with the required properties to enable Microsoft Azure Government and the ServiceNow instance to communicate correctly.
- Use the Microsoft Azure Service Principal credentials to discover associated service accounts and cloud resources.

- In the base system Terraform template, set the *environment* variable to `usgovernment`. For information on where to find the template to modify, see Discover the Infrastructure as Code (IaC) templates.

```
provider "azurerm" {
     subscription_id = "<REPLACE_WITH_ACCOUNT_ID»"
     client_id = "<REPLACE_WITH_APP_ID»"
     client_secret = "<REPLACE_WITH_PASSWORD»"
     tenant_id = "<REPLACE_WITH_TENANT_ID»"
     environment = "usgovernment"
  }
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin

Note: ServiceNow® ITOM products are not officially certified or supported for Germany and China regions.

About this task

Be sure to perform the procedures in order. After you have performed the Day 1 setup, you can perform the optional Day 2 setup and configuration procedures as needed and in any order. Detailed instructions for each procedure follow this overview.

Procedure

1. Create a Microsoft Azure service principal on the Microsoft Azure Government portal.

   Note: Navigate to **App registrations** > **Register an application** > **Supported account types** section. Then, select the **Accounts in this organizational directory only (Single tenant)** option to ensure that the Active Directory that is configured for the Microsoft Azure Government cloud is used for authentication.

   For more information, see Create an Azure AD application and service principal that can access resources.

2. Configure MID Servers to connect Cloud Provisioning and Governance to a Microsoft Azure Government cloud account.

   a. Navigate to **MID server** > **Capabilities**.

   b. Click **Edit** and add the following options:

   - **Azure**
   - **Cloud Management**

   c. Click **Save**.

   d. **Optional:** Navigate to **MID server** > **Supported Applications**.
   Skip the next step if **ALL** applications are selected.

   e. **Optional:** Click **Edit** and add the following applications to the **Supported Applications List**:

   - **Discovery**
   - **Cloud Management**
   - **Orchestration**

   f. Click **Save**.

   Note: For more information, see:

   - MID server selection
   - MID server capabilities

3. Navigate to **MID server** > **Properties**.

   a. Click **New** and add new records for each of the following name-value pairs:
   MID Server properties

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Value
mid.cmp.azure_api.base_endpoint	https://management.usgovcloudapi.net/${apiPath}
mid.cmp.azure_api.storage.container.list	https://${accountName}.blob.core.usgovcloudapi.net/?comp=list
mid.property.azure_token_management_endpoint	https://login.microsoftonline.us/${tenantId}/oauth2/token
mid.property.azure_token_management_payload	grant_type=client_credentials&client_id=${clientId}&resource=https%3A%2F%2Fmanagement.usgovcloudapi.net%2F&client_secret=${clientSecret}
mid.cmp.azure_api.storage.blob.snapshot	https://${accountName}.blob.core.usgovcloudapi.net/${containerName}?restype=container&include=snapshots&comp=list

b. Click **Submit**.

4. Store the Azure service principal credentials in the instance.
5. In the **Cloud Admin Portal**, navigate to **Manage** > **Credentials**, and select the Azure Service principal record.
6. To discover service accounts associated to the Azure subscriptions that are available to this Service principal, click the **Discover Subscriptions** related link.

Result
The **Azure Subscriptions** related list displays all subscriptions that are associated with the Azure service principal.

What to do next
Define the schedule for downloading Microsoft Azure billing data.

Assign roles to Azure users of Cloud Provisioning and Governance

You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

Before you begin
Role required: user_admin or admin

Procedure

Assign the following roles to groups and users as appropriate:

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
Root administrator `[sn_cmp.cmp_root_admin]`	Highest level of application access for Cloud Provisioning and Governance.	All
Cloud administrator	Configures the Cloud Provisioning and Governance	Cloud infrastructure:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
[sn_cmp.cloud_admin]	application and sets up the cloud infrastructure.	Service accounts and cloud account Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • IBM Cloud Update • VMware Events
Governor [sn_cmp.cloud_governor]	Monitors overall cloud usage and enforces compliance of the organization's rules, quotas, and policies. Also manages tags and permissions to various objects.	Governance: • Policies • Pools • Quotas • Permissions
Service Designer [sn.cmp.cloud_service_designer]	Creates blueprints, ARM and CloudFormation templates, and catalog items.	Cloud Service Design: • Cloud templates • Blueprints • Blueprint catalog items • Resource blocks
Cloud user [sn_cmp.cloud_service_user]	Requests and manages stacks and resources.	Cloud User Portal. When you assign the role to a group, all members of the group share quota limitations and ownership of certain resources. You can access all the task records, assigned to you or otherwise, if you have the cloud service user role.
Cloud operator [sn_cmp.cloud_operator]	Monitors and troubleshoots the Cloud Provisioning and Governance application.	Dashboards and reports: • Cloud Operations Dashboard • Cloud Root Cause Analysis

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
		• Cloud Orchestration Trail • Cloud API Trail
Cloud Event Integration `[sn_cmp.cloud_event_integration]` **Note:** Not supported by IBM Cloud Connector.	Authorizes access to the instance for external services that auto-update the CMDB when cloud events occur. This role gives the access to the cloud event REST endpoint.	Cloud infrastructure: • Service accounts and cloud account • Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • VMware Events
Cloud Group administrator `[sn_cmp.cloud_group_admin]`	Grants admin access to any group that you belong to.	

Related topics

Assign a role to a group

Assign a role to a user

Create a Microsoft Azure service principal

To securely access resource and billing data on your Microsoft Azure account, the Discovery process must present appropriate Microsoft Azure account credentials. You create a special programmatic account — a Microsoft Azure service principal — to generate the required credentials.

Before you begin

Roles required:

- discovery_admin, service_mapping_admin, sn_cmp.cloud_admin roles in Cloud Provisioning and Governance or sn_cloud_ops_ws.cloud_ops_admin role in Cloud Discovery Workspace.
- Operations on the Microsoft Azure portal require one of the following roles:
  - Azure or Azure AD (Active Directory) Administrator
  - Application Administrator
  - Application Developer
  - Cloud Application Administrator

  and the Resource Policy Contributor role to create or modify resource policies.
- Enable internal network connection between the MID Servers and the Azure Cloud API endpoints:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The US GovCloud URL is `https://management.usgovcloudapi.net/`.
- The commercial Azure Cloud URL is *`management.azure.com`*.

Note: It isn't necessary when adding a credential if the account being added is already a GovCloud account.

Procedure

1. Log in to the Azure portal and navigate to **Azure Active Directory**.
2. Navigate to the **App registrations** section and click **New application registration**. Enter the following information for your application:

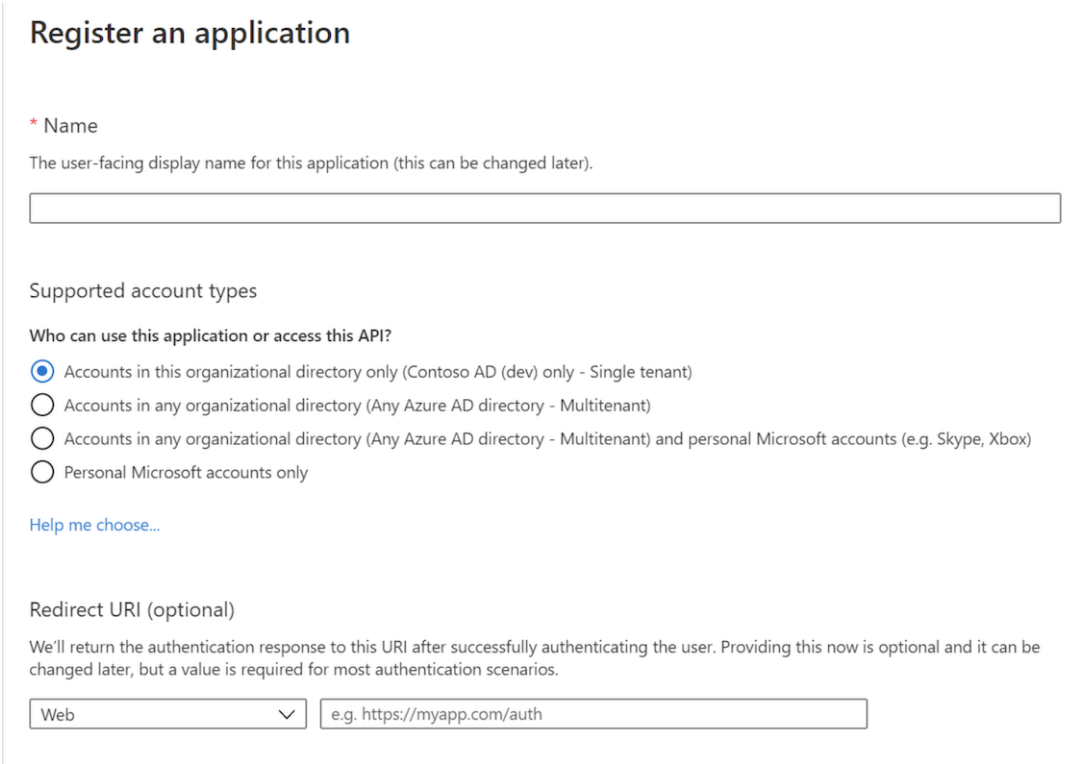

Field	Description
Name	Unique name for the application and its integration credentials. For example, `ServiceNow Integration`.
Supported account types	Specify who can use the application.
Redirect URI (Optional)	URL that will access Azure. Typically the URL of the ServiceNow instance.

3. Select **Register** to complete the app registration.
4. When registration completes, copy the Application (client) ID and Directory (tenant) ID values, and paste them in the text editor.
5. Label the values **Application ID** and **Directory ID** respectively.
6. In the Azure portal, navigate to the **Certificates & secrets** section and **New client secret** then specify the following values:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Key description	Description for the key.
Duration	Expiration for the key. **Note:** Your organization may apply policies to restrict key durability. Select the appropriate duration.

7. Click **Add**.
8. Copy and paste the key value into the text editor and label the value **Application key**.
9. To enable the service principal to work with various Azure subscriptions, navigate to **Subscriptions**.
   To manage multiple subscriptions, you must perform the following procedure for each subscription:

   a. Paste the subscription ID into the text editor and label it **Subscription ID**.
   The text file that you generate during this procedure might look something like this:

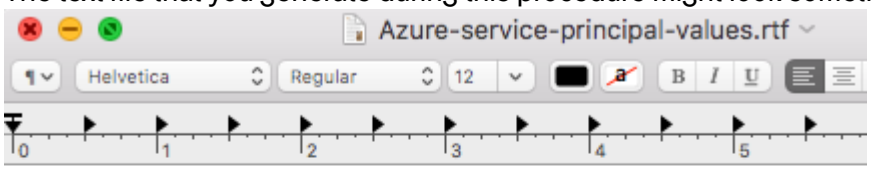

Directory ID: 8bcdd120-9979-491e-8683-d9a4550bdd

Application ID: e6821ea9-e5c5-428a-b50c-7ea62993b2d

Application key: d348050-f893ccb7-hg84yg8-rg200hd

Subscription ID: 4803be48c3ee89-96d2-4fed-bedb-22e49013ee89

   b. Navigate to the subscription and select **Access Control (IAM)** from the menu.

   c. Click **+ Add** at the top of the screen then **Add role assignment**.

   d. Select the value **reader** from the **Role** field.
   Let the default value **User, group, or service principal** remain as is in the **Assign access to** field.

   Note: The **Resource policy contributor** role is only required for provisioning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. Select the name you created in step 2 in the **Select** field and click **Save**.

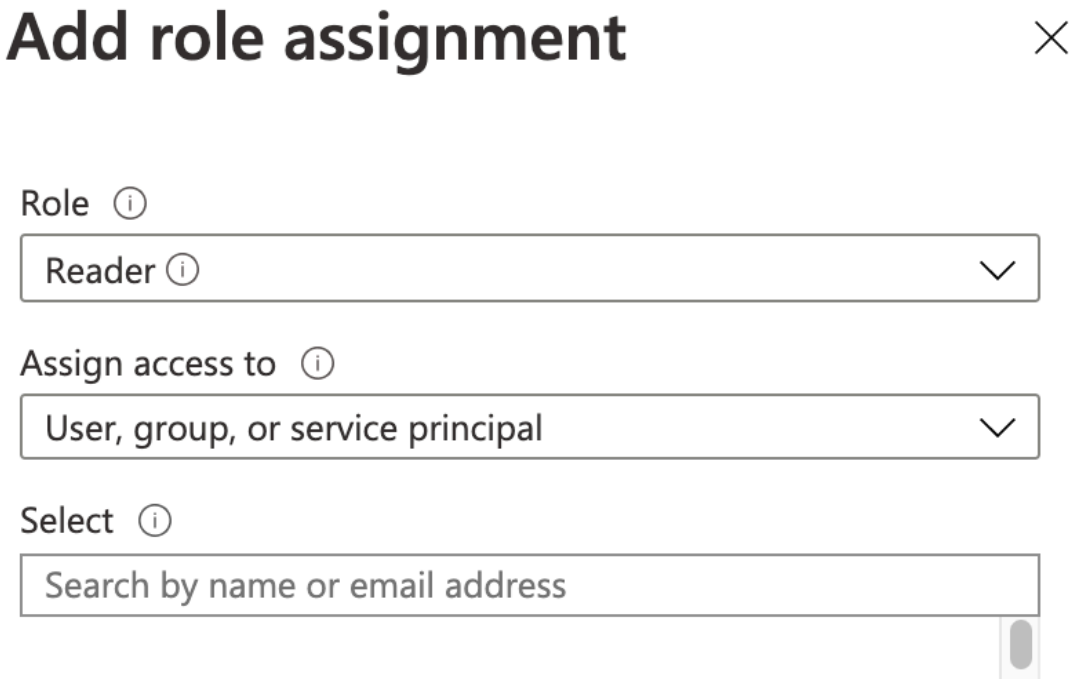

10. Perform the appropriate action.

- If you are not using Cloud Discovery through Cloud Discovery Workspace, do the following:

  In the Discovery Manager, click the plus icon (**+**) and then select **Azure Service Principal** from the list.

- If you are using Cloud Discovery through Cloud Discovery Workspace, do the following:

  a. Navigate to **All** > **Connections & Credentials** > **Credentials**.

  b. Select **New**.

  c. Select **Azure Service Principal**.

11. Specify the following values on the Azure Service Principal form:

Field	Value
Name	Name of the service principal to register with the instance. For example, `Azure service principal credentials`.
Authentication Method	Select **Client secret**. The **Secret key** field appears when you select **Client secret**. **Note:** **Client assertion** is not supported.

12. Copy and paste values from the temporary text file into the remaining fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

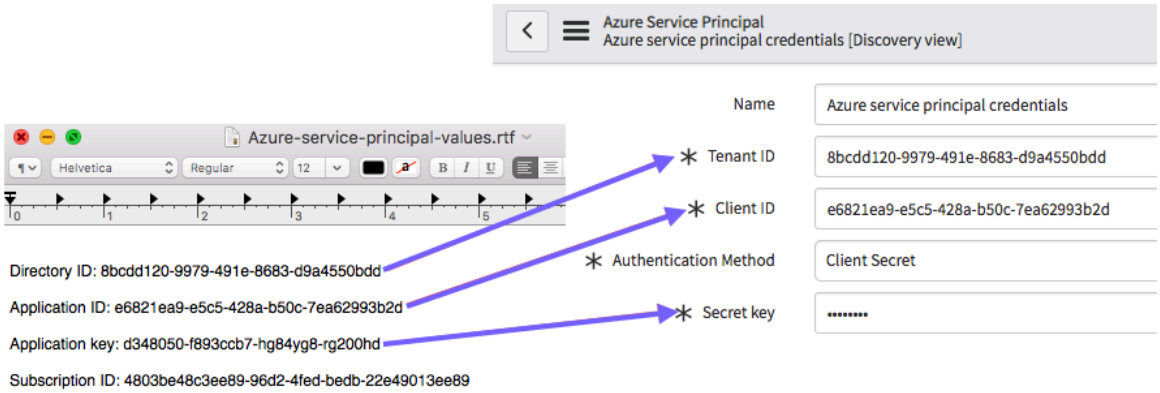

Credentials form field	Azure Service Principal value
Tenant ID	Azure **Directory ID** value from the text file.
Client ID	Azure **Application ID** value from the text file.
Secret key	Azure **Application key** value from the text file.

13. Click **Save** to create the Azure service principal.
14. Click the **Discover Subscriptions** related link to find all subscriptions for the Azure service principal.
    The instance creates a service account for each discovered subscription. The **Azure Subscriptions** related list displays all subscriptions for the Azure service principal.
15. Click a subscription to view the service account created for the subscription.
16. Click a Discovery status entry in the **Credential Discovery Status** list to view the Discovery log. Each time you click **Discover Subscription**, the instance generates a new Discovery status and displays it in the **Credential Discovery Status** list.

What to do next

Cloud Provisioning and Governance only: Create a record of the service principal credentials on the ServiceNow instance so that Cloud Provisioning and Governance processes can access Microsoft Azure data. See Store the Azure service principal credentials in the instance.

Store the Azure service principal credentials in the instance

To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery, you first create Azure service principal credentials in the Azure Portal. You then securely store the credentials in a service account in your instance.

Before you begin

Role required:

- Azure Portal Active Directory (AD) administrator
- Cloud Provisioning and Governance: admin or sn_cmp.cloud_admin, discovery_admin

Procedure

1. Create a Microsoft Azure service principal and open the text file that you created during the procedure.
2. In the Cloud Admin Portal, navigate to **Manage** > **Credentials**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click **New** and then select **Azure Service Principal**.
4. Specify the following values on the Azure Service Principal form:

Field	Value
Name	Name of the service principal to register with the instance. For example, `Azure service principal credentials`.
Authentication Method	Select **Client Secret**. The **Secret key** field appears when you select **Client Secret**. **Note:** **Client Assertion** is not supported.

5. Copy/paste values from the `Azure-Credentials.txt` text file into the remaining fields.

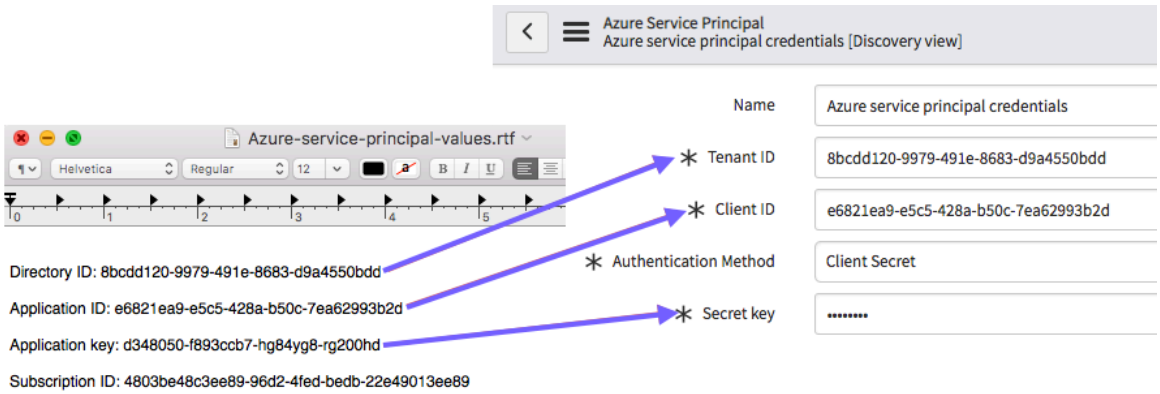

6. Select the appropriate **EA credential** from the list, select the **Active** check box, and then click **Save** to create the record.
7. Click the **Discover Subscriptions** related link to find all subscriptions that are associated with the Azure service principal.
   The instance creates a service account for each discovered subscription.The **Azure Subscriptions** related list displays all subscriptions that are associated with the Azure service principal.
8. Click a subscription to view the service account that was created for the subscription.
9. Click a discovery status entry in the **Credential Discovery Status** list to view the associated discovery log.
   Each time that you click **Discover Subscriptions**, the instance generates a new discovery status and lists the status in the **Credential Discovery Status** list.

Add an Azure service account

During Cloud Provisioning Day 1 setup, you added one service account to the cloud account. To compartmentalize your infrastructure or to include different datacenters, you can add another service account. A particular datacenter, however, cannot be selected in more than one service account in a cloud account.

Before you begin
Role required: sn_cmp.cloud_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open the `Azure-Credentials.txt` text file that you created during the Create a Microsoft Azure service principal procedure.
2. Navigate to **Cloud Admin Portal** > **Manage** > **Service Accounts**.
3. Click **New**, enter a unique and meaningful **Name**, and then fill in the form.

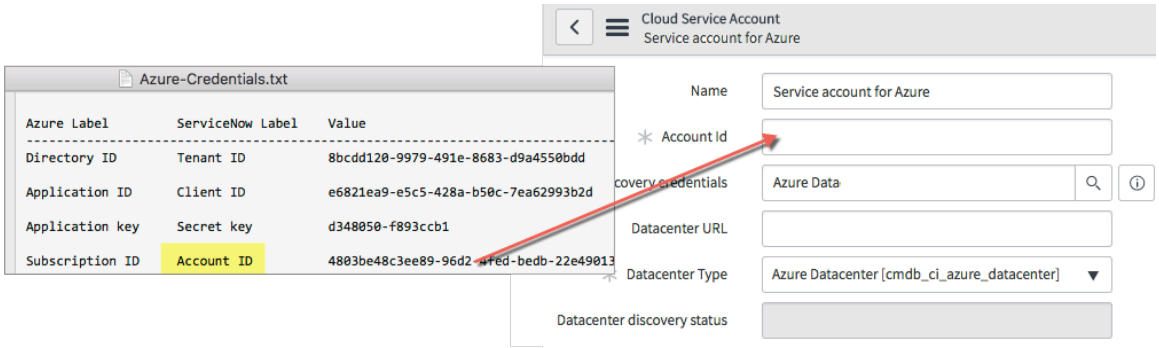

Field	Description
Account ID	The Azure **Subscription ID** value that you copied from the Azure Portal into the text file. See Discovery for Microsoft Azure to learn Management Groups subscription IDs.
Discovery credentials	Name of the credentials that you created in the Store the Azure service principal credentials in the instance procedure. In the example, you used the name `Azure service principal credentials`.
Datacenter URL	URL of the datacenter. This field is required only for the government and regional accounts. Discovery and Service Mapping Patterns supports the US, Germany, and China Microsoft Azure government clouds.
Datacenter type	Select **Azure Datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.
Should pull events	Option to enable the ServiceNow AI Platform® to collect resource change information from the Microsoft Azure cloud. The ServiceNow AI Platform® uses the resource change information to update the Configuration Management Database (CMDB). For more information on the supported Azure resource types, see Azure change processing. If you select the Should pull events check box for a management group, the ServiceNow AI Platform enables Azure change processing for all the subscriptions in the management group.

4. Click **Update** or **Submit**.
   The system creates the service account and displays the list of all discovered datacenters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

After you complete all procedures in this initial "Day 1" setup, you can create additional cloud accounts and service accounts to organize and compartmentalize your cloud infrastructure.

Note:

If you are on a domain separated instance, only those events that are updated to the CMDB and belong to your domain are visible. Events create configuration items (CI) in the same domain as the cloud service account they are mapped to.

Related topics

Discovery for Microsoft Azure

Set up a cloud account for Azure

A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

Before you begin

Role required: sn_cmp.cloud_admin

About this task

See Cloud accounts for more information about cloud accounts.

Note: Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings. In Azure, the region-specific containers for virtual resources are called regions. In Cloud Provisioning and Governance, regions are called datacenters or logical datacenters (LDCs). The term logical is used to reinforce the idea that Cloud Provisioning is provider-agnostic. All infrastructure or applications that are deployed using Cloud Provisioning are associated with a datacenter. .

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Cloud Accounts**.

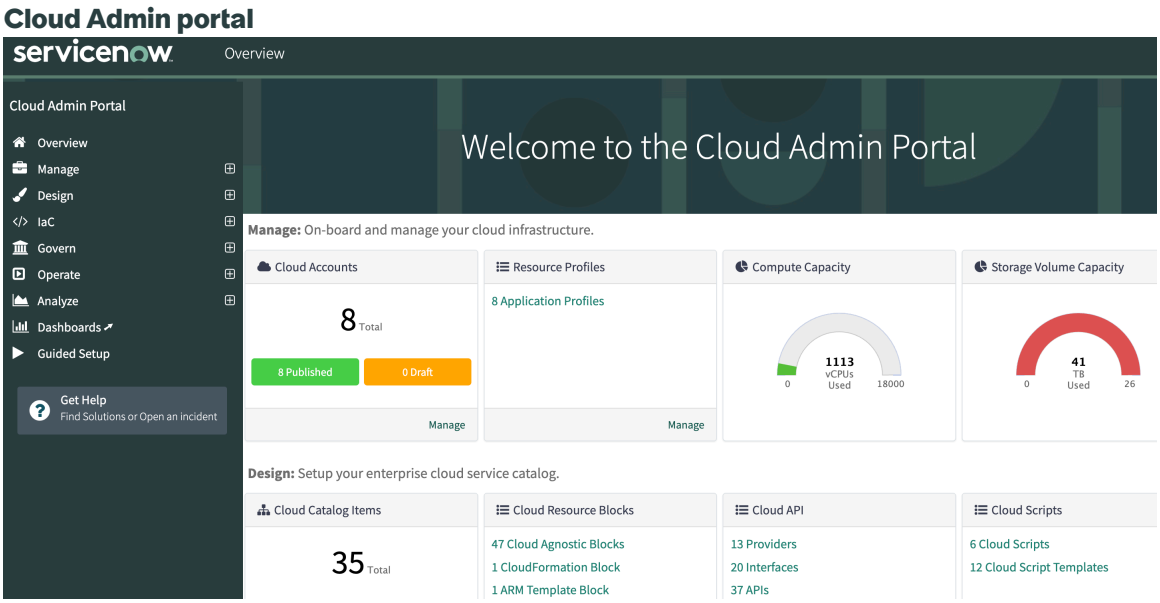

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Click **New**, enter a unique and meaningful **Name** and **Description**, and then select the cloud **Provider** for the account (Azure).

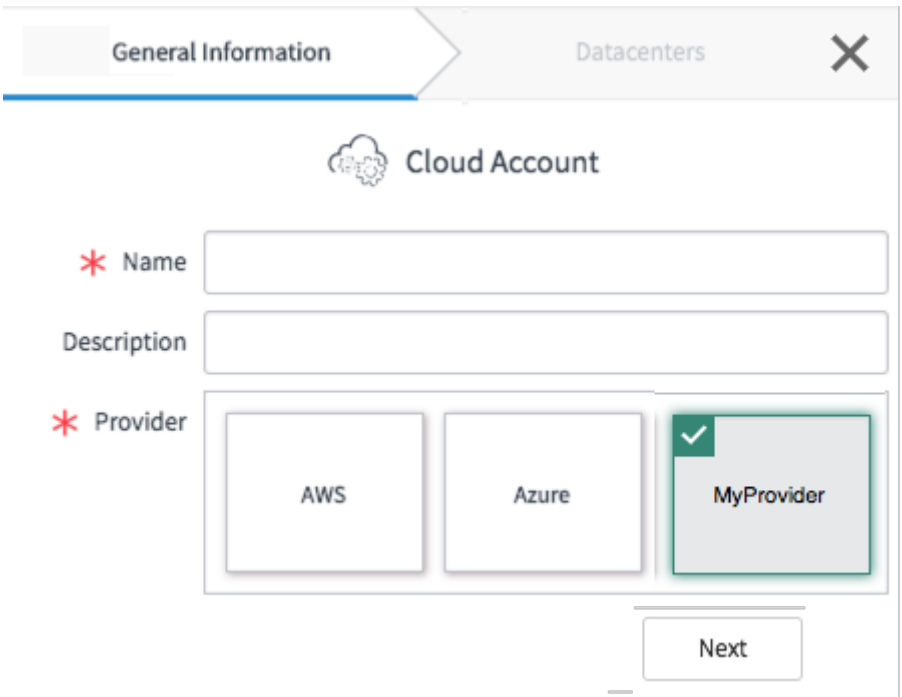

3. Click **Next**.
   Based on the provider you select, the **Service Account** field is populated and a list of discovered datacenters appears. If no datacenters appear, click **Discover Datacenters** to find them. See Discover all datacenters in a service account on-demand for details.
4. When the datacenters appear, select one or more datacenters to include in the cloud account and then click **Save**.

   Note:

   - Select only those LDCs/Regions where your infrastructure resources are present. If you don't have any resources under these new data centers, exclude them while you run the full discovery.
   - To discover resources in certain regions or LDCs, raise an access request to the required endpoints with GCP Support. These regions include North America (US East, US West, and Canada Central), South America, Europe, Africa, Middle East, and Asia Pacific.

   The Cloud Account dashboard appears. The datacenters that you selected appear on the **Datacenters** tab.
5. Click a datacenter.
   The following lists appear:
   - **Resources**: Cloud resources for the current datacenter. Run Discovery to populate the CMDB for the datacenter and populate the tab. See Discover all datacenters in a service account on-demand.
   - **Capacity Limits**: Limits on virtual CPUs, virtual networks, storage volume size, and other services. See Set capacity limits on user requests for resources.
   - **Discovery Log**: The process creates a log record for each action associated with a discovery status.
6. Repeat the process to add as many service accounts as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create Microsoft Azure credentials for billing download

Define the scheduled job that regularly uses a MID Server to download billing data from the provider. Cloud Provisioning and Governance saves the data in a cost table and uses the information to generate reports.

Before you begin

Important:

Starting with the Vancouver release, the Billing dashboard is no longer available if you have downloaded and activated the ServiceNow Store Cloud Cost Management application. The following changes occur:

- You're redirected to the Cloud Cost Management home page by default.
- The View Dashboard widget in the Cloud User portal is replaced by the View Resources widget.
- The Current Month Spend widget and the Budget widget on the Cloud User portal don't show any data if Cloud Cost Management is activated on the instance.

If you have activated the Cloud Cost Management application, you can only navigate to the Billing Dashboard when you're using Cloud Provisioning and Governance on a domain-separated instance.

You must have an API Access Key credential for an Enterprise Agreement (EA) for all Azure accounts for which you want billing information.

Role required: sn_cmp.cloud_governor

About this task

During the billing download process, all the resources are pulled into the system. Azure databases are placed in the Database [cmdb_ci_database] table.

Procedure

1. Create Client Secret by navigating to the Overview page of the previously deployed Service Principal application.

   a. On the left pane, select **Certificates & secrets**.

   b. In the Client secrets section, select **New client secret**.

   c. On the form, fill in the fields.

Field	Description
Description	Description of the Client secret. For example, `Client secret for MyServicePrincipal`.
Expires	Expiration period that suits your requirements. For example, `6 months`, `12 months`, `24 months`, or `Never`.

   d. Select **Add**.

   After adding the Client Secret, a value gets generated and displayed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Copy this value immediately as it wouldn't be shown again, which is your Client Secret (access key).

2. Copy the Application (client) ID and Directory (tenant) ID for later use.
3. Assign a role to the Service Principal.
   a. On the left pane, select **Azure Active Directory**.
   b. Select **Roles and administrators**.
   c. Select **All roles** and find the role that you want to assign.
   d. Select the role and then select **Add assignment**.
   e. Search for and select your Service principal by its name and then select **Add**.
4. Obtain a Subscription ID by selecting **Subscriptions**.
   a. Select the subscription that you want to use.
   b. Copy the Subscription ID from the Subscription Overview page.

Configure the download size of Microsoft Azure billing data

A single set of cloud billing data can be large. The MID Server, therefore, sends the data to the ECC queue in manageable chunks. You can optionally configure a system property to limit the size of each chunk to avoid performance issues caused by large data transfers.

Before you begin

Role required: admin

Procedure

1. Open the System properties table by entering `sys_properties.list` in the application filter.
2. Filter the list to open the `sn_cmp.billing.page_size` property.
   The default value is 3 MB.
3. Modify the value as needed.
   The maximum is 5 MB.
4. Save the property.

Day 1 setup guide for Google Cloud through Cloud Services Catalog Terraform Connector

To set up Google Cloud for the first time, follow the steps in this Day 1 setup guide in the order they are presented. After completing the Day 1 setup, you can proceed with the optional Day 2 configuration steps as needed, in any order. Detailed instructions for each step are provided in the sections below.

Note: The ServiceNow® Cloud Provisioning and Governance: Google Cloud Connector is no longer deployed, enhanced, or supported. For details, see the Deprecation Process [KB0867184] article in the Now Support knowledge base.

Roles required to set up Google Cloud Platform

- Operations in the Google Cloud Console require the Google administrator role.
- Operations in Cloud Provisioning and Governance require the sn_cmp.cloud_admin role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About terms that Cloud Provisioning and Governance uses

Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings. In the Google Cloud Platform, the region-specific containers for virtual resources are called regions. In ServiceNow cloud data model, regions are called datacenters or logical datacenters (LDCs). The term logical is used to reinforce the idea that Cloud Provisioning is provider-agnostic. All infrastructure or applications that are deployed using Cloud Provisioning are associated with a datacenter.

Quick overview of the setup process

Your Day 1 setup process includes these tasks:

- If needed: Request the Cloud Provisioning and Governance application.
- Download the Connector app and supporting store apps on the ServiceNow Store and activate the store applications.
- Assign appropriate roles to cloud admins and end users.
- On the provider portal, collect your account settings and the credentials that the Discovery process will use (through a MID Server) to programmatically access your provider accounts. Securely associate the account settings and credentials with a service account in Cloud Provisioning.
- Set up the MID Servers that will handle secure communications with the provider API endpoints.
- Set up a cloud account to represent your entire managed cloud infrastructure and set up a service account that works with one of your provider accounts. You specify which datacenters in the service account should be included in the cloud account. (Later, on "Day 2", you can set up additional cloud accounts and service accounts from the same or other providers.)
- To populate the CMDB with resource data for all datacenters, you manually run the Discovery process on each datacenter in the service account. Then, to ensure that the data continues to be updated, you configure Discovery to run on a regular schedule. Your cloud account might look like this:

Structure of a cloud account on Day 1

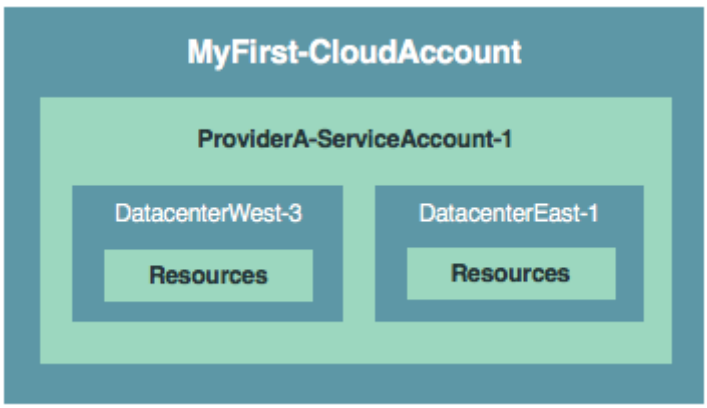

- Providers offer services that can auto-update the CMDB whenever a create/modify/terminate life-cycle change or configuration update occurs to a resource. You can configure the service to integrate directly with Cloud Provisioning and Governance.

What you will do to integrate your Google Cloud Platform cloud accounts

Detailed instructions for each procedure follow this overview.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Get the CSC Terraform Connector app on the ServiceNow Store
2. Assign roles to Google Cloud Platform users

   You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

3. Install and configure MID Servers to access cloud environments

   To ensure secure and reliable communications, the Discovery process communicates with your cloud provider accounts and cloud resources through one or more MID Servers. You can set up the MID Servers on your network or in one of your cloud networks.

   Note: Data is encrypted to the MID Server and between the MID Server and the API endpoint. To ensure high performance and security, you should configure one or more MID Server for each datacenter under management. Configure the MID Server even if you have already configured other MID Servers while setting up Cloud Provisioning and Governance for another cloud provider.

4. Specify the credentials that CSC Terraform Connector uses to access Google Cloud Platform data

   To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery, you open the Google Cloud Console to identify the Google Cloud Platform project that will have programmatic access to your Google Cloud Platform data. You then securely store the credentials in a service account in your instance.

5. Set up a cloud account and service account for Google Cloud Platform

   A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter.

Next steps

When you have finished all Day-1 and Day-2 procedures in this setup guide, see the Cloud Provisioning and Governance administration guide for information on using the Cloud Provisioning and Governance application in your organization.

Assign roles to Google Cloud Platform users

You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

Before you begin

Role required: user_admin or admin

Procedure

Assign the following roles to groups and users as appropriate:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
Root administrator `[sn_cmp.cmp_root_admin]`	Highest level of application access for Cloud Provisioning and Governance.	All
Cloud administrator `[sn_cmp.cloud_admin]`	Configures the Cloud Provisioning and Governance application and sets up the cloud infrastructure.	Cloud infrastructure: Service accounts and cloud account Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • IBM Cloud Update • VMware Events
Governor `[sn_cmp.cloud_governor]`	Monitors overall cloud usage and enforces compliance of the organization's rules, quotas, and policies. Also manages tags and permissions to various objects.	Governance: • Policies • Pools • Quotas • Permissions
Service Designer `[sn.cmp.cloud_service_designer]`	Creates blueprints, ARM and CloudFormation templates, and catalog items.	Cloud Service Design: • Cloud templates • Blueprints • Blueprint catalog items • Resource blocks
Cloud user `[sn_cmp.cloud_service_user]`	Requests and manages stacks and resources.	Cloud User Portal. When you assign the role to a group, all members of the group share quota limitations and ownership of certain resources. You can access all the task records, assigned to you or otherwise, if you have the cloud service user role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
Cloud operator `[sn_cmp.cloud_operator]`	Monitors and troubleshoots the Cloud Provisioning and Governance application.	Dashboards and reports: • Cloud Operations Dashboard • Cloud Root Cause Analysis • Cloud Orchestration Trail • Cloud API Trail
Cloud Event Integration `[sn_cmp.cloud_event_integration]` **Note:** Not supported by IBM Cloud Connector.	Authorizes access to the instance for external services that auto-update the CMDB when cloud events occur. This role gives the access to the cloud event REST endpoint.	Cloud infrastructure: • Service accounts and cloud account • Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • VMware Events
Cloud Group administrator `[sn_cmp.cloud_group_admin]`	Grants admin access to any group that you belong to.	

Related topics

Assign a role to a group

Assign a role to a user

Specify the credentials that CSC Terraform Connector uses to access Google Cloud Platform data

To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery, you open the Google Cloud Console to identify the Google Cloud Platform project that will have programmatic access to your Google Cloud Platform data. You then securely store the credentials in a service account in your instance.

Before you begin

Roles required:

- Operations in the Google Cloud Console require the Google administrator role.
- Operations in Cloud Provisioning and Governance require the sn_cmp.cloud_admin role.

About this task

Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. If you have already generated the credentials that is used by the Google Cloud, skip to the final step.
   Otherwise, continue with the next step.
2. On the Google Cloud Console dashboard, select the project.
3. Navigate to **APIs & Services** > **Credentials**.
4. On the **Credentials** tab, click **Create credentials** > **Service account key**, and then create the key.
   Specify the **JSON** key type.
5. Identify the credentials that Google Cloud, Discovery, and the Cloud Services Catalog (CSC) Terraform Connector should use on your instance to access data from the Google Cloud Platform.

   a. Open the JSON key file in a text editor.

   b. In the Cloud Admin Portal, navigate to **Manage** > **Credentials**, click **New** and then select **Google Credentials**.

   c. Specify the **Name** of the credentials to register with the instance.
   For example, enter `Google credentials-1` and then copy/paste values from the JSON key file into the form fields. Paste the `private_key` value into the **Secret key** field and the `client_email` value into the **EMail** field.

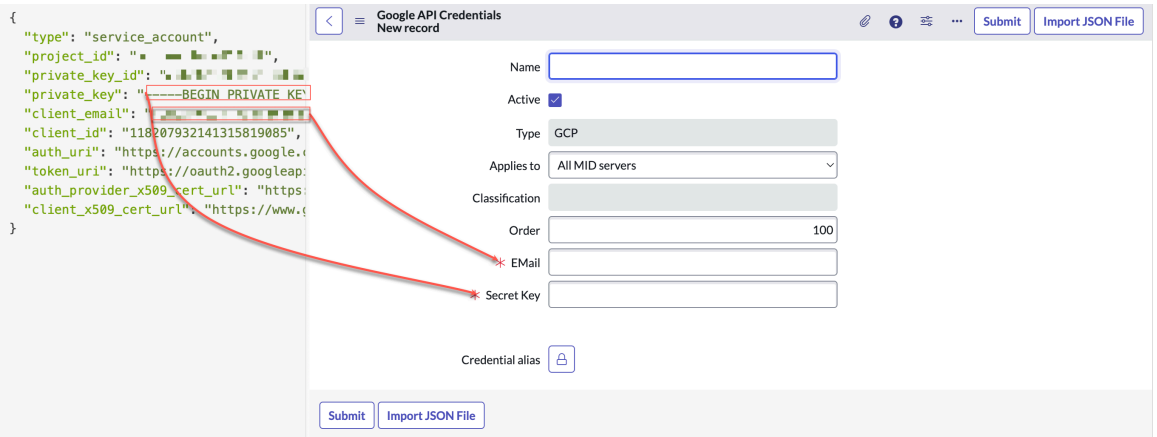

   d. Select the **Active** check box to use the credentials and then click **Submit**.

Set up a cloud account and service account for Google Cloud Platform

A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter. A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

Before you begin

Role required: none

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Download and activate the CSC Terraform Connector from the ServiceNow® Store. To try out the Out Of Box Catalogs using Cloud Services Catalog for Google, install the CSC content pack.
- Operations in the Google Cloud Console require the Google administrator role.
- Operations in Cloud Provisioning and Governance require the sn_cmp.cloud_admin role.

About this task

When you finish the Day-1 setup process, your cloud account might look like this:

Structure of a cloud account on Day 1

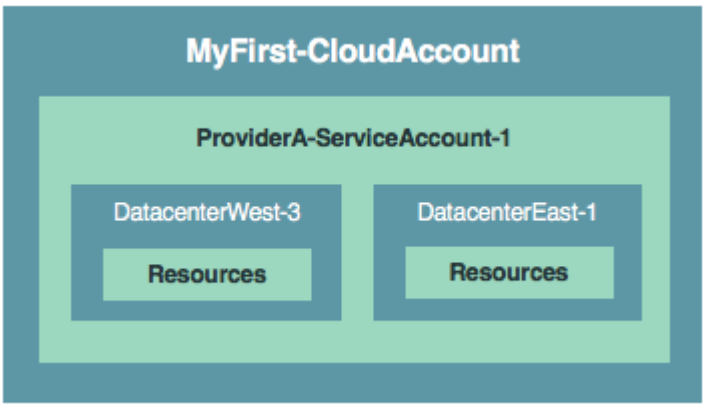

Procedure

1. On the Cloud Admin Portal, navigate to **Manage** > **Cloud Accounts**.
2. Select or create the cloud account:
   - Select an existing cloud account:
     a. On the Cloud Account page, click **Configure**.
     b. On the General Information page, click **Next** and then continue with the next step.
   - Create a cloud account:
     a. On the General Information page, click **New** and then enter a unique and meaningful **Name** and **Description** for the cloud account.
     b. Select the provider.
     c. Click **Next** and then continue with the next step.
3. On the Datacenters page, create a service account: Click the **+** next to the **Service Account** field, and then fill in the form for the service account.
4. From the JSON key file that is associated with the service account, populate the **project_id** value into the **Account ID** field.

   Note: You can use the import JSON file option to populate. For more details, see Configure the Google API credentials for Cloud Discovery.

5. In the **Discovery credentials** field, select the appropriate credentials for the service account. These are the credentials that you generated in the Specify the credentials that CSC Terraform Connector uses to access Google Cloud Platform data procedure.
6. Leave the **Datacenter URL** field blank.
7. Fill in the remaining fields:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Datacenter type	Select **Google Cloud Platform Datacenter**.
Datacenter discovery status	Select **Google Cloud Platform Datacenter**.
Should pull events	Enables Cloud Provisioning and Governance to collect events from resource types in Google Project LDCs in the `sn_cmp_cloud_event_list` .

8. Click **Submit** to create the service account.
9. Click **Discover Datacenters**.
10. When the datacenters appear, select one or more datacenters to include in the cloud account and then click **Save**.

    Note:

    - Select only those LDCs/Regions where your infrastructure resources are present. If you don't have any resources under these new data centers, exclude them while you run the full discovery.
    - To discover resources in certain regions or LDCs, raise an access request to the required endpoints with GCP Support. These regions include North America (US East, US West, and Canada Central), South America, Europe, Africa, Middle East, and Asia Pacific.

    The Cloud Account dashboard appears. The datacenters that you selected appear on the **Datacenters** tab.
11. Click a datacenter.
    The following lists appear:
    - **Resources**: Cloud resources for the current datacenter. Run Discovery to populate the CMDB for the datacenter and populate the tab. See Discover all datacenters in a service account on-demand.
    - **Capacity Limits**: Limits on virtual CPUs, virtual networks, storage volume size, and other services. See Set capacity limits on user requests for resources.
    - **Discovery Log**: The process creates a log record for each action associated with a discovery status.
12. Repeat the process to add as many service accounts as needed.

What to do next

After you complete all procedures in this initial "Day 1" setup, you can create additional cloud accounts and service accounts to organize and compartmentalize your cloud infrastructure.

Related topics

Set up an additional cloud account

(Optional) Add a Google Cloud Platform service account to the cloud account

Discovery for Google Cloud Platform

(Optional) Add a Google Cloud Platform service account to the cloud account

During Cloud Provisioning and Governance Day 1 setup, you added one service account to the cloud account. To compartmentalize your infrastructure or to include different datacenters, you

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

can add another service account. A particular datacenter, however, cannot be selected in more than one service account in a cloud account.

Before you begin
Role required: sn_cmp.cloud_admin

About this task
A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter.

In this example, you added the service account named **ProviderB-ServiceAccount-1** and selected three datacenters to include in the cloud account:

A service account with three selected datacenters

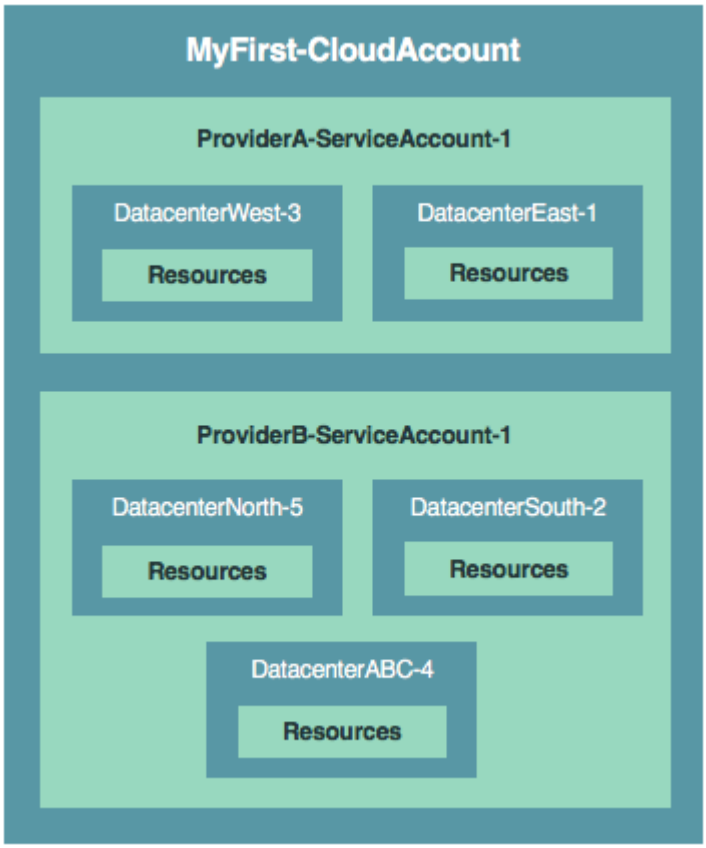

Important: In a cloud account, you cannot select a particular datacenter in two different service accounts.

Procedure
1. Navigate to **Cloud Admin Portal** > **Service Accounts**.
2. Click **New**, enter a unique and meaningful **Name**, and then fill in the form.
3. Fill in the remaining fields:

Field	Description
Account ID	From the JSON key file that is associated with the service account, copy/paste the **pr** into the **Account ID** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Google Cloud Platform (GCP) project ID 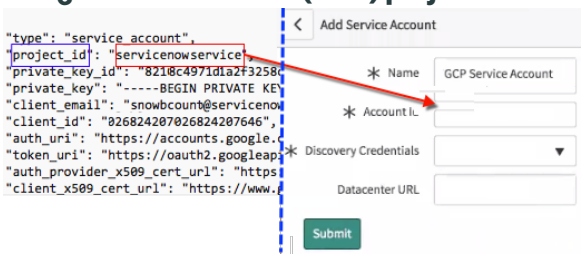
Discovery credentials	Select the appropriate credentials for the service account For information on genera the credentials, see Specify the credentials that CSC Terraform Connector uses to a Cloud Platform data.
Datacenter URL	Leave the **Datacenter URL** field blank.
Datacenter type	Select **Google Cloud Platform Datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on th

4. Click **Update** or **Submit**.
   The system creates the service account and displays the list of all discovered datacenters.
5. Repeat the process to add as many service accounts as needed.

Day 1 setup guide for VMware on Cloud Provisioning and Governance

To set up Cloud Provisioning and Governance for the very first time, you perform the procedures in this "Day 1" setup guide. Be sure to perform the procedures in order. After you have performed Day 1 setup, you can perform optional Day 2 setup and configuration procedures as needed and in any order. Detailed instructions for each procedure follow this overview.

Request the Cloud Provisioning and Governance application

The Cloud Provisioning and Governance application is available as a separate subscription and requires the Cloud Provisioning and Governance plugin (com.snc.cloud.mgmt). See Request the Cloud Provisioning and Governance application.

Roles required to set up Cloud Provisioning and Governance

- Operations in the VMware Console require the VMware administrator role.
- Operations in Cloud Provisioning and Governance require the sn_cmp.cloud_admin role.

About terms that Cloud Management uses

Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings.

Quick overview of the setup process

Your setup process includes these tasks:

- If needed: Request the Cloud Provisioning and Governance application.
- Assign appropriate roles to cloud users.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Set up the MID Servers that will handle secure communications with the provider API endpoints.
- On the provider portal, collect your account settings and the credentials that the Discovery process will use (through a MID Server) to programatically access your provider accounts. Securely associate the account settings and credentials with a service account in Cloud Provisioning and Governance.
- Set up a cloud account to represent your entire managed cloud infrastructure and set up a service account that works with one of your provider accounts. You specify which datacenters in the service account should be included in the cloud account. (Later, on "Day 2", you can set up additional cloud accounts and service accounts from the same or other providers.)
- Cloud Provisioning uses CMDB data to help users request and manage cloud resources and to help you manage your cloud infrastructure. To populate the CMDB with resource data for all datacenters, you manually run the Discovery process on each datacenter in the service account. Then, to ensure that the data continues to be updated, you configure Discovery to run on a regular schedule. Your cloud account might look like this:

Structure of a cloud account on Day 1

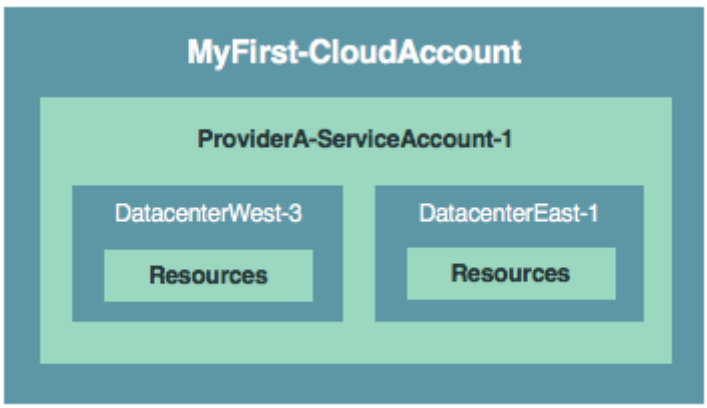

- Providers offer services that can auto-update the CMDB whenever a create/modify/terminate life-cycle change or configuration update occurs to a resource. You can configure the service to integrate directly with Cloud Provisioning and Governance.

What you will do to integrate Cloud Provisioning and Governance with your VMware cloud accounts

Detailed instructions for each procedure follow this overview.

1. Assign roles to VMware users of Cloud Provisioning and Governance

You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

2. Create the credential and service account that will access your VMware data

To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery, you first create an account in the VMware Console. You then securely store the credentials in a service account in your instance. A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter.

3. Discover all datacenters in a service account on-demand

Discovery uses the information in the service account to identify all logical datacenters associated with the provider account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Set up cloud accounts for VMware

A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

5. Set capacity limits on user requests for resources

Capacity limits place restrictions on the attributes of cloud resources such as the number of virtual machines, virtual CPUs, or aggregate storage. You can set limits on resources separately for each logical datacenter in a cloud account.

vSphere

- Enterprise or Enterprise Plus Edition 5.5 or newer.
- Enable the vSphere environment for Distributed Resource Scheduling (DRS).
- Place datastores under Storage DRS.

When provisioning Windows virtual machines, read and modify the named customization specifications.

ESXi hosts

If the IP address of a virtual machine is set manually using a script or post-init scripts run on the VM, then you must allow HTTPS traffic between the MID Server and the ESXi host for the VM.

Enable NTP time sync. Enable the NTP client for the hypervisor on the vSphere web client. See the vSphere documentation for more information.

vCenter server requirements

vCenter 5.5 or 6.0 Standard or Foundation edition.

vCenter access requirements

Minimum permissions for the user that connects to vSphere through the instance:

- Datastore: All
- Datastore Cluster: All
- Folder: Create, Delete
- Global: Cancel task, License
- Host > Local operations: Create/Delete/Reconfigure virtual machines
- Profile-driven storage: All
- Network: Assign network
- Resource: Apply recommendation; Create/Modify/Assign/Remove virtual machines.
- Virtual machines: All

vCenter credentials

- Username
- Password

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- URL or IP address of the vCenter server that the instance access through the MID Server.
- vCenter instance UUID. See Create the credential and service account that will access your VMware data for instructions on how to obtain the UUID.

Next steps

When you have finished all Day-1 and Day-2 procedures in this setup guide, see the Cloud Provisioning and Governance administration guide for information on using the Cloud Provisioning and Governance application in your organization.

Assign roles to VMware users of Cloud Provisioning and Governance

You assign Cloud Provisioning and Governance roles to user groups and to individual users based on user activities and responsibilities.

Before you begin

Role required: user_admin or admin

Procedure

Assign the following roles to groups and users as appropriate:

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
Root administrator `[sn_cmp.cmp_root_admin]`	Highest level of application access for Cloud Provisioning and Governance.	All
Cloud administrator `[sn_cmp.cloud_admin]`	Configures the Cloud Provisioning and Governance application and sets up the cloud infrastructure.	Cloud infrastructure: Service accounts and cloud account Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • IBM Cloud Update • VMware Events
Governor `[sn_cmp.cloud_governor]`	Monitors overall cloud usage and enforces compliance of the organization's rules, quotas, and policies. Also manages tags and permissions to various objects.	Governance: • Policies • Pools • Quotas • Permissions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Description and tasks	Access rights in Cloud Provisioning and Governance
Service Designer `[sn.cmp.cloud_service_designer]`	Creates blueprints, ARM and CloudFormation templates, and catalog items.	Cloud Service Design: • Cloud templates • Blueprints • Blueprint catalog items • Resource blocks
Cloud user `[sn_cmp.cloud_service_user]`	Requests and manages stacks and resources.	Cloud User Portal. When you assign the role to a group, all members of the group share quota limitations and ownership of certain resources. You can access all the task records, assigned to you or otherwise, if you have the cloud service user role.
Cloud operator `[sn_cmp.cloud_operator]`	Monitors and troubleshoots the Cloud Provisioning and Governance application.	Dashboards and reports: • Cloud Operations Dashboard • Cloud Root Cause Analysis • Cloud Orchestration Trail • Cloud API Trail
Cloud Event Integration `[sn_cmp.cloud_event_integration]` **Note:** Not supported by IBM Cloud Connector.	Authorizes access to the instance for external services that auto-update the CMDB when cloud events occur. This role gives the access to the cloud event REST endpoint.	Cloud infrastructure: • Service accounts and cloud account • Networks and IPAM Settings for provider services that auto-update the CMDB: • AWS Config • Azure Alert • Google Cloud Logging • VMware Events
Cloud Group administrator `[sn_cmp.cloud_group_admin]`	Grants admin access to any group that you belong to.	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Assign a role to a group

Assign a role to a user

Create the credential and service account that will access your VMware data

To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery, you first create an account in the VMware Console. You then securely store the credentials in a service account in your instance.

Before you begin

Roles required:

- VMware Console administrator
- Cloud Provisioning and Governance: admin or sn_cmp.cloud_admin

About this task

Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Credentials**.
2. On the **Credentials** tab, click **New**, select **VMware Credentials**, fill in the form, and then click **Submit**.

VMware Credentials form fields

Field	Input value
Name	Unique and meaningful name for the credentials (for example, `Cloud Management VMware Credentials`).
Active	Select the check box to enable the credential for use.
User name	User name for access to your vSphere account.
Password	Password for access to your vSphere account.

3. Navigate to **Cloud Admin Portal** > **Manage** > **Service Accounts**, click **New**, enter a unique and meaningful **Name** (for example, `Cloud Provisioning Service Account`), and then fill in the form.

Service Account for VMware form

Account ID	There are two methods for determining the value: ○ If you have access to the managed object browser for the vCenter (MOB): The Account ID is the instance UUID. Open the MOB and then click **Content** > **About**. Paste the value of the **instanceUuid** property in the **Account ID** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

	○ If you do not have access to the MOB, enter a temporary alphanumeric text such as `5ca3033a-1290-4a02-8221-2c5f2f29be8b` in the **Account ID** field. In a step that follows this table, you create the actual value and then return to this form to replace the temporary text with the actual value.
Discovery credentials	Name of the credentials that you created in the preceding step. `Cloud Provisioning VMware Credentials` in the example.
Datacenter URL	Full URL of the vCenter, for example: `https://<IP address>/sdk`.
Datacenter Type	Select **VMware vCenter datacenter [cmdb_ci_vcenter_datacenter]**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.
Should pull events	Enables Cloud Provisioning and Governance to collect events from resource types in VMware vCenter datacenter in the `sn_cmp_cloud_event_list` .

4. Click the **Discover Datacenters** related link.
   Using the credential, Discovery accesses the vCenter at the specified URL to discover VMware datacenters.
5. If you used the instance UUID value in the **Account ID** field, click **Submit** and skip the remaining steps in this procedure.
   The system creates the service account and displays the list of all service accounts.
6. On the Cloud Admin Portal, navigate to **Operate** > **Trails**.
7. On the **Cloud API Trail** tab, find the entry with the **Method name** value of **ListDatacenters** and **Provider name** value of **vmware-compute** (typically the most-recent entry).
   Open the entry.
8. On the **CAPI Trail Logs** related list, open the **route_result** entry.
9. In the **Log value** text box, copy the **locationId** value.
   This value is the **Account ID** that you will paste into the form in the next step.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

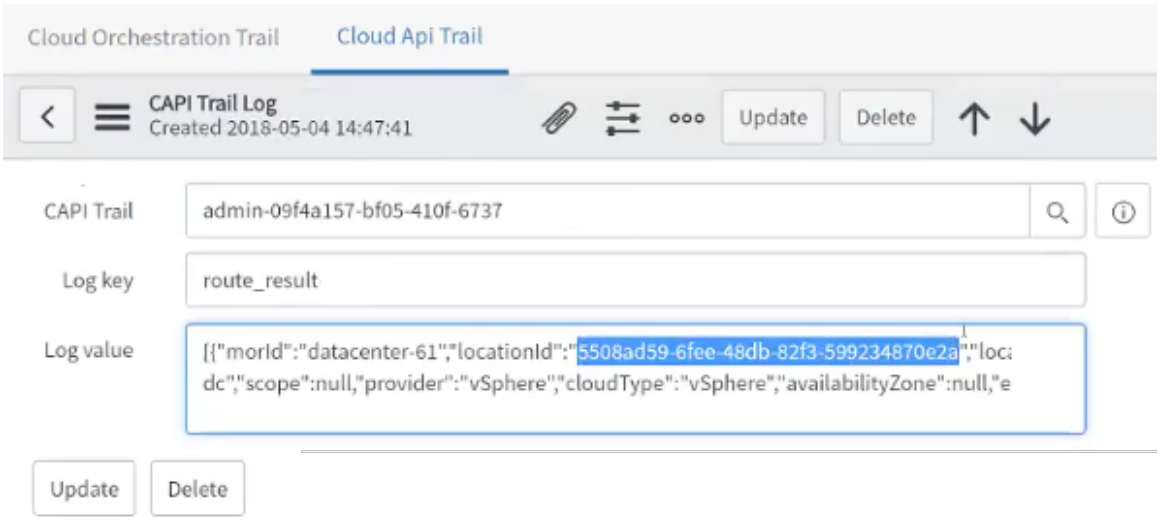

10. On the instance, navigate to **Cloud Admin Portal** > **Manage** > **Service Accounts** and then click the service account that you created earlier (`Cloud Provisioning Service Account` in the example).
11. Paste the **locationId** value (that you copied from the **CAPI Trail Log** form) in the **Account ID** field on the form and then click **Update**.
12. Click **Submit**.
    The system creates the service account and displays the list of all discovered datacenters.

Related topics

Discovery for VMware

Set up cloud accounts for VMware

A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

Before you begin

Role required: sn_cmp.cloud_admin

About this task

See Cloud accounts for detailed information.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Cloud Accounts**.
2. Click **New**, enter a unique and meaningful **Name** and **Description**, and then select the cloud **Provider** for the account (VMware vCenter).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

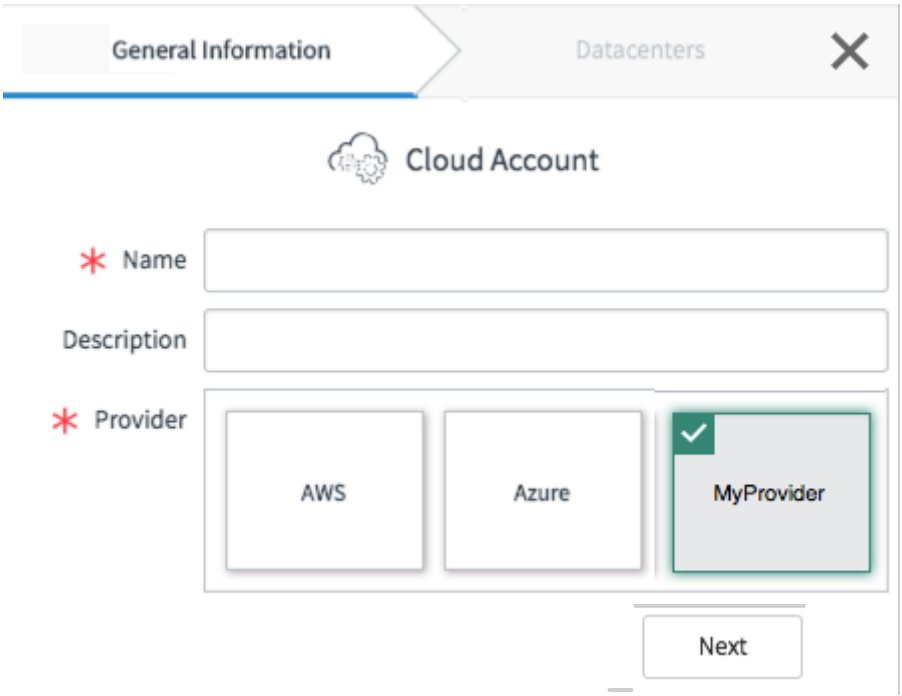

3. Click **Next**.
   Based on the provider you select, the **Service Account** field is populated and a list of discovered datacenters appears. If no datacenters appear, click **Discover Datacenters** to find them. See Discover all datacenters in a service account on-demand for details.
4. Select one or more datacenters to associate with the cloud account, click **Save**, and then close the popup.
   The Cloud Account dashboard appears. The datacenters that you selected appear on the **Datacenters** tab.
5. Click a datacenter.
   The following lists appear:
   - **Resources**: Cloud resources for the current datacenter. Run Discovery to populate the CMDB for the datacenter and populate the tab. See Discover all datacenters in a service account on-demand.
   - **Capacity Limits**: Limits on virtual CPUs, virtual networks, storage volume size, and other services. Set capacity limits on user requests for resources
   - **Discovery Log**: The process creates a log record for each action in an datacenter.

   Note:

   If you are on a domain separated instance, only those events that are updated to the CMDB and belong to your domain are visible. Events create configuration items (CI) in the same domain as the cloud service account they are mapped to.

What to do next

Repeat the process to create additional cloud accounts as needed. Run Discovery and set capacity limits for each datacenter.

Add a datacenter to a cloud account

At any time, you can add a logical datacenter to the cloud infrastructure that is represented by a cloud account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_cmp.cloud_admin

Procedure
1. In the Cloud Admin Portal, navigate to **Manage** > **Cloud Accounts**.
2. Open a cloud account, set the state to **Draft**, and then click **Configure** in the form header.
3. Click the **Datacenters** tab and then select the **Service Account** that holds the credentials that enable access to the provider account that includes the datacenter to add.

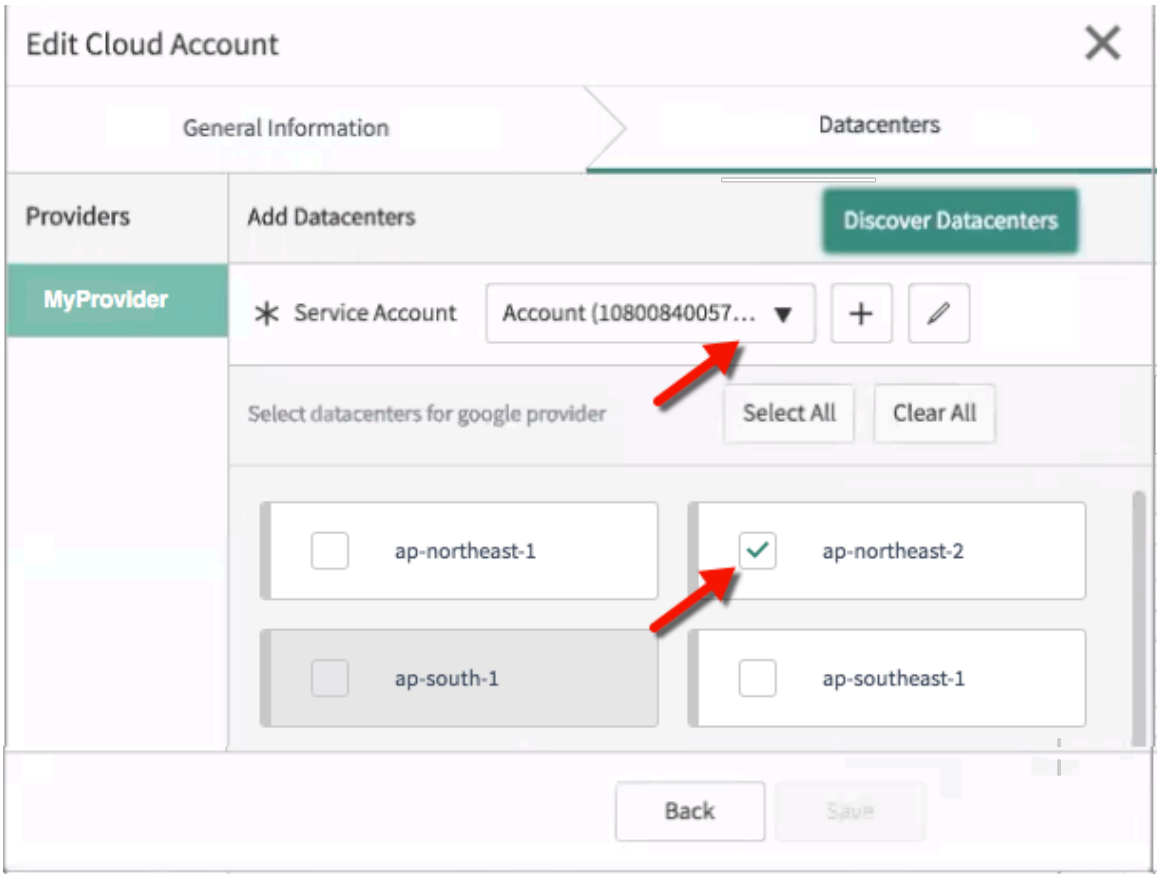

When you select the service account, the list of discovered datacenters for the service account appears.

Note: If the expected datacenters do not appear, click **Discover Datacenters** to update the list. Discovery runs and displays all datacenters associated with the service account.

4. Select the datacenters to add and then click **Save**.
   In the example, you select **ap-northeast-2**.

Result
The datacenters are added to the cloud account and appear on the **Datacenters** tab. When Discovery runs, the resources in the datacenter appear on the Resources tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Datacenter added to the cloud account

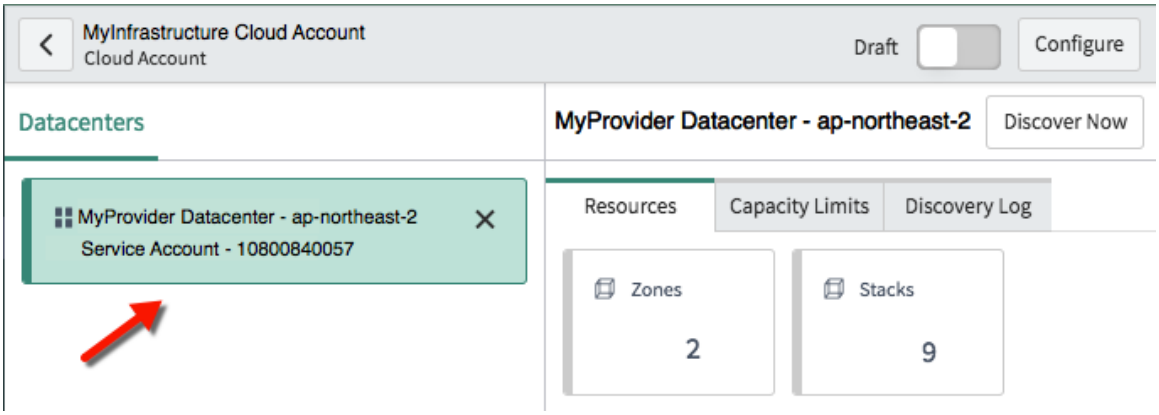

Additional Cloud Provisioning and Governance setup on day 2

After you have performed Day 1 setup, you can perform optional setup and configuration procedures as needed and in any order. Detailed instructions for each procedure follow this overview.

Additional setup and configuration operations

Detailed instructions for each procedure follow this overview.

Set up an additional cloud account

A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

Schedule discovery of all resources in a service account

Add a custom schedule for Discovery to discover all resources in each datacenter in a service account and then update the CMDB with the new information.

You can use the discovery schedules as offered on the cloud discovery schedule page.

Note:

Starting Xanadu release, any new user will run Pattern-based discovery, by default. In scenarios where the latest patterns application is not installed, discovery fails and the System Admin is prompted to install or migrate from CAPI-based discovery.

Installing Cloud Discovery Workspace triggers automatic installation of the latest patterns application.

For more information, see Migrate from CAPI to Patterns.

Discover all resources in a datacenter on-demand

At any time, you can run Discovery on a datacenter to update the CMDB.

Set capacity limits on user requests for resources

Capacity limits place restrictions on the attributes of cloud resources such as the number of virtual machines, virtual CPUs, or aggregate storage. You can set limits on resources separately for each logical datacenter in a cloud account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Next steps

When you have finished all Day-1 and Day-2 procedures in this setup guide, see the Cloud Provisioning and Governance administration guide for information on using the Cloud Provisioning and Governance application in your organization.

Set up an additional cloud account

During initial installation, you set up one cloud account. To organize and compartmentalize your infrastructure, you can set up additional cloud accounts to include different providers or service accounts or datacenters.

Before you begin

Role required: sn_cmp.cloud_admin

About this task

A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account. See Cloud accounts for detailed information.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Cloud Accounts**.
2. Click **New**, enter a unique and meaningful **Name** and **Description**, select the cloud **Provider** for the account (**MyProvider** in the example), and then click **Next**.

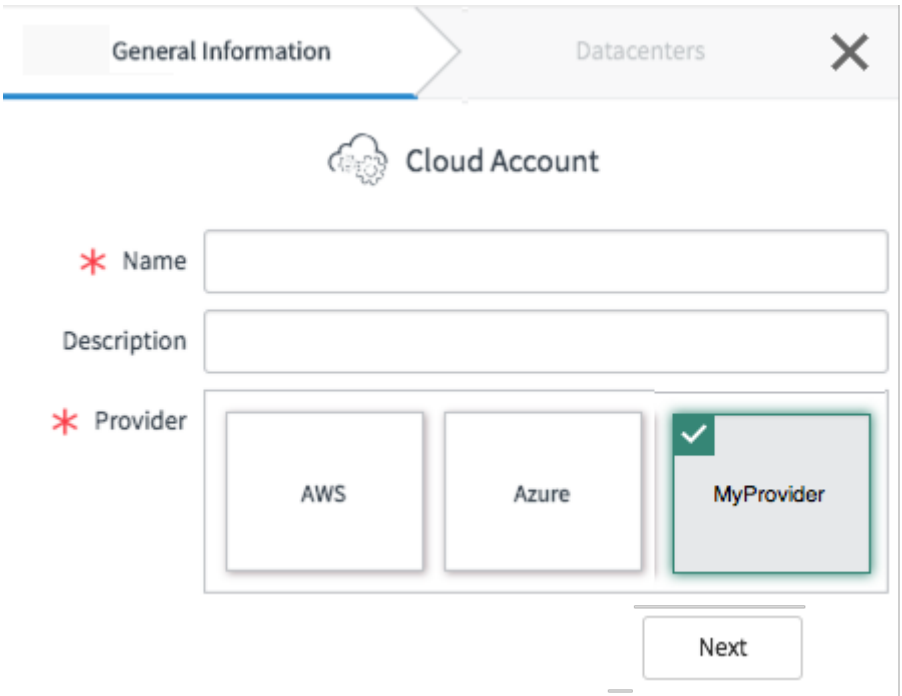

   Based on the provider you select, the **Service Account** field is populated.
3. Select the appropriate **Service Account**.
   A list of discovered datacenters appears. If no datacenters appear, click **Discover Now**. See Discover all datacenters in a service account on-demand for details.
4. Select each of the datacenters that you want to associate with the cloud account, click **Save**, and then close the popup.

   The datacenters appear on the **Datacenters** tab for the cloud account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

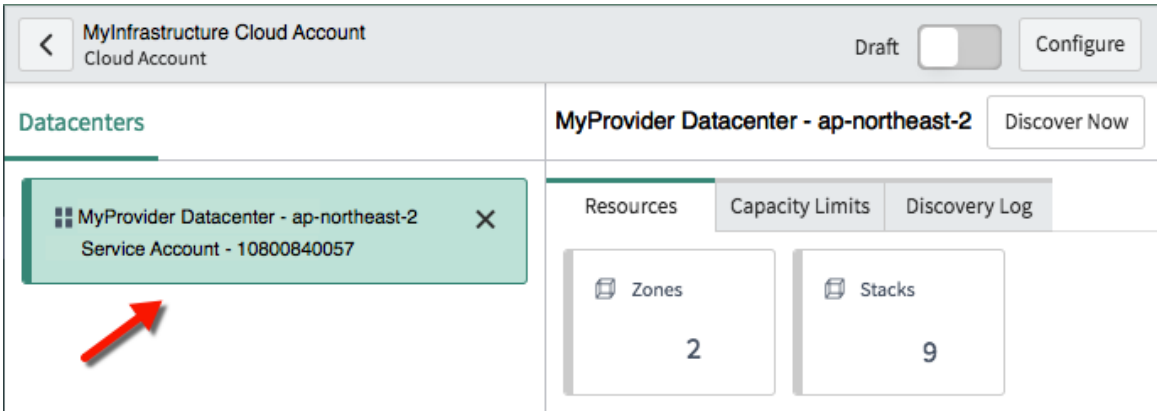

5. When the datacenters appear, select one or more datacenters to include in the cloud account and then click **Save**.

   Note:

   - Select only those LDCs/Regions where your infrastructure resources are present. If you don't have any resources under these new data centers, exclude them while you run the full discovery.
   - To discover resources in certain regions or LDCs, raise an access request to the required endpoints with GCP Support. These regions include North America (US East, US West, and Canada Central), South America, Europe, Africa, Middle East, and Asia Pacific.

   The Cloud Account dashboard appears. The datacenters that you selected appear on the **Datacenters** tab.

6. Click a datacenter.
   The following lists appear:

   - **Resources**: Cloud resources for the current datacenter. Run Discovery to populate the CMDB for the datacenter and populate the tab. See Discover all datacenters in a service account on-demand.
   - **Capacity Limits**: Limits on virtual CPUs, virtual networks, storage volume size, and other services. See Set capacity limits on user requests for resources.
   - **Discovery Log**: The process creates a log record for each action associated with a discovery status.

7. Repeat the process to add as many service accounts as needed.

What to do next

Repeat the process to create additional cloud accounts as needed. Run Discovery and set capacity limits for each datacenter.

Discover using Cloud Operations Workspace and Patterns

You can now schedule and perform discovery using Cloud Discovery Workspace and Discovery and Service Mapping Patterns.

Before you begin

Role required: System admin

About this task

Cloud Operations Workspace is now known as Cloud Discovery Workspace or CDW.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Procedure

1. Check if *Cloud Discovery Workspace* (CDW) is installed.
2. Navigate to **Cloud Admin Portal** > **Manage** > **Discovery Schedules** > **ManageDiscovery Schedules**.

   Follow these scenarios for next steps.

Scenario	Next steps
Cloud Discovery schedules don't dis play for System Admin login	○ *Cloud Discovery Workspace* isn't installed. The page notifying that *Cloud Discovery Workspace* must be installed, displays. ○ The page prompts to **Go to Applications Manager** to allow installation
Cloud Discovery schedules don't dis play for sn_cmp Cloud Admin login	○ *Cloud Discovery Workspace* isn't installed. The page notifying that *Cloud Discovery Workspace* must be installed, displays. ○ The page requests to contact the system admin for *Cloud Discovery Workspace* installation.
Cloud Discovery displays	Cloud Discovery schedules display in a new tab of your browser.

3. Ensure that the sn_cmp cloud admin role contains the role 'sn_cloud_ops_ws.cloud_ops_admin'.
4. Navigate to **Cloud Admin Portal** > **Manage** > **Discovery Schedules**.
5. Click and confirm when Discovery Schedules open in a new tab.
6. Ensure that the sn_cmp Cloud Admin is able to view the Discovery Schedules page by default.

What to do next
Install and use the latest Discovery and Service Mapping Patterns, at least version 1.15, as the default mechanism for performing Discovery.

Discover all datacenters in a service account on-demand

After you set up a service account, run Discovery to populate the datacenters in the account.

Before you begin

- A cloud account must exist with service accounts and associated datacenters. You must run on-demand Discovery to discover datacenters before you can schedule Discovery.
- Use the Cloud Discovery Workspace, and at least version 1.15 of Discovery and Service Mapping Patterns from the ServiceNow Store.

Role required: sn_cmp.cloud_admin

About this task
Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings. In the Google Cloud Platform, the region-specific containers for virtual resources are called regions. In ServiceNow cloud data model, regions are called

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

datacenters or logical datacenters (LDCs). The term logical is used to reinforce the idea that Cloud Provisioning is provider-agnostic. All infrastructure or applications that are deployed using Cloud Provisioning are associated with a datacenter.

You can run Discovery on-demand.

Important: Use Cloud Discovery Workspace (CDW) to run the discovery schedule. The discover datacenter UI action is legacy and the preferred approach is to setup schedules from CDW.

Procedure

1. Navigate to the Cloud Service Accounts [cmdb_ci_cloud_service_account_list] table.
   For more information on directly navigating to a table, see Navigate directly to a table.
2. Select the service account whose LDCs you want to discover.
3. Select the **Create Discovery Schedule** related link.
4. Select **Next**.
5. Select **Refresh Datacenters**.
6. Select **Cancel**.
7. Select the service account whose LDCs you tried to discover.
   On successful datacenter discovery, the Datacenter Discovery status field is set to `YYYY-MM-DD HH:MM:SS: Completed … Reload this form to see an updated datacenter list`.

Result

The discovered datacenters appear in the Logical Datacenters related list.

Related topics

Logs for horizontal discovery

Schedule discovery of all resources in a service account

Add a custom schedule for Discovery to discover all resources in each datacenter in a service account and then update the CMDB with the new information.

Before you begin

- A cloud account must exist with service accounts and associated datacenters. You must run on-demand Discovery to discover datacenters before you can schedule Discovery.
- Use the Cloud Discovery Workspace, and at least version 1.15 of Discovery and Service Mapping Patterns from the ServiceNow Store.

Role required: sn_cmp.cloud_admin

About this task

- Scheduled Cloud Discovery discovers only datacenters that are part of the service account when the schedule is created. If datacenters have been added since the time you created the schedule, Discovery does not discover those datacenters. Delete the old schedule and create a schedule every time the datacenters in a service account change.
- Discovery populates the CMDB with the discovered resources and updates the resource information on the **Resources** tab. You can view Discovery results on the **Discovery Log** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: Use the Cloud Discovery Workspace to run the discovery schedule. The discover datacenter UI action is legacy and the preferred approach is to setup schedules from CDW.

Procedure

1. In the **Cloud Admin Portal** navigate to **Manage** > **Discovery Schedules**.
2. Select **Add Cloud Schedule**.
3. On the Add Account page, select the cloud **Provider** and enter a unique, descriptive **Name** for the schedule.
4. In the **Service Account** section, select **Select Account**, specify the service account, and then click **Next**.

   Note: To create a service account, select **Add Account** and then fill in the form. For details, see the process of setting up a service account in the Day 1 setup guide.

5. On the Select Datacenters page, specify which datacenters to discover.

   Move individual datacenters to the **Selected** list, or select **Discover all datacenters listed below** to move all datacenters.

6. Click **Next**.
7. On the Create Schedule page, you specify the time and frequency of discovery.

   Note: **Run time** setting uses the system time zone.

Discovery schedule run options

Frequency	Description
On Demand	Doesn't run on a schedule. Select **Discover now** to run Discovery. This is the defau
Daily	Runs every day. Use the **Start time** field to specify the time of day.
Weekly	Runs on one designated day of each week. Use the **Start time** field to specify the t
Monthly	Runs on one designated day of each month. Use the **Day** field to select the day of to specify the time of day. If the designated day doesn't occur in the month, the sc example, if you designate day 30, the schedule doesn't run in February.
Periodically	Runs every designated time. Use the **Repeat Interval** field to define the time in day The first Discovery runs at the point in time defined in the **Starting** field. The subse **Repeat Interval** period passes.
Once	Run one time as designated by the date and time defined in the **Starting** field.
Weekdays	Runs every Monday, Tuesday, Wednesday, Thursday, and Friday. Use the **Start tim**
Weekends	Runs every Saturday and Sunday. Use the **Start time** field to select the time of day.
Month Last Day	Run the last day of every month. Use the **Start time** field to select the time of day.
Calendar Quarter End	Runs on March 31, June 30, September 30, and December 31. Use the **Start time** fi change the dates, modify the *`DiscoveryScheduleRunType`* script include.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Frequency	Description
After Discovery	Allows you to sequentially stagger the schedule. Use this option to run this schedu in the **Run after** field finishes. Select the **Even if canceled** check box to designate if the **Run after** Discovery is canceled before it finishes. ○ This option isn't valid when the Discovery is started via DiscoverNow, or when us ○ You can't designate an inactive Discovery schedule. ○ You can't create a loop by designating the run after Discovery to be the same Dis ○ This Discovery doesn't run if the **Run after** Discovery doesn't finish, with the exc check box is selected and the Discovery is canceled.
Cancel Discovery if longer than	Time limit for running the Discovery. When the specified time elapses, the remainir canceled even if the Discovery isn't finished. If the check box isn't selected, the Di setting constrains system load to a desirable time window.

8. Select **Finish** and **Run Now**.

Result

Discovery runs immediately. Discovery runs at the specified times and adds discovered datacenters to the Logical Datacenters related list.

If you configured a Discovery schedule, it appears in the Discovery Config - Logical Datacenters related list. To run Discovery on a single datacenter, add the datacenter to the Discovery schedule in the related list, then run Discovery again.

Open a datacenter record to see the relationships between the datacenter and the resources in the datacenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resources in an AWS datacenter

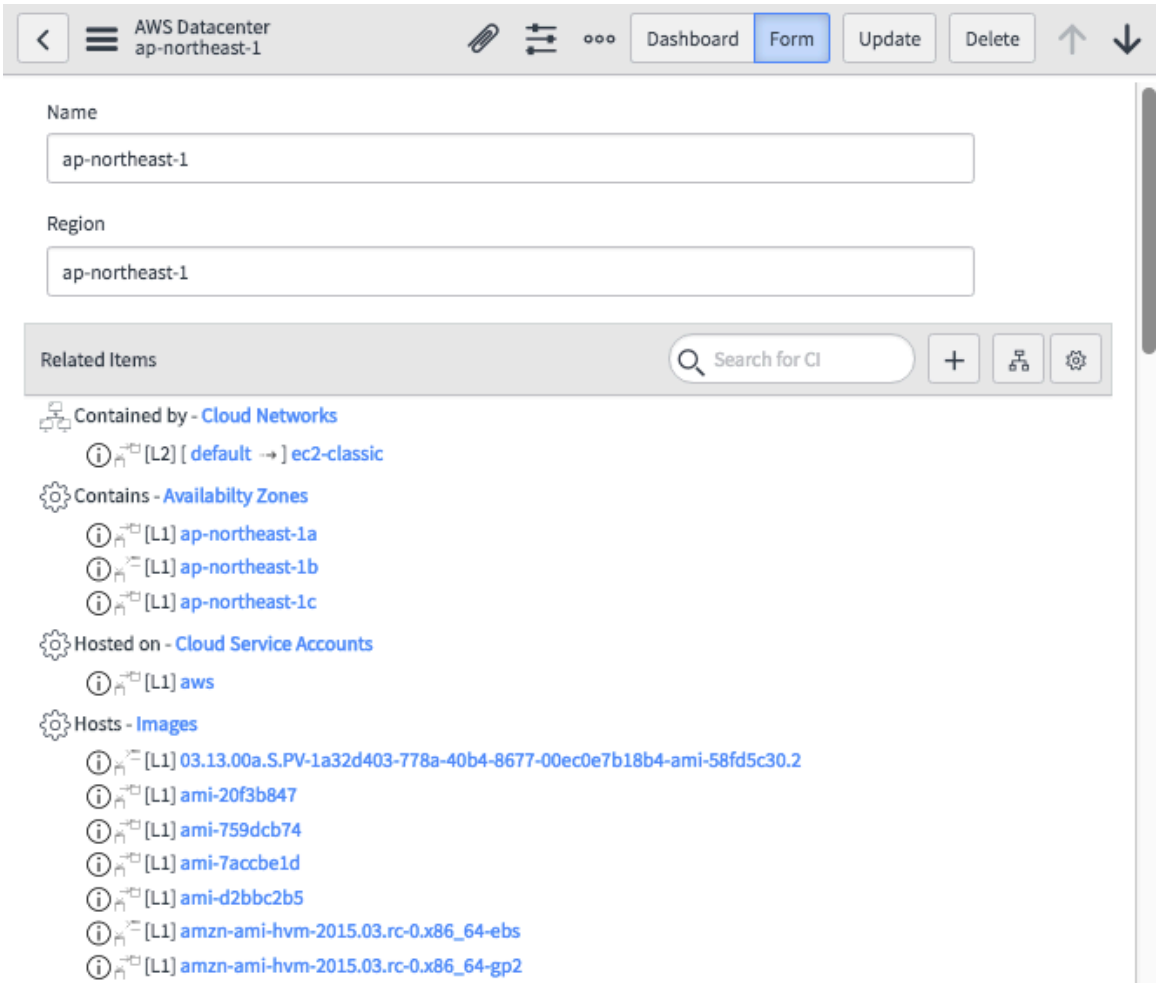

Note: By default, Discovery stores the discovered VM hardware types in the Hardware Type [cmdb_ci_compute_template] table. If you notice an unusually large number of records in this table, use the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table to store the discovered VM hardware types. For more information, see Enable the Cloud Hardware Type class extension. All Cloud Provisioning and Governance features, such as provisioning, quotas, and so on continue to work as is after you switch to the Cloud Hardware Type table.

What to do next

- To view discovery status, discovery logs, discovered devices, and the ECC queue, click **Advanced View**.
- If Discovery fails, check the Cloud Discovery results on the Discovery Home page.

Discover all resources in a datacenter on-demand

At any time, you can run Discovery on a datacenter to update the CMDB.

Before you begin

- A cloud account must exist with service accounts and associated datacenters. You must run on-demand Discovery to discover datacenters before you can schedule Discovery.
- Use the Cloud Discovery Workspace, and at least version 1.15 of Discovery and Service Mapping Patterns from the ServiceNow Store.

Role required: sn_cmp.cloud_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Cloud Accounts**.
2. Click a cloud account and, on the **Datacenters** tab, select a datacenter.
   The **Resources** tab lists the resources in the CMDB.

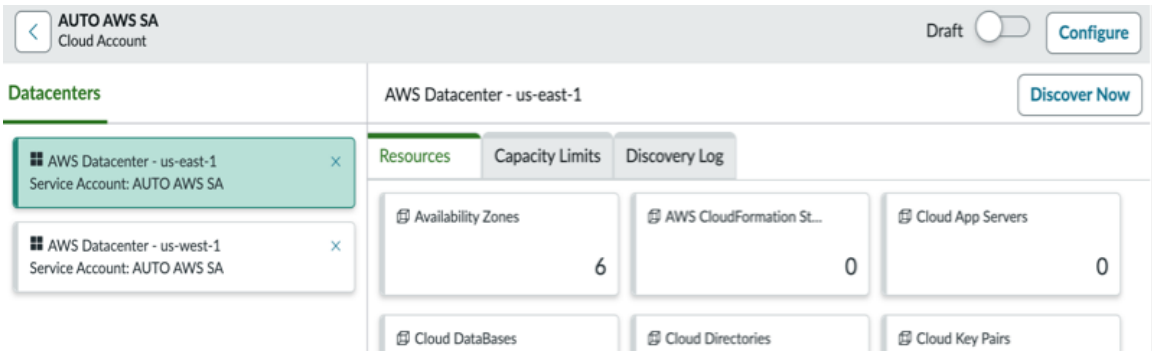

3. Change the state to **Draft** and then click **Discover Now**.
   Note that **Discover Now** is disabled if neither CAPI or Patterns is installed. Follow these scenarios to proceed.

   Note: **Discover now** will also be disabled with an message suggesting you to install the latest Discovery and Service Mapping Patterns. This can be fixed by installing version 1.15.1 at least or by triggering the manual migration step.

Scenario	Next steps
Discover Now is disabled for System Admin login	○ The page suggests that you must install the latest patterns application to proceed. ○ The banner at the top of the page prompts you to visit **Go to Applications Manager** for installation
Discover Now is disabled for sn_cmp Cloud Admin login	○ The page suggests that you must install the latest patterns application to proceed. ○ The page requests to contact the system admin for installation of patterns application.
Latest patterns application is in stalled	Discover Now is activated. **Note:** Discovery will now use Patterns as the default mechanism.

4. Run **Discover Now** and wait for Discovery to finish.
5. Change the state of the cloud account to **Published**.

   Discover now will be disabled with the same message shown. This can be fixed by installing 1.15.1+ version or by triggering the manual migration step

Result

Discovery populates the CMDB with the discovered resources and updates the information on the **Resources** tab. You can view Discovery results in the **Discovery Log** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: By default, Discovery stores the discovered VM hardware types in the Hardware Type [cmdb_ci_compute_template] table. If you notice an unusually large number of records in this table, use the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table to store the discovered VM hardware types. For more information, see Enable the Cloud Hardware Type class extension. All Cloud Provisioning and Governance features, such as provisioning, quotas, and so on continue to work as is after you switch to the Cloud Hardware Type table.

Related topics

View the Discovery Log

Set capacity limits on user requests for resources

Capacity limits place restrictions on the attributes of cloud resources such as the number of virtual machines, virtual CPUs, or aggregate storage. You can set limits on resources separately for each logical datacenter in a cloud account.

Before you begin

- A cloud account must exist with service accounts and associated datacenters. You must run on-demand Discovery to discover datacenters before you can schedule Discovery.
- Use the Cloud Discovery Workspace, and at least version 1.15 of Discovery and Service Mapping Patterns from the ServiceNow Store.

Role required: sn_cmp.cloud_admin

About this task

Note: You can also use quotas to set per-user and total capacity limitations on resources. See Quotas and resource order controls.

Procedure

1. Navigate to **All** > **Cloud Admin Portal** > **Manage** > **Cloud Accounts**.
2. Click a cloud account, and then, on the **Datacenters** tab, select a datacenter.
3. Set the state to **Draft**.
4. On the **Capacity Limits** tab, set the **Capacity Limit** values as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

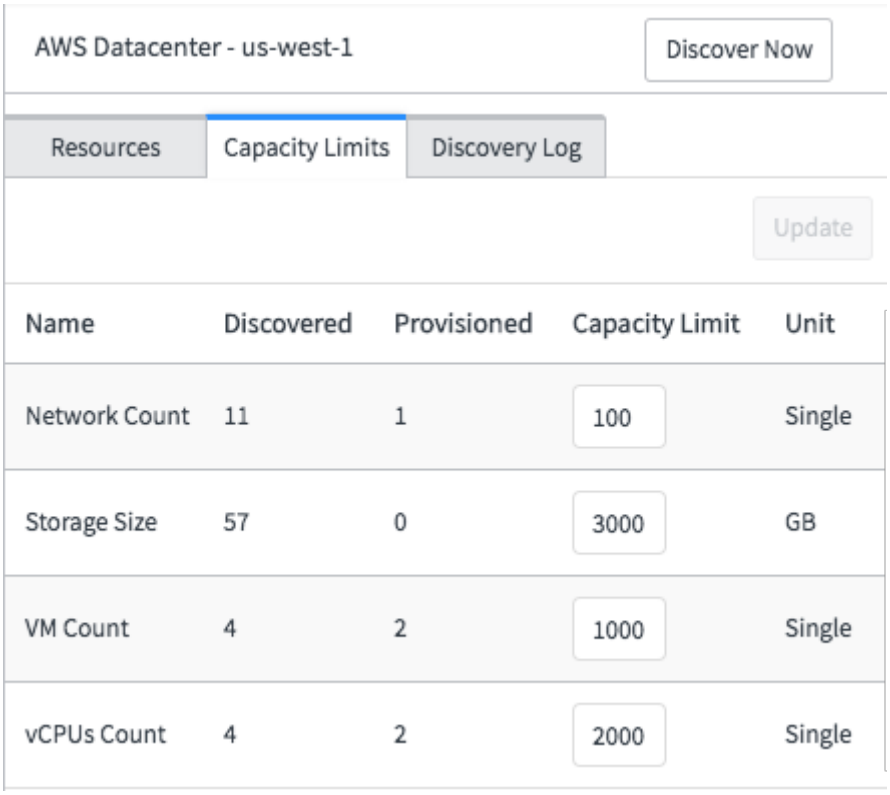

Value	Description
Discovered	Total value discovered by the Discovery process for all resources in the selected service account.
Provisioned	Total value provisioned for the selected service account.
Capacity Limit	Maximum allowed value calculated as the sum for all resources in the selected service account.
Unit	Unit of measure for the specified value. For example, storage is measured in units of GB.

5. Click **Update**.
6. Set the state to **Published**.

What to do next
Click **Discover Now** to run Discovery on the logical datacenters.

Use owner and assignment groups

You must set up and use the user groups in Cloud Provisioning and Governance as a part of the day-2 task.

Before you begin
Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Owner group- This is used in quota management. During any stack provision, the user must select the "User group" on the order form. A quota check is performed based on the user group selected on the form, and this user group is later added as the stack 'Owner group' in the sn_cmp_stack table. Quota checks can be defined at the user group level. For more details, see Quotas and resource order controls.

Note: If the stack is deleted or the owner group field changes for the stack, the quota is recalculated. The quota is recalculated in the business rule: Recalc Consumed Quota on Stack Update.

Assign group- The discovery schedule runs bring the resources created from other instances into your instance. These resources are unmanaged. Similarly, the 'Fill in tag assignments' scheduled job creates unmanaged stacks. To manage these stacks and resources, you can assign a change group to them.

Assign a change group to an unmanaged stack or resource, to make it visible in the **Cloud User Portal**, enabling you to perform day 2 operations on this resource.

Procedure

1. Navigate to **Cloud User Portal**.
2. Click **View Stacks** and navigate to **Discovered Stacks** for unmanaged stacks.
3. Click **View Resources** and navigate to **Virtual Machines**for unmanaged resources.
4. Change the view to **My Groups**.
5. Select any of the virtual machines and perform day 2 operation.

Cloud Provisioning and Governance administration guide

After you set up the Cloud Provisioning and Governance application on your instance, there are many items that you must configure before your users can request and manage any cloud resources.

Requirements

Role required to use the Cloud Provisioning and Governance administration guide is cloud_admin.

Before you begin

Set up Cloud Provisioning and Governance

See the appropriate setup guide for your cloud provider:

- Cloud Provisioning and Governance setup guide for AWS
- Cloud Provisioning and Governance setup guide for Azure
- Cloud Provisioning and Governance setup guide for VMware

.

What to do after setup

Create resource profiles

Resource profiles contain the attributes that a resource can have. Users select values for these attributes when they order a cloud resource. You can configure these types of resource profiles: Compute (such as the size of a VM), Operating

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

System (such as the type and version), Storage (such as backup), and Security Groups (such as whether to use HTTPS). See Resource Profiles for more information.

Configure governance

Governance tools that you can use to restrict the provisioning of cloud resources include quotas and policies. Examples of things you can configure include the approval subflow that is used after a user requests a cloud resource, whether an email is sent after a cloud resource is stopped, and which datacenter should be automatically selected when a user selects a resource from a specific cloud provider. See Cloud Governance for more information.

Populate the Cloud Service Catalog

After you apply governance quotas and policies, create Service Catalog items. You can create a catalog item in the following ways:

- Import a template from AWS or Microsoft Azure
- Use or modify a base-system catalog item
- Use a Terraform template

Domain separation and Cloud Provisioning and Governance

Domain separation is supported in Cloud Provisioning and Governance. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

Cloud Provisioning and Governance domain separation overview

All tables in Cloud Provisioning and Governance are not domain separated. Delegated domain separation is not supported.

Domain separation for Cloud Provisioning and Governance supports:

Service Providers (SPs) using the application to provide data separation.

In this scenario, SPs can provide data separation to multiple customers, where domains are necessary to contain all relevant customer data and processes. For example, an SP provides support to customers who typically use Cloud Provisioning and Governance to manage their IT infrastructure on the cloud. SPs can provide catalogs, template profiles, resource pools, and filter, resource profiles, quotas,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

and permissions as domain-separated offerings to their customers. They can also provide IP address management (IPAM), lease and business hours scheduling, and a view to billing.

How domain separation works in Cloud Provisioning and Governance

Domain separation for Cloud Provisioning and Governance aligns one or more companies to a domain. To use domain separation with the application, assign all user accounts to a specific company associated with that domain.

All entities that are related to the company, such as cloud accounts and service accounts, are created in the same domain as the company. When a new company is created, create a domain with a unique name and assign it to the company. All related entities for an account, such as contacts and cases, must reside in the same domain. When you create a related entity for a domain-separated account, the entity is assigned to the company domain.

Members of a domain can only view the data that is contained within their domain or child domains that are lower in the domain hierarchy. By default, all users and all records are members of the global domain unless you assign them to a particular domain. After you assign a user or a record to a domain, the instance compares the domains. The instance uses this comparison to determine whether the user can view the record.

Service Providers (SPs) use domain separation to segregate data for each customer. Users in a given domain can only view the data in their own domains or in child domains. SPs typically control the top-level domain, which allows them to view data that is associated with all domains. Don't delegate administration to cloud admin users of the child domains in Cloud Provisioning and Governance.

Set up domain separation for Cloud Provisioning and Governance

Ensure that you activate the following plugins:

Domain Support - Domain Extensions Installer plugin (com.glide.domain.msp_extensions.installer) to enable domain separation in Cloud Provisioning and Governance

Service Catalog - Domain Separation plugin (glideapp.servicecatalog.domain_separation) to enable separation of catalog items in different domains in Cloud Provisioning and Governance

Changes to Cloud Provisioning and Governance tables

Domain separation for Cloud Provisioning and Governance adds the **Domain** and **Domain Path** fields to the list views. These fields are not exposed by default. As a domain admin you can customize lists and forms to view these fields. Not all tables in Cloud Provisioning and Governance are domain separated. While some top-level tables are domain separated, several child tables are not domain separated. However, this does not impact how the Cloud Provisioning and Governance application works in a domain-separated context.

Account domains and related entities

When you create related entities for an account, the domain for the related entities is set to the account domain.

Domain visibility for cloud administrators and users

Manually assign users with the Cloud User Portal (sn_cmp.cloud_service_user) roles and Cloud Admin Portal (sn_cmp.cmp_root_admin) roles for each domain to the **TOP/MSP/Default/Company** or leaf domain. Domain administrators and users in Cloud Provisioning and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Governance can only view data in the domain that they are created in, until they are assigned to the TOP domain. The Top domain represents a single common parent domain, which acts as a single parent node, for the Service Provider domains.

Next Steps

For more information on creating, implementing, and maintaining domain separation for Cloud Provisioning and Governance services in the instance you are setting up for your customers, see Domain separation in Cloud Provisioning and Governance - considerations for service providers.

Related topics

Domain separation for service providers

Domain separation in Cloud Provisioning and Governance - considerations for service providers

Review the following considerations to effectively create, implement, and maintain domain separation for Cloud Provisioning and Governance services in the instance that you are setting up for your customers.

- Includes all aspects of **Basic** level support.
- Application properties are domain-aware as needed.
- Business logic: The service provider (SP) creates or modifies processes per customer. The use cases reflect proper use of the application by multiple SP customers in a single instance.
- The instance owner must configure the minimum viable product (MVP) business logic and data parameters per tenant as expected for the specific application.

How domain separation works in Cloud Provisioning and Governance

Domain separation for Cloud Provisioning and Governance aligns one or more companies to a domain. To use domain separation with the application, assign all user accounts to a specific company associated with that domain.

All entities that are related to the company, such as cloud accounts and service accounts, are created in the same domain as the company. When a new company is created, create a domain with a unique name and assign it to the company. All related entities for an account, such as contacts and cases, must reside in the same domain. When you create a related entity for a domain-separated account, the entity is assigned to the company domain.

Members of a domain can only view the data that is contained within their domain or child domains that are lower in the domain hierarchy. By default, all users and all records are members of the global domain unless you assign them to a particular domain. After you assign a user or a record to a domain, the instance compares the domains. The instance uses this comparison to determine whether the user can view the record.

Service Providers (SPs) use domain separation to segregate data for each customer. Users in a given domain can only view the data in their own domains or in child domains. SPs typically control the top-level domain, which allows them to view data that is associated with all domains. Don't delegate administration to cloud admin users of the child domains in Cloud Provisioning and Governance.

Pool-Filter, Resource profile, Billing, Quota, Permissions, IPAM, Budget, and Notifications are domain separated. Not all tables in Cloud Provisioning and Governance are domain separated. While most top-level tables are domain separated, several child tables are not domain separated. However, the fact that all tables are not domain separated, doesn’t impact the working of the application in a domain-separated context.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Warning: Do not modify the blueprint-related tables directly in the table records.

Instance setup considerations for Service Providers

While domain separation provides multi-tenancy support, multi-tenancy is still contained within a single instance. Some global properties, data, and processes are shared across all domains. The following image illustrates domain separation hierarchy for unrestricted and restricted use cases.

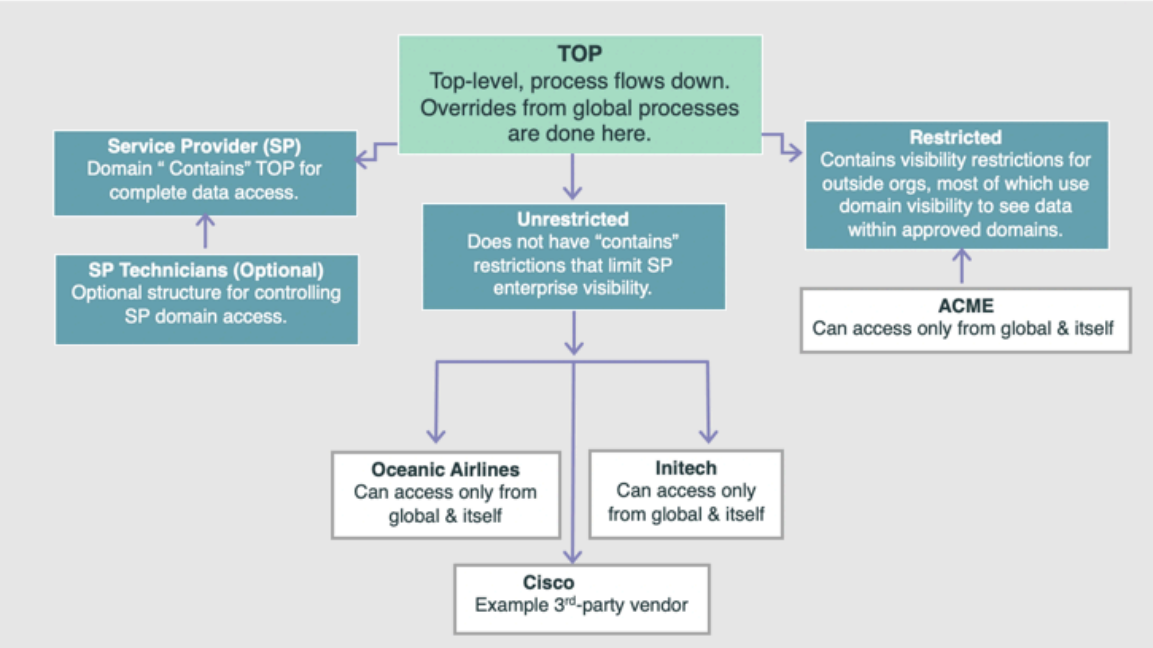

Perform the following tasks to set up your domain-separated instance:

- **Create a domain.**

    - You can create as many domains as needed, but make sure that you don't create unnecessary domains on the instance. A large number of domains on the instance can impact performance. Before you create new domains, check the domain hierarchy under **Domain Admin** > **Domain Map**. Check whether you actually need to create a new domain, or an existing domain hierarchy can work instead.

      Select a primary domain to indicate the top-level domain in the domain map. The Top domain represents a single common parent domain, which acts as a single parent node, for the Service Provider domains. For more information, see Select a primary domain.

    - Create the SP domain under the top domain and set the domain as default. Keeping your domains organized is a crucial part of the domain separation process. Ensure that you are aware of the importance of setting a default. The system automatically assigns task and user records that are not already assigned to a domain to the default domain. For more information, see Make a domain the default for instructions on how to make a default domain.

    - Create a domain hierarchy. For more information, see Create Domain hierarchy.

      Create a domain hierarchy that mirrors organization structures for grouping under the top domain. Because you can perform discovery and provisioning

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

only from the leaf domain, set up leaf domains in the hierarchy that represent concrete entities and organizations.

Based on the domain hierarchy, users have access to data in their home domain and any child domains. The process flows down and the data bubbles up. You can assign user records with a domain value that represents the user's home domain. While users have no access to data in parent domains, peer domains, or domains in other branches of the hierarchy. A process that is setup in a parent domain is applicable to all child domains. For more information, see

- Service provider reference architecture
- Domain separation architecture data

Customer data is affected based on domain hierarchies:

- Parent/child: Process and data affected
  - Design that is based on a process flow.
  - Remember that the parent domains can access all data in the child domains.
- "Contains" domain: Only data is affected. For example, making SP in the diagram contain TOP does not make processes in SP run in the TOP domain and downward.
  - Grants data access rights to individuals in groups that require dedicated access to certain domains.
  - Contains causes or conditions to be added to database queries that can cause performance issues with large domain and data sets.
- Visibility: Hierarchy that is always visible to users once you provide access. Only the data is affected, not the processes.
  - Grants data access of a domain to another domain that did not have that access when the parent-and-child hierarchy was built.
  - Enables users to see all the data in the domains that they have visibility access for, all the time, regardless of the record they are working on.

    Note: Use sparingly, as Visibility can allow complete access that you may not intend.

- Add a new company

You can add companies that represent vendors, manufacturers, or customers with whom you do business. These companies provide a way to categorize users, groups, and assets.

Create companies and leaf domains for each customer that you are supporting on a single instance. Companies help you associate users to the domain. Company and domain must have either one-to-one mapping or many-to-one mapping. This makes assigning users to a domain easier and manageable for SPs.

Important: When you create a company for each domain, use a prefix with a unique identifier, such as "COMPANY_" or "ORG_".

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To customize the ServiceNow instance for each company, you can enter contact phone numbers, street addresses, and additional notes. You can also customize the company logo and banner text that your end users see at the top of each page. For more information see Create a new company profile.

- Create user groups and assign roles and domain to the users of each new company. You assign users to companies to associate them with domains. Associate a company with a user to map the user to a domain.

  1. Navigate to **Organizations** > **Company,** and click **New** to create new users. For more information on creating users, see, Create a user

     Note: Don't use the manual-managed domain option to choose the domain for users.

  2. Create groups and assign roles to the groups. Users that are assigned to the group inherit the roles. You can create cloud user portal users for a company in the leaf domain. For more information on creating groups, see Create a user group.
  3. Create these users in each of the leaf domains:
     - MID Server user for each leaf domain and company. Assign the mid_server role to MID Server users.
     - Domain admin for each leaf domain and company. Assign the sn_cmp.cmp_root_admin role to cloud (domain) administrators.
     - Domain user for each leaf domain. Assign the sn_cmp.cloud_service_user role to cloud users.
     - (Optional) Create users for additional roles in leaf domains (roles like designer, governor).

       To provide limited visibility for certain domains to certain users, add them via groups. For more information, see Grant visibility domains to an individual user.
  4. Create a "contains" domain relationship between the SP domain and the TOP domain.

     Note: By creating a "contains" relationship, you allow all the administrators in the Service Provider domain to have access to all the domain data. Users of global domain are unable to see data from card layout as by default expand domain scope is not allowed.

     For more information on how to control what specific users and what an entire domain of users can view, see Visibility domains and Contains domains.

Next Steps

For more information on onboarding customers or companies in a domain-separated instance for Cloud Provisioning and Governance services, see On-board a Company.

Related topics

Domain Separation

Understanding Domain Separation

On-board a Company

Review the following considerations to on-board customers or companies in a domain separated instance, for Cloud Provisioning and Governance services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Domain separation for Cloud Provisioning and Governance overview

Domain separation for Cloud Provisioning and Governance is designed to give Service Providers a multi-tenant instance architecture that delivers offerings efficiently and securely to their clients. Strong universal process standards, data-driven process design, strict governance, and centralized administration maximize these benefits.

Customer On-boarding by Service Providers

- Create the MID Server user and grant the role.

  Domain Separation in Cloud Provisioning and Governance is achieved by installing and configuring one or more dedicated MID Servers for each company/ domain that you are managing in your instance. Specifying a MID Server for each domain/company, ensures that Config Items (CIs) are assigned to the intended domain. Each domain has one or more MID Servers setup for cloud discovery, provisioning, and other orchestration.

  Note: Create a MID Server in the global domain so that you can explore targets in other domains. For more information, see Set up domain separation for MID servers

  The Management, Instrumentation, and Discovery (MID) Server is a Java application that runs as a Windows service or UNIX daemon on a server in your local network. The MID Server enables communication and the movement of data between a ServiceNow® instance and external applications, data sources, and services.

  To communicate with the instance, MID Servers need a user ID and the appropriate role. Create the user ID for a MID Server and grant the ID the **mid_server** role. To prevent MID Server validation failure, the system runs business rules to monitor the settings selected for MID Server users.

- Setting up the MID Server.

  For more information on setting up the MID Server, see:

  - Installing the MID Server
  - Configuring MID Servers
- Resource Blocks in Cloud Provisioning and Governance

  Create Resource Blocks in global domain as there is a 1:1 relationship between CI and resource block. Do not create resource blocks in any of the child or leaf domains. Resource blocks act as a process record, so any template of a parent domain is visible and works in the child domain.

  Resource blocks act as a type of middleware component between catalog items, the Cloud API (CAPI), and the CMDB.

  - If the blocks in the base system do not provide the cloud resources that you need to provision, you can create a custom resource block.
  - In the Cloud Provisioning and Governance application, script execution is divided into cloud scripts and cloud script templates. Use scripts in catalog items, resource blocks, OS profiles, and use policy scripts to set request form attributes. Policy scripts cannot override user data. For more information, see Cloud scripts and cloud script templates.
  - Map a script to an OS profile

    To execute scripted actions during VM provisioning, you can map a script to an OS profile. The script runs on VMs that are created based on the image template in the OS profile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Define Resource Profile

  Create resource profile definitions in the global domain, and create mappings to these definitions in the child (leaf) domains. Do not delete resource profiles without checking with all impacted companies. You can create domain-specific resource pools and pools filters for all supported cloud providers. As a Domain admin you can see global, domain and custom resource pools specific to sub-domains. A resource pool is a query or script that filters a table. You configure a resource pool to limit the values that are available to users when they request a catalog item. For domain separated resource pools only records of the relevant domain are returned.

  Resource profiles are cloud provider-agnostic definitions that specify the allowed attribute values for a resource. Resource profiles enable you to control the choices that the user sees when requesting a cloud resource. As a result, you do not need to define a unique blueprint for each variation of the resource.

- Administer and set up Cloud Provisioning and Governance for Quota, Policies, Pools, Permissions. Governance refers to the limitations that you can set on available cloud resources. As an SP you could manage process separation in Cloud Provisioning and Governance policies, and permissions for each company or domain you are managing in your instance. Policies are recommended to be set up in global domain only. Process separation at a domain level is achievable for governance and policies. For customized policy enforcement, use company field in the policy rule condition to target a policy towards the particular domain.

Type	Description
Quotas	Quotas are limitations or requirements for groups and users on CIs or cloud resources. Use quotas to prevent wasteful resource usage in your cloud environment.
Policies	A cloud policy can override a property value set by a user, create an approval task, reserve an IP address, pre-populate or hide form fields, execute custom scripts, call the Cloud API, or start or abort subflows. A cloud policy gives you system-wide control over approvals, resource operations, blueprint operations, or catalog item settings.
Pools	A resource pool is a query or script that filters a table. You configure a resource pool to limit the values that are available to users when they request a catalog item.
Permissions	Permissions are user group-level access rights to features in the Cloud Provisioning and Governance application and to specific records in the instance, such as blueprints or cloud accounts.

  Permissions are user group-level access rights to features in the Cloud Provisioning and Governance application and to specific records in the instance, such as catalog items or cloud accounts. You could also set up basic governance policies, quotas, pools, and allow the cloud admin user to further customize them based on their needs. Create a user-group representing the domain users and associate permission to this user group, to set up permissions at a domain level. For more information on managing permissions for Cloud Provisioning and Governance roles, see Manage Cloud Permissions.

- Configure and manage policies and manage lease for provisioned stack. A policy rule is a collection of conditions and actions. If all conditions evaluate to true, the policy engine

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

performs the actions. If any condition evaluates to false, the policy engine does not perform the actions.

Set up policies in global domain only when you want the policies to apply to all domains. You can customize Cloud Provisioning and Governance and set policies specific to a customer or a domain, if you don't want global policies to apply to a domain. Use the domain field in policy rule condition to target policy towards the particular domain, to enable customized policy enforcement. When global and domain policies co-exist, both policies are executed in the sequence of the order configured for each policy. In case the values set in the order columns are same, the sequence in which the policy is executed is undetermined. Ensure that policy actions are domain separated, that is all action rules are within a single domain, and do not depend on rules in other domains.

For more information, see:

- Policy governance in Cloud Provisioning and Governance
- Configure a cloud policy rule
- Create an action for an 'on Lease end' policy
- Cloud policy example

- Create a cloud catalog item

Create a cloud catalog item for provisioning, based on a template and publish the catalog item to provide a service.

You can create common or global catalogs items based on templates, at the SP domain, making it available for all domains you're managing. You can also create custom catalog items for each domain/company based on your requirements. A user with domain admin/cloud designer roles, can update or delete cloud scripts and cloud script templates belongs to domain and sub-domains. You cannot use common catalogs with different pool-filter for each domain. Create a custom catalog for each domain instead.

Deployment, post-provisioning, and Day-2 operations are domain separated. For example, if a cloud portal user from a specific company orders a VM, the stack, request, virtual machine, and all components are stored in its respective domain. Subsequent operations for a domain are also handled in the same manner.

You can generate a catalog item based on a Terraform configuration management template. You can also use Azure Resource Manager and AWS CloudFormation templates.

- Create Cloud Templates

Do not add private information such as CIDR values, IP addresses, or credentials to the template content, if setting up the catalog item as a global catalog item. Other cloud admins can access this information.

Create a cloud template and associate the template with a catalog item. Once you've created a template, you can reuse the template to create additional catalog items for the services you want to provision.

- Set capacity limits on user requests for resources

Capacity limits place restrictions on the attributes of cloud resources such as the number of virtual machines, virtual CPUs, or aggregate storage. You can set limits on resources separately for each logical datacenter in a cloud account.

- Create a cloud quota definition

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can set up separate quota definitions for each domain you're managing in an instance.

You apply quotas to a resource, such as a virtual server, datacenter, or a generic cloud resource, and to users of a particular group. You can set a per-user limit, and a total limit for the group for the resource. A cloud quota definition specifies the limitations on a resource block. You can use templates or manually create different types of configurations for each quota definition.

- Using the IPAM integration

  IP address management (IPAM) in Cloud Provisioning and Governance supports domain separation.

  If you use an IP address management (IPAM) tool such as Infoblox, you can manage cloud IP addresses, networks, and subnets in your cloud catalog offerings.

- Business hours scheduling

  Set up business hour scheduling on all virtual machines in a stack. Business hour scheduling can be set up on stacks at the time of provisioning or on existing stacks. Schedules can be set for when a stack should be started, stopped, and/or deprovisioned by defining parameters such as business hours, holidays, long weekends.

  You map a schedule profile to an instance schedule. The schedule profile applies to all newly provisioned resources that use the profile. For example, a schedule profile can specify the days of the week and times of day when a stack should start and stop.

  Note: Global Business schedules apply to all domains. You cannot set up domain-specific schedules if an active global schedule exists.

- Create a schedule profile

  To set up domain-specific lease schedules ensure that no global lease schedules exist. You can only create a domain-specific lease or schedule when no active global policies exist. While you can have multiple leases or schedules in draft state, only one lease or schedule can be active.

Next Steps

For more information on provisioning cloud resources, and making Cloud Provisioning and Governance life-cycle operations available for each domain that you are managing in your instance for, see Domain admin considerations.

Domain admin considerations

Before configuring domain separation for customers, ensure that you review the following considerations ahead of provisioning cloud resources for each domain you are managing in the Cloud Provisioning and Governance application.

Domain administrator activities in a Service Provider's organization

The following section lists activities that a domain's administrator performs for the company that a Service Provider manages. Set up service accounts and cloud accounts using the Cloud Admin Portal. Create pre-provisioning operations that you can configure for cloud resources before your users provision them or perform life-cycle operations.

Strong universal process standards, data-driven process design, strict governance, and centralized administration maximize the benefits of domain separation in Cloud Provisioning

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

and Governance in a single instance. The domain admin role must strictly be restricted to users in the Service Provider's organization, and not be assigned to the Cloud Admin users from the customer's organization. This restriction enables the SP to ensure that a customer does not get full access to other domain's data. Since data is often shared across multiple clients, as domain admins, don't expose or provide permissions that could result in data leakage.

Provisioning Cloud resources

- Ensure that you map relevant service accounts to cloud accounts in each domain. For instance, discovering a primary subscription from Azure cloud creates one or more child service accounts and are created in the same domain. This means that all service accounts in a master subscription from Azure cloud must belong to a single company and domain alone.

  Log in using a user with a root_admin or cloud_admin role, while setting up cloud and service accounts and running discovery for any domain. Use the domain picker while performing actions such as discovery or cloud and service account creation and mapping.

  Note:
  - Cloud Provisioning and Governance does not support the expand and collapse domain scope capability.
  - For more information on what a user can and cannot access, see Domain Scope

  A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter. A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

  Ensure that management keys and service account credentials are unique to each domain and are not shared. To set up cloud accounts and service accounts for multiple cloud providers, perform Day 1 Setup actions:
  - Day 1 setup guide for Amazon Web Services on Cloud Provisioning and Governance
  - Day 1 setup guide for Azure on Cloud Provisioning and Governance
  - Day 1 setup guide for Google Cloud Connector on Cloud Provisioning and Governance
  - Day 1 setup guide for VMware on Cloud Provisioning and Governance
- Cloud API (CAPI) and Cloud scripts and cloud script templates

  CAPI is not domain separated in Cloud Provisioning and Governance. CAPI is set up at a global domain and shared across leaf domains. Ensure that scripts do not contain hard-coded sensitive information such as account details, credentials, or names even in comments or annotations.

  CAPI enables you to integrate Cloud Provisioning and Governance with cloud providers using REST APIs. Cloud scripts are simple java scripts that use platform features. In the Cloud Provisioning and Governance application, script execution is divided into cloud scripts and cloud script templates. Use scripts in templates, resource blocks, OS profiles, and use policy scripts to set request form attributes. Policy scripts cannot override user data. Cloud script templates are actual executables which are passed to target a virtual machine for execution. Create a cloud template first and then associate it with a cloud script.

- Cloud Discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Providers (SPs) use domain separation to segregate data for each customer. Users in a given domain have visibility only to the data in their own domains or in child domains. SPs typically control the top-level domain, which gives them visibility to data associated with all domains. Though domain separation support for Discovery is considered Level 2, there is no delegated administration to the child domains in Cloud Provisioning and Governance. The SPs must retain administrative control. As a SP, always run discovery from a leaf domain by logging in as or impersonating a domain administrator to discover your cloud resources.

Cloud Discovery provides a wizard that allows you to create and run cloud schedules in a single interface. When you create a schedule with the Discovery Manager, you select the accounts to discover and the credentials for accessing these accounts. You also select the MID Servers to scan the resources. You can then view the results in the Discovery Home page and track any errors that might have occurred.

For more information see,

- Discovery Manager
- Running discoveries in your network
  - Microsoft Azure Cloud Discovery
  - VMware Cloud Discovery
  - Google Cloud Platform Discovery
  - IBM Cloud Discovery

- Set up and configure event management to receive external events and generates alerts based on event and alert management rules. The visibility of events depends on the associated service account's domain. Only users belonging that domain can see the event details for processed events. Events that are not associated to a service account are visible to all domains.

  Monitor the health of business services and infrastructure using a single management console and respond appropriately to any issues that come up. Event Management provides intelligent event and alert analysis to ensure continuity of your business service performance. Event Management receives and processes events via the MID Server.

  For more information, see

  - Exploring Event Management
  - Event Management Setup
    - AWS events-driven discovery
    - Configure the Microsoft Azure Alert service to auto-update the CMDB
    - Configure the Google Cloud Logging service to auto-update the CMDB
    - Configure the VMware Events service to auto-update the CMDB

- The Cloud Provisioning and Governance application supports integration with continuous delivery solutions (also known as configuration management). Create an Ansible or Terraform configuration management provider and then run Discovery on the provider to find its resources. For more information, see Support for continuous delivery (configuration management) and Create a workload provider type for each new configuration management provider. This information appears in the order catalog form as management attributes that your users can select when provisioning a virtual resource through a configuration management provider.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you are creating catalogs based on terraform templates and sharing the catalog with multiple domains. Perform module listing (config discovery) in the global domain. The MID Server should be created in the global domain and should only be assigned with Terraform capability for its discovery. This allows SPs to share catalogs with multiple domains.

  ⚠ Warning:

  Don't create a global MID Server with any other capability other than configuration management for Terraform.

  Create a common terraform catalog and share it with multiple customers:

  - Create MID Server in global domain as a global admin.
  - Create a Terraform Open Source config provider.

    Note: Add only “Terraform” as the config provider.

  - Create a catalog item based on a Terraform template.

- Cloud Provisioning and Governance supports using ServiceNow AI Platform with subflows. Rules are collections of conditions and actions. If all conditions of a rule evaluate to true, the system performs the actions. If any condition evaluates to false, the system does not perform the actions. Creating rules helps you track activities and more quickly respond to and resolve issues. Take advantage of the flow designer subflow to automate your Day 2 operations. Quickly write a subflow that communicates with a Cloud API or a particular resource. Use SSH, PowerShell, or a similar tool, to access and then extend the subflow capabilities. For more information, see Day 2 operations using subflows.

- Budget-based notification and approval

  Domain Separation in Cloud Provisioning and Governance supports data separation of budgets. You can assign domain-specific budgets. Assign a budget for a group and a user within the group. When the user or group reaches the budget limit threshold, notifications are sent alerting them about it. For more information, see Configure budgets

  - Create Tags for cloud resources

    Tags categorize cloud resources to provide richer and more detailed tracking and billing report data. Tag keys are not domain separated and are shown to users of other domains. Tag visibility depends on the associated CI domain or associated billing record domain. Tags that are not associated to any CI or billing record are visible to all domain admins. When a tag is created, it shows up in any new catalog that you create. The cloud admin has to choose the tags they want for their catalog.

  - Billing data can be viewed separately for each domain. Billing jobs use the domain of service account configured. Configure the Billing setup to enable cloud admins and cloud users to view reports like cloud billing data and cloud tag usage on the billing dashboard.

    Define the scheduled job that regularly uses a MID Server to download billing data from the provider. Cloud Provisioning and Governance saves the data in a cost table and uses the information to generate reports.

    For more information on setting up billing schedules and downloading billing reports, see:

    - Define the schedule for downloading AWS billing data
    - Define the schedule for downloading Azure billing data

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Next Steps

Important:

While creating custom operations invoking workflows on resource blocks, ensure that you change the context to the appropriate domain.

To change the context to the right domain, edit the **Workflow Activity** to call the following script from within the workflow. Ensure that an appropriate user is obtained from the Service Account creator, or the order user. Invoking this impersonation script enables the workflow to be executed in the correct context and choose the right MID Server belonging to the associated domain.

```
//var orderContext =
 json.decode(workflow.inputs.ordercontext) ;
new
 CMPDomainSeparationUtil() .impersonateUser(current.request.
requested_for);
```

For more information on performing Cloud Provisioning and Governance life-cycle operations that are available for each domain that you are managing in your instance, see Cloud User Portal.

Related topics

Additional Cloud Provisioning and Governance setup on Day 2

Discovery patterns used by ITOM Visibility

Cloud Provisioning and Governance Recipes

Multi-cloud recipes provide ready content for typical cloud deployment and common operations scenarios across cloud platforms.

Recipes for Enterprise Workloads and Day-2 Operations

Cloud Provisioning and Governance recipes are ready-to-use catalog items built from sample CFT and ARM templates to support real-world cloud deployment use-cases. These recipes are ready for use in Proof of Concept as well as non-production environments. For production environments, ensure that the recipes are well tested and internally evaluated. You can avail the recipes, as update sets, to be downloaded from the SHARE in the developer portal under the Cloud Provisioning Recipes category.

List of Update Sets

Name	Description
Catalog Items for Amazon Web Service	
CFT Single Linux VM	Catalog Item for a single Linux Virtual Machine using CFT.
CFT Single Windows VM	Catalog Item for a single Windows Virtual Machine using CFT.
CFT Linux VM with Multiple NICs	Catalog Item for a Linux Virtual Machine with multiple network interfaces using CFT.
CFT Linux VM with LB and Storage	Catalog Item for a Linux Virtual Machine with load balancing and storage using CFT.
CFT Linux VM with Multi VM Multi Disk	Cloud Formation Template for a Linux VM with multiple VMs and multiple virtual disk interfaces.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
CFT Linux Install Apache	Catalog Item to install an Apache server on a Linux VM using CFT.
CFT Simple VM with Approval Policy	Catalog Item for a simple VM with an approval policy, using CFT.
CFT Windows VM Multi Disk and LB	Catalog Item for a Windows VM with multiple virtual disk interfaces and load balancing, using CFT.
Catalog Items for Azure	
ARM Windows VM with Multi NICs	Catalog Item for a Windows VM with multiple network interfaces using ARM template.
ARM Windows Single VM Multi Disk	Catalog Item for a single VM with multiple virtual disk interfaces, using ARM template.
ARM Windows Stack Multi Disk	Catalog Item for a Windows stack with multiple virtual disk interfaces, using ARM template.
ARM Windows Stack Multi NIC	Catalog Item for a Windows Stack VM with multiple network interfaces, using ARM template.
ARM Linux VM Multi NIC Multi Disk	Catalog Item for a Linux VM with multiple virtual disk and network interfaces, using ARM template.
ARM Linux VM Install Apache	Catalog Item to install an Apache server on a Linux VM, using ARM template.
ARM Linux VM Install MySQL	Catalog Item to install a MySQL server on a Linux VM, using ARM template.
ARM WebApp MySQL	Catalog Item to build a Web app with Azure database for MySQL, using ARM template.
ARM VNET Multi Subnet with Approval Policy	Catalog Item for multi subnet VNET with approval policy, using ARM template.
ARM VNIC Multi Subnet with Naming Policy	Catalog Item for multi subnet VNIC with a naming policy, using ARM template.

Day 2 Operations for AWS VM

An update set is provided for some of the important Day 2 operations on AWS cloud provider. Broadly, Day 2 operations are the type of operations that are executed on workloads after they are provisioned. Day 2 operations can include complex life-cycle processes or simple utilities. In this update set, you have access to the following set of operations:

Name	Description
AttachVolume	Attaches and exposes an Amazon (Elastic Block Store) EBS volume to the instance with the specified device name.
CreateSnapshot [AWS datacenter level discovery operation]	Creates a snapshot of your Amazon (Elastic Block Store) EBS volume.
DescribeSnapshots	Describes the specified Amazon (Elastic Block Store) EBS snapshots available to you or all EBS snapshots available to you.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
DetachVolume	Detaches an EBS volume from an instance.
ModifyInstanceAttribute [Only InstanceType]	Modify the specified attribute of the specified instance.
ModifyVolume [Only Size]	Modify parameters of an existing Amazon (Elastic Block Store) EBS volume.
Day 2 Operations for AWS VM – Set 1	Catalog Item for Day 2 operations for an AWS VM, Set 1.

Cloud accounts

A cloud account is the logical representation in Cloud Provisioning and Governance of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts — even service accounts from different providers. For each service account, you specify which datacenters to include in the cloud account.

From a cloud account, you can:

- Add LDCs. A logical datacenter (LDC) is a region-specific virtual cloud that is associated with a service account. The datacenter hosts your cloud resources.
- Run Discovery on LDCs to update the CMDB with configuration changes or life cycle changes (create/modify/terminate) for each resource in each logical datacenter (LDC) that is associated with the cloud account.
- Set capacity limits on cloud services like virtual machines, virtual CPUs, virtual networks, aggregate storage volume size, and others. Set limits to help ensure that cloud resources are provisioned at appropriate scales. Limits are especially important for capacity-constrained environments like private clouds.
- Publish a cloud account to enable business teams to deploy stacks.

Structure of a cloud account

An advantage of cloud accounts is that you can group specified service accounts with only the regions (datacenters) that you want to allow cloud resources in. For example, your AWS account could have more than a dozen datacenters globally. However, you might use only one or two regions. To conform with your security and compliance rules, you might not want any resources provisioned to additional regions around the world.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

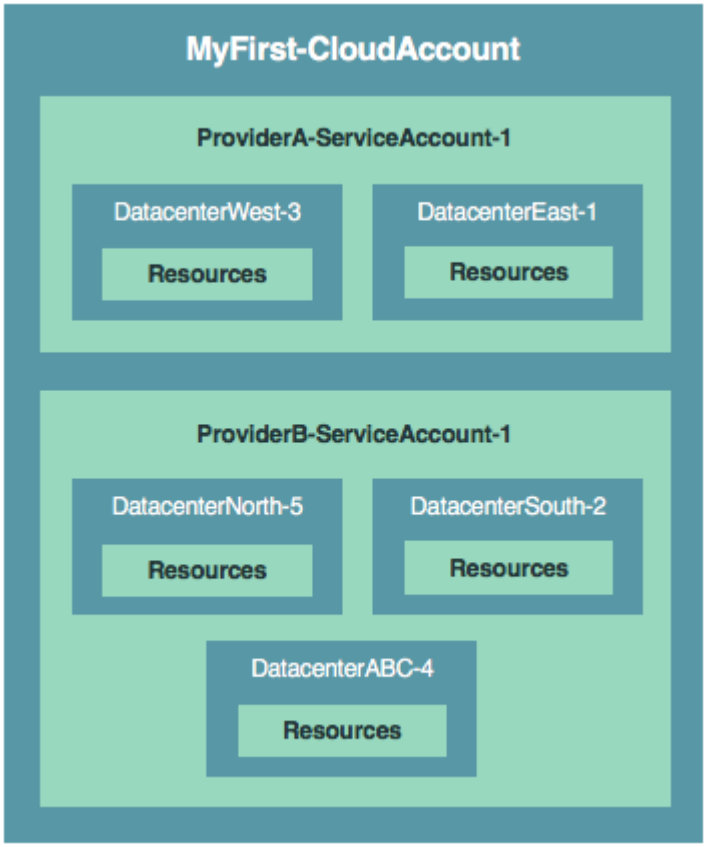

A cloud account can contain as many service accounts as necessary in the following scenarios:

Possible cloud account structures

Structure	Supported	Not supported
Add the same service account as many times as necessary, provided each entry is associated with a different LDC.	X	
Add the same LDC from different service accounts in one cloud account.		X
Add different LDCs from different service accounts to a particular cloud account, provided the regions are not the same. Example: You have us-west-1 in serviceaccount1 and us-west-2 in serviceaccount2. Both can be added to cloudaccount1. However, both us-west-1 in serviceaccount1 and us-west-1 in serviceaccount2 cannot be added to cloudaccount1.	X	

Update a cloud account

You keep a cloud account in the **Draft** state until you are ready to make the datacenters and capacity limits available for use — in blueprints and the Cloud User Portal, for example. You can change the state to **Published** after you run Discovery.

Set up an additional cloud account

See Set up an additional cloud account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add a datacenter to a cloud account

At any time, you can add a logical datacenter to the cloud infrastructure that is represented by a cloud account. See Add a datacenter to a cloud account.

Set capacity limits on user requests for resources

Capacity limits place restrictions on the attributes of cloud resources such as the number of virtual machines, virtual CPUs, or aggregate storage. You can set limits on resources separately for each logical datacenter in a cloud account.

Set capacity limits on user requests for resources

Service categories discovered in AWS

The Discovery process discovers AWS services.

Services categories discovered in AWS

Service Category	Name	CI Class	Pattern	Prod
Compute	AWS::CloudFormation::Stack	cmdb_ci_aws_cloudformation_stack		AWS Clou
Database	AWS::Datacenter	cmdb_ci_aws_datacenter		
Security	AWS::DirectoryService			AWS Serv
Compute	AWS::EC2::EIP	cmdb_ci_cloud_ip_address		AWS Com
Compute	AWS::EC2::Image	cmdb_ci_os_template		AWS Com
Compute	AWS::EC2::Instance	cmdb_ci_vm_instance		AWS Com
Compute	AWS::EC2::InstanceType	cmdb_ci_compute_template		AWS Com
Security	AWS::EC2::Keypair	cmdb_ci_cloud_key_pair		AWS Man Serv
Network	AWS::EC2::NetworkInterface	cmdb_ci_nic		AWS Priva
Compute	AWS::EC2::SecurityGroup	cmdb_ci_compute_security_group		AWS Com
Network	AWS::EC2::Subnet	cmdb_ci_cloud_subnet		AWS Priva
Storage	AWS::EC2::Volume	cmdb_ci_storage_volume		AWS Bloc
Network	AWS::EC2::VPC	cmdb_ci_network		AWS Priva
Load Balancer	AWS::ElasticLoadBalancing::LoadBalancer	cmdb_ci_cloud_load_balancer	Amazon AWS	AWS Load

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Category	Name	CI Class	Pattern	Prod
			Elastic Load Balancer Service	
Load Balancer	AWS::ElasticLoadBalancingV2::LoadBalancer	cmdb_ci_cloud_load_balancer	Amazon AWS Elastic Load Balancer Service	AWS Load
Database	AWS::RDS::DBEvent			AWS Data Serv
Database	AWS::RDS::DBInstance	cmdb_ci_cloud_database	Amazon AWS Relational Database Service	AWS Data Serv
	Availability zone	cmdb_ci_availability_zone		AWS Zone

Note: The full list of classes/resources matching can be found in the Cloud Resource Types [sn_capi_resource_type] table.

Service categories discovered in Azure

The Discovery process discovers Microsoft Azure services.

Services categories discovered in Azure

Note:

The CAPI-based discovery method is deprecated. Follow Migrate from CAPI to Patterns for more information.

Service Category	Name	CI Class	Pattern	Product
	Azure Datacenter	cmdb_ci_azure_datacenter		
	Azure Subscription	cmdb_ci_cloud_service_account		
	Cloud Resource	cmdb_ci_cmp_resource		
Application Services	Microsoft (Unknown Resource Type)			Microsoft (Unknown Product)
Application Services	Microsoft.ApiManagement/service			Microsoft.ApiManagem
Application Services	Microsoft.Cache			Microsoft.Cache

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Category	Name	CI Class	Pattern	Product
Application Services	Microsoft.Web/ serverFarms			Microsoft.Web
Application Services	Microsoft.Web/sites	cmdb_ci_cloud_webserver		Microsoft.Web
Compute	Microsoft.Compute/ availabilitySets	cmdb_ci_availability_zone		Microsoft.Compute
Compute	Microsoft.Compute/images	cmdb_ci_os_template		Microsoft.Compute
Compute	Microsoft.Compute/ virtualMachines	cmdb_ci_vm_instance		Microsoft.Compute
Compute	Microsoft.Compute/ virtualMachineScaleSets			Microsoft.Compute
Compute	Microsoft.ServiceFabric/ clusters			Microsoft.Compute
Compute	Microsoft.Compute/vmSize	cmdb_ci_compute_template		Microsoft.Compute
Compute	Microsoft.Resources/ deployments	cmdb_ci_azure_deployment		Microsoft.Compute
Compute	Microsoft.Resources/ resourceGroups	cmdb_ci_resource_group		Microsoft.Compute
Database	Microsoft.DocumentDB			Microsoft.DocumentDB
Database	Microsoft.Sql/servers	cmdb_ci_cloud_database	Azure Database	Microsoft.Sql
Database	microsoft.sql/ managedInstances	cmdb_ci_cloud_database	Azure Database	Microsoft.Sql
Database	microsoft.dbformysql	cmdb_ci_cloud_database	Azure Database	Microsoft.Sql
Database	microsoft.documentdb/ databaseaccounts	cmdb_ci_cloud_database	Azure Database	Microsoft.Sql
Database	microsoft.cache/redis	cmdb_ci_cloud_database	Azure Database	Microsoft.Sql
Database	microsoft.dbforpostgresql/ servers	cmdb_ci_cloud_database	Azure Database	Microsoft.Sql
Database	microsoft.dbformysql/ servers	cmdb_ci_cloud_database	Azure Database	Microsoft.Sql
Database	Microsoft.Sql/servers/ databases	cmdb_ci_database	Azure Database	Microsoft.Sql
Load Balancer	Microsoft.Network/ loadBalancers	cmdb_ci_cloud_load_balancer		Microsoft.LoadBalance
Network	Microsoft.Network/ applicationGateways			Microsoft.Network
Network	Microsoft.Network/ networkInterfaces	cmdb_ci_nic		Microsoft.Network

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Category	Name	CI Class	Pattern	Product
Network	Microsoft.Network/ networkSecurityGroups	cmdb_ci_compute_security_group		Microsoft.Network
Network	Microsoft.Network/ publicIPAddresses	cmdb_ci_cloud_public_ipaddress		Microsoft.Network
Network	Microsoft.Network/ routeTables	cmdb_ci_route_table		Microsoft.Network
Network	Microsoft.Network/subnets	cmdb_ci_cloud_subnet		Microsoft.Network
Network	Microsoft.Network/ virtualNetworks	cmdb_ci_network		Microsoft.Network
Security	Microsoft.Security			Microsoft.Security
Storage	Microsoft.Storage/blobs	cmdb_ci_storage_volume		Microsoft.Storage
Storage	Microsoft.Storage/ storageAccounts	cmdb_ci_cloud_storage_account		Microsoft.Storage
Tools	Microsoft.Insights/ components			Microsoft.Insights
Tools	Microsoft.ServiceBus			Microsoft.ServiceBus

Cloud Admin Portal

The Cloud Admin Portal is a role-based portal. You can manage, design, govern, operate, and analyze all your cloud resources from a unified base.

Depending on the role you are logged in with, the Cloud Admin Portal displays different metrics. As an example, if you are logged in to the Cloud Provisioning and Governance application as a Cloud Service Designer, the Cloud Admin Portal displays only the **Design** metrics.

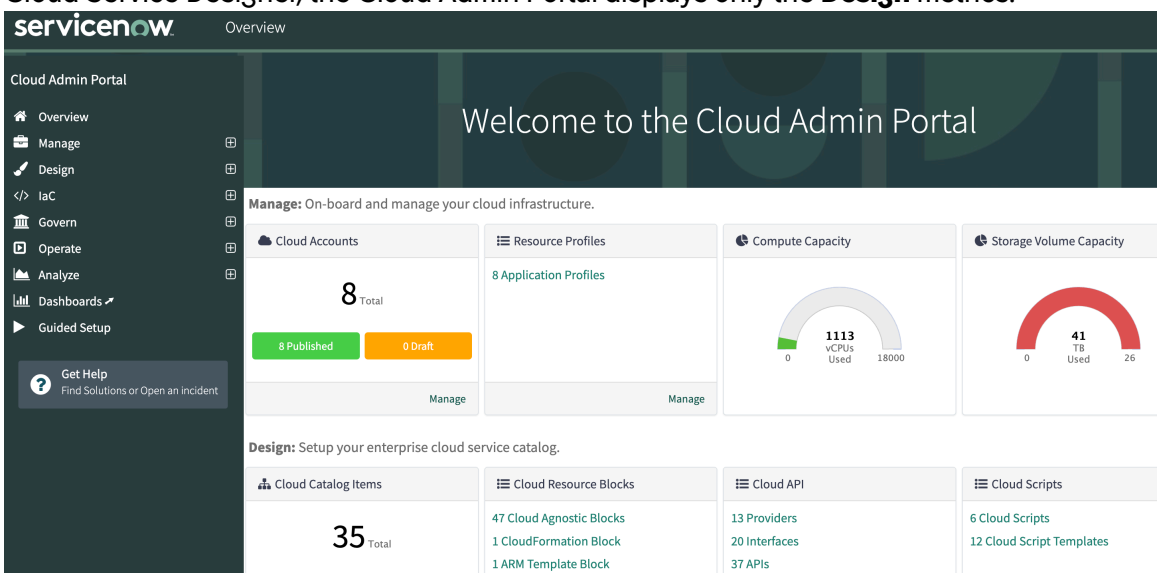

Starting this release, you can now use the **Dashboards** link on the left navigation pane of Cloud Admin Portal to access these Next Experience dashboard pages:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

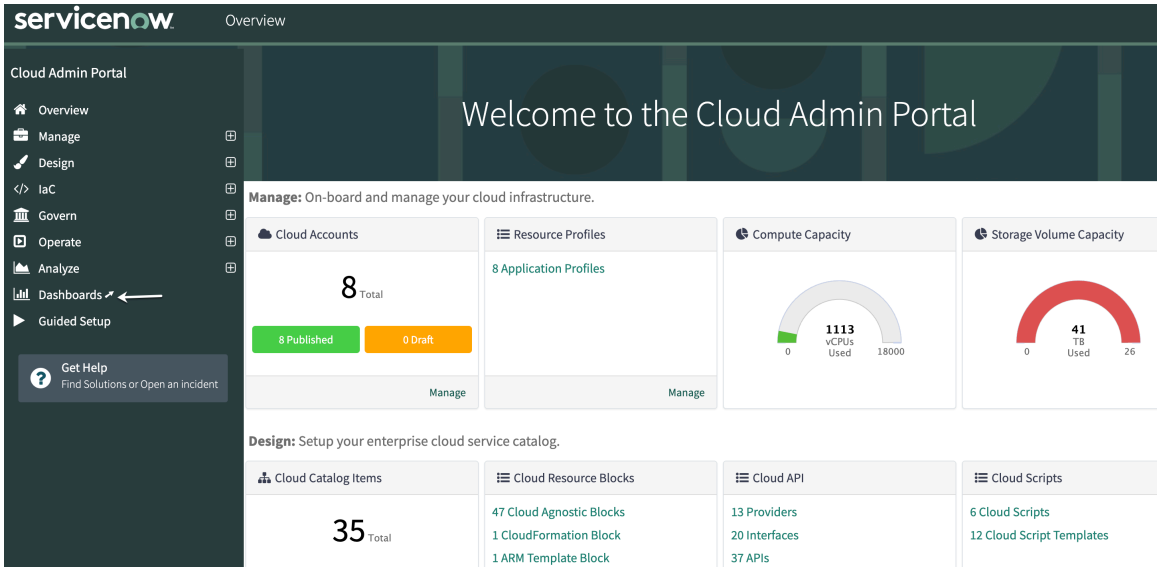

Find the link to these dashboards on a new tab based on your user role:

- Cloud root cause analysis dashboard
- Cloud operations dashboard
- Cloud tag values dashboard

Follow Request the Cloud Provisioning and Governance application to learn more about installing Cloud Provisioning and Governance to access the admin and user pages. Log in to the portal with the following Cloud Provisioning and Governance roles to view or manage the associated metrics:

Descriptive name and role name	Access to the metrics in Cloud Admin Portal
Root administrator `[sn_cmp.cmp_root_admin]`	All the metrics. • Overview • Manage • Design • IaC • Govern • Operate • Analyze • Guided Setup
Cloud administrator `[sn_cmp.cloud_admin]`	Managethe metrics. • Overview • Manage
System administrator `[admin]`	All the metrics. • Overview • Manage • Design

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Descriptive name and role name	Access to the metrics in Cloud Admin Portal
	• IaC • Govern • Operate • Analyze • Guided Setup

Governor
`[sn_cmp.cloud_governor]`
`[sn_capi.cloud_developer]`
	• Govern
	• Analyze
	• Design

Service Designer `[sn.cmp.cloud_service_designer]`	Design
Operator `[sn_cmp.cloud_operator]`	Operate
CAPI Developer `sn_capi.cloud_developer`	Design

Only the Root administrator, the Cloud administrator, and the System administrator roles can view and manage all the metrics in the Cloud Admin Portal.

Click **Guided Setup** to configure the Cloud Provisioning and Governance application on your ServiceNow instance. You are guided through completing various configuration tasks. The percentage of the tasks that you have completed is displayed for you to monitor your progress.

Click **Get Help** to ask a question, open an incident, go to the product documentation or the Knowledge base website.

Resource Profiles

Resource profiles are cloud provider-agnostic definitions that specify the allowed attribute values for a resource. Resource profiles enable you to control the choices that the user sees when requesting a cloud resource. As a result, you do not need to define a unique blueprint for each variation of the resource.

Note: In addition to the cloud-agnostic profiles, you can add provider public images to the Cloud Provisioning and Governance catalog. See Add an AWS public image to Cloud Provisioning and Governance and Add an Azure public image to Cloud Provisioning and Governance.

Example: Compute profile

You might create a compute profile named *`Micro`* and map it to a "micro" hardware type, with specified values for memory and CPU. Later, when a user selects a compute resource while requesting a virtual server, **Micro** appears as an option in the list of available compute resources. The user does not need to specify values for memory and CPU — the values come from the profile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Compute profile on the request form

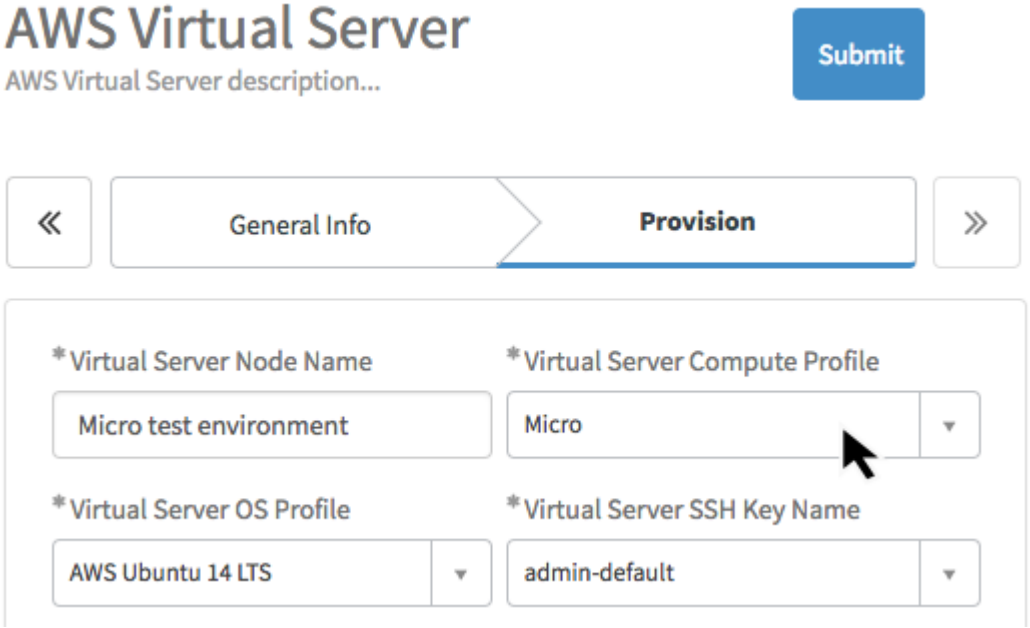

Resource Profile mappings

Resource profiles create mapping associations between the following items:

- A cloud account.
- A logical datacenter in the cloud account.
- A specific resource type in the CMDB that provides the attributes.
- Optional: A pricing value that appears to users when they request a resource that uses the resource profile.

By default for AWS and Azure, Cloud Provisioning and Governance maps profiles to templates after Discovery runs. For other providers, such as Google cloud, you must manually associate the profile with the correct template and datacenter.

Resource types

Each resource profile has associated attributes. The attributes depend on the resource types that are populated after you run Discovery on a cloud account.

Resource profile type	Description and attributes	Resource type and template
Application profile	An application profile specifies application software to install on newly-provisioned resources. Users can select applications when they request a stack. Use application profiles when you integrate with configuration management (continuous delivery) providers such as Ansible playbooks. Attributes: Applications running on virtual machines.	Application Template [sn_cmp_application_template]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource profile type	Description and attributes	Resource type and template
Compute profile	A compute profile specifies the hardware to use for newly-provisioned virtual machines. Attributes: The size of computing resources, including the virtual CPUs, memory, and local storage.	Hardware Type [cmdb_ci_compute_template]
Compute Security Group profile	A compute security group profile applies specified security rules to newly-provisioned resources. You map a compute security group profile to a cloud account, a datacenter, a Compute Security Group template, and security rules for the template. Attributes: Firewall rules, such as enabling HTTP and HTTPS.	Compute Security Group Template [cmdb_ci_security_grp_template]
OS profile	An OS profile installs a specified image on a newly-provisioned virtual machine. You map an OS profile to a cloud account, a location (datacenter), an image template, and a cloud script. Attributes: Operating system images, including the OS type and version, the root device type, and the image source.	Image [cmdb_ci_os_template]
Schedule profile	You map a schedule profile to an instance schedule. The schedule profile applies to all newly-provisioned resources that use the profile. For example, a schedule profile can specify the days of the week and times of day when a stack should start and stop. Attributes: Schedule attributes such as when a stack should be started, stopped, or deprovisioned.	N/A

Related topics

Discover all datacenters in a service account on-demand

Create an application profile

An application profile specifies application software to install on newly-provisioned resources. Users can select applications when they request a stack. Use application profiles when you integrate with configuration management (continuous delivery) providers such as Ansible playbooks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- You must have a cloud account with datacenters. You must run Discovery on the service accounts to populate the datacenters.
- Role required: sn_cmp.cloud_admin

About this task

This example shows an application profile mapping for a Tomcat server on Ansible.

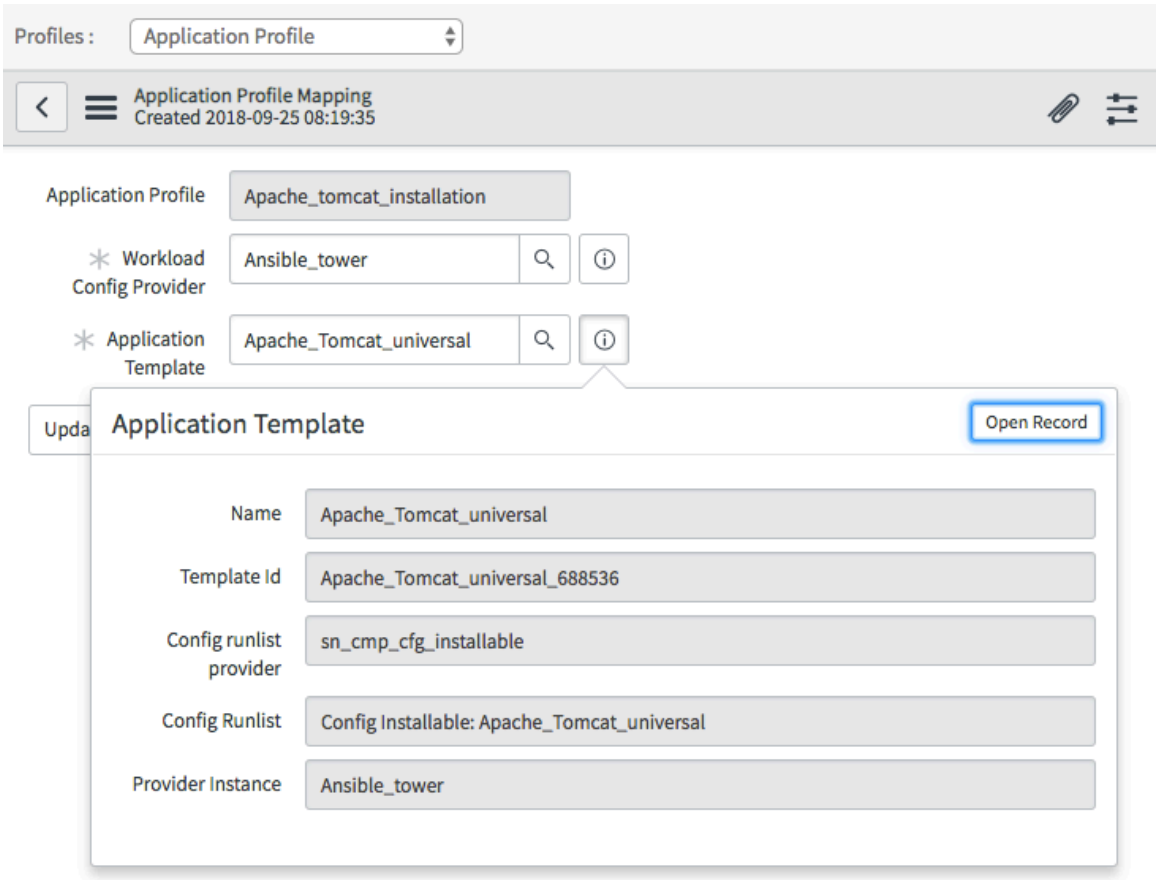

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Resource Profiles**.
2. In the **Profiles** list, select **Application Profile** and then click **New**.

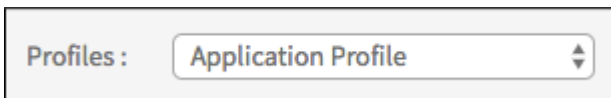

3. Enter a unique and descriptive **Name** and **Description** for the profile and then click **Submit**. The profile is created.
4. Map the profile to a template.
   a. In the list, click the profile that you created.
   b. In the **Application Profile Mappings** related list, click **New**, fill in the form, and then click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Workload Config Provider	Select a configuration management provider to associate with the profile.
Application Template [sn_cmp_application_template]	Select a template that the profile should be mapped to. If you ran Discovery on the Configuration Provider, the templates are already populated. The resource type associated with an application profile is `sn_cmp_application_template`. Click the reference icon ( ) to view the details of the template. To create a new template: **i.** Click the reference icon ( ) to open the Application Templates list. **ii.** Click **New** and then fill in the Application Template form: ▪ **Name**: Enter a descriptive name for the template. ▪ **Template ID**: Enter an ID to use for the template. ▪ **Config runlist provider**: Click the list, select a table from the **Table name** list, and then select records for the table from the **Document** list. ▪ **Config runlist**: Specify a configuration runlist. ▪ **Provider instance**: Select a provider instance.

5. Create as many mappings as needed.

Related topics

Discover all datacenters in a service account on-demand

Create a compute profile

A compute profile specifies the hardware to use for newly-provisioned virtual machines. A compute profile maps to a cloud account, a datacenter, and a hardware template.

Before you begin

Role required: sn_cmp.cloud_admin

You must have a cloud account with datacenters. You must run Discovery on the service accounts to populate the datacenters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Resource Profiles**.
2. In the **Profiles** list, select **Compute Profile** and then click **New**.
3. Enter a unique and descriptive **Name** and **Description** for the profile and then click **Submit**. The profile is created.
4. Click the resource profile that you created.
5. Map the profile to a template.

   a. In the list, click the profile that you created.

   b. In the **Profile Mappings** related list, click **New**, fill in the form, and then click **Submit**.

Example compute profile

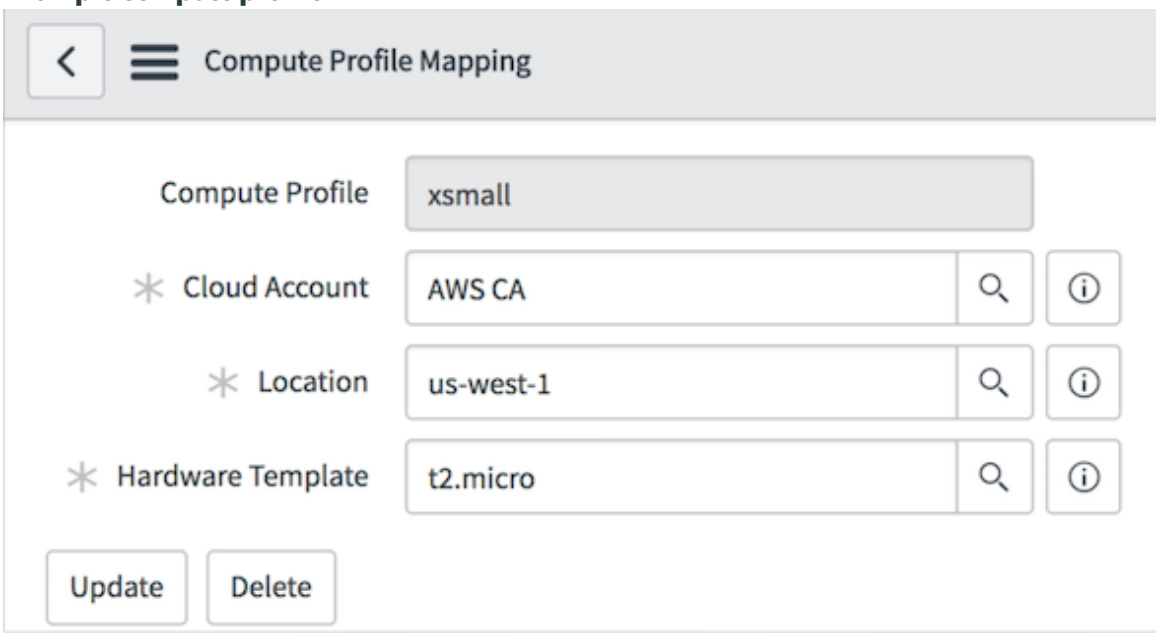

Field	Description
Cloud Account	Select a cloud account that the profile is valid for.
Location	Select the datacenter that the profile is valid for.
Hardware Template [cmdb_ci_compute_template]	Select the hardware type that the profile should be mapped to. Click the reference icon ( ) to view the details of the template.

Related topics

Discover all datacenters in a service account on-demand

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a compute security group profile

A compute security group profile applies specified security rules to newly-provisioned resources. You map a compute security group profile to a cloud account, a datacenter, a Compute Security Group template, and security rules for the template.

Before you begin

You must have a cloud account with datacenters. You must run Discovery on the service accounts to populate the datacenters.

Role required: sn_cmp.cloud_admin

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Resource Profiles**.
2. In the **Profiles** list, select **Compute Security Group Profile** and then click **New**.
3. Enter a unique and descriptive **Name** and **Description** for the profile and then click **Submit**. The profile is created.
4. Map the profile to a template.

   a. In the list, click the profile that you created.

   b. In the **Compute Security Group Profile Mappings** related list, click **New**, fill in the form, and then click **Submit**.

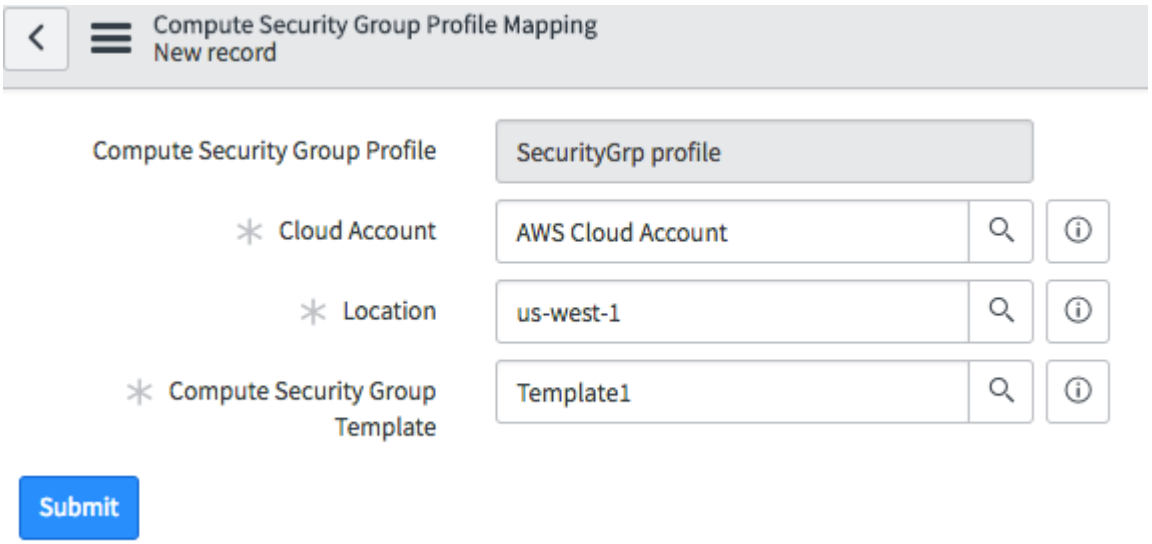

Field	Description
Cloud Account	Select a cloud account for the profile.
Location	Select the datacenter that belongs to the cloud account.
Compute Security Group Template [cmdb_ci_security_grp_template]	Select or create a template that the profile should be mapped to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	To create a new template, click the list icon and then click **New**. Enter a name and a template ID and click **Submit**. Click the reference icon ( ) to view the details of the template.

5. Add rules to the template.
   You can create more than one rule per template.

   a. In the **Security Group Profile Mappings** related list, under the **Compute Security Group Template** heading, click the template name.

   b. In the **Compute Security Group Rule Template** related list, click **New**, fill in the form, and then click **Submit**.

Field	Description
CIDR Range	Specify a CIDR range. For example, 10.0.0.0/24.
IP Protocol	Select an IP protocol. In Azure, the protocols supported are TDP and UDP. ICMP is not supported.
From Port	Specify the source port number.
Is Outbound	Enter `0` for inbound and `1` for outbound.
Name	Specify a name for the rule.
Network Type	Select the type of network.
To Port	Specify the destination port number.

Related topics

Discover all datacenters in a service account on-demand

Create an OS profile

An OS profile installs a specified image on a newly-provisioned virtual machine. You map an OS profile to a cloud account, a location (datacenter), an image template, and a cloud script. OS profiles are provider-agnostic and you can use the same profile for multiple cloud accounts.

Before you begin

- You must have a cloud account with datacenters. You must run Discovery on the service accounts to populate the datacenters.
- Role required: sn_cmp.cloud_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Note: Profile mappings that specify more details run first. For example, a mapping that specifies a blueprint, OS profile, and resource alias takes precedence over a mapping that specifies only an OS profile.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Resource Profiles**.
2. In the **Profiles** list, select **OS Profile** and then click **New**.
3. Enter a unique and descriptive **Name** and **Description** for the profile and then click **Submit**. The profile is created.
4. Map the profile to a template.

   a. In the list, click the profile that you created.

   b. In the **OS Profile Mappings** related list, click **New**, fill in the form, and then click **Submit**.

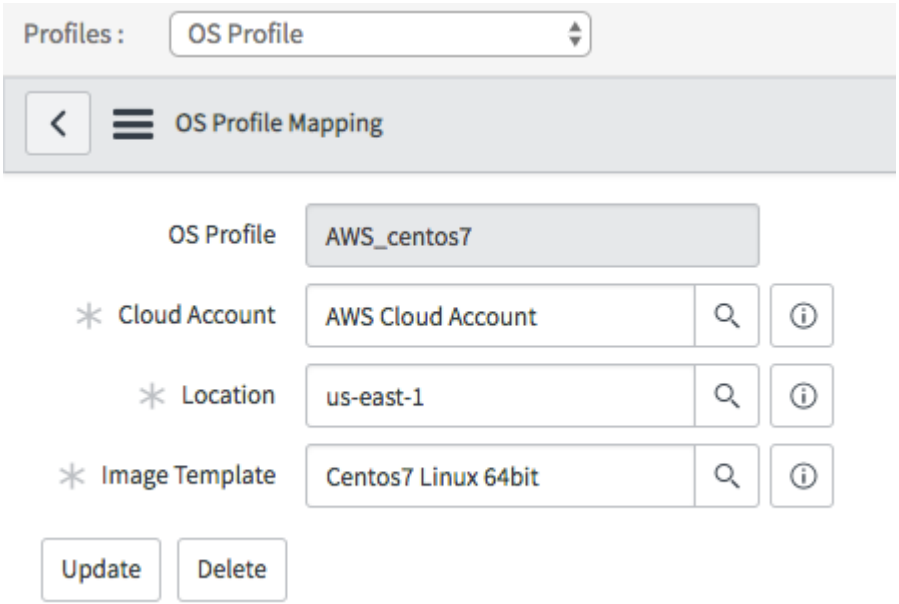

Field	Description
Cloud Account	Select the cloud account to map the OS profile to.
Location	Select a datacenter in the cloud account.
Image Template [cmdb_ci_os_template]	Select an image template that the profile should be mapped to. Click the reference icon ( ) to view the details of the template. **Important:** Do not create a new image template. Image templates must be discovered from the Cloud Accounts page on the Cloud Admin Portal (**Manage** > **Cloud Accounts**).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	When you add credentials to an image template, the credentials are inherited by all VMs that are provisioned using the template. See Add credentials to an image template.

5. Map the profile to a cloud script. In this procedure, you specify an existing script. See Create cloud initialization script templates and a script.

   a. On the Cloud Script OS Profile Mappings related list, click **New** and then fill in the form.

Field	Description
Blueprint [Optional]	Select a blueprint to limit the script to a specific resource block used in the blueprint.
Cloud script	Select a cloud script to map the OS profile to.
Active	Select the check box if the cloud script should be run after the virtual machine is provisioned.
Application	Cloud Provisioning and Governance is auto-selected.
OS profile	If you specify a blueprint, the cloud script is run when the blueprint is provisioned.
Resource Alias [Optional]	If you specify a resource alias for the blueprint, then the cloud script is executed when the blueprint with the specified resource alias is provisioned.

   b. Right-click the form header and select **Save**.
   The **Name** attribute `[scriptName]` appears in the OS Profile Mapping Overrides list.

   c. In the OS Profile Mapping Overrides list, enter an attribute name and value to use when the resource is provisioned.

Create a schedule profile

You map a schedule profile to an instance schedule. The schedule profile applies to all newly-provisioned resources that use the profile. For example, a schedule profile can specify the days of the week and times of day when a stack should start and stop.

Before you begin

Role required: sn_cmp.cloud_admin

You must have a cloud account with datacenters. You must run Discovery on the service accounts to populate the datacenters.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Resource Profiles**.
2. In the **Profiles** list, select **Schedule Profile**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click **New**, enter a unique and descriptive **Name** and **Description**, fill in the form, and then click **Submit**.

Field	Description
Schedule	Select an existing schedule or define a custom schedule. A schedule is a rule that includes or excludes specified time periods for actions. To define a schedule, use the **System Scheduler** module. See Define a schedule.
Multiplier	Auto-populated value between 0 and 1: The fraction of time that a stack with this schedule is expected to be powered on.

Related topics

Discover all datacenters in a service account on-demand

Add credentials to an image template

When you add credentials to an image template, the credentials are inherited by all VMs that are provisioned using the template.

Before you begin

- Role required: cloud_admin
- A set of basic authentication credentials must exist.

Procedure

1. Perform one of the following actions:
   - If an OS profile is mapped to the virtual machine template:
     a. Open the OS profile (**Manage** > **Resource Profiles**).
     b. In the OS Profile Mappings related list, click the **Name** of the template in the **Image Template** column.
   - For VMware templates, navigate to **All** > **Configuration** > **VMware** > **Virtual Machine Templates**.
2. On the Image form, select the basic authentication **Credentials**.
3. Click **Update**.

Add an AWS public image to Cloud Provisioning and Governance

You can make any AWS image selectable in an OS profile.

Before you begin

Role required: sn_cmp.cloud_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the AWS console, select the EC2 service and then click **IMAGES** > **AMIs**.

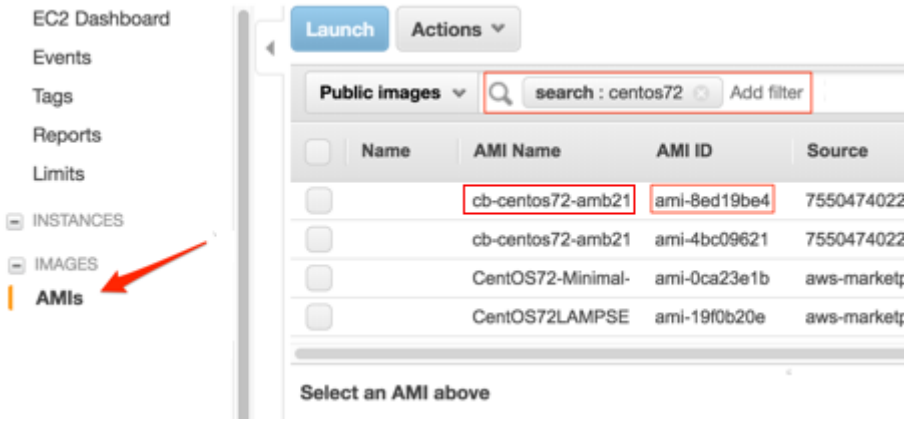

2. Filter the list of images and then take note of the **AMI Name** and **AMI ID** values for the image of interest.
3. On the instance, enter `cmdb_ci_os_template.list` in the **Filter navigator** text box to open the Images [cmdb_ci_os_template] table and then click **New**.
4. On the **Image** form, enter the following values:

Field	Value
Object ID	**AMI ID** value from AWS.
Name	Meaningful name that is based on the **AMI Name** value from AWS.
Guest OS	Operating system.

5. On the **Credentials** field, click the magnifying glass and then click **New** on the Credentials page.
6. In the list, click **Basic Auth Credentials**.
7. On the Basic Auth Credentials form, enter a meaningful **Name** for the credential.
8. Enter a **User name** to use for instances that are provisioned from this image.

   AWS uses SSH keys, so a password is not required.

   Note: If the name `ec2-user` or `root` does not work, check with your AMI provider.

Username guidelines

Type	Name
Amazon Linux	`ec2-user`
RHEL	`ec2-user` or `root`
Ubuntu	`ubuntu` or `root`
Centos	`centos`
Fedora	`ec2-user`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Name
SUSE	`ec2-user` or `root`

9. Click **Submit** to close the Basic Auth Credentials form.
10. Right-click in the header and select **Save** to save the current configuration of the new image record.
    In the Related Items list, you can see that the image is associated with the specified datacenter.
11. Click **Update** to make the public image selectable in an OS profile.
    This step also creates a relationship between the public image and the logical datacenter (LDC) in the CMDB.

Add an Azure public image to Cloud Provisioning and Governance

You can make any Azure image selectable in an OS profile.

Before you begin

Role required: sn_cmp.cloud_admin

Procedure

1. In the Azure command-line interface (CLI), make note of the values of the **Offer**, **Publisher**, **Sku**, and **Version**.

   On the command line, you can search by offer or by publisher or by both, as shown in the example. Each line item is an image of a particular version of CentOS from OpenLogic.

   You can use the following sample command on the CLI - `az vm image list -f Windows --output table` to fetch the values.

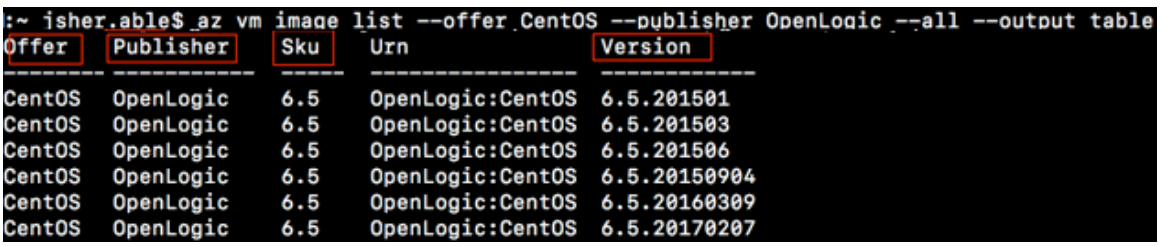

2. On the instance, enter `cmdb_ci_os_template.list` in the **Filter navigator** text box to open the Images [cmdb_ci_os_template] table and then click **New**.
3. Configure the form layout to show **Offer**, **Serial number**, **Version**, and **Vendor**.
4. On the Image form, enter the following values:

Field	Value
Object ID	Enter `/publisher/offer/sku/version`. You can get these values from the Azure CLI. As an example, `/OpenLogic/CentOS/6.5/6.5.201501`.
Name	Meaningful name that represents the version of the OS. `CentOS 6.5` in the example.
Guest OS	Operating system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Offer	**Offer** value from Azure.
Serial Number	**Sku** value from Azure.
Version	**Version** value from Azure.
Vendor	**a.** Click the magnifying glass and then click **New**. **b.** In the **Name** field, enter the **Publisher** value from Azure. **c.** Click **Submit**.

5. On the **Credentials** field, click the magnifying glass and then click **New** on the Credentials page.
6. In the list, click **Basic Auth Credentials**.
7. On the Basic Auth Credentials form, enter a meaningful **Name** for the credential.
8. Enter a **User name** and **Password** to use for instances that are provisioned from this image.
9. Click **Submit** to close the Basic Auth Credentials form.
10. Right-click in the header and select **Save** to save the current configuration of the new image record.
    In the Related Items list, you can see that the image is associated with the specified datacenter.
11. Click **Update** to make the public image selectable in an OS profile.
    This step also creates a relationship between the public image and the logical datacenter (LDC) in the CMDB.

IPAM integration

Use an IP address management (IPAM) tool like Infoblox to manage cloud IP addresses, networks, and subnets within your cloud catalog offerings. However, it's important to note that this functionality is not supported with our template-based cloud catalogs.

Infoblox integration

By default, Cloud Provisioning and Governance supports Infoblox integration. You can create Infoblox credentials and then set up an Infoblox instance record. See Register an Infoblox server for instructions.

Note: IPAM integration is not supported with template based cloud catalog items.

Note: This integration is supported with legacy blueprints only. Refer to this Knowledge Base article to learn more.

Subnets

If you have subnets in your AWS and Azure clouds, Cloud Provisioning and Governance finds them when Discovery runs on a service or cloud account.

Note: If you have networks and subnets in a VMware vSphere platform, Discovery does not create subnet records in the CMDB and cannot map the subnet to the appropriate network. You must create the records manually.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Register an Infoblox server

To integrate with Infoblox, you must create an Infoblox record and associate it with the appropriate credentials to register your Infoblox server with the instance.

Before you begin

- Role required: sn_cmp.cloud_admin
- Create Infoblox credentials.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Networks & IPAM** > **IPAM Providers** > **Infoblox**.
2. Click **New**, enter a unique and descriptive **Name**, fill in the form, and then click **Submit**.

Field	Description
URL	Enter the URL of the Infoblox server.
Credential	Select the Infoblox credential.

Create an IP pool

Create an IP pool to associate a cloud subnet with an IPAM provider instance.

Before you begin

- Role required: sn_cmp.cloud_admin
- Create Infoblox credentials and register an Infoblox server.
- Run Discovery on the cloud account to find subnets or manually create cloud subnets.

Note: This functionality is not supported with our template-based cloud catalogs.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Networks & IPAM**.
2. On the **IP Pools** tab, click **New**, enter a unique and descriptive **Name**, and then fill in the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

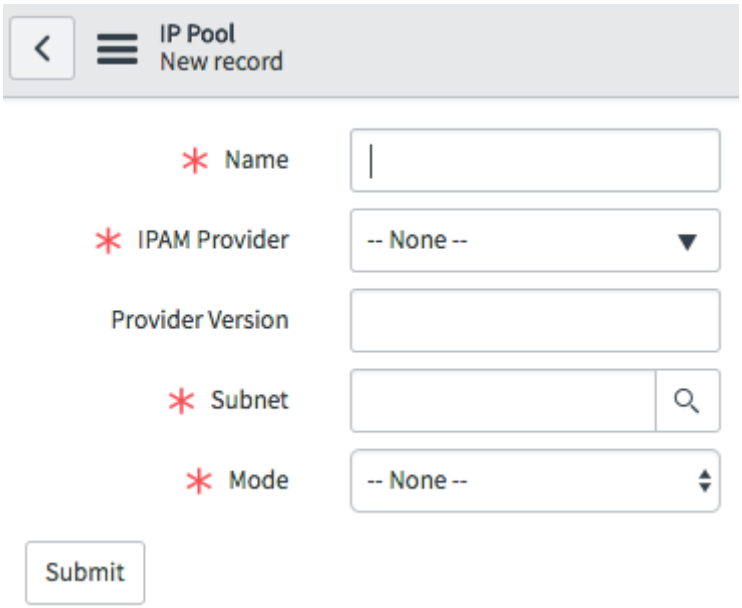

Field	Description
IPAM Provider	Select an IPAM provider.
Provider Instance	Select an instance for the provider.
Provider Version	Enter a version for the provider.
Subnet	Select a subnet. For AWS or Azure, select a discovered subnet that has been configured in AWS or Azure. For VMware, copy the CIDR range from the IPAM tool and paste the CIDR range in this field. See Create a cloud subnet.
Mode	Select a mode based on the provider that you are creating this IP pool for. The model varies based on the provider.

3. Click **Submit**.

Result
The related subnet appears in the **Subnet** related list at the bottom of the form.

Create a cloud subnet

If you have subnets in your VMware vSphere account, create a cloud subnet record in the instance.

Before you begin
Role required: sn_cmp.cloud_admin

Procedure
1. In the Cloud Admin Portal, navigate to **Manage** > **Networks & IPAM**.
2. On the **IP Pools** tab, click **New** to create a new record.
3. In the **Subnet** field, click the lookup icon.
   The Cloud Subnets list appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Click **New**.
5. Fill out the form fields (see table).

Field	Description
Name	Provide a descriptive name.
Status	Leave the status value as **Installed**.
CIDR/ Gateway/ Subnet Mask	Enter the CIDR associated with the subnet, the IP address of the gateway, and the subnet mask. For VMware, the CIDR, Gateway, and Subnet Mask fields are mandatory.
Primary and Secondary DNS	Enter the DNS values. For VMware, the Primary DNS field is mandatory.

6. Click **Submit**.

Create a relationship between a VMware network and subnet

If you discover VMware networks and subnets, you must manually create a relationship between the two.

Before you begin
Role required: sn_cmp.cloud_admin

About this task
If no relationship exists and a user must select a virtual network and subnet in the Cloud User Portal, the user can encounter an error.

Procedure
1. In the Cloud Admin Portal, navigate to **Manage** > **Service Accounts**.
2. Open the datacenter that contains the VMware cloud network.

   Note: VMware cloud networks are saved in the VMware vCenter Network [cmdb_ci_vcenter_network] table. AWS and Azure cloud networks are saved in the Cloud Networks `[cmdb_ci_network]` table.

3. Under **Related items**, click the plus icon to add a new relationship.
4. Make sure the **Use suggested relationships** check box is selected, and then select the **Contains (Parent)** relationship.
5. Under **Filter**, create the filter necessary to find the subnet.
6. Under **Configuration Items**, select the check box next to the subnet.
7. In the **Relationships** section, click the plus icon.
   The relationship is added to the list.
8. Click **Save and Exit**.

Reserve IP addresses for VMware vSphere VMs in InfoBlox

Create a policy to reserve IP addresses for VMware vSphere virtual machines in Infoblox, at the time of provisioning the virtual machines.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_cmp.cloud_admin.

About this task
You can create a policy that is invoked when a VMware virtual machine is provisioned to get an IP address from Infoblox. Before the virtual machine is provisioned, the instance makes a call to Infoblox that creates a host record with the IP address.

Procedure
1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Click **New** and then fill in the form.

Field	Description
Policy Name	Enter a descriptive name.
Policy Trigger	Select **on Resource operation**.
Resource Block	Select **ResourceBlock Virtual Server**.
Operation	Select **Virtual Server interface provision**.
Moment	Select **Pre Operation**.

3. Click **Submit** to publish the policy.
4. Create a policy rule for the policy.
5. Specify a condition in the condition builder for locating the vSphere virtual machine.
6. Select **Request from** > **starts with** and enter `vSphere`.
7. Click **Update**.
8. In the **Policy Rules Action** related list, click **New** to create a policy action.
9. Select **IP Address Management** in the **Create Policy Rule Action** dialog box.
10. Enter a name in the **Action Name** field, select **Infoblox** in the **Provider Type** list, and select **Reserve IP address** in the **IPAM Method Name** list.
11. Click **Submit**.
12. In the **Policy Rule Action Attributes** section, enter the domain name of your network in the **DNSSuffix** field.
    As an example, `corp.servicenow.com`. Values for all the other properties are already populated.
13. Click **Update** and republish the policy.

Register IP addresses for AWS and Azure VMs in Infoblox

Create a policy to register IP addresses for AWS and Azure virtual machines in Infoblox, once these virtual machines are provisioned.

Before you begin
Role required: sn_cmp.cloud_admin.

About this task
When AWS and Azure virtual machines are provisioned, IP addresses are automatically allocated to them. You can create a policy and the policy is invoked, once these virtual machines are provisioned, to get their IP addresses and register these IP addresses in Infoblox.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: This functionality is not supported with our template-based cloud catalogs.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Click **New** and then fill in the form.

Field	Description
Policy Name	Enter a descriptive name.
Policy Trigger	Select **on Resource operation**.
Resource Block	Select **ResourceBlock Virtual Server**.
Operation	Select **Virtual Server interface provision**.
Moment	Select **Pre Operation**.

3. Click **Submit** to publish the policy.
4. Create a policy rule for the policy.
   Specify two conditions in the condition builder for locating the AWS and the Azure virtual machines. Select **Request form** > **starts with** and enter AWS. Specify another condition for Azure.
5. Click **Update**.
6. In the **Policy Rules Action** related list, click **New** to create a policy action.
7. Select **IP Address Management** in the **Create Policy Rule Action** dialog box.
   Enter a name in the **Action Name** field, select **Infoblox** in the **Provider Type** list, and select **Register IP address** in the **IPAM Method Name** list.
8. Click **Submit**.
9. In the **Policy Rule Action Attributes** section, enter the domain name of your company's network in the **DNSSuffix** field.
   As an example, `corp.servicenow.com`. Values for all the other properties are already populated.
10. Click **Update** and republish the policy.

Delete a host record in Infoblox

When a virtual machine is de-provisioned, delete the host record from Infoblox. Deleting the host record enables that IP address to be available and be reused for other virtual machines.

Before you begin

Role required: sn_cmp.cloud_admin

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Click **New**.
3. Fill out the fields on the Policy form (see table).

Field	Value
Policy Name	Enter a descriptive name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Policy Trigger	Select **on Resource operation**.
Resource Block	Select **ResourceBlock Virtual Server**.
Operation	Select **Virtual Server interface Deprovision**.
Moment	Select **Post Operation**.

4. Click **Submit** to publish the policy.
5. Create a policy rule for the policy and click **Update**.
6. In the **Policy Rules Action** related list, click **New** to create a policy action.
7. Select **IP Address Management** in the **Create Policy Rule Action** dialog box.
   Enter a name in the **Action Name** field, select **Infoblox** in the **Provider Type** list, and select **Release IP address** in the **IPAM Method Name** list.
8. Click **Submit**.
9. In the **Policy Rule Action Attributes** section, enter the domain name of your company network in the **DNSSuffix** field.
   As an example, `corp.servicenow.com`. Values for all the other properties are already populated.
10. Click **Update** and republish the policy.

Support for continuous delivery (configuration management)

The Cloud Provisioning and Governance application supports integration with continuous delivery solutions (also known as configuration management). Ansible is supported as the default config management provider.

Important: Starting with the Brazil release, Cloud Config Management is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

The Cloud Provisioning and Governance application supports configuration management providers in two ways:

- **Application profile based approach**: Create an application profile for an entity and in the order catalog form, select that profile to deploy that entity. Each profile has a config installable. You can create multiple application profiles- as an example, one for PostgreSQL and one for Apache. Each profile can be mapped to a configuration management provider. In the Cloud User Portal, based on the profile you select, a configuration grid appears populated with attributes and values for that specific configuration management provider. You can modify any values in the order form and provision the entity.
- **Resource block based approach**: Use a resource block, like Apache2, in a blueprint to represent an entity that a configuration management provider manages. The same resource block can be used to support more than one entity of the same kind. As an example, an Apache2 resource block can be used to support multiple Apache2 servers. In the order catalog form, select a provider type and then select a provider. Based on the specific provider, configuration management attributes and values appear. Once you provision the entity, you can save the configuration of the resource block and store it in the CMDB for future use.
- **Post-provisioning actions during catalog deployment**: Users can call ansible job as a post-provisioning step in any catalog which deploys a VM. This will be executed as a day-2

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

operation once the VM installation is completed. For more information on the processes, see https://www.servicenow.com/community/itom-blog/cpg-calling-ansible-job-as-post-provision-step-for-cloud-catalog/ba-p/2271552 and https://www.servicenow.com/community/itom-blog/cpg-calling-ansible-job-as-day2-operation/ba-p/2270969.

Important: Starting with the Orlando release, the cloud provisioning blueprints are available on instances upgraded from a previous release but you cannot create new blueprints. Existing blueprints and catalog items from those blueprints remain unaffected and continue to work.

What to do

Step	Goal	See these topics
1. Create credentials for Ansible.	Enable Cloud Provisioning and Governance to access the configuration provider with the necessary credentials.	Configure Ansible Tower user name and password
2. Create the workload provider type, and then create the provider.	Configure Cloud Provisioning and Governance to work with the configuration management provider of your choice. Then run discovery on the provider, using the credentials you provide, to find the resources that the provider already owns.	Create a workload provider type and Create an Ansible configuration management provider and run Discovery
3. Create an application profile with a mapping	Create the profile that defines the application that the configuration management provider manages. The profile is based on the discovered inventories in the provider.	Create an application profile
4. Create a blueprint	Create a blueprint with *BootstrapNode*, *Register Node*, and *ExecuteConfigPackages* operations on the virtual resource that the configuration provider manages. You can also customize the input parameters on the form to allow the user to select important inputs, like the application profile template, organization, and credential ID.	
5. Provision a resource from the Cloud User Portal	The resource should provision, with the stack status indicating the *BootstrapNode*, *Register Node*, and *ExecuteConfigPackages* steps.	

Configure Ansible Tower user name and password

Configure Ansible Tower credentials to access the Ansible Automation Engine through your instance.

Before you begin

Role required: sn_cmp.cloud_service_designer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Orchestration** > **Credentials**, and then click **New**.
2. Click **Ansible Tower Credentials**.
3. Fill out the form fields (see table):

   To integrate Cloud Provisioning and Governance with the Ansible configuration management account, you must configure the user name and password for the administrator account in Ansible.

Form fields for Ansible Tower credentials

Field	Description
Name	Provide a descriptive name.
User name/ Password	Enter the authentication credentials for the Ansible Tower user with administrator rights.

Note: You do not need to configure the other fields.

Create a workload provider type

Create a workload provider type for each new configuration management provider. This information appears in the order catalog form as management attributes that your users can select when provisioning a virtual resource through a configuration management provider.

Before you begin

- For information on Ansible naming conventions, refer to KB3136715
- Role required: cloud_admin

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Config Management**.
2. Select **Workload Config Provider Types**, and then select **New**.
3. Fill in the form fields (as shown in the table).

Workload Config Provider Type form

Field	Description
Name	Name for the workload provider type.
Product Type	A cloud product type from the list.
Config CI	Configuration table created for the provider that displays the discovered resources from the list.
Credential Resolver	A Credential Resolver translates a unique identifier into the actual user credentials.
Server Type	Ansible server type used by the provider. The available options are: ○ Ansible Community ○ Ansible Tower (Legacy)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ AWX ○ Ansible Automation Platform
Version	Enter the version of the provider. For any new config provider type, a **Version** field needs to be added to the corresponding Name **[config_ci]** table.
Credential Type	Types of credentials stored for this provider.

Example Ansible workload provider type

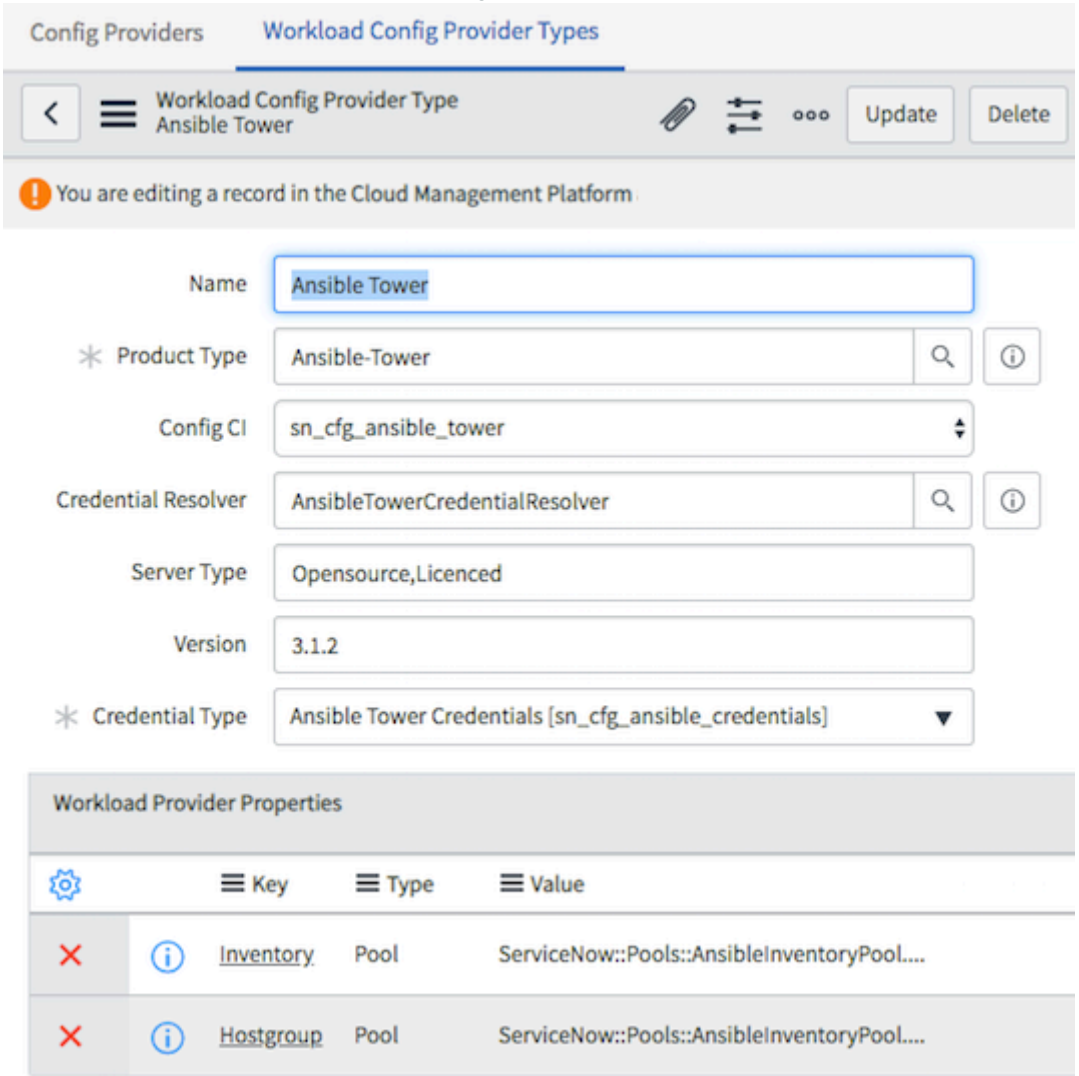

4. Add properties for the workload type in the **Workload Provider Properties** section.
   In addition to the existing properties, you can add more properties. Workload provider properties are displayed in the order catalog form (in the Cloud User Portal) as management attributes.

   For example, for an Ansible provider type, **Inventory**, and **Hostgroup** are required. The values for these properties come from Resource Pools.

5. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create an Ansible configuration management provider and run Discovery

Create an Ansible configuration management provider, and then run Discovery on the provider to find its resources.

Before you begin

- Ensure to have an Ansible server and Ansible credentials.
- If you want to use Ansible version 3.6.x or higher, ensure to set the mid.cmp.ansible.api_version property to V2. You can access this property under the Properties section of the Mid Server module.
- For information on Ansible naming conventions, refer to KB3136715
- Role required: cloud_admin

About this task

Important: Starting with the Orlando release, the cloud provisioning blueprints are available on instances upgraded from a previous release, but you cannot create a blueprint. Existing blueprints and catalog items created from those blueprints remain unaffected and continue to work.

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Config Management**.
2. Select **New**.
3. On the form fill in the fields.

Note: Most information required to create an Ansible configuration management provider comes from the Ansible server settings.

Field	Description
Name	Descriptive name for the configuration provider.
Organization	Ansible organization for access control You can obtain this information from the configuration management provider console.
URL	URL of the Ansible config. Enter the full URL of the Ansible config including the `https://<IP>` protocol.
Provider	Configuration provider type. Select **Ansible** from the drop-down list.
Server Type	Ansible server type used by the provider. The available options are: ○ Ansible Community ○ Ansible (Legacy) ○ AWX ○ Ansible Automation Platform
Credential	Credentials to access the Ansible server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	For more information on creating credential Ansible credentials
Version	Version of the configuration provider you're creating. **Note:** For Ansible versions higher than 3.1.2, select 3.4.0 from the list.

4. Select **Submit**.
5. Select the Ansible configuration provider card.

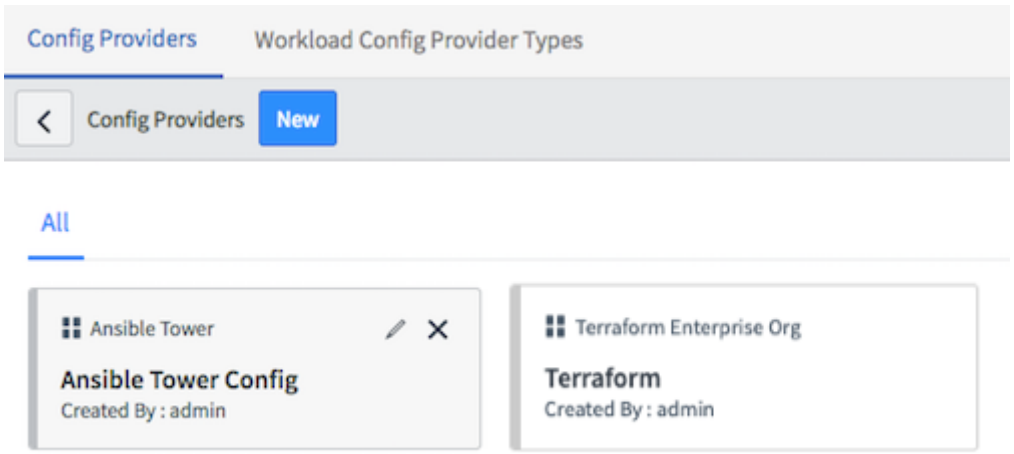

6. Select **Discover Now** to find the resources in the Ansible configuration management provider.

   Note: If you provision or otherwise change resources through the configuration management provider interface, you must manually run discovery again through the Config Providers page. You cannot create a scheduled discovery for configuration management providers.

   The discovered resources appear under **Entities**.

   The following Ansible resources are discovered:

   - **Ansible Inventory**: Displays the discovered applications and virtual resources.
   - **Cfg Installable**: Displays theactual components that run when the applications install. For Ansible, the items in this list are job templates, which are also referred to as runlists in this form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example Ansible Inventory

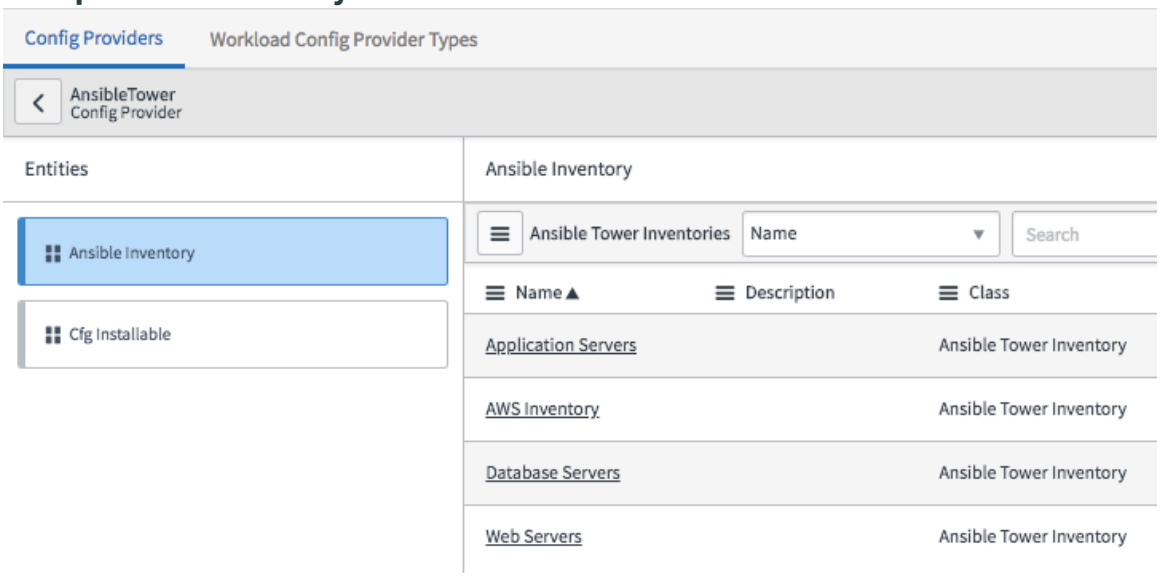

Ansible config installables (job templates)

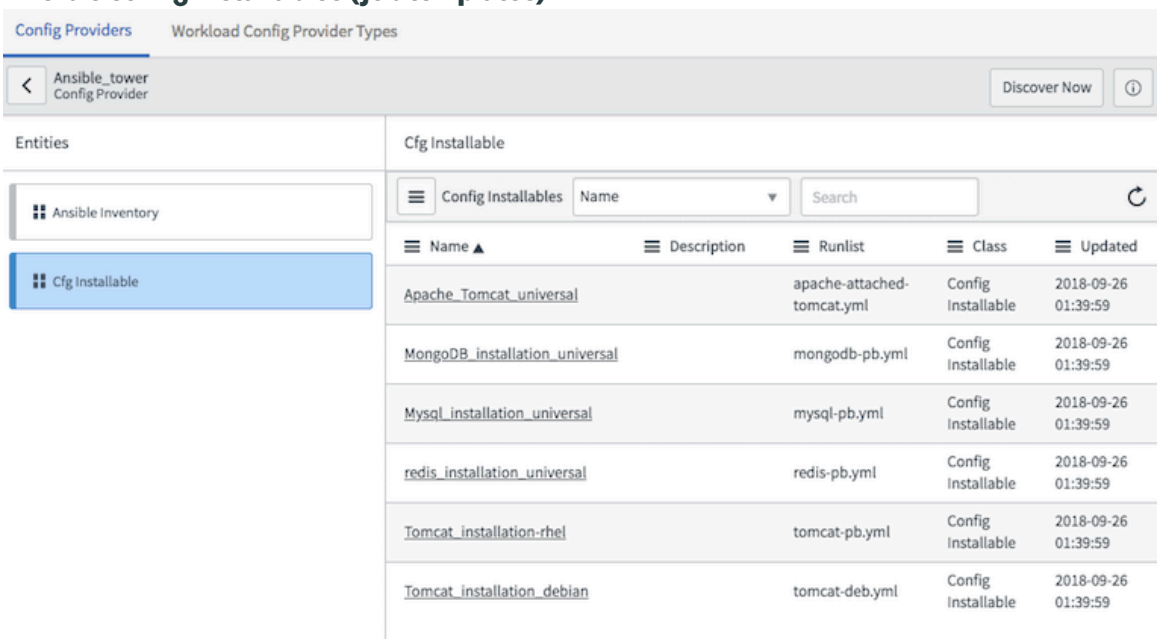

7. Select **Ansible Inventory** to explore the contents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Communicate the supported applications to users so they can select the correct application from the **Hostgroup** field in the Cloud User Portal. The **Hostgroup** field shows all possible applications (called Groups in Ansible), not just the ones that the configuration management provider supports. Therefore, the provisioning fails when the user selects an unsupported application.

Ansible groups within an inventory

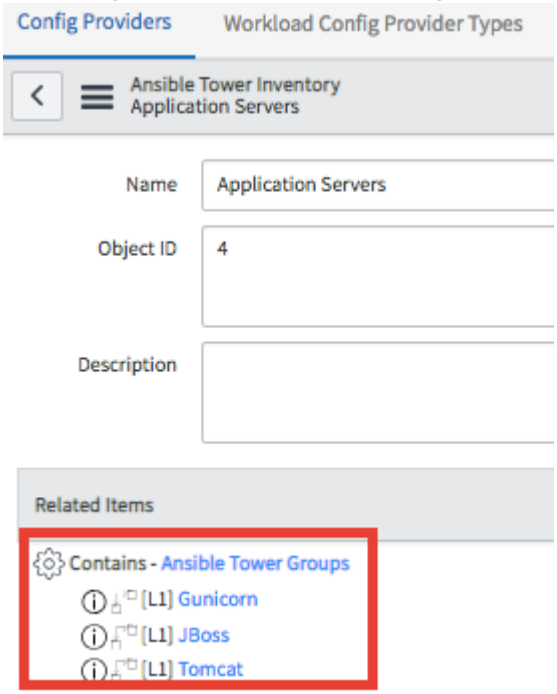

What to do next

After a user provisions, the Stack, the Stack Status indicates how the system runs through the Create node, Bootstrap, and ExecuteConfigPackage steps. You can obtain the IP address of a virtual machine in the User Portal by navigating to **Stacks** > **{category}** and selecting the new virtual machine. Open the configuration management provider server to see the newly provisioned resource on the node that the user specified.

Cloud Services Catalog Terraform Connector

Use the ServiceNow® Cloud Services Catalog Terraform Connector to generate catalog items from the Terraform templates and use them to provision and manage resources in various clouds.

Important: Starting with the Brazil release, Cloud Services Catalog Terraform Connector is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore Learn about the Cloud Services Catalog Terraform Connector application.	Configure Plan and configure your implementation.
Use Use the Cloud Services Catalog Terraform Connector app to discover resources from the Terraform environment, create catalog items, and maintain them.	Reference Get details about components like fields, tables, and properties.

Troubleshoot and get help

- Cloud Provisioning and Governance forum on the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Cloud Services Catalog Terraform Connector

Whether you're starting or expanding your implementation of Cloud Services Catalog Terraform Connector, learn more about the available features that help in provisioning and managing cloud resources across various public and private clouds.

The Cloud Services Catalog Terraform Connector integrates with the Cloud Provisioning and Governance application and helps you to create or update catalog items from the Terraform templates. Then, you can use the catalog items to provision the cloud resources.

Key features of Cloud Services Catalog Terraform Connector

- Discover the Terraform templates and workspaces from the Terraform environment.
- Create catalog items from the Terraform templates.
- Detect changes in the Terraform template and use the IaC change task workflow to update the catalog item.
- Utilize the multi-repo support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Specify metadata to populate the catalog item form fields and values.
- Reduce the Terraform Cloud license cost by using the discovered Terraform workspaces to provision the cloud infrastructure.

Use Terraform as a configuration provider

Use the Terraform Open Source, Terraform Enterprise, and Terraform Cloud config provider to provision resources in the following clouds.

Starting this release, you can use OpenTofu with Cloud Provisioning and Governance. You must have admin rights to configure the alias name for Terraform Open Source at the vm host level.

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud Provider
- VMware vSphere
- IBM Cloud Platform (Deprecated)

In addition, you can use the Cloud Services Catalog Terraform Connector to provision resources to the Oracle Cloud Infrastructure (OCI) through the Cloud Provisioning and Governance: Oracle Cloud (OCI) Add-On for Terraform Connector app. For more information, see Cloud Provisioning and Governance: Oracle Cloud (OCI) Add-On for Terraform Connector.

Starting with Cloud Provisioning and Governance: Terraform Connector version 1.5.2, the app doesn't support provisioning resources to the IBM Cloud Platform.

Discover the Infrastructure as Code (IaC) templates

Use Cloud Services Catalog Terraform Connector to discover the IaC (Terraform) templates from the Terraform environment. Terraform Enterprise or Terraform Cloud organizations use various Version Control Systems to store the Terraform templates. Cloud Services Catalog Terraform Connector can discover the templates stored in the following VCS:

- Public and self-hosted GitHub
- Public and self-hosted GitLab
- Public and self-hosted Microsoft Azure DevOps

For more information, see Run the IaC Discovery.

Utilize the hierarchical structure of multiple repositories with the support of Terraform Connector IaC discovery capabilities.

Create catalog items from the Terraform templates

Use the Cloud Services Catalog Terraform Connector app to create catalog items from the discovered Terraform templates and use them to provision resources in various clouds. For more information, see Create a catalog item from the Terraform template.

Detect changes in the Terraform templates

Use IaC Discovery to detect changes in the Terraform templates that are used in existing catalog items. When Cloud Services Catalog Terraform Connector detects a change in the Terraform template, it creates an IaC change task. You can use the IaC change task to update the affected catalog item or deactivate it, and then create another catalog item from the updated Terraform

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

template. For more information on the catalog item change task, see Manage the Terraform template-based catalog items.

Reduce Terraform Cloud license cost

Use the discovered Terraform workspaces to provision the catalog items from the Terraform templates. Reusing workspace helps you to manage the Terraform Cloud license cost by avoiding the creation of workspaces.

Configuring Cloud Services Catalog Terraform Connector

Understand the high-level workflow for configuring the Cloud Services Catalog Terraform Connector application.

Perform the following steps to configure the Cloud Services Catalog Terraform Connector application:

1. Install Cloud Services Catalog Terraform Connector. For more information, see Install Cloud Services Catalog Terraform Connector.
2. If you want to use a minor or patch version of Terraform Open Source, Terraform Enterprise, or Terraform Cloud, add the minor or patch version to the workload config provider record. For more information, see Add support for minor or patch versions of the Terraform.
3. To use Terraform Open Source as a config provider for Cloud Provisioning and Governance, do the following:
   a. Assign Terraform capability to the MID Server. For more information, see Configure MID Server capabilities.
   b. Create a Terraform Open Source config provider. For more information, see Create a Terraform Open Source config provider.
4. To use Terraform Enterprise or Terraform Cloud as a config provider for Cloud Provisioning and Governance, do the following:
   a. Assign Terraform Enterprise and REST capabilities to the MID Server. For more information, see Configure MID Server capabilities.
   b. Create a Terraform Enterprise or Terraform Cloud config provider. For more information, see Create a Terraform Enterprise or Terraform Cloud config provider.
5. Discover resources from the Terraform environment. For more information, see Run the IaC Discovery.

Install Cloud Services Catalog Terraform Connector

You can install the Cloud Services Catalog Terraform Connector application (com.sn_cmp_terraform) if you have the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Cloud Services Catalog: Terraform Connector application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Ensure that Cloud Provisioning and Governance (com.snc.cloud.mgmt) application is installed and activated in the ServiceNow AI Platform. For more information, see Request the Cloud Provisioning and Governance application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Cloud Services Catalog Terraform Connector application (com.sn_cmp_terraform) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

What to do next

- If you want to use Cloud Services Catalog Terraform Connector with Terraform Open Source, Create a Terraform Open Source config provider.
- If you want to use Cloud Services Catalog Terraform Connector with Terraform Enterprise, Create a Terraform Enterprise or Terraform Cloud config provider.

Add support for minor or patch versions of the Terraform

Add minor or patch versions of the Terraform (Terraform Open Source, Terraform Enterprise, or Terraform Cloud) to the workload provider type record.

Before you begin
Role required: admin

About this task
The base system Cloud Services Catalog Terraform Connector supports the following Terraform versions:

- Terraform Open Source versions 1.1.0 through 1.1.9 and 1.2.0 or higher.
- Terraform Enterprise versions 1.1.9 and 1.2.0 or higher.
- Terraform Cloud versions 1.1.9 and 1.2.0 or higher.

To use templates that are supported by a minor or patch version, for example version 0.12.29, add the version number to the Workload provider type record. Then, use the minor or patch versions to create configuration providers in Cloud Provisioning and Governance.

Procedure

1. Navigate to `https://<instance-name>/sn_cmp_workload_provider_type_list.do`
2. Select the Terraform type whose minor or patch version that you want to use.
3. Change the application scope to **Cloud Provisioning and Governance: Terraform Connector**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the **Version** field of the Workload Config Provider Type form, enter the minor or patch versions as comma-separated values.
5. Select **Submit**.

Create a Terraform Open Source config provider

Create a Terraform Open Source config provider in Cloud Provisioning and Governance. The Terraform Open Source config provider enables Cloud Provisioning and Governance to discover the Terraform Open Source config installables (Terraform templates) and detect changes in them.

About this task

Starting this release, you can use OpenTofu with Cloud Provisioning and Governance. You must have admin rights to configure the alias name for at the vm host level.Ensure that server is running on Windows or Linux operating system.

Before you begin

- Ensure that a supported Terraform Open Source version is installed in the server. Cloud Services Catalog Terraform Connector supports Terraform Open Source versions 1.1.0 through 1.1.9 and 1.2.0 or higher
- If you want to use a minor or patch version of the Terraform Open Source CLI, ensure that the minor or patch version is added to the workload config provider. For more information, see Add support for minor or patch versions of the Terraform.
- Ensure that the MID Server can access the Terraform server.
- Ensure that the MID Server has Terraform capability.
- As a new user, add Terraform MID capability to provision resources using Terraform Open-source host (VM). Add cloud provider specific mid capability (AWS,Azure,VMware, Google) to discover resources from cloud after provisioning Terraform Open-source catalog items. For example: If you are provisioning VM using Terraform Open-Source in Google, you would need mid with Terraform and Google capability. Both capabilities can be set on the same or different mids. Mid with Terraform capability should be able to connect to Terraform Host (VM) and that with Google capability must be able to reach google cloud. For more information, see Configure MID Server capabilities.
- Ensure that all the templates are stored in separate directories and all the template directories are in the same base directory.
- Ensure that Terraform Connector has access to the Terraform registry (registry.hashicorp.com).

Role required: cloud_admin

Procedure

1. **Optional:** If you have deployed the Terraform Open Source on a Linux server, create SSH Private Key Credentials for the server in the ServiceNow AI Platform.

   a. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.
   The Cloud Admin Portal opens in a new browser tab.

   b. On the Cloud Admin Portal, navigate to **Manage** > **Credentials**.

   c. Select **New**.

   d. Select **SSH Private Key Credentials**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. On the form, fill in the fields.
For a description of the form fields, see Cloud Services Catalog Terraform Connector SSH Private Key Credentials form reference.

f. Select the Test Credentials related link.

g. On the form, fill in the fields.

Test Credentials form

Field	Description
Target	IP address of the Terraform server for which you are creating the credential.
MID Server	MID Server that can be used to connect with the Terraform server.

h. Select **OK**.

i. Select **Submit**.

2. **Optional:** If you have deployed the Terraform Open Source on a Windows server, create Windows Credentials for the server in the ServiceNow AI Platform.

a. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.
The Cloud Admin Portal opens in a new browser tab.

b. On the Cloud Admin Portal, navigate to **Manage** > **Credentials**.

c. Select **New**.

d. Select **Windows Credentials**.

e. On the form, fill in the fields.
For a description of the form fields, see Cloud Services Catalog Terraform Connector Windows Credentials form reference.

f. Select the Test Credentials related link.

g. On the form, fill in the fields.

Test Credentials form

Field	Description
Target	IP address of the Terraform server for which you are creating the credential.
MID Server	MID Server that can be used to connect with the Terraform server.

h. Select **OK**.

i. Select **Submit**.

3. Create a config provider for Terraform Open Source.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. On the Cloud Admin Portal, navigate to **Manage** > **Config Management**.

b. Select **New**.

c. On the form, fill in the fields.
For a description of the form fields, see Cloud Services Catalog Terraform Connector Create Config Provider form reference.

d. Select **Submit**.

What to do next

Discover the Terraform templates from the Terraform Open Source server. For more information, see Run the IaC Discovery.

Create a Terraform Enterprise or Terraform Cloud config provider

Create Terraform Enterprise or Terraform Cloud config provider in Cloud Provisioning and Governance. The config provider enables Cloud Provisioning and Governance to discover the workspaces and Version Control System (VCS) repositories attached to the Terraform organization.

Before you begin

- Ensure that a supported Terraform Enterprise or Terraform Cloud version is in use. Cloud Services Catalog Terraform Connector supports the following Terraform versions:
  - Terraform Enterprise versions 1.1.9 and 1.2.0 or higher
  - Terraform Cloud versions 1.1.9 and 1.2.0 or higher
- To use a minor or patch version of the Terraform Enterprise or Terraform Cloud, ensure that the minor or patch version is added to the workload config provider. For more information, see Add support for minor or patch versions of the Terraform.
- Ensure that the following properties are correctly configured in the MID Server. The MID Server uses these properties to connect with the Terraform Enterprise or Terraform Cloud.
  - *glide.http.proxy_host*
  - *glide.http.proxy_port*
  - *glide.http.proxy_username*
  - *glide.http.proxy_password*

  For more information, see the Knowledge Base article KB0754247.
- Ensure that the MID Server has Terraform Enterprise and REST capabilities. For more information, see Configure MID Server capabilities.
- Ensure that the MID Server can communicate with the VCS accounts provisioned in the Terraform Enterprise or Terraform Cloud organization.

Role required: admin

Procedure

1. Create an API key credential for Terraform Enterprise or Terraform Cloud account.
   For more information, see Create API key credentials for Terraform and VCS accounts account.
2. Create a config provider for Terraform Enterprise or Terraform Cloud.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. On the Cloud Admin Portal, navigate to **Manage** > **Config Management**.

b. Select **New**.

c. On the form, fill in the fields.
For a description of the form fields, see Cloud Services Catalog Terraform Connector Create Config Provider form reference.

d. Select **Submit**.

What to do next
Discover the workspaces and Version Control System (VCS) repositories attached to the Terraform organization specified in the config provider. For more information, see Run the IaC Discovery.

Create API key credentials for Terraform and VCS accounts account

Create API key credential for Terraform account and supported Version Control System (VCS) accounts that contain the Terraform template repositories.

Before you begin
Role required: admin

About this task
An API key credential is a secured record in the ServiceNow AI Platform that stores an API key such as Terraform Enterprise, Terraform Cloud, or VCS account API key. Cloud Services Catalog Terraform Connector uses the API key credential to connect with the account and discover its resources.

Procedure
1. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.
   The Cloud Admin Portal opens in a new browser tab.
2. On the Cloud Admin Portal, navigate to **Manage** > **Credentials**.
3. Select **New**.
4. Select **API Key Credentials**.
5. On the form, fill in the fields.
   For description of the form fields, see Cloud Services Catalog Terraform Connector API Key Credentials form reference.
6. Select **Submit**.

Use multiple repositories structure with Terraform Connector app

You can now discover multiple Terraform configurations from one repository, at any folder level with Terraform Enterprise Config Provider.

Before you begin
Role required: admin

Perform the Run the IaC Discovery steps and specify what repository, branches and folder paths have to be considered for discovering Terraform configurations with these steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Add repositories and branches that you want to consider for discovering Terraform configurations.
2. Use **<Repository-Name>::<Branch-Name>** format to include a branch from a repository.
3. To specify multiple repository and branch mappings, use comma (,).

   Example: Repo1::Branch1, Repo2::Branch2.

   Leave the field bank if you want to discovery all the repositories and branches in that VCS.
4. Select **Do you wish to specify sub-folder paths for Discovery** if you want to discover Terraform configurations from specific folder paths of the repository.
5. In the **Folder Path List**, add the list of folder paths separated by comma(,) to discover terraform configurations from it.
   - Folder path filter applies to all Repository Branches included in 'Repo and Branch inclusion List'.
   - All Terraform configurations you want to discover from a specific folder path needs to be in subfolders just one level below in folder hierarchy. Example:**/gcp/vm** folder path has one Terraform configuration for virtual machine,
   - Specify **/gcp** in folder path list to discover vm Terraform configuration.

Run the IaC Discovery

Run the Infrastructure as Code (IaC) discovery to identify the resources of the Terraform environment.

Before you begin

- Ensure that you have the appropriate API credentials for the Version Control System (VCS) accounts provisioned in the Terraform Enterprise or Terraform Cloud organizations. For more information, see Create API key credentials for Terraform and VCS accounts account.
- Ensure that Terraform Open Source folders which contain the Terraform templates have unique names. If several Terraform template folders have the same name, multiple change tasks are created even though all the templates mayn't contain variable changes.

Role required: sn.cmp.cloud_service_designer

About this task

Depending on the configuration provider, Cloud Services Catalog Terraform Connector discovers the following resources from the Terraform environment:

Terraform Open Source

The IaC Discovery discovers the config installables (Terraform templates) from the Terraform Open Source server.

Terraform Enterprise or Terraform Cloud

The IaC Discovery discovers the following items:

- Workspaces: Terraform workspaces provisioned in the Terraform organization.
- VCS: VCS repositories and branches provisioned in the Terraform organization.
- Agent pools:

  Terraform Cloud uses agents and agent pools to communicate with on-premise infrastructure. Deploy agents in the same network as your infrastructure to

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

establish a connection between the infrastructure and Terraform Cloud. For more information on agents and agent pools, see Terraform Cloud agents.

Cloud Services Catalog Terraform Connector discovers the agent pools created in the Terraform Cloud. You can view the discovered agents in the Config Management module for your config provider. On the catalog order form, an agent pool filter enables you select an active agent to communicate with the private VMware infrastructure.

Note: You can see agent pools option in the config management form and the catalog order form when the following conditions are met:

- You have created the VMware agent pools in the Terraform Cloud organization.
- You have VMware templates for Terraform Enterprise.

- Cfg installables: Terraform templates stored in the repositories of the VCS account provisioned the Terraform organization.

Note: To discover the Terraform templates from the VCS, Cloud Services Catalog Terraform Connector requires the API key credentials of the VCS account.

Procedure

1. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.
   The Cloud Admin Portal opens in a new browser tab.
2. On the Cloud Admin Portal, navigate to **Manage** > **Config Management**.
3. Select the Terraform provider card.
4. Select **Discover Now**.

   Note: You can't cancel the IaC Discovery job after triggering it.

5. **Optional:** To use Cloud Services Catalog Terraform Connector with Terraform Enterprise or Terraform Cloud, discover the Terraform templates stored in the provisioned VCS.
   a. Select the **Tfe Vcs** tab.
   b. From the Tfe Vcs list, select the VCS account that contains the Terraform templates you want to use for Cloud Provisioning and Governance catalog item creation.
   c. From the Credentials drop-down list, select the API credential of the selected VCS account.
   d. From the Related Links, select **Find Branches and Repositories**.
   e. Repeat step 5.b to 5.d for all VCS accounts that contain the Terraform templates you want to use for Cloud Provisioning and Governance catalog item creation.
6. Create an Infrastructure as Code (Iac) discovery schedule.
   IaC discovery schedule helps to discover new Terraform templates, workspaces, agent pools, and VCS branches for the specified config provider. It also detects changes in the already discovered Terraform templates. When Cloud Services Catalog Terraform Connector detects a change in the template variable or the template metadata, it automatically creates IaC change tasks for the affected catalog items.

   a. On the Cloud Admin portal, navigate to **Manage** > **IaC Discovery**.

   b. Select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. On the form, fill in the fields.

IaC Discovery Schedulers form

Field	Description
Name	Unique and descriptive name of the IaC Discovery schedule.
Provider	Config management provider for which you are creating the discovery schedule.
Active	Option to activate or deactivate the schedule.
Run	Frequency of schedule execution. Configure the schedule execution frequency in the Run field and its associated form fields.

d. Select **Submit**.

What to do next

- Use multiple repositories structure with Terraform Connector app with IaC Discovery
- If Cloud Services Catalog Terraform Connector discovers new Terraform templates that you want to use for cloud resource provisioning, create a catalog item from the discovered template. For more information, see Create a catalog item from the Terraform template.
- If Cloud Services Catalog Terraform Connector has discovered changes in an existing template, take one of the following actions to resolve the associated IaC change task:
  - Update the affected catalog item per the updated Terraform template.
  - Deactivate the affected catalog item, and then create a catalog item per the updated template.

  For more information, see Manage the Terraform template-based catalog items.

Domain separation and Cloud Services Catalog Terraform Connector

Domain separation is supported for Cloud Services Catalog Terraform Connector. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Cloud Services Catalog Terraform Connector integrates with Cloud Provisioning and Governance and it complies with the Cloud Provisioning and Governance domain separation. For more information, see Domain separation and Cloud Provisioning and Governance.

Upgrading the Terraform environment

Understand the high-level workflows for upgrading the Terraform environment.

Upgrading Terraform Open Source

Upgrade methods for Terraform Open Source

Upgrade method	Description
Upgrade the existing Terraform Open Source installation to a newer supported version	**Note:** If you choose to upgrade the existing Terraform Open Source installation, you may experience a longer upgrade time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Upgrade methods for Terraform Open Source (continued)

Upgrade method	Description
	Perform the following actions to upgrade the existing Terraform Open Source installation: **1.** Upgrade the Terraform Open Source server to a version that is supported by the Cloud Services Catalog Terraform Connector. For more information on upgrading Terraform, see Upgrade Guides in Terraform documentation. Cloud Services Catalog Terraform Connector supports Terraform Open Source versions 1.1.0 through 1.1.9 and 1.2.0 or higher **2.** Test and modify the existing Terraform templates to ensure compatibility with the target Terraform Open Source version. **3.** Update the Terraform Open Source version mentioned in the Terraform Open Source config provider.
Set up a new Terraform server with the desired Terraform Open Source version	Perform the following actions to set up a new Terraform server with the desired Terraform Open Source version: **1.** Create a Linux or Windows server for Terraform Open Source installation. **2.** Install a Terraform Open Source version that is supported by the Cloud Services Catalog Terraform Connector. Cloud Services Catalog Terraform Connector supports Terraform Open Source versions 1.1.0 through 1.1.9 and 1.2.0 or higher **3.** Copy the existing Terraform templates to the new Terraform server. **Note:** Ensure that all the templates are stored in separate directories and all the template directories are in the same base directory. **4.** Test and modify the copied Terraform templates to ensure compatibility with the target Terraform Open Source version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Upgrade methods for Terraform Open Source (continued)

Upgrade method	Description
	5. Create a Terraform Open Source config provider. For more information, see Create a Terraform Open Source config provider. 6. Remove the old catalog items per the business strategy of your organization, and then retire the old Terraform server.

Upgrading Terraform Enterprise and Terraform Cloud

Test and modify the existing Terraform templates to ensure compatibility with the new Terraform Enterprise or Terraform Cloud version.

Note: The stacks and Terraform workspaces are tied to the Terraform version through which they were provisioned. Therefore, the same Terraform version is used to provision and deprovision a stack or Terraform workspace.

Using Cloud Services Catalog Terraform Connector

Understand the high-level use cases and usage workflows for the Cloud Services Catalog Terraform Connector application.

Associate the Terraform input parameters with the Cloud Provisioning and Governance catalog item

Create a JSON file to store Terraform template variables and map them to discovered resources, pools, filters, and catalog item provision form fields.

Before you begin

Role required: cloud_designer

About this task

Cloud Services Catalog Terraform Connector supports this capability for Terraform Open Source, Terraform Enterprise, and Terraform Cloud.

Important: When you use Terraform templates to create catalog items, the Location value specifies the datacenter or region to provision the configuration items and stacks. Use the metadata file to map other variables based on discovered resources, networks, profiles, pools and filters, and so on.

To auto-populate form fields in the catalog creation form while using Terraform templates, do the following:

- Specify or define the key value pairs in the JSON file for variables that are defined in the Terraform template.
- Configure variables to auto-populate fields and data in the **Cloud User Portal** > **Cloud Catalog** > **New** > **Provision** page.

For more information on binding a parameter to a resource pool, see Bind a parameter to a resource pool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the directory or repository that is the basis for the catalog item, create a JSON file.

   Important: Name the JSON file `metadata.snc` and store it in the same directory or repository as the template (.tf) file. This naming convention enables discovery and parsing of the `metadata.snc` JSON file as part of the config installable.

2. Specify the equivalent pool and filter as the data type for the Terraform variable in the `metadata.snc` JSON file.

```
 "Metadata": {
   "SNC::Parameter::Metadata": {
     "rgName": {
       "datasource": "ServiceNow::Pools::
ResourceGroupPool.getByLDC"
     },
     "Network": {
       "datasource":
"ServiceNow::Pools::NetworkPool.getObjectsByLDC"
     },
     "Subnet": {
       "datasource": "ServiceNow::Pools::
SubnetPool.getObjectsByNetwork",
       "datasourceFilter": { "Network": "Network" }
     }
   }
 }
```

   The value that you specify in the variable field of the `metadata.snc` file is populated in the Provision section of the of the catalog item form.

   The following image shows the fields in the Provision tab of the catalog form. The values are auto-populate based on the variables defined in the `metadata.snc` file.

   Example: Auto-populated values on the Provision section of the catalog creation form

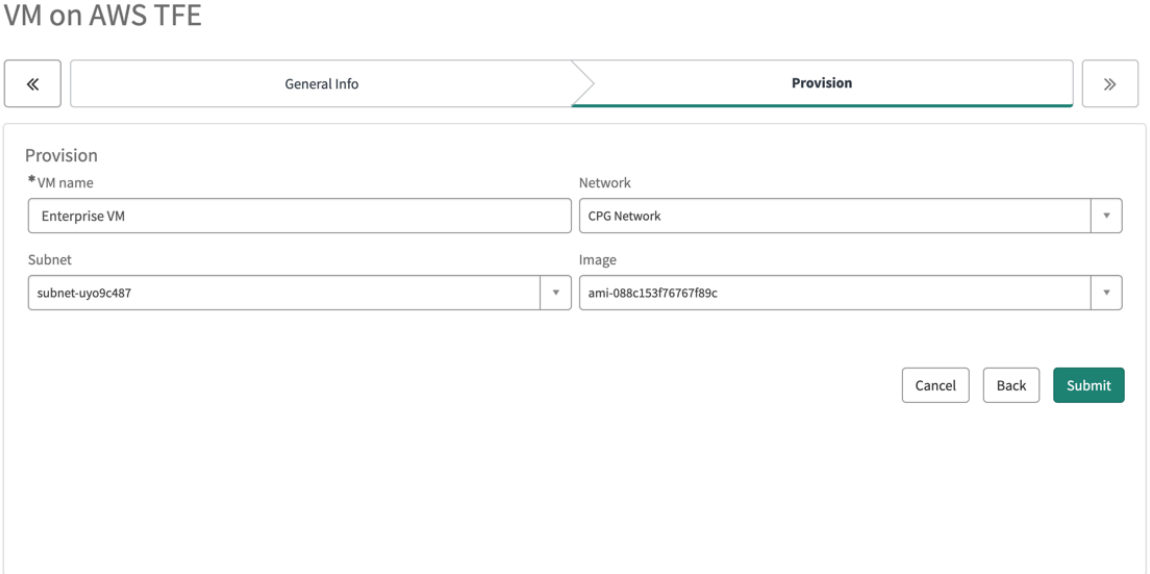

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Populate the datacenter in Terraform-based catalog items

In any request for a catalog item using Terraform templates, the required Location value specifies the datacenter or region where to provision the configuration items and stacks. You can use a variable named *region* in the Terraform template to map to the Location field on the request form.

Before you begin

Role required: cloud_designer

Procedure

In the template file that is the basis for the catalog item, define a variable named *region*.

```
variable "region" {
                type="string"
              }
```

- The field defined for the *region* variable does not appear on the Provision section of the request form for the catalog item (the Location field serves that purpose).
- The value that the user specifies for the Location field on the request form configures the datacenter or region for the resource.

Create a catalog item from the Terraform template

Create a catalog item from the Terraform template to request cloud resource provisioning. Activated catalog items appear in the cloud user portal.

Before you begin

- Ensure that the Terraform template contains only one provider. Templates with multiple providers or provider aliases will fail validation.
- Ensure that an appropriate Terraform config provider is available. For more information on creating the config provider, see Create a Terraform Open Source config provider and Create a Terraform Enterprise or Terraform Cloud config provider.
- Ensure that the Terraform template is working (that is, you can directly use the template with the Terraform system to provision the specified cloud resources). For more information on Terraform templates, see the Terraform language documentation.
- Ensure that the provider block of the Terraform template adheres to the following syntax:
  - Provider block syntax for the Microsoft Azure cloud:

```
provider "azurerm" {

  subscription_id = "${var.subscriptionId}"
  client_id       = "${var.clientId}"
  client_secret   = "${var.clientSecret}"
  tenant_id       = "${var.tenantId}"
}
```

  - Provider block syntax for the Amazon Web Services cloud:

> **Important:** By default, the ServiceNow® Cloud Provisioning and Governance (CPG) Terraform Connector does not support AWS IAM role-based authentication when using Terraform templates.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
provider "aws" {
  access_key      = "${var.access_key}"
  secret_key      = "${var.secret_key}"
  region          = "${var.region}"
}
```

- Provider block syntax for Google Cloud Platform

```
terraform {
  required_providers {
    google = {
      source  = "hashicorp/google"
      version = "5.10.0"
    }
  }
}

provider "google" {
  credentials = var.credentials
  project     = var.project
  region      = var.region
}
```

- Provider block syntax for VMware Cloud:

```
provider "vsphere" {
  user                 = "${var.user}"
  password             = "${var.password}"
  vsphere_server       = "${var.vsphere_server}"
}
```

Note: Do not hard-code any credential in the Terraform template. Use appropriate Terraform variables to fetch the credentials from the ServiceNow AI Platform.

- Ensure that the variable definition file (`variable.tf`, `.tfvars`, or `.tfvars.json`) contains the definition of the variables used in the provider block.
- If you are using Terraform Enterprise or Terraform Cloud, ensure that the MID Server can connect with Version Control System (VCS) accounts provisioned in the Terraform organization.

  To verify connectivity with the VCS accounts provisioned in the Terraform organization, execute the appropriate curl command on the MID Server.

  GitHub

  ```
  curl -H "Authorization: <API credential>
   https://<github-api-host>/user/repos?per_page=100&pag
  e=1"
  ```

  GitLab

  ```
  curl -H "Authorization: <API credential>"
   https://<gitlab-api-host>/api/v4/projects?visibility=
  private&per_page=100&page=1
  ```

  Microsoft Azure DevOps

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
curl -H "Authorization: <API credential>
 https://<azure-devops-api-host>/<organization>/<proje
ct>/_apis/git/repositories?api-version=6.0"
```

Where,

- *organization* is the name of the Microsoft Azure DevOps organization.
- *project* is the name of the Microsoft Azure DevOps project.

For more information on the API credentials, see Cloud Services Catalog Terraform Connector API Key Credentials form reference.

Role required: sn_cmp.cloud_service_designer

About this task

The Cloud Services Catalog Terraform Connector app supports the following Terraform data types:

- string
- number
- bool
- list
- tuple
- map
- object

Procedure

1. Create a cloud catalog item.

   a. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.
   The Cloud Admin Portal opens in a new browser tab.

   b. Navigate to **Design** > **Cloud Catalog Items**.

   c. Select **New**.

   d. On the form, fill in the fields.

Cloud Catalog Item form

Field	Description
Name	Unique and descriptive name of the catalog item.
Source	Source of the catalog item. Select **Configuration Management Template** from the drop-down list.
Provider Type	Terraform config provider type that you want to use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	▪ To use Terraform Enterprise or Terraform Cloud, select **Terraform Enterprise**. ▪ To use a Linux-based Terraform Open Source server, select **Terraform Environment**. ▪ To use a Windows-based Terraform Open Source server, select **Terraform Windows Environment**.
Provider	Name of the Terraform config provider that you want to use.

e. Save the Cloud Catalog Item form.

2. Associate a Terraform template with the catalog item.

   a. From the Cloud Templates related list, select **New**.

   b. From the Configuration installable drop-down list, select the Terraform template you want to use.

   c. Save the ServiceNow Cloud Templates Versions form.

   Note: A catalog item can include multiple versions of the cloud template, but only one version can be active.

   d. Select **Activate** to activate the cloud template.

3. **Optional:** To add or remove catalog item form fields, edit the variable sets associated with the catalog item.
4. Select the **Active** check box to activate the catalog item.
5. Select **Update**.

What to do next
Verify that the catalog item appears in the Cloud User Portal.

Manage the Terraform template-based catalog items

Use an Infrastructure as Code (IaC) change task to create or update the Terraform template-based catalog items. The IaC change task helps ensure that the catalog items are aligned with the latest infrastructure specifications coded in the Terraform template.

Before you begin
Role required: admin

About this task
When the Cloud Services Catalog Terraform Connector app detects a change in the variables on the Terraform template (.TF) file or the metadata (`metadata.snc`) file, it creates change tasks. The change tasks are assigned to users who had created the catalog items (user with [sn.cmp.cloud_service_designer] role). Based on the change detected on the template variables or the metadata, the catalog item task prompts you take one of the following actions:

- Create a catalog item from the modified template version.

  A new catalog item is created with the existing catalog item name appended by a number. For example, if the existing catalog item is called `Linux VM` the new catalog item will be created with the name `Linux VM1`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Update the existing catalog item per the updated template version.

  When you choose to update the existing catalog item per the updated template version, a new version of the template is created and activated. The existing template version is retired and rendered inactive.

Note: The name of the Terraform config provider should match between the environments, while exporting Terraform template based catalog items. For example, if the config provider is called TFE in dev environment, it should be named the same, TFE, in the target environment (UAT, prod, etc.) too.

Until you perform one of the above mentioned actions, the catalog item is deactivated. When the catalog item is in deactivated state, cloud users can't access it.

The Cloud Services Catalog Terraform Connector creates an IaC change task for a Terraform template-based catalog item when any one of the following conditions is met:

- One or more variables are added to the source Terraform template.
- One or more variables are deleted from the source Terraform template.
- One or more variables used in the source Terraform template undergo a data type change.
- Default value of one or more variables changes.

Note: Deprovisioning the impacted stack triggers de-provisioning at the cloud provider's end.

Procedure

1. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.
   The Cloud Admin Portal opens in a new browser tab.
2. Navigate to **IaC** > **Change Tasks**.
3. Open the appropriate change task.
4. **Optional:** Customize the new or updated variables detected in the Terraform template.

   a. Select the **Customize Variables** related link.

   b. Edit the default value of the new or updated Terraform variables from the Terraform Task Metadata related list.

   Note: The default value mentioned in the variable record of the Terraform Task Metadata related list overrides the default value mentioned in the Terraform template.

5. **Optional:** Select the **changes** link on the recommendation message to view a difference comparison in the template content, as detected by IaC Discovery.
6. On the form, fill in the fields.
   For more information, see Cloud Services Catalog Terraform Connector Terraform Catalog Item Task form reference.
7. Select **Update**.

Map credentials between Cloud Provisioning and Governance and Terraform

Map credential types and service account fields between Cloud Provisioning and Governance and Terraform using Credential Mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

The latest Cloud Services Catalog Terraform Connector must be installed from the ServiceNow Store.

Role required: sn_cmp.cloud_admin, admin

About this task

Use the Credential Mapping module to perform credential and service account mappings for datacenter credentials used by the Terraform provider to provision datacenter resources and Cloud Provisioning and Governance for discovery.

Procedure

1. Navigate to **All** > **Design** > **Credential Mapping**.
2. Select **New**.

Terraform Datacenter Mapping form

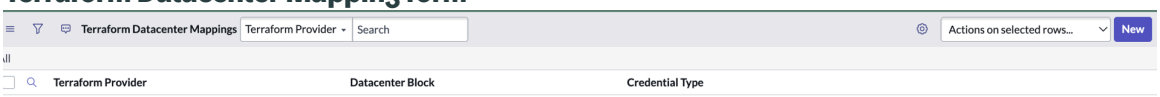

3. On the form, fill in the fields.

Terraform Datacenter Mappings form

Field	Description
Terraform Provider	Name of the Terraform provider.
Datacenter Block	The resource block whose credentials you are mapping from the lookup table.
Credential Type	The appropriate credential type from the lookup table. For example, Azure Service Principal [azure_service_principal].
Credential Resolver	An appropriate MID Server script include from the lookup table. For example, Search and select the TerraformProviderCredentialResolver script.

4. Select **Submit**.
   The Terraform provider credentials record is created and displayed in the Terraform Datacenter Mappings list.
5. On the record you created, navigate to the **Credentials Mappings** related list, and then select **New**.

   a. Complete the **Credential Mapping** form.

Field	Description
Provider Authentication Key	Authentication key (variable) field name as used by the provider. For example, the Terraform provider block for (Amazon Web Services) AWS uses the field `secret_key`.
Discovery Credential Key	The corresponding field in the Discovery credentials from the lookup table. For example, `access_key` field in Terraform corresponds to `Secret key` field in the ServiceNow instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Terraform Datacenter Mapping	Datacenter for which you are mapping the credentials.

b. Select **Submit**.

6. Select the **Service Account Mappings** tab, and then select **New**.

a. Complete the **Service Account Mapping** form.

Field	Description
Provider Authentication Key	Authentication key (variable) field name as used by the provider. For example, the Terraform provider block for AWS uses the field `secret_key`.
Service Account Field	Service Account (variable) field name from the lookup table, whose credentials you are mapping to the Provider Authentication Key.
Terraform Datacenter Mapping	Datacenter for which you are mapping the credentials.

b. Select **Submit**.

You have mapped the field values that you added under the Credential Mapping and Service Account Mapping related lists to the Terraform provider.

Associate the Terraform environment variable with a resource block

Associate the Terraform environment variable with the Cloud Provisioning and Governance datacenter resource block. After associating the environment variable, you can use it to pass custom inputs parameters during resource deployment through Terraform Enterprise.

Before you begin

Role required: admin

About this task

To set the Terraform environment variable conditionally, use a scripted expression. For more information, see Using expressions in Cloud Provisioning and Governance.

Procedure

1. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal**.
   The Cloud Admin Portal opens in a new browser tab.
2. Select **Design** > **Resource Blocks**.
3. Select the appropriate logical datacenter card.
4. Set the logical datacenter resource block to the **Draft** state.
5. Select the **Operations** tab.
6. From the Interface drop-down list, select **Terraform Orchestration Interface**.
7. Select **Steps** > **Flow.Subflow.TFEProvision**.
8. Set the **terraform_variables** text box to the environment variable key value pair.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example: **{"USER_AGENT_PROVIDER_NAME":"AgentValue"}**

9. Set the logical datacenter resource block to the **Published** state.

Cloud Services Catalog Terraform Connector reference

Reference topics provide additional information about the Cloud Services Catalog Terraform Connector application.

Cloud Services Catalog Terraform Connector Create Config Provider form reference

Cloud Provisioning and Governance uses the config provider to interact with the Terraform deployment.

Field	Description
Name	Name of the config provider.
Provider	Config provider type. • Terraform Enterprise: Select this value to create a Terraform Enterprise config provider. • Terraform Environment: Select this value to create a Linux based Terraform Open Source config provider. • Terraform Windows Environment: Select this value to create a Windows based Terraform Open Source config provider.
Organization	Terraform organization that contains the Terraform workspaces and template repositories. This field appears when you select Terraform Enterprise from the Provider drop-down list.
URL	Base URL of the Terraform API server or IP address of the Terraform server. • To create a config provider for Terraform Open Source, enter the IP address or hostname of the machine where Terraform CLI is installed or accessed via MID Server. • To create a config provider for Terraform Enterprise, enter the base URL of the Terraform Enterprise API server. • To create a config provider for Terraform Cloud, enter `https://app.terraform.io/api/v2`
Base Directory	Directory where all the individual Terraform template directories are located. This field appears when you select Terraform Environment or Terraform Windows Environment from the Provider drop-down list.
Server Type	Server type of the Terraform deployment.
Credential	Credential for accessing Terraform.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• For the Linux based Terraform Open Source server, select the appropriate SSH private key credential. • For the Windows based Terraform Open Source server, select the appropriate Remote Desktop Protocol (RDP) credential. • For Terraform Enterprise or Terraform Cloud, select the appropriate Terraform API key credential.
Version	Supported Terraform Open Source CLI version installed in the Terraform server. For more information on the supported Terraform Open Source versions, see the prerequisites section of Create a Terraform Open Source config provider. For more information on the supported Terraform Enterprise and Terraform Cloud versions, see the prerequisites section of Create a Terraform Enterprise or Terraform Cloud config provider. Starting with Cloud Provisioning and Governance: Terraform Connector version 1.5.2, version 1.0.0 is deprecated. The deprecated version represents Terraform Open Source CLI version 0.11 and 0.12.

Cloud Services Catalog Terraform Connector SSH Private Key Credentials form reference

Cloud Services Catalog Terraform Connector uses the SSH private key credential to connect with the Linux server on which Terraform Open Source is running.

The Linux server supports the following authentication modes:

- Authentication with username and password.
- Authentication using the private key only.
- Authentication using the private key and SSH passphrase.

Field	Description
Name	Name of the credential.
Username	SSH username of the Linux server on which Terraform is running.
Password	SSH password of the Linux server on which Terraform is running.
SSH passphrase	SSH passphrase of the Linux server on which Terraform is running.
SSH private key	SSH private key of the Linux server on which Terraform is running.

Cloud Services Catalog Terraform Connector Windows Credentials form reference

Cloud Services Catalog Terraform Connector uses the Windows credential to connect with the Windows server on which Terraform Open Source is running.

Field	Description
Name	Name of the credential.
Username	Remote Desktop Protocol (RDP) username for the Windows server on which Terraform is running.
Password	RDP password for the Windows server on which Terraform is running.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Services Catalog Terraform Connector API Key Credentials form reference

The Cloud Services Catalog Terraform Connector uses API key credentials to authenticate the requests sent to the Terraform Enterprise and Terraform Cloud and the version control systems provisioned in them.

To learn more about API key permissions and tokens, see https://developer.hashicorp.com/terraform/cloud-docs/users-teams-organizations/api-tokens.

Field	Description
Name	Unique and descriptive name of the credential.
Active	Option for enabling this credential for discovery. Select this check box to enable discovery.

API Key	API key that can be used to access the account.
	• Terraform Enterprise or Terraform Cloud API key: Bearer token for accessing Terraform Enterprise or Terraform Cloud.
	Ensure that the API key adheres to the following syntax:
	`Bearer <API key>`
	• GitHub API key: API token for accessing the GitHub enterprise or community account.
	Ensure that the token adheres to the following syntax:
	`Token <API key>`
	• GitLab API key: Bearer token for accessing the GitLab enterprise or community account.
	Ensure that the token adheres to the following syntax:
	`Bearer <API key>`
	• Microsoft Azure DevOps API key: base-64 encoded token for accessing the Azure DevOps enterprise or community account.
	Important: Prefix the API token with a colon (:), and then encode the entire string to base-64.
	Ensure that the token adheres to the following syntax:
	`Basic <base-64 encoded API key>`

Applies to	MID Server to which you apply the credentials. Select **All MID servers**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Order	Order in which the platform runs this credential as it attempts to log in to devices. A smaller number indicates that the credential appears higher in the list. Establish the credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number, or none, the instance tries the credentials in a random order. The default value is 100.
Credential alias	Credential alias that can resolve the credential.

Cloud Services Catalog Terraform Connector Terraform Catalog Item Task form reference

Use the Terraform Catalog Item Task form to specify the action you want to take to resolve the catalog item change task.

Terraform Catalog Item Task form

Field	Description
Catalog Item Task Action	Action that you want to take. • Update catalog item with this template version Select this option to update the catalog item with this template version. A new version of the template is created and activated while the existing template version is retired and rendered inactive. • Create catalog item with this template Select this option to create a catalog item with the existing catalog item name and append the catalog item name with a number. For example, if your catalog item is named `Linux VM`, the new catalog is created with the name `Linux VM 1`. The catalog item is automatically activated and available for provisioning from the Cloud User portal.
Assignment Group	The assignment group to which you're assigning the catalog item task.
Assigned to	The user to whom you're assigning the task.
Short Description	Brief description of the task.
Approval	The approval status of the task.
Work Notes	Any notes that you want to enter for the task record.

Day 2 operations using Workflow Studio subflow

Take advantage of the flow designer to automate your Day 2 operations. Quickly write a subflow that communicates with a Cloud API or a particular resource.

Day 1 operations provision resources in a cloud as part of ordering from the ServiceNow catalog or the cloud catalog. These operations result in a stack that contains a list of provisioned resources. Day 2 operations can be carried out on the resources that are part of the stack or the resources discovered by the system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Previously, to execute Day 2 operations in Cloud Provisioning and Governance, you had to interact directly with the Cloud API (CAPI). This frequently meant that you had to write new APIs to support a new operation. You can now use subflows to map Day 2 operations and create REST endpoints that call cloud providers and trigger platform-specific actions.

Learn more

To learn about flows, subflows, and actions in general, see:

- Exploring flows
- Exploring subflows
- Exploring actions.

Credential handling

Extend the workflow engine to manage processes and automate things outside of an instance with Orchestration. Use the appropriate credentials required by Orchestration SSH and PowerShell activity elements: SSH requires SSH and PowerShell requires Windows.

The Cloud Provisioning and Governance application generates credentials (node credentials) after a VM is provisioned. Node credentials are used to connect to a VM. In addition to node credentials, the system generates SSH/Windows credentials. These credentials are required for Orchestration, which provides the automation for workflows.

Pre-existing VMs

Cloud Provisioning and Governance now generates SSH private key credentials and Windows credentials for all active pre-existing VMs provisioned through Cloud Provisioning and Governance.

Delete credentials

When you deprovision or delete a VM, all associated credentials are deleted. This includes the node credential, SSH private key credentials, Windows credentials, and the credentials alias tag.

Create a response action for Cloud Provisioning and Governance

Create a subflow with the required input parameters that notifies Cloud Provisioning and Governance whether the flow completes successfully or not.

Before you begin

- Ensure that you are familiar with the concepts of , Subflows, and Action Designer.
- Create a subflow that you intend to invoke for pre or post-provisioning.

Role required: flow_designer, action_designer, or admin

About this task

Inform the Cloud Provisioning and Governance application, the execution status of steps in the subflow, by populating a response for each step.

- Create a action with relevant input parameters.
- Add a script to the action.
- Pass a correlation id to the action in the subflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The response action communicates to the Cloud Provisioning and Governance application whether a step in the subflow, executes successfully or the step's condition fails and results in an error.

Procedure

1. Navigate to **All** > **Flow Designer** > **Actions** tab, select **New Action**.
   For more information, see Create an action.
2. Fill in the Action Properties.
3. Create and define action inputs to make data available to the action steps and outputs.
   Create the following input variables: error, log, output, and correlation id.

Sample response action to Cloud Provisioning and Governance

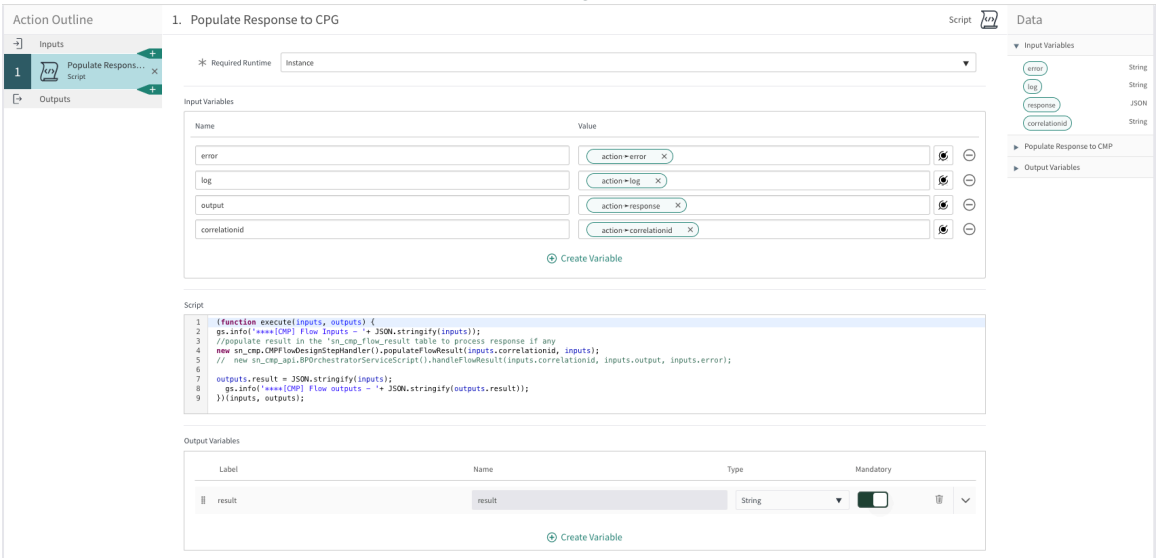

4. Select and add a script action step to perform an operation on the action inputs.
5. **Save** and **Publish** the action.
6. In the **Flow Designer**, open the subflow you are invoking.

> **Note:** Ensure that the subflow input `flowcorrelationid` exists in the subflow.

7. On the same subflow page, click **Actions** to expand and find **Actions**, **Subflow** or **Flow Logic**.
8. Select Cloud Provisioning and Governance and, search and add the response action from the drop down.
9. Select relevant values from the **Data Pill Picker**, for the inputs in the action.

> **Note:** Reuse the `flowcorrelationid` input from the **Data Panel** for every action in the subflow that has a logical end step.

10. **Optional:** If you are using the Terraform Connector for Terraform Enterprise, follow these steps.
    a. To add actions, flows, subflows, or flow logic, click **Select to add an Action, Flow Logic, or Subflow**.
    b. Select **Action** > **Cloud Management: Terraform Connector**, and select a default action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

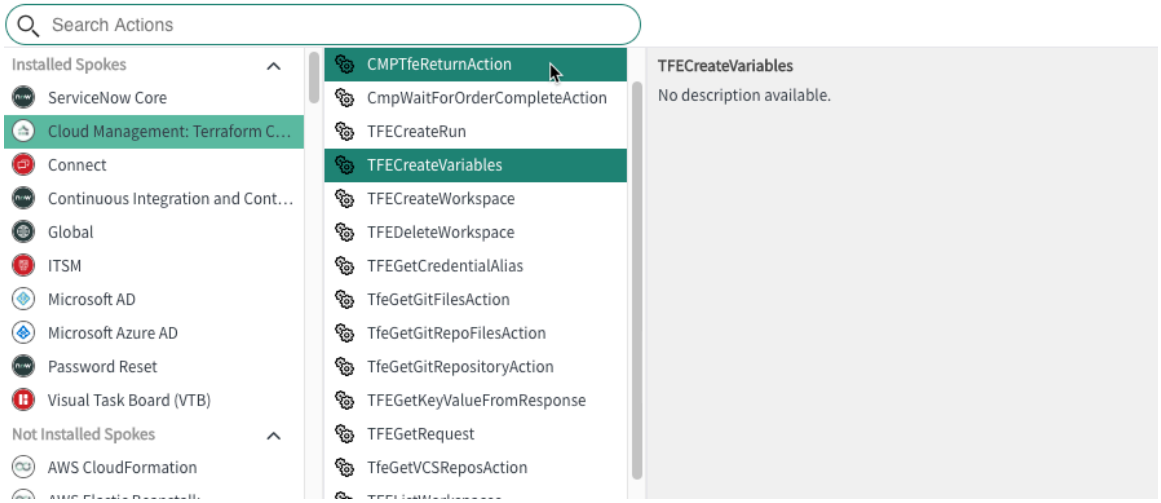

The action you created sends a response to the Cloud Provisioning and Governance application, updating the order with the status of the subflow's step.

What to do next

Test the subflow, and publish it when it is ready to be added to a flow or called from a script.

Note: You can only test or publish subflows that contain at least one action.

Create a cloud catalog item

Create a cloud catalog item for provisioning, based on a template or an existing blueprint, and publish the catalog item to provide a service.

Before you begin

Role required: sn_cmp.cloud_service_designer.

About this task

While creating a catalog item, you can associate it with a template native to the cloud provider (for example, CFT, ARM, GDM) or a Terraform IaC template. After you generate a catalog item, the system automatically generates components such as variable sets, lifecycle operations, rules, and pre- and post-provision operations. You can manage all these components from the Create Catalog Item page.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Catalog items**.
   All the catalog items that you've created are listed in a card layout view. You can click an inactive catalog item to provision it.
2. Click **New**.
   The Cloud Catalog Items screen appears.
3. On the form, fill in the fields.

Cloud Catalog Item form

Field	Description
Name	Name for the catalog item.
Active	Option for making the catalog item available. This check box can be selected only after you've selected a blueprint.
Icon	Image to appear as an icon beside the item name in the catalog. You can upload a 16

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	x16 pixel image. If no image is uploaded, the default icon appears beside this item. The uploaded image overwrites the default image. The image you upload appears in the service catalog; not in the ServiceNow® Cloud Provisioning and Governance application.
Picture	Image for the catalog item. The image that you select is displayed for the catalog item in the Cloud User Portal (after you click **Launch a Stack**).
Catalogs	Names of catalogs.
Category	Category for the item appears by default. You can change the category. Catalog searches find only items that are assigned to a category. The categories appear in the service catalog. The catalog item can appear in multiple categories. To add categories, click the **Categories** tab and then click **New** or **Edit**.
Source	Source type (Cloud Template or Custom) to associate the catalog item with.
Template type	○ Select **Cloud Template**, then select the type of template (**ARM Template** or **CloudFormation Template**, to import cloud-native templates) ○ Select **Configuration Management Template** to use Config Management provider templates from Anisible or Terraform.
Blueprint	If you selected **Custom** in the **Source** field, then select a blueprint to associate with the catalog item.
Blueprint Operation	The value in this field is automatically populated after you generate the catalog item. This field appears only when you select a blueprint.
Short description	Brief description of the catalog item. This description appears in the Cloud User Portal next to the catalog item.
Description	Full description of the catalog item. This description appears in the service catalog; not in the Cloud Provisioning and Governance application. You can embed videos, images, links to internal knowledge base (KB) articles, and links to external sources of information and instruction documentation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Click **Submit**.
   The catalog item and all associated components such as variable sets, lifecycle operations, and rules are generated and appear as tabs in the Cloud Catalog Item page.

   Note: If you selected **Template** from the Source type list, you must create a template and associate the template with the catalog item.

   The related links are actions that you can perform on the catalog item.

Action	Description
Export Catalog Item	Exports all the components (such as blueprints, resource blocks, rules) of the catalog item from one environment to the other. See Moving Cloud Provisioning and Governance content across environments.
Item Diagnostic	Identifies possible issues with the catalog item data and configuration. See Item diagnostic report
Manage Attributes	See Manage operation attributes.
Reset Catalog Item	Resets (to the last change you made) any customizations that you've made to the catalog item.

What to do next
Create a cloud template-based catalog item.

Related topics
Create a catalog item from a Terraform template

Publish a blueprint as a cloud catalog item

You can create and publish a catalog item directly from a blueprint.

Before you begin
Role required: sn_cmp.cloud_service_designer

- Starting with Orlando, Cloud Provisioning blueprints are available on instances upgraded from a previous release but you cannot create new blueprints. Resource profiles and custom-created blueprints will no longer be supported starting with the Australia release.
- Use Cloud Provisioning cloud templates to create catalog items in place of blueprints. Cloud Provisioning cloud templates allow you to ingest Azure ARM, AWS CFT, Google Deployment Manager (GDM) and Terraform specification syntax in cloud catalog items to run your cloud deployment orchestration.

Procedure
1. In the Cloud Admin Portal, navigate to **Design** > **Blueprints**.
2. Open a blueprint (in draft mode), click the **Catalog** tab, and then click **Create Catalog Item**.
   The Create Blueprint Catalog dialog box appears.
3. Enter a name for the catalog item in the **Name** field and click **Save**.
   The catalog item is generated and appears in the **Catalog** tab.
4. Click the catalog item to make any updates to the catalog item.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Cloud Catalog Item screen appears.

5. Click **Update** to save the changes you made.

Create a cloud template-based catalog item

Create a cloud template and associate the template with a catalog item. Once you've created a template, you can reuse the template to create additional catalog items for the services you want to provision.

Before you begin

Role required: sn_cmp.cloud_service_designer.

About this task

You can create a template by uploading a generic template file, importing a template from a public URL, or by copying and pasting the contents of an existing template. AWS CloudFormation (CFT) templates support JSON and YAML. Azure Resource Manager (ARM) templates support JSON. The following public template URL content types are supported:

- Text/Plain
- Application/JSON
- YAML MIME

Note: Catalog items that were based on templates in earlier releases are now treated as if they are based on blueprints.

You can update a template as often as needed. With each update, a new version of the template is created. A template can be in one of three states: **draft**, **active**, or **history**. When you create a template and do not activate it, the template is in a draft state. Once you activate the template, the state changes to active. If you create another template and activate it, the state of the previous active template turns into history. At any time, only one template is in the active state. Each time you activate a template, the previous active template becomes history.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Catalog Items** and then create a catalog item or open an existing catalog item.
2. Click **Cloud Templates** > **New** and then enter a meaningful short description in the **Short description** field.
3. Specify the **Ingestion method**.

Option	Description
Import from URL	Import a template by specifying a public URL where the template resides. Select this option and then click the lock icon ( ) to unlock the **Cloud template URL** field. Enter the pub lic URL in this field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	Note: For the public URL, we support only HTTP and HTTPS protocols, and do not support FTP. Ensure that the size of the template does not exceed the default value of 3 MB. You can change the default value of the file size by changing the value in the *sn_cmp.template_content_size_supported_* property in the sys_properties table. Enter the file size in bytes. There's also a default time-out of five minutes for an HTTP request. You can change the time-out value in the *sn_cmp.template_url_import_http_timeout* property in the sys_properties table. Enter the new value in milliseconds.
Upload a file	Upload a template from your local workstation.
Use template body	Paste the contents of the template file in the **Body** field.

After you specify the **Ingestion method**, the template is validated and the results appear in the **Validation status** and **Validation message** fields.

4. Click **Submit**.
   The template is created and the Cloud Catalog Item page opens.
5. If this is the first version of template, skip this step else on the **Cloud Templates** tab, open the template (currently in Draft state) to view the template version parameters.

   Note: The first version of a template that you create and publish has no conflict issues. If you update the template, a conflict can arise. If a parameter **Action Type** has the value **Update**, then specify a value for **Decision**.

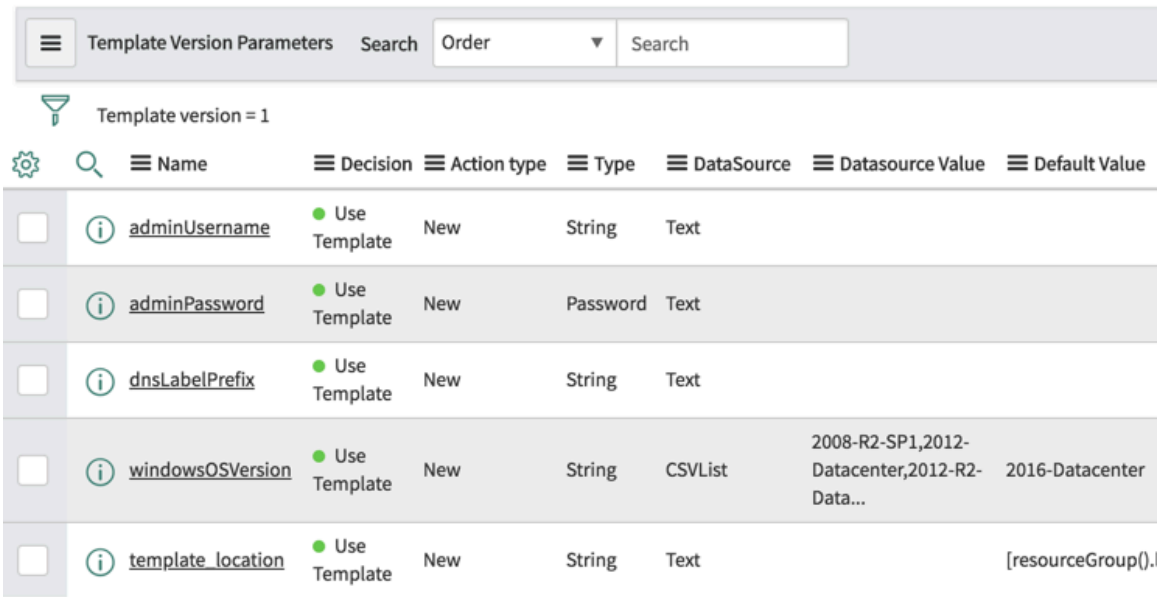

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Open each parameter with a **Decision** value of **Pending**, select a value, and then click the check mark .

Option	Description
Skip Update	Discards the updates to the template.
Use template	Applies the updates to the template.

7. After you update each **Pending** decision, click **Activate**.
   The state of the template changes from **Draft** to **Active** and the system generates the catalog item.

Sample CloudFormation JSON to create an AWS EC2 instance

Reference code sample to deploy an Amazon AWS Cloud EC2 instance.

Deploying an AWS EC2 instance using CloudFormation

- While designing a catalog item, review the cloud formation template to identify the services that are being provisioned. Check that appropriate permissions for the services can be granted, to ensure that no failures occur from lack of permissions when the stack is ordered. For more information, see Configure permissions on the AWS console
- Make necessary modifications using this sample as your reference.

```
{
   "AWSTemplateFormatVersion":"2010-09-09",
   "Description":"AWS CloudFormation Sample Template
EC2InstanceWithSecurityGroupSample: Create an Amazon EC2
instance running the Amazon Linux AMI. The AMI is chosen based
on the region in which the stack is run. This example creates
an EC2 security group for the instance to give you SSH access.
**WARNING** This template creates an Amazon EC2 instance. You
will be billed for the AWS resources used if you create a stack
from this template.",
   "Parameters":{
      "KeyName":{
         "Description":"Name of an existing EC2 KeyPair to
enable SSH access to the instance",
         "Type":"AWS::EC2::KeyPair::KeyName",
         "Default":"admin-Default",
         "ConstraintDescription":"must be the name of an
existing EC2 KeyPair."
      },
      "InstanceType":{
         "Description":"WebServer EC2 instance type",
         "Type":"String",
         "Default":"t2.micro",
         "AllowedValues":[
            "t1.micro",
            "t2.nano",
            "t2.micro",
            "t2.small",
            "t2.medium",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
            "t2.large",
            "m1.small",
            "m1.medium",
            "m1.large",
            "m1.xlarge",
            "m2.xlarge",
            "m2.2xlarge",
            "m2.4xlarge",
            "m3.medium",
            "m3.large",
            "m3.xlarge",
            "m3.2xlarge",
            "m4.large",
            "m4.xlarge",
            "m4.2xlarge",
            "m4.4xlarge",
            "m4.10xlarge",
            "c1.medium",
            "c1.xlarge",
            "c3.large",
            "c3.xlarge",
            "c3.2xlarge",
            "c3.4xlarge",
            "c3.8xlarge",
            "c4.large",
            "c4.xlarge",
            "c4.2xlarge",
            "c4.4xlarge",
            "c4.8xlarge",
            "g2.2xlarge",
            "g2.8xlarge",
            "r3.large",
            "r3.xlarge",
            "r3.2xlarge",
            "r3.4xlarge",
            "r3.8xlarge",
            "i2.xlarge",
            "i2.2xlarge",
            "i2.4xlarge",
            "i2.8xlarge",
            "d2.xlarge",
            "d2.2xlarge",
            "d2.4xlarge",
            "d2.8xlarge",
            "hi1.4xlarge",
            "hs1.8xlarge",
            "cr1.8xlarge",
            "cc2.8xlarge",
            "cg1.4xlarge"
         ],
         "ConstraintDescription":"must be a valid EC2 instance
type."
      },
      "SSHLocation":{
         "Description":"The IP address range that can be used to
SSH to the EC2 instances",
         "Type":"String",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
         "MinLength":"9",
         "MaxLength":"18",
         "Default":"0.0.0.0/0",

 "AllowedPattern":"(\\d{1,3})\\.(\\d{1,3})\\.(\\d{1,3})\\.(\
\d{1,3})/(\\d{1,2})",
         "ConstraintDescription":"must be a valid IP CIDR range
 of the form x.x.x.x/x."
      },
      "Network":{
         "Description":"The Network to the EC2 instances",
         "Type":"String"
      },
      "AvailabilityZone":{
         "Description":"The AZ of the EC2 instance",
         "Type":"String"
      },
      "SubnetPool":{
         "Description":"The IP address range that can be used to
 SSH to the EC2 instances",
         "Type":"String"
      }
   },
   "Metadata":{
      "SNC::Parameter::Metadata":{
         "SubnetPool":{

 "datasource":"ServiceNow::Pools::SubnetPool.getObjectsByNetwor
k",
            "datasourceFilter":{
               "Network":"Network"
            }
         },
         "Network":{

 "datasource":"ServiceNow::Pools::NetworkPool.getObjectsByLDC"
         },
         "KeyName":{

 "datasource":"ServiceNow::Pools::CloudKeyPairPool.getObjectsByL
DC"
         }
      }
   },
   "Mappings":{
      "AWSInstanceType2Arch":{
         "t1.micro":{
            "Arch":"PV64"
         },
         "t2.nano":{
            "Arch":"HVM64"
         },
         "t2.micro":{
            "Arch":"HVM64"
         },
         "t2.small":{
            "Arch":"HVM64"
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
        },
        "t2.medium":{
           "Arch":"HVM64"
        },
        "t2.large":{
           "Arch":"HVM64"
        },
        "m1.small":{
           "Arch":"PV64"
        },
        "m1.medium":{
           "Arch":"PV64"
        },
        "m1.large":{
           "Arch":"PV64"
        },
        "m1.xlarge":{
           "Arch":"PV64"
        },
        "m2.xlarge":{
           "Arch":"PV64"
        },
        "m2.2xlarge":{
           "Arch":"PV64"
        },
        "m2.4xlarge":{
           "Arch":"PV64"
        },
        "m3.medium":{
           "Arch":"HVM64"
        },
        "m3.large":{
           "Arch":"HVM64"
        },
        "m3.xlarge":{
           "Arch":"HVM64"
        },
        "m3.2xlarge":{
           "Arch":"HVM64"
        },
        "m4.large":{
           "Arch":"HVM64"
        },
        "m4.xlarge":{
           "Arch":"HVM64"
        },
        "m4.2xlarge":{
           "Arch":"HVM64"
        },
        "m4.4xlarge":{
           "Arch":"HVM64"
        },
        "m4.10xlarge":{
           "Arch":"HVM64"
        },
        "c1.medium":{
           "Arch":"PV64"
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
        },
        "c1.xlarge":{
           "Arch":"PV64"
        },
        "c3.large":{
           "Arch":"HVM64"
        },
        "c3.xlarge":{
           "Arch":"HVM64"
        },
        "c3.2xlarge":{
           "Arch":"HVM64"
        },
        "c3.4xlarge":{
           "Arch":"HVM64"
        },
        "c3.8xlarge":{
           "Arch":"HVM64"
        },
        "c4.large":{
           "Arch":"HVM64"
        },
        "c4.xlarge":{
           "Arch":"HVM64"
        },
        "c4.2xlarge":{
           "Arch":"HVM64"
        },
        "c4.4xlarge":{
           "Arch":"HVM64"
        },
        "c4.8xlarge":{
           "Arch":"HVM64"
        },
        "g2.2xlarge":{
           "Arch":"HVMG2"
        },
        "g2.8xlarge":{
           "Arch":"HVMG2"
        },
        "r3.large":{
           "Arch":"HVM64"
        },
        "r3.xlarge":{
           "Arch":"HVM64"
        },
        "r3.2xlarge":{
           "Arch":"HVM64"
        },
        "r3.4xlarge":{
           "Arch":"HVM64"
        },
        "r3.8xlarge":{
           "Arch":"HVM64"
        },
        "i2.xlarge":{
           "Arch":"HVM64"
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
        },
        "i2.2xlarge":{
           "Arch":"HVM64"
        },
        "i2.4xlarge":{
           "Arch":"HVM64"
        },
        "i2.8xlarge":{
           "Arch":"HVM64"
        },
        "d2.xlarge":{
           "Arch":"HVM64"
        },
        "d2.2xlarge":{
           "Arch":"HVM64"
        },
        "d2.4xlarge":{
           "Arch":"HVM64"
        },
        "d2.8xlarge":{
           "Arch":"HVM64"
        },
        "hi1.4xlarge":{
           "Arch":"HVM64"
        },
        "hs1.8xlarge":{
           "Arch":"HVM64"
        },
        "cr1.8xlarge":{
           "Arch":"HVM64"
        },
        "cc2.8xlarge":{
           "Arch":"HVM64"
        }
     },
     "AWSRegionArch2AMI":{
        "us-east-1":{
           "PV64":"ami-2a69aa47",
           "HVM64":"ami-6869aa05",
           "HVMG2":"ami-a41a3fb3"
        },
        "us-west-2":{
           "PV64":"ami-7f77b31f",
           "HVM64":"ami-7172b611",
           "HVMG2":"ami-caf253aa"
        },
        "us-west-1":{
           "PV64":"ami-a2490dc2",
           "HVM64":"ami-31490d51",
           "HVMG2":"ami-00347e60"
        },
        "eu-west-1":{
           "PV64":"ami-4cdd453f",
           "HVM64":"ami-f9dd458a",
           "HVMG2":"ami-e2f7bd91"
        },
        "eu-central-1":{
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
            "PV64":"ami-6527cf0a",
            "HVM64":"ami-ea26ce85",
            "HVMG2":"ami-d2ff04bd"
         },
         "ap-northeast-1":{
            "PV64":"ami-3e42b65f",
            "HVM64":"ami-374db956",
            "HVMG2":"ami-4c78d52d"
         },
         "ap-northeast-2":{
            "PV64":"NOT_SUPPORTED",
            "HVM64":"ami-2b408b45",
            "HVMG2":"NOT_SUPPORTED"
         },
         "ap-southeast-1":{
            "PV64":"ami-df9e4cbc",
            "HVM64":"ami-a59b49c6",
            "HVMG2":"ami-f3f95990"
         },
         "ap-southeast-2":{
            "PV64":"ami-63351d00",
            "HVM64":"ami-dc361ebf",
            "HVMG2":"ami-3a122e59"
         },
         "ap-south-1":{
            "PV64":"NOT_SUPPORTED",
            "HVM64":"ami-ffbdd790",
            "HVMG2":"ami-21a7d34e"
         },
         "us-east-2":{
            "PV64":"NOT_SUPPORTED",
            "HVM64":"ami-f6035893",
            "HVMG2":"NOT_SUPPORTED"
         },
         "sa-east-1":{
            "PV64":"ami-1ad34676",
            "HVM64":"ami-6dd04501",
            "HVMG2":"NOT_SUPPORTED"
         },
         "cn-north-1":{
            "PV64":"ami-77559f1a",
            "HVM64":"ami-8e6aa0e3",
            "HVMG2":"NOT_SUPPORTED"
         }
      }
   },
   "Resources":{
      "EC2Instance":{
         "Type":"AWS::EC2::Instance",
         "Properties":{
            "InstanceType":{
               "Ref":"InstanceType"
            },
            "KeyName":{
               "Ref":"KeyName"
            },
            "SubnetId":{
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                "Ref":"SubnetPool"
             },
             "ImageId":{
                "Fn::FindInMap":[
                   "AWSRegionArch2AMI",
                   {
                      "Ref":"AWS::Region"
                   },
                   {
                      "Fn::FindInMap":[
                         "AWSInstanceType2Arch",
                         {
                            "Ref":"InstanceType"
                         },
                         "Arch"
                      ]
                   }
                ]
             }
          }
       },
       "InstanceSecurityGroup":{
          "Type":"AWS::EC2::SecurityGroup",
          "Properties":{
             "GroupDescription":"Enable SSH access via port 22",
             "SecurityGroupIngress":[
                {
                   "IpProtocol":"tcp",
                   "FromPort":"22",
                   "ToPort":"22",
                   "CidrIp":{
                      "Ref":"SSHLocation"
                   }
                }
             ]
          }
       }
    },
    "Outputs":{
       "InstanceId":{
          "Description":"InstanceId of the newly created EC2
instance",
          "Value":{
             "Ref":"EC2Instance"
          }
       },
       "AZ":{
          "Description":"Availability Zone of the newly created
EC2 instance",
          "Value":{
             "Fn::GetAtt":[
                "EC2Instance",
                "AvailabilityZone"
             ]
          }
       }
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    }
}
```

Create pre-provisioning operations

Pre-provisioning operations streamline catalog item deployment by automating tasks like notifications, workflows, or scripts before provisioning. Configuring these operations eliminates manual steps and saves time. The system automatically executes specified actions when a provisioning request is submitted, such as notifying users about the impending item provisioning process.

Before you begin

Role required: sn_cmp.cloud_service_designer.

About this task

You can create pre-provisioning tasks only for catalog items that are based on templates.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Catalog Items**.
2. Open a catalog item, click the **Pre Provisioning Operation** tab, click **New**, and then fill in the form.

Field	Description
Step Type	Pre-provisioning operation. Based on the type of pre-provisioning operation you select, additional fields appear. ○ CAPI: ▪ CAPI Provider ▪ CAPI Method ▪ CAPI Version ○ InstanceScript: Script: Enter the name of the script. ○ Resource: ▪ Resource Interface ▪ Interface Operation ○ Flow: Select a subflow. All active s in the instance are listed. **Note:** For more information, see Subflows.
Order	Order of pre-provisioning operation. Enter the order that pre-provisioning operation should be invoked. Lower numbers are invoked first. For example, if you enter 2 for CAPI and 3 for InstanceScript, then CAPI is invoked first and then the InstanceScript gets invoked. If you do not enter any value in this field, the system

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	automatically enters a value in increments of 20.
Enabled	Option for enabling the pre-provision operation you want invoked.

3. Click **Submit**.
   Parameters for the pre-provisioning operations are created as operation attributes in the Manage Attributes page.
4. To add values for the pre-provision operation attributes, click **Manage Attributes**.
   The Manage Attributes screen appears.
5. Click **PreProvision**.
   All attributes for the pre-provision operation are listed.

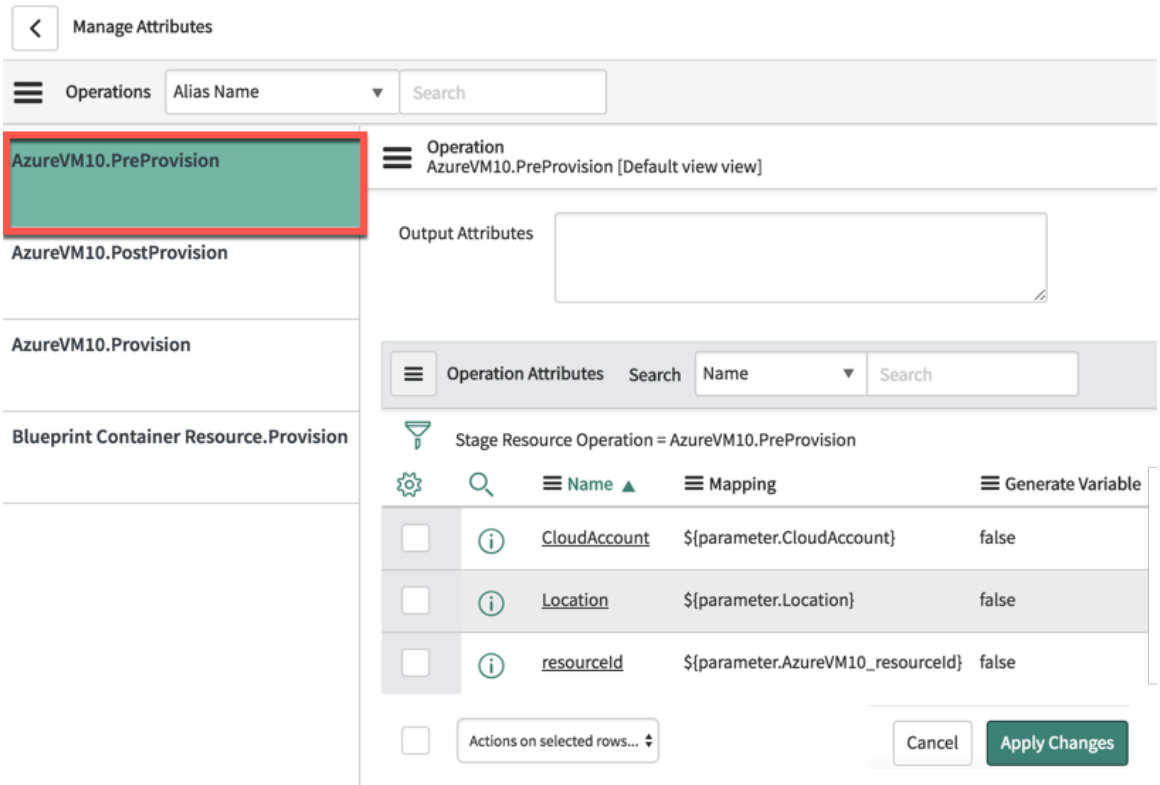

6. Either create expression mappings for an attribute (in the **Mapping** field) or generate variables for an attribute.
   See Manage operation attributes for more information.
7. Click **Apply Changes**.
   All the configurations you made to the attributes are saved and you are returned to the Cloud Catalog Items screen.

Create post-provision operations

Create post-provision tasks for newly provisioned catalog items by defining subflows, scripts, CAPI calls, or resource operations. For efficiency, set up a post-provision operation, eliminating the need for manual invocation. This approach enhances time management, ensuring seamless execution of desired actions on the catalog item after resource provisioning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_cmp.cloud_service_designer

About this task
You can create post-provision operations only for catalog items that are based on templates.

Procedure
1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Catalog Items**.
2. Open the catalog item, click the **Post Provisioning Operation** tab, and then click **New**.
3. On the form, fill in the fields.

Operation Setup form

Field	Description
Step Type	Post-provision operation. Based on the post-operation you select, additional fields appear. ○ CAPI: ▪ CAPI Provider ▪ CAPI Interface ▪ CAPI Method ▪ CAPI Version ○ InstanceScript: Script: Enter the name of the script. ○ Resource: ▪ Resource Interface ▪ Interface Operation ○ Flow: Select a subflow. All active subflows in the instance are listed. **Note:** For more information, see Subflows.
Order	Order of post-provision operations. Enter the order (in numerals) in which you want the post-provision operation to be invoked. For example, if you enter the value 2 for CAPI and the value 3 for InstanceScript, then CAPI is invoked first and then the InstanceScript gets invoked. If you do not enter any value in this field, the system automatically enters a value in increments of 20.
Enabled	Option for enabling the post-provision operation you want invoked.

4. Click **Submit**.
   Parameters for the post-provision operation are automatically created as operation attributes on the Manage Attributes form.
5. To set attribute values for a post-provision operation:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Click the **Manage Attributes** related link.

The Manage Attributes form opens.

b. Select the post-provision operation to update.

All attributes for the selected operation appear in the Operations list. You can either create expression mappings (shown in the **Mapping** column) for an attribute or can you generate variables for an attribute. See Manage operation attributes for more information.

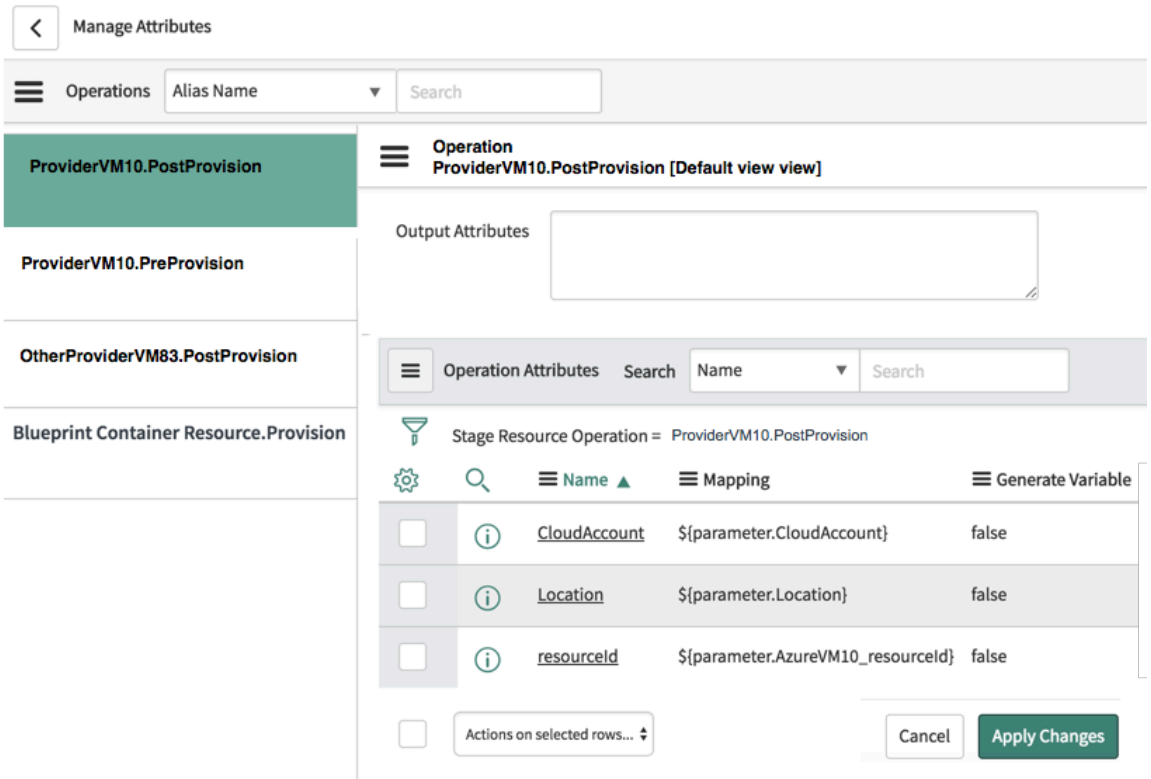

c. On the Operation Attributes form, edit the values in the **Mapping** field as needed.

d. To set a value and not show the field to the cloud user, set **Generate Variable** to false.

For VMware post-provision operations, you can set the following attribute values:

Post-provision operation attributes for VMware

Attribute name	<Value field> to set in the Mapping expression
NodeCredential	$(capiResolver.NodeCredentialResolver#nodeCredentialId=<SysID>) **Note:** You must create the credentials manually to be able to specify the sys_id in the expression
NodeAddress	$(Script:CMPVMUtils.getReachableIp[arg=$(Stack.items[<VM Name>].attributes[sys_id])])

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attribute name	<Value field> to set in the Mapping expression
ScriptExpression	${CloudScript.scripts.<Script Name>}

e. Click **Apply Changes**.
All configurations are saved and the Cloud Catalog Items page opens.

Add pre- or post-provisioning operations to a catalog item

Create a post-provisioning operation using subflows on a template-based catalog item.

Before you begin

Verify that the Orchestration application is installed.

Role required: Cloud user, designer, admin, sn_cmp.cloud_service_designer

Procedure

1. In the Cloud Admin portal, navigate to **Design** > **Cloud Catalog Items**.
2. Open a catalog item.
3. Select the **Pre Provisioning Operation** or **Post Provisioning Operation** tab.
4. Select **New**.
5. From the **Step Type** list, select **Flow**.
6. Select the Search icon on the **Flow** field.
7. From the **Subflow** look-up list, select an active subflow.
The look-up table lists all active flows and subflows in the instance.

   Note: You cannot invoke flows while adding a resource operation step. This is a known limitation.

8. In the **Order** field, specify an order for this operation.
9. Enable the operations by selecting the **Enabled** check box.
10. Select **Submit**.
The subflow operation step is attached to the catalog item and appears on the Pre-Provision Operation or Post-provision Operation related list.
11. Select the subflow operation step from the relevant related list to add or modify input parameters.

    Note:

    When creating a post-provision action, include **CMPTfeReturnAction** as the final step in the flow.

    Verify the following steps:

    - Verify that control returns to the main order so that data and responses flow back correctly and the process closes successfully.
    - Configure parameters to indicate whether the step resulted in an error.
    - Send an error from this action if the step is designed to fail.
    - Include the correlation ID. The flow will not function without it.

    Any input parameters associated with the subflow you selected are auto-populated on the Key Values related list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

- Verify that the `flowcorrelationid` input parameter is created in the subflow with the following expression string:

Key	Value
`flowcorrelationid`	`$(Script:CMPFlowStepHandler.generateCorrelationId)`

- To check whether the flow completes successfully, see Create a response action for Cloud Provisioning and Governance.

Make catalog item visible

You can use the catalog items created in cloud management, in standard service portal.

Before you begin

Role required: admin

Procedure

1. Navigate to a cloud catalog item.
2. Add **Service Catalog** or your specific catalog to the **Catalog item**.
3. Save it.
4. Select **categories** from Related list items.
5. Edit the categories and add a category from **service catalog** or your specific catalog.
6. Visit the service catalog page or your specific catalog page.
   The new catalog item shows up.

Create a variable set for Cloud Provisioning and Governance

Reduce the number of steps required to create and manage multiple catalog items and order guides by creating variable sets. With variable sets, you don't have to create variables individually for each catalog item.

Before you begin

Role required: sn_cmp.cloud_service_designer

About this task

Create a collection of structured variables that can be reused across multiple catalog items and order guides using variable sets.

Add a new variable set when you need specific request form variables for running actions such as post-provisioning. Variable sets are also useful for collecting additional information such as custom tags.

You can modify the variable set and the changes are reflected across all the catalog items that are associated with the variable set. Variable sets also allow you to define catalog client scripts and UI policies that are applicable to the variables in the set.

When you generate a catalog item, two variable sets, **General Info** and **Provision**, are automatically created. Create as many variable sets as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open a cloud catalog item and in the **Variable Sets** tab, click **New**.
   The Variable Set screen appears.
2. Fill in the form details (as shown in the table).

Field	Description
Title	Unique name for the variable set
Internal Name	A (unique) variable set name for internal use. **Note:** If this field is empty, its value is auto-populated based on the **Title** field for all variable types except break, container split, and container end.
Layout	Layout display selections from: **1 Column Wide**, **2 Columns Wide alternating sides**, or **2 Columns Wide, one side, then the other**.
Order	Order number for the variable set.
Description	Description for the variable set.

3. Click **Submit**.
   The new variable set is created and appears in the **Variable Sets** tab. Open the new variable set and click **New** in the **Cloud Variables** tab to create variables for the set.

Manage operation attributes

You can specify attribute values for catalog items help cloud users to correctly provision catalog items. You can configure and manage attributes for provision, pre-provision, and post-provision operations.

Before you begin

Role required: sn_cmp.cloud_service_designer.

About this task

The attributes for provision, pre-provision, and post-provision operations are generated automatically. You can generate variables for these attributes and also create new attributes for the resource operation. You can also use the Pre-Provision Operation, Post-Provision Operation or Lifecycle Operations related lists to create variables for Orchestration attributes.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Catalog Items**.
2. Open a catalog item and click the Pre-Provision Operation or Post-Provision Operation related list.
   The Operations Step page lists the operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

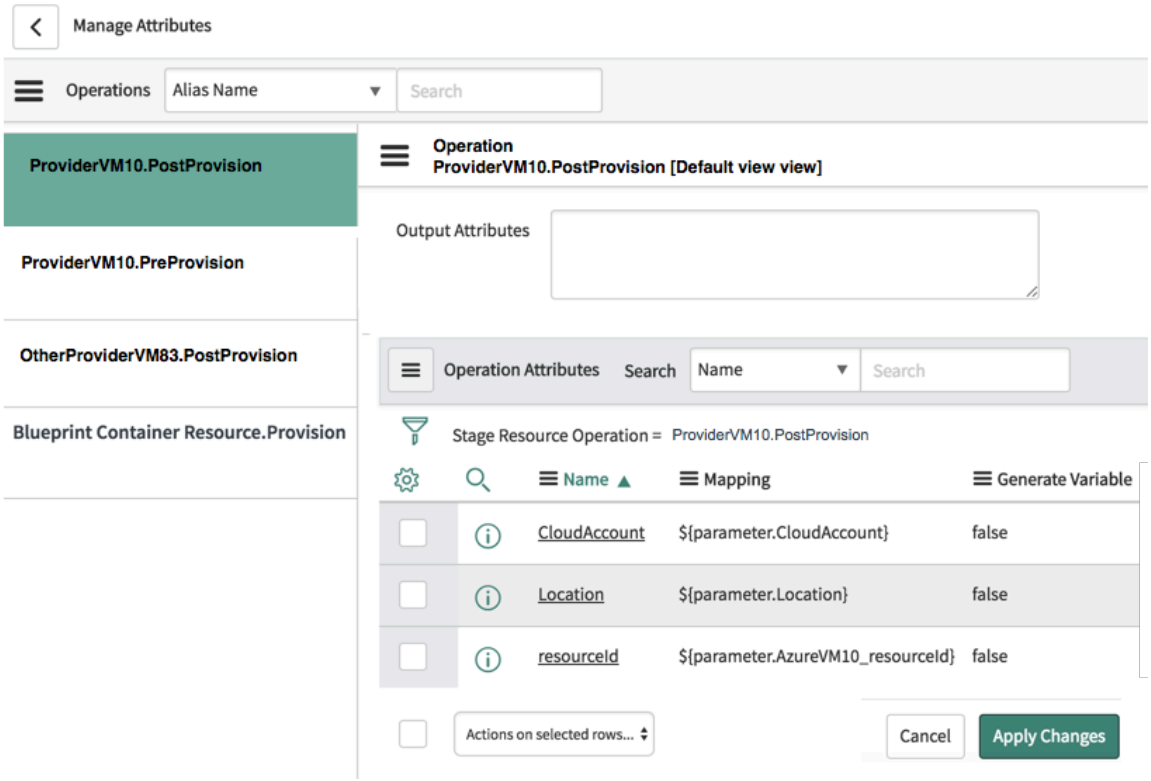

3. Click an operation to view its corresponding attributes.
4. Configure the attributes for each operation.
   - **Mapping**: View the auto-generated mapping values for each attribute. Mappings specify where to pull the information in the system. To modify the expression mappings, double-click the mapping value and enter new values. Then click the check mark icon (✓). Expressions can hold the values for information that's dynamically generated in the system, such as the values for the selections that users make in the Cloud User Portal. See Using expressions in Cloud Provisioning and Governance.
   - **Generate Variable**: Generates a variable for an attribute. Double-click the value (*false*), select **True**, and click the check mark icon (✓).
5. Click **Apply Changes** after you've made all the required changes.
   All the configurations you made to the attributes are saved and you're returned to the Cloud Catalog Items screen.

Cloud Provisioning Blueprints

A blueprint is a specialized catalog item template for offering cloud services, or stacks, to cloud users. Blueprints work with any cloud service provider, such as Amazon AWS Cloud or Microsoft Azure Cloud. Blueprints are in restricted usage from the Orlando release.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Blueprint components

Note:

- Starting with Orlando, Cloud Provisioning blueprints are available on instances upgraded from a previous release but you cannot create new blueprints. Resource profiles and custom-created blueprints will no longer be supported starting with the Australia release.
- Use Cloud Provisioning cloud templates to create catalog items in place of blueprints. Cloud Provisioning cloud templates allow you to ingest Azure ARM, AWS CFT, Google Deployment Manager (GDM) and Terraform specification syntax in cloud catalog items to run your cloud deployment orchestration.

The blueprint is the mechanism that is passed to the cloud provider to establish the cloud assets. Requesting a cloud catalog item is called launching a stack. After the cloud asset is obtained, it can then be discovered, managed, and maintained.

Blueprint deployment model components

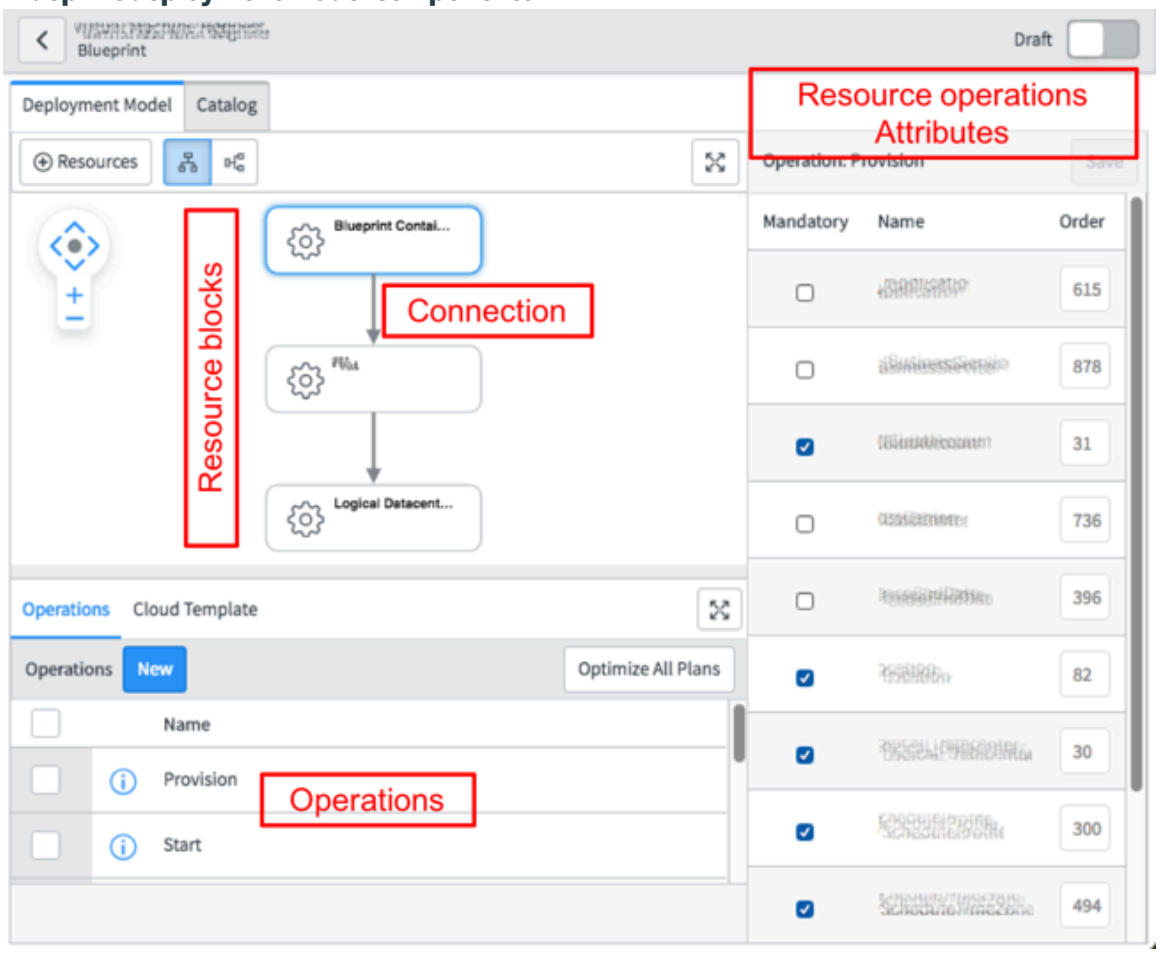

Resource blocks

The cloud resources included in the blueprint (catalog item). For example, a resource block can contain a virtual machine, an app that runs on the VM, and information about the datacenter that hosts the VM. Also referred to as resources.

Connections

The relationship between resources. For example, the VM has a host relationship with the app. There are two types of connections, containment and attached to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Containment: a connection where resource block contains another, visually depicted as a top-down relationship as shown in the image. For example, the blueprint container must contain at least one other resource block.
- Attached to: a connection between resources that require binding for direct communication to each other, visually depicted as a side-by-side relationship. For example, the binding between a virtual machine and storage forms an attached to relationship.

Operations

The actions that are performed on resource blocks or on the resource stack to manage the resources. For example, the cloud user can provision, start, and stop a resource.

Resource operations

The steps to perform the operation on the resource. For example, to provision the catalog item, the system must provision the VM and the app that runs on the VM. Click an operation to drill into the resource operations.

Attributes

The details of the resource. Attributes can appear on the catalog item form as catalog properties. For example, the blueprint can provide the version of the app to provision to the VM.

Constraints

An attribute option to constrain a cloud-agnostic resource block, such as a logical datacenter, to be used with a specific cloud provider. For example, a logical datacenter can be constrained to Amazon AWS Cloud. Later, if you switch or add Microsoft Azure Cloud or another cloud provider, you can change the constraint as needed.

Planning blueprints

Before creating blueprints, plan the stacks you want to offer in the service catalog. Prepare the following information before you begin development.

Resource blocks

Plan the resource blocks you want to use in blueprints. Ensure that the appropriate resource blocks are available to use in blueprints, and add them if necessary.

Service offerings (stacks)

To plan a service offering or stack, diagram the resource blocks that it needs, and how they are connected.

Cloud accounts and datacenters

Obtain a list of the cloud accounts and datacenters that the cloud administrator has configured. Determine whether any are constrained.

Service catalog variables and variable sets

Ensure that you have an understanding of how service catalog items use variables and variable sets. Understanding service catalog items and processes help with your understanding of blueprints.

Start here

Build a Cloud Provisioning blueprint

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Build a Cloud Provisioning blueprint

Use the blueprint designer to create custom blueprints for the Cloud Provisioning and Governance application. Blueprints are deprecated in the Orlando release.

Before you begin

Role required: sn_cmp.cloud_service_designer

Note:

- Starting with Orlando, Cloud Provisioning blueprints are available on instances upgraded from a previous release but you cannot create new blueprints. Resource profiles and custom-created blueprints will no longer be supported starting with the Australia release.
- Use Cloud Provisioning cloud templates to create catalog items in place of blueprints. Cloud Provisioning cloud templates allow you to ingest Azure ARM, AWS CFT, Google Deployment Manager (GDM) and Terraform specification syntax in cloud catalog items to run your cloud deployment orchestration.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Blueprints**.
2. Click **New**, enter a unique and meaningful **Name** and **Description**, and then specify the service **Category**.
   Several categories are provided by default.
3. Click **Submit**.
   The system creates the blueprint record and displays the Deployment Model tab of the blueprint designer. The Deployment Model tab is a visual canvas where you drop resource blocks. Each resource block represents a CI in a multi-tier application stack. The Deployment Model specifies the resources and connections between resources for the cloud items that you are defining.

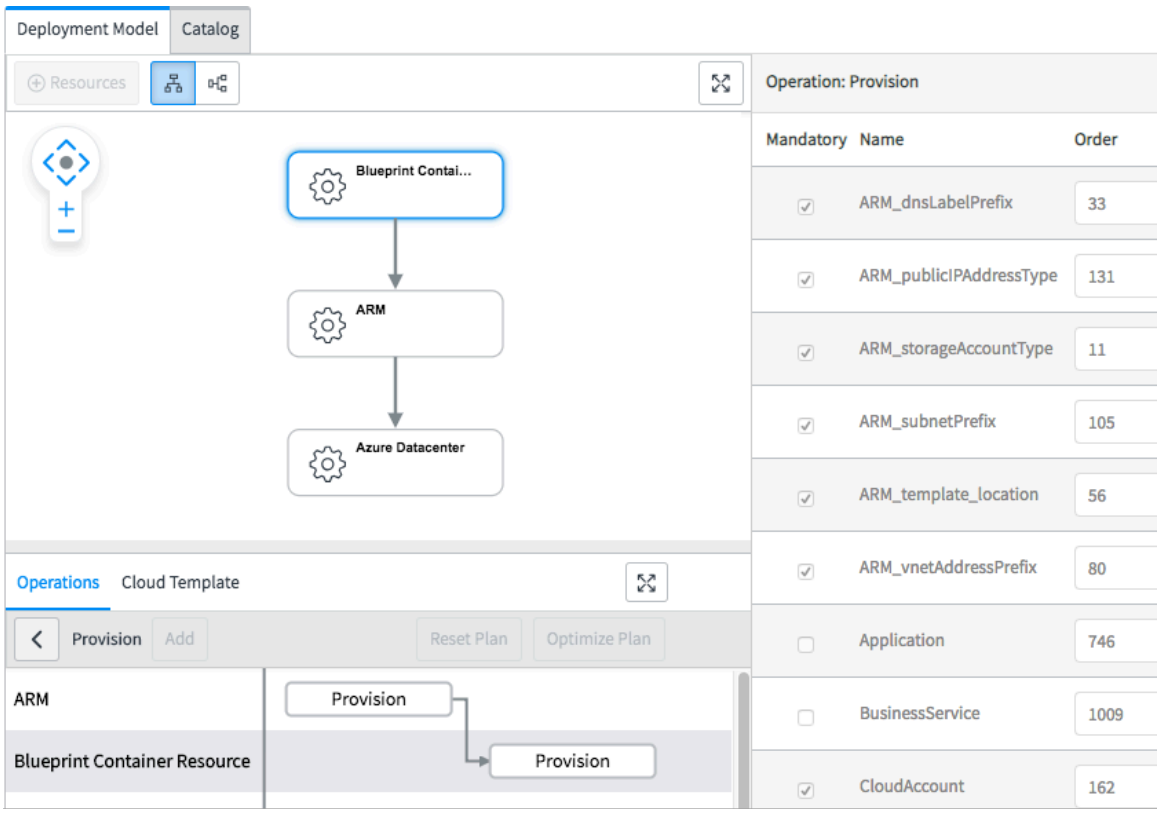

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. To add a resource to the deployment model, click **Resources** and then select the resource from the list.
   The resources appear on the canvas. Move the resource to an appropriate location for the model that you have in mind. Add as many resources as needed.
5. Connect the container resource to the first resource to deploy.
   To create a connection between resources, select the parent resource, click the connection pad in the resource, and then drop the connection in any resource with a dotted outline. The system does not allow certain connections (for example, an AWS datacenter cannot host an ARM resource).
6. Add and connect all resources that make up the blueprint.
   You must add and connect all resources before you configure operations and attributes in the **Operations** section. As you update the resources in the Deployment Model tab, the system generates the Cloud Template: a complete representation of the blueprint in YAML script. You can export the Cloud Template representation of the deployment model and later import it into the Cloud Provisioning and Governance application on another instance.
7. Configure the operations that the system performs on the resources and the attributes to use while performing the operations.
   By default, **Provision**, **Stop**, **Start** and **Deprovision** operations are provided. You can either select an existing operation or create a new blueprint operation.
   - Select an existing operation: In the **Operations** section, select the **Provision** operation check box and click **Provision** to view a graph. You can rearrange the graph by deleting and adding operations. You can click **Optimize All Plans** for the Cloud Provisioning and Governance application to suggest an optimized plan for you. The Cloud Provisioning and Governance application looks at the blueprint you created and based on the attributes, the mappings, and the dependencies of each resource, suggests an optimized plan to execute. All the attributes for the **Provision** operation appear in the window. You can select or de-select any attribute. Once you configured the **Provision** operation, you can select the next operation and configure its attributes as needed.
   - Create a new blueprint operation: Click **New** in the **Operations** section. The New Blueprint Operation dialog box appears. Enter a name, select an operation type, enter input parameters for the new operation, and click **Save**. The new blueprint operation appears in the **Operations** section.
8. If you have completed the design of the blueprint and you do not want to configure the blueprint any further, click the **Draft** toggle button to publish the blueprint.
   A message appears saying that your blueprint has been published. To continue configuring the blueprint, then skip this step.
9. To configure parameters such as variable sets, rules, life cycle operations for the blueprint catalog item, click the **Catalog** tab and then click the catalog.
   The Cloud Catalog Item screen appears.
10. Click **Update** if you made any changes to the form.
11. Click **Publish** to publish the blueprint.
    To view all published catalog items, click **Published** in the Blueprints landing page.

Add a resource operation to the execution plan

The base system includes a large set of predefined resource operations. Rarely, you might add a custom operation for a resource.

Before you begin

Role required: sn_cmp.cloud_service_designer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. While working in the blueprint designer, click the **Operations** tab and select the appropriate blueprint operation.
   As an example, you may select the **Provision** operation.
2. Click **Add**.
   The Add Resource Operation dialog box appears.
3. Fill in the form fields (as shown in the table).

Resource	Select a resource to perform the operation on.
Interface	Select an interface for the resource.
Operation	Select an operation to perform for the step. By default, **Provision**, **Stop**, **Start** and **Deprovision** are provided. You at least need a provision operation for a user to order a cloud catalog item.

4. Click **Submit**.
   The new operation is added to the operation graph. The system adds the appropriate parameters for the operation.
5. Click the **Create Form Parameter** check box to select a parameter.
   Parameters differ for each operation. Enter a mapping expression for the selected parameter in the **Mapping** field. See Using expressions in Cloud Provisioning and Governance.
6. Click **Save**.

Execute blueprint operations in parallel

You can execute your blueprint resource operations to run in a linear or in a parallel way.

Before you begin

Role required: sn.cmp.cloud_service_designer.

About this task

By default, the blueprint resource operations are listed serially and follow an order. You can run operations in parallel by arranging them in branches. Identify the operations specified in a blueprint that have no dependencies and can run in parallel. Arrange such operations in a single branch. The Blueprint Designer runs these operations together. If an error occurs in a branch, the operations in the affected branch cease to execute. However, the operations in the other branches, that are unaffected by the error, continue to execute. The parallel way of executing operations hence cuts down time in provisioning stacks.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Blueprints**.
2. Click a blueprint, click the **Deployment Model** tab and then click **Operations**.
   All the operations for the blueprint are listed.
3. Click an operation to view a graph.
   You can rearrange the graph by deleting and adding operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Click **Optimize Plan** for the Cloud Provisioning and Governance application to suggest an optimized plan for you.
   The Cloud Provisioning and Governance application looks at the blueprint that you created. Based on the attributes, mappings, and dependencies of each resource, the application suggests an optimized plan to execute.

Blueprint attributes

When you add a blueprint operation to a resource, the system adds the appropriate attributes. You can choose whether or not to make the attributes available as catalog properties on the form when the system performs the operation.

Common attributes for the provisioning operation for virtual machines

Attribute	Description
VmFolderName	The name of the virtual machine folder.
ResourceGroupName	The name of the resource group to which the resource belongs.
StackName	The name of the stack that the user is provisioning.
CloudAccount	The cloud account to which the resource stack belongs.
OSProfile	The resource profile, in this case the Operating System profile, upon which the cloud resource is based.
DiscoverAndAttachIP	Finds a free elastic IP and assigns it to a virtual machine on AWS. By default, if a VM is deployed on a VPC in AWS, it will not get a public IP. If elastic IPs are not available, no IP can be assigned to the VM.

⚠ **Warning:** It is recommended that you do not use *`KeyPreference`*. If you use the attribute, do not set *`KeyPreference`* to `user`.

Configure a blueprint rule

Configure a rule that contains actions the system can take on blueprint attributes.

Before you begin

Role required: sn_cmp.cloud_service_designer

Procedure

1. Open a blueprint that has the attribute for which you want to create a rule.
2. On the **Catalog** tab, click the operation, such as the provision operation, that contains the attribute.

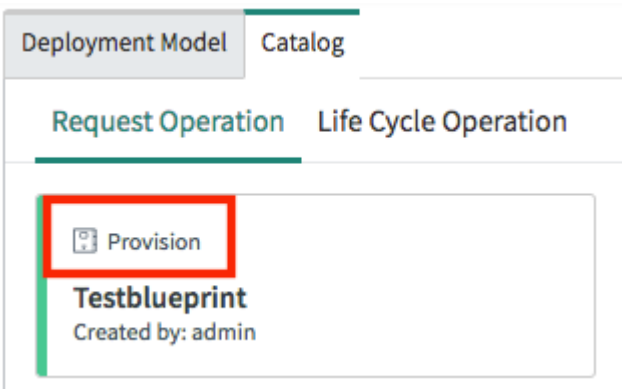

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the Variable Sets related list on the Cloud Catalog Item form, click the name of the variable set that contains the variable.
4. In the Cloud Variables related list on the Variable Set form, click the name of the variable for which you want to crate the rule.
5. In the Rules related list, click **New** and then fill in the form.

Field	Description
Name Description	Enter a name and description.
Order	Enter a number that tells the system which record to evaluate first in relation to others. The record with the lowest number runs first.

6. Right-click the form header and select **Save**.

Example:
This example rule is based on the *CloudAccount* property and contains an action that tells the system to update the list of locations when the user changes the cloud account.

Example blueprint rule

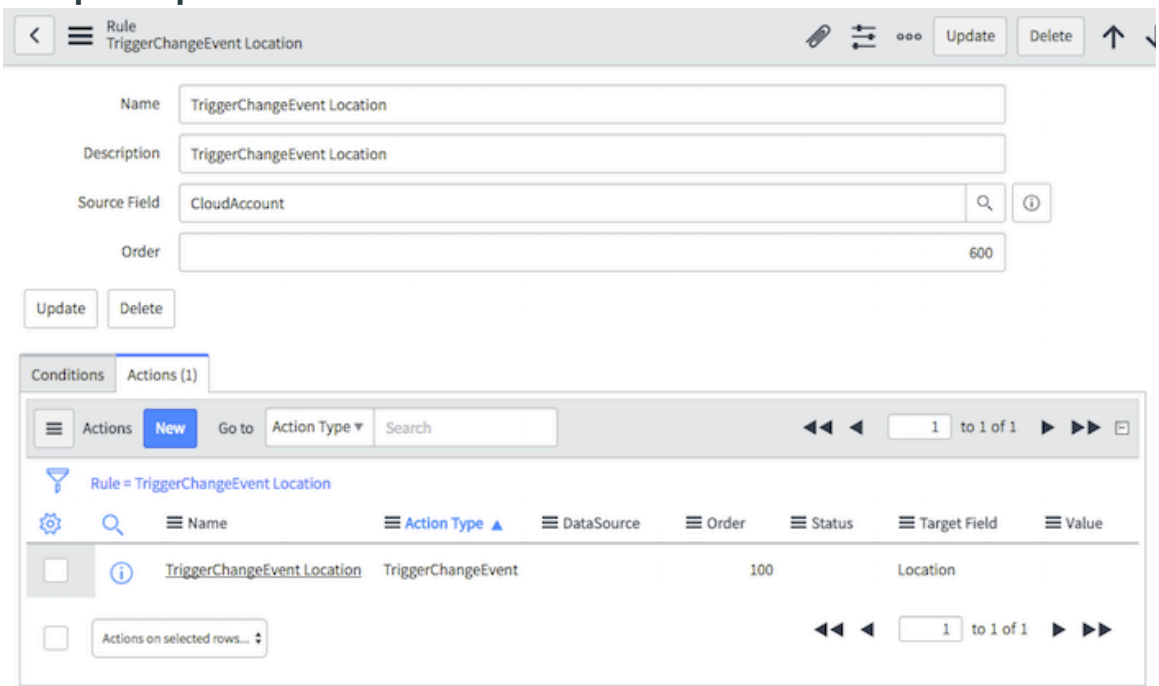

What to do next
Configure an optional condition and a mandatory action.

Configure a blueprint rule condition

Configure the condition on a blueprint form event rule that determines when the actions in the rule can run. Conditions are optional.

Before you begin
Role required: sn_cmp.cloud_service_designer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open a rule for a variable on a blueprint.
2. On the Conditions related list, click **New**.
3. Fill out the form fields (see table).

Field	Description
Operand	Enter the property against which you can perform the condiontal evaluation. For example, if your blueprint has a property which holds the value of a datacenter, enter the property name here.
Operator	Select an operator that evaluates the operand.
Value	Enter a value that the evaluated event rule must match.
Condition Type	If you configure multiple conditions, select the logical operator that determines how the system evaluates the condition in relation to the next condition (the next condition is determined by the **Order** field): ○ **AND**: this condition the next condition must meet the criteria before the policy is eligible for execution. ○ **OR**: this condition or the next condition that matches the criteria can allow the policy to be eligible for execution. If you do not select an option, **AND** is used by default.
Order	Enter a number that tells the system which record to evaluate first in relation to others. The record with the lowest number runs first.
Status	Keep the status active to have the system process the condition. Change the status to inactive when you do not want the system to process the condition.

4. Click **Submit**.

Example:

This example condition applies when the user selects the AWS Datacenter for the Logical_Datacenter property, which is a field on the catalog item form:

Example blueprint condition

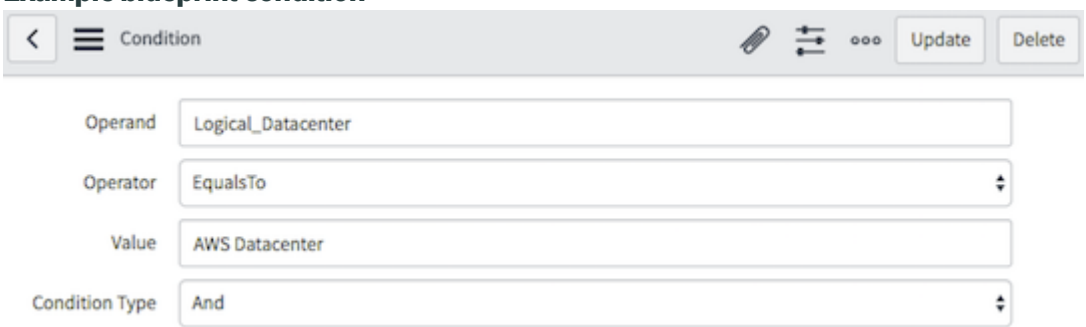

What to do next

Configure a blueprint action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure a blueprint action

Configure actions that the system runs when blueprint form attribute actions or event actions are triggered. You can configure actions like making a field mandatory, showing or hiding a field, or automatically populating a field.

Before you begin

Role required: sn_cmp.cloud_service_designer

Procedure

1. Open a rule for a property or an event on a blueprint.
2. On the **Actions** related list, click **New** and then fill in the form.

Field	Description
Name	Unique and descriptive name.
Action Type	Select an action type: ○ **Reload**: Select a variable to reload ○ **Show**: Select a variable to make visible. ○ **Hide**: Select a variable to hide. ○ **MakeMandatory**: Select a variable to make mandatory. ○ **MakeNonMandatory**: Select a variable to make optional. ○ **TriggerChangeEvent**: Select a variable on which to trigger a change event. All rules on the selected variable are triggered. ○ **SetValue**: Set the value of a variable.
Datasource	Select an action type: ○ **Reload Cloud Account**: ○ **Reload Cloud Location**: ○ **Check Quota**: ○ **Reload Resource Profile**: ○ **Reload From Profile**:
Target Field	Enter the name of the field on the form that the action applies to. For example, if you selected **TriggerChangeEvent** for the action type, the value of the field changes when something changes on the form, as specified in the blueprint rule.
Value	Enter an optional value that you want to assign to the target field. For example, if you selected **SetValue** for the **Action Type** field and you want to update the **Resource Group Name**, you can use the value `${ResourceGroup}`, which pulls the names of the resource groups from the blueprint.
Order	Enter a number that tells the system which record to evaluate first in relation to others. The record with the lowest number runs first.

3. Right-click the header and select **Save**.
4. **Optional:** Enter a **Key** and **Value** in the KeyValues related list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The value in the **Value** field is applied to the attribute in the **Key** field.

5. Create as many actions as necessary.

Example:
This example action applies to the blueprint rule example, which is based on the CloudAccount attribute. The action tells the system to update the list of locations when the user changes the cloud account.

Example blueprint action

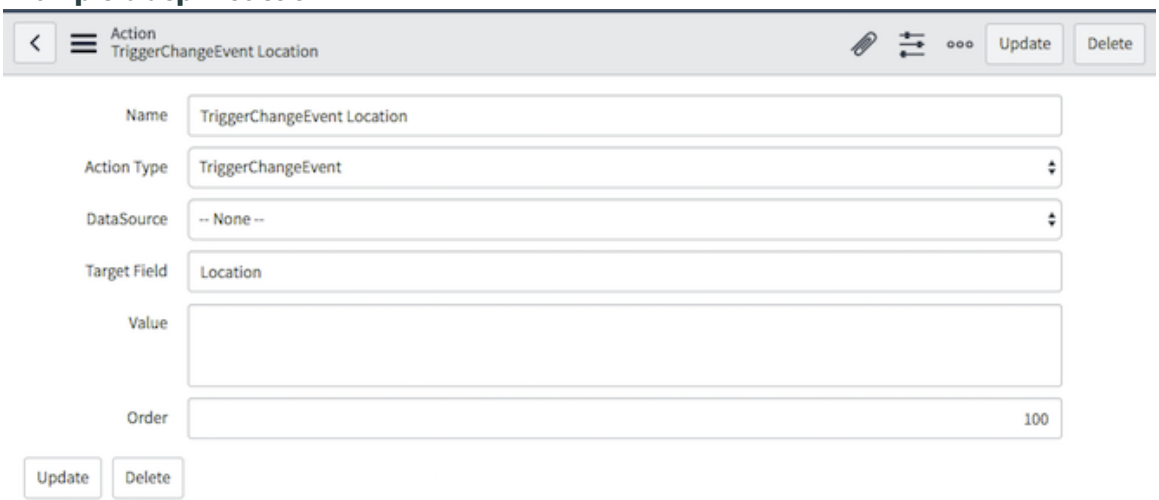

What to do next
Publish a blueprint as a cloud catalog item.

Form configuration with blueprints

Blueprints control what the user sees on the catalog order form. You can configure the form to show or hide fields, populate fields with default values, and otherwise create a unique form for the user based on criteria.

Note:

- Starting with Orlando, Cloud Provisioning blueprints are available on instances upgraded from a previous release but you cannot create new blueprints. Resource profiles and custom-created blueprints will no longer be supported starting with the Australia release.
- Use Cloud Provisioning cloud templates to create catalog items in place of blueprints. Cloud Provisioning cloud templates allow you to ingest Azure ARM, AWS CFT, Google Deployment Manager (GDM) and Terraform specification syntax in cloud catalog items to run your cloud deployment orchestration.

You can also change the values in form fields using policies. See Create a cloud policy and Create a policy action script for examples.

Components of blueprint form configuration

Blueprints use these components to control how the system uses catalog form fields:

- Catalog variables define the basic structure of each field on catalog order forms.
- Rules specify an action that the system takes on a field when a condition is met.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Scripts query the system for information that you need to show the catalog requester, populate form fields, and process the data that users submit.
- UI groups to put related fields together in a tab. Users can progress through tabs to fill out all the form fields.

Start here

These tasks show you the basics of form configuration with blueprints:

- Add fields to a cloud catalog item order form
- Control visibility default values for Cloud catalog items
- Create a variable set for Cloud Provisioning and Governance

Examples

These examples guide you through form configuration based on a specific goal: Cloud catalog form configuration examples.

Add fields to a cloud catalog item order form

You can add fields to a cloud catalog form if you want the user to enter additional information beyond what a default blueprint provides.

Before you begin

- You need a blueprint that you published to the catalog. The blueprint should be in the draft state. See Build a Cloud Provisioning blueprint if you did not complete this prerequisite.
- By default, cloud items appear in the Cloud Catalog, rather than the standard service catalog that comes with your instance. However, Blueprints utilize several features from the standard service catalog. It is suggested that you familiarize yourself with how the service catalog works. See Service Catalog to get started.
- Role required: sn_cmp.cloud_service_designer

About this task

Fields are represented as variables in blueprints. To add a field, you add a variable to an operation in the blueprint.

Procedure

1. Open a blueprint that you already created.
2. Click the **Catalog** tab, and then click the operation, such as Provision, that you want to add the variable to.

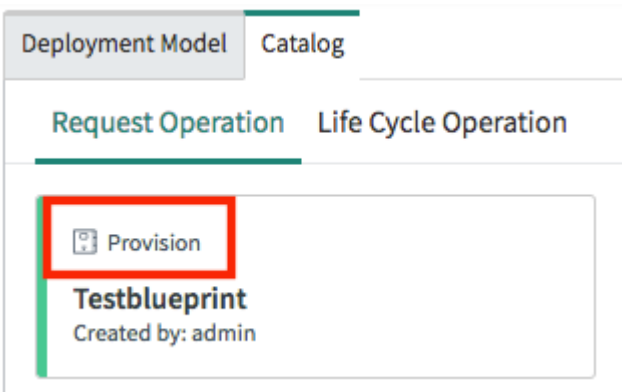

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Cloud Catalog Item form opens.

3. In the Variable Sets related list at the bottom, click the set in which you want the variable to belong.
   For example, click General Info, which is a default variable set for all blueprints, to have the field appear in the first tab on the catalog order form.
   The Variable Set form opens.
4. In the Cloud Variables related list at the bottom, click **New**.
5. Fill in the important form fields, such as:

Field	Description
Type	Select the type of variable, such as **Single Line Text** for a text field, **CheckBox**, or **Date**, for example.
Question tab	In **Question**, which is the value the user sees for the field on the catalog order form, enter a name. The **Name** field, which is the actual name of the variable, is automatically populated with the same value in lower case.
Other tabs	Fill in other fields on the remaining tabs to control the field's visibility, form help text, default value, and access rights.

Note: This variable will not be available for orchestration. If you intend to use this variable for orchestration, please add it through Cloud Template or Blueprint instead.

An example variable for an email address

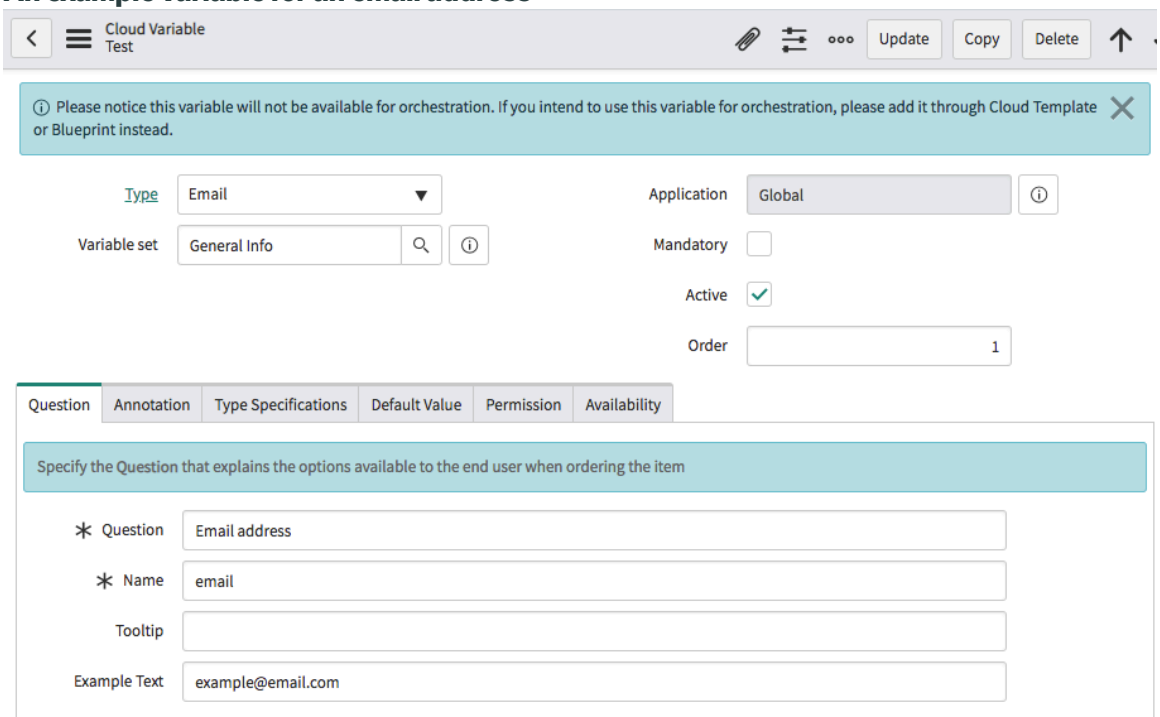

6. Click **Submit**.

   The Variable Set form opens, showing the new variable in the Cloud Variables related list.
7. To control the position of the field on the form, change the value in the **Order** column.

Note: A lower order value moves the field higher up on the form, and vice versa.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

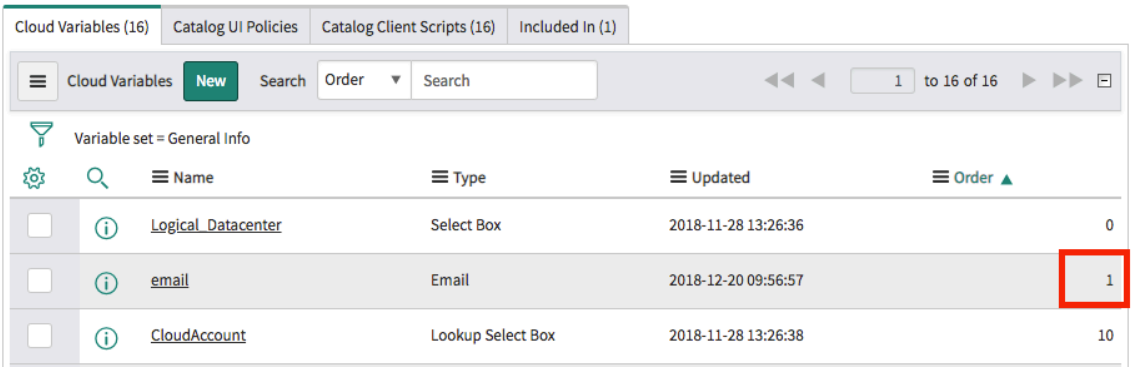

In this example, the order is set to **1**, which means the Email address field appears after the Logical Datacenter field and before the Cloud Account field.

8. On the form banner in the upper-left, click **Try It** to preview the order form as it would look to a user from the Cloud Catalog.
   The form appears with the new field that you created.

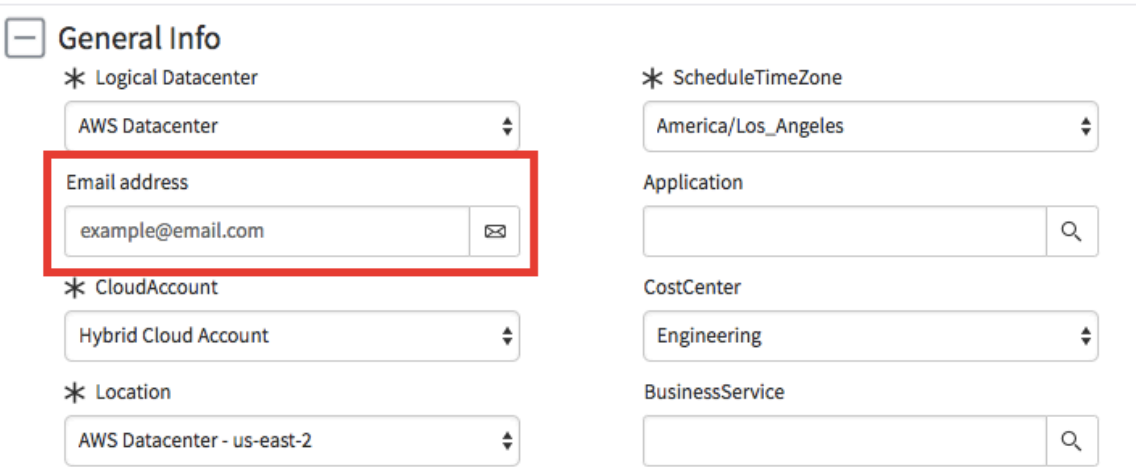

9. Click the back arrow on the upper-left of the form.
10. Toggle the catalog item to **Published**.

Control visibility default values for Cloud catalog items

Through blueprints, you can control how form fields appear to users in the cloud catalog. For example, you can control where the fields appear to users on the catalog item form, and whether they are read only or mandatory.

Before you begin

Role required: admin

Procedure

1. Open the blueprint to modify, and then toggle it from **Published** to **Draft**.
2. Click the **Catalog** tab.
   The Request Operation sub-tab opens and displays the published catalog items for the Provision operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click the **Provision** tile for the catalog item.
   The settings in the Provision operation control the form that the user sees upon initial request of the item, before it is provisioned.

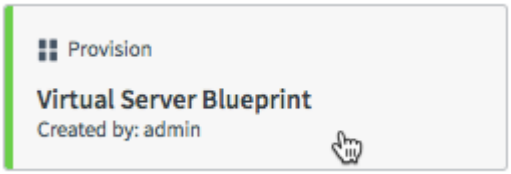

   The Cloud Catalog Item form opens.

4. In the Variable Sets related list at the bottom, click the set in which you want the variable to belong.
   For example, click **General Info**, which is a default variable set for all blueprints, to have the field appear in the first tab on the catalog order form.
5. In the Variable Set form, select the variable that you want to modify.
6. In the Cloud Variable form, complete any of the following actions.

To	Action
Make the field mandatory	Select the **Mandatory** check box. Clear the check box to make the field optional.
Make the field read-only	Click the **Availability** tab, and then select the **Read Only** check box. Clear the check box to make the field editable. Use this option to automatically populate a val ue in the field that the user cannot change.
Show or hide the field	Click the **Availability** tab, and then select or clear the **Visible** check box.
Move the field to another tab	Select a different **Variable set**. **General Info** and **Provision** are provided in the base sys tem, and you can configure additional form UI groups. To add a new tab for organizing fields in the catalog item, see Create a variable set for Cloud Provisioning and Governance.
Set a default value	Click the **Default Value** tab and enter the val ue in **Default value**. If the field is not read-on ly, users can change the value on the cloud catalog form. Reference fields must use the Sys ID of the record that represents the default value.
Add help for the user	Do any of the following: ○ Click the **Question** tab, and add a value in the **Tooltip** field to show a bubble with the text when the user hovers the cursor over the field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To	Action
	○ Click the **Annotation** tab, select **Show** help, and then enter values in the **Help text** or **In structions** fields to display additional text to the user above the field.

7. Click **Update**.
8. Toggle the catalog item to **Published**.

Example:
In this example, the Business Service field:

- Is populated with the Email business service record from the cmdb_ci_service table based on the Sys ID of the record.
- Is read-only so users cannot change it.
- Is mandatory because it is considered necessary before the system can provision the virtual resource.
- Contains help text to tell the user that the default selection is Email.

On the catalog order form, the field appears like this graphic:

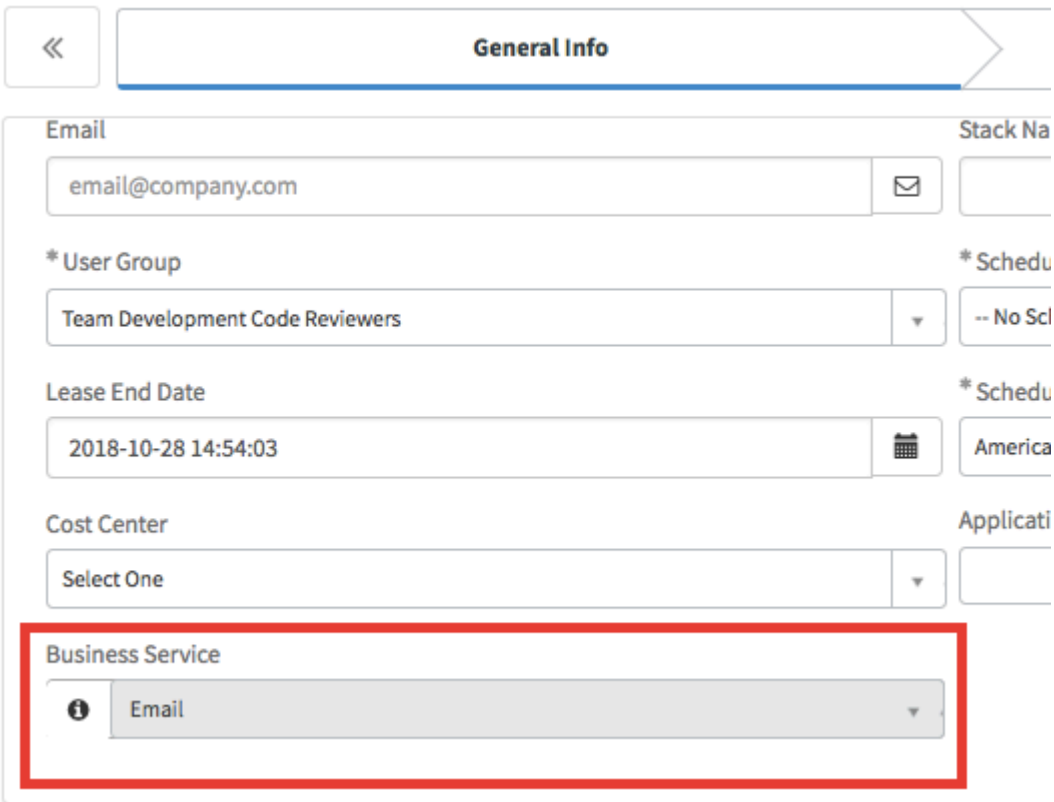

Cloud catalog form configuration examples

Several examples are provided to help you understand how to manipulate Cloud catalog order forms. These example cover blueprint rules, blueprint variables, and blueprint form UI groups.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example	Goal	Blueprint components used
Populate the location based on the cloud account	Use existing blueprint rules to illustrate how the Location field is populated based on the Cloud Account that the user selects.	This example illustrates rules.
Remove a field from a cloud catalog form	Hide the Cost Center field on the catalog form when the catalog requester belongs to the Database group.	This example illustrates: • Rules • Resource pools
Populate stack name with a script	Give the catalog requester a pre-defined stack name that includes the ID of the catalog requester. The field remains editable to the user can change it if necessary.	This example illustrates: • Rules • Resource pools • Scripted values
Add fields to a cloud catalog item order form	Create a catalog variable and property, and use it to publish a new filed on a form.	This example illustrates: • Catalog variables • Field order • Rules • Resource pools • Scripted values
Set default values for a cloud catalog form fields	Set a default value on a property and on a blueprint rule.	This example illustrates: • Default variables • Rules
Add a UI group with additional fields	Group catalog form fields into tabs to better organize the form and guide the user from one set of settings to another.	This example illustrates UI groups (tabs).

Populate the location based on the cloud account

This example walks you through an existing set of blueprint rules that Cloud Provisioning and Governance provides by default.

Before you begin

Before you can follow this example, you should be familiar with creating blueprints and blueprint rules. See these resources:

- Cloud Provisioning Blueprints
- Build a Cloud Provisioning blueprint
- Configure a blueprint rule

Role required: sn_cmp.cloud_service_designer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The goal of this example is to show you how a default set of blueprint rules accomplishes the task of populating the Location field based on the Cloud Account. Because this example walks you through default data, you do not need to create any records if you are following these steps.

This blueprint, which includes a virtual server that can be provisioned in an AWS datacenter, is used to illustrate the example:

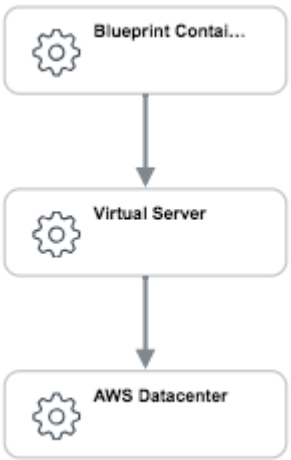

Procedure

1. Open a similar blueprint with a virtual server or similar.
2. On the blueprint form, click the **Catalog** tab, and then click the tile that represents the **Provision** operation.

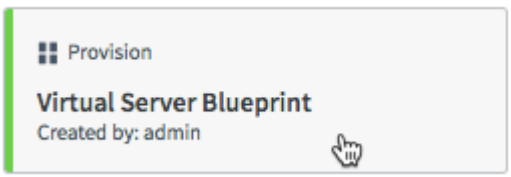

3. Click the **General Info** variable set in the Variable Sets related list.
4. Click the **CloudAccount** variable to open the Cloud Variable form.

   Notice these items:

   - In the Rules related list at the bottom, two rules are provided:
     - ReloadLocation reloads the value in the relocation field. This rule runs first because it has a lower order number than the other rule.
     - TriggerChangeEvent Location tells the system to trigger the actions for the rules that are already configured for the Location variable.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The **Lookup Select Box** field specifies that the Cloud Account field provides a selection for the user to make. The options that the user can select comes from the CloudAccountPool, which is on the **Type Specifications** field.

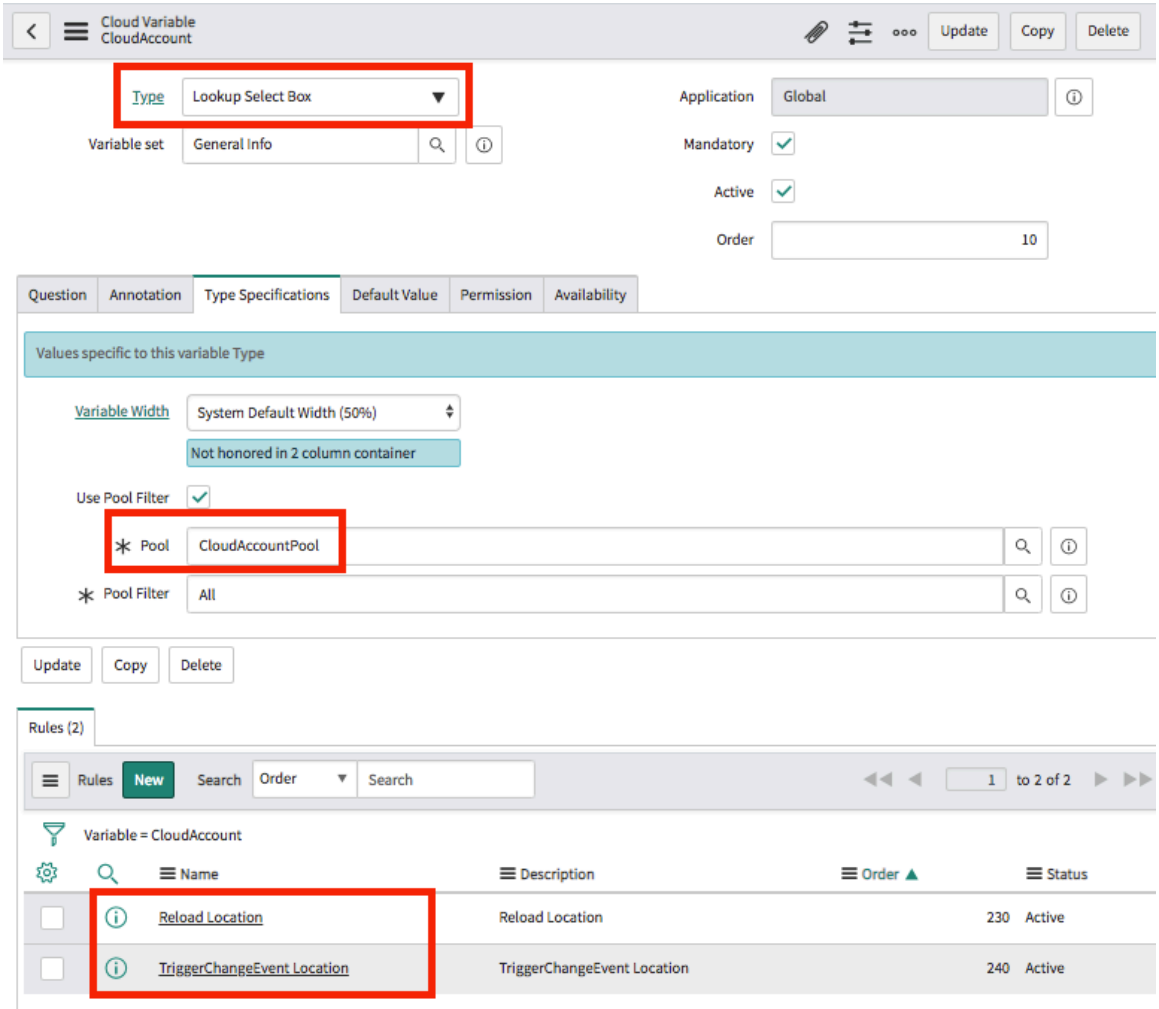

5. Click the **Reload Location** rule, and then click the **Reload Location** action from the Actions related list.

   Notice these items:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The **Action Type** is **Reload**, which tells the system to reload the catalog item form so that a new value for the field specified in the **Target Variable** (the **Location**) can be populated.

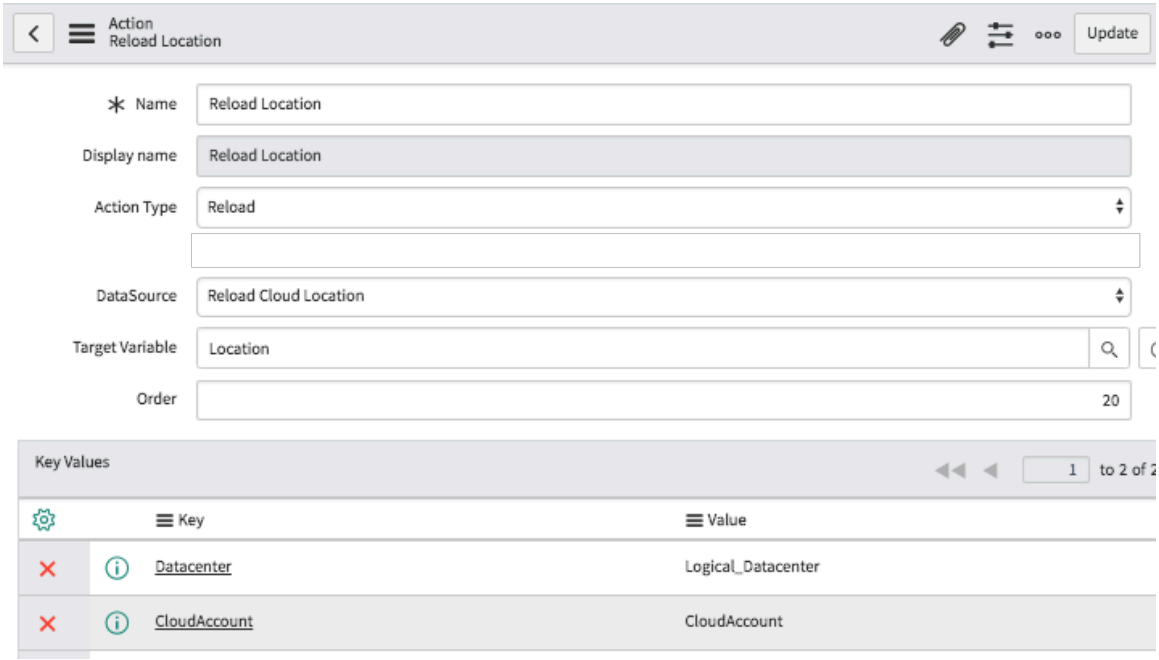

- The **DataSource** field tells the system to reload the location.
- The **KeyValues** list, which is only used when you specify a **Reload** action type, specifies that the user is passing in the cloud account via the form. The cloud account is necessary to reload the correct value for the location because the location is loaded from a pool. That pool has a filter that requires the input of the cloud account to return the correct location options.

6. Navigate back to the CloudAccount catalog property form and open the other rule: **TriggerChangeEvent Location**.
7. Click the **Actions** related list, open the **TriggerChangeEvent Location** action, and then click the **TriggerChangeEvent Location** action from the Actions related list.

   Notice these items:

   - The **Action Type** is **TriggerChangeEvent**, which tells the system to trigger the actions for the rules that are already configured for the Location property. By default, several rules are already configured for the Location property. These include rules that reload other fields, like **OS Profile**, **Network**, and so on.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

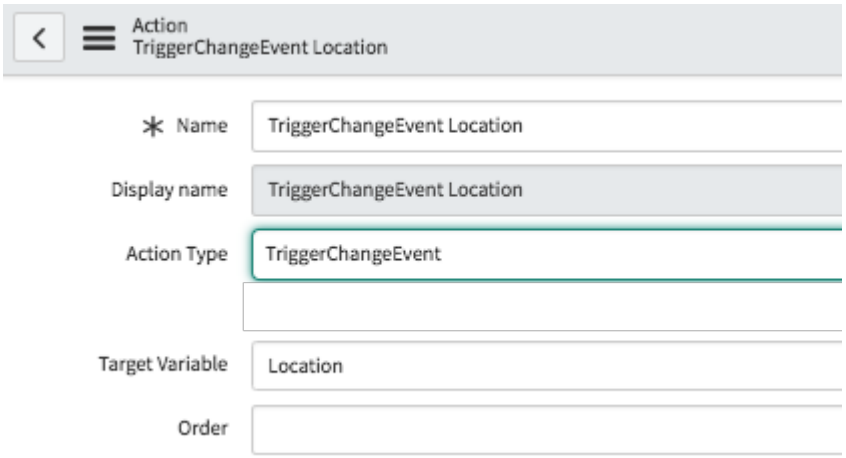

- No **KeyValues** are required because this action is not a reload action.

The result of these two rules on the CloudAccount attribute is illustrated in the following example:

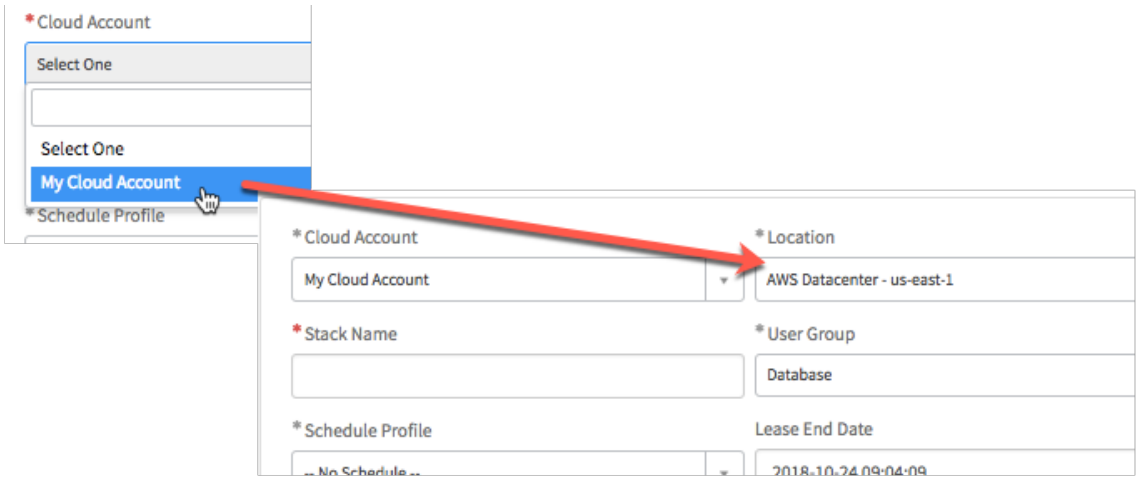

8. Perform these steps to see why the system populated the value **AWS Datacenter - us-east-1** in the **Location** field and how the system determines which location is related to the **My Cloud Account** value:

   a. In the same blueprint, open the **General Info** variable set, and then open the **Location** parameter.

   b. Click the **Type Specifications** tab.

      Notice the values that can be populated for the **Location** field from the *`CloudLocationPool`* resource pool.

   c. Navigate to **Manage** > **Resource Pools**, open **CloudLocationPool**, and notice these items:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

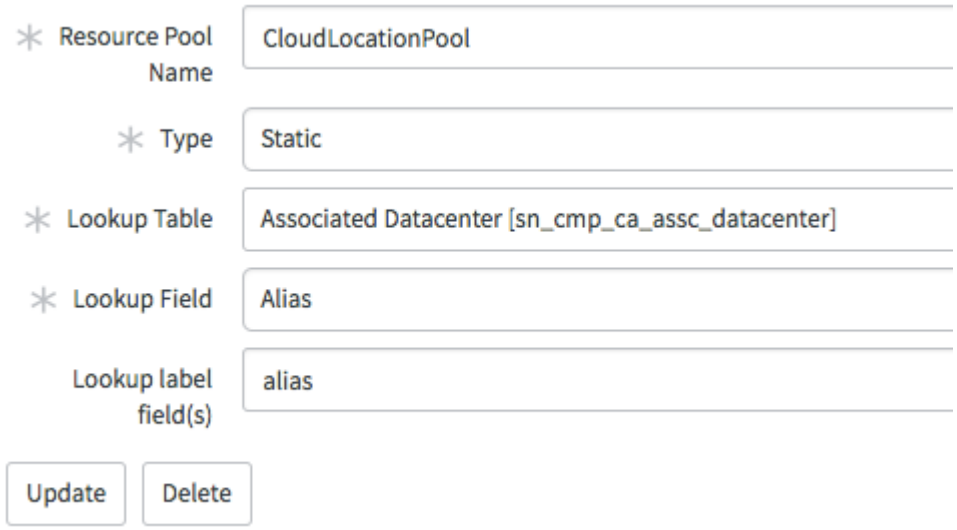

- The locations come from the records in the **Lookup Table**, which in this example is **Associated Datacenter [sn_cmp_ca_assc_datacenter]**.
- The actual value to display to users in the Location field on the catalog item form comes from the **Lookup Field**. In this example, the **Lookup Field** specifies the Alias column from the Associated Datacenter [sn_cmp_ca_assc_datacenter] table. The **Alias** column holds the actual, user-readable names for all locations.

d. Go to the standard interface of your instance, not the Admin Portal, and enter `sn_cmp_ca_assc_datacenter.list` in the navigation filter.

The list of datacenters that are already associated with cloud accounts appears.

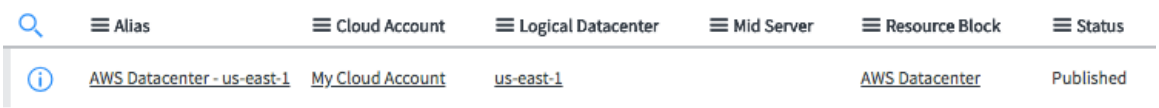

Notice that the **My Cloud Account** cloud account is already associated with **AWS Datacenter - us-east-1**, which is the value that the user sees in the Location field. The system knew to display that value because of the association in this Associated Datacenter table. The system created this association when your Cloud Provisioning and Governance administrator created the cloud account for AWS, and then ran discovery on the us-east-1 datacenter in that cloud account.

It is possible that multiple datacenters could be associated with **My Cloud Account**. In that case, the Location field reloads all the associated datacenters and in the Location field. Since the Location field is of type Select, the field allows the user to select which location to choose.

Remove a field from a cloud catalog form

Remove a field from the Cloud Provisioning and Governance catalog form when the user belongs to a specific group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: sn_cmp.cloud_service_designer

Before you can follow this example, you should be familiar with creating blueprints and blueprint rules. See these resources:

- Cloud Provisioning Blueprints
- Build a Cloud Provisioning blueprint
- Configure a blueprint rule

About this task

The goal of this example is to hide the Cost Center field on the catalog form when the catalog requester belongs to the Database group. You can follow an example like this one when you want to show or hide catalog item fields based on values in other fields.

Procedure

1. Create a blueprint with a virtual server.

   For this example, a blueprint with a virtual server on an AWS datacenter is used.

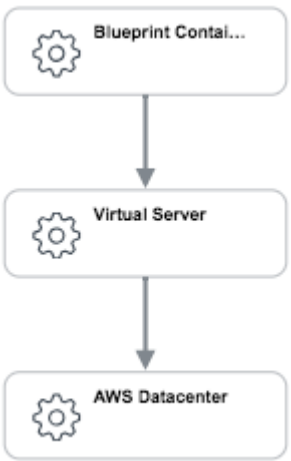

2. On the blueprint form, click the **Catalog** tab, and then click the tile that represents the **Provision** operation.

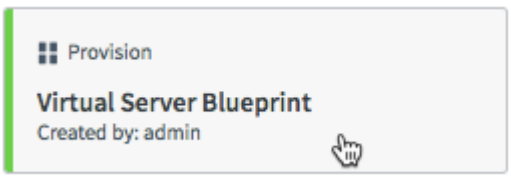

3. Click the **General Info** variable set in the Variable Sets related list.
4. Click the **CostCenter** variable to open the Cloud Variable form.
5. Configure the rule that removes the cost center based on the user group:
   a. In the **Rules** related list, click **New**.
   b. Fill out the **Name** and description **Fields**.
   c. Right-click the form header and select **Save**.
6. Configure the condition that requires the User Group to equal Database:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Conditions** related list, click **New**.

b. On the form, fill in the fields.

Field	Value to enter or select	Description
Operand	UserGroup	This value is the User Group field that the user selects.
Operator	--None--	Keep this value **--None--** because this condition is the only condition you need to configure. If you add more conditions, you can select another operator to determine how the system evaluates this condition in relation to the others.
Value	287ee6fea9fe198100ada7950d0b1b73	This value is the Sys ID of the Database user group, which is provided as part of the demo data in your instance. See Sys ID for more information.

c. Click **Submit**.

7. To understand why the Sys ID of the user group is required for the system, perform these steps:

a. Open the blueprint again and click the Provision operation again.

b. Click the **General Info** from the **Variable Sets** related list.

c. Click the **UserGroup** variable.

Notice that on the **Type Specifications** tab, the system populates user groups from **UserGroupPool**. The system runs the scripted pool filter named **GroupsUserBelongsTo**.

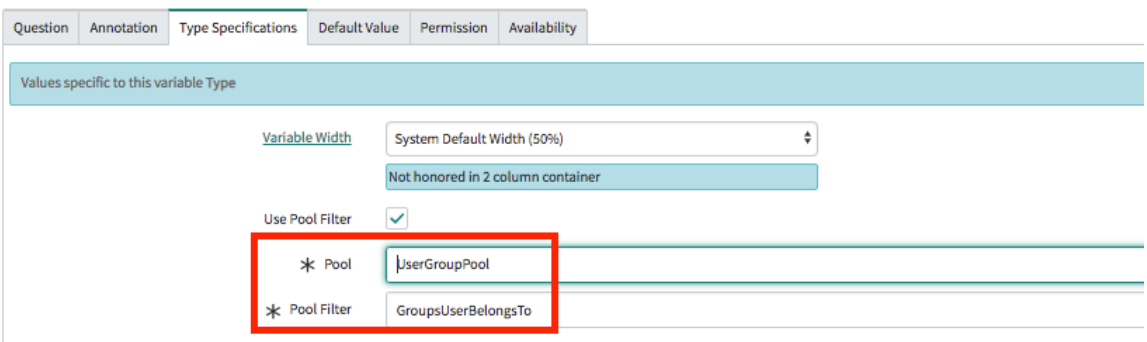

d. In the Cloud Admin Portal, navigate to **Manage** > **Resource Pools**.

e. Find and open **UserGroupPool**.

Notice that the values for the pool are derived from the table in the **Lookup Table** field, which in this case is the Group [sys_user_group] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Also notice that there are two queries in the Resource Pool Filters related list:

- **All** is a generic query that returns the list of all records from the **Lookup Table**. The value it returns from that table is from the **Lookup Field**, which in this case is the Sys ID. However, the system ignores the *A11* query in this case because the scripted query is present.
- **GroupsUserBelongsTo** is a scripted query that returns the Sys ID of records in the **Lookup Table**.

Important: When a scripted query is present, the system uses it instead of the generic *ALL* query.

Both the generic query and the scripted query return the Sys ID of the record in the lookup table. But what is actually happening here is that the system ignores the **All** query and the value specified in the **Lookup Field**, and it uses the **GroupsUserBelongsTo** script.

A scripted query

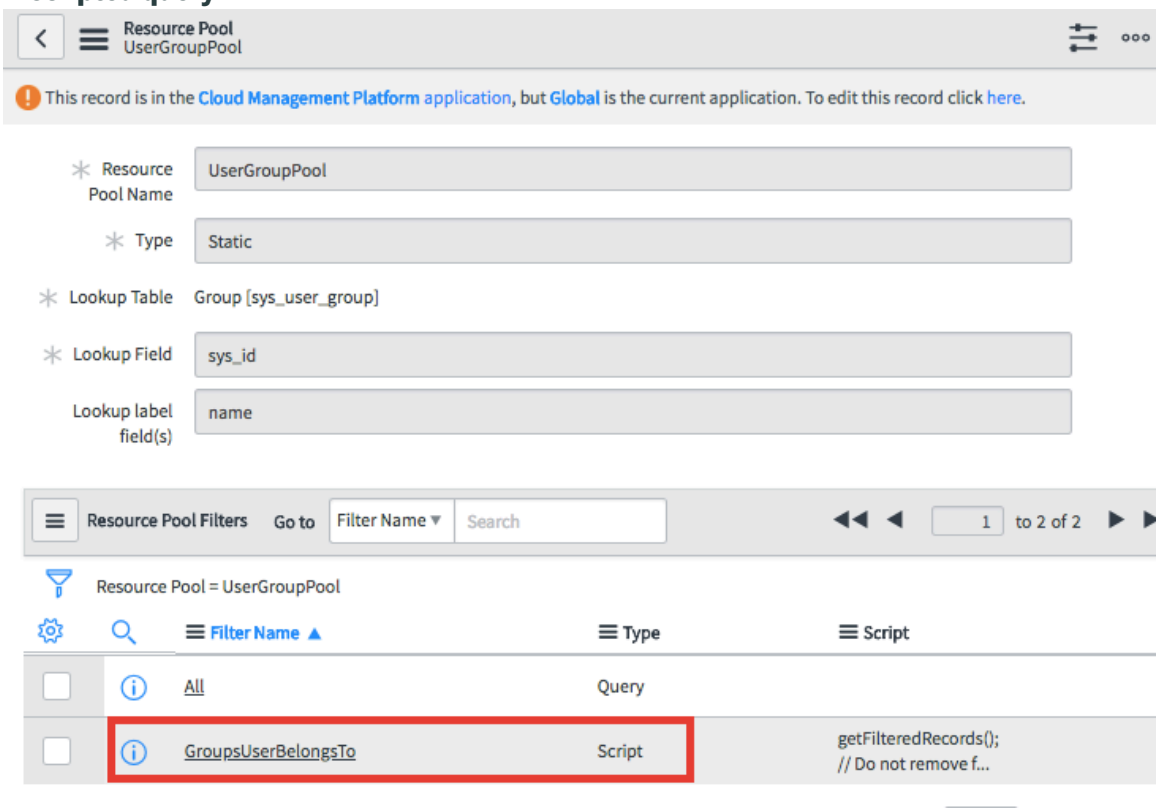

Populate stack name with a script

This example illustrates how to populate the name of the stack that a user requests from the Cloud Provisioning and Governance catalog by utilizing the user's ID. The example uses a script include that you can call from a blueprint rule action.

Before you begin

Role required: sn_cmp.cloud_service_designer

Before you can follow this example, you should be familiar with creating blueprints and blueprint rules. See these resources:

- Cloud Provisioning Blueprints
- Build a Cloud Provisioning blueprint
- Configure a blueprint rule

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This example uses a script include. If you want to create similar script includes, familiarize yourself with the Glide System APIs that are part of your instance. See the Developer site for a comprehensive list of all APIs. See Script Includes for more information about how the system uses these scripts.

About this task

The goal of this example is to take the ID of the catalog requester and populate it in the Stack Name field. The field remains editable. Follow an example like this one when you want to loosely enforce a naming or labeling convention on certain values. This approach still allows the user to enter the text that they want.

Procedure

1. Create a script include that returns the user ID of the person requesting the catalog item:

   a. In the standard interface for the instance (not in the Cloud Admin Portal), click the gear icon ( ) in the upper left, and then click **Developer**.

   b. In the **Application** field, select **Cloud Provisioning and Governance**.

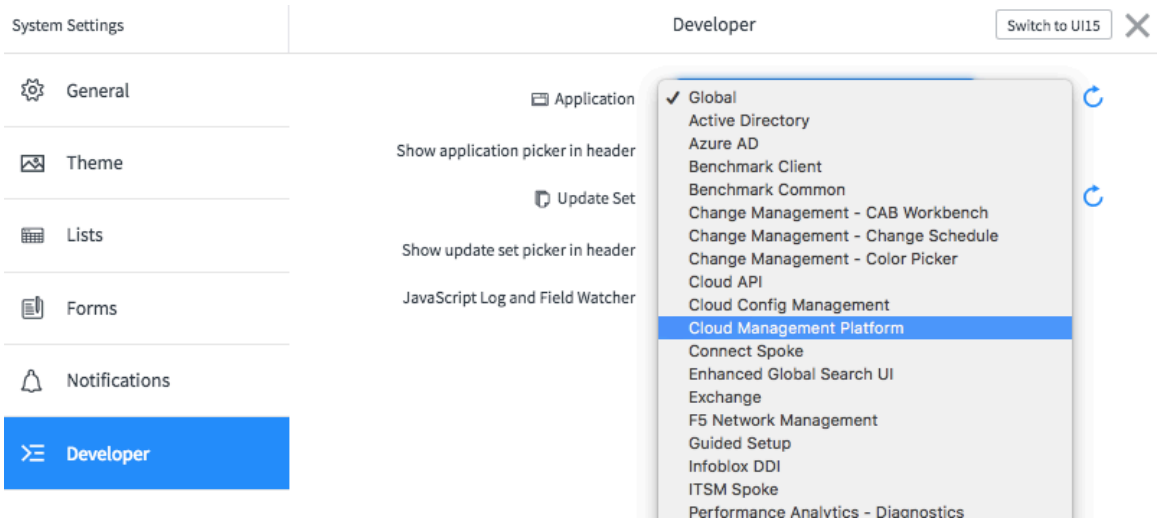

   The application scope of your instance changes to**Cloud Provisioning and Governance**. This selection means that the Cloud Provisioning and Governance application can find the script include.

   c. Close the System Settings window.

   d. Navigate to **System Definition** > **Script Includes**.

   e. Click **New**.

   f. Fill in the form fields (see table):

Field	Description
Name	Provide a descriptive name without spaces.
Accessible from	Select **All application scopes**.
Script	Enter this script:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description

	``
	var SetFields = Class.create();
	SetFields.prototype = {
	 initialize: function() {
	 },
	 setStackName: function(){
	 return gs.getUserName();
	 },
	 type: 'HideFormAttributes'
	};
	``

The function `setStackName()` returns the user name of the catalog item requester. You can call this function from a blueprint rule action.	

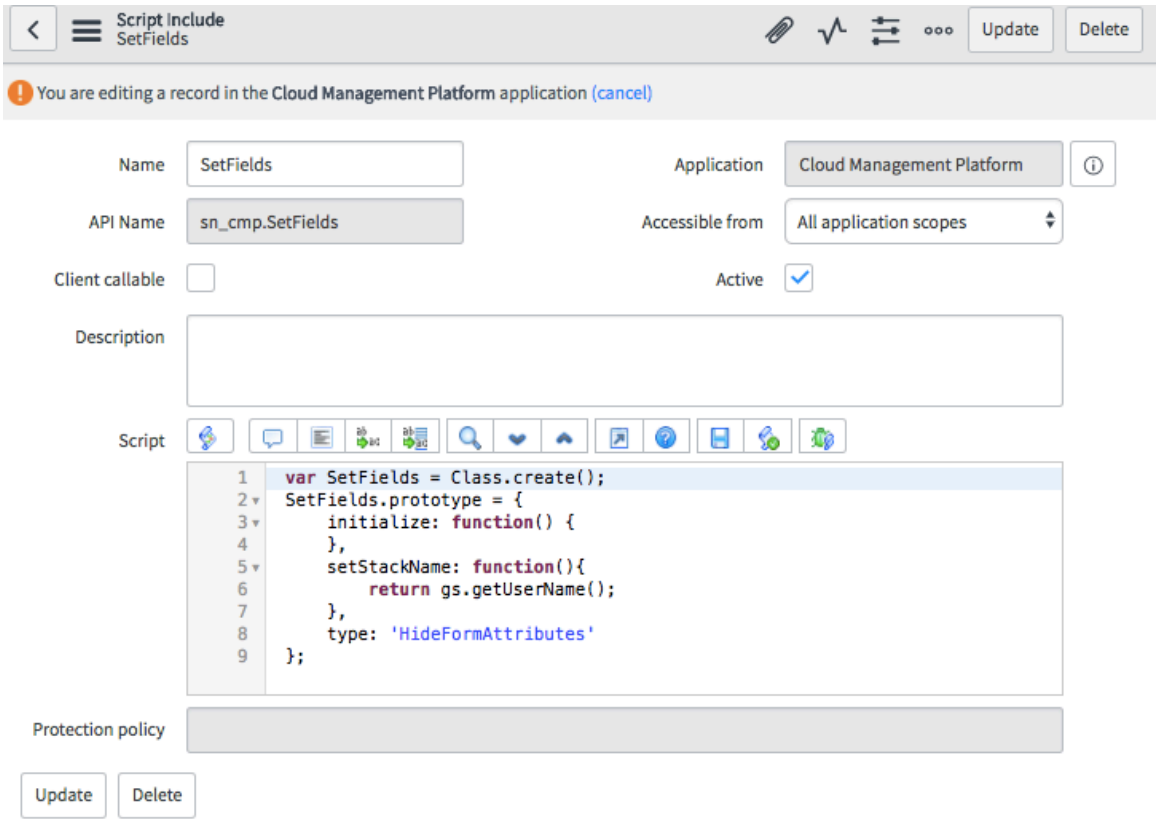

g. Click **Submit**.

h. Click the gear icon ( ) in the upper left again, and then change the **Application** back to **Global**.

2. In the Cloud Admin portal, navigate to **Design** > **Blueprints** and open the relevant blueprint.
3. On the blueprint form, click the **Catalog** tab, and then click the tile that represents the **Provision** operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

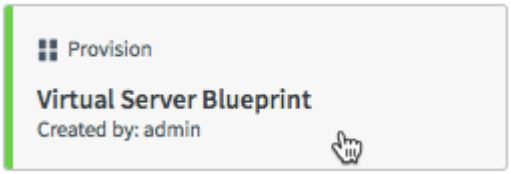

4. Click the **Stack Name** variable set in the Variable Sets related list.

   Notice that Stack Name is a single line text field that takes text input.

5. Create a rule to populate the stack name:
   a. Scroll down to the bottom of the form and click **New** in the **Rules** related list.
   b. Provide a name and description.
   c. Click **Submit**.
6. Configure the action as follows:

   a. From the Rule form, scroll down and click **New** in the **Action** related list.

   b. Configure the rule as follows:

Field	Description
Name	Provide a meaningful name and description, such as `Set Stack Name to User ID`.
Action Type	Select **SetValueByScript** to tell the system to call the script include that you specify in the **Value** field.
Target Variable	Select **StackName** to tell the system to set the value of the Stack Name field.
Value	Call the method in the script include that returns the user ID. The call must be in the following syntax: `$(Script:ScriptIncludeName.methodName[])` In this example, use `$(Script:SetFields.setStackName[])`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Click **Submit**.

Result

In the Cloud User Portal, the catalog order form for the Provision operation shows the user ID of the logged-in user. This example uses the user Able Tutor, which comes in the base system as demo data. The value in the **Stack Name** field is the same value as in the User ID field on the User [sys_user] table. The profile icon for the same logged in user is also displayed in the upper right of the Cloud User Portal.

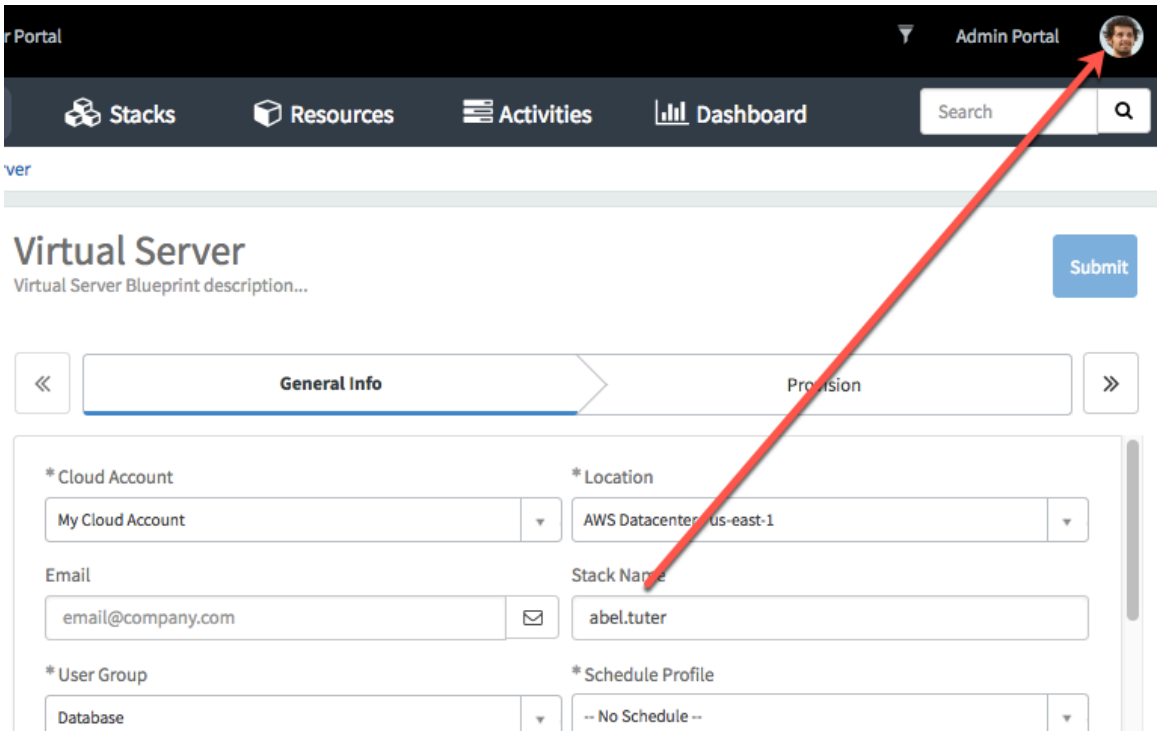

Set default values for a cloud catalog form fields

You can set default values for all users or you can populate fields with default values based on conditions. This example shows you how to set a default schedule profile and business service.

Before you begin

Role required: sn_cmp.cloud_service_designer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you can follow this example, you should be familiar with creating blueprints and blueprint rules. See these resources:

- Cloud Provisioning Blueprints
- Build a Cloud Provisioning blueprint
- Configure a blueprint rule

About this task

The goal of this example is to display a default value for both the schedule profile and the business service field. Each field derives its values differently:

- The schedule profile derives its values from a pool filter that returns a text value. Therefore, you must set the default value by applying the correct, matching text for one of the schedule profile options that appear to the user in a list.
- The business service derives its values from a reference to the Service [cmdb_ci_service] table. Therefore, you must set the value by specifying the Sys ID of the record that you want to display.

The example in this topic sets the values in different ways:

- The schedule profile default value is set directly on the Schedule Profile property and applies to all users without any conditions. You might want to do this to limit the use of virtual resources to specific schedule, such as during business hours only, for all users.
- The business service default value is set using an action on a blueprint rule. You might want to use a rule to set a default value when you want the value to be the default only when certain conditions are met. For example, if the user belongs to a user group that deals with only one business service, set the value for the user. This prevents them from sorting through a long list of all the business services in your organization.

Procedure

1. Create a blueprint with a virtual server.

   For this example, a blueprint with a virtual server on an AWS datacenter is used.

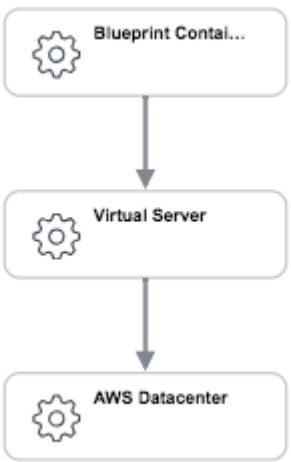

2. On the blueprint form, click the **Catalog** tab, and then click the tile that represents the **Provision** operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

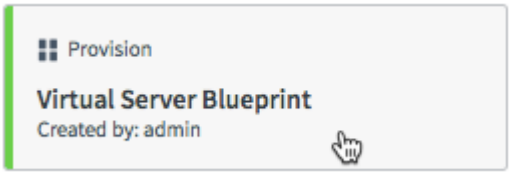

3. Click the **General Info** variable set in the Variable Sets related list.
4. Click the **ScheduleProfile** variable.

   Notice that the **Pool** field on the **Type Specifications** tab specifies the pool from which the values for the schedule profile are derived: ScheduleProfilePool.

   The **ScheduleProfilePool** pool specifies the **All** filter, which returns the values in the Name field from the Schedule Profile [sn_cmp_schedule_profile] table. These values are static text values. You can specify the exact text value for the name of the schedule profile that you want to make the default.

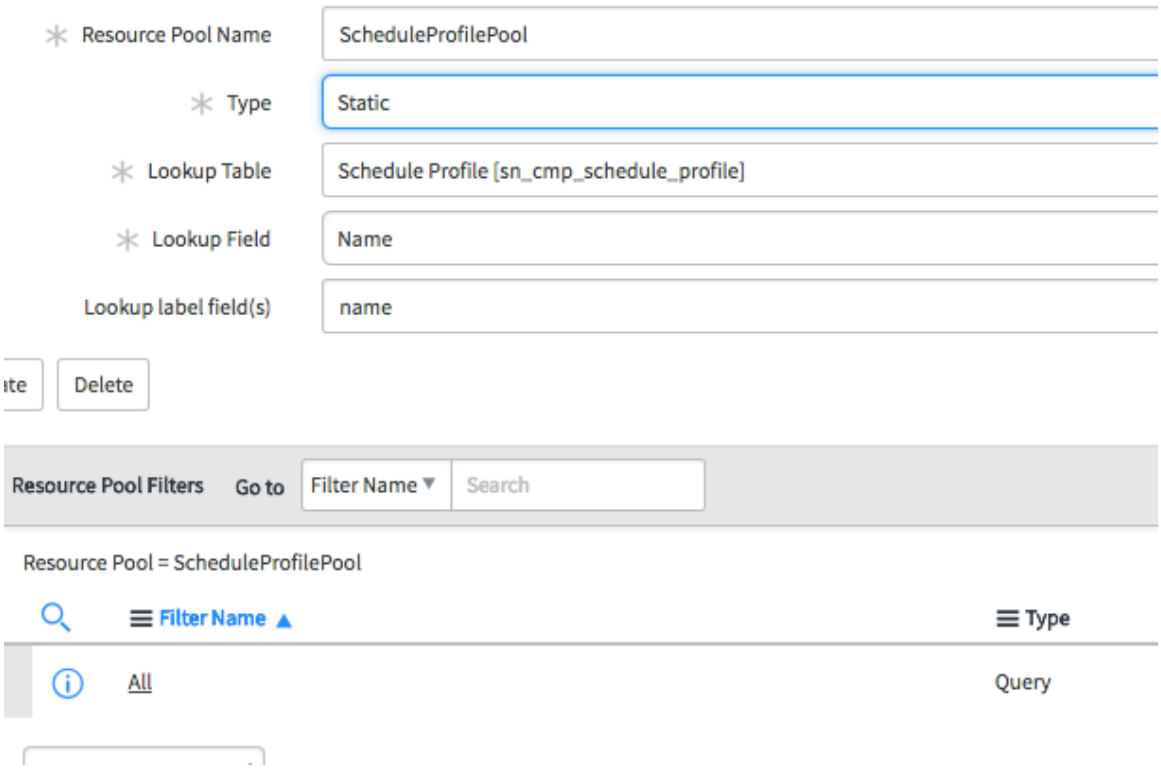

5. On the **Default Value** tab of the ScheduleProfile property form, enter the text that matches the value in the Name field from the Schedule Profile [sn_cmp_schedule_profile] table.

   In this example, that value is `8-5 weekdays`.

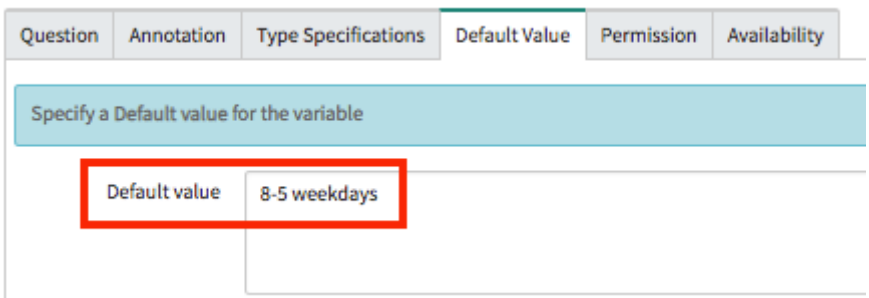

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Click **Update**.
7. Click the **BusinessService** variable.

   Notice that the **Reference** field on the **Type Specifications** tab specifies the Service [cmdb_ci_service] table, from which the values for the business service are derived.

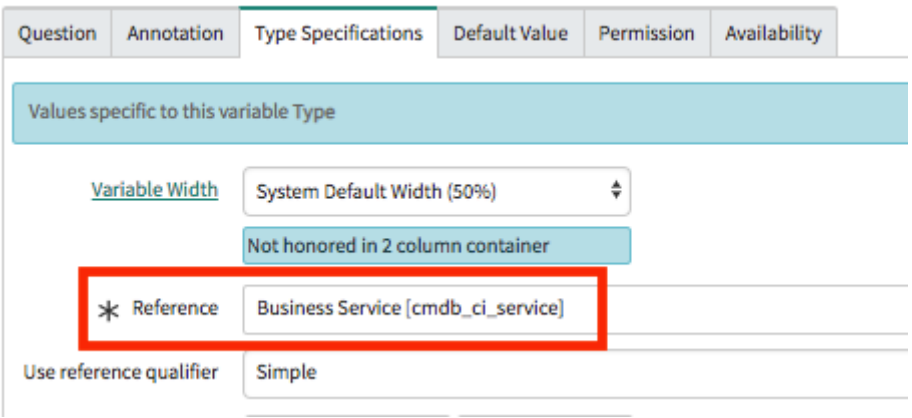

   By default, to refer to a specific record in a table, you must know the Sys ID of the record.

8. To obtain the Sys ID of a business service:

   a. In the main interface for the instance (not the Cloud Admin Portal interface), enter `cmdb_ci_service.list` in the application filter.

   b. Find the business service that you want to make the default.

   c. Right-click the record, and then select **Copy sys_id**.

      This example uses **Workday Enterprise Services**.

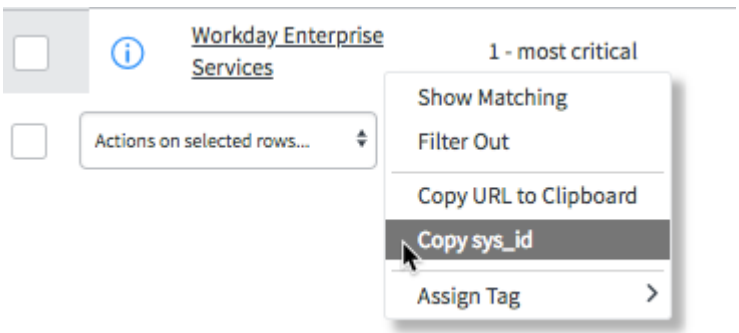

9. On the **BusinessService** variable form, create a rule, action, and condition with these values:

   a. In the **Rule** related list, click **New**.

   b. Provide a name and a description.

   c. Right-click the form header and select **Save**.

   d. In the **Action** related list for the rule, click **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. Provide a name, and then configure these fields as specified:

Field	Value
Action Type	**SetValue** This option tells the system to change the field to the value you specify. Selecting this option is the same as entering a default value on the property form.
Target Field	**BusinessService**
Value	The Sys ID of the business service that you copied.

In this example, the business service is set to Workday Enterprise Services by default, which is represented in the system with the Sys ID that you already copied.

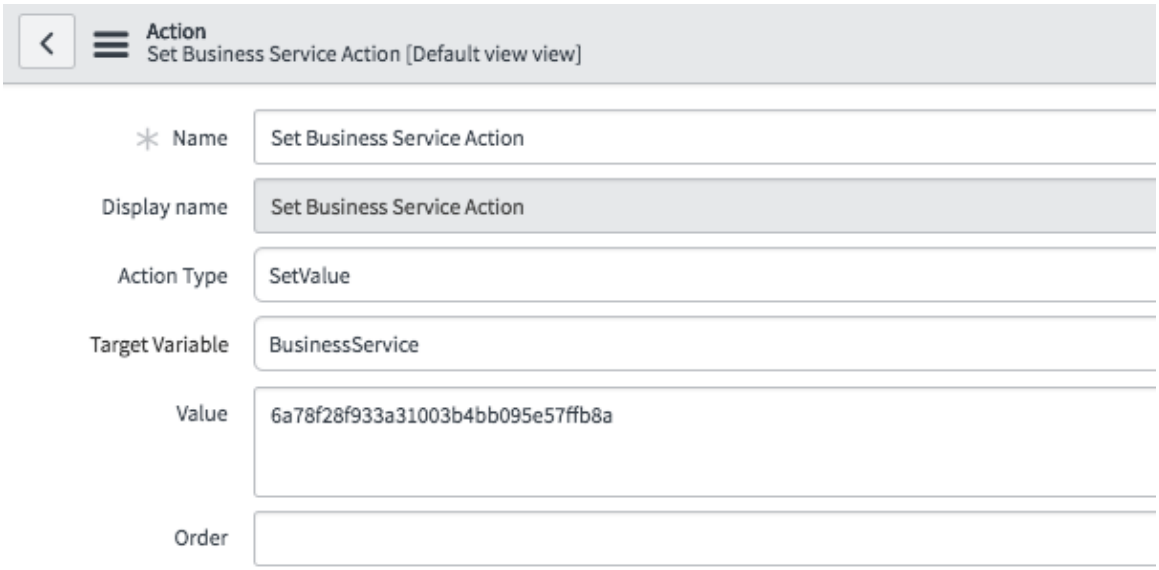

f. Click **Submit**.

Result

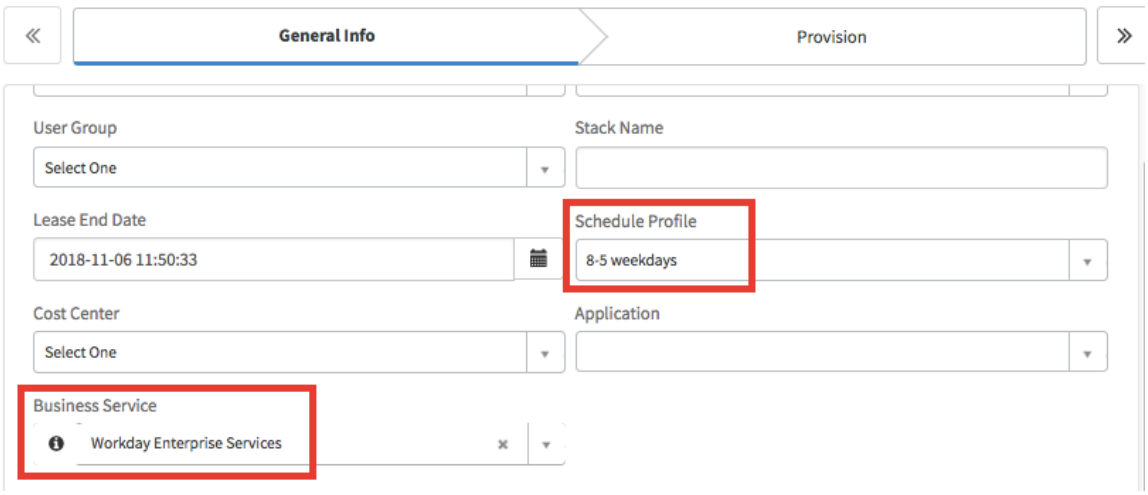

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can create a condition or multiple conditions that must be met before the system sets the default value for the business service. For example, you might want to change the default value of business service to Workday Enterprise Services. This change applies only when the group that the user belongs to is the Workday Admins group.

Resource blocks in Cloud Provisioning and Governance

Resource blocks act as a type of middleware component between catalog items, the Cloud API (CAPI), and the CMDB.

The following graphic illustrates how resource blocks fit between other components in the system:

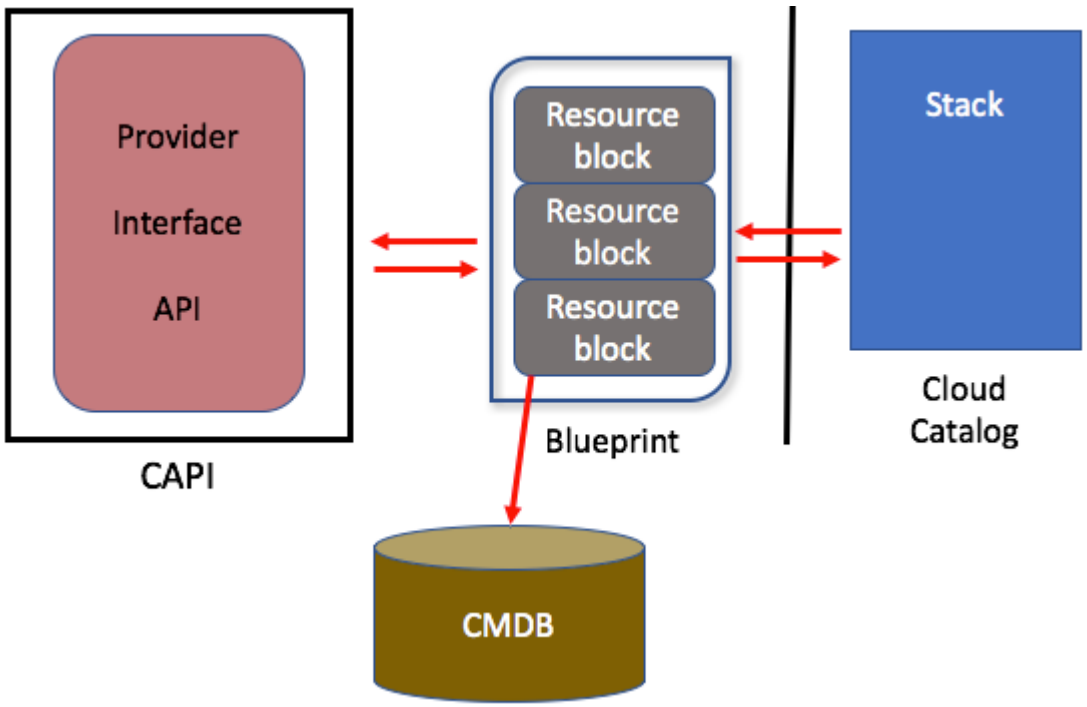

Blueprints

Each resource block becomes a building block within a blueprint. The blueprint eventually becomes a catalog item (also called a stack) that a user provisions from the Cloud Catalog.

Cloud API (CAPI)

Each resource block defines allowable operations, such as Provision and Deprovision. These operations call various components within CAPI to carry out the operation that is received from the stack.

The CMDB

Each resource block is based on a CI type from the CMDB. Resource blocks also have a response processor that handles responses from the cloud provider. The response processor takes action in the CMDB, such as creating or updating a CI.

How resource blocks are connected

Resource blocks are connected to each other in blueprints. For example, in this blueprint, a virtual server resource block is connected to storage. The virtual server and the storage are both connected to an AWS datacenter resource block.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Blueprint server with storage

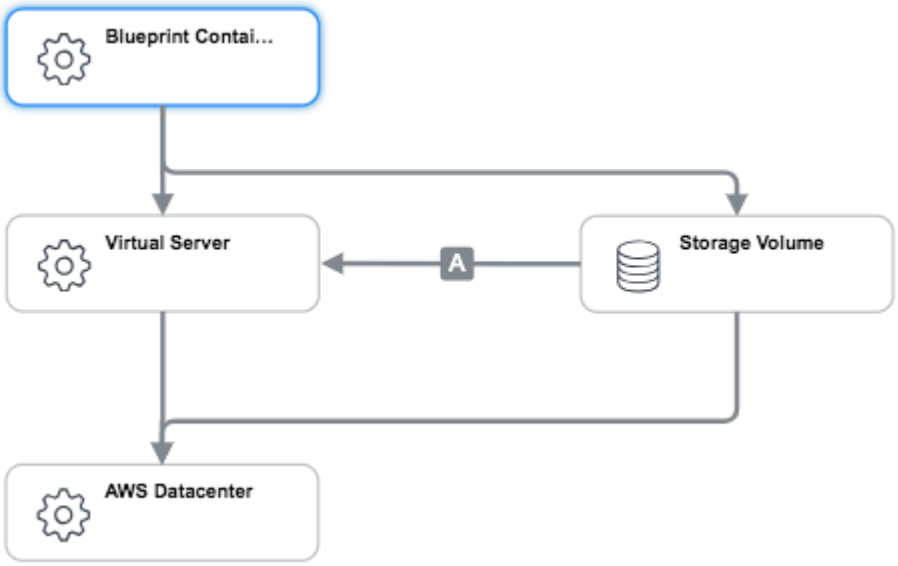

For these connections to be successful, each resource block needs to specify a type of interface. This graphic illustrates these interfaces:

Types of interfaces for resource blocks

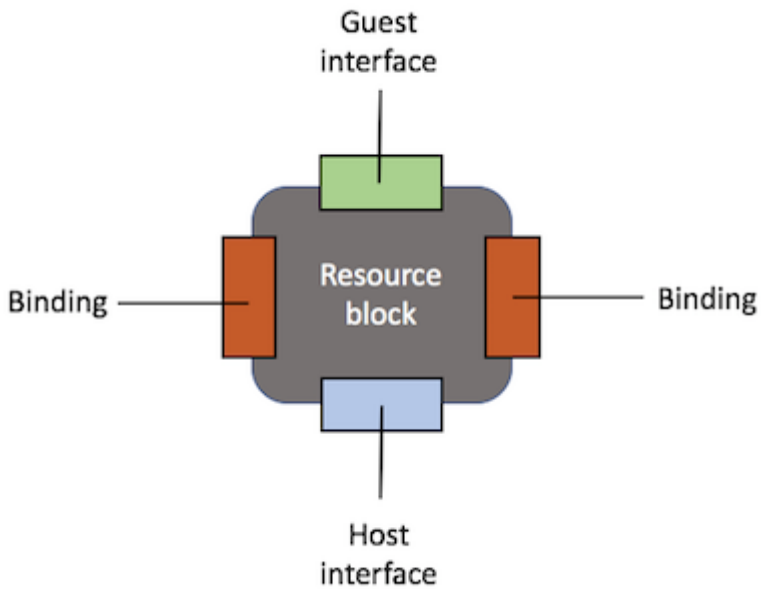

- The guest interface connects to the resource that is above it. The guest interface contains the operations (also called operation signatures), which allows users to take action on their resource.
- The host interface connects to the resource below the resource block in the blueprint. The host interface also contains operation signatures.
- The bindings connect adjacent resources.

Resource blocks can have as many guest interfaces, host interfaces, and bindings as necessary. Connections can be made in different directions:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Vertical connections require the top resource block to have the same host interface as the bottom resource block guest interface.

  Note: In blueprints, this vertical connection is called a containment connection.

- Horizontal connections bind one resource to another, but both resources have to reside on top of the same resource block.

  Note: In blueprints, this horizontal connection is called an attached-to connection.

The same example blueprint with a virtual server with attached storage on an AWS datacenter, the resource blocks are configured as this graphic illustrates:

Interfaces for resource blocks

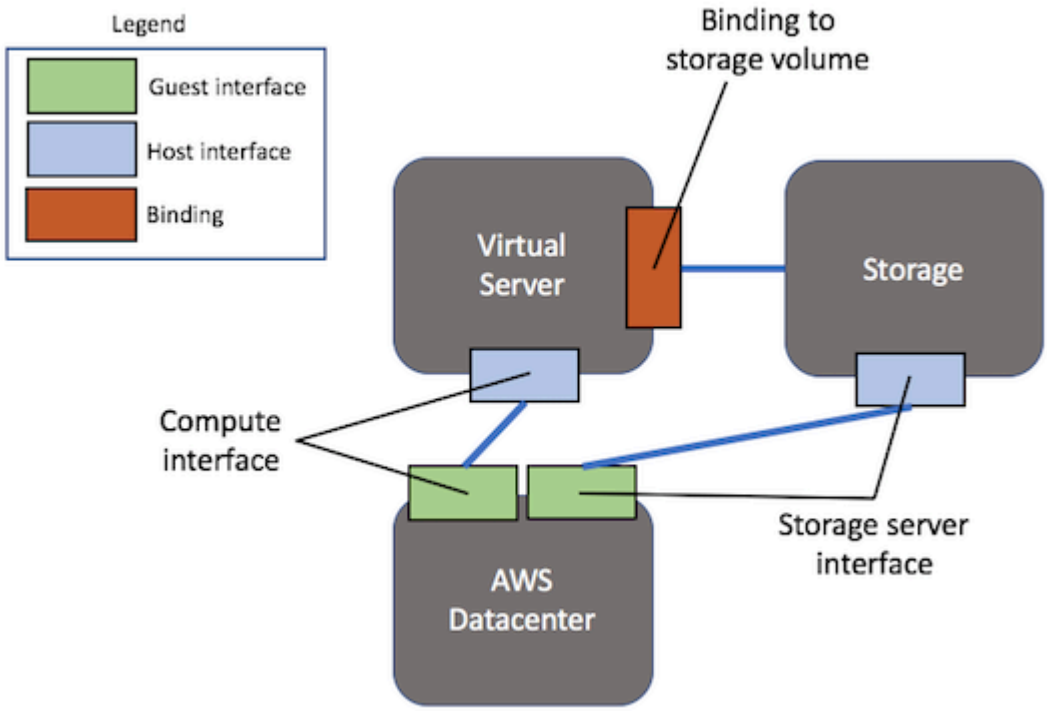

Resource block	Interfaces and bindings
Virtual server	• Host interface: Compute • Binding: to Storage Volume
AWS Datacenter	• Guest interface: Compute • Guest interface: Storage Server interface
Storage	Host interface: Storage Server interface The binding between the virtual server and the storage is unidirectional. You configure the binding on the virtual storage resource block only by specifying

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource block	Interfaces and bindings
	the storage resource block. You do not need to configure another binding on the storage resource block.

Guest interfaces and operations

Each guest interface provides a default set of operations (also called operation signatures) that allow the user to choose what they want to do with a virtual resource. The most common operation is Provision, which means that a virtual resource is created. Many guest interfaces and corresponding operations are provided by default with the Cloud Provisioning and Governance application. You can use these default guest interfaces for many of your resource blocks.

If you do need to create a new guest interface, these operations are provided by default:

- Start
- Stop
- Provision
- Deprovision
- Cleanup
- ModifySchedule
- ModifyLease

Resource block component details

Each resource block contains these components:

Layers

The logical layer in the stack that this resource block belongs to. By default, Cloud Provisioning and Governance provides the layers. When you configure resource blocks, you can choose the one that relates to you.

Operations for each guest interface

As mentioned, each guest interface provides operations. For example, the most common operations for a virtual server are **Provision**, **Deprovision**, **Start**, **Stop**.

Input parameters for each operation

The input parameters hold values that the cloud provider uses, via CAPI, when operations are run. These parameters hold the information that the user selects while performing an operation, such as provisioning, on a stack.

For example, a user typically must specify the location, which is a datacenter, that a stack belongs in. The value that the user selects in the **Location** field is held in the Location parameter. The system takes this parameter, which is available by default for the virtual server resource block, and passes it to the *`CreateNode`* CAPI interface operation. The *`CreateNode`* interface operation tells the cloud provider to create (or provision) the actual virtual server.

Steps for each operation

Each operation requires a step to do the actual work. Each step calls a specific interface and method in CAPI, and passes along the input parameters required by that method. For example, the **Provision** operation on a virtual server provides one step to first connect and create a virtual machine. Another step creates the node

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

in the cloud provider. Each step provides different input parameters. To connect and create a virtual machine, the cloud account, among many other parameters, is required. To create the node, the cloud account is not required.

You can add as many steps as needed for each operation. You can also use output attributes to pass values in attributes from one step to another step. You can pass values from steps in one operation to steps in a subsequent operation. For example, the **Provision** operation for a virtual server can provide an output attribute for the network interface ID and another for the node ID. When the cloud provider sends the response, the actual ID of the network and node are parsed out of the response and saved in the output attributes. Subsequent steps can then use the IDs to perform actions.

Response processors

Response processors tell the system what to do in the CMDB. When the cloud provider sends a response after an operation is run, the processor handles the response.

Note: These processors are scripts that are system includes. They are not MID Server script includes like the CAPI scripts that you can create to integrate with the cloud provider.

How components work together

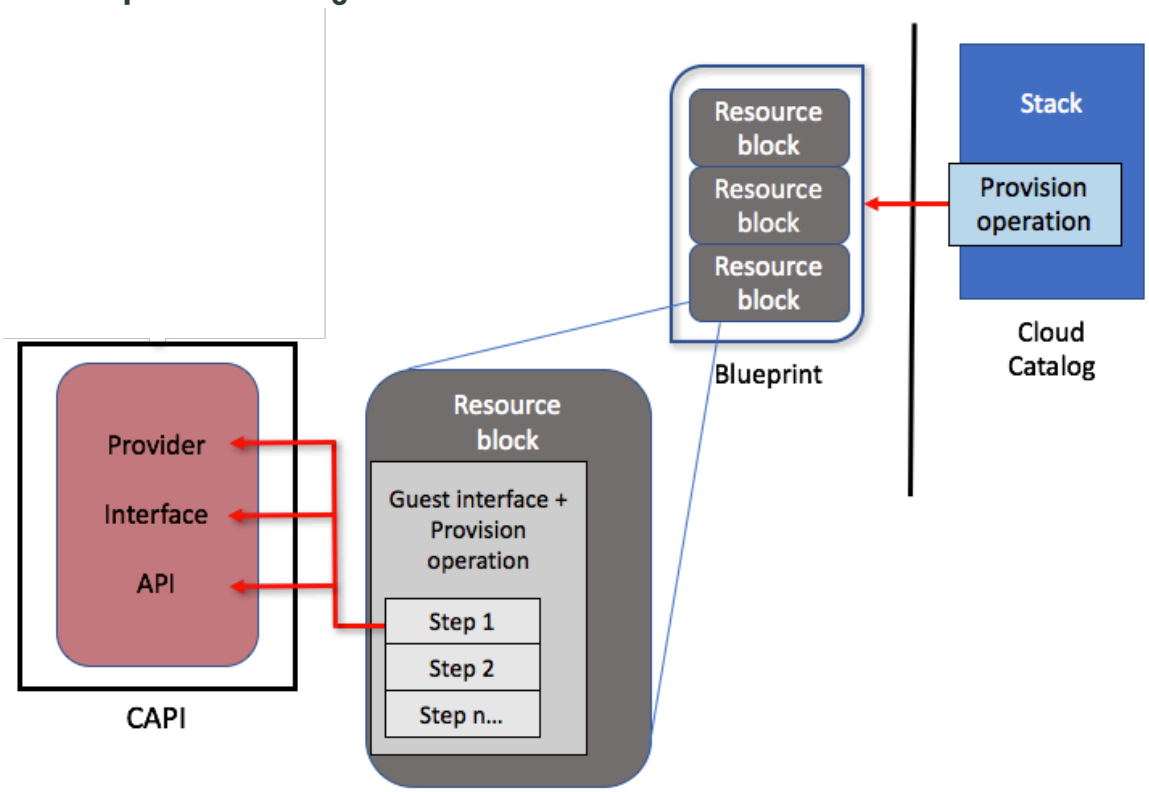

Users can see catalog items, also called stacks, in the Cloud Catalog. A catalog item is based on a blueprint, which in turn is composed of resource blocks that are connected together.

1. The user kicks off the process by requesting a stack. The request is the **Provision** operation in the system.
2. The system runs the steps in order for the **Provision** operation for the resource blocks that comprise the stack.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Each step calls CAPI to determine which method to run, while also passing in values for necessary parameters.
4. CAPI interfaces with the API of the cloud provider to provision the actual resource.

Process flow when a resource is created in the cloud provider

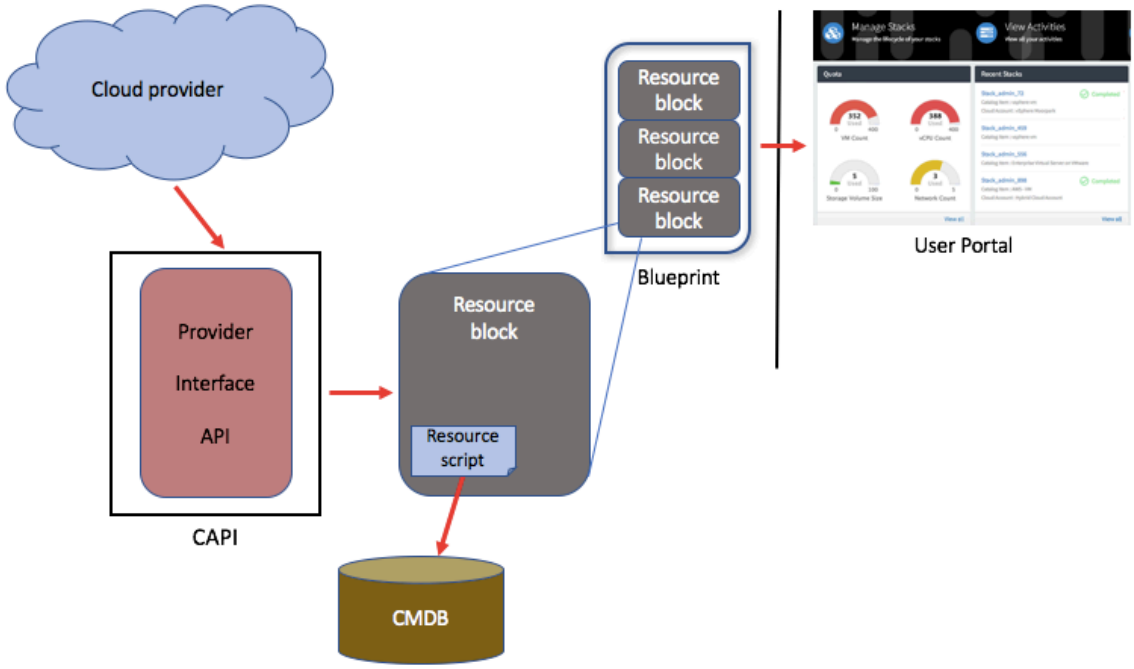

1. The cloud provider provisions the resource and responds with JSON.
2. The system parses the response inside the resource script in the resource block.
3. The system updates the CMDB. When provisioning a new resource, the system creates a CI in the correct table.
4. The new stack appears in the cloud user portal for the user to manage.

The Resource Blocks page

Manage resource blocks in the Resource Blocks page. Open the Cloud Admin portal, and then navigate to **Design** > **Resource Blocks**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource Blocks page components

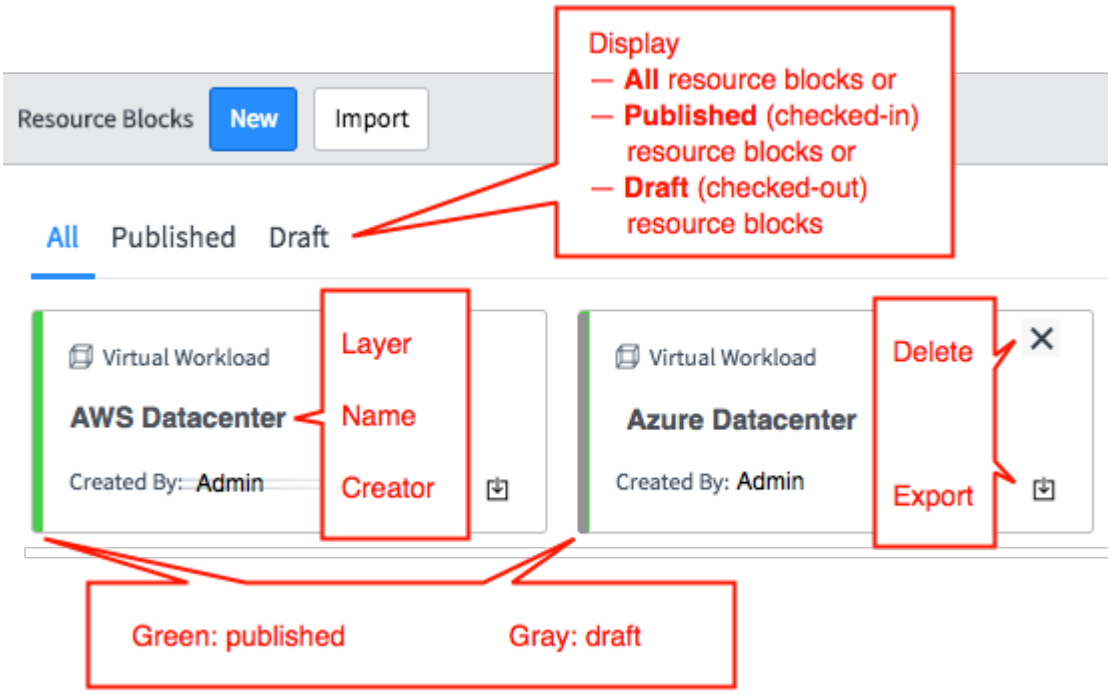

Create a custom resource block

If the blocks in the base system do not provide the cloud resources that you need to provision, you can create a custom resource block.

Before you begin

Role required: sn_cmp.cloud_service_designer

Procedure

1. Navigate to **All** > **Cloud Provisioning and Governance** > **Cloud Admin Portal** > **Design** > **Resource Blocks** and click **New**.
2. Enter a unique and meaningful **Name** and **Description** for the resource block, and then enter the following settings:

Layer	Select the layer that represents the type of resource block: ○ **Application**: Resources that are applications. For example, the 'PetStore' application might exist on a Tomcat or WebLogic platform. ○ **Platform**: Resources that are typically hosted on a virtual server like MySQL, Tomcat, or Apache. ○ **Logical DataCenter**: Resources that are datacenters like AWS, Azure, or vSphere. ○ **Virtual Infrastructure**: Resources that are virtual and still provide infrastructure for the virtual world like NetApp SVM, NetApp CDOT, CISCO ICF Cloud, or VMware NSX. ○ **Virtual Workload**: Resources that are hosted on a datacenter. For example, Virtual Server, Storage Volume, Network, or Gateway.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

	Note: ○ Applications are deployed on Platforms. ○ Platforms are deployed on Logical Datacenters or Virtual Infrastructures.
CI type	Select the type of CI in your CMDB that the resource block represents. When the user launches a stack, the cloud resources are provisioned and actual CIs that for the resources are created in the CMDB in this CI type. For example, if you select Storage Volume, the system creates a record in the Storage Volume [cmdb_ci_storage_volume] table when the user provisions a new volume. **Note:** Each resource block is based on a CI type from the CMDB. Ensure that you do not have more than one resource block bound to a CI type.
Default Guest Interface Create interface	Select the guest interface that contains the operations you need for the resource. Later you can further define and select the operations to use for this default guest interface. The system gives you a set of default operations for each interface. If another resource block specifies a **Host interface** that matches what you specify here, that resource block can be a potential host for this resource block. See the Host interface field description. If you select the **Create interface** check box, enter a name for the guest interface. The system creates the guest interface, opens the resource block that you are configuring in the resource block form, and lists the new interface in the **Resource Guest Interface** related list. See the description of the Resource Guest Interface in the following steps. If you create a new guest interface, also keep in mind that other resource blocks can use the interface.
Host interface	Select the interface that the potential host of this resource block must offer up. Any other resource block that offers up a default guest interface of the same type you specify here can become a host of this resource. Later on, you can select one or more actual host resources. For example, by default the AWS Datacenter resource block specifies the **Compute** interface as a default guest interface. That means the AWS Datacenter is offering the **Compute** host interface to other resource blocks through which to connect to the AWS Datacenter. Specify the **Compute** interface as the host interface of the resource block. The resource block can then look for any other resource block using the **Compute** interface as its guest interface. The AWS Datacenter uses the **Compute** interface as its guest interface. You can later select the AWS Datacenter as the actual host in the Host Resource related list. Other resource blocks that offer the **Compute** interface by specifying **Compute** as the guest interface, are also potential hosts for this resource block. You must know which interfaces the potential host resources offer as guest interfaces. And determine which operations are part of those guest interfaces.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click **Submit**.
   Continue updating settings on the related lists, as needed.
4. **Optional:** Complete the following steps to add another guest interface.
   You can specify an additional guest interface to offer to other resource blocks. For example, a datacenter resource block might offer a **Compute** interface, a **Network Service Container interface**, and a **Firewall** interface for other resource blocks to use.

   a. Under **Guest Interface**, click **New**.

   b. In the **Guest interface** field, select an interface.

(Optional)

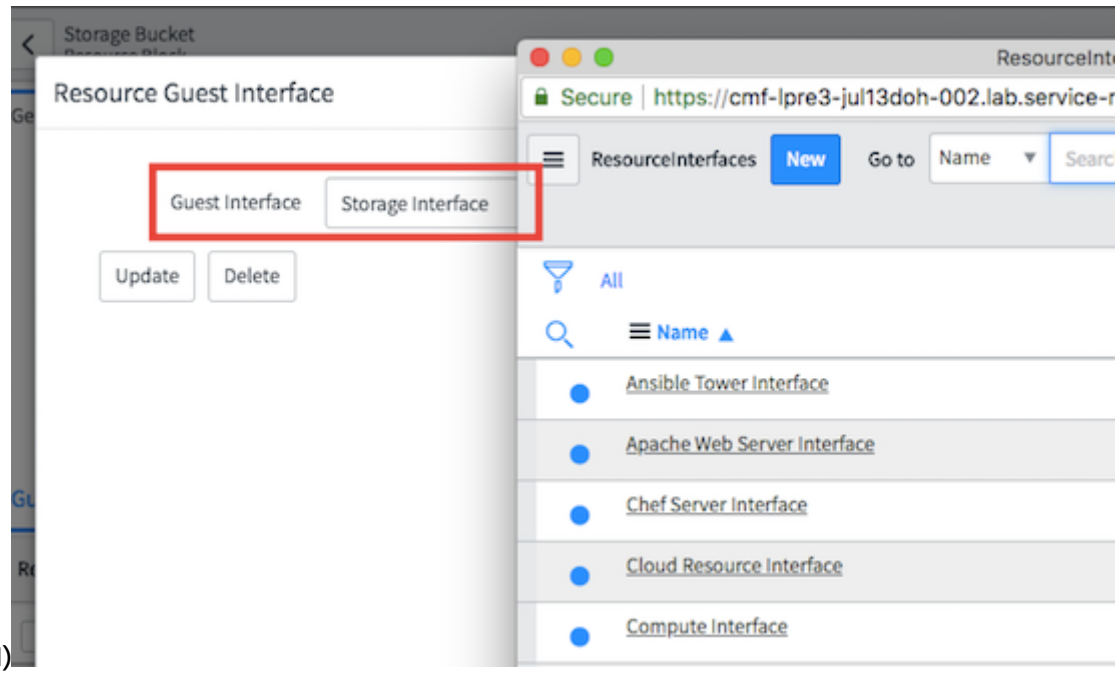

   c. Click **Submit** or **Update**.

5. In Host Resource, specify the resource that can host this resource block, typically the one that offers the interface you specified in the **Host Interface** field.
   The selected host resource block handles the provisioning of this resource block. For example, the host resource is a virtual machine resource block that can host an Apache Web server. The host resource selection limits the options that are presented to users when they select a host.

   Example

   For example select **vSphere Datacenter** in the **Host Resource** field for this resource to be hosted in a VMWare cloud.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

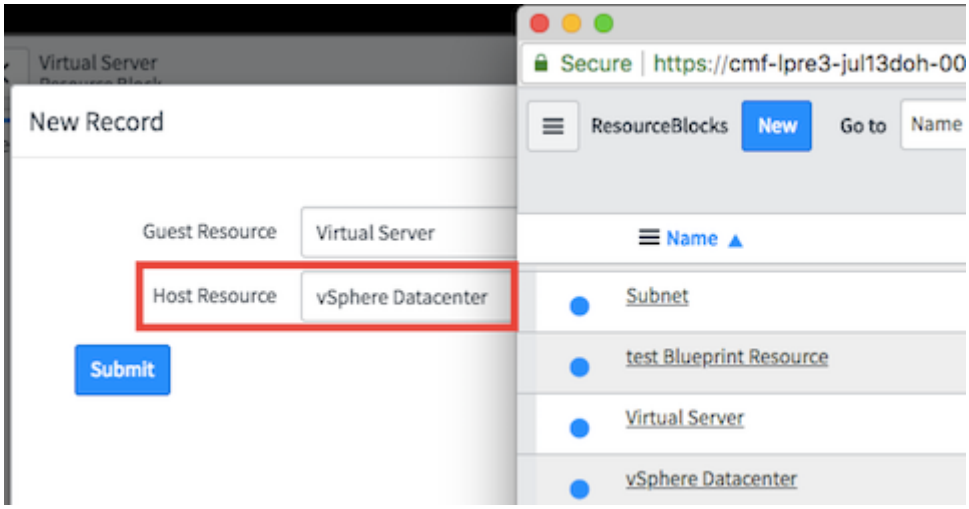

Note: You can select only the host resources that support the **Host interface** previously specified for this resource block.

6. **Optional:** Complete the following steps to add bindings.
   A binding is another resource block that this resource block can connect to. Select resource blocks that support the same **Guest interface** configured for this resource block.

   When you create a blueprint, you can add an attach connection between two resource blocks that are connected with this binding. Understanding the direction of the binding is essential to configuring the resource block. You attach a virtual server to a network, or you attach storage to a virtual server. You cannot reverse these directions.

   a. Under **Bindings**, click **New**.

   b. In the **Host Resource** field, select a resource in the **Guest Resource** field.
      The guest resource is the resource block that you are attaching to the current resource block.

   c. Click **Submit** or **Update**.

   Example
   For example, a virtual server can connect to virtual storage. By default, the base system provides a Virtual Server resource block with a binding to the Storage Volume resource block to enable this connection.

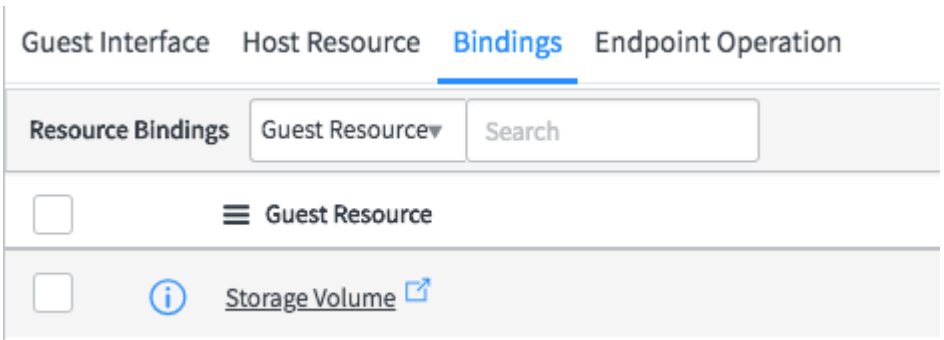

(Optional)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Complete the following steps to configure endpoint operations between the two resources to bind.
   The endpoints call operations during binding. For example, a storage volume can specify an endpoint with required conditions and the operations to execute to make a connection to a virtual server. The virtual server must also support a binding interface to the storage volume.

   Implement an endpoint type, such as **Block EP**, for the virtual server to consume the endpoint. Later, when you create blueprints, create the binding between the storage volume and the virtual server.

Example of binding in a blueprint

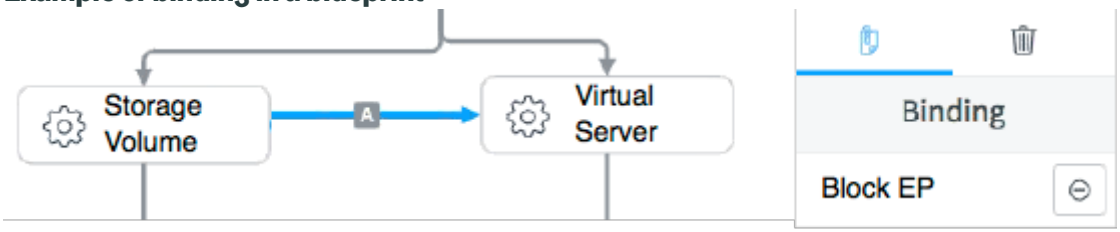

   a. Under **Endpoint Operations**, click **New**.

   b. In **Binding Resource**, select the resource block that consumes the endpoint from this resource block.

   c. Select the type of endpoint that the **Source Resource** presents.

   d. In **Operation implementation**, select the implementation that performs the operation you are defining in this endpoint mapping.

   e. Select the **Type** of operation to perform for the mapping.
   Only operations that are appropriate for the specified **Endpoint** appear in the list.

   Do not change the settings in **Source Resource**.

   f. Click **Submit** or **Update**.

Specify a host resource for a resource block

Hosts that support the **Host interface** of a resource block are potential hosts for the resource block. You use the **Host interface** setting to further limit the options that are presented to the stack requester while selecting a host type.

Before you begin

Role required: admin

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.
2. Select the resource block and then select **Draft** on the **Resource Blocks** form:

3. In the **Host Resource** related list, click **New**.
   On the **Resource Host** form, the **Application** and the **Guest Resource** (the resource block that you are currently configuring) are already specified.
4. Specify the **Host Resource**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, for a virtual server, you might select **AWS Datacenter** or **Azure Datacenter** as potential host types.

5. Click **Submit**.

Specify the bindings for resource blocks

Bindings represent endpoint relationships. For example, a storage volume might implement an endpoint type of Block EP (cmdb_ci_endpoint_block). A virtual server might consume an endpoint of that type. Bindings must support the Guest interface that is specified for the resource block.

Before you begin
Role required: admin

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.
2. Select the resource block and then select **Draft** on the **Resource Blocks** form:

3. In the **Resource Bindings** related list, click **New**.
4. On the **Resource Binding** form, the **Application** and the **Guest Resource** (the resource block that you are currently configuring) are already specified.
   Specify the **Binding Resource**.
5. Click **Submit**.

Configure endpoint operation mapping

You can configure endpoint mapping on a storage volume to make a connection to a resource.

Before you begin
Role required: admin

About this task
In the following example, the endpoint mapping that is configured for a storage volume specifies the conditions that must be met and the operation to execute to make a connection to a resource (a virtual server in the example) that supports a binding interface to a storage volume. The storage volume implements the endpoint (type Block EP in the example) and the virtual server consumes the endpoint:

Endpoint mapping for a storage volume

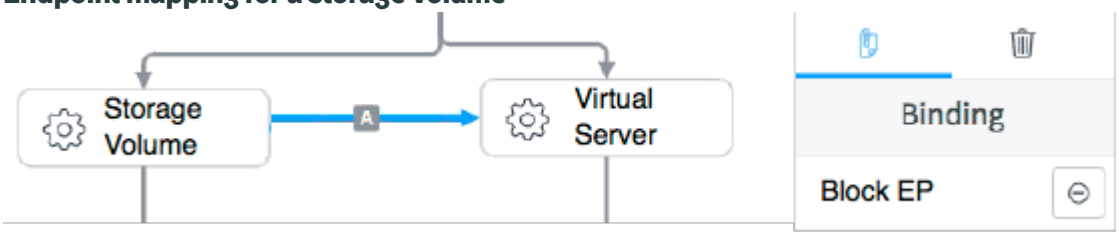

The settings on the example **Resource Endpoint Mappings** related list specifies the conditions and an action:

- For a resource block that supports a binding interface to a **Storage Volume**
- For a **Source** resource with **CI type** of **Storage Volume**
- With an **Endpoint** type of **Block EP** (cmdb_ci_endpoint_block)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the **Attach** operation to implement the connection. Use the specified **Operation Implementation** to perform the **Attach** operation.

Storage volume for a resource block

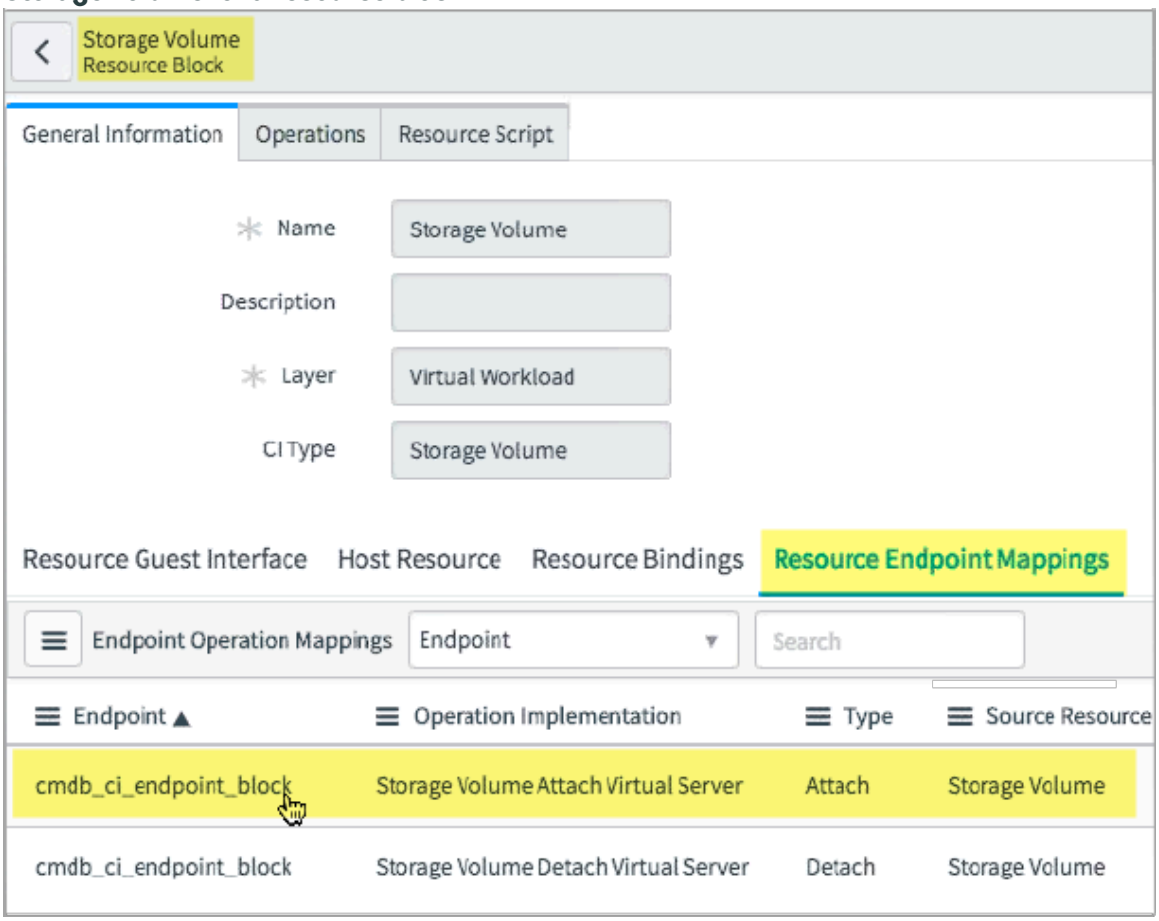

Procedure

1. While configuring a resource block, open the **Resource Endpoint Mapping** related list, click **New**, and then specify the following settings:

Binding Resource	Resource block that consumes the **Endpoint** that the **Source Resource** presents.
Endpoint	Type of endpoint that the **Source Resource** presents.
Type	Type of operation to perform for the mapping. Only operations that are appropriate for the specified **Endpoint** appear in the list.
Operation implementation	Implementation that is called to perform the operation that you are defining in this mapping.
Application	Read-only. Cloud Provisioning and Governance application.
Source Resource	Resource block that can present an **Endpoint** of the type that you specified. You typically do not change this setting.

2. Repeat the process for as many mappings as are needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure resource block input parameters

You can configure operations for each interface of a resource block by specifying input parameters, steps, and output attributes. Input parameters hold values that the system requires to identify and manage a virtual resource, such as the datacenter and resource group that the resource is in.

Before you begin

Role required: sn_cmp.cloud_service_designer

About this task

The system automatically associates default operations from the Cloud API (CAPI) to the resource block when you add interfaces to it. It also provides default input parameters for each operation.

Procedure

1. Open a resource block, and then click the **Operations** tab.

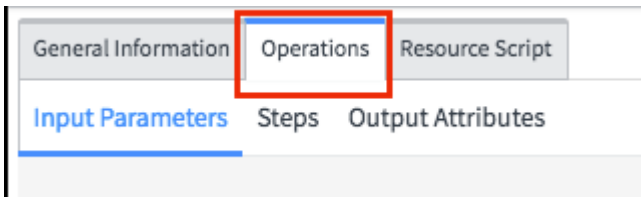

   The **Input Parameters** sub-tab appears by default.

2. Select the interface from the **Interface** choice list.

   The interfaces in this choice list are derived from the **Resource Guest Interface** list on the **General Information** tab that you already configured when you created a resource block.

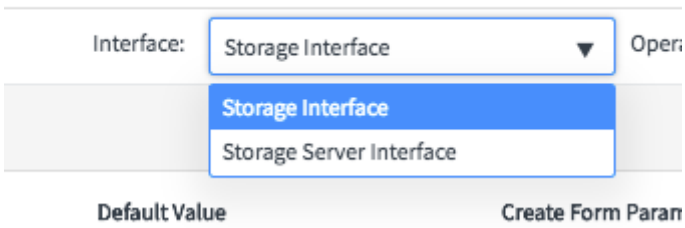

3. Select the operation on which you want to modify the input parameters. You can also create a new operation.

   a. Click the **Add Operation** icon.

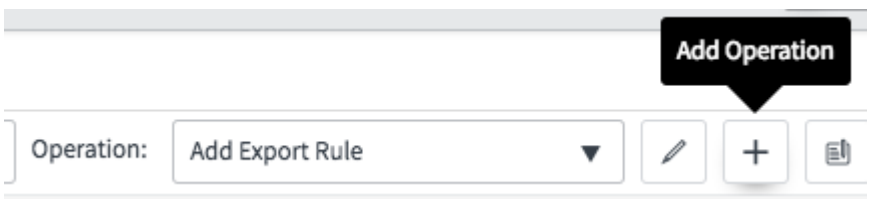

   b. Fill out the form (see table):

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	Enter a descriptive name.
Operation type	Select the type of operation from the choice list.
Access type	Select **Public** to make the resource block operation available from the catalog order form when a user manages a stack that they own. The actual operations appear in a list on the catalog order form: 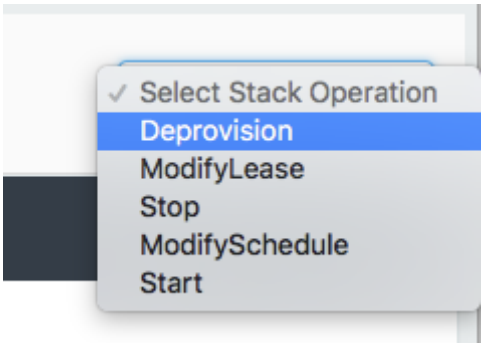 Select **Private** to prevent the operation from being used when a user launches a stack. This setting applies at the catalog-item level, not the whole-stack level. You can control the access type on the whole stack when you configure blueprints. As an example, you might want to mark the `Deprovision` operation as private for an individual VM, but not for the whole stack. This allows the user to deprovision everything at once and prevents the user from just deprovisioning the VM.

c. Click **Submit**.

4. Configure the input parameters to provide a way to input the items you need to run the selected operation.

For example, if you add a Provision operation for an AWS resource, you need, at minimum, the CloudAccount parameter and the Location parameter to perform the operation. You need to know the providers API to know what to pass to it. For example, if you want to add a provision operation for an S3 bucket on AWS, you need to know the allowable AWS operations on buckets and the Request Elements. For creating a bucket, see the PUT Bucket operation. You can typically find these resources in the AWS and Azure documentation.

a. Click the **Add Input Parameters** icon.

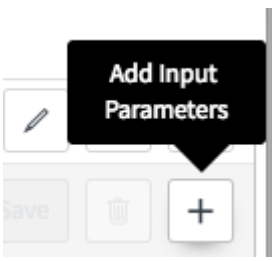

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Add the **Name** of the parameter that you obtained from the cloud provider API.
In this example, you could add `CreateBucketConfiguration` for an AWS S3 bucket.

c. In the **Datasource** column, select the type of input parameter you are configuring.
This value specifies the source or location where the parameter value comes from:

- **Text**: a text value that the user passes.
- **Pools**: the query that you want to run against a pool to get a value back from it. The pool **Datasource Value** should follow this convention: `ServiceNow::Pools::PoolName.Pool-Filter-Name`. See Pools and Filters for Cloud Provisioning for more information on pools.
- **CSVList**: values from a comma-separated value list that you enter into the **Datasource Value**.
- **Array**: an array of values. You can programmatically send it an array in JavaScript. The format is as follows: `['abc','def']`.
- **FormData**: data submitted from the form configured in the blueprint.
- **ResourceScript**: an expression that passes a value. The expression should follow this convention in the **Mapping** field:

  ```
  $(Script:script-name.function[arg1=${value-to-pass},arg2=
  ${value-to-pass}])
  ```

- **TLR**: the value inherited from the top-level resource container above this resource block in the blueprint.

d. Click **Save**.

e. In the **Mapping** column, you can view the auto-generated mapping values.
Mappings specify where to pull the information in the system. You do not have to change the mapping values except for **ResourceScript** type parameters, where you must enter the expression.

f. In the **Default Value** column, specify a value to use when no value is passed or entered in **Datasource Value**.

g. Select **Create Form** to make this attribute available in the order form that the user sees when trying to provision the stack item.
You can also select **Mandatory** to force the user to enter a value.

h. In the **Restricted To** column, click the restricted icon ( ), and then select a resource block to restrict the attribute to be passed to a specific host.

In this example, you are configuring a resource block for a virtual server on vSphere, which requires the name of the VM folder (represented by the `VmFolderName` attribute). Other datacenter types do not need the name of the VM folder. Therefore, you can restrict `VmFolderName` to the vSphere datacenter resource block.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This image illustrates how the VMFolderName attribute is restricted to the vSphere Datacenter.

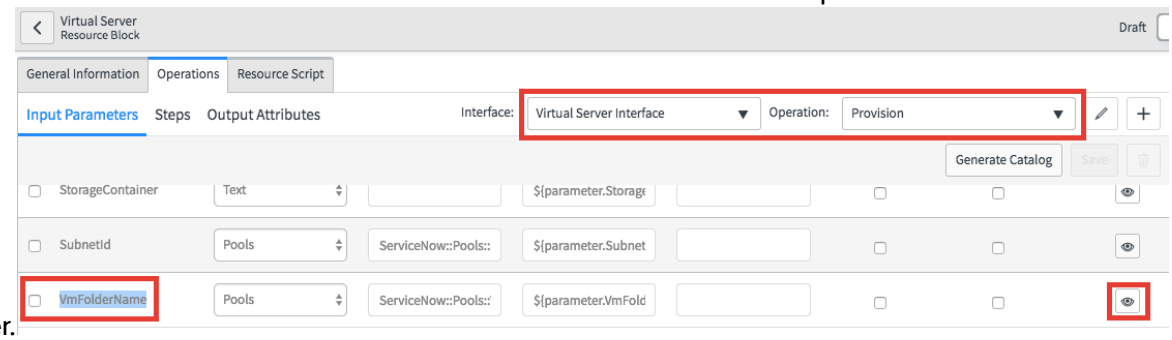

Add operation steps to a resource block

After you configure operations and input parameters for your resource block, add steps for each operation to tell the system which Cloud API (CAPI) to call. Each step is a separate call to CAPI.

Before you begin

Role required: sn_cmp.cloud_service_designer

Procedure

1. On the resource block, click the **Steps** sub-tab under **Operations**.
2. Click **Add Step**.
3. Fill in the form fields.

Field	Description
Operation Type	Select the type of operation: ○ **Invoke Cloud API**: Select this option to select a provider, interface, and method from CAPI. ○ **Invoke Host Resource Operation**: Select this option to use an operation (also called an operation signature) from the host interface that you already specified for this resource block. ○ **Invoke Script**: Select this option to invoke a method from an existing MID Server script include. ○ **Invoke Flow**: Select this option to invoke subflows from a list of active subflows.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
API Provider [Operation Type = Invoke Cloud API]	Select the API that provides the interfaces you need. The providers are part of the Cloud API (CAPI). For example, CAPI supplies the AWS provider, which has products, some of which have provider names. This value appears if you select the Invoke Cloud API operation type.
API Interface [Operation Type = Invoke Cloud API]	Select the interface that has the methods you want to use. For example, the **Blob Storage Interface** provides the **CreateContainer** interface, which the system uses during provisioning.
API Method [Operation Type = Invoke Cloud API]	Select the method to perform the operation. In this example, **CreateContainer** is required to provision storage.
CAPI Version [Operation Type = Invoke Cloud API]	Enter the version of the CAPI API.
Host Interface [Operation Type = Invoke Host Resource Operation]	Select the interface from which to call an operation. You can only select the same interface that you configured on the **General Information** tab of the resource block. To change the interface, navigate back to the **General Information** tab and select another value in the **Host Interface** field.
Interface Operation [Operation Type = Invoke Host Resource Operation]	Select the operation from the host interface. For example, if you select **Storage Server** Interface, you can call all the operations in that interface, such as *`CreateStorageVolume`*.
Scripts [Operation Type = Invoke Script]	Enter the script you want to invoke, along with a method in the script. Use this syntax: `script-name.method`.
Flow [Operation Type = Invoke Flow > Subflow]	Select the subflow from the Subflows list. For example, if you select a Terraform Enterprise Provision subflow, you can simply reuse the existing subflow to populate the Inputs&Outputs
Condition	Enter an optional condition that must be met before the system uses this step.

If you look at the Azure API and the CreateNode method mapper, you can see that the endpoint operation is a script. It calls the MID Server script include, which shows you the *`createNode()`* function. This function uses *`AzureComputeVirtualMachine`* script include. Take time to review and understand this script include.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Click **Submit**.

   The **Input** and **Response Processor** tabs appear for the step.

5. On the **Input** tab, configure the CAPI operation inputs.

   On the CAPI interface, which in this case is **Blob Storage** Interface, open `CreateContainer` method. The inputs from CAPI are shown on the Input tab, where you create parameters.

   Remember that the parameters on the **Input Parameters** tab is where users enter values. You can map those input parameters to the parameters on the **Input** tab. To do this, copy the **Mapping** value from the **Input Parameters** tab into the **Value** field on the **Steps** > **Input** tab.

   If this step is invoking a script, and that script requires arguments to return values, you can add a new input parameter for each argument. Use this syntax in the **Value** field: `${parameter.arg1}`.

6. If any subsequent steps in this operation or subsequent operations require parameters values from this step to be passed to them, configure output attributes.

   For example, the Provision operation for a virtual server can provide an output attribute for the network interface ID and another for the node ID. When the cloud provider sends the response, the actual ID of the network and node are parsed out of the response and saved in the output attributes. Subsequent steps use the IDs to perform actions.

   a. Click the **Output Attributes** tab.

   b. Enter code in the following notation:

   ```
   {
     "Attribute": "${Value}"
   }
   ```

   Where the value comes from a mapping to the interface and the operation.

   For example, in the following output attribute:

   ```
   {
     "NodeId": "${Compute
    Interface.CreateNode.Output.resp.nodeId}"
   }
   ```

   - `NodeId` is the attribute.
   - `Compute Interface` specifies the interface.
   - `CreateNode` is the name of the method that creates the node.
   - `Output.resp` is a mandatory part of the syntax.
   - `nodeId`

Add a resource operation step to invoke a subflow

Invoke a subflow by adding an operation step to a resource and linking it to a new or existing subflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Make sure the Orchestration application is installed.

Role required: Cloud user, designer, or admin

Procedure
1. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.
2. Set the resource block to **Draft** state by toggling the publish and then select **Operations** > **Steps** > **Add Step**.
   The Add Operation Steps form appears.
3. From the **Operation Type** list, select **Invoke Flow**.
4. From the **Content Type** list, select **Subflow**.

   Note: You cannot invoke flows while adding a resource operation step. This is a known limitation.

5. Select an active subflow from the **Flow/Subflow** lookup list.
6. **Optional:** Specify a condition in the **Condition** field.
7. Select **Submit**.
   The subflow operation step is attached to the resource block and appears on the page. Any input parameters associated with the subflow you selected are auto-populated on the **Input** tab.
8. In the **Inputs** subtab, select the + icon to add a new step parameter.

   a. Add the following key value pair to the subflow.

Key	Value
flowcorrelationid	$(Script:CMPFlowStepHandler.generateCorrel

9. Select the **Response Processor** tab and then select the plus icon.
   The Add Response Processor dialog box appears.
10. In the **Script Name** list, select a script for the response processor.
    For a script to appear in the **Script Name** list, the script should already have been created in the **Resource Script** tab. For more information, see Configure a response processor.
11. Select **Submit**.
12. Set the resource block to **Publish** state to incorporate your changes and publish the resource block.

What to do next
Create a response action for Cloud Provisioning and Governance

Extend Cloud Provisioning and Governance resource blocks with an override operation

You can extend your existing guest interface to create a new operation. You can also create a new guest interface with its own operations, parameters, and inputs.

Before you begin
Become familiar with resource block input parameters, operations, and steps. See Configure resource block input parameters and Add operation steps to a resource block for instructions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: sn_cmp.cloud_admin

About this task

The Cloud Provisioning and Governance application is bundled with several useful resource blocks. Create extensions to the existing resource block interfaces to add more functionality to these resource blocks. After creating an interface extension, use it in a blueprint to execute a customized operation. For example, if you do not like the way a virtual server is provisioned, extend the virtual server resource block interface by adding operations or overriding existing operations.

Note: Operations in resource blocks are also called operation signatures.

Procedure

1. Open a resource block.
2. Change the status of the resource block from **Published** to **Draft**.
3. Click the **Operations** tab.
4. Click the **Override Operation** icon.

5. Select an existing guest interface or select the **Create new** check box to create a new guest interface.
6. **Optional:** If you create an interface, enter a name in the **Interface** field.
7. Click **Submit** or **Update**.
   The new operation appears in the **Operation** list. If you create a new interface, it appears in the **Interface** list.
8. Configure the necessary input parameters, steps, and output attributes for your new guest interface.
9. To edit the operation or to change the access type:

   a. Click the **Edit Operation Signature** icon.

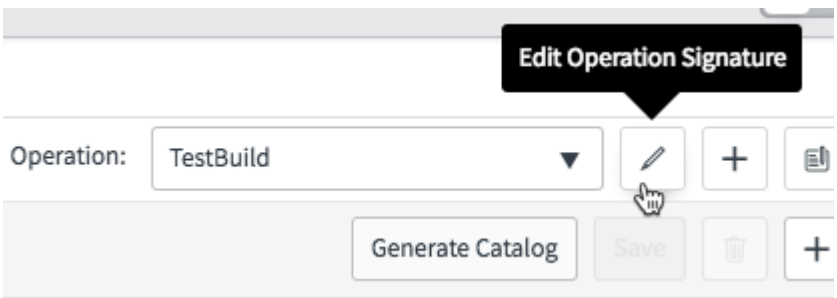

   b. Change the **Operation Type** if necessary, and add, edit, or remove input parameters at the bottom.

   c. If you want the blueprint designer to perform the operation directly on the resource only when it is instantiated, select **Private** for the **Access Type**.

   d. Click **Update**.

10. After completing the customizations, toggle the switch to **Published**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Navigate to **Cloud Admin Portal** > **Design** > **Blueprints**.
12. Follow the steps in Build a Cloud Provisioning blueprint to add a blueprint.
13. Click **Operations** and click the operation to modify in the resource block.
    The default resource block interface appears in the right column.
14. Click the **Edit Resource Operation** (pencil) icon to apply the interface extension you created.
15. Click **Submit**.
    The blueprint gets executed with the customized interface extension.

Result

After you add this resource block to a blueprint, the parameters you added are available for you to customize.

Add a resource script to a resource block

A resource script operates on a resource during deployment or returns data to the CMDB after a resource is deployed.

Before you begin

Role required: sn_cmp.cloud_service_designer

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.
2. Open the resource block and then set it to **Draft**

state.
3. On the **Resource Script** tab, specify a unique and meaningful **Name** and then define the script:

Script	Script code.
Type	Type of script: ○ Instance Script: Javascript that performs an operation on the instance (for example, a subflow). ○ PowerShell Script ○ Python Script ○ Response Processor: Converts return data from the external cloud API to the operation implementation into a format that the CMDB can accept. ○ Shell Script

4. Click **Submit**.
   Add as many scripts as are needed.

Configure a response processor

Configure a response processor to handle responses from datacenters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- The CI class types for your resources must already exist in the CMDB. If you created custom products or providers in CAPI, you must also create custom CI types so the response processor knows how to identify CIs. See Create a CI class for a virtual cloud resource for information on creating CI types.

Role required: sn_cmp.cloud_service_designer

About this task

After each step is completed, the datacenter sends a response to the ECC Queue of the MID Server. For example, if a user provisions a virtual server on an AWS datacenter, AWS sends information about the new virtual server to the ECC Queue. The system needs to create a CI in the CMDB for the virtual server. It does this using a response processor.

Alternatively, you can use pattern-based targeted discovery to handle responses from the datacenters. For more information, see Pattern-based targeted discovery.

Note: The resource type for which you are adding the response processor script must exist in the [sn_capi_resource_type] table.

Procedure

1. Open a resource block, and click the **Resource Script** tab.
   Ensure that the resource block is not in the published state.
2. Click **New**.
3. In the **Name** field, enter the name for the resource script.
4. In the **Script** field, enter a script.
5. Click **Submit**.
6. Click the **Operations** tab on the resource block.
7. Click the **Steps** sub tab.
8. Click a step.
9. Click the **Response Processor** tab for the step.
10. Click the add icon (+).
11. In the **Script Name** field, select the processor you just created.
12. Click **Submit**.
    The response processor script can process the response from the cloud provider for the CI type of the resource block.

Execute response processor for subflow

Execute a response processor for a subflow to get the subflow data back into a configuration item (CI). The response processor picks up the data, sends the data to the CMDB, which in turn puts the data in a CI.

Before you begin

Role required: sn_cmp.cloud_service_designer

About this task

Before you execute a response processor for a subflow, you must create a subflow, attach the subflow to a resource block operation step, and then generate the catalog. To return a response from the subflow to the Cloud Provisioning and Governance application, the subflow must create a record in the sn_cmp_flow_result table with the appropriate details as follows.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
flow_type           : 'sys_flow_context'
flow_correlation_id : flowcorrelationid (This variable must
be present in the Resource Block as mentioned here Add a resource
operation step to invoke a subflow .
flow_output         : String that will be processed by Response
Processor
flow_error          : Error string if there is any error
error_detail        : Error detail string if there is any error
```

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.
2. Open a resource block that is in a draft state and navigate to **Operations Steps**.
3. Select **Add Step**.
   The Add Operation Steps dialog box appears.
4. Add a workflow operation step.
   See Add operation steps to a resource block.
   The workflow operation step gets attached to the resource block and appears on the page. Any input parameters associated with the workflow appear on the **Input** tab.
5. Select **Generate Catalog**.
   The workflow appears in the Cloud User Portal as an operation. Select the workflow from the **Select Operation** picker to execute the operation. The status of the operation is visible in the Track operation subtab.
6. Select the **Response Processor** tab and then select the plus icon.
   The Add Response Processor dialog box appears.

7. In the **Script Name** list, select a script for the response processor.
   For a script to appear in the **Script Name** list, the script should already have been created in the **Resource Script** tab.
8. Select **Submit**.
   The script appears in the **Response Processor** tab. You can open the script and modify the script.

Virtual server response processor example

The *Create_Virtual_Server_Response_Processor* script, which is available by default in Cloud Provisioning and Governance, is the response processor that handles the creation of AWS virtual machine CIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create Virtual Server

The Create_Virtual_Server_Response_Processor resource processor script is available by default on the Virtual Server resource block. Its job is to create a virtual server record in the Virtual Machine Instance [cmdb_ci_vm_instance] table when a new virtual server is provisioned.

All response processors have this first line with these common parameters:

```
function processResponse(response, cloudServiceAccountId, ldc,
 correlationId,step, requestorContext) {
```

This brings in the response from the cloud provider and the important information, such as the account ID, that is required for the new CI that the system can create. All of these parameters are required for all response processors.

Line 10 parses the response into JSON so that the system can process it. The information is held in the vmResponse variable:

```
var vmResponse = global.JSON.parse(response);
```

Whenever you create or edit a response processor, you must know which inputs are required for the CI type. Line 11 handles one of the necessary inputs, the hardware ID, that the CMDB record requires:

```
var hardwareId = vmResponse.hardwareId;
```

Line 39 shows the information that is required for the system to identify the new virtual server and related CIs, so the information can be put into the CMDB. In this case, the service account object ID identifies the account associated with the virtual server, the datacenter object ID identifies the datacenter in which the virtual server lives, and the virtual machine instance object ID identifies the virtual server itself. This identification code block prevents the creation of duplicate CIs.

```
var vmInfo = {
    "cmdb_ci_vm_instance": {
      "validator": "virtual_machine_create_update_validator",
      "validator_overrides": {},
      "identification": {
        "cmdb_ci_cloud_service_account": {
          "criterion": {
            "object_id": cloudServiceAccountId
          }
        },
        "cmdb_ci_aws_datacenter": {
          "criterion": {
            "object_id": ldc
          }
        },
        "cmdb_ci_vm_instance": {
          "criterion": {
            "object_id": vmResponse.nodeId
          }
        }
      },
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attributes are populated into the fields on in the cmdb_ci_vm_instance table. These attributes are defined in line 61:

```
"attributes": {
  "name": vmResponse.nodeName,
  "object_id": vmResponse.nodeId,
  "state": status_map[vmResponse.state],
  "dns_suffix": vmResponse.dnsSuffix,
  "cpus": vmCPUs,
  "memory": vmMemory,
  "disks": vmResponse.volumes.length,
  "disks_size": "",
  "nics": vmResponse.networkInterfaces.length,
  "terminated_on": "",
  "termination_protection": "",
  "ip_address": vmPubIPAddr,
  "assigned_to": reqContext.userId,
  "assignment_group": reqContext.groupId
},
```

References to other CIs can also be included in the response processor. In this case, the OS template that the virtual server is based on is identified by first identifying the object ID of the service account and the datacenter along with the actual OS template.

```
"references": {
  "cmdb_ci_os_template": {
    "identification": {
      "cmdb_ci_cloud_service_account": {
        "criterion": {
          "object_id": cloudServiceAccountId
        }
      },
      "cmdb_ci_aws_datacenter": {
        "criterion": {
          "object_id": ldc
        }
      },
      "cmdb_ci_os_template": {
        "criterion": {
          "object_id": imageIdTrim
        }
      }
    },
```

The following code block adds the object ID of the OS image to the attributes list so that this information can be populated into the virtual server CMDB record.

```
"attributes": {
  "object_id": imageIdTrim
}
```

This code block performs additional identification on the Compute template (the hardware type) and then add it to the attributes:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
"cmdb_ci_compute_template": {
  "identification": {
    "cmdb_ci_cloud_service_account": {
      "criterion": {
        "object_id": cloudServiceAccountId
      }
    },
    "cmdb_ci_aws_datacenter": {
      "criterion": {
        "object_id": ldc
      }
    },
    "cmdb_ci_compute_template": {
      "criterion": {
        "object_id": vmResponse.hardwareId
      }
    }
  },
  "attributes": {
    "object_id": vmResponse.hardwareId,
    "name": vmResponse.hardwareId
  }
}
```

Additional code sections make the relationship with network interfaces and identify any storage attached to the virtual server.

This mandatory code block pushes the data to the CMDB and returns the JSON string:

```
cloudModelString.push(vmInfo);
return global.JSON.stringify(cloudModelString);
```

Add Terraform Orchestration interface on a resource block

After you have activated the Terraform store app for Cloud Provisioning and Governance, to support cloud-based operations for cloud providers using Terraform Open Source Edition. Add the Terraform Orchestration interface on resource block, to execute cloud-based operations via terraform templates.

Before you begin

Role required: sn_cmp.cmp_root_admin, sn_cmp.cloud_service_designer

Procedure

1. On the Cloud Admin portal, navigate to **Design** > **Resource Block**.
2. Click the resource block, for example **Azure Datacenter**.
3. Move the state of the resource block from **Published** to **Draft** to make it editable.
4. Navigate to **General Information** > **Guest Interface**, and then click **New**.
5. Search for and select the **Terraform Orchestration Interface** from the Guest Interface list, and click **Submit**.
6. Click the **Operations** tab.
7. In the **Input Parameters**tab, navigate to **Interface** and select **Terraform Orchestration Interface**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. On the **Input Parameters** tab, configure the CAPI operation inputs.

   On the CAPI interface, which in this case is **Terraform Orchestration Interface**, open *`CreateStack`* method. The inputs from CAPI display on the Input tab, where you create parameters.

   The **Input Parameters** tab is where users enter values. You can map those input parameters to the parameters on the **Input** tab. To do this, copy the **Mapping** value from the **Input Parameters** tab into the **Value** field on the **Steps** > **Input** tab.

   If this step is invoking a script, and that script requires arguments to return values, you can add a new input parameter for each argument. Use this syntax in the **Value** field: `${parameter.arg1}`.

9. Click the **Steps** sub-tab under **Operations**.
10. Click **Add Step**.
11. Fill in the form fields.

Field	Description
Operation Type	Select the type of operation. **Invoke Cloud API**: Select this option to select a provider, interface, and method from CAPI.
API Provider	Select the API that provides the interfaces you need. The providers are part of the Cloud API (CAPI). For example, CAPI supplies the **Terraform OpenSource**provider, which provides the **Configuration Orchestration Interface**. This value appears if you select the Invoke Cloud API operation type.
API Interface	Select the interface that has the methods you want to use. For this example, the **Configuration Orchestration Interface** provides multiple interfaces or methods, which the system uses during provisioning.
API Method	Select the method to perform the operation. In this example, **ExecuteCommand** is required to provision terraform templates.
CAPI Version	Enter the version of the CAPI API.
Condition	Enter an optional condition that must be met before the system uses this step.

12. Click **Submit**.

    The **Input** and **Response Processor** tabs appear for the step.

13. On the **Input** tab, configure the CAPI operation inputs.

    On the CAPI interface, which in this case is **Configuration Orchestration Interface**, open *`ExecuteCommand`* method. The inputs from CAPI show on the Input tab, where you create parameters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remember that the parameters on the **Input Parameters** tab is where users enter values. You can map those input parameters to the parameters on the **Input** tab. Copy the **Mapping** value from the **Input Parameters** tab into the **Value** field on the **Steps** > **Input** tab.

If this step is invoking a script, and that script requires arguments to return values, you can add a new input parameter for each argument. Use this syntax in the **Value** field: `${parameter.arg1}`.

Note: Map the values as mentioned in the following table for the following parameters:

Parameter	Value
AdditionalParameters	Enter the following script: `{"ServerType":"$(Script:sn_cmp_terraform.TerraformCommandUtils.getConfigProviderServerTypeByName[arg=${parameter.WorkloadConfigProvider}])"}`
ConfigMgmtProviderInfo	`$(capiResolver.NodeCredentialResolver#nodeCredentialId=$(ci.sn_cmp_wl_cfg_mgmt_provider[name=${parameter.WorkloadConfigProvider}].credential))`
Node Address	`$(ci.cmdb_ci_workload_config_provider[name=${parameter.WorkloadConfigProvider}].url)`
Script For CreateStack operation	`${CloudScript.scripts.TerraformDeploy}`
Script For DeleteStack operation	`${CloudScript.scripts.DestroyTerraform}`

14. On the **Response Processor** tab, configure the response processors by attaching resource scripts.
    You would have added Resource Scripts to the resource block, while creating the Resource Block.
15. From the Operation list, select the **CreateStack** or **DeleteStack** operation.
16. Click the add icon + for **Add Response Processor**.
17. Click the **Resource Script** sub tab, and then click **New**.
18. Complete the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New Record

Field	Description
Name	Enter a name (For example, **Create_AzureTF_Stack_Response_Processor**) for the resource script.
Type	Select a type (**Response Processor**) for the script.
Script Resource Script for Create	Enter the following script: `function processResponse (response, cloudServiceAccountId, ldc, correlationId, step, requestorContext, stackId) { var responseObject = global.JSON.parse(response); var tfState = global.JSON.parse(responseObject.terraform.state); var processor = new sn_cmp_terraform.TerraformResponseProcessor(); var processedResponses = processor.processResponse(tfState, cloudServiceAccountId, ldc, correlationId, step, requestorContext, stackId); return global.JSON.stringify(processedResponses); }`
Script Resource Script for Delete	Enter the following script: `function processResponse (response, cloudServiceAccountId, ldc, correlationId, step, requestorContext, stackId) { var responseObject = global.JSON.parse(response); var tfState = global.JSON.parse(responseObject.terraform.state); var processor = new sn_cmp_terraform.TerraformResponseProcessor(); var processedResponses = processor.processResponse(tfState, cloudServiceAccountId, ldc,`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	correlationId, step, requestorContext, stackId); return global.JSON.stringify(proces sedResponses); }

19. From the **Script Name** list, select the script corresponding to the selected operation, and click **Submit**.
20. Move the state of the resource block from **Draft** to **Published**.
    The datacenter resource block moves from the draft to the published state and you cannot make further changes to the resource block. You can only modify a resource block after you set the status to **Draft**.

Create a resource block for Microsoft Azure Cloud

Resource blocks are the building blocks of cloud catalog items. Create a resource block for the Microsoft Azure Cloud Platform provider, based on a CI type from the CMDB.

Before you begin

Role required: sn_cmp.cmp_root_admin

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.
2. Click **New** to open the Create Resource Block dialog box and then fill in the form.

   For a detailed description of each field, see Create a custom resource block.

Field	Description
Name	Enter a name (**Azure Datacenter**) for the resource block.
Description	Enter a valid description
Layer	Select the type of resource block (**Logical Datacenter**).
CI Type	Select the CI type (**Azure Datacenter**) that the resource block represents.
Guest Interface	Select a default guest interface (**Compute Interface**) to use when you add a resource block to a template.

3. Click **Submit**.
   The resource block information appears.
4. Click the **Resource Script** sub tab and then click **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	Enter a name (For example, **Create_AzureTF_Stack_Response_Processor**) for the resource script.
Type	Select a type (**Response Processor**) for the script.
Script Resource Script for Create	Enter the following script: `function processResponse (response, cloudServiceAccountId, ldc, correlationId, step, requestorContext, stackId) { var responseObject = global.JSON.parse(response); var tfState = global.JSON.parse(responseObject.terraform.state); var processor = new sn_cmp_terraform.TerraformResponseProcessor(); var processedResponses = processor.processResponse(tfState, cloudServiceAccountId, ldc, correlationId, step, requestorContext, stackId); return global.JSON.stringify(processedResponses); }`
Script Resource Script for Delete	Enter the following script: `function processResponse (response, cloudServiceAccountId, ldc, correlationId, step, requestorContext, stackId) { var responseObject = global.JSON.parse(response); var tfState = global.JSON.parse(responseObject.terraform.state); var processor = new sn_cmp_terraform.TerraformResponseProcessor(); var processedResponses = processor.processResponse(tfState, cloudServiceAccountId, ldc,`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description

	``
	correlationId,
	step, requestorContext,
	stackId);
	return
	global.JSON.stringify(proces
	sedResponses);
	}
	``

5. Click **Submit**.
6. Click the **Operations** sub tab and then click **Steps**.
7. Select **CreateStack** from the **Operations** list.
8. Click the **Add a Step** plus icon to open the Add Operation Steps dialog box and then fill in the form.

Field	Description
Operation Type	Select the type of operation. **Invoke Cloud API**: Select this option to select a provider, interface, and method from CAPI.
API Provider [Operation Type = Invoke Cloud API]	Select the API that provides the interfaces you need. The providers are part of the Cloud API (CAPI). For example, CAPI supplies the **Terraform OpenSource**provider, which provides the **Configuration Orchestration Interface**. This value appears if you select the Invoke Cloud API operation type.
API Interface [Operation Type = Invoke Cloud API]	Select the interface that has the methods you want to use. For this example, the **Configuration Orchestration Interface** provides multiple interfaces or methods, which the system uses during provisioning.
API Method [Operation Type = Invoke Cloud API]	Select the method to perform the operation. In this example, **ExecuteCommand** is required to provision terraform templates.
CAPI Version [Operation Type = Invoke Cloud API]	Enter the version of the CAPI API.
Condition	Enter an optional condition that must be met before the system uses this step.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. Click **Submit**.
10. Click **Response Processor** and then click the plus icon.
    The Add Response Processor dialog box appears.
11. Select **Create_AzureTF_Stack_Response_Processor** from the **Script Name** list and click **Submit**.
12. Move the state of the resource block from **Draft** to **Published**.
    The Azure Datacenter resource block moves from the draft to the published state and no further changes can be made to the resource block. Changes to a resource block can be made only if the status is set to **Draft**.

Cloud scripts and cloud script templates

In the Cloud Provisioning and Governance application, script execution is divided into cloud scripts and cloud script templates. Use scripts in blueprints, resource blocks, OS profiles, and use policy scripts to set request form attributes. Policy scripts cannot override user data.

Cloud scripts

Cloud scripts are simple java scripts that use platform features. An example of a sample cloud script:

```
function evaluateTemplate() {
 // Template Resolver is a helper function which we will use to
 fetch a template and replace its parameters,
 // var templateResolver=new TemplateResolver(); this comes
 inbuild
 // templateAttributes below will passed from the user.
 var
 listDir=templateResolver.getTemplate('Install',templateAttribut
es);
 return listDir;
}
```

Each cloud script contains certain parameters. Parameters are the attributes that you want to access. A parameter contains a name, a default value, and an override value. The default value can be a string literal, a resource block parameter or a complex script expression used in resource blocks.

A cloud script, by default, has access to the *templateResolver* object. The *templateResolver* object has, by default, a getTemplate method, whose first parameter is the cloud script template name. *templateAttributes* are the cloud script parameters created as part of a cloud script. If needed, customizations can be done in *templateAttributes*. As an example, see the following code snippet if the value of *DevName* parameter needs to be changed from John to Emily.

```
function evaluateTemplate() {
 // Template Resolver is a helper function which we will use to
 fetch a template and replace its parameters,
 // var templateResolver=new TemplateResolver(); this comes
 inbuild
 // templateAttributes below will passed from the user.
 templateAttributes['DevName'] = 'Emily';
 var customNodeName;
 var now_GR = new GlideRecord('some_custom_table_to_query');
 gr.addQuery('some_parameter','some_value');
 gr.query();
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 if(gr.next())
   customNodeName=gr.getValue('some_custom_node_name');

 templateAttributes['NodeName'] = customNodeName;
 var
 listDir=templateResolver.getTemplate('Install',templateAttribut
es);
 return listDir;
}
evaluateTemplate();
```

Cloud script templates

Cloud script templates are actual executables which are passed to target a virtual machine for execution. Cloud templates can be of any type (Shell/PowerShell/CloudInit) depending on the execution context. You need to create a cloud template first and then associate it with a cloud script.

Cloud scripts as CloudInit

Cloud providers provide a way to execute scripts at boottime for virtual machines. In the Cloud Management application, you can specify a cloudinit script at various levels:

- Image level: At the time of creating a cloud script, you can specify an image against which this script should be executed. Once an OS profile which contains this specific image gets selected, the default cloud script against this image is executed. In the example below, an Apache cloud script is associated with a particular image.

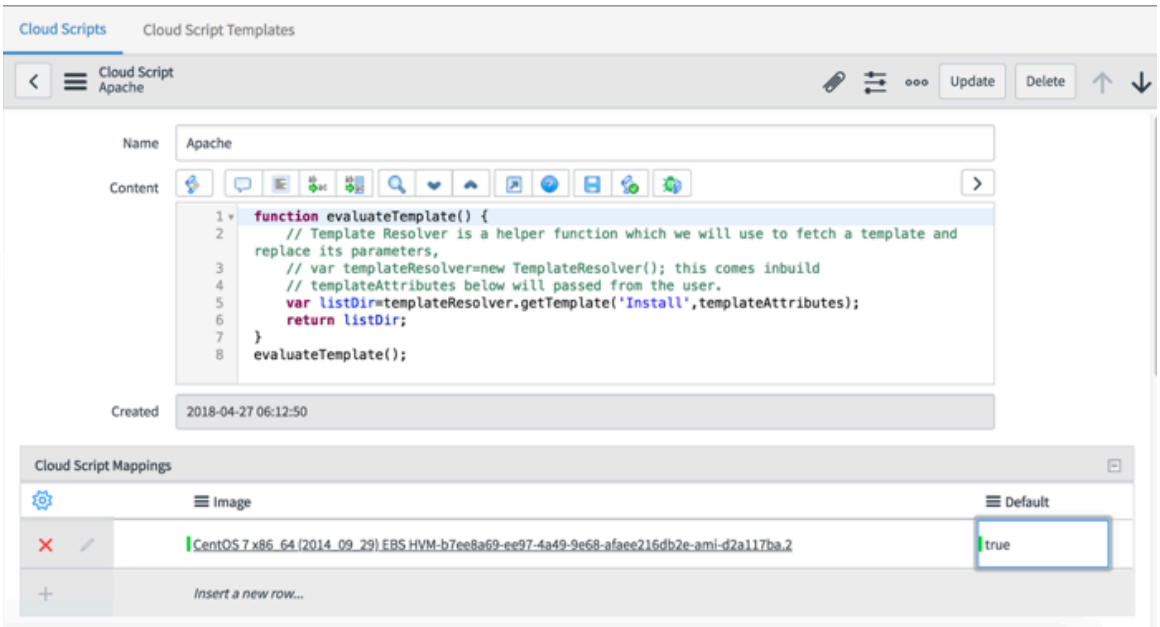

- OS profile level: You can run a cloudinit against a specific cloud provider and against a specific AMI/image. You can also be generic and associate a cloudinit to an OS profile, irrespective of the underlying image. In the following example, multiple images are associated with the OS profile. You can map any cloud script to the OS profile in the Cloud Script OS Profile Mappings tab. The mapping of a cloud script with an OS profile takes precedence over the image-level cloud script association.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

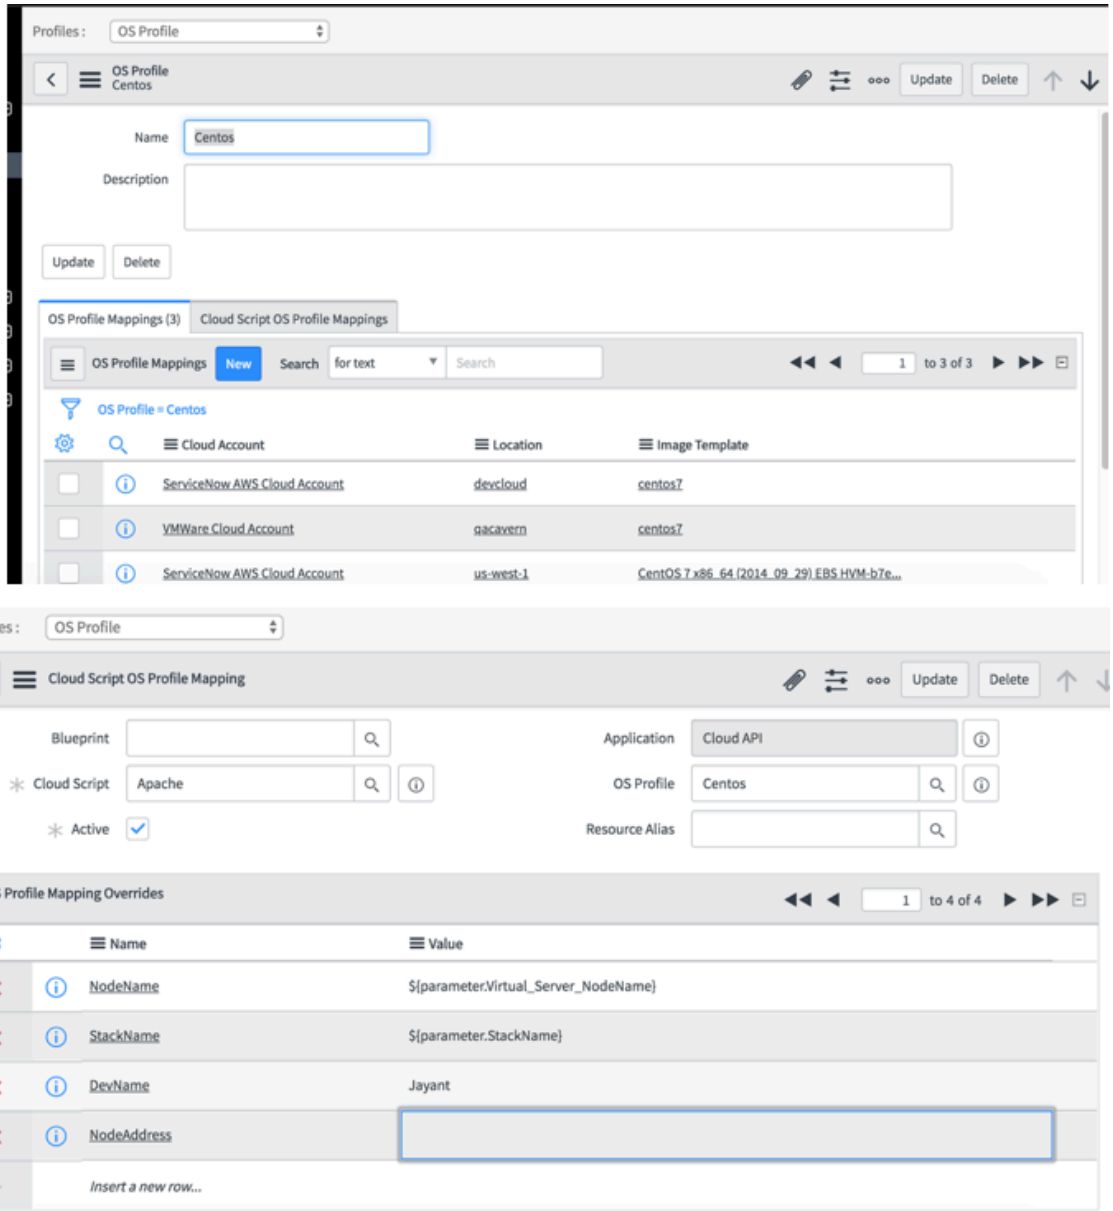

In this example, the Apache cloud script is mapped to the Centos OS profile. If needed, you can override the script parameters in the OS Profile Mapping Overrides section.

- Resource blocks: While mapping a cloud script to an OS profile, you can choose whether to execute the cloud script for a specific resource block used in a blueprint and against a particular resource alias. See Create an OS profile.

Cloud scripts as Postinit

You can use cloud scripts as Postinit scripts in a resource block. Once you create a resource block with a virtual machine to be provisioned, you can add an ExecuteScript operation. The ExecuteScript operation takes a script parameter (a pool of cloud scripts) and you can select any script to execute. Use the *`ScriptParameters`* attribute in the ExecuteScript operation to override any script parameter. You can provide a JSON map of script parameter and its override value. All expressions that are supported by Cloud Management work in the *`ScriptParamters`* attribute.

Create cloud initialization script templates and a script

Create cloud initialization templates and a script to use during VM provisioning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: cloud_admin

Procedure
1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Scripts** > **Cloud Script Templates**.
2. Click **New**, enter the name and content of the script, and then click **Submit**.

   Create as many script templates as needed.
3. In the Cloud Admin Portal, navigate to **Design** > **Cloud Scripts**.
4. Click **New**.
5. Enter a descriptive name and create script that uses the template name in the following line of code (*`scriptName`* is the name of the template):

   ```
   templateResolver.getTemplate(scriptName);
   ```

   Note: For more information, see Cloud scripts and cloud script templates.
6. Click **Submit**.

What to do next
To run a script at the time a virtual resource is provisioned, map the script to an OS profile. To run a post-provisioning script, create a blueprint and specify the script in the Execute Script operation. See Example: Post-provision cloud script for an example.

Map a script to an OS profile

To execute scripted actions during VM provisioning, you can map a script to an OS profile. The script runs on VMs that are created based on the image template in the OS profile.

Before you begin
Role required: sn_cmp.cloud_admin

You must have a cloud account with datacenters. You must run Discovery on the service accounts to populate the datacenters.

About this task
In this procedure, you specify an existing script. See Create cloud initialization script templates and a script.

Note: Profile mappings that specify more details run first. For example, a mapping that specifies a blueprint, OS profile, and resource alias takes precedence over a mapping that specifies only an OS profile.

Procedure
1. In the Cloud Admin Portal, navigate to **Manage** > **Resource Profiles**.
2. In the **Profiles** list, select **OS Profile** and then open the profile.
3. On the Cloud Script OS Profile Mappings related list, click **New** and then fill in the form.

Field	Description
Blueprint	Select a blueprint to limit the script to a specific resource block used in the blueprint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
[Optional]	
Cloud script	Select a cloud script to map the OS profile to.
Active	Select the check box if the cloud script should be run after the virtual machine is provisioned.
Application	Cloud Provisioning and Governance is auto-selected.
OS profile	If you specify a blueprint, the cloud script is run when the blueprint is provisioned.
Resource Alias [Optional]	If you specify a resource alias for the blueprint, then the cloud script is executed when the blueprint with the specified resource alias is provisioned.

4. Right-click the form header and select **Save**.
   The **Name** attribute `[scriptName]` appears in the OS Profile Mapping Overrides list.
5. In the OS Profile Mapping Overrides list, enter an attribute name and value to use when the resource is provisioned.

Related topics

Create an OS profile

Example: Post-provision cloud script

This example shows a cloud script that runs after a user provisions a virtual machine.

Before you begin

- Role required: cloud_admin
- OS resource profile that is mapped to an image template. The user provisions a resource using the profile. You must add basic authentication credentials to the image.
- Compute resource profile that is mapped to a hardware type.

About this task

In this example, you create a cloud init script, customize a resource block, and provision a VM from the Cloud User Portal.

Procedure

1. Create a cloud script template and a cloud script that refers to the template.
2. Open the virtual server resource block, and then follow these steps:

   a. Set the **Draft/Published** switch to **Draft**.

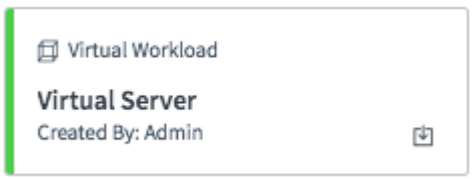

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. On the **Operations** tab, select **Virtual Server Interface** from the **Interface** list.

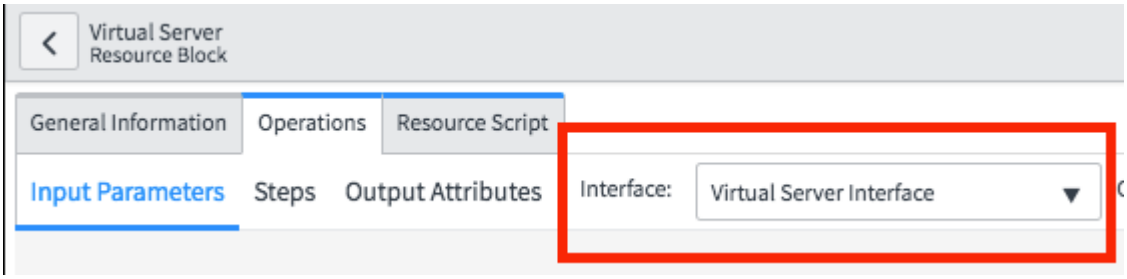

c. In the **Operation** field, verify that these operations exist:

- **Block Until Ready**
- **Execute Script**: This operation tells the blueprint to run a script.

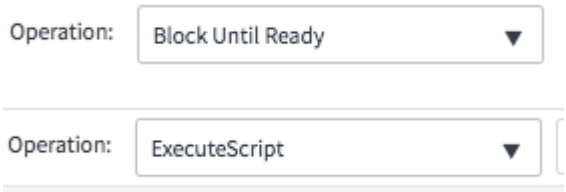

d. Click the **Input Parameters** tab to view the inputs for each operation.

e. On the **Steps** tab, add steps as needed.

f. For the **Script** parameter on each step, clear the **Script** check box and change the expression in the **Mapping** column to the following text:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

`$(CloudScript.scripts.scriptname)`, where `scriptname` is the name of the script, not the script template, that you created earlier.

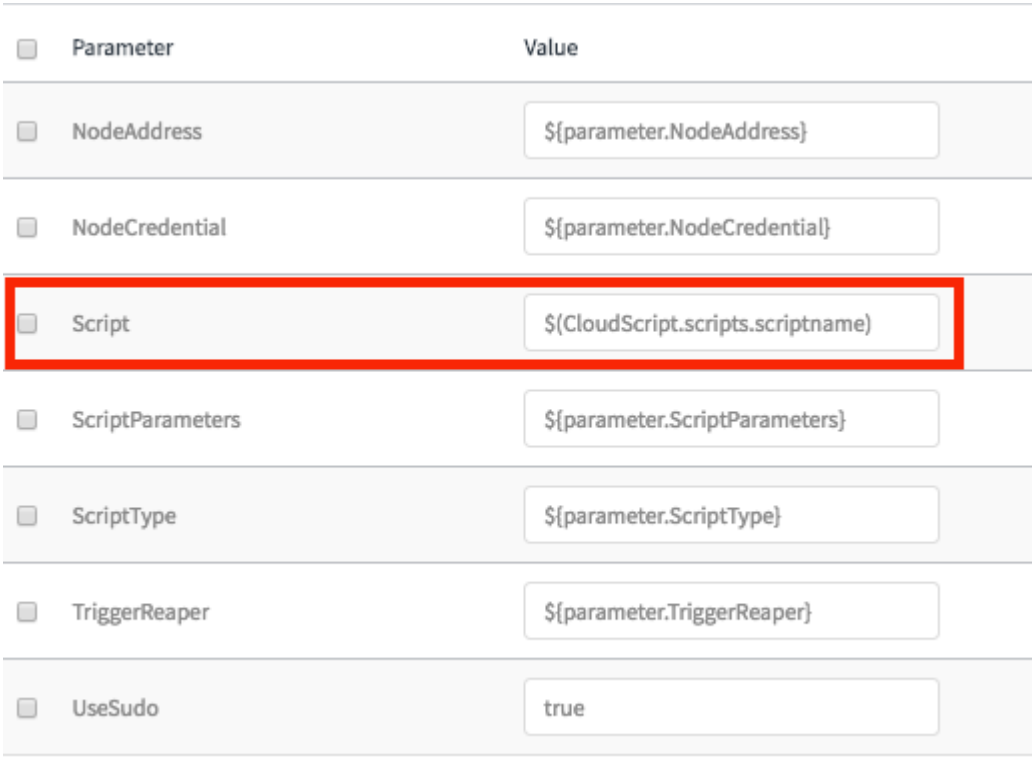

3. Publish the virtual server resource block, and use it in a blueprint that in turn is published to a catalog item.
4. Launch a new stack to provision a VM and specify the OS profile and Compute profile that you created as a prerequisite.

Result
The Stack Status page runs through the *`CreateNode`*, *`BlockUntilReady`*, and *`ExecuteShellScript`* operations. The last operation kicked off the script you created.

Response Mappings

Response Mapping is the mechanism CMP uses to write back all the events and operations into CMDB.

Response Mapping mechanism

CMP (Cloud Management Platform) is designed to actively engage with the cloud environment at various critical stages, including Discovery, Provisioning, Day-2 operations, and Lifecycle events. Throughout these processes, CMP is required to effectively record and store data within the Configuration Management Database (CMDB). This data could encompass diverse information ranging from system states to comprehensive sets of data.

To achieve this, CMP leverages its Response Processing mechanism to ensure the seamless and accurate transfer of data back into the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Features of Response Mapping

- Maps external data format into ServiceNow format for IRE consumption.
- Supports simple, complex, scripted, binding and relationship mappings.
- Significant to add new/change mappings.
- Performant, used in CMP for the past few releases.

See https://www.servicenow.com/community/itom-articles/cmp-using-response-mapping-how-to-populate-cmdb-using-cmp/ta-p/2323213 to learn more.

Validate a parameter value with the AllowedPattern property

The *AllowedPattern* property in an Azure Resource Manager (ARM) and CloudFormation (CF) template performs input validation before sending a provisioning request.

Before you begin

Role required: sn.cmp.cloud_service_designer

About this task

You can provide rules to validate a parameter value before sending the provisioning request to the cloud. To be notified with an error message if the validation fails, specify the *ConstraintDescription* property in the template. You can customize the error message in the *ConstraintDescription* property to make the message clear and explicit. Usually, the CF templates already have the *AllowedPattern* and *ConstraintDescription* properties. If these properties are not present, you can add them as mentioned in the following procedure:

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Templates**.
2. Select a template or click **New** to create a new cloud template.
   If you select an existing template, rename the template.
3. In the **Body** text box, make the following changes in an ARM and a CF template.

   a. In an ARM template, make these changes:

   - In the parameters metadata section, enter *SNC::Parameter::Metadata*.
   - Add the *AllowedPattern* key and enter a valid value range.
   - Add the *ConstraintDescription* key and enter an error message.

```
"parameters": {
        "newStorageAccountName": {
            "type": "string",
             "metadata": {
                 "description": "Unique DNS Name for the
Storage Account where the Virtual Machine's disks will be
placed.",
                  "SNC::Parameter::Metadata": {
                      "allowedPattern": "[0-9a-z]{1,11}",
                       "ConstraintDescription": "must be an
alphanumeric "

                          }}}
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. In a CF template, make these changes:

- Add the *AllowedPattern* key and enter a valid value range.
- Add the *ConstraintDescription* key and enter an error message.

```
"Parameters" : {
                "SSHLocation" : {
                "Description" : "The IP address range that
 can be used to SSH to the EC2 instances",
                "Type": "String",
                "MinLength": "9",
                "MaxLength": "18",
                "Default": "0.0.0.0/0",
                 "AllowedPattern":
 "(\\d{1,3})\\.(\\d{1,3})\\.(\\d{1,3})\\.(\\d{1,3})/(\
\d{1,2})",
                 "ConstraintDescription": "must be a valid IP
 CIDR range of the form x.x.x.x/x."

    }
  }
```

4. Click **Submit**.

Cloud Governance

Governance refers to the limitations that you can set on available cloud resources.

Types of governance

Type	Description
Quotas	Quotas are limitations or requirements for groups and users on CIs or cloud resources. Use quotas to prevent wasteful resource usage in your cloud environment.
Policies	A cloud policy can override a property value set by a user, create an approval task, reserve an IP address, pre-populate or hide form fields, execute custom scripts, call the Cloud API, or start or abort subflows. A cloud policy gives you system-wide control over approvals, resource operations, blueprint operations, or catalog item settings.
Pools	A resource pool is a query or script that filters a table. You configure a resource pool to limit the values that are available to users when they request a catalog item.
Permissions	Permissions are user group-level access rights to features in the Cloud Provisioning and Governance application and to specific records in the instance, such as blueprints or cloud accounts.

Permissions management for Cloud Governance roles

Permissions are user group-level access rights to features in the Cloud Provisioning and Governance application and to specific records in the instance, such as blueprints or cloud accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can refine the actions that are allowed or prohibited for users based on user group. By default, each role includes access rights, but not all users with the same role can see or edit each others records. For example, by default users with the cloud_designer role have full read and write access to their own blueprints, but they do not have read or write access to blueprints created by other cloud designers. If you want them to have access, you must put the users into a group that has the cloud_designer role, then give that group read and write access to particular blueprints.

Permission types

Access [Read]

Users can see the record, but not edit or delete.

Manage [Create & Update]

Users can see the record and create and update new records in the same table, but cannot delete any records.

Delete

Users can delete specified records or all records in the table.

Execute

Users can execute an action on records in the table.

All

Users have permissions on all records in the table.

Permissions suggestions

Consider granting these permissions in your organization:

Suggested group and role	Type of users in the group	Suggest permissions	Description
Catalog user group Cloud user [sn_cmp.cloud_service_user]	Users who order similar items from the catalog in the Cloud User Portal.	Access [read] or Execute	Catalog items, even after they are published, cannot be seen by users in the Cloud User Portal until you grant a user group to which that user belongs read access to the items. Grant Access and Execute permissions to cloud users on the Blueprint Catalog Item table and the Cloud Account table.
Blueprint designers Service Designer [sn.cmp.cloud_service_designer]	Users who design blueprints.	Manage [Create and update]	Blueprint designers cannot see or edit other blueprints or catalog items by default. To collaborate or reuse existing blueprints and catalog items, blueprint

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Suggested group and role	Type of users in the group	Suggest permissions	Description
			designers need access to each others' blueprints through the Manage permission.
Cloud admins Cloud administrator `[sn_cmp.cloud_admin]`	Users who create and manage cloud accounts.	Manage [Create and update]	Cloud admins must map templates to appropriate resource profiles. To collaborate or reuse existing resource profiles and templates.

Assign a cloud permission

Assign a permission to refine the actions that are allowed or prohibited for users based on the user group they belong to.

Before you begin

- Role required: sn_cmp.cloud_governor
- The user group to which you want the permission applied.

Procedure

1. Navigate to **All** > **Cloud Admin Portal** > **Governance** > **Permission**.
2. Fill out the form fields (see table).

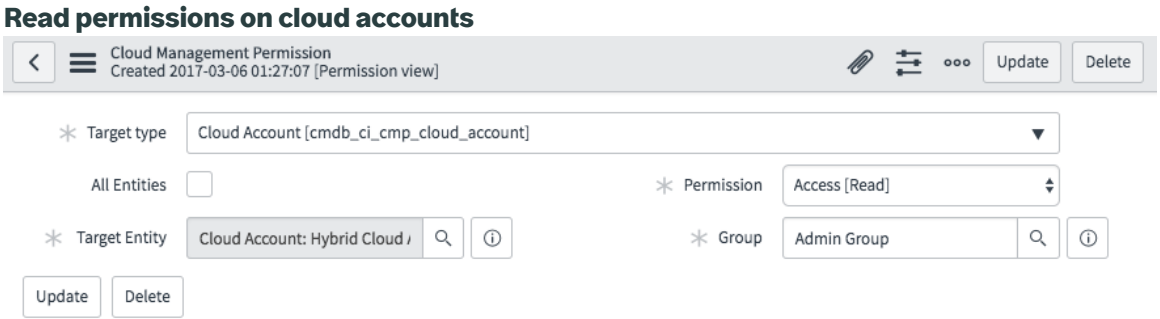

Field	Description
Target type	Select the cloud table in which the target record belongs.
All Entities	Select this option to apply the permission to all records in the table.
Permission	Select the permission type.
Target Entity	Select the record that the permission is based upon.
Group	Select the user group.

3. Click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Policies for Cloud Provisioning

A cloud policy can override a property value set by a user, create an approval task, reserve an IP address, pre-populate or hide form fields, execute custom scripts, call the Cloud API, or start or abort subflows. A cloud policy gives you system-wide control over approvals, resource operations, blueprint operations, or catalog item settings.

Example: A user requests a stack that triggers an approval policy

1. On the Cloud User Portal, a user submits a request to provision a particular blueprint. The process of provisioning the blueprint is the trigger that causes the policy engine to apply an approval policy. A cloud approval policy specifies the users who must approve a specified cloud activity before the activity can proceed.
2. The policy engine determines that the request meets the condition specified in the rule for the policy. In this example, the condition evaluates to true whenever a particular blueprint is being provisioned.
3. Because the condition is met, the policy engine performs the action that is also specified in the rule. In this example, the action is to create an approval action for the manager.
4. While the approver (the manager) reviews the approval request, the user sees a "waiting for approval" status message on the Cloud User Portal.
5. After the manager approves, the blueprint is provisioned.

User waiting for approval

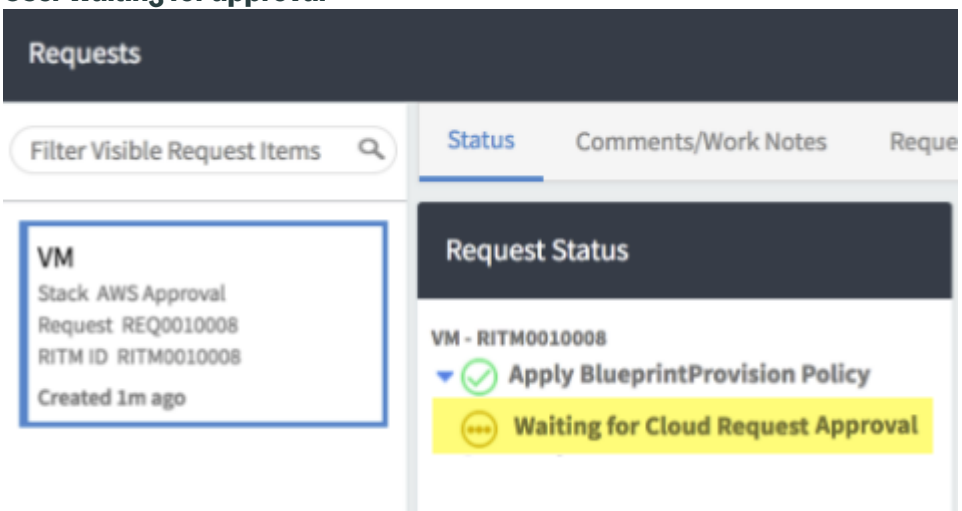

Note: You cannot invoke flows while adding a resource operation step. This is a known limitation.

About policy triggers

Triggers are events that set the policy engine in motion. For example, the *on Catalog item request end* trigger fires after a user submits a request form. When you think about defining a policy, you might first consider how the policy will be triggered. Examples:

- A user requests a *Stop* operation on a virtual server (the *on Stack resource operation* trigger fires)
- A resource reaches the end of its lease (the *on Lease End* trigger fires)
- A user requests a particular stack (the *on Blueprint provision* trigger fires)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You typically refer to a policy by the name of the trigger for the policy. For example, you might refer to a policy that is triggered by the *on Lease end* trigger as a "Lease end policy." For more detail on the types of trigger that you can implement, see Triggers for cloud policies.

How policies work

For each policy, you typically configure a trigger that starts the policy engine. The policy engine enforces the policy by running all rules that are defined for the policy. A policy rule is a collection of conditions and actions. If all conditions evaluate to true, the policy engine performs the actions. If any condition evaluates to false, the policy engine does not perform the actions.

- Some policy types apply only to particular types of cloud operation, like start, stop, provision, or deprovision, or to a particular target, like 'on blueprint123 provision operation' or 'on catalog item ABC launch'.
- You can configure a policy that does not specify a target, for example, 'on any blueprint provision operation' or 'on any catalog item launch'. Errors on policies that apply to any object are ignored.
- If multiple policies apply, you can specify the order that the policies are applied (with an exception that is described in a following section).
- Policies can work with dynamic forms to allow you to show or hide form fields from end users. Users see only information that you decide they need for them to understand and complete their tasks.

Policy groups

A cloud policy group is a container for related policies. Consider grouping policies that are often used together or should be considered together. Grouping policies can help you to apply policies consistently across your organization.

Order of execution when multiple policies apply

Order of execution when multiple "approval" policies apply:

When multiple "approval" policies apply, the policies are applied in the following order. (The approval policies are *on Blueprint provision (approval)*, *on Stack operation (approval)*, *on Stack resource operation (approval)*, and *on Task remediation*):

1. Only the first successful approval policy is applied and no other approval policies are applied.
2. If the applied approval policy has multiple rules, only the first successful rule is used.
3. If a rule has multiple actions, only the first successful action is performed.
4. If the applied approval policy includes both a custom approval and a Service Now approval, only the custom approval process is performed.

Order of execution for all other policy types:

- Policies are applied in the order that is specified by the **Order of Execution** property setting.
- If multiple policies have the same **Order of Execution** setting, then order is not guaranteed.

About operations in policies

Triggers are often based on user requests and the operations (start, stop, provision, or de-provision) that can run on a blueprint, a catalog item, a resource, or a stack. Some trigger types do not specify a cloud operation. For example, the *on Lease End* trigger fires independently of any operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About policy rules

A policy rule is a collection of conditions and actions. If all conditions evaluate to true, the policy engine performs the actions. If any condition evaluates to false, the policy engine does not perform the actions.

- Conditions: Conditions can consider request form data, resource activity, or user activity. Examples:
  - Is the size of the requested CPU greater than 32?
  - Is this a `Stop` operation?
  - Is the lease for this resource ending in the next 7 days?
- Actions: If all conditions evaluate to true, the policy engine runs the actions that the rule specifies. Expressions in policy actions can set or override values. Examples:
  - Override the value that the user specified by changing the CPU size to 16 and then start the approval process. (because the condition was met that the requested CPU is greater than 32).
  - Create an approval task for the manager. (because the condition was met that a `Stop` operation was requested).
  - Send notifications to every user in the ABC group. (because the condition was met that the lease ends in 7 days).

Policy action scripts

- Use policy action scripts to get, update, or set values in cloud requests.
- Your instance tracks tagged resources for billing and reporting. Policy action scripts can add and modify resource tags.

Cloud policy example

A base system cloud policy, `Lease End ServiceNow`, uses the `on Lease end` trigger that fires when a virtual resource is near the lease end date. If the conditions in the policy rule are met, the policy engine sends a notification to the owner of the resource and performs operations on the resource.

Before you begin

Role required: sn_cmp.cloud_governor or admin

About this task

In this example, you work with the `Lease End ServiceNow` policy. The policy performs the following actions when the conditions in the policy rule are met:

- Send a notification to the owner of a cloud resource when the resource lease is about to expire.
- Run a `Stop` operation on the resource.
- Run a `Deprovision` operation on the resource.

You can change any base-system policy by modifying settings or making the policy active or inactive.

Note: Scheduled Start/Stop operations are not supported on the Terraform template and the Google Deployment Manager (GDM) stacks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Search for and open the **Lease End ServiceNow** policy.
3. Notice that the trigger is **on Lease end**, which means that the policy is triggered when the lease for the resource is near expiration.

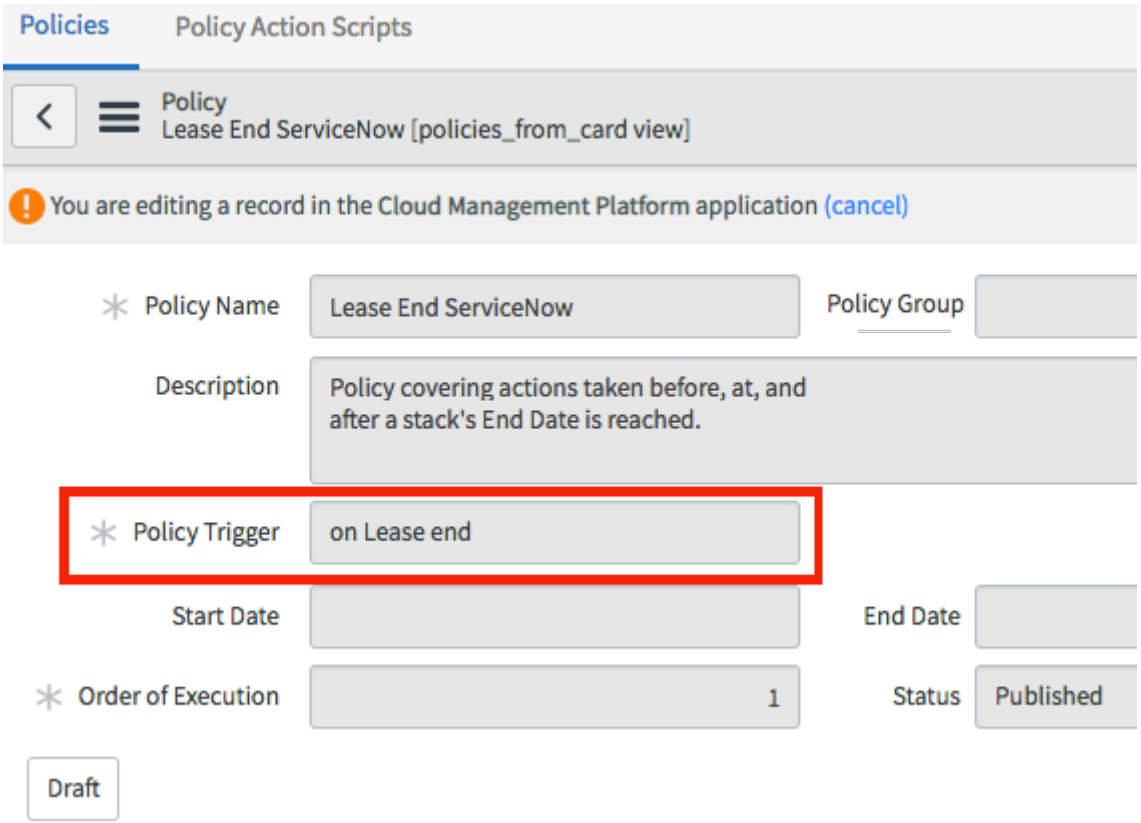

4. Sort the Rules related list by **Order of Execution**.
   The rule with the lowest number runs first.

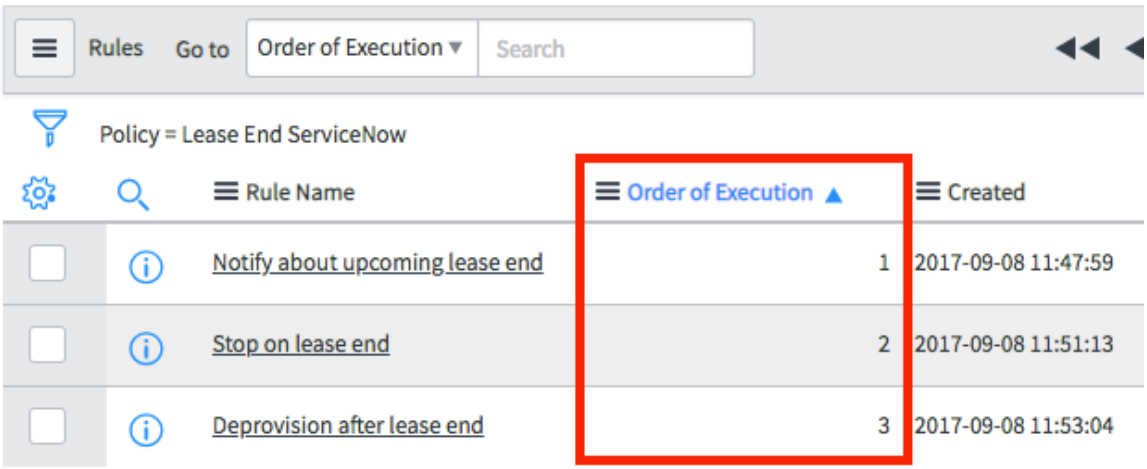

   Notice that a notification rule runs first, followed by the `Stop` and `Deprovision` operations on the resource.

5. Click the **Notify about upcoming lease end** rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Notice that the rule specifies a condition: If it is one day before the end date, return `true`.

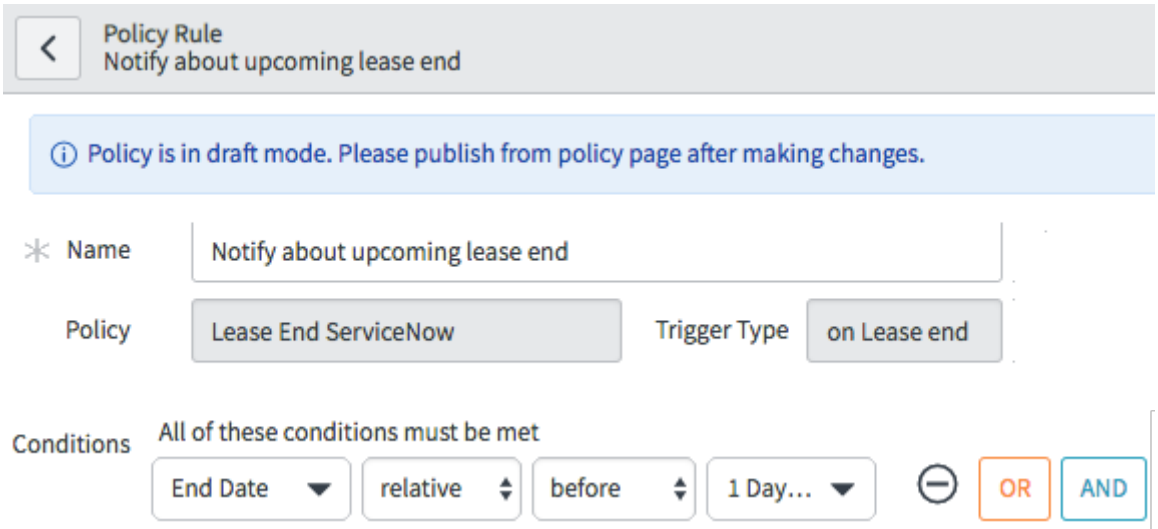

7. In the Policy Rule Actions related list, click the (empty) **Name** of the action.
8. On the Update Action pop-up window, enter an **Action Name**.

   Note the value of the **Notification** field for use in the next step.

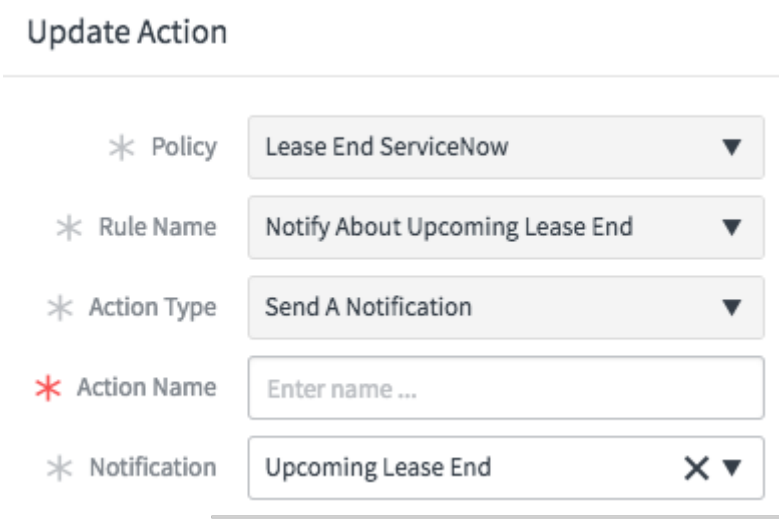

9. Review the notification message other settings:

   a. On the instance (not the Cloud Admin Portal), navigate to **System Notification** > **Email** > **Notifications**.

   b. Search for and open the notification that is specified in the action by entering the filter text `*Upcoming lease end` for **Name**:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

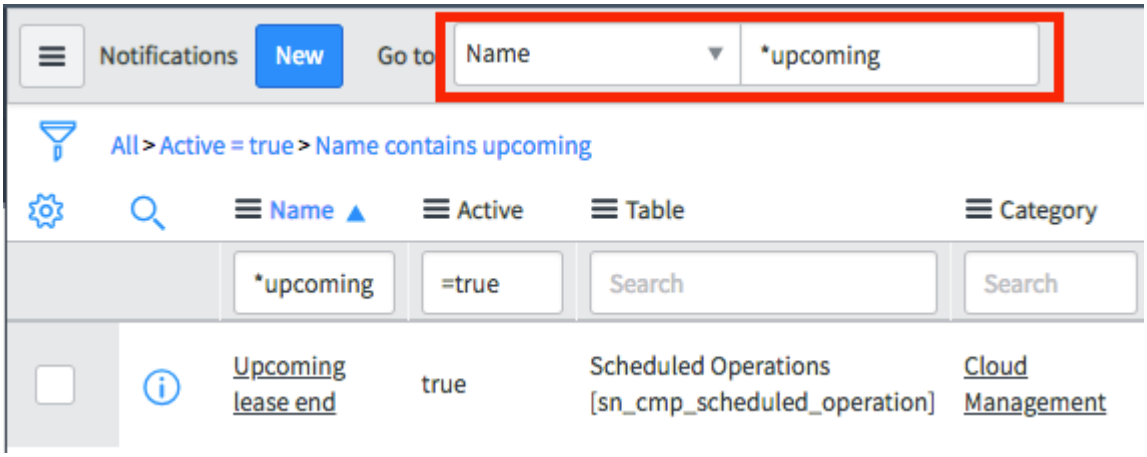

c. Click **Upcoming lease end** to open the notification form.

Notice that the **Scheduled Operations [sn_cmp_scheduled_operation]** table is specified.

Note: You base all notifications for Cloud Provisioning and Governance on this table.

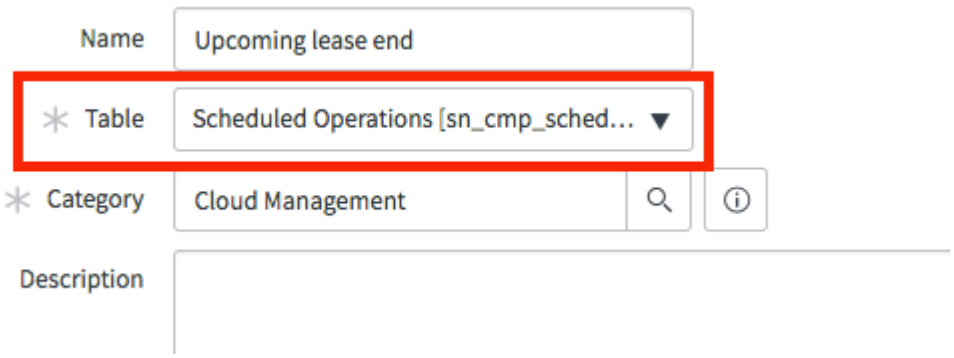

d. On the **When to send** tab, notice the conditions that apply:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

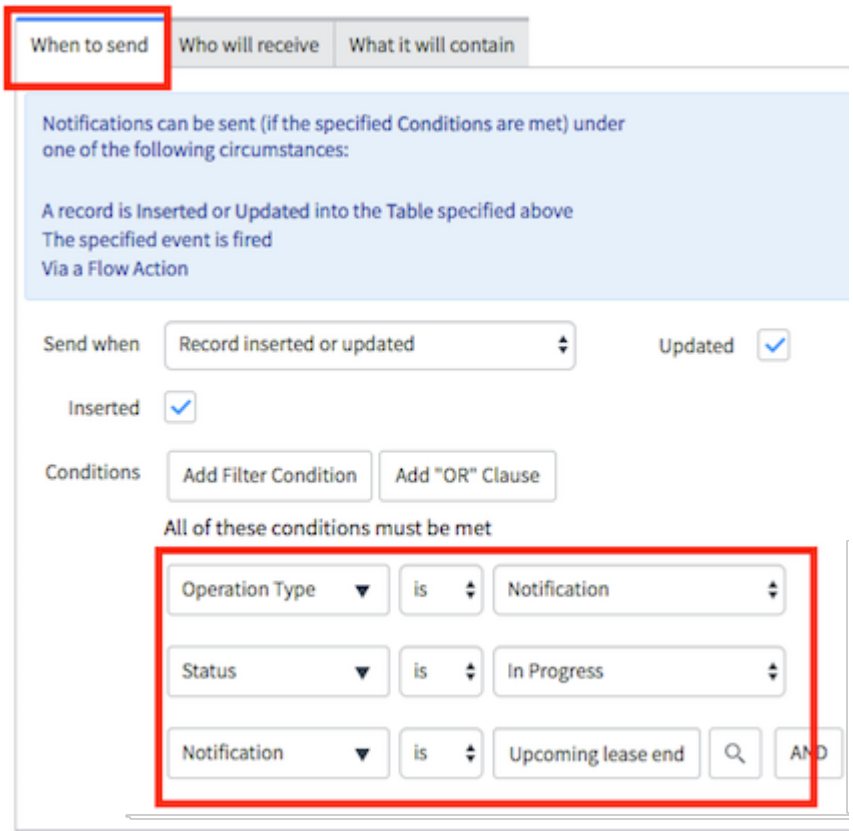

Condition	Description
Operation Type is Notification	Specifies that the **Operation Type** field on the Scheduled Operations table must be **Notification**.
Status	Specifies that the system should trigger the notification when the status of the scheduled operation record is in progress, rather than waiting until it is complete.
Notification	Specifies that the *`Upcoming lease end`* notification must be specified.

Note: When you create a custom notification, use the default conditions, but change the **Notification** condition to the name of your notification.

e. Click the **Who will receive** tab.

Notice that the **Users/Groups in fields** value is set to `Target.Assigned to`, `Target.Owned by`. You should use these settings to send the notification to the owner of the virtual resource.

10. In the Cloud Admin Portal, open the Lease End ServiceNow policy again, and then open the **Stop on lease end** rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

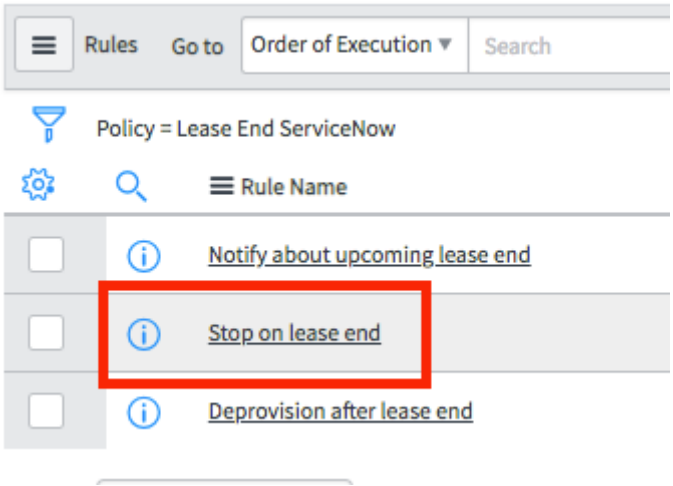

The rule has two actions — one action to send a notification and the other action to run the *`Stop`* operation on the virtual resource.

11. Click the (empty) **Action Name** of the action for the *`Stop`* operation.

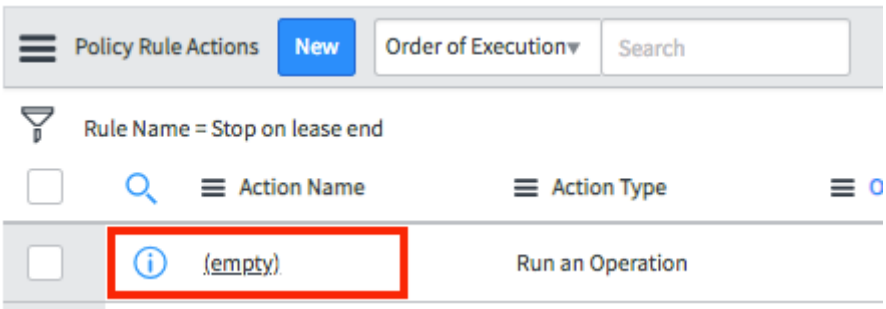

12. On the Update Action pop-up window, notice the settings:

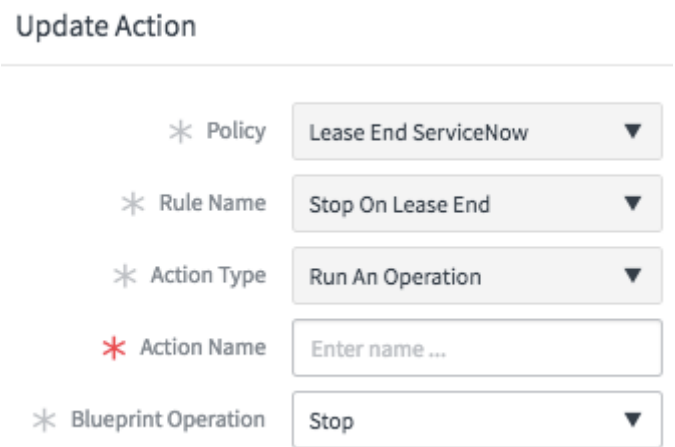

Field	Description
Policy and Rule Name	The policy and the rule to which this action belongs are specified in these fields by default.
Action Type	**Run An Operation** tells the policy engine to trigger a standard operation (*`Start`*, *`Stop`*, *`Deprovision`*, and so on) on the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Blueprint Operation	The system runs the *STOP* operation on the resource.

The other rule, **Deprovision after lease end**, performs the *Deprovision* operation on the resource so that it does not stay active after the lease end date.

What to do next

This example uses the settings available through the form configuration. To see example scripts that you could use in addition, see Create a policy action script.

Triggers for cloud policies

Triggers are events that set the policy engine in motion. For example, the *on Catalog item request end* trigger fires after a user submits a request form. When the trigger for a policy fires, the policy engine tests the conditions specified in the policy rule and performs the actions specified in the rule, if the conditions are met.

About triggers

- You typically refer to a policy by the name of the trigger for the policy. For example, you might refer to a policy that is triggered by the *on Lease end* trigger as a "Lease end policy."
- Triggers are often based on user requests and the operations (start, stop, provision, or de-provision) that can run on a blueprint, a catalog item, a resource, or a stack. Some trigger types do not specify a cloud operation. For example, the *on Lease End* trigger fires independently of any operation.
- To optimize performance, limit the number of policies with general triggers like the *on Catalog item* triggers.
- A trigger that does not specify a target (a blueprint, catalog item, stack, or resource) is always executed. To optimize performance, therefore, minimize the use of such policies.

Important: Blueprints are deprecated. The trigger descriptions following the time-line schema below, are applicable for creating cloud template based catalogs as well.

Provisioning trigger time-line

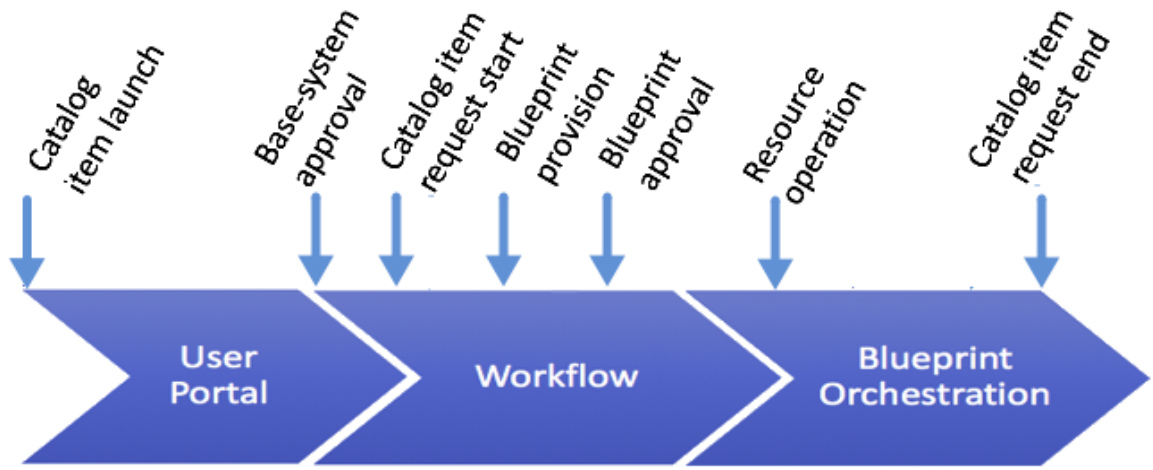

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Policy triggers

Types of policy trigger and the actions that are available for the associated policy rules

Trigger name and actions	Description
on Blueprint provision Actions: • Execute a script • Property override • Abort process	The *on Blueprint provision* trigger fires after execution of *on Catalog item request start* policies. A policy that is triggered by the *on Blueprint provision* trigger can run a script, override a user-requested attribute value, or abort and send a message about the provision operation. Use this trigger to override a value that the user enters. For example, when a user chooses a value for an attribute like the stack name, a policy with this trigger can change the stack name. In addition, another action can change the name again when the user finally provisions the resource. The user does not see the final value on the catalog item form because the change is made at provision time. Create an action for an 'on Blueprint provision' policy
Approval triggers • **on Stack operation (approval)**: Triggered during any stack operation on the Cloud User Portal. • **on Stack resource operation (approval)**: Triggered during any resource operation (start, stop, provision, and so on) on the Cloud User Portal. • **on Task remediation**: Triggered when a user resubmits a failed request. Actions: • ServiceNow Approval • Custom Approval	A cloud approval policy specifies the users who must approve a specified cloud activity before the activity can proceed. Approvers can include the manager of the user making a request, a specified user or group, or users with a specified role. You can specify multiple approvers. Approvals occur in the order that you specify. A policy that is triggered by one of the approval triggers can start approval subflows. The targeted approval policies complement the base-system approval operations. **Note:** The approval process is performed after properties are set because property values that were overridden could change costs. *on Blueprint provision (approval)* is applied before the blueprint is provisioned. Because the provisioning process can alter request data (and possibly change costs), approval processes run after the blueprint is provisioned. Use *on Stack operation (approval)* to run an approval workflow when an operation is performed on a stack. By default, a change request is generated when an operation is performed on a stack, but it does not require an approval. This trigger can launch a mandatory approval. Use *on Stack resource operation (approval)* to run an approval workflow when an operation is performed on a single resource that is part of a stack. By default, a change request is generated when an operation is performed on a stack, but it does not require an approval. This trigger can launch a mandatory approval.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Types of policy trigger and the actions that are available for the associated policy rules (continued)

Trigger name and actions	Description
	A policy that is triggered by the *on Task Remediation* trigger can start approval subflows. Create an action for an approval policy
on Catalog item launch Actions: • Execute a script • Property override	The *on Catalog item launch* trigger fires when an order form (stack request form) is launched for a catalog item. A policy that is triggered by the *on Catalog item launch* trigger can run a script or override a user-requested value (text values only). Use this trigger to control what the user sees in the form when it first opens in the Cloud User Portal. For example, you can override a default value that first appears to the user. The user can see this value on the catalog item form. When both a policy rule and a form rule overwrite a value, the value in the form rule is used. Create an action for an 'on Catalog item launch' policy
on Catalog item request start Actions: • Execute a script • Execute a workflow	The *on Catalog item request start* trigger fires after the user opens a request form. A policy that is triggered by the *on Catalog item request start* or *on Catalog item request end* trigger can run a script or execute a subflow. You can use this trigger to run a custom script or workflow to fulfill enterprise processes like custom approval before the catalog item request is processed. Create action for 'on Catalog item request start/end'
on Catalog item request end Actions: • Execute a script • Execute a workflow	The *on Catalog item request end* trigger fires after a user submits a request form. A policy that is triggered by the *on Catalog item request start* or *on Catalog item request end* trigger can run a script or execute a subflow. Use this trigger to launch a workflow after a catalog item request is processed. Consider this trigger a post-provisioning step. For example, you could launch a workflow to install MySQL on the provisioned resource. Create action for 'on Catalog item request start/end'
on Lease end Actions:	A policy that is triggered by the *on Lease end* trigger can send a notification or perform a *Start*, *Stop*, or *Deprovision* life cycle operation. Create an action for an 'on Lease end' policy

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Types of policy trigger and the actions that are available for the associated policy rules (continued)

Trigger name and actions	Description
• Run an Operation • Send a Notification	
on Resource operation launch Actions: • Execute a script • Property override	The *on Resource operation launch* trigger fires before the catalog for a resource operation is loaded from the Cloud User Portal. A policy that is triggered by the *on Resource operation launch* trigger can run a script or can override a user-requested value (text values only). When both a policy rule and a form rule overwrite a value, the value in the form rule is used. Create action for 'on Resource operation launch' policy
on Resource operation request start Actions: • Execute a script • Execute a workflow	The *on Resource operation request start* trigger fires after a user submits a resource operation request (*Start*, *Stop*, *Deprovision*). A policy that is triggered by the *on Resource operation request start* or *on Resource operation request end* trigger can run a script or override a user-requested attribute value. Create action for 'on Resource operation request start/end' policy
on Resource operation request end Actions: • Execute a script • Execute a workflow	The *on Resource operation request end* trigger fires before completion of a life cycle operation on a resource (*Start*, *Stop*, *Deprovision*). A policy that is triggered by the *on Resource operation request start* or *on Resource operation request end* trigger can run a script or override a user-requested attribute value. Create action for 'on Resource operation request start/end' policy
on Resource operation Actions: • Property override • Execute a script • Call Cloud API • IP Address Management	The *on Resource operation* trigger fires during the Orchestration process when a user performs a *Start*, *Stop*, or *Deprovision* life cycle operation on a specific resource. A policy that is triggered by the *on Resource operation* trigger can override a user-requested attribute value, run a script, call a Cloud API, or perform an IP address management operation. Create an action for an 'on Resource operation' policy

Create a cloud policy

A cloud policy can override a property value set by a user, create an approval task, reserve an IP address, pre-populate or hide form fields, execute custom scripts, call the Cloud API, or

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

start or abort subflows. A cloud policy gives you system-wide control over approvals, resource operations, blueprint operations, or catalog item settings.

Before you begin

- Optional: Create one or more cloud policy groups.
- Role required: sn_cmp.cloud_governor or admin

About this task

This procedure describes every policy type except approval policies (*on Blueprint provision (approval)*, *on Stack operation (approval)*, *on Stack resource operation (approval)*, and *on Task remediation*). See Create a cloud approval policy for instructions on creating a policy with an approval trigger.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Click **New** and then fill in the form.

Example policy form

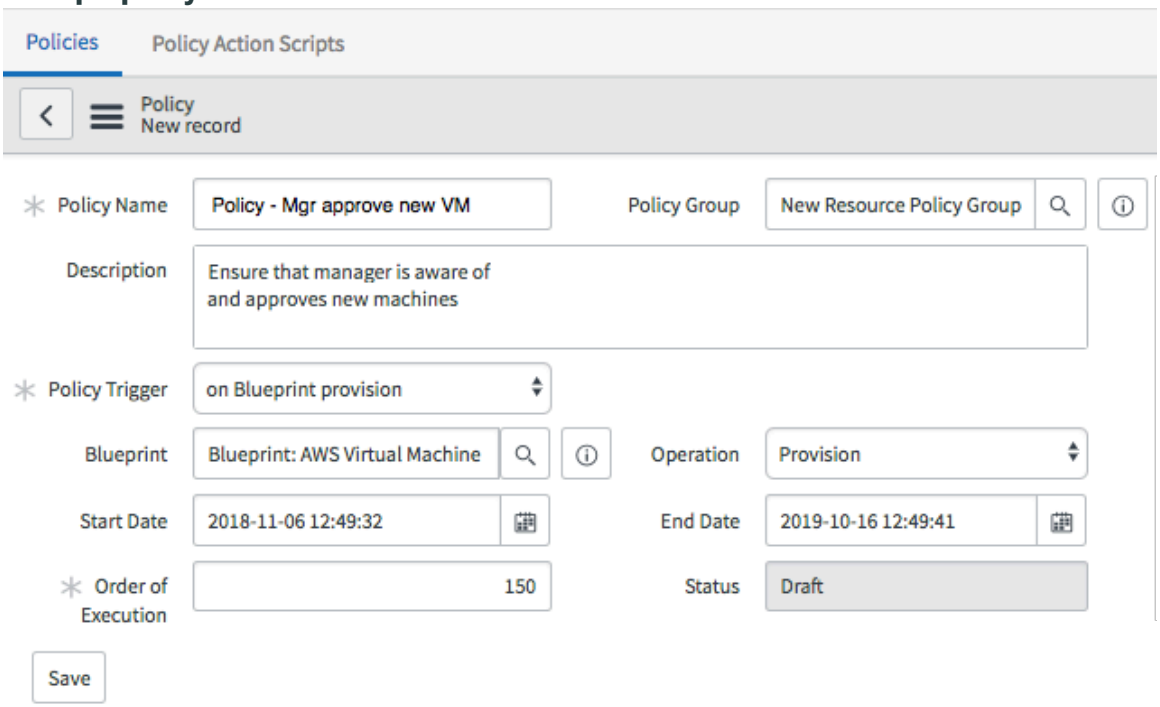

Field	Description
Policy Name	A descriptive name that includes the word `Policy`. Do not start the name with a number.
Description	A description of the intent of the policy.
Policy Group	A policy group to which the policy belongs.
Policy Trigger	A trigger that specifies when the policy should be applied. For example, the **on Lease end** trigger applies the policy when the lease for a stack ends. See Triggers for cloud policies.
Resource Block	The resource block that the policy applies to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
[appears when the *on Resource Operation* trigger is selected]	
Operation [appears when the *on Blueprint provision*, *on Blueprint provision (approval)* or *on Resource Operation* trigger is selected]	Select the operation that the policy applies to. For example, a policy can apply to the *Deprovision* operation only or to all operations on the blueprint or catalog item. **Note:** If no operation is specified, then the policy applies for every operation. This condition can decrease performance. ○ **All**: Any operation executes. ○ **Start**: The resource starts. ○ **Stop**: The resource stops. ○ **Provision**: The resource is provisioned. ○ **Deprovision**: The resource is no longer available to users. ○ **Execute Script**: A script runs on the resource.
Moment [appears when the *on Resource Operation* trigger is selected]	Specify when the policy should be enforced: ○ **Pre-operation**: Before the specified operation starts. ○ **Post-operation**: After the specified operation finishes. **Note:** If you are integrating with Infoblox, use **Pre-operation** for a vSphere virtual machine. Use **Post-operation** for AWS and Azure clouds because AWS and Azure control the allocation of IP addresses. You can register the IP address that is provided with Infoblox.
Catalog item [appears when an *on Catalog item launch*, *on Catalog item request start*, or *on Catalog item request end* trigger is selected]	Select the catalog item that the policy applies to.
Start Date / End Date	Specify the start date when the policy should be considered and the end date when the policy should no longer be considered.
Order of Execution	Specify a number that represents the order in which the policy is applied. A policy with a lower number runs before a policy with a higher number. For example, a policy with **Order** of **100** runs before a policy with an **Order** of **200**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: The **Order of Execution** property does not apply for *on Blueprint provision (approval)*, *on Stack operation (approval)*, *on Stack resource operation (approval)*, and *on Task remediation* policies. See Create a cloud approval policy for details.
Status	○ **Published** policies are enforced. You cannot edit a policy in the **Published** state. To edit a policy that is in the **Published** state, click **Draft** on the form header. ○ You can edit **Draft** policies. **Draft** policies are not enforced. To enforce a policy, click **Publish** on the form header.

3. Right-click in the header and select **Save**.

What to do next
Configure a cloud policy rule for the policy.

Configure a cloud policy rule

A policy rule is a collection of conditions and actions. If all conditions evaluate to true, the policy engine performs the actions. If any condition evaluates to false, the policy engine does not perform the actions.

Before you begin
Optional: Create one or more cloud policy groups.

Role required: sn_cmp.cloud_governor or admin

About this task
- A policy can include multiple rules. You can configure the order that rules are applied.
- Every rule is executed to completion.
- Any rule error stops execution. Exception: Policies that do not specify an operation.

Procedure
1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. On the Rules related list, click **New** and then enter a descriptive **Name** for the rule.
   - The rule name cannot be the same as the policy name.
   - The rule name should not start with a number.
4. Using the base-system condition builder, specify the **Conditions** that must be met for the policy engine to perform the actions.

   The following example rule is applied for a blueprint provision approval policy. The condition evaluates to true if the **OS profile** field on the request form has the value “Windows”.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

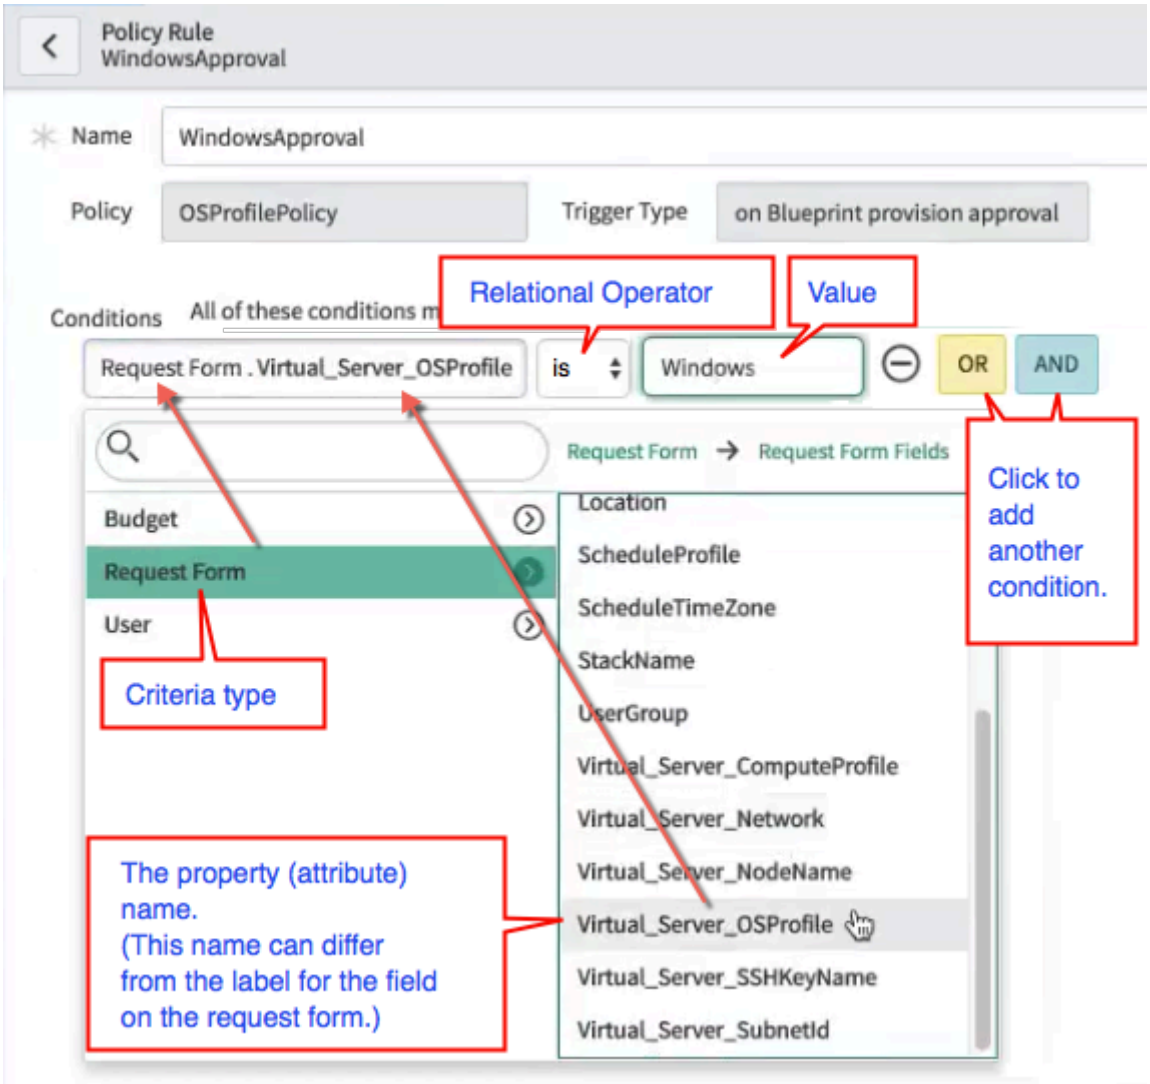

Field	Description
[Criteria type]	Select one of the following options: ○ **Request Form**: Base the condition on a property (attribute) value in the request form in the Cloud User Portal. You specify which property to test. In the example, the condition tests whether the OS profile field on the request form has the value "Windows." ○ **User**: Base the condition on a user role or on the group that the user belongs to. For example, to create a condition that applies when the user belongs to the Marketing group, select: ▪ [Criteria Type]: User ▪ [User Entity]: Group ▪ [Relational Operator]: Equals ▪ [Group]: Marketing

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: You use different criteria to build a condition for policies that are triggered by the `on Lease end` trigger. See the example that appears after the table.
[Relational Operator]	Select an appropriate operator.
[Value]	Depending on the selected [Criteria type], enter or select a value. ○ For **Request Form**, enter a value for the attribute. ○ For **User** > **Group**, select a group. ○ For **User** > **Role**, select a role.
[Logical operator]	When defining multiple conditions, select a logical operator **OR** or **AND** that specifies the logical relationship with the next condition.

5. Click **Submit**.

Example:
For policies that are triggered by the `on Lease end` trigger, use the base-system date condition builder to configure the condition to test. In addition to the condition, the following example shows the actions that the policy engine performs if the condition is met.

Condition and actions for a lease end policy

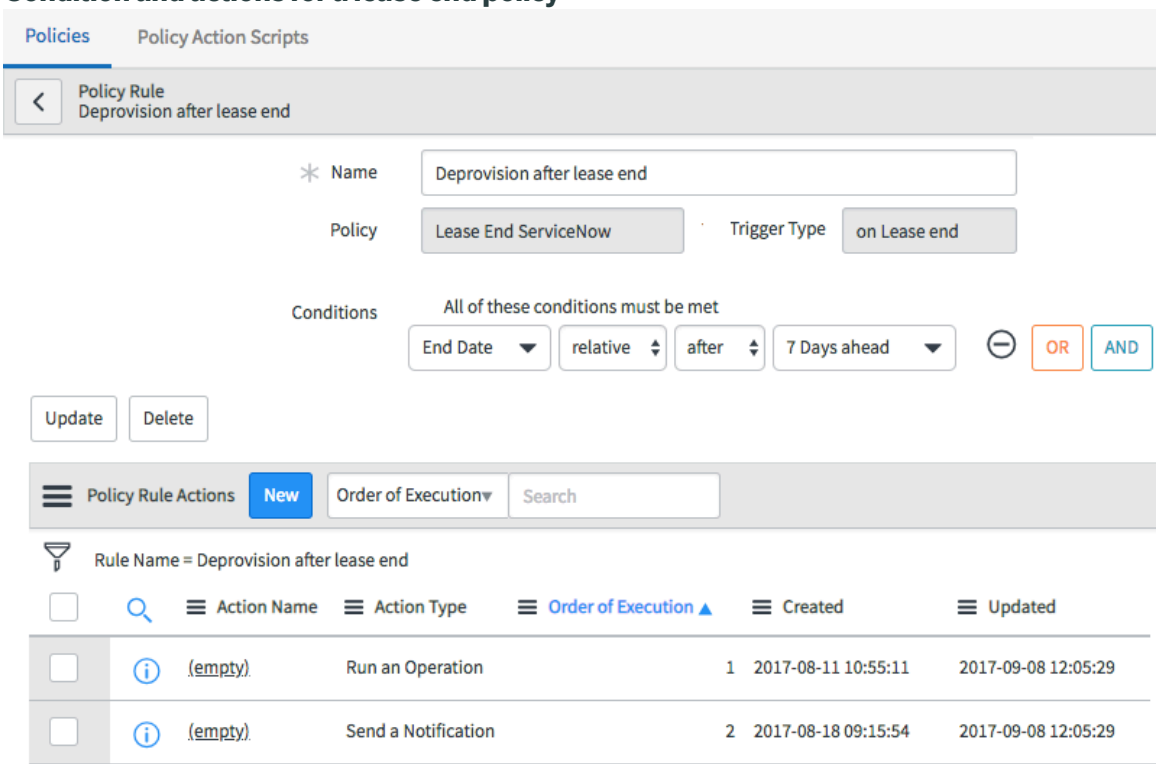

What to do next
Create the policy actions that should run when all conditions of the rule evaluate to true.

Using expressions in Cloud Provisioning and Governance

Expressions in policy actions can set or override values. Expressions in blueprints can access attributes of resources and can map values to request form fields. Expressions are available in

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

resource blocks, blueprints, policies, and anywhere that Cloud Provisioning and Governance allows scripts.

Using expressions

Expressions can hold the values for information that is dynamically generated in the system, such as the values for the selections users make on the Cloud User Portal. Expressions are available in resource blocks, blueprints, policies, and anywhere that Cloud Provisioning and Governance allows scripts.

Using expressions in resource blocks

Each resource block operation uses an expression to hold a value for each parameter. The expression can use hard-coded values, data from the stack that the user provisions in the Cloud User Portal, data in the CMDB, and data that is derived from scripts. By default, Cloud Provisioning and Governance generates a set of parameters and their expressions for each resource block operation.

For example, the parameter *Location*, which holds the datacenter that a resource belongs to, uses the expression `${parameter.Location}`

Expressions in a resource block

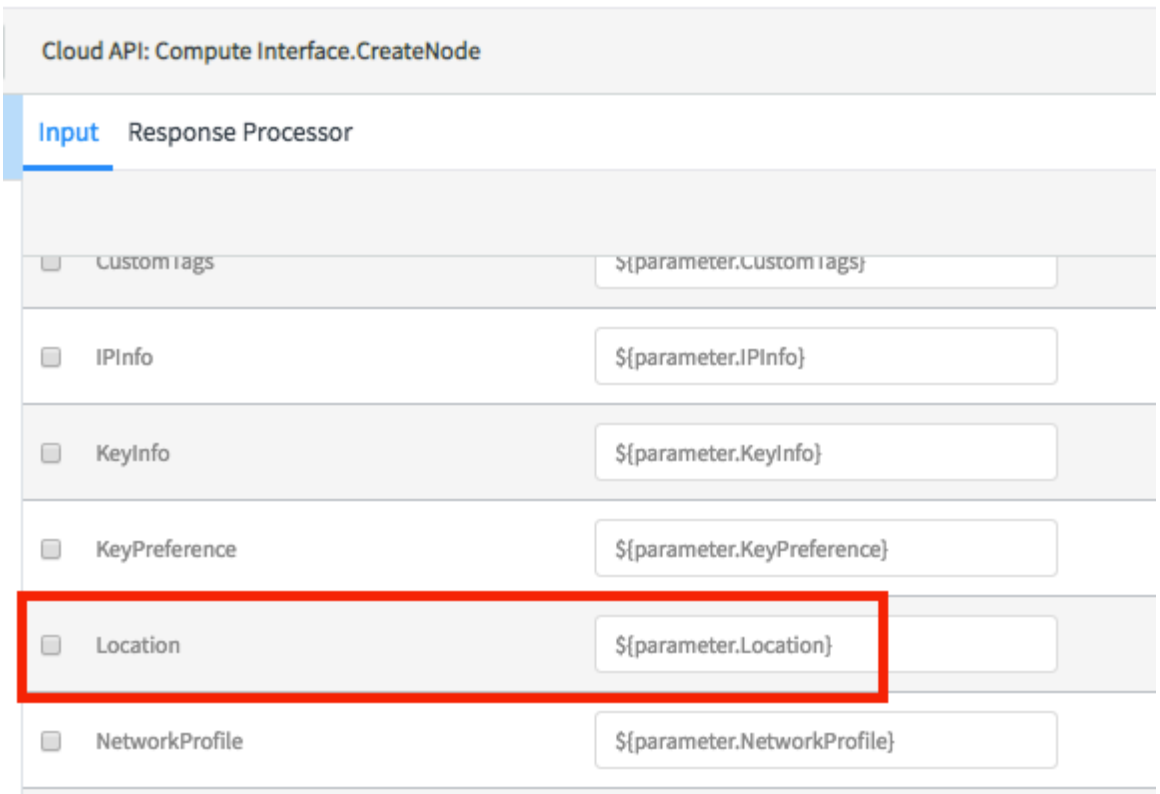

Using expressions in blueprints

Blueprints can use expressions to map input parameters for each operation on a resource block. By default, the blueprint displays the same parameters and values that are specified in the resource block.

You can access blueprint attributes with expressions. The expression in the **Mapping** column defines blueprint attributes for each operation in a step. For example, the **Location** attribute, which appears on the stack request form, is accessible through the `{$parameter.Location}` expression. The Location

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

parameter with the `${parameter.Location}` value appears on the **Inputs** tab when you select a blueprint operation.

Expressions in a blueprint

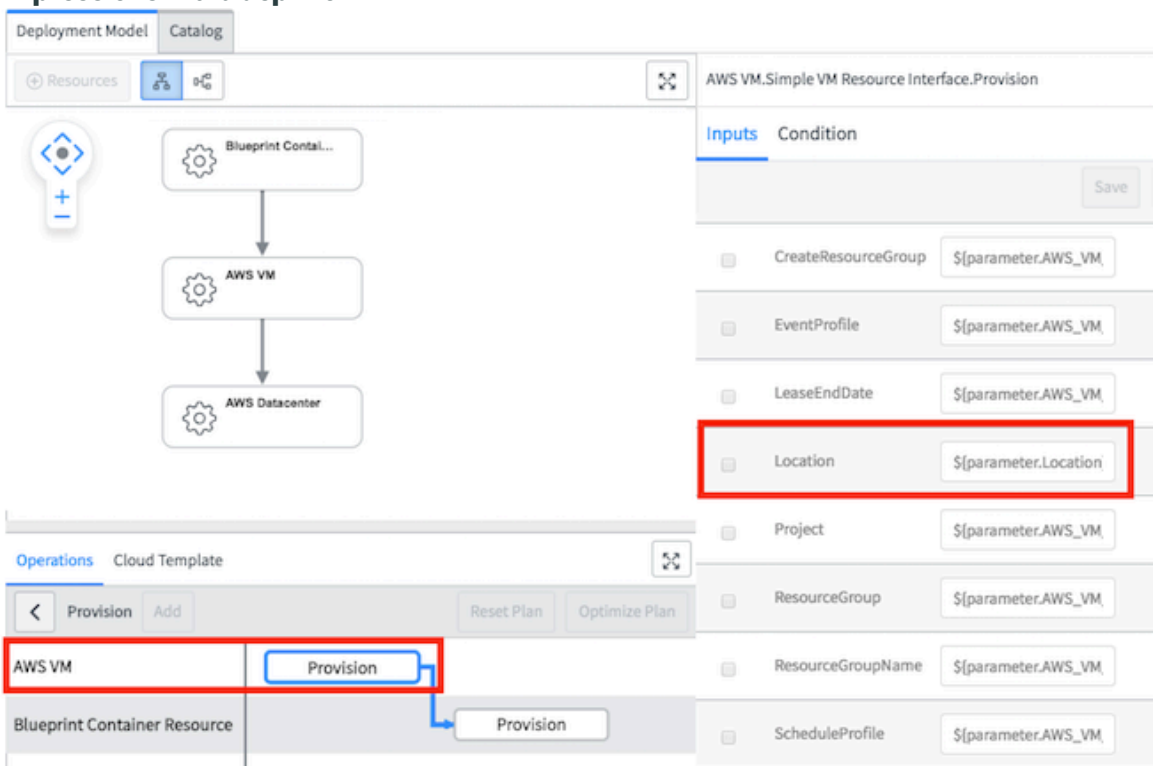

Using expressions in policy actions

Policy actions can use expressions to override parameter values that users submit in a request form on the Cloud User Portal. You can also access and override user data in stack operations. For example, use the following expression to rename a stack: `formData.Stackname = "your-naming-convention";`.

Expressions in a policy script

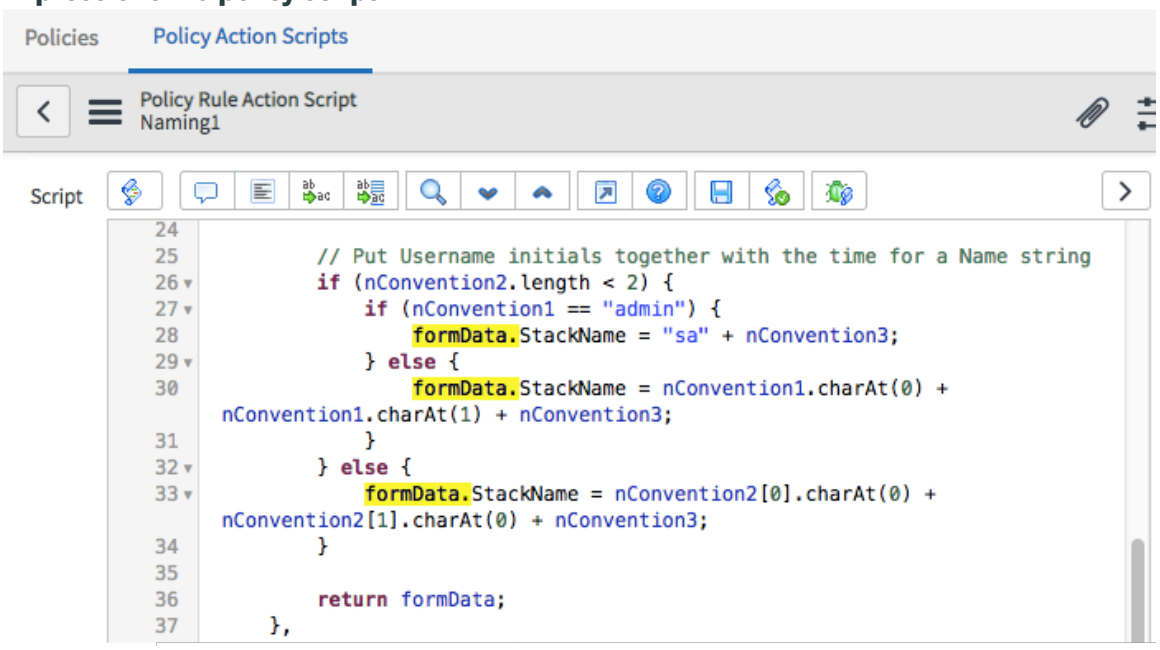

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can access blueprint attributes with expressions. The expression in the **Mapping** column defines blueprint attributes for each operation in a step. For example, the **Location** attribute, which appears on the request form for a virtual machine, is accessed in the following expression: `{$parameter.Location}`.

Expression types

Definitional expressions

Definitional expressions are a form of early binding (or static binding). These bindings refer to compile-time binding and are evaluated when the user submits the stack request form (when stack or resource provisioning starts). Definitional expressions do not have access to the data that is generated during Orchestration. Definitional expressions are enclosed in curly braces. For example: `${parameter.CloudAccount}`

Runtime expressions

Runtime expressions are a form of late binding (or dynamic binding). These bindings refer to runtime binding and have access to the data that is created during Orchestration (for example, stack items). Runtime expressions are evaluated when called during an Orchestration run. Runtime expressions are typically used for dot-walking to data in tables. Runtime expressions are enclosed in parentheses. For example: `$(Stack.items[VM1].attributes[node_id])`

Definitional/Runtime expressions

Some expressions include both definitional and runtime expressions.

Complex expressions

You can nest expressions of any type within other expressions.

Definition expression syntax and examples

Allowable syntax uses a dollar sign and curly braces. These types are available:

- `${parameter.}`. Use this kind of expression to retrieve values from input attributes of processes like blueprint provisions. Data is not fetched from tables.
- `${Stack.items[]}`. Use this syntax to access attributes of specific items in a stack.
- `${randomNumber}`. Use this syntax to generate a random number. For example, you can provision a VM with a random node name or stack name.

Definition expression examples

Example	Description
${parameter.BillingCode}	This expression shows how a billing code parameter appears in a blueprint.
${parameter.formData.CatalogAttributeType}	This expression takes a catalog attribute that is submitted by a user while the user fills out the form for a catalog item in the Cloud User Portal.
${parameter.userData.userId}	This expression takes the ID of a user while the user fills out the form for a catalog item in the Cloud User Portal.
${Stack.items[Virtual Server].attributes[sys_id]}	This expression gets the sys_id of a virtual server that is apart of a stack. Virtual Server is the alias of a resource block used in the stack.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Runtime expression syntax and examples

Allowable syntax uses a dollar sign and parenthesis. These types are available:

- `$(ci.tablename)` where `tablename` is a table in the system, usually a CI table in the CMDB. Use this syntax to access values of fields in the table.
- `$(Stack.items[])`. Use this syntax to access attributes of specific items in a stack.

Runtime expression examples

Example	Description
`$(ci.cmdb_ci_cloud_subnet[ sys_id=12231231231231231231].cidr)`	This expression dot walks to the Cloud Subnet table, finds the specific record with the given sys_id, and takes the value from the `cidr` column.
`$(ci.sn_cmp _ip_pool[subnet=${parameter.formData.Subnet Id}])`	This expression combines runtime and definition type expressions. The equal sign = is used to evaluate a value for a match. The expression dot-walks to the IP Pool table and looks for the subnet that has the subnet ID that the user submitted.
`$(ci.cmdb_ci_cloud_subnet[${parameter.formData.SubnetId}].cidr)`	This expression combines runtime and definition type

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Runtime expression examples (continued)

Example	Description
	expressions. The expression takes the value of the **cidr** field from a subnet that the user chose during provisioning. The square brackets [ ] indicate that the expression dot-walks to the Cloud Subnet table and then looks at the subnet value that the user submitted on a blueprint provision. The expression grabs the **cidr** field value and then walks to the value in the **cidr** field of the sys_id of the subnet.
$(Stack.items[Virtual Server].attributes[sys_id])	As in the definition expression example, this expression takes the sys_id of a virtual server that is apart of a stack.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example expression

```
$(ci.cmdb_ci_nic[$(Script:CMPVMNICs.getNICs[arg=
$(Stack.items[Virtual Server].attributes[sys_id])])].private_ip)
```

- *$(ci*: Runtime expression to retrieve data from table.
- *cmdb_ci_nic*: CI for NIC (Network Interface Card).
- *$(Script*: Script-based expression.
- *CMPVMNICs*: Script include.
- *getNICs*: Function inside a script include.
- *arg*: Arguments to script include function. Arguments are separated by "," when there are multiple attributes.
- *$(Stack.items*: Runtime stack expression to retrieve stack item from a stack. Argument is the alias specified in blueprint.
- *Virtual Server*: Alias of resource used in blueprint.
- *$(Stack.items[Virtual Server].attributes[sys_id])*: Retrieve sys_id of stack item ("Virtual Server") resource instance in a stack.
- *private_ip*: Attribute from cmdb_ci_nic. Replace with public IP if required.

Expressions

Simple parameter mapping expression

This type of expression retrieves values from input attributes of processes, such as blueprint operations, resource blocks, and policies. Data is not fetched from tables. Maps values from one layer to another, such as from a Blueprint to a Resource to the Cloud API.

- Type: definitional
- Syntax: `${parameter.}`
- Examples:
  - `${parameter.BillingCode}` returns the billing code.
  - `${parameter.formData.CatalogAttributeType}` takes a catalog attribute that is submitted on the request form in the Cloud User Portal.
  - `${parameter.userData.userId}` takes the ID of a user working on the request form in the Cloud User Portal.

Stack item expression

A CI instance in the CMDB represents each stack item. Use Stack Item expressions to look up first-level properties on the CI that back the stack item or on the stack item itself.

- Type: definitional/runtime
- Syntax: `${Stack.items[]}` or `$(Stack.items[])`
- Examples:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

`${Stack.items[VirtualServer1].attributes[sys_id]}` is a definitional type expression that gets the sys_id of a virtual server in a stack. `VirtualServer1` is the alias of a resource block that is used in the stack.

`$(Stack.items[VirtualServer2].attributes[sys_id])` is a runtime type expression that takes the sys_id of a virtual server that is a part of a stack.

`${Stack.items[vm1].attributes[node_id]}` reads the `node_id` attribute of the CI that was created for the VM. `vm1` is the stack item name (or the alias of the resource in the blueprint).

`$(Stack.items[vm1].status)` reads the status of the stack item.

Property override expressions in policies

In policies, you can override properties by pulling a value from the system or by using a random number. You can use data from both the forms in the Cloud User Portal and from the user who performed the operation on the form.

- Type: definitional/runtime
- Syntax: `${parameter.formData.xyz}` or `${this.parameter.userData.xyz}`
- Examples:

  The following value can set the stack name to *Stack_Bob.Smith@company.com*:
  `Stack_${parameter.userData.userId}`

  Set stack or table values using runtime expressions by taking the subnet from the IP Pools table:

  ```
  $(ci.sn_cmp _ip_pool[subnet=${parameter.formData.Subnet Id}])
  ```

  In scripts, you can assign values as follows:

  ```
  formData.App_Server_NodeName = "MyNodeName";
  ```

  ```
  this.parameter.userData.userId == 'servicenowuserId';
  ```

- See also Create a policy action script.

Script expression

In the example, `VMProperties` is a script include with a function called `getIP`. A script expression is also an example of a complex expression (the expression is nested).

- Type: runtime
- Syntax: `$(Script:scriptName.function[])`
- Examples:

  ```
  $(Script:VMPropertiesUtil.getIP[ arg=
  $(Stack.items[VM1].attributes[object_id])])
  ```

  Expression for Private IP:
  ```
  $(ci.cmdb_ci_nic[$(Script:CMPVMNICs.getNICs[arg=
  $(Stack.items[Virtual
  Server].attributes[sys_id])])].private_ip)
  ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Expression for Public IP:
`$(ci.cmdb_ci_nic[$(Script:CMPVMNICs.getNICs[arg=$(Stack.items[Virtual Server].attributes[sys_id])])].public_ip)`

Expression to get the credential alias:
`$(Script:CMPVMUtils.getCredentialAlias[arg=${Stack.items[Virtual Server].attributes[sys_id]}])`

Expression to get the IP address: `$(Script:CMPVMUtils.getReachableIp[arg=$(Stack.items[Virtual Server].attributes[sys_id])])`

CI lookup expression

Accesses values of fields in a table, usually a CI table in the CMDB.

- Type: definitional/runtime
- Syntax: `$(ci.tableName)`
- Examples:

  `$(ci.sn_cmp _ip_pool[subnet=${parameter.formData.Subnet Id}])` combines runtime and definition type expressions. The = operator evaluates a value for a match. The expression dot-walks to the IP Pool table and looks for the subnet that has the subnet ID that the user submitted.

  `$(ci.cmdb_ci_cloud_subnet[${parameter.formData.SubnetId}].cidr)` takes the value of the `cidr` field from a subnet that the user specified during provisioning. The expression dot-walks to the Cloud Subnet table, looks at the subnet value that the user submitted on a blueprint provision, extracts the `cidr` field value, and then walks to the value in the `cidr` field of the sys_id of the subnet.

Random number expression

Generates a random number. For example, you can provision a VM with a random node name or stack name.

- Type: runtime
- Syntax and example: `${randomNumber}`

Scratchpad expression/Resource operation output attribute expression

Reads output attributes from one operation into another operation.

- Type: runtime
- Syntax: varies
- Example: To set/expose outputs from one operation:

  `${Compute Interface.CreateNode.Output.resp.nodeId}`

  To read the output attributes (where `VM1.Provision` is the operation whose output attributes are read):

  `$(Outputs[VM1.Provision].NodeId)`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Conditions in expressions

You can use conditional expressions in blueprint steps and resource operation steps to conditionally execute or skip the step. The conditional expressions are Javascript expressions and they support expression substitutions.

- Type: definitional/runtime
- Syntax: varies
- Examples:

  '${parameter.CloudAccount}'=='Amazon Cloud Account'

  '$(Stack.items[vm1].attributes[node_id]) '== 'VM1'

Complex expression

You can nest expressions. In the example, `${parameter.ServerID}` maps the sys_id of the CI and is replaced before the outer expression is consumed.

- Type: any
- Syntax: varies
- Example: `$(ci.cmdb_ci_vm_instance[${parameter.ServerID}].name)`

Order context

This expression is useful for life cycle operations and enables you to dot-walk on the order attributes as part of the *sn_cmp_order* table.

- Type: runtime
- Syntax: `$(context.order.column_name)`
- Examples:
  - Request Item: `$(context.order.sc_req_item)`
  - Dot-walk on request item object: `$(context.order.sc_req_item.number)`
  - Mixed Expression (Constant + Expression): `"ram $(context.order.sc_req_item.number)"`

Create an action for an 'on Blueprint provision' policy

The *on Blueprint provision* trigger fires after execution of *on Catalog item request start* policies. A policy that is triggered by the *on Blueprint provision* trigger can run a script, override a user-requested attribute value, or abort and send a message about the provision operation.

Before you begin

Optional: Create one or more cloud policy groups.

Configure a cloud policy rule

Role required: sn_cmp.cloud_governor or admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform, enter a unique and meaningful **Action Name**, and then fill in the form for the action.

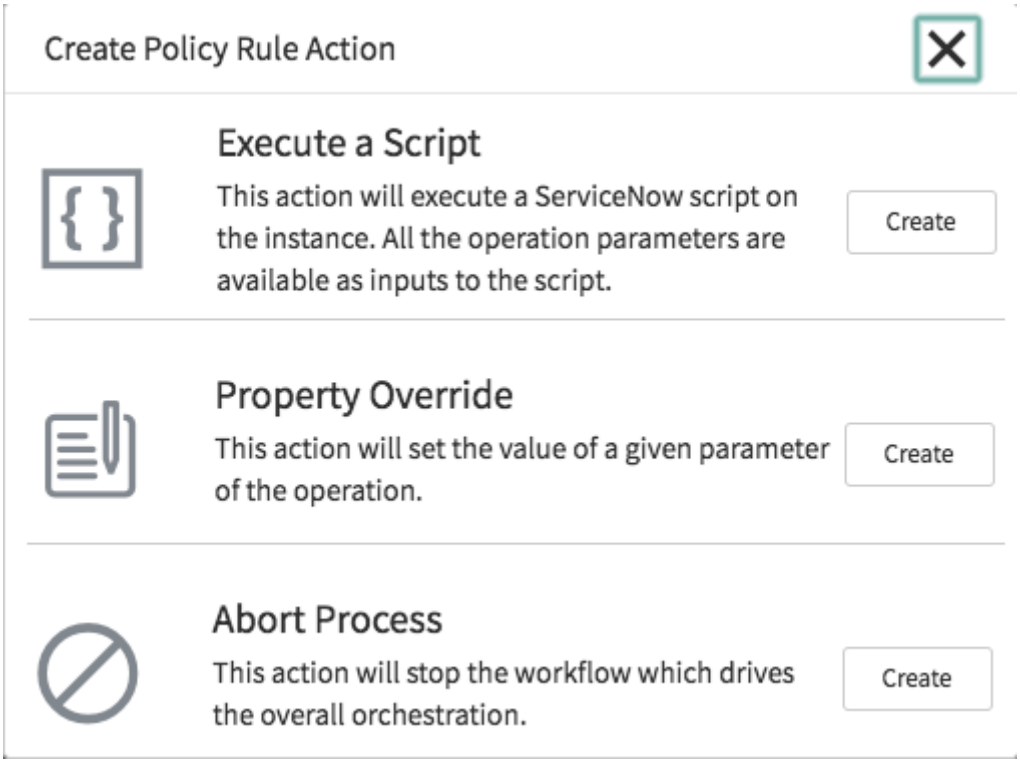

Settings for the 'Execute a Script' action

Field	Description
Action Script Category	Select a category.
Action Script Name	Specify a unique and meaningful name for the script.
Action Script	Create the script in the text box. See Create a policy action script for details.

Settings for the 'Property Override' action

Field	Description
Property	Specify the name of the property (attribute) on the user request form to override.
Value	Enter a value that overrides the value in the **Property** field. You can override text values only. You can specify a static value, an expression, or both. The example action, named **SetTheCostCenter**, specifies the value **Marketing** for the **CostCenter** property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Configure a Property Override action 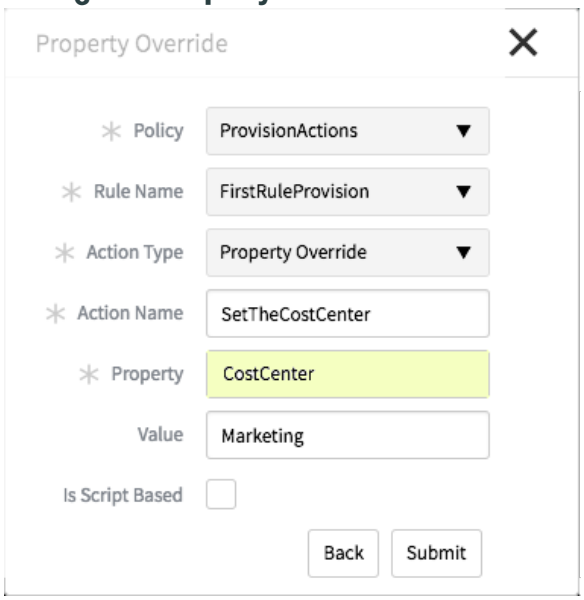 **Note:** When both a policy rule and a form rule overwrite a value, the value in the form rule is used. Expressions can perform the following actions (see Using expressions in Cloud Provisioning and Governance for details): ○ Set form data values using definition expressions. For example: `${parameter.formData.CatalogAttributeType}` ○ Assign user data values using definition expressions. For example: `${parameter.userData.userId}` For example, the following value can set the stack name to **Stack_Bob.Smith@company.com**: `Stack_${parameter.userData.userId}` ○ Set stack or table values using runtime expressions. For example: `$(ci.sn_cmp_ip_pool[subnet=${parameter.formData.Subnet Id}])` takes the subnet from the IP Pools table. ○ Associate a random number with a field using static expressions. Use: `${randomNumber}`
Is Script Based	Select the check box to display the **Script** text box and then specify the script. You can use the following example script snippet to override a stack name. The `function( formData)` section of the script modifies the values for fields on the form. `MyStack` is the stack name in this example. `customScript : function( formData){ // Manipulation of form parameter is only supported here. // Change in any other attributes will be ignored`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description

	`// data available for manipulation are`
	`// Form Data - Ex. StackName can be accessed`
	`// through formData.StackName`
	`// formData.StackName = "MyStack";`
	`// User Data - Ex. User Id can be accessed`
	`// through this.parameters.userData`
	`// if(this.parameter.userData.userId == 'servicenowuserId')`
	`formData.StackName = "MyStack";`
	`return formData;`
	`},`

Setting for the 'Abort Process' action

Field	Description
Message	Enter the message to present to the requester when the process aborts.

Create an action for an approval policy

A policy that is triggered by one of the approval triggers can start approval subflows. The approval triggers are (*on Blueprint provision (approval)*, *on Stack operation (approval)*, *on Stack resource operation (approval)*, and *on Task remediation*.

Before you begin

Optional: Create one or more cloud policy groups.

Configure a cloud policy rule and associated conditions.

Role required: sn_cmp.cloud_governor or admin

About this task

When multiple "approval" policies apply, the policies are applied in the following order. (The approval policies are *on Blueprint provision (approval)*, *on Stack operation (approval)*, *on Stack resource operation (approval)*, and *on Task remediation*):

1. Only the first successful approval policy is applied and no other approval policies are applied.
2. If the applied approval policy has multiple rules, only the first successful rule is used.
3. If a rule has multiple actions, only the first successful action is performed.
4. If the applied approval policy includes both a custom approval and a Service Now approval, only the custom approval process is performed.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

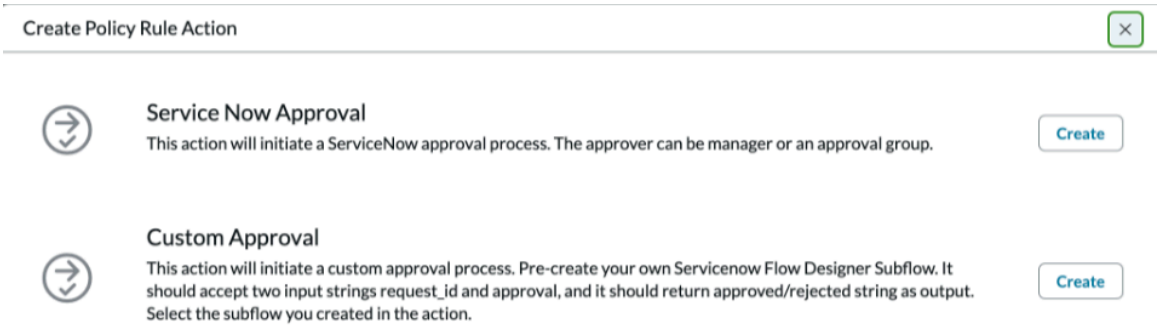

5. Select any of the following actions.
   - **Custom Approval** runs the subflow that you specify. The subflow should return a value of **approved** to complete the operation. Follow **Service Now Approval** for use case.
   - **Service Now Approval** supports differing approvers based on the trigger for the policy:

Trigger	Subflow	Available approvers
on Blueprint provision (approval)	Cloud Approval Subflow	▪ Manager Approval ▪ Assignment Group ▪ User (lock)
on Stack operation (approval)	Cloud Operation Change Request Subflow	Assignment Group
on Stack resource operation (approval)	Cloud Operation Change Request Subflow	Assignment Group

6. On the Approval form, specify a unique and meaningful **Action Name**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action in the rule

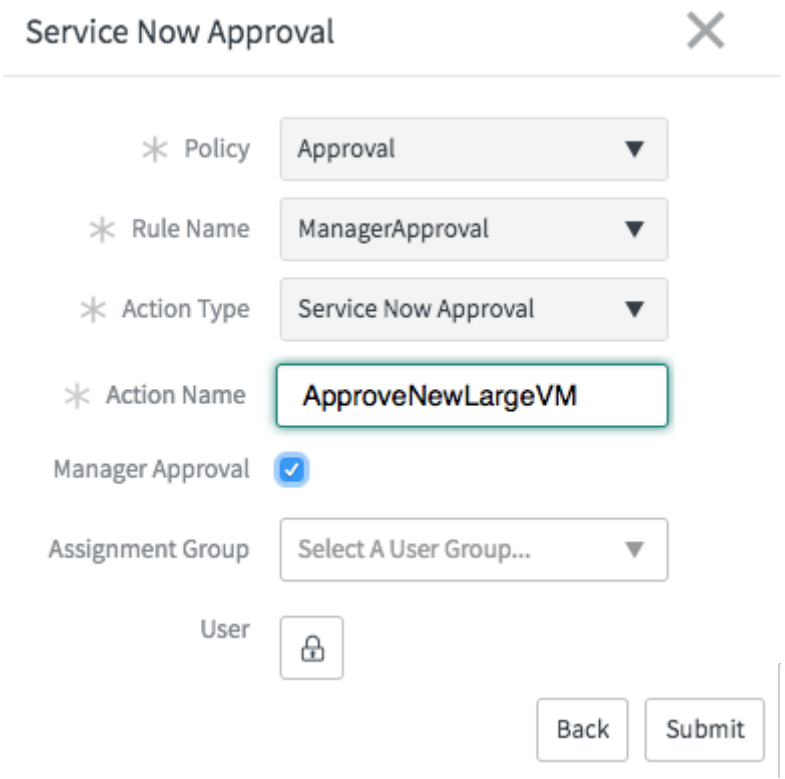

7. If you select **Custom Approval**, specify the subflow and then click **Submit** and if you select **ServiceNow Approval**, specify who should approve the cloud activity.

Approval fields

Field	Description
Manager Approval	One or all of the following: ○ Select the **Manager Approval** check box to require the manager of the approver to also approve the request. The default approval subflow goes to the manager first, then to the group, and finally to the user. ○ Select a user group from the **Assignment Group** list. ○ Unlock the **User** lock, select a user from the list, then close the lock.
Assignment group	Select the assignment group that can approve the request. Any user in the group can approve the action using the default subflow. The approval then goes to the users in the **User** field. **Note:** Users in an assignment group that do not have the itil role still receive an approval record and notification by default, but they cannot perform the approval. Assign the itil role to any users or user groups that must make operational and provisioning approvals.
User	Select one or more users to whom the approval action applies. All selected users must approve the action using the default subflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create an action for an 'on Catalog item launch' policy

The `on Catalog item launch` trigger fires when an order form (stack request form) is launched for a catalog item. A policy that is triggered by the `on Catalog item launch` trigger can run a script or override a user-requested value (text values only).

Before you begin

- Role required: sn_cmp.cloud_governor or admin
- Optional: Create one or more cloud policy groups.
- Configure a cloud policy rule and associated conditions.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform, enter a unique and meaningful **Action Name**, and then fill in the form for the action.

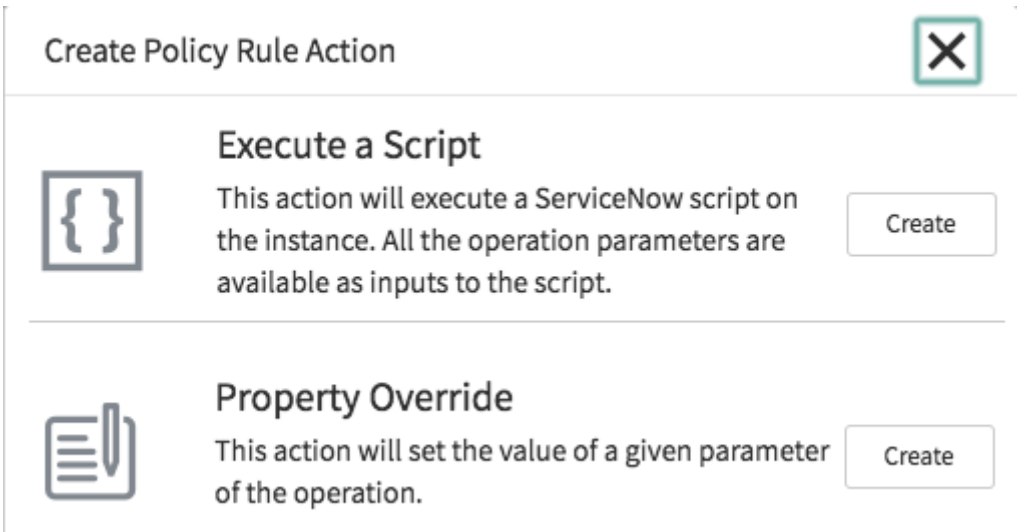

Settings for the 'Execute a Script' action

Field	Description
Action Script Category	Select a category.
Action Script Name	Specify a unique and meaningful name for the script.
Action Script	Create the script in the text box. See Create a policy action script for details.

Settings for the 'Property Override' action

Field	Description
Property	Specify the name of the property (attribute) on the user request form to override.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Value	Enter a value that overrides the value in the **Property** field. You can override text values only. You can specify a static value, an expression, or both. The example action, named **SetTheCostCenter**, specifies the value **Marketing** for the **CostCenter** property. **Configure a Property Override action** 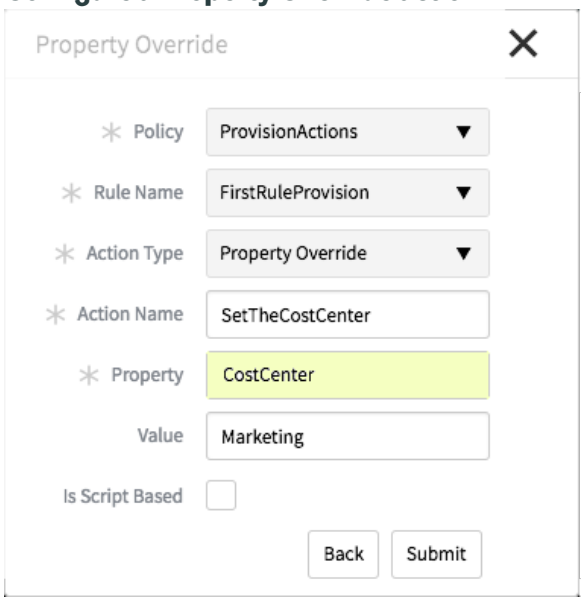 **Note:** When both a policy rule and a form rule overwrite a value, the value in the form rule is used. Expressions can perform the following actions (see Using expressions in Cloud Provisioning and Governance for details): ○ Set form data values using definition expressions. For example: `${parameter.formData.CatalogAttributeType}` ○ Assign user data values using definition expressions. For example: `${parameter.userData.userId}` For example, the following value can set the stack name to **Stack_Bob.Smith@company.com**: `Stack_${parameter.userData.userId}` ○ Set stack or table values using runtime expressions. For example: `$(ci.sn_cmp_ip_pool[subnet=${parameter.formData.Subnet Id}])` takes the subnet from the IP Pools table. ○ Associate a random number with a field using static expressions. Use: `${randomNumber}`
Is Script Based	Select the check box to display the **Script** text box and then specify the script. You can use the following example script snippet to override a stack name. The `function( formData)` section of the script modifies the values for fields on the form. `MyStack` is the stack name in this example.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	`customScript : function(formData){ // Manipulation of form parameter is only supported here. // Change in any other attributes will be ignored // data available for manipulation are // Form Data - Ex. StackName can be accessed // through formData.StackName // formData.StackName = "MyStack"; // User Data - Ex. User Id can be accessed // through this.parameters.userData // if(this.parameter.userData.userId == 'servicenowuserId') formData.StackName = "MyStack"; return formData; },`

Create action for 'on Catalog item request start/end'

The *on Catalog item request start* trigger fires after the user opens a request form. The *on Catalog item request end* trigger fires after a user submits a request form. A policy that is triggered by the *on Catalog item request start* or *on Catalog item request end* trigger can run a script or execute a subflow.

Before you begin

- Role required: sn_cmp.cloud_governor or admin
- Optional: Create one or more cloud policy groups.
- Configure a cloud policy rule and associated conditions.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform, enter a unique and meaningful **Action Name**, and then fill in the form for the action.

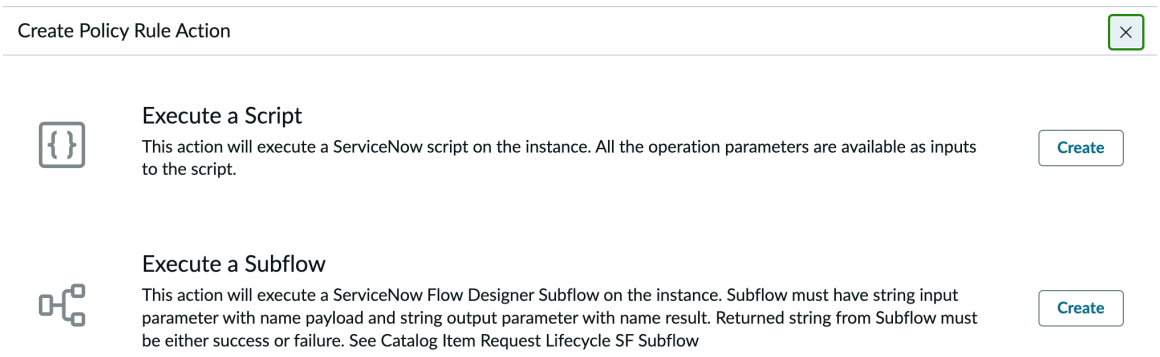

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Settings for the 'Execute a Script' action

Field	Description
Action Script Category	Select a category.
Action Script Name	Specify a unique and meaningful name for the script.
Action Script	Create the script in the text box. See Create a policy action script for details.

Configuring an 'Execute a Subflow' action to run the base-system 'Change Request' subflow

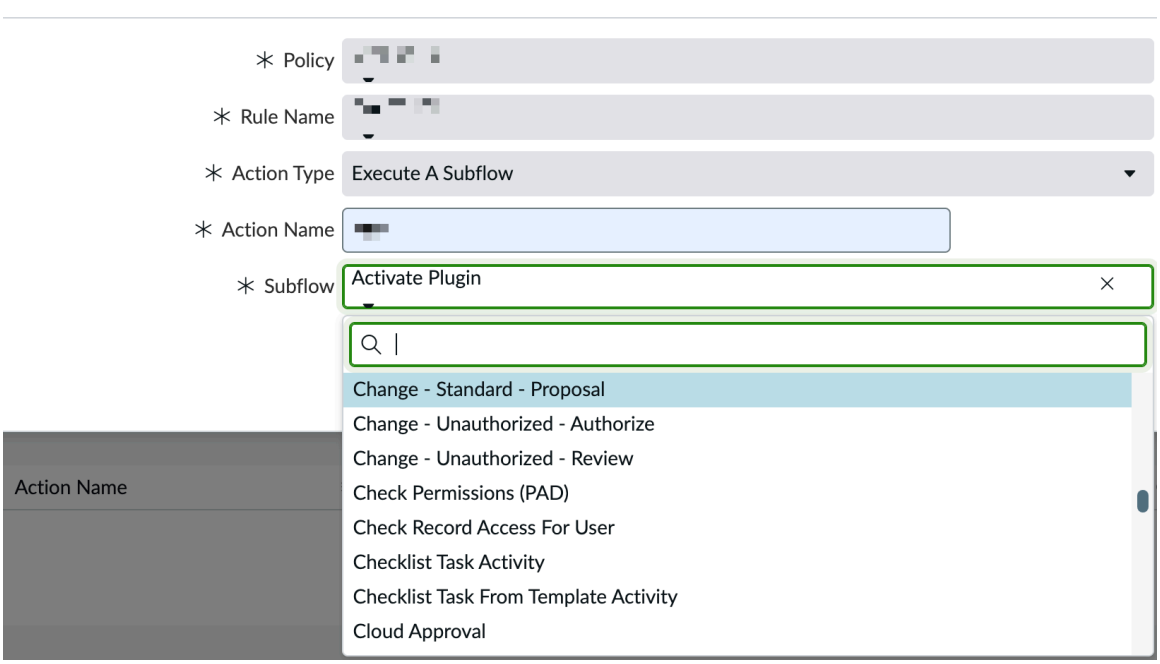

Settings for the 'Execute a Subflow' action

Field	Description
Action Name	Specify a unique and meaningful name for the action that starts the subflow.
Subflow	Select the subflow to execute.

Create an action for an 'on Lease end' policy

A policy that is triggered by the `on Lease end` trigger can send a notification or perform a `Start`, `Stop`, or `Deprovision` life cycle operation.

Before you begin

- Role required: sn_cmp.cloud_governor or admin
- Optional: Create one or more cloud policy groups.
- Configure a cloud policy rule and associated conditions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform, enter a unique and meaningful **Action Name**, and then fill in the form for the action.

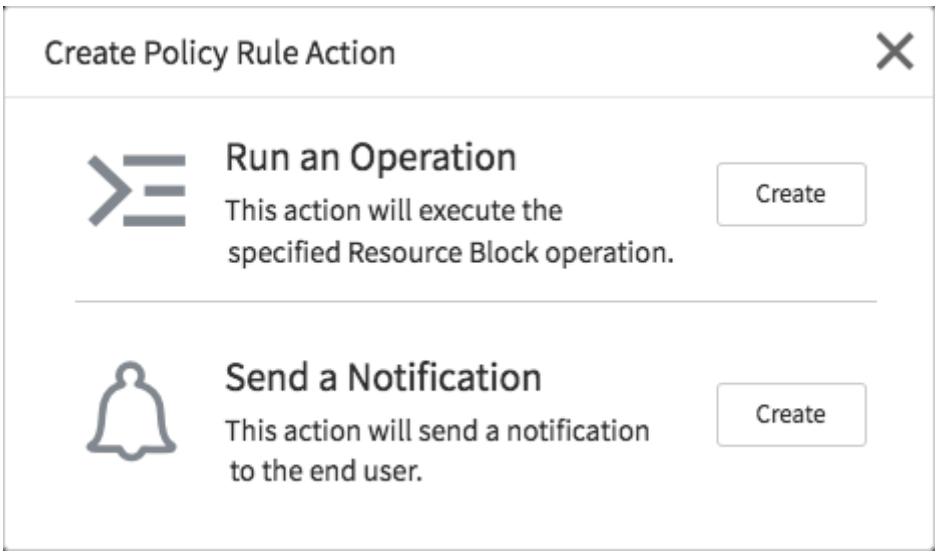

Settings for the 'Run an Operation' action

Field	Description
Blueprint Operation	Specify **Start**, **Stop**, or **Deprovision**.

Settings for the 'Send a Notification' action

Field	Description
Notification	Select the text for the notification: ○ Lease End Date Reached ○ Stack Decommissioned (Lease End) ○ Upcoming Lease End

Create an action for an 'on Resource operation' policy

The `on Resource operation` trigger fires during the Orchestration process when a user performs a `Start`, `Stop`, or `Deprovision` life cycle operation on a specific resource. A policy that is triggered by the `on Resource operation` trigger can override a user-requested attribute value, run a script, call a Cloud API, or perform an IP address management operation.

Before you begin

- Role required: sn_cmp.cloud_governor or admin
- Optional: Create one or more cloud policy groups.
- Configure a cloud policy rule and associated conditions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform, enter a unique and meaningful **Action Name**, and then fill in the form for the action.

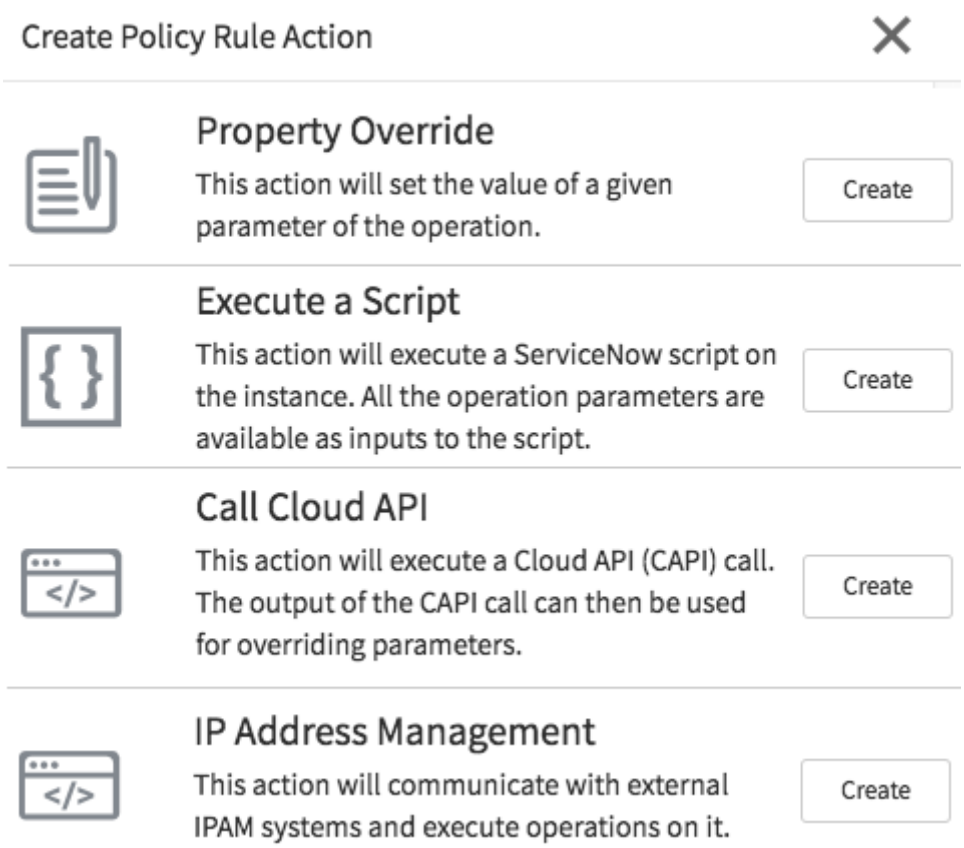

Settings for the 'Property Override' action

Field	Description
Property	Specify the name of the property (attribute) on the user request form to override.
Value	Enter a value that overrides the value in the **Property** field. You can override text values only. You can specify a static value, an expression, or both. The example action, named **SetTheCostCenter**, specifies the value **Marketing** for the **CostCenter** property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Configure a Property Override action 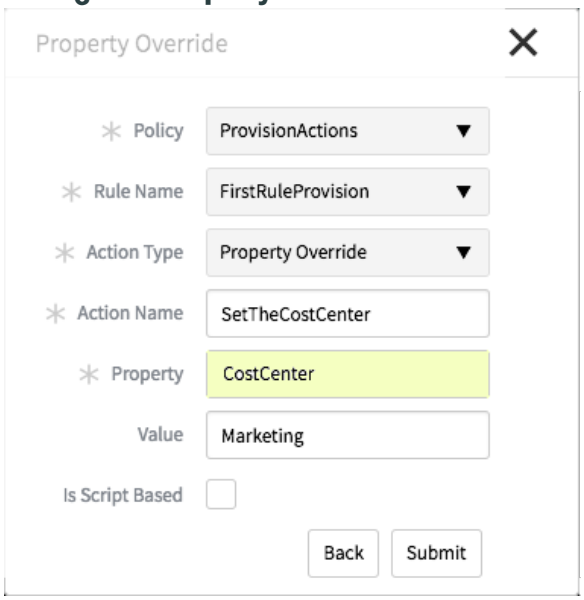 **Note:** When both a policy rule and a form rule overwrite a value, the value in the form rule is used. Expressions can perform the following actions (see Using expressions in Cloud Provisioning and Governance for details): ○ Set form data values using definition expressions. For example: `${parameter.formData.CatalogAttributeType}` ○ Assign user data values using definition expressions. For example: `${parameter.userData.userId}` For example, the following value can set the stack name to **Stack_Bob.Smith@company.com**: `Stack_${parameter.userData.userId}` ○ Set stack or table values using runtime expressions. For example: `$(ci.sn_cmp_ip_pool[subnet=${parameter.formData.Subnet Id}])` takes the subnet from the IP Pools table. ○ Associate a random number with a field using static expressions. Use: `${randomNumber}`
Is Script Based	Select the check box to display the **Script** text box and then specify the script. You can use the following example script snippet to override a stack name. The `function( formData)` section of the script modifies the values for fields on the form. `MyStack` is the stack name in this example. `customScript : function( formData){ // Manipulation of form parameter is only supported here. // Change in any other attributes will be ignored`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description

	``
	// data available for manipulation are
	// Form Data - Ex. StackName can be accessed
	// through formData.StackName
	// formData.StackName = "MyStack";
	// User Data - Ex. User Id can be accessed
	// through this.parameters.userData
	// if(this.parameter.userData.userId == 'servicenowuserId')
	formData.StackName = "MyStack";
	return formData;
	},
	``

Settings for the 'Execute a Script' action

Field	Description
Action Script Category	Select a category.
Action Script Name	Specify a unique and meaningful name for the script.
Action Script	Create the script in the text box. See Create a policy action script for details.

If you select the **Call Cloud API** action type and save the record, the Policy Rule Action Attributes related list is populated with the attributes to pass to the provider. If you are integrating with Infoblox, you must configure the **DNSSuffix** field. When you create a host, for example, the value in the **DNSSuffix** field appears in Infoblox for the newly created virtual machine.

Settings for the 'Call Cloud API' action

Field	Description
Action Name	Enter a descriptive name for the action. In subsequent policy actions for the same policy, you can reference this name in a **Value** field. It refers to the JSON payload that is received from the cloud provider after an API call. For example, to reference an allocated IP address that was returned by the action that is named *RefAction*, you can use the following expression in the **Value** field of another policy action: `${parameter.RefAction.Allocated_IP}`.
Provider	Enter the cloud provider from the API library.
Version	Enter the version from the API library.
Interface	Select a CAPI interface. For Infoblox, select **IPAM interface**.
Operation	Select a CAPI interface operation to perform. For Infoblox, select **CreateHostRecord** or **DeleteHostRecord**.
Credentials	Enter the sys_ID of the necessary credentials. You can also use an expression like:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	`$(ci.sn_cmp_infoblox_server[$(ci.sn_cmp_ip_pool[subnet=${parameter.formData.SubnetId}].provider_instance)].infoblox_server_credential)`
Endpoint URL	Enter the endpoint URL of the connector from the API library. For Infoblox, it is the IP address of the Infoblox server.

Settings for the 'IP Address Management' action

Field	Description
IPAM Method Name	Select the IPAM Method Name. For Infoblox, select one of the following options: ○ **Release IP Address**: Release an IP address from Infoblox when a machine is deprovisioned. ○ **Register IP Address**: Register a new virtual machine to Infoblox. This option is used by AWS and Azure clouds. ○ **Reserve IP Address**: Reserve a new IP address. This option is used by a vSphere cloud. See Register IP addresses for AWS and Azure VMs in Infoblox, Reserve IP addresses for VMware vSphere VMs in InfoBlox, and IPAM integration.

Create action for 'on Resource operation launch' policy

The *on Resource operation launch* trigger fires before the catalog for a resource operation is loaded from the Cloud User Portal. A policy that is triggered by the *on Resource operation launch* trigger can run a script or can override a user-requested value (text values only).

Before you begin

- Role required: sn_cmp.cloud_governor or admin
- Optional: Create one or more cloud policy groups.
- Configure a cloud policy rule and associated conditions.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform, enter a unique and meaningful **Action Name**, and then fill in the form for the action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

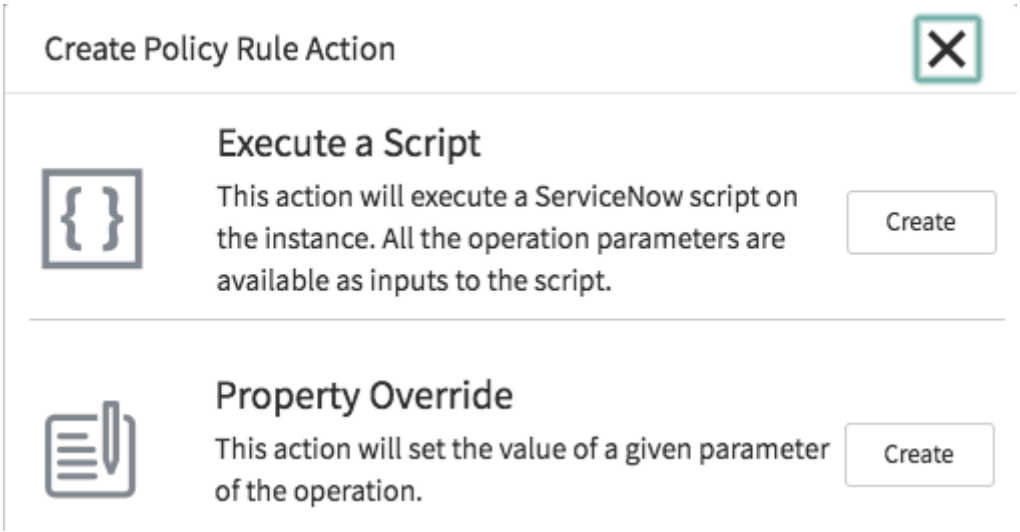

Settings for the 'Execute a Script' action

Field	Description
Action Script Category	Select a category.
Action Script Name	Specify a unique and meaningful name for the script.
Action Script	Create the script in the text box. See Create a policy action script for details.

Settings for the 'Property Override' action

Field	Description
Property	Specify the name of the property (attribute) on the user request form to override.
Value	Enter a value that overrides the value in the **Property** field. You can override text values only. You can specify a static value, an expression, or both. The example action, named **SetTheCostCenter**, specifies the value **Marketing** for the **CostCenter** property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Configure a Property Override action 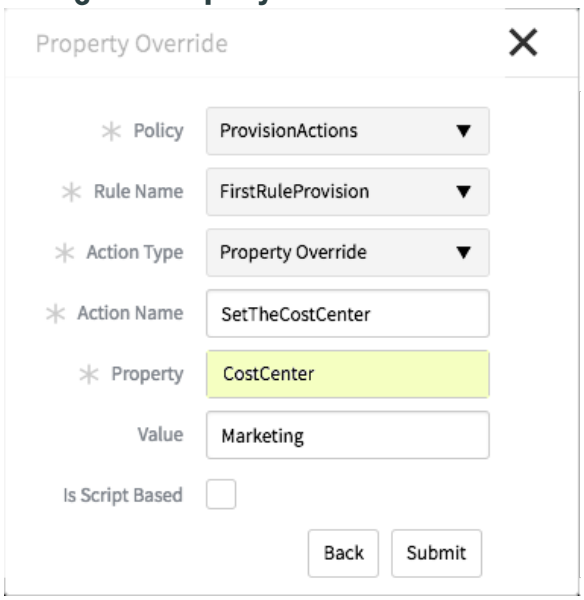 **Note:** When both a policy rule and a form rule overwrite a value, the value in the form rule is used. Expressions can perform the following actions (see Using expressions in Cloud Provisioning and Governance for details): ○ Set form data values using definition expressions. For example: `${parameter.formData.CatalogAttributeType}` ○ Assign user data values using definition expressions. For example: `${parameter.userData.userId}` For example, the following value can set the stack name to **Stack_Bob.Smith@company.com**: `Stack_${parameter.userData.userId}` ○ Set stack or table values using runtime expressions. For example: `$(ci.sn_cmp_ip_pool[subnet=${parameter.formData.Subnet Id}])` takes the subnet from the IP Pools table. ○ Associate a random number with a field using static expressions. Use: `${randomNumber}`
Is Script Based	Select the check box to display the **Script** text box and then specify the script. You can use the following example script snippet to override a stack name. The `function( formData)` section of the script modifies the values for fields on the form. `MyStack` is the stack name in this example. `customScript : function( formData){ // Manipulation of form parameter is only supported here. // Change in any other attributes will be ignored`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description

	`// data available for manipulation are`
	`// Form Data - Ex. StackName can be accessed`
	`// through formData.StackName`
	`// formData.StackName = "MyStack";`
	`// User Data - Ex. User Id can be accessed`
	`// through this.parameters.userData`
	`// if(this.parameter.userData.userId == 'servicenowuserId')`
	`formData.StackName = "MyStack";`
	`return formData;`
	`},`

Create action for 'on Resource operation request start/end' policy

A policy that is triggered by the *on Resource operation request start* or *on Resource operation request end* trigger can run a script or override a user-requested attribute value.

Before you begin

- Role required: sn_cmp.cloud_governor or admin
- Optional: Create one or more cloud policy groups.
- Configure a cloud policy rule and associated conditions.

About this task

- The *on Resource operation request start* trigger fires after a user submits a resource operation request (*Start*, *Stop*, *Deprovision*).
- The *on Resource operation request end* trigger fires before completion of a life cycle operation on a resource (*Start*, *Stop*, *Deprovision*).

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform, enter a unique and meaningful **Action Name**, and then fill in the form for the action.

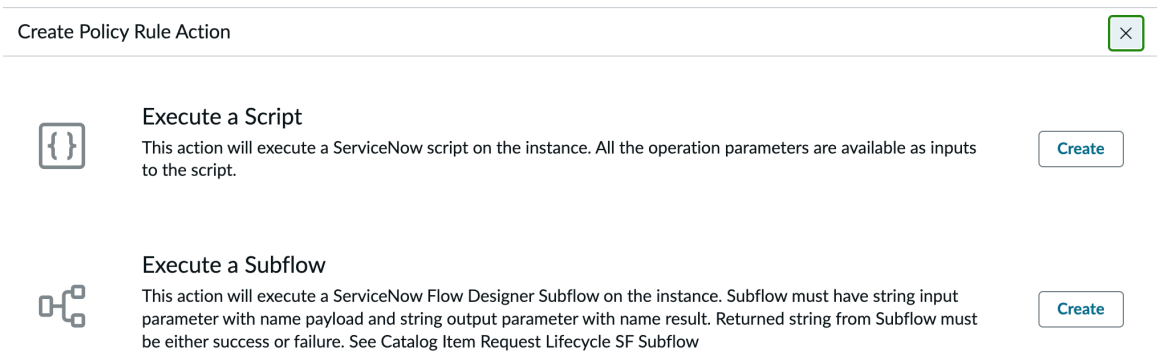

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Settings for the 'Execute a Script' action

Field	Description
Action Script Category	Select a category.
Action Script Name	Specify a unique and meaningful name for the script.
Action Script	Create the script in the text box. See Create a policy action script for details.

Settings for the 'Execute a Subflow' action

Field	Description
Action Name	Specify a unique and meaningful name for the action that starts the subflow.
Subflow	Select the subflow to execute.

Configuring an 'Execute a Subflow' action to run the base-system 'Change Request' subflow

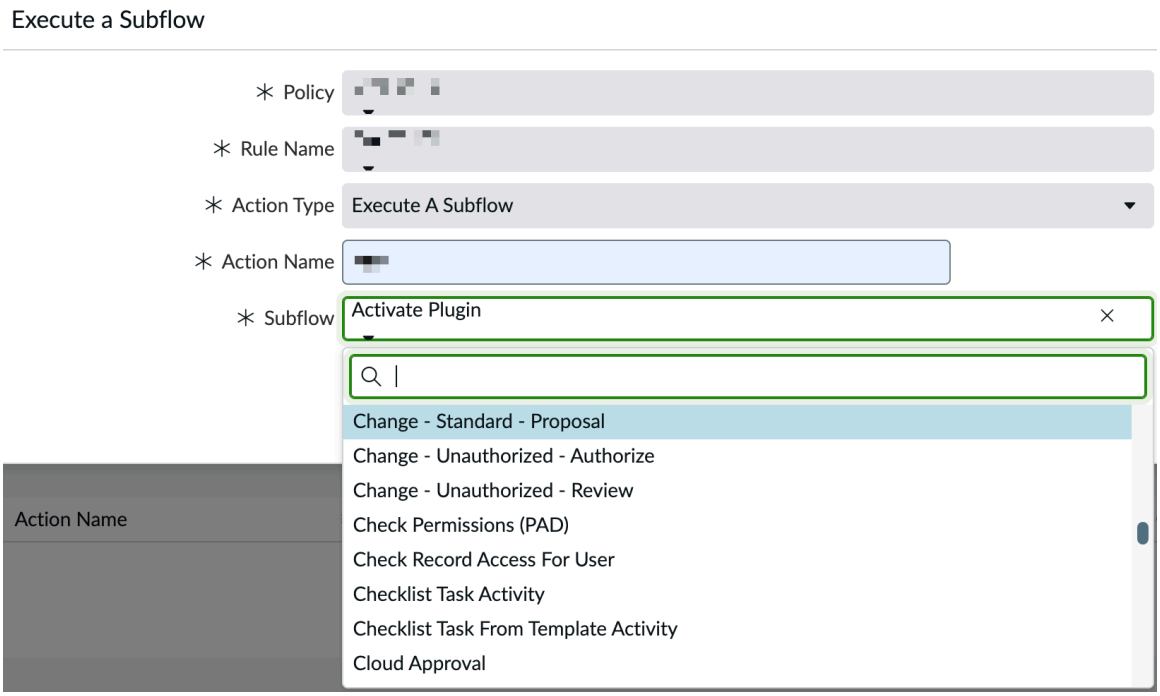

Create an action for an 'on Resource Limit exceeded' policy

Set the *on Resource limits exceeded* policy to automatically run approval subflows or send notifications when a template-based catalog item order exceeds defined resource limits.

Before you begin

You must have configured the following items:

- One or more cloud policy groups
- Cloud policy
- Approval policy

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Email notification
- Create a subflow in Workflow Studio

Role required: sn_cmp.cloud_governor or cloud_admin

About this task

The *On resource limit exceeded* base system policy action has a default ServiceNow approval policy rule configured by default. You can trigger other actions such as notifications or a custom approval subflow when the resource limit is exceeded. Modify the base system policy action or create a custom policy for your instance.

Procedure

1. Modify or create a cloud policy action to use the *on Resource limit exceeded* trigger.
   a. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
   b. Modify a base system policy or create a new policy record.
   - a. Open the base system **Resource Order Control** policy.
     b. Modify the policy name, start, and end dates and order of execution.
     c. Modify the policy rules conditions.
   - Click **New** to create a new policy record.

     For more information, see Create a cloud policy.
2. In the Rules related list, modify the base system rule or create a new rule.
   - Open the base system *approveresourcelimit* rule that should perform the action and customize it.
   - Create a policy rule.

     For more information, see Configure a cloud policy rule.
3. Specify a policy action.
   a. In the Policy Rule Actions related list, click **New**.
   b. Click **Create** next to the action you want to occur when the quota is exceeded.

Option	Steps
ServiceNow Approval	**a.** Enter a unique name for the action in the **Action Name** field. **b.** Select the **Manager Approval** option to au tomatically send the request for manager's approval. **c.** Select the **Assignment group** to which this action applies.
Custom Approval	**a.** Enter a unique name for the action in the **Action Name** field. **b.** Select an existing custom workflow in the **Subflow** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Steps
Send aNotification	**a.** Enter a unique name for the action in the **Action Name** field. **b.** Select a notification template in the **Notifi cation** field. The base system notification email for the 'on Resource Limit exceeded' policy action appears in the lookup table if you're using the base system policy.

4. Click **Submit**.

Create an action for an 'on Task Remediation' policy

The `on Task remediation` trigger fires when a user resubmits a failed request. A policy that is triggered by the `on Task Remediation` trigger can start approval subflows.

Before you begin

- Role required: sn_cmp.cloud_governor or admin
- Optional: Create one or more cloud policy groups.
- Configure a cloud policy rule and associated conditions.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and set the policy to the **Draft** state if needed.
3. Open the rule that should perform the action and then click **New** on the Policy Rule Actions related list.
4. On the popup, click **Create** for the type of action to perform.

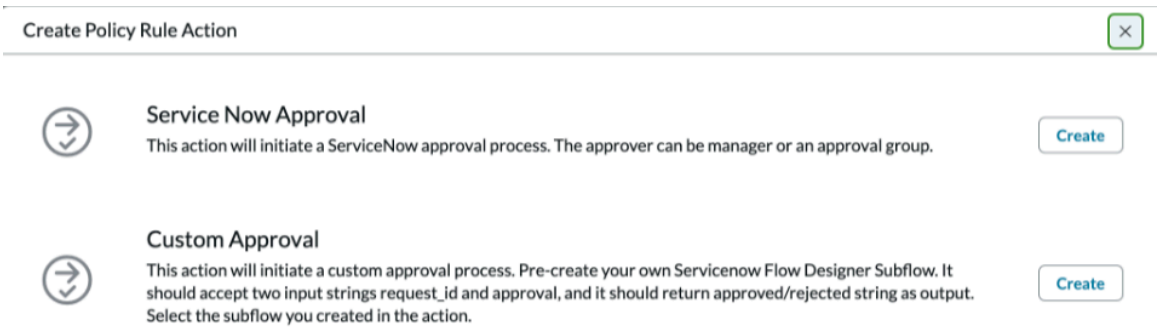

5. Select one of the following options.
   - **Custom Approval** runs the subflow that you specify. The subflow should return a value of **approved** to complete the operation.
   - **Service Now Approval** supports differing approvers based on the trigger for the policy:
     - Manager Approval
     - Assignment Group
     - User (lock)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. On the Approval form, specify a unique and meaningful **Action Name**.

Action in the rule

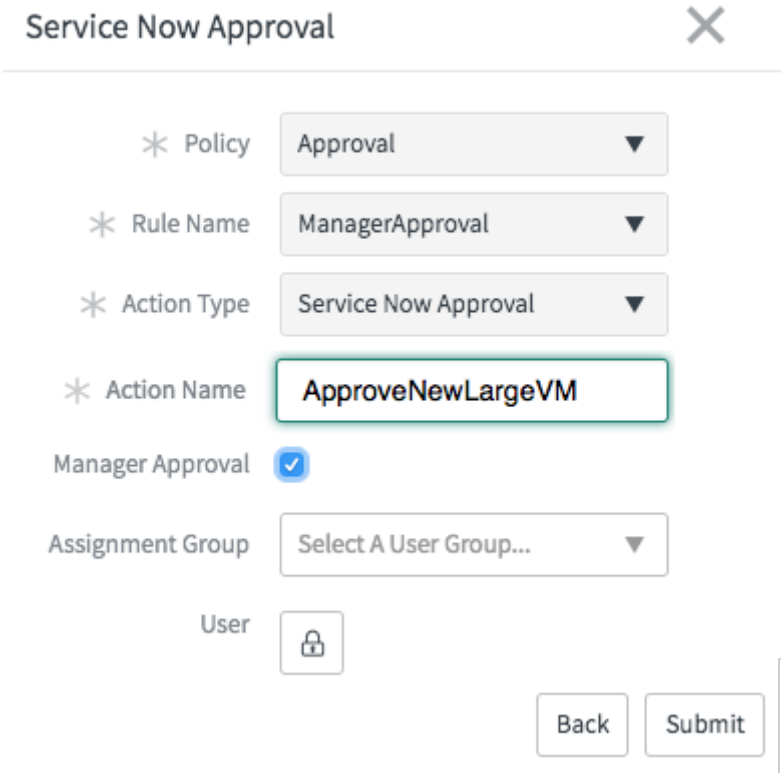

7. If you select **Custom Approval**, specify the subflow and then click **Submit** and if you select **ServiceNow Approval**, specify who should approve the cloud activity.

Approval fields

Field	Description
Manager Approval	One or all of the following: ○ Select the **Manager Approval** check box to require the manager of the approver to also approve the request. The default approval subflow goes to the manager first, then to the group, and finally to the user. ○ Select a user group from the **Assignment Group** list. ○ Unlock the **User** lock, select a user from the list, then close the lock.
Assignment group	Select the assignment group that can approve the request. Any user in the group can approve the action using the default subflow. The approval then goes to the users in the **User** field. **Note:** Users in an assignment group that do not have the itil role still receive an approval record and notification by default, but they cannot perform the approval. Assign the itil role to any users or user groups that must make operational and provisioning approvals.
User	Select one or more users to whom the approval action applies. All selected users must approve the action using the default subflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a policy action script

Use policy action scripts to get, update, or set values in cloud requests.

Before you begin

Role required: sn_cmp.cloud_governor or admin.

About this task

For more information on using scripts, see Cloud scripts and cloud script templates.

Procedure

1. On the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. On the **Policy Action Scripts** tab, click **New**, enter a unique and descriptive **Name** and **Description**, fill in the form, and then click **Submit**.
   After you enter a name, a script appears with a variable that uses the name that you entered.

Field	Description
Action Script Category	Choose another category or keep the one you created in an earlier step.
Script	Modify the `customScripts` function in the script. This is the only part of the script that you need to modify.

Example:

The following example script determines whether the user belongs to a user group and then overrides the user-specified application server node name value with the value **MyNodeName**.

```
customScript : function(formData){
  var groups = this.parameters.userData.groups;
    if( groups.length > 0){
      formData.App_Server_NodeName = "MyNodeName";
    }
  return formData;
},
```

This example returns a random number:

```
customScript : function(formData){
  var min = 10000;
  var max = 99999;
  var num = Math.floor(Math.random() * (max - min + 1)) + min;
  formData.randomNumber = num + "";
  return formData;
},
```

This example script adds three hours to the lease date using the base-system GlideDateTime class.

```
var SetLeaseShort = Class.create();
SetLeaseShort.prototype =
 Object.extendsObject(sn_cmp.PolicyExecutionBase, {
 customScript : function(formData){
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
  // Manipulation of form parameters are supported. Changes in
 any other attributes
             // will be ignored.
  // The data available for manipulation is as follows:
  // Form Data -  for example: StackName can be accessed through
 formData.StackName
  // formData.StackName = "MyStack";
  // User Data  - for example: User Id can be accessed through
 this.parameters.userData
  // if(this.parameter.userData.userId == 'servicenowuserId')
  // this.info("id = " + formData.StackName);

  // var leaseValue = now();
  // leaseValue = leaseValue.setHours(leaseValue.getHours()+4);
  // formData.LeaseEndDate = leaseValue;

  if (formData.hasOwnProperty("LeaseEndDate")) {
   gdt = new GlideDateTime();
   // 10800 = 3 hours
   gdt.addSeconds(10800);
   formData.LeaseEndDate = gdt.getDisplayValue();
  }
  return formData;
 },
 execute: function() {
  if(this.parameters !=  null && this.parameters.formData !=
 null){
        var inputData =
 JSON.parse(JSON.stringify(this.parameters.formData));
     var outputFormData = this.customScript(inputData);
           if( outputFormData != null){
               this.parameters.formData = outputFormData;
           }
       }
  var output = {};
  output.answer = this.parameters;
  return JSON.stringify(output);
 },
 type : 'SetLeaseShort'
});

var SetLeaseShortObj = new SetLeaseShort(inputAttributes);
var outputParams = SetLeaseShortObj.execute();
gs.info( "Output of script is = " + outputParams);
```

This example script sets the user-specified **Azure Wordpress DNS** name field from the form to the same name as the stack name:

```
var WordpressFill = Class.create();
WordpressFill.prototype =
 Object.extendsObject(sn_cmp.PolicyExecutionBase, {
 customScript : function(formData){
  // Manipulation of form parameters are supported. Changes in
 any other attributes
        // will be ignored.
  // The data available for manipulation is as follows:
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 // Form Data -  for example: StackName can be accessed through
formData.StackName
 // formData.StackName = "MyStack";
 // User Data  - for example: User Id can be accessed through
this.parameters.userData
 // if(this.parameter.userData.userId == 'servicenowuserId')
 //this.info("id = " + formData.StackName);

 formData.Azure_Wordpress_v6_vmDnsName = formData.StackName;

 return formData;
},
execute: function() {
 if(this.parameters !=  null && this.parameters.formData !=
null){
        var inputData =
JSON.parse(JSON.stringify(this.parameters.formData));
     var outputFormData = this.customScript(inputData);
           if( outputFormData != null){
               this.parameters.formData = outputFormData;
           }
       }
 var output = {};
 output.answer = this.parameters;
 return JSON.stringify(output);
},
type : 'WordpressFill'
});

var WordpressFillObj = new WordpressFill(inputAttributes);
var outputParams = WordpressFillObj.execute();
gs.info( "Output of script is = " + outputParams);
```

Related topics

Example policy action script that tags resources

Example policy action script that tags resources

Your instance tracks tagged resources for billing and reporting. Policy action scripts can add and modify resource tags. This example tags a storage volume resource.

Before you begin

Role required: sn_cmp.cloud_governor or admin

About this task

See Create a custom tag for cloud resources for information about how to create and use tags in Cloud Provisioning and Governance.

Procedure

1. On the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. On the **Policy Action Scripts** tab, click **New** and then enter a unique and descriptive **Name**.
3. Click the lookup icon for **Action Script Category**, click **New** to create a new category, and then click **Submit**.
4. On the Policy Rule Action Script page, enter the script for the resource you want to tag.
5. Click **Submit** to save the script.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The script appears in the list.

6. On the Policies tab, click **New**, enter a unique and descriptive **Name** for the policy, fill in the form with the following example settings, and then click **Save**.

Policy Trigger	on Resource operation
Resource Block	Storage Volume
Operation	Storage Volume Interface.Provision
Moment	Pre Operation

7. On the Rules related list, click **New**, enter **Add Custom Tags for Storage** for the **Name**, accept the default settings, and then click **Submit**.
8. On the Policy Rule Actions related list, click **New**.
9. At Execute a Script, click **Create**, fill in the form with the following example settings, and then click **Submit**.

Action Name	CreateCustomTagsForStorage
Action Script Category	ScriptsForCustomTags
Action Script	AddCustomTagsForStorage

10. Click **Update** on the Policy Rule page.
11. Return to the Policy page and then click **Publish**.

Result
On the Cloud Admin Portal, launch the stack that is associated with the storage resource. You can then confirm that the tags are applied.

Create a cloud approval policy

A cloud approval policy specifies the users who must approve a specified cloud activity before the activity can proceed. Approvers can include the manager of the user making a request, a specified user or group, or users with a specified role. You can specify multiple approvers. Approvals occur in the order that you specify.

Before you begin
Optional: Create one or more cloud policy groups.

Role required: sn_cmp.cloud_governor or admin

About this task
In this example, a user requests a stack that triggers an approval policy.

1. On the Cloud User Portal, a user submits a request to provision a particular blueprint. The process of provisioning the blueprint is the trigger that causes the policy engine to apply an approval policy. A cloud approval policy specifies the users who must approve a specified cloud activity before the activity can proceed.
2. The policy engine determines that the request meets the condition specified in the rule for the policy. In this example, the condition evaluates to true whenever a particular blueprint is being provisioned.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Because the condition is met, the policy engine performs the action that is also specified in the rule. In this example, the action is to create an approval action for the manager.
4. While the approver (the manager) reviews the approval request, the user sees a "waiting for approval" status message on the Cloud User Portal.
5. After the manager approves, the blueprint is provisioned.

User waiting for approval

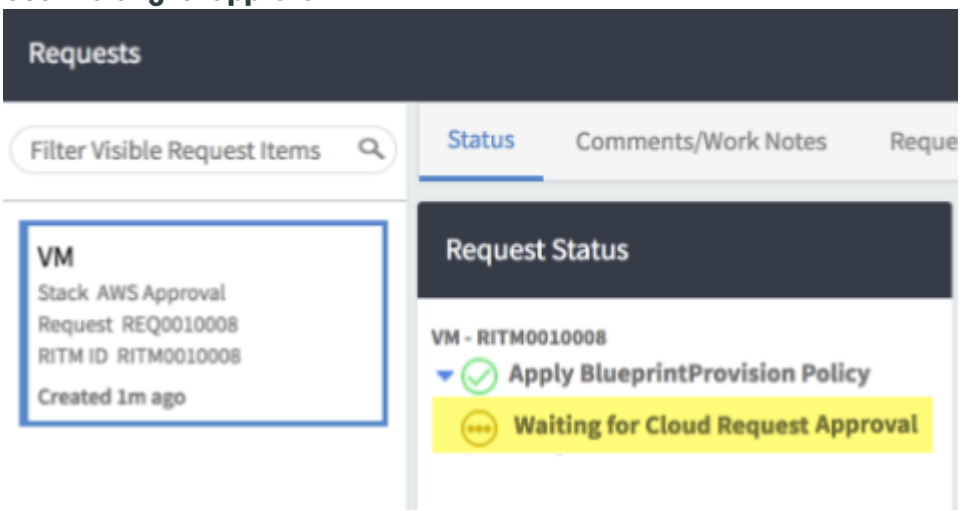

The targeted approval policies complement application-wide approval.

- **on Stack operation (approval)**: Triggered during any stack operation on the Cloud User Portal.
- **on Stack resource operation (approval)**: Triggered during any resource operation (start, stop, provision, and so on) on the Cloud User Portal.
- **on Task remediation**: Triggered when a user resubmits a failed request.

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Click **New** and then fill in the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example approval policy

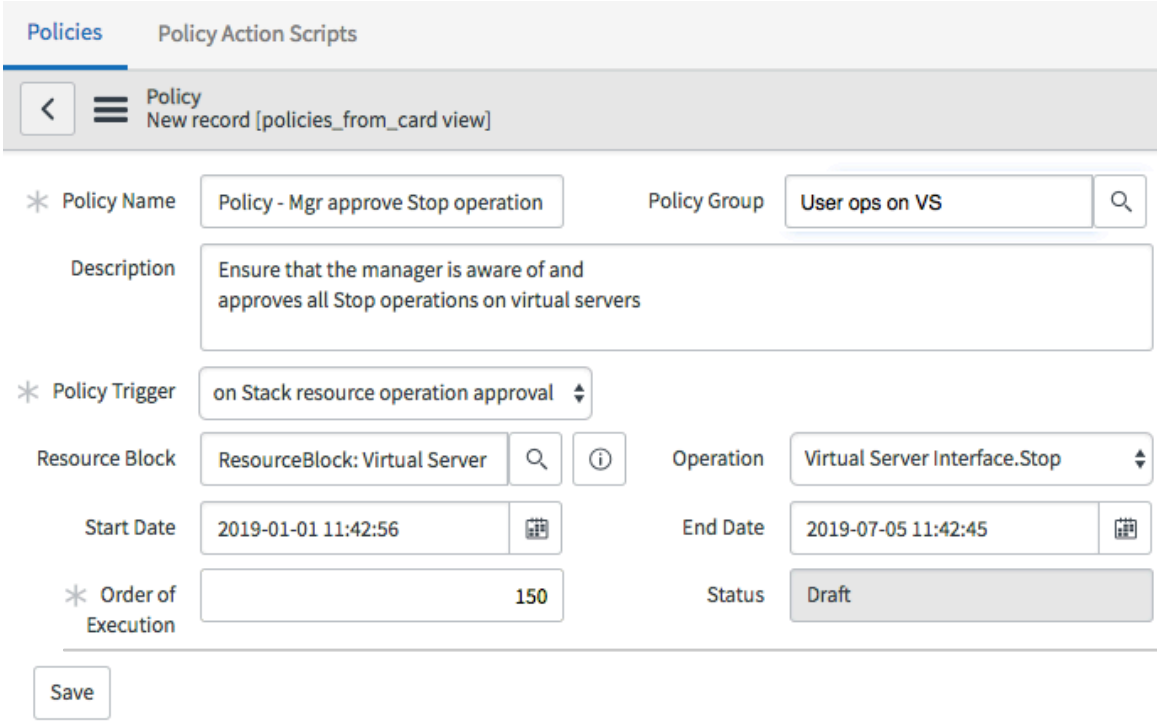

Field	Description
Policy name	Enter a descriptive name that includes the word `Policy`. Do not start the name with a number.
Description	Enter a description of the intent of the policy.
Policy group	Optional: Select a policy group. Each policy in the group is enforced.
Policy trigger	Select a trigger that specifies when the policy should be applied. The following triggers can start approval policies: ○ on Blueprint provision (approval) ○ on Stack operation (approval) ○ on Stack resource operation (approval) ○ on Task remediation
Blueprint [appears when the **on Blueprint provision (approval)** or **on Stack operation (approval)** trigger is selected]	Select the blueprint that the policy applies to. ○ If no blueprint is specified, then the policy applies for every blueprint. This setting can decrease performance. ○ You can assign multiple policies to a blueprint. Multiple policies running simultaneously, however, might decrease performance. Publish a policy only when it should be enforced.
Operation [appears when the **on Stack resource operation (approval)** trigger is selected]	Select the operation that the policy applies to. For example, a policy can apply to the *`Deprovision`* operation only or to all operations on the blueprint or catalog item.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Reviewers: Does this also appear for on Stack operation (approval)?	**Note:** If no operation is specified, then the policy applies for every operation. This condition can decrease performance. ○ **All**: Any operation executes. ○ **Start**: The resource starts. ○ **Stop**: The resource stops. ○ **Provision**: The resource is provisioned. ○ **Deprovision**: The resource is no longer available to users. ○ **Execute Script**: A script runs on the resource.
Start Date / End Date	Specify the start date when the policy should be enforced and the end date when the policy should no longer be enforced.
Order of Execution	Specify a number that represents the order in which the policy is applied. A policy with a lower number runs before a policy with a higher number. For example, a policy with **Order** of **100** runs before a policy with an **Order** of **200**.
Status	A new policy is in **Draft** state. Click **Publish** on the form header to enforce the policy. After a policy is published, you must set it to the **Draft** state to update it.

3. Right-click in the header and select **Save**.
4. When you are ready to enforce the policy, click **Publish**.

What to do next
Configure one or more cloud policy rules.

Create a cloud policy group

A cloud policy group is a container for related policies. Consider grouping policies that are often used together or should be considered together. Grouping policies can help you to apply policies consistently across your organization.

Before you begin
Role required: sn_cmp.cloud_governor or admin

Procedure
1. In the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Open a cloud policy and click the lookup icon for the **Policy Group**.
3. Select the group to add the policy to.
   - To add the policy to an existing group, select the group and then click **Submit**.
   - To create a policy group to add the policy to, click **New** on the Policy Groups list, specify a unique and meaningful **Name**, and then click **Submit**. The policy is added to the new group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Export or import a cloud policy

To back up, move, or restore a policy, you can export and import the policy as an update set. The update set includes rules, conditions, actions, scripts, and script categories.

Before you begin

Role required: sn_cmp.cloud_admin or admin.

Procedure

1. On the Cloud Admin Portal, navigate to **Govern** > **Policies**.
2. Perform the following action to export a policy:
   a. Click the export icon for the policy.

   The instance generates an update set.
3. To import a policy, install the update set as follows:
   a. On the instance to import to, navigate to **System Update Sets** > **Retrieved Update Sets**.
   b. Click the **Import Update Set from XML** related link.
   c. In the Import XML popup, click **Choose File**, select the file, and then click **Upload**.
   If a policy that you exported from the source environment exists in the target environment and both policies have the same sys_id, then the policy in the target environment is updated.

Trouble?

Note: If the policy that you exported from the source environment exists in the target environment, and the policies have different sys_ids, then an error occurs during the import process. You must export a different policy from the source environment or delete the policy with the same name in the target environment.

Related topics

System update sets

Pools and Filters for Cloud Provisioning

A resource pool is a query or script that filters a table. You configure a resource pool to limit the values that are available to users when they request a catalog item.

Example use of a pool

To limit a user to only the networks in a particular logical datacenter, use the resource pool called *`NetworkPool`* that runs against the Cloud Network [cmdb_ci_network] table. The resource pool uses a script to filter networks based on the datacenter that the network belongs to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

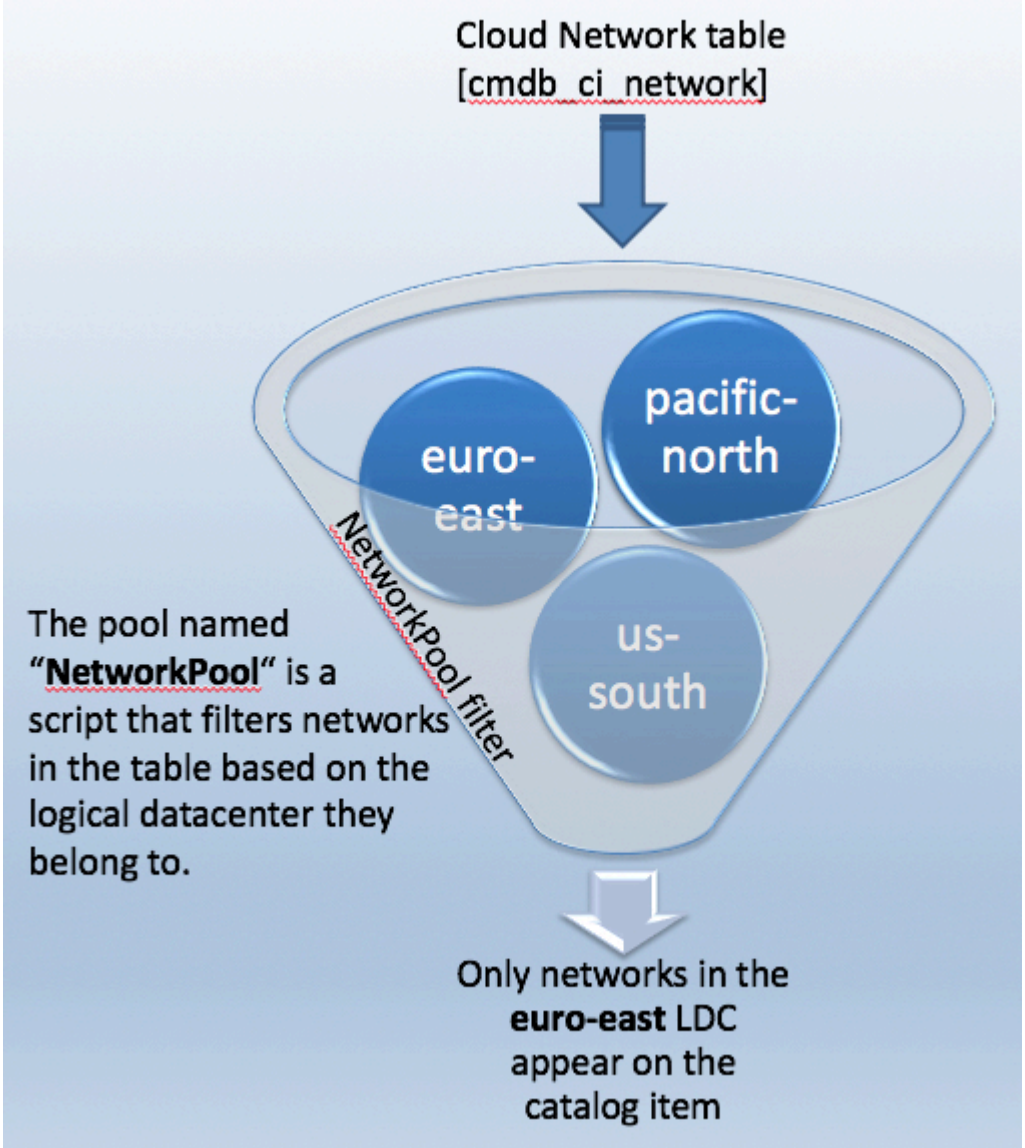

How pools work

The naming convention for pools is:

```
get<thisItem>By<condition>
```

Where the `condition` is the top-level entity that forms the relationship between the return values. For example, `getObjectIdByServiceAccount` filters all ObjectIDs that are hosted on the specified service account.
Filters query tables using only the specified condition. Any record that matches the condition is therefore returned.

Pools and blueprints

When you configure blueprints, you specify variables for form behavior. You can specify resource pools as the data source for each variable and select which resource pools to use. The variable then uses the filtered values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource Pool in the base system

Pool	Based on this table	Filter type	Filter name and description
AnsibleInventoryPool	Ansible Tower Inventory [sn_cfg_ansible_inventory]	Script	Filters the name of Ansible Tower Inventory.
ApplicationPool	Application	Query	Filters the names of applications.
ApplicationProfilePool	Application Profile [sn_cmp_application_profile]	Script	Filters the name of application profiles.
AvailabilityZonePool	Availability zone [cmdb_ci_availability_zone]	Script	Filters availability zones based on the logical datacenter they belong to. getNameByLDC: Filters availability zones based on the logical datacenter they belong to. The Terraform template for IBM requires the name of the availability zone as input for provisioning.
AzureDevOpsPipelineConfigPool	Config Installable [sn_cmp_cfg_installable]	Script	Filters the installable configs.
AzureDevOpsProjectPool	Azure DevOps Project [sn_itom_csc_cp_azure_devops_project]	Script	Filters projects based on config provider.
BusinessServicePool	Service [cmdb_ci_service]	Query	Filters the names of business services.
ChefServerPool	Chef Server [sn_cfg_chef_server]	Script	Filters chef server credentials by Chef server.
CloudAccountPool	Cloud Account [cmdb_ci_cmp_cloud_account]	Query	Filters the names of cloud accounts.
CloudKeyPairPool	[cmdb_ci_cloud_key_pair]	Script	getObjectIdByServiceAccount: Filters cloud keypairs based on the service account they belong to.
ComputeProfilePool	Compute Profile [sn_cmp_compute_profile]	Script	Filters compute profiles by the logical datacenter and the cloud account it belongs to.
ConfigMgmtPool		Script	Filters config management workload provider.
CostCenterPool	Cost Center [cmn_cost_center]	Script	Filters to list only the cost centers that the user belongs to.
DatastorePool	VMware vCenter datastore [cmdb_ci_vcenter_datastore]	Script	Filters datastores based on the logical datacenter they belong to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pool	Based on this table	Filter type	Filter name and description
HardwareTypePool	[cmdb_ci_compute_template]	Script	getObjectIdByAvailabilityZone: Filters hardware templates based on availability zone they belong to.
ImagePool	[cmdb_ci_os_template]	Script	getObjectIdByServiceAccount: Filters the resourceId of OS Image by service account. In IBM Cloud, OS Images are not specific to a datacenter, so they are hosted at the service account level. Terraform requires the resourceId as input for provisioning.
IPAddressPool	Cloud IP Address [cmdb_ci_cloud_ip_address]	Query	Filters the IP address.
LaunchConfigurationPool	Server Array Launch Configuration [cmdb_ci_sa_launch_config]	Query	Filters the configuration pool.
NetworkInterfacePool	Cloud Mgmt Network Interface [cmdb_ci_nic]	Script	Filters the network interface pool.
NetworkPool	Cloud Network [cmdb_ci_network] table	Script	• Filters networks based on the logical datacenter they belong to. • getObjectIdByLDC: Filters network IDs based on the logical datacenter they belong to.
OSProfilePool	OS Profile [sn_cmp_os_profile]	Script	Filters compute profiles by the logical datacenter and the cloud account it belongs to.
ProjectsPool	Project [pm_project]	Query	Filters the names of projects.
ResourceGroupPool	Resource group [cmdb_ci_resource_group]	Script	Filters resource group based on the logical datacenter they belong to.
ScheduleTimeZonePool	Choice [sys_choice]	Script	Filters scheduled time zones.
SecurityGroupPool	Compute Security Group [cmdb_ci_compute_security_group]	Script	• getByNetwork: Filters the security group by the network it belongs to. • getObjectIdByServiceAccount: Filters security group IDs based on the service account they belong to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pool	Based on this table	Filter type	Filter name and description
SecurityGroupProfilePool	Compute Security Group Profile [sn_cmp_security_grp_profile]	Query	Filters the names of security group profiles.
SSHKeyPool	CMP SSH Key Pair [sn_cmp_ssh_credentials]	Script	Filters user keys by user.
StorageAccountPool	Cloud Storage Account [cmdb_ci_cloud_storage_account]	Script	Filters cloud storage accounts based on the logical datacenter they belong to.
StorageVolumePool	Storage Volume [cmdb_ci_storage_volume]	Query	Filters the names of storage volumes.
SubnetPool	Cloud Subnet [cmdb_ci_cloud_subnet]	Script	• Filters the subnet by the network it belongs to. • getObjectIdByNetwork: Filters subnets based on the network they belong to.
UserGroupPool	Group [sys_user_group]	Script	Returns only the groups that the user belongs to.
VirtualMachinePool	[sn_cmp_resource_pool]	Script	getByAvailabilityZone: Filters virtual machine IDs based on the availability zone that they belong to.
VmFolderPool	VMware vCenter folder [cmdb_ci_vcenter_folder]	Script	Filters VM folders based on the logical datacenter they belong to.

Create a resource pool

Based on blueprint settings, resource pools control the values that a user sees in a catalog item when they request a resource. Only values that pass the pool filter or script appear as options on the catalog item request form.

Before you begin

Role required: sn_cmp.cloud_governor

Procedure

1. In the Cloud Admin Portal, navigate to **Manage** > **Resource Pools**.
2. Click **New**, enter a unique and descriptive **Resource Pool Name**, and then fill in the form.

Field	Description
Type	Select the type that determines whether the total count of available values remains the same (Static) or decreases by one each time a user selects the option (Diminishing).
Lookup Table	Select the table that contains the records that should appear as options in the catalog item for the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Lookup Field [Not necessary if using a script. This field is ignored if you configure a script filter.]	Select the table field that contains the values to display to the user. The display value for the lookup field is not used on the Cloud User Portal. Instead, use the **Lookup label field(s)**.
Lookup label fields [Not necessary if using a script. This field is ignored if you configure a script filter.]	Enter the lookup field label that should appear to the user.

3. Right-click in the header and select **Save**.
4. On the Resource Pool Filters related list, click **New**, enter a unique and descriptive **Filter Name**, and then fill in the form.

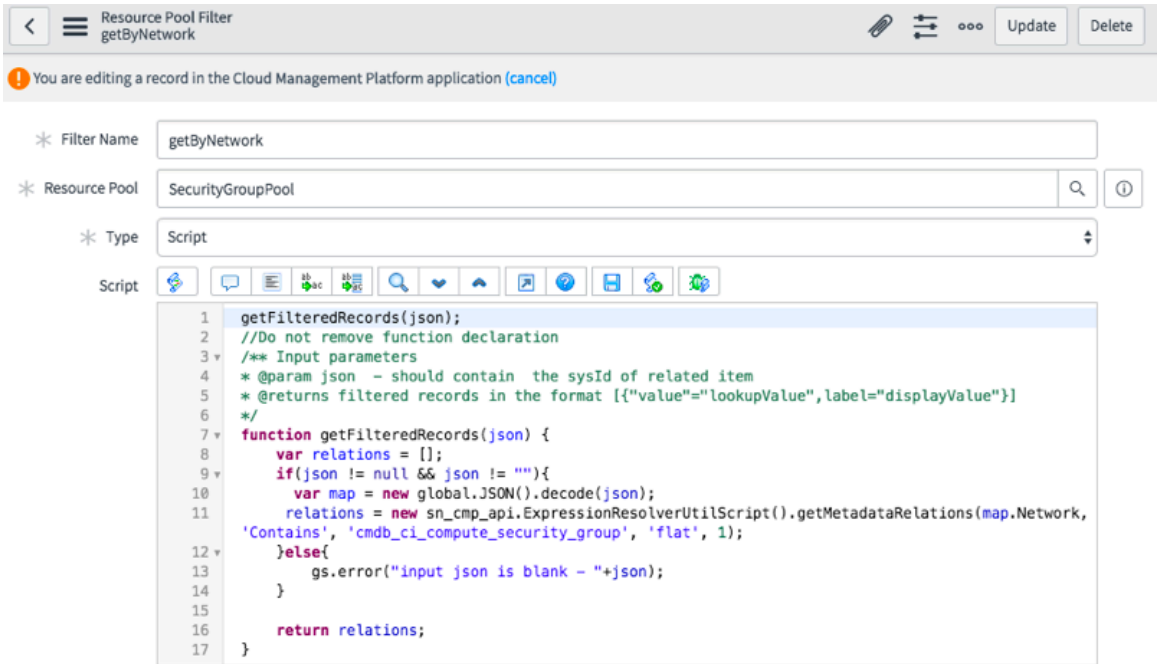

Field	Description
Resource Pool	Select the resource pool that the filter belongs to.
Type	Select the type of filter: ○ **Query**: Perform a query on records in the table that you specified for the **Lookup Table** based on the **Lookup Field** that you specified. ○ **Script**: Run a script to filter the records. Enter the code in the **Script** text box.

The resource pool filter is created.

5. Right-click the form header and select **Save**.
If you must return a filter value to the pool and that value is mapped to another table, create a record in the Resource Pool Filter Values related list.
6. On the Resource Pool Filter Values related list, click **New**, fill in the form, and then click **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Resource Pool Filter	Enter the name of the filter that the filter value belongs to.
Type	Enter the type of filter value: ○ **Query**: ○ **Metadata Rule**:
Metadata Rule Name [Appears if you select the Metadata Rule type]	Enter the metadata rule name. This field appears only when **Metadata Rule** is selected from the **Type** field.
Operator	Select either: ○ **AND** ○ **OR**
Field	Enter the field that is in the **Lookup Table**. For example, **CloudAccount**.
Value	Enter the value that you are passing back to the resource pool in the format `${field}`, where `field` is the value in the **Field** field. For example, **${CloudAccount}**.
Order	Enter an order number that determines when the value applies, relative to other resource pool filter values. Lower values are applied first.

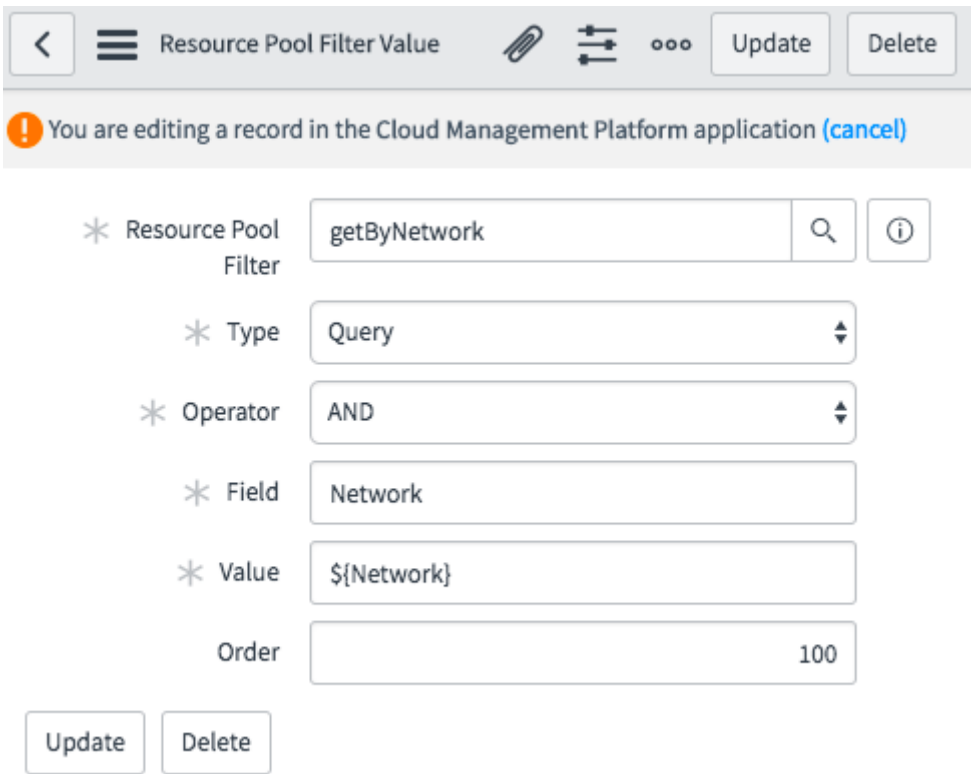

The resource pool filter is created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Bind a parameter to a resource pool

To make catalog ordering less error-prone, you can bind a parameter to an resource pool (a pool provided in the base system). Parameters that are based on an resource pool list only specified values from existing tables on the catalog order form in the Cloud User Portal.

Before you begin

Role required: sn.cmp.cloud_service_designer

About this task

The catalog order form lists only parameters values that have the **datasource** specified. For example, if you use the *`Resource Group Name`* parameter in your template, you can bind the **Resource Group Name** field to an existing SN pool (*`ResourceGroupPool`*) in the metadata section.

When you create a resource block using a cloud template, you can bind a parameter to an SN pool in the metadata section of the Azure Resource Manager (ARM) or CloudFormation (CF) template. All SN pools are listed in the procedure.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud Catalog Items**.
2. Select an existing template or click **New** to create a new cloud template.

   If you select an existing template, rename the template.

   In the next step, you modify the template definition for the appropriate template type (choose ARM or CloudFormation). You specify the SN pool filter as the *`datasource`* value.

SN pools to use as the datasource value

SN Pool	Filter
AvailabilityZonePool	getAllObjects: Gets all availability zones. getObjectsByLDC: Gets availability zone depending on LDC.
HardwareTypePool	
ImagePool	
LoadBalancerPool	All: Gets all load balancers.
NetworkPool	getAllObjects: Gets all networks. getObjectsByLDC: Gets networks depending on LDC.
ResourceGroupPool	getObjectsByLDC: Gets resource groups depending on LDC.
SecurityGroupPool	getObjectsByNetwork: Gets all security groups.
SecurityGroupProfilePool	All: Gets all security group profiles.
SSHKeyPool	UserKeys: Gets all SSH keys.
StorageAccountPool	getObjectsByLDC: Gets storage accounts depending on LDC.
StorageVolumePool	All: Gets all storage volumes.
SubnetPool	getObjectsByNetwork:Gets all subnets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SN Pool	Filter
VirtualMachinePool	getByAvailabilityZone: Gets all virtual machines.

3. Modify an ARM template definition:

   a. Add the attribute *SNC::Parameter::Metadata* in the *parameters* metadata section.

   b. For the *SNC::Parameter::Metadata* attribute, add *datasource* as the key and enter an SN pool filter for the value, such as `StorageAccountPool.getObjectsByLDC` as shown in this example.

```
"parameters": {
    "newStorageAccountName": {
      "type": "string",
      "metadata": {
        "description": "Unique DNS Name for the Storage
 Account where the Virtual Machine's disks will be placed.",
         "SNC::Parameter::Metadata": {
         "allowedPattern": "[0-9a-z]{1,11}",
         "ConstraintDescription": "must be a alphanumeric ",

 "datasource":"ServiceNow::Pools::StorageAccountPool.getObjec
tsByLDC",

        }
      }
    }
```

4. Modify a CloudFormation template definition:

   a. Add the attribute *SNC::Parameter::Metadata* and define the custom attribute name (*VpcId*) with a *datasource*.
   For the *datasource* value, enter an SN pool filter such as `NetworkPool.getObjectsByLDC` in this example.

   b. Define the *datasourceFilter* mapping attribute to bind the custom name (*VpcId*) to the actual attribute name (*Network*).

```
"Parameters" : {

     "VpcId" : {
     "Type" : "AWS::EC2::VPC::Id",
     "Description" : "VpcId of your existing Virtual Private
 Cloud (VPC)",
     "ConstraintDescription" : "must be the VPC Id of an
 existing Virtual Private Cloud."
   },

   "SubnetId" : {
     "Type" : "AWS::EC2::Subnet::Id",
     "Description" : "SubnetId of an existing subnet (for
 the primary network) in your Virtual Private Cloud (VPC)",
     "ConstraintDescription" : "must be an existing subnet
 in the selected Virtual Private Cloud."
   }
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
  },
  "Metadata": {
          "SNC::Parameter::Metadata": {
                "VpcId":{

 "datasource":"ServiceNow::Pools::NetworkPool.getObjectsByLD
C",
                    "datasourceFilter":{"Network":"VpcId"}
                },
                "SubnetId":{

 "datasource":"ServiceNow::Pools::SubnetPool.getObjectsByNetw
ork",
                    "datasourceFilter":{"Network":"VpcId"}
                }

            }
    }
```

5. Click **Submit**.
6. After the resource block, blueprint, and catalog items are created, the cloud service user sees only the list of values of the pool data for the parameter in the catalog order form.

Example resource pool that limits choices to cost center

You can use resource pools with blueprints to limit the choices on the cloud catalog request form.

Use case: Restrict cost center selection

In this example, the cost of the cloud asset is charged against the budget of the cost center of the user. The base-system **UserCostCenter** resource pool ensures that a user can select only resources in their cost center.

Assumptions

- The Cost Management [con.snc.cost_management] plugin is active.
- Cost centers are defined and users are associated with the cost centers.
- At least one blueprint is defined. This example uses a blueprint named **AWS Virtual Server**.
- You are assigned the sn_cmp_cloud_admin role and know JavaScript and JSON scripting.

Components

Review resource pool filter

1. On the Cloud Admin portal navigate to **Manage** > **Resource Pools**.
2. Open the **CostCenterPool** and review the related Resource Pool Filters.
   - **All** is a query filter that returns all cost centers in the table.
   - **UserCostCenter** is a script filter that looks up the cost center associated with the user who is ordering the item.

   Here is the script for the **UserCostCenter** filter:

```
getFilteredRecords();
//Do not remove function declaration
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
/**
* @returns filtered records in the format
 [{"value"="lookupValue",label="displayValue"}]
*/
function getFilteredRecords() {
 var filteredRecords = [];
 var userId = gs.getUserID();
 var userGr = new GlideRecord('sys_user');
 if (userGr.get(userId)){
  var costCenterId = userGr.getValue('cost_center');
  if (costCenterId){
   var costCenterGr = new
 GlideRecord('cmn_cost_center');
   if (costCenterGr.get(costCenterId)){
    var costCenter = {};
    costCenter.value = costCenterGr.getUniqueValue();
    costCenter.label = costCenterGr.getValue('name');
    filteredRecords.push(costCenter);
   }
  }
 }

 //force to string
 return new global.JSON().encode(filteredRecords);
}
```

Blueprint catalog form parameters

1. Navigate to **Design** > **Blueprints**, and then click the tile for the blueprint you want to open.
2. With the blueprint in **Draft** state, click the **Provision** operation tile on the **Catalog** > **Request Operation** tab.

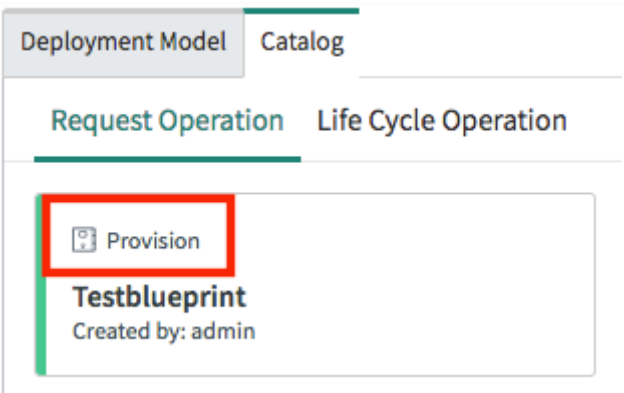

3. In the Variable sets related list, click the **General Info** variable set. By default, the CostCenter variable is in this variable set.
4. In the Cloud Variables related list on the Variable Set form, click the **CostCenter** variable.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

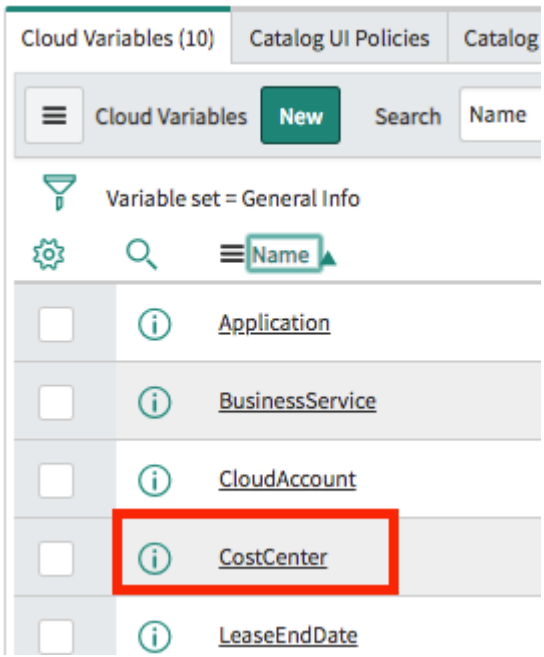

5. On the Cloud Variable form, click the **Type Specifications** tab.
6. Look at the **Pool** and **Pool Filter** fields that refer to the resource pool and filter.
   - *`CostCenterPool`* is the name of the resource pool.
   - *`UserCostCenter`* is the filter script that pulls in the cost center options for the user to select from.

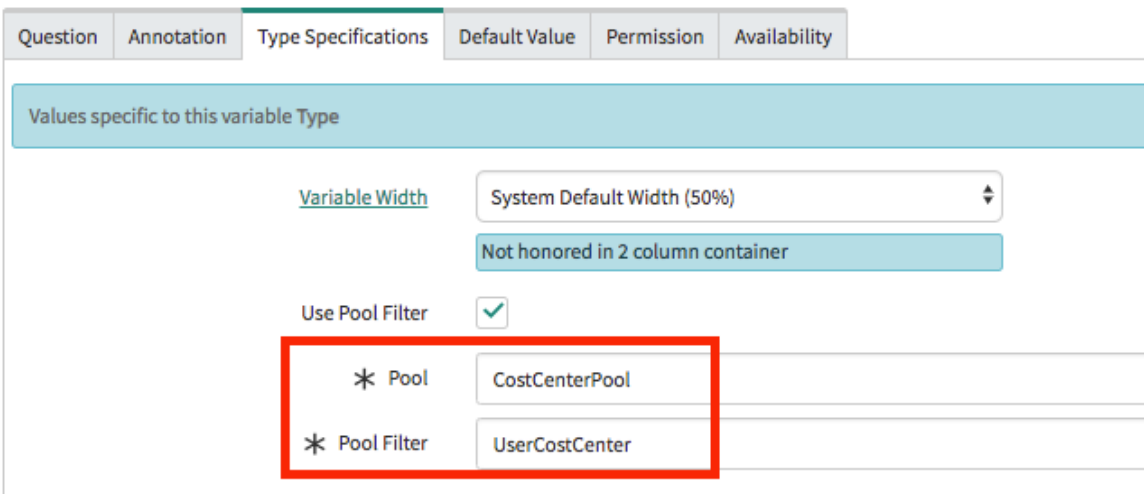

7. Set the blueprint to **Published**.

Cost center user

Identify a user who is a member of a cost center and who has access to the Cloud User Portal.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

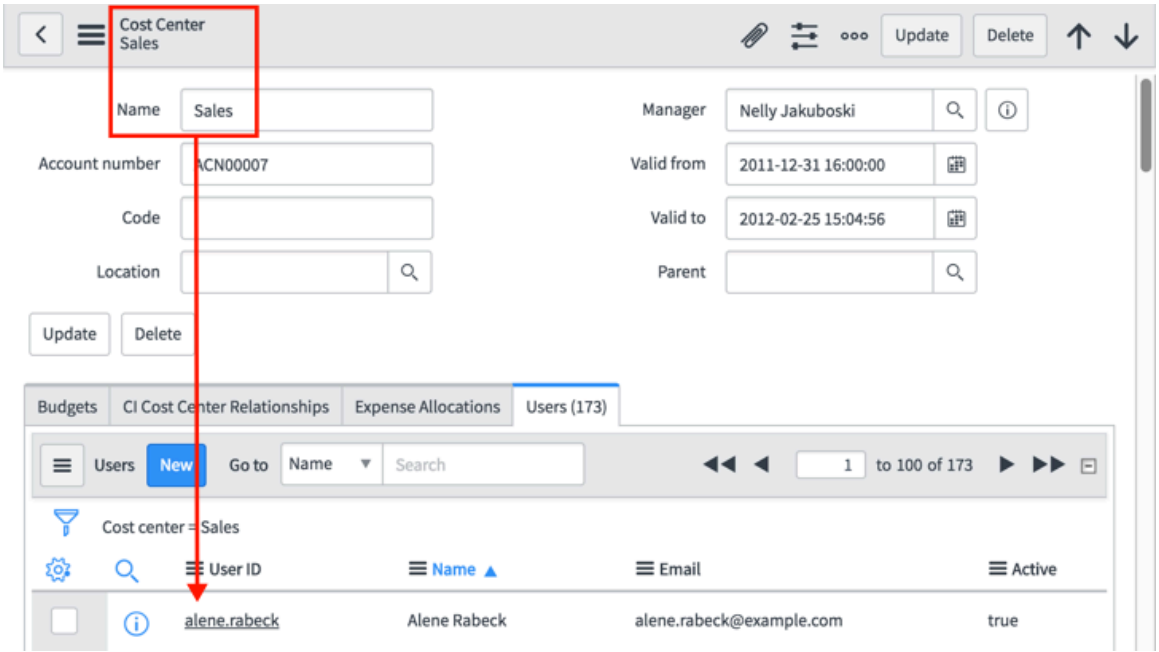

Testing the resource pool filter

After reviewing the components that comprise this use case, test the cloud catalog item to verify that users can select only their cost center.

1. Impersonate the user, **Alene Rabeck** in this example.
2. On the Cloud User Portal, click **Launch a Stack**, and then select the cloud catalog item (**AWS Virtual Server** in this example).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Review the selections in the **Cost Center** list.

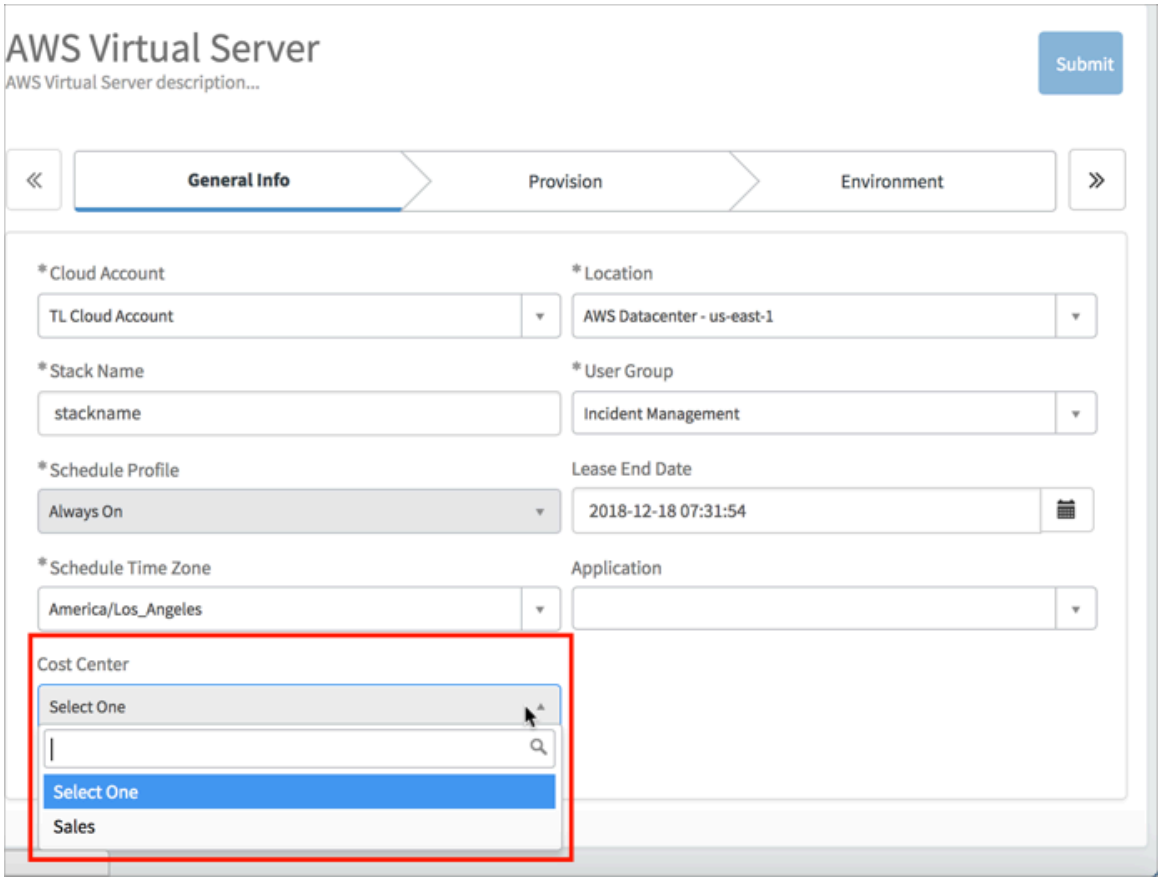

With the *CostCenterPool::UserCostCenter* datasource value for this catalog item, the only option for the **Cost Center** is the cost center the user is a member of.

Changing the resource pool filter

Test that the resource pool filter is controlling the behavior of the **Cost Center** field by changing it and viewing the results.

1. On the Cloud Admin Portal, navigate to **Design** > **Blueprints** and then click **AWS Virtual Server**.
2. Click the **Provision** operation tile.
3. In the Variable sets related list, click the **General Info** variable set. By default, the CostCenter variable is in this variable set.
4. In the Cloud Variables related list on the Variable Set form, click the **CostCenter** variable.
5. On the Cloud Variable form, click the **Type Specifications** tab.
6. Edit the **Pool filter** field to change the filter from *UserCostCenter* to *All*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

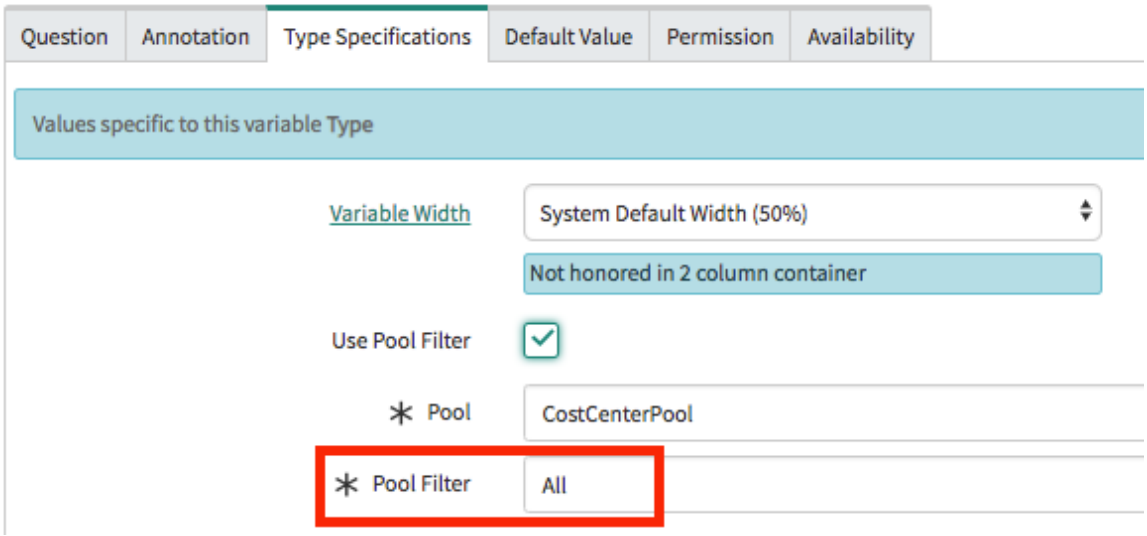

7. Click **Update**, then click**Publish.**.
8. Impersonate the user, **Alene Rabeck** in this example.
9. On the Cloud User Portal, launch a stack, and then select **AWS Virtual Server**.
10. Verify that all cost centers are listed.

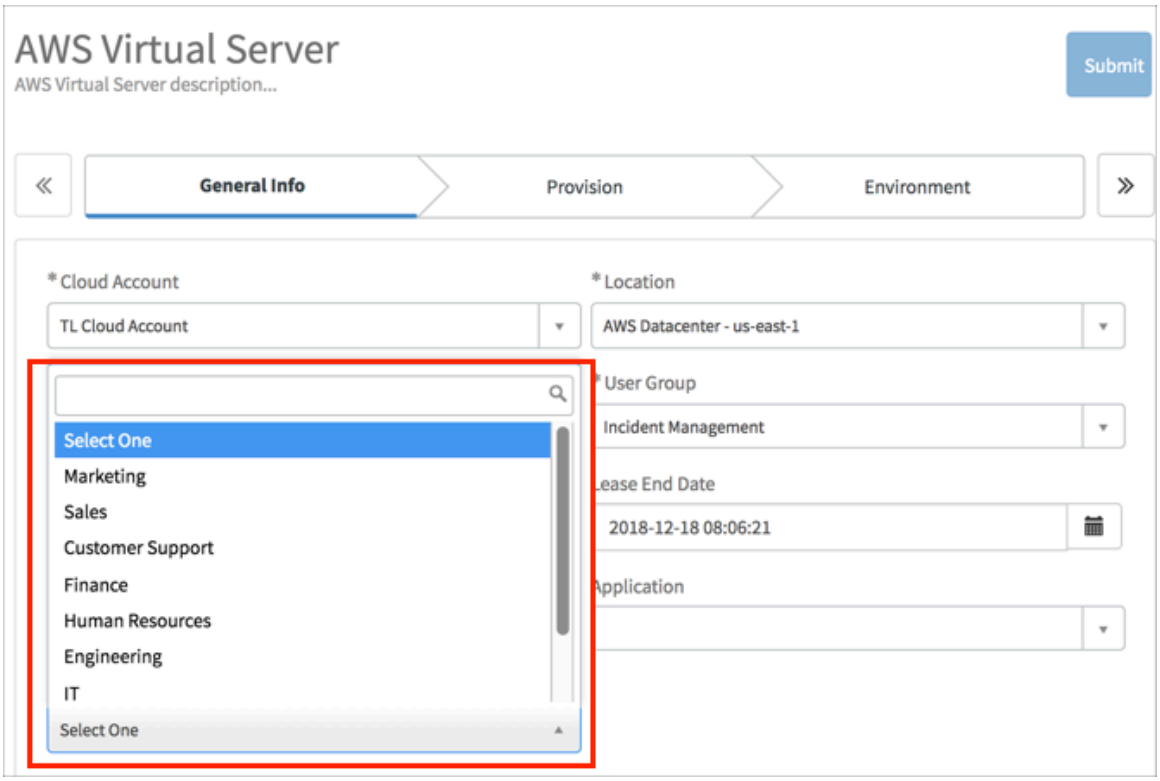

Quotas and resource order controls

Quotas are limitations or requirements for groups and users on CIs or cloud resources. Use quotas to prevent wasteful resource usage in your cloud environment.

You apply quotas to a resource, such as a virtual server, datacenter, or a generic cloud resource, and to users of a particular group. You can set a per-user limit, and a total limit for the group for

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

the resource. For example, a quota might limit the number of virtual servers allowed for users in the **Database** group.

After this kind of quota is set, users who try to provision additional virtual servers will receive an error message and they will not be able to proceed with provisioning virtual servers until the resource count is under the quota limit. If a user belongs to more than one group, they receive the total quota limit from both groups combined.

Note: You can also set per-datacenter limitations on resources by using capacity limits in cloud accounts. See Set capacity limits on user requests for resources.

Resource block types, quota types, and attributes

You can assign quotas to users or groups for the following resource block types:

- Virtual Servers
- Storage Volume Size
- Network
- Stacks

Quota types include:

Sum

The sum of all attributes for a resource, such as the sum of all virtual CPUs.

Count

The total number of attributes for a resource, such as the number of virtual machines.

Create a cloud quota definition

A cloud quota definition specifies the limitations on a resource block. You can use templates or manually create different types of configurations for each quota definition.

Before you begin

Role required: sn_cmp.cloud_governor

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Quota**.
2. Click **New**.
3. On the form, fill in the fields.

Quota Definition form

Field	Description
Quota Definition	A unique identifier for the quota definition. Typically it is the `resource name_count` or `resource name_size`. For example: Network_Count, VM_Count, or StorageVolume_Size.
Quota Type	Type of quota restriction to set: ○ Count ○ Size
Resource Name	The Resource Block type to apply the quota to. Select from the lookup list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	For example, select Cloud Resource for any template-based cloud resource. For more information, see Resource blocks in Cloud Provisioning and Governance.
Calculation Unit	The unit this quota is calculated in: ○ GB ○ Single
CI Class	The CMDB CI class to be selected from the lookup list.
Attribute	The column name from the underlying CI type on the resource block that holds the value to restrict with the quota. For example, if your quota limits the CPU count on a virtual server, the table of the underlying CI type is cmdb_ci_vm_instance. The **cpus** column holds that CPU data. So you would enter `cpus`. You can also enter `instance`.
Template [optional]	The table name of a resource type, for example Hardware Type [cmdb_ci_compute_template] or Image [cmdb_ci_os_template].
Template Attribute [optional]	A column from the table you specified in the Template field. This it the attribute on which the quota is set. For example, if you selected the [cmdb_ci_compute_template] template and you want to put a quota on the virtual CPUs, you would enter `vcpus`. That column stores the virtual CPU count for all records in the [cmdb_ci_compute_template] table.

4. Right-click the header and select **Save**.
   The Quota Configurations, Default Quotas, and Cloud Template Resources related lists appear.
5. In the Quota Configurations related list, click **New**.
6. On the form, fill in the fields:

Quota Configuration form

Field	Description
Group Name	A unique name for the configuration group.
Group Maximum Limit	The maximum limit for the group's quota.
Quota Definition Type	The quota definition name.
Per User Limit	The maximum quota limit for each user in the group.

7. Click **Submit**.
8. **Optional:** You can specify optional default quotas that will be set on the resource block by clicking **New** in the Default Quotas related list, and then filling in the form.

Default Quotas form

Field	Description
Group Maximum Limit	The maximum quota limit for the group.
Quota Definition	The quota definition for which you are setting the default values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Per User Limit	The quota limit applicable to each user in this group.

9. Click **Submit**.
10. Map catalog item to which the quota you defined applies by clicking **New** in the Cloud Template Resources related list.

    Note: If you do not map quota to a catalog, quota is not calculated for that catalog.

11. On the form, fill in the fields.

Cloud Template Resources form

Field	Description
Domain	The domain to which this quota applies.
Quota Definition	The name of the quota definition you are applying.
Catalog Item	The catalog item or cloud resource you are mapping to the quota.

12. Click **Submit**.

Result
You have specified the cloud quota definitions and quota limits for users and user groups.

Related topics

Resource blocks in Cloud Provisioning and Governance

Resource order controls for template-based catalog items

Use resource order controls to perform quota checks for template-based catalog items using quota definitions and policies. Quota limits enable you to control provisioning or ordering resources for users and groups. Configure policy actions to trigger notifications, an approval subflow, or both.

Resource Order Controls

Resource order controls enable you to control the quota limit of resources provisioned through template-based catalogs. Quota checks for cloud template-based catalog items are available beginning with the Brazil release.

You can now map quota limits to template-based catalogs. You can trigger notifications or an approval subflow when the resource limit or quota definitions you specify for a user or user group is exceeded. For more information, see Resource Quota.

Use resource controls to perform the following actions:

- Calculate quota and capacity for template-based catalog items.
- Set limits for stack quota during provisioning for the cloud user.
- Create a policy to be triggered when a quota limit is exceeded.
- Add a policy notification and request for approval when a quota limit is exceeded.
- Show user and user group quotas for template-based catalogs in the Cloud User portal.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How resource order controls work

Resource order control applies only when the quota has exceeded. If the quota exceeds during catalog provisioning for the first time after resource order controls are configured, the cloud resources continue to be provisioned. However, if the cloud user tries to provision a resource that breaches the configured quota limits again, the Resource order limit exceeded policy action set up for the Resource Order Control policy is triggered.

Note: If quota is not mapped to a template-based catalog, quota calculations for that catalog are not performed, even if users exceed their quota limits.

For example, say you specify a quota limit of 20 GB for a storage volume for a cloud user and the user provisions a 40 GB storage volume. The first time, the storage volume is provisioned even though it exceeds the quota limit. However, provisioning will fail the next time the user attempts to provision a cloud resource over 20 GB.

Stack count quota

The stack count quota introduces a default base system quota for template-based cloud catalog items. The Stack count quota counts all active stacks (identified by the *`Stack`* CI class) provisioned by users and user groups. The stack count quota applies to all template-based cloud resources and you do not need to map catalog items to quota definitions. Once the stack quota is enabled, whenever a user from the group mapped to the stack quota provisions a cloud resource, the quota is calculated. When users from the mapped group try to provision resources beyond the limit defined in the quota definition for stacks, the Resource Order Control policy and *`on Resource Limit exceeded`* policy action is triggered.

Note: The base system stack count quota does not have predefined default quotas to preserve stacks that might have been created prior to upgrading to Brazil. You must specify default quota limits for users and user groups to enable the Stack count quota.

Resource Order Control Policy

The Resource Order Control policy is a base system policy that is in the Draft state by default. You can customize the default policy to meet your requirements and publish it, or create a new policy and define quotas, rules and actions to enable resource order controls. The 'on Resource Limits exceeded' policy action triggers the Resource Order Control policy when any resource exceeds the defined quota limit. You can determine whether to initiate an approval subflow, a notification, or a combination of both when the resource limit exceeded policy is triggered. For more information, see Create an action for an 'on Resource Limit exceeded' policy.

Note: You cannot provision catalog items once the user or user group quota is exceeded without the policy being published because the existing quota behavior disables further provisioning and does not process the request further.

Next Steps

Set up resource order controls

Create quota definitions and policies mapped to cloud-template based catalog items. Use policy rules and actions to trigger an approval subflow or notification when the quota limit defined for the group or user exceeds.

Before you begin
You must have the following items:

- One or more cloud policy groups.
- Cloud template-based catalog items.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Ensure that your CMDB is updated through event-driven discovery for precise quota enforcement.
- Ensure that you map quota to a catalog in the Cloud Template Resources related list, when you configure quota definitions.

Role required: sn_cmp.cloud_governor

About this task

Consider the following points before you begin setting up resource order controls:

- Quota is calculated only for the resource blocks that are bound to a CI class on the CMDB. If you create custom resource blocks that do not map to CI class on the CMDB, no quota calculation is performed for those resource types.
- If you have configured both Resource Order Control policy and a blueprint provision approval policy, the blueprint provisioning policy takes precedence over the Resource Order Control approval policy.

Procedure

1. Define the quota or use a base system quota definition for a resource block.

   a. In the Cloud Admin Portal, navigate to **Govern** > **Quota** > **Quota Definition**.

   b. Specify quota definitions.

   The Resource Order Control base system policy has the default ServiceNow approval policy rule configured by default. If you want to trigger other actions such as notifications or a custom approval subflow when the resource limit is exceeded, you have two options. You can either modify the base system policy action or create a custom policy for your instance.

   - Modify a base system quota definition by specifying the default quota values per user and user group and the user groups to which the defined quota applies.

     The following predefined base system quotas have default quota definitions:

     - *VM_Count*
     - *vCPUs Count*
     - *StorageVolume_Size*
     - *Stack_Count*
     - *Network_Count*

     Important: *Stack_Count* is applicable only to template-based cloud resources.

   - Create a quota definition.

     For more information see, Create a cloud quota definition.

2. Map the quota to a template-based cloud catalog item.

   Important: If quota is not mapped to template-based catalogs when you configure quota definitions, quota calculations for those catalogs are not performed, even if user-define quotas exceed. For more information, see Create a cloud quota definition

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the Cloud Template Resources related list, click **New**.

b. In the **Catalog Item** field, select the catalog item you want to map this quota definition to.

c. Click **Submit**.

3. Configure a policy rule.

a. Navigate to **Govern** > **Policies**.

b. Open the base system Resource Order Control policy.
By default, the base system resource order control policy is in Draft state.

c. Navigate to the Rules related list and either modify a base system rule or configure a new rule if you have numerous customizations.

- Select the base system *approveresourcelimit* rule and modify it to include a policy action triggering a custom subflow or a notification.

  Note: By default, only the ServiceNow approval action is configured for the *approveresourcelimit* rule.

- Click **New** and configure a new policy rule.

  For more information, see Configure a cloud policy rule.

4. Configure a policy rule action.
   - Select the base system approve policy rule action in the Policy Rule Actions related list to include a notification or a custom approval subflow.
   - Create a cloud policy rule action in the Policy Rule Actions related list, to trigger custom approval subflows or notifications, if you created a policy rule in the previous step.

     For more information, see Create an action for 'on Resource Limit exceeded' policy.

   Note: In the **Policy Group** field, ensure that the policy group you want this policy to apply to is selected.

5. Publish the policy to activate resource order controls by clicking **Update** and then **Publish**.

Result

Based on the Policy Rule action you configured for the Resource order control policy, either an email notification or an approval subflow or both will be triggered if the specified quota is exceeded.
If you have configured an approval policy, the requester receives a notification when the specified quota is exceeded. The notification states that the approval request has been sent to a designated approver. Typically, the assigned approver designed the template used to create the cloud resource and has the *sn_cmp.cloud_service_designer* role.

Use the Cloud Operations Dashboard

The Cloud Operations Dashboard breaks down cloud service requests from your end users and cloud stacks that you offer in the Cloud User Portal.

Before you begin

Role required: sn_cmp.cloud_operator or sn_cmp.cloud_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Operate** > **Cloud Operations Dashboard**.

   The Cloud Operation Dashboard appears, showing you Cloud Service Requests and Stacks. Requests are also broken down by requester in the bar chart below. Stacks are broken down by datacenters.

An example Cloud Operations Dashboard

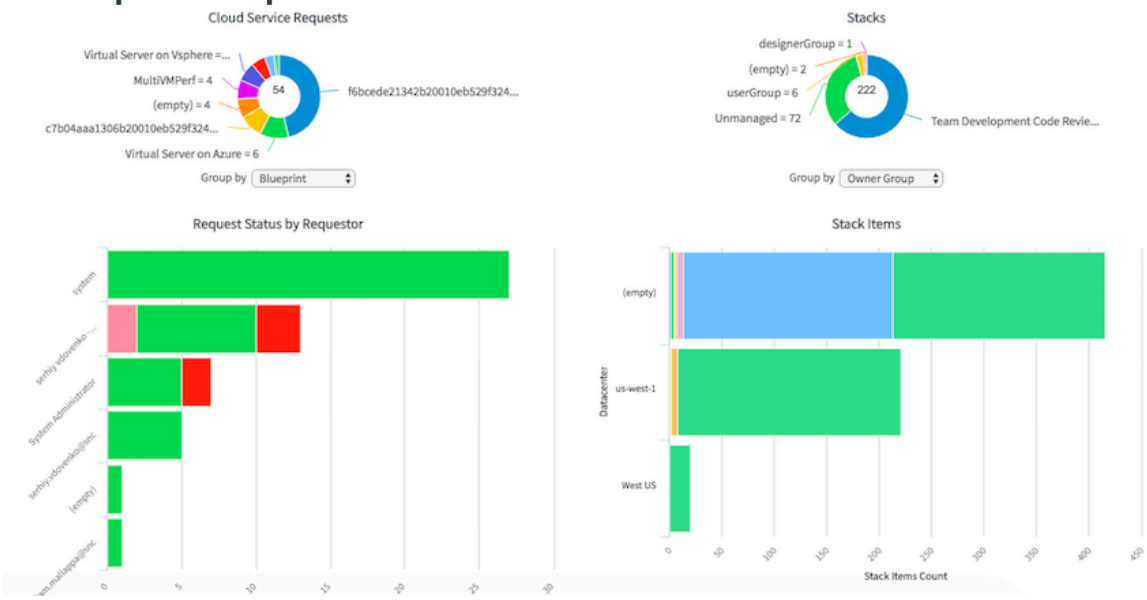

2. Do any of the following to obtain the tag data you want on the report.

Goal	Do this
See data grouped by another attribute	Select a value in the **Group by** choice list for either chart.
See updated data	Point to the top of any of the charts until the refresh icon ( ) appears, and then click the icon.
Save an image of a chart	Point to any of the charts until the options icon ( ) appears, and then select **Save as PNG** or **Save as JPEG**.

Use Cloud Root Cause Analysis reports

Root Cause Analysis reports help you troubleshoot issues with Cloud Provisioning and Governance. Use the reports to view details about the Cloud Orchestration Trail and the Cloud API Trail.

Before you begin

Role required: sn_cmp.cloud_operator or sn_cmp.cloud_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Operate** > **Root Cause Analysis Dashboard**.

   The Root Cause Analysis dashboard appears, with the **Cloud Root Cause Analysis Dashboard** displayed by default. This gives you a graphical summary of the Cloud Orchestration Trail with trail records under the reports. These charts group records by different attributes:

   - **Trails**: groups data by the template, which is the component, such as a resource block or blueprint, on which a Cloud Provisioning and Governance action was applied.
   - **Stages**: groups data by the stage, which identifies when the transaction took place.
   - **Components**: groups data by the component in the system that was involved in the transaction.
   - **Steps**: groups data by the step, which is the descriptive step associated with a stage and component.

   See The Cloud Orchestration Trail for a list of possible steps.

   Requests are also broken down by requester in the bar chart below. Stacks are broken down by datacenters.

   The Cloud Root Cause Analysis dashboard

   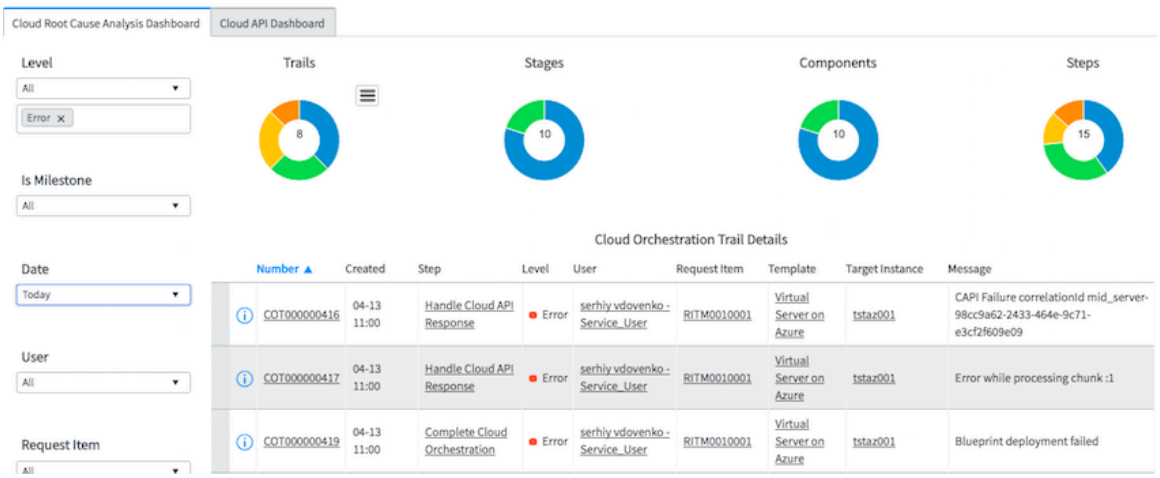

2. Select an option to obtain the tag data that you want on the report.

Goal	Do this
Filter the data	Select from the options, such as the **Level**, **Date**, and the **Request Item**.
See updated data	Point to the top of any of the charts until the refresh icon ( ) appears, and then click the icon.
Save an image of a chart	Point to any of the charts until the options icon ( ) appears, and then select **Save as PNG** or **Save as JPEG**.

3. Click any Cloud Orchestration API **Number** to view that specific record.
   See The Cloud Orchestration Trail for field descriptions.
4. To view reports about the Cloud API Trail, click the **Cloud API Dashboard** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This dashboard shows you several reports for the Cloud API Trail grouped by provider, status, and connectors. You can also see information the methods that were invoked and the API invocations today.

Note:

Starting with Brazil the Cloud API dashboard will be hidden and no longer activated on new instances but will continue to be supported. The Cloud Services Catalog application provides this functionality now. The application is being prepared for future deprecation.

The Cloud API dashboard

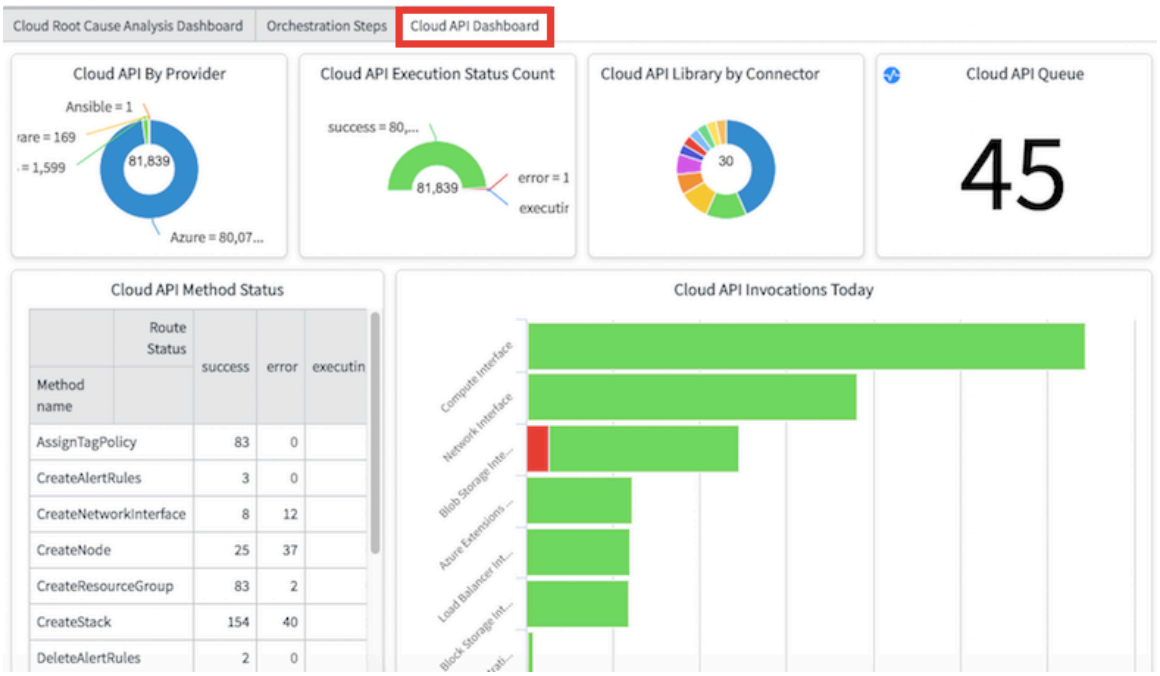

5. Click any the data in the graphs to open the matching Cloud API Trail records. See The Cloud API Trail for descriptions of the Cloud API Trail form fields.

The Cloud API Trail

The Cloud API Trail is an activity log for all activity that uses the Cloud API and goes through the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud API Trail contents

CAPI trail form

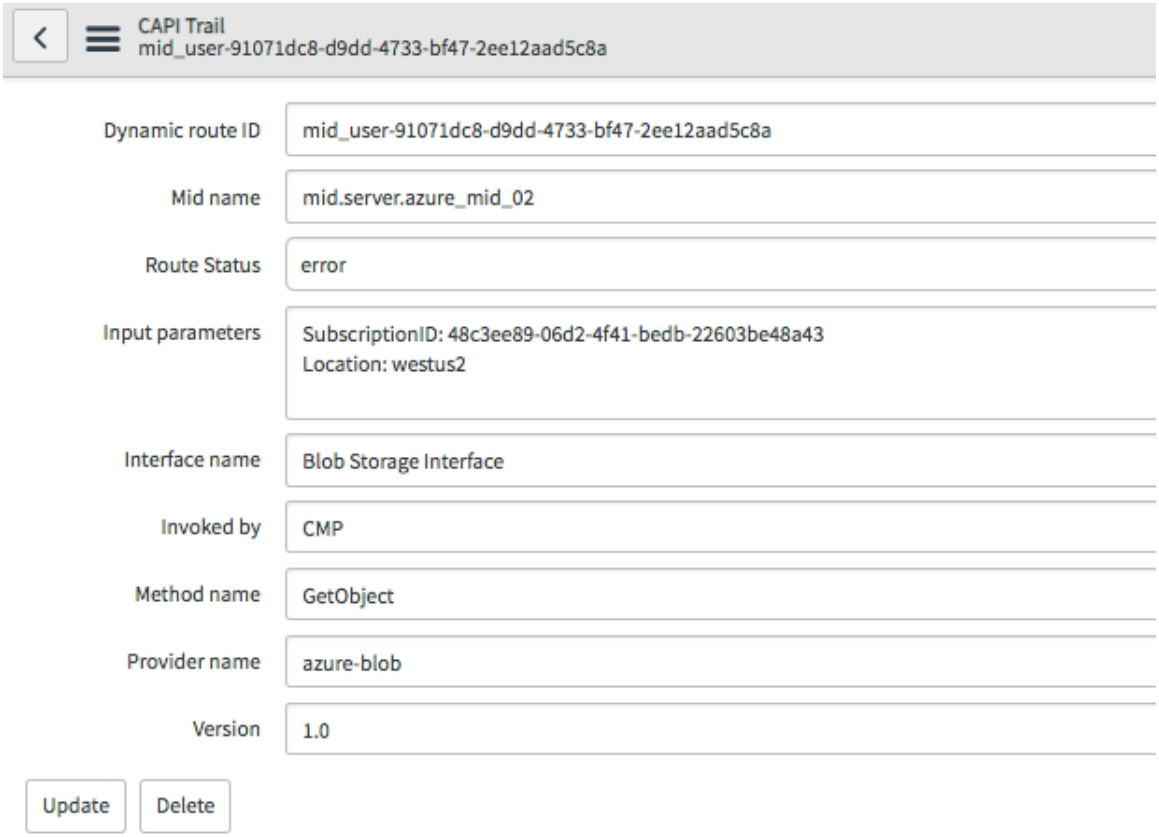

Field	Description
Dynamic route ID	An auto-generated ID number for the entry.
Mid name	The name of the MID Server through which the Discovery was performed.
Route status	Whether or not the Discovery operation run by the API was successful. Possible values are: • **success** • **error** • **executing**
Input parameters	The input parameter that generated the API trail record. This value is usually the datacenter in which the Discovery was run.
Interface name	
Invoked by	This value is always **CMP** when running Discovery.
Method name	The interface operation from the Cloud API that processed this record.
Provider name	The cloud provider.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Version	The version specified in the Cloud API.

CAPI Trail Logs

The CAPI Trail Logs related list provides more details about Cloud API trail entry. The following types of log keys are available:

Log key	Description
route_data	Information about the Cloud API calls.
dynamic_route	Information about the actual route the data took, including URIs.
route_result	The payload received by the instance, or a description of the result of the data transfer. The payload
chunk_number	The number of data chunks that the instance received.
route_status	Whether the route connection and payload transfer was successful.
route_error	The error that occurred. For example, the error `Failed to list loadbalancer Failed : HTTP error code : 403` means that your credentials were incorrect and Discovery could not access the cloud resource.
error_detail	More details about the error, including the Cloud Provisioning and Governance API and connector that was used in the attempted Discovery, and the errors that the cloud provider threw.

An example of a route_error is as follows:

```
Failed to list loadbalancer Failed : HTTP error code : 403
```

An example of the error_detail entry for the same error is as follows:

```
com.snc.cmp.connector.cloud.loadbalancer.component.LoadBalancerE
xception: Failed to list loadbalancer Failed : HTTP error code :
 403
       at
 com.snc.cmp.connector.cloud.loadbalancer.customizer.impl.AWSLoa
dBalancerCustomizer.listLoadBalancers(AWSLoadBalancerCustomizer.
java:56)
 at
 com.snc.cmp.connector.cloud.loadbalancer.component.LoadBalancer
Producer.process(LoadBalancerProducer.java:46)
 at
 org.apache.camel.util.AsyncProcessorConverterHelper$ProcessorTo
AsyncProcessorBridge.process(AsyncProcessorConverterHelper.java:
61)
       at
 org.apache.camel.processor.SendProcessor.process(SendProcessor.
java:145)
       ...
```

These two entries indicate that the credentials were incorrect, and Discovery could not access the cloud resource. The load balancer interface throws the first error because the load balancer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

device is the first device that allows access to the cloud resource. The `org.apache.camel` errors indicated routing errors on the Amazon Web Services side.

The corresponding error on the instance side is captured in the Cloud Orchestration Trail.

Related topics

The Cloud Orchestration Trail

Open the Cloud API Trail

Open the Cloud API Trail to debug and troubleshoot issues like a failed policy or failed Discovery of cloud resources.

Before you begin
Role required: sn_cmp.cloud_operator or sn_cmp.cloud_admin

Procedure

1. In the Cloud Admin Portal, navigate to **Operate** > **Trails**.
2. On the **Cloud Api Trail** tab, filter and sort the list of Cloud API Trail records as needed.
   If you are looking for something like a failed Discovery, filter the list so the **Route Status** column shows only entries with **error**.
3. Click a link in the Created column to open the Cloud API Trail record.
4. In the CAPI Trail Logs related list, open the log record that displays the information you want.
   For example, open **route_error** or **error_detail** to debug a failed operation.

Cloud Provisioning and Governance dashboards and reports

Dashboards enable cloud admins and cloud users to view reports like cloud billing data and cloud tag usage.

Follow this short video to learn more about Cloud Provisioning and Governance dashboards.

Important: Starting with the Brazil release, the Billing dashboard is no longer available if you have downloaded the ServiceNow Store Cloud Cost Management app. The following changes occur:

- You are redirected to the Cloud Cost Management home page by default.
- The View Dashboard widget in the Cloud User portal is replaced by the View Resources widget.
- The Current Month Spend widget and the Budget widget on the Cloud User portal do not show any data if Cloud Cost Management is activated on the instance.

View cloud costs on the Billing dashboard

The Billing dashboard provides rich summary information on cloud usage, cost trends, and cost aggregates.

Before you begin
Role required: sn_cmp.cloud_admin

Important: You can only access the **Billing Dashboard**, if you have not downloaded and activated the ServiceNow Store Cloud Cost Management app. If you have activated the Cloud Cost Management app, you can only navigate to the Billing Dashboard, if you are using Cloud Provisioning and Governance on a domain separated instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Analyze** > **Billing** > **Billing Dashboard**.

   The Billing dashboard shows the **Cloud Cost** tab by default. This tab displays these two graphs:

   - **Cost Trend**: A line graph showing daily costs over time for the selected **Usage Date**.
   - **Cost Aggregate**: Total costs for each tag value for the selected **Group by** tag.

   Cloud Cost tab on the Billing dashboard

   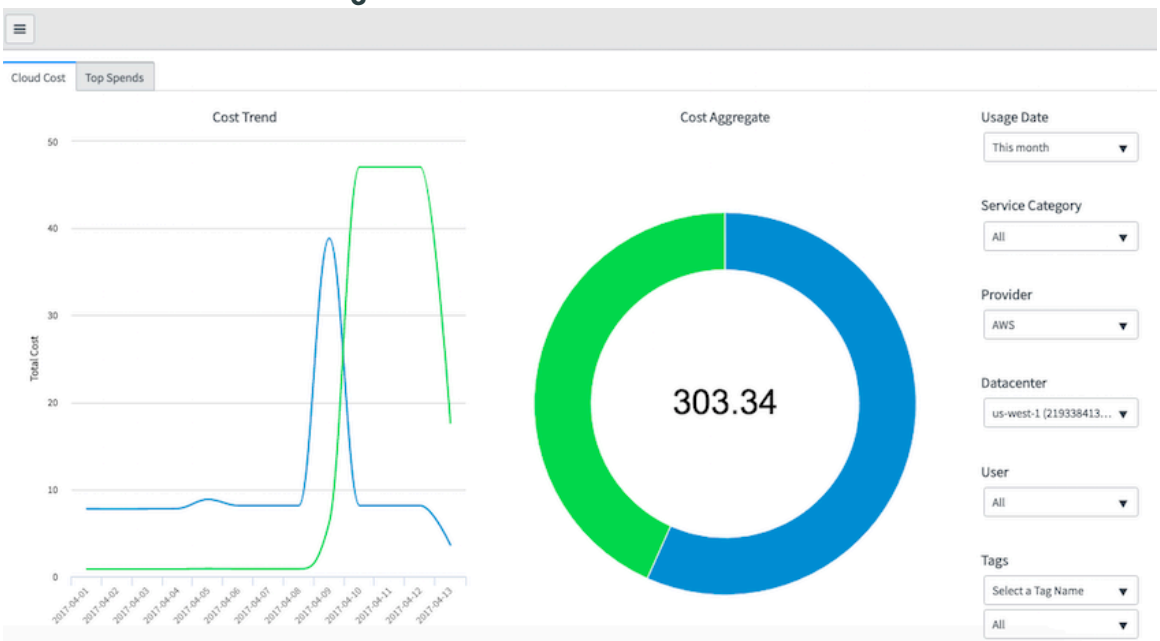

2. Perform any of the following actions to obtain the tag data you want on the report.

Goal	Action
Update the billing period	Select a new option from the Usage Data choice list.
See data grouped by another tag	Select a value in the **Group by** choice list un der either chart.
Filter by tag values	Select a tag value from the list.
Save an image of a chart	Point to either donut chart until the options icon ( ) appears, and then select **Save as PNG** or **Save as JPEG**.

3. To view the areas with the highest level of spending, click the **Top Spends** tab.
   Bar charts showing the total cost for each tag display by default. You can further filter the data by **Service Category**, **Provider**, and **Datacenter**.
4. Click any part of the charts to see cost records for the selected criteria.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The list view of the cost records appear. You can view information such as usage quantity and specific cost per usage date.

Example cost records

Usage date	Service Category	Resource	Usage profile	Usage quantity	Cost
2017-04-08	Compute	flow-vm	Data Transfer ($0 / GB)	0.020323	0
2017-04-08	Compute	flow-vm	Data Transfer ($0.08199078 / GB)	0.02444	0.00200385
2017-04-08	Compute	flow-vm	Hourly ($0.06172149 / Hour)	24	1.48131575

View tagged resources on the Tag dashboard

The Tag dashboard shows all tagged resources. Use the Tag dashboard to see specific tag values for a group of resources such as stacks or virtual machines.

Before you begin

Role required: sn_cmp.cloud_admin

Procedure

1. In the Cloud Admin Portal, navigate to **Analyze** > **Tag**.

   The **Tag Dashboard** tab displays the Tag dashboard where all tagged resources are broken down by tag.

   The Tag dashboard

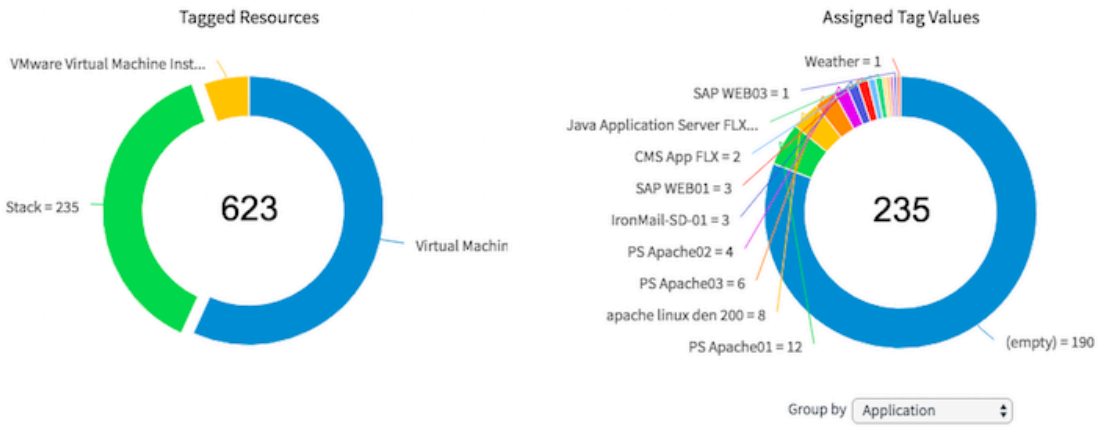

2. Perform any of the following actions to obtain the tag data you want on the report.

Goal	Action
See tag values for a specific resource group	Click a section of the **Tagged Resources** donut chart.
See data grouped by another tag	Select a value in the **Group by** choice list un der the **Assigned Tag Values** donut chart.
See updated data	Point to the top of either chart until the re fresh icon ( ) appears, and then click the icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Goal	Action
Save an image of a chart	Point to either chart until the options icon ( ) appears, and then select **Save as PNG** or **Save as JPEG**.

3. To view a specific tag history record, click any section in the **Assigned Tag Values** chart.
A list of tag histories appears that matches the CI class for the selected tag value.

A tag history record

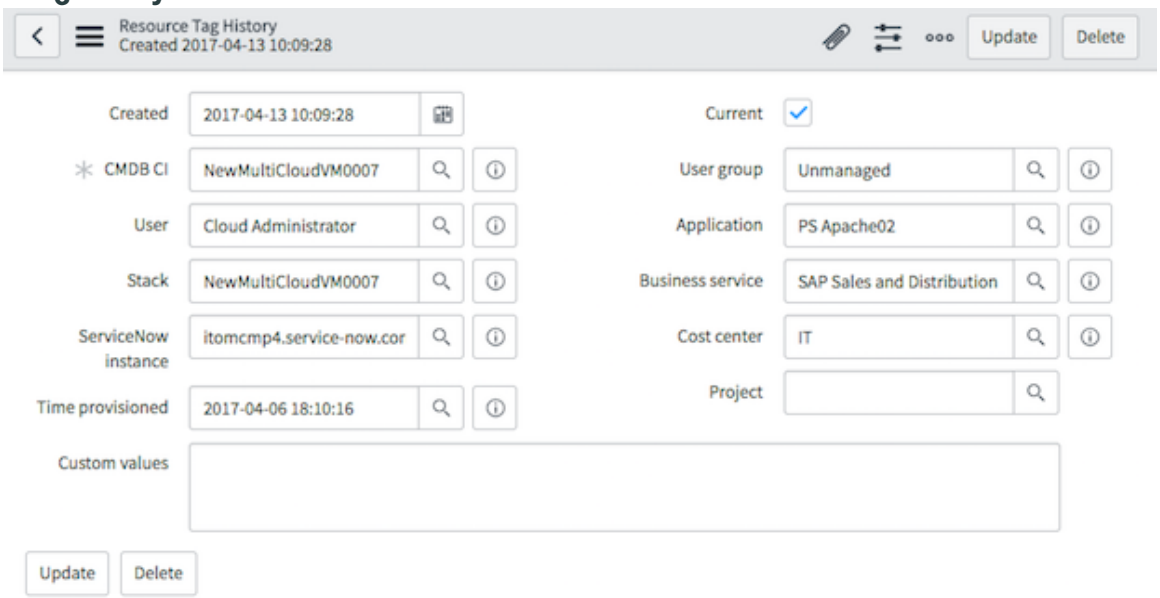

Field	Description
Created	The date the tag history was created.
Current	If selected, indicates if the current record is the must current.
CMDB CI	The actual CI in the CMDB. Click the record icon to see the CI form.
User group	The group tag associated with this CI.
User	The user tag associated with this CI.
Application	The application tag associated with this CI.
Stack	The stack tag associated with this CI.
Business service	The business service tag associated with this CI.
ServiceNow instance	The instance tag associated with this CI.
Cost center	The cost center tag associated with this CI.
Time provisioned	The date the machine was provisioned.
Project	The project tag associated with this CI.
Custom values	Any custom values tagged to the CI.

Budget-based notification and approval

Administrators can assign a budget for a group and a user within the group. When the user or group reaches the budget limit threshold, notifications are sent alerting them about it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

- The budget-based notification and approval feature is no longer available if you are using the Cloud Cost Management app for cloud billing.
- You can only continue using the Budget Consumption feature in two scenarios. You are using Cloud Provisioning and Governance on a domain separated instance, or you have switched back to the native Cloud Provisioning and Governance billing feature.

Administrators can assign a budget (in USD) on a weekly, monthly, quarterly, or yearly basis. The budget set for a user and for a group are independent of each other. For example, a group consisting of five users can have a budget of $100 and each user in that group can be assigned a limit of 25 dollars. An organization decides on the frequency of the budget and all the groups in that organization follow the same frequency. For example, if an organization decides on a monthly budget, then all the groups and users in that organization follow a monthly budget. See Configure budgets.

A default budget is given to each new group and new user. A new group gets a default budget of $1000 and a new user gets a default budget of 100 dollars.

Administrators can increase or decrease the budget at any given time. Notifications are sent to the user and the group when the budget limit reaches its threshold limits. Notifications trigger as part of the billing process. Billing discovery is scheduled for each user and group. At the end of the billing discovery, the budget limit is compared to the actual cost. If the threshold is reached or exceeded, notifications are triggered.

Administrators can set up a policy for budget limits. When the budget limit reaches a particular threshold or exceeds the limit, the administrator can approve continued resource use for the user or group. For example, the administrator can create a policy for a group that requires approval when the group's budget threshold reaches 90%. If the administrator does not set up a policy for the budget, the user or the group can continue using the resources.

You can view the budget details on the at **Analyze** > **Budget Consumption**. You can view the budget consumption details for a user as well as for groups.

Configure budgets

You can configure budgets for groups and users within that group. You can set up a budget period for the group, allocate a budget limit to a group and to each user within that group.

Before you begin

Role required: admin

Procedure

1. In the Cloud Admin Portal, navigate to **Govern** > **Budget**.
2. Click **New** next to **Budget Configuration**.
3. Fill in the form fields (as shown in the table)

Field	Description
Group Name	Select a group from the list.
Budget Period	Select a budget period for the group: weekly, monthly, quarterly, or yearly.
Group Max Limit	Enter a budget limit for the group. By default, the budget limit for each new group is 1000 dollars.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Per User Limit	Enter a budget limit for each user in that group. By default, the budget limit for each new user is 100 dollars.

4. Click **Submit**.

Tags for cloud resources

Tags categorize cloud resources to provide richer and more detailed tracking and billing report data.

Create tags to enable Cloud Provisioning and Governance to automatically categorize cloud resources after they are provisioned. Each tag associates a specific `key` : `value` pair with a requested cloud resource, such as `Application`=HRMS.

After the cloud resource is provisioned, the key:value information in tags become part of the resource's metadata and are saved along with all other parameters about the resource. Because a cloud resource can have multiple tags, you can view cost items from multiple perspectives.

Tag definitions

Tag definitions are records that associate a tag name with a value type. These value types are available:

Value type	Description
Variable	A value that user enters or selects from a choice list. This value is stored as a service catalog variable. If a table is associated with a variable tag, the values in the field marked as the **Display** field are displayed to the user when they request the instance. In most cases, this is the Name field. For example, the names of the business services in the Service table display to the user. If no table is specified in the tag definition, the user is presented with a text field.
Script	A script that obtains a value from the instance, such as the current time on the instance. The script must obtain the value via the base-system API and return the value to the `answer` variable. For example, to obtain the current time on the instance, the script would be: `answer = new GlideDateTime().getValue();`

Default tags

These tags are provided by default.

Important: Do not deactivate the default tags, as this action will delete all the corresponding catalog item variables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Default Cloud Tags

Label	Name	Value type	Table that provides choices
Application	application	variable	Application [cmdb_ci_appl]
Business Service	business_service	variable	Service [cmdb_ci_service]
Cost Center	cost_center	variable	Cost Center [cmn_cost_center]
Project	project	variable	n/a
ServiceNow instance	sn_instance	script	n/a
Stack	stack	variable	Stack [sn_cmp_stack]
Time Provisioned	time_provisioned	script	n/a
User	user	order field	User [sys_user]
User Group	user_group	variable	Group [sys_user_group]

Cloud administrators can also create custom tags to suit the specific needs of the organization. For example, you can create a Location tag can to track the location of the resource.

AWS tags

Tags generated for Amazon Web Services (AWS) resources are sent to AWS and are also available from your AWS console.

Related topics

View tagged resources on the Tag dashboard

Create a custom tag for cloud resources

Cloud Administrators can create custom tags in addition to the default tags to categorize cloud resources.

Before you begin

Role required: sn_cmp.cloud_admin

About this task

If a resource was provisioned prior to the existence of a particular tag, then that tag must be manually added to the record for that resource.

Procedure

1. In the Cloud Admin Portal, navigate to **Analyze** > **Tag** > **Tag Management** > **Tag Definitions**.
2. Select **New**.
3. On the form, fill in the fields.

Field	Description
Label	Name to display on screen for the cloud user.
Name	Name of the key. The name is stored in all tables, along with the value.
Value type	The type of value for this tag.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Variable name	The name of the variable. This option appears if you select the **Variable** value type.
Table [optional]	The table that provides the choices for the **Variable** value type field. Select a table from the system. If you do not select a table, cloud users are presented with an empty text field for the variable.
Order field	The field on the request form. Click the down arrow next to **Click to select**, and then select the field from the Orders [sn_cmp_order] table.
Active	An indicator that specifies if the tag is in use. Select this option to apply the tag to all new requests.

Related topics

View tagged resources on the Tag dashboard

Group billing data using custom tags from provisioned resources

Configure user-defined tags to view and analyze billing data using custom tags existing on provisioned resources.

Before you begin

- Perform cloud discovery in Cloud Provisioning and Governance for provisioned resources.
- Run a billing download job for the associated service account.

Role required: Resource Policy Contributor role in Azure Service account and sn_cmp.cloud_admin in the instance.

About this task

Provisioned resources with existing tags are discovered as part of Cloud discovery. Existing tags are considered user-defined tags in Tag Management. By default, the imported (user-defined) tags are not displayed in the Billing dashboard, restricting you from filtering and analyzing billing data using imported (user-defined) tags.

Important: Starting with the Rome release, billing and budgeting features are no longer supported for Cloud Provisioning and Governance running on a non-domain-separated instance. However, if you are using Cloud Provisioning and Governance on a domain-separated instance, there is no change in functionality for the billing and budgeting features.

Procedure

1. In the **Cloud Admin portal**, navigate to **Analyze** > **Tag** > **Tag Management**.

   The **Tag Definitions** list displays discovered resources provisioned with existing tags labeled as User-defined.

2. Open a tag record of type **User-defined**.
3. In the Tag Definition form, change the tag from User-defined to ServiceNow (Custom):
   a. In the **Tag type** field, select **ServiceNow (Custom)**.
   b. Copy the **Label** field value to the **Variable** field.
   c. Select the **Active** and **Visible** check boxes.
   d. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Repeat Steps 2 and 3 for all the user-defined tags that you wish to enable in the Billing Dashboard.
5. Perform Cloud Discovery again to view updated ServiceNow (Custom) tags.

   For more information, see Cloud Discovery.
6. Run the Billing schedule.
   a. In the **Cloud Admin Portal**, navigate to **Analyze** > **Billing** > **Billing Schedule**.
   b. Create a schedule or open an existing schedule.
   c. Select **Execute Now**.
   d. On the **Retrieve Billing Data** dialog box that appears, select a date range and then click **Download**.

   The ServiceNow (Custom) tags now show up in the Tags field on the Billing Dashboard.

What to do next

View cloud costs on the Billing dashboard.

Perform life-cycle operations on discovered resources

Configure life-cycle operations on cloud resources that are not provisioned using ServiceNow® Cloud Provisioning and Governance. Provision resources and modify resource block operations to perform day-2 operations on cloud resources that are discovered using Cloud Discovery but not provisioned using Cloud Provisioning and Governance.

Before you begin

- Assign the configuration item (CI) to a user who performs the operation.
- Role required: sn_cmp.cloud_service_user, root_admin, admin, or sn_cmp.cloud_admin

About this task

If a resource is provisioned from the Cloud Provisioning and Governance catalog, you can perform base system life-cycle operations on a stack or on a resource. However, if a resource is discovered as part of Cloud Discovery and is not provisioned from Cloud Provisioning and Governance, the same Start/Stop, Deprovision operations are not available, by default. The following procedure is applicable for a Stop operation on a virtual machine that is not provisioned from Cloud Provisioning and Governance.

Procedure

1. Navigate to **All** > **Cloud Admin Portal** > **Design** > **Resource Blocks**.
2. Search and open the resource/CI you want to perform life-cycle operations on.
3. Set toggle switch to **Draft**, to make the resource block editable.
4. In the selected **Resource Block** > **Operations** tab, select the required **Interface** and **Operation** from their respective lists.

   Note: The **Interface** and **Operations** fields are only available for resources that are discovered but not added or generated to catalog.
5. Click the **Override Operation** icon.
6. In the **Override Operation** dialog, specify a name for the Interface-operation.
   The **Create New** check box is selected, by default.
7. Click **Submit**.
   You have enabled the selected operation for the interface you created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. In the **Input Parameters** form, click the **Add Input Parameters** icon.
9. Specify `resourceID` as the parameter **Name** and click **Save**.
   By default, a system assigned expression is generated and appears in the **Mapping** field.
10. For the `ServerID` parameter, replace the default expression in the **Mapping** field with `${parameter.resourceId}` and click **Save**.
11. Navigate to **Operations** > **Steps** subtab.
12. In the **Input** form, change the `NodeId` parameter value to `${parameter.resourceId}`.
13. Navigate back to the **Input Parameters** form, and click **Generate Catalog**.
14. Set toggle switch to **Published**, to publish the resource block again.
    You have configured the resource operation for the specific resource/CI.
15. **Optional:** Navigate to **Cloud User Portal** > **Resources** > **Virtual Machine Instance**, and select the virtual machine instance for which you created the resource operation.
16. **Optional:** Select the new interface operation you created from the **Select Resource Operations** list.

What to do next

Navigate to the **Cloud User Portal** > **Resources** > **Virtual Machine Instance** and select the new interface operation you created from the **Select Resource Operations** list.

Pattern-based targeted discovery

Automatically discover the newly provisioned cloud resources and store their information in the ServiceNow® ServiceNow® Configuration Management Database (CMDB) by using pattern-based targeted discovery.

After provisioning the resources mentioned in the resource template, such as Azure Resource Manager (ARM) or CloudFormation Template (CFT), it sends a response to the Cloud Provisioning and Governance application. The response contains information about the provisioned resources. Cloud Provisioning and Governance uses this information to launch targeted discovery. Targeted discovery discovers only the newly provisioned or modified resources. It updates their attributes in the appropriate CMDB CI class.

The pattern-based targeted discovery offers an easier alternative to the existing CAPI-based approach. Cloud Provisioning and Governance uses single-resource-discovery patterns to discover only the newly provisioned or modified resources. Contrary to the other patterns, the single-resource-discovery pattern doesn't scan the entire cloud landscape for changes. It only discovers the newly provisioned or modified resource of a given type.

The Cloud Resource Types [sn_capi_resource_type] table stores the one-to-one association between the resource type and the single-resource-discovery pattern used to discover the resource. Cloud Provisioning and Governance uses the sn_capi_resource_type table to identify the appropriate pattern for the given resource type. For more information on patterns, see Discovery patterns used by ITOM Visibility.

In the base system, Cloud Provisioning and Governance supports pattern-based targeted discovery for a wide variety of resource types. If required, you can use custom patterns that you created to perform targeted discovery. For more information, see Associate a resource type with a single-resource-discovery pattern.

Cloud Provisioning and Governance checks the following conditions to invoke the pattern-based targeted discovery:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The CAPI-based targeted discovery isn't implemented for the resource type.
- The *sn_cmp.targeted.disco.patterns* system property is set to **True**.
- The sn_capi_resource_type table contains a one-to-one association between the resource type and the single-resource-discovery pattern.

For example, if you use CFT to provision an Amazon S3 bucket, Cloud Provisioning and Governance reads the resource type from the cloud response (AWS::S3::Bucket). It then identifies the appropriate pattern (AWS S3) from the sn_capi_resource_type table and invokes the pattern to discover the newly provisioned Amazon S3 bucket only. The pattern updates the bucket's attributes in the appropriate CMDB CI class (Cloud Object Storage [cmbd_ci_cloud_object_storage] table).

If both CAPI-based targeted discovery and pattern-based targeted discovery are implemented for the resource type, then the system uses CAPI-based targeted discovery.

If neither CAPI-based nor pattern-based targeted discovery is implemented for the resource type, the resource type is classified as a generic cloud resource. The resource appears in the Cloud Resources [cmdb_ci_cmp_resource] table. In such cases, you can create a single-resource-discovery pattern and associate it with the resource type. For more information, see Create or customize patterns and Associate a resource type with a single-resource-discovery pattern.

Related topics

Configure a response processor

Create or extend a CAPI interface

Associate a resource type with a single-resource-discovery pattern

Define a one-to-one association between the resource type and the single-resource-discovery pattern in the Cloud Resource Types [sn_capi_resource_type] table. The pattern-based targeted discovery feature uses this information to invoke the appropriate pattern and discover the provisioned resource.

Before you begin

- Ensure that the Cloud Provisioning and Governance application is deployed.
- Ensure that the latest Discovery and Service Mapping is installed.
- Ensure that the application scope is Cloud API.

Role required: sn_cmp.cloud_service_designer or sn_cmp.cloud_admin

Procedure

1. Navigate to the sn_capi_resource_type table.
2. Select the appropriate resource type record.
3. On the form, fill in the fields.

   The example values provided in this table establish a one-to-one relationship between the AWS::S3::Bucket resource type and its corresponding single-resource-discovery pattern (Amazon AWS S3).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	Cloud resource type that you want to discover Example: AWS::S3::Bucket
Product	(Optional) Product category Example: AWS Simple Storage Service
CI Class	CMDB CI table name. The system writes the discovered information into this table. Example: Cloud Object Storage [cmdb_ci_cloud_object_storage]
Pattern	A single-resource-discovery pattern to discover the resource Example: Amazon AWS S3
Dependency sequence	(Optional) Order for triggering the pattern.

4. Select **Update**.

Troubleshooting tools for Cloud Provisioning and Governance

The Cloud Provisioning and Governance application provides several tools to help you troubleshoot errors. You might encounter errors during the Discovery of cloud accounts, the provisioning and managing of cloud resources, and the updating of the CMDB.

These troubleshooting tools are available:

Cloud Orchestration Trail

The Cloud Orchestration Trail is an activity log for all cloud resource activity on the instance. Use the Cloud Orchestration Trail to find errors that occur with cloud resources, such as credential errors during Discovery or API execution errors.

See Open the Cloud Orchestration Trail for instructions.

Cloud API Trail

The Cloud API Trail is an activity log for all activity that uses the Cloud API and goes through the MID Server. Use the Cloud API Trail to see API invocations and errors related to the route data, specific API route errors, and Java runtime exceptions.

See Open the Cloud API Trail for instructions.

Root Cause Analysis Dashboard

The Root Cause Analysis Dashboard brings together records from the cloud orchestration trail and cloud API trail and presents them in useful, filterable, lists and charts.

See Use Cloud Root Cause Analysis reports for instructions.

Cloud Orchestrations

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Orchestrations are the orders that your instance processed for each attempted operation on a stack. Use Cloud Orchestration records to view status messages for operations that are run on cloud resources and for details about each API step.

See Open cloud orchestrations for instructions.

The Cloud Orchestration Trail

The Cloud Orchestration Trail is an activity log for all cloud resource activity on the instance.

Cloud Orchestration Trail contents

Cloud orchestration trail form

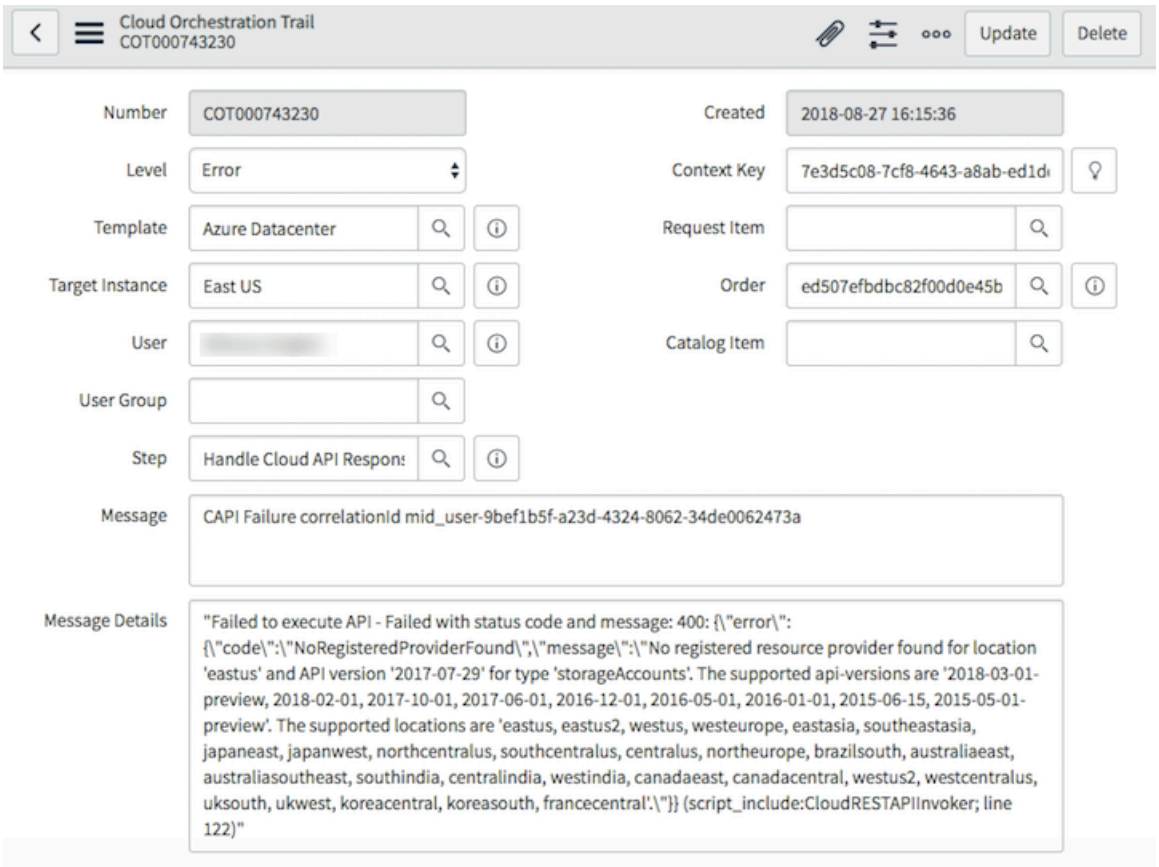

Field	Description
Number	An auto-generated number for the trail entry.
Level	One of the following: • Debug • Info • Warn • Error
Template	The component, such as a resource block or blueprint, on which a Cloud Provisioning and Governance action was applied.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Template Instance	The datacenter in which the cloud resource is located.
User	The user involved in the transaction.
User group	The user group that the user belongs to.
Created	The date and time that the system created the entry.
Milestone	If this transaction is a milestone. You can configure the form to see this field.
Request Item	The service catalog request generated for the transaction.
Step	The step that identifies additional information about the transaction, including the component and the stage. You can configure the form to see this field.
Catalog Item	The catalog item, if any, that was involved in the execution.
Message	The error message or successful completion message. This could include items like the ID of the corresponding message in the Cloud API Trail record or the Cloud API that was invoked.
Message details	Details about the error, such as an HTTP error.

An example of an error message detail is as follows:

```
Failed to list loadbalancer Failed : HTTP error code : 403
```

This indicate that the credentials were incorrect, and Discovery could not access the cloud resource. The load balancer interface throws the first error because the load balancer device is the first device that allows access to the cloud resource.

A corresponding error on the MID- Server is captured in the Cloud API Trail.

Steps, stages, and components

Each Cloud Orchestration entry has identifying attributes that explain what the triggering transaction was about. Steps group these attributes together and identify the component in the system that was involved in the transaction and the stage at which the transaction took place.

You can configure the Cloud Orchestration Trail form to show the **Step** field if it is not already visible.

Step	Component	Stage
Apply API Parameter	CAPI Probe	Cloud API Orchestration
Apply Config Parameter	CAPI Probe	Cloud API Orchestration
Apply Endpoint Binding	CAPI Probe	Cloud API Orchestration
Apply Policy	Policy Engine	Workflow and Policy Execution
Apply Policy Rule Action	Policy Engine	Workflow and Policy Execution
Call Cloud API	BP DSL Executor	Cloud Orchestration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Step	Component	Stage
Call Cloud API Connector	CAPI Probe	Cloud API Orchestration
Check Cloud Match	BP DSL Validator	Cloud Orchestration
Check Datacenter Capacity	BP DSL Validator	Cloud Orchestration
Check Resource Block Policies	BP DSL Executor	Cloud Orchestration
Check User Permissions	BP DSL Validator	Cloud Orchestration
Check User Quota	BP DSL Validator	Cloud Orchestration
Complete Cloud Orchestration	BP DSL Executor	Cloud Orchestration
Execute Cat Item Policy	Dynamic Forms	Catalog Item Request
Execute Form Rules	Dynamic Forms	Catalog Item Request
Execute Resource Block Operation	BP DSL Executor	Cloud Orchestration
Execute Resource Instance Script	BP DSL Executor	Cloud Orchestration
Get Quota Details	Dynamic Forms	Catalog Item Request
Get Response	CAPI Orchestrator	Cloud API Orchestration
Handle Cloud API Response	BP DSL Executor	Cloud Orchestration
Handle Policy Execution Response	BP DSL Executor	Cloud Orchestration
Handle Response and Chunking	CAPI Probe	Cloud API Orchestration
Invoke Cloud Orchestration	Workflow Engine	Workflow and Policy Execution
Invoke Policy Engine	Workflow Engine	Workflow and Policy Execution
Invoke Pre-Process Workflow	Workflow Engine	Workflow and Policy Execution
Launch Cat Item	Catalog Item	Catalog Item Request
Notify Caller	CAPI Orchestrator	Cloud API Orchestration
Prepare BP Execution Plan	BP DSL Executor	Cloud Orchestration
Prepare Runnable API	CAPI Orchestrator	Cloud API Orchestration
Publish API Route Request	CAPI Orchestrator	Cloud API Orchestration
Send Cloud API Response	CAPI Probe	Cloud API Orchestration
Show Form with Filters	Dynamic Forms	Catalog Item Request
Submit Form	Catalog Item	Catalog Item Request
Validate Inputs	BP DSL Validator	Cloud Orchestration

See Use Cloud Root Cause Analysis reports for instructions on how to access reports related to the Cloud Orchestration Trail.

Related topics

The Cloud API Trail

Open the Cloud Orchestration Trail

Open the Cloud Orchestration Trail to debug and troubleshoot issues like a failed policy or failed Discovery of cloud resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_cmp.cloud_operator or sn_cmp.cloud_admin

Procedure

1. In the Cloud Admin Portal, navigate to **Operate** > **Trails**.
2. On the **Cloud Orchestration Trail** tab, filter and sort the list of Orchestration API Trail records as needed.
   If you are looking for something like a failed Discovery, filter the list so the **Level** column shows only entries with **Error**.
3. Click a link in the **Created** column to open the Orchestration API Trail record.
4. In the CAPI Trail Logs related list, open the log record that displays the information you want. For example, open **route_error** or **error_detail** to debug a failed operation.

Open cloud orchestrations

Cloud orchestration records show you the orders that your instance processed for each attempted operation on a stack. They also show you the values of the fields that the user submitted through the Cloud User Portal. Use cloud orchestrations to troubleshoot issues that occur when a user provisions a cloud resource or runs another operation on an existing cloud resource.

Before you begin
Role required: sn_cmp.cloud_operator or sn_cmp.cloud_admin

Procedure

1. On the instance (not the Cloud User Portal), enter the following text into the application filter:

   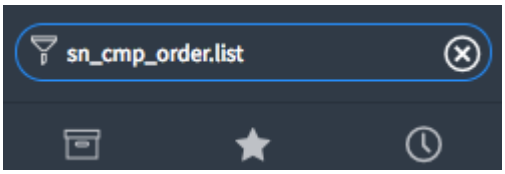

   The list of Cloud Orchestrations appear.
2. Sort and filter the list to find the error or other message you are looking for.

   For example, to sort the list by date, click the **Order Date** column. To show only errors that occurred, right-click **Error** in a cell and select **Show Matching**.
3. Open a cloud orchestration record by clicking the **Order Date** field.
4. Review the form fields:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

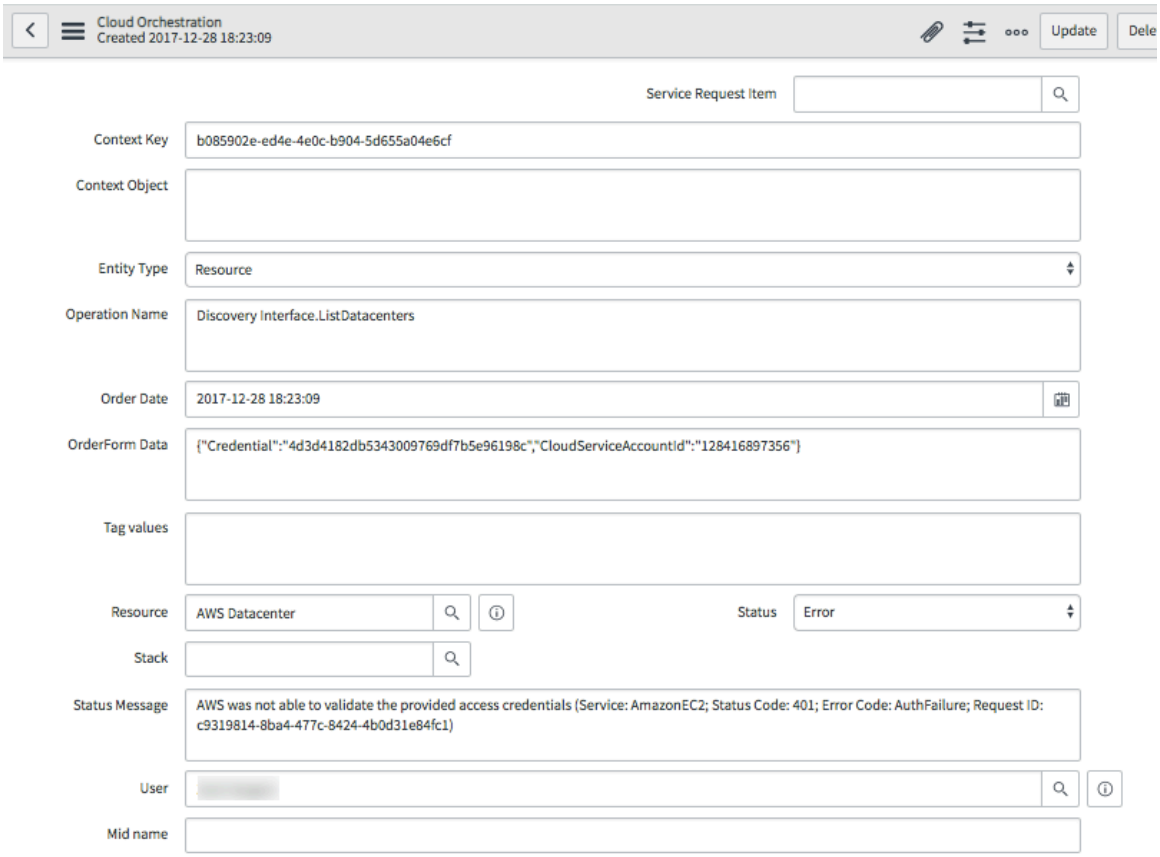

The example form shows you the type of message you can see when an operation fails due to the incorrect credentials. The OrderForm Data field shows the Sys ID of the credential record and the service account ID.

Field	Description
Service Request Item	A service catalog request, if any, associated with this transaction.
Context Key	A key identifier for this message.
Context Object	This field is typically empty.
Context Instance	The blueprint used in transactions on a stack or the resource block used in transactions on a cloud resource.
Owner Instance Owner Table	The Owner Instance is the resource block on which the transaction took place. This field is only populated when the **Entity Type** is **Resource**. The owner table is always Stack [sn_cmp_stack] for the provisioning of a stack. Otherwise, the table is the CI type of the resource, such as Virtual Machine [sn_cmp_vm_instance].
Entity Type	Whether the transaction occurred on **Blueprint** or a **Resource** (resource block).
Operation Name	The interface and operation that is triggered.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Order Date	The date of the transaction.
OrderForm Data	The data submitted through the user form.
Tag values	Any tag values involved in the transaction.
Resource	The resource block on which the operation took place. This field is populated only when the value of the Entity type field is **Resource**.
Status	The status of the message. If you are viewing an error, the **Error** option is listed.
Stack	The stack that the resource belongs to.
Status Message	The status message that explains the issue. Search the cloud provider documentation for the status codes and messages that appear in this field. The cloud provider might provide solutions available to you.
User	The user who triggered the operation.
Mid name	This field can show the name of the MID Server involved in the transaction. It is often left blank.

Moving Cloud Provisioning and Governance content across environments

You can move the content in the Cloud Provisioning and Governance application from one environment to the other. For example, develop an application in an environment and then move the application and all its dependencies to another environment for testing or production.

Content entities that can be moved

The Cloud Provisioning and Governance application consists of code, resources, content, and configurations.

You can move the following content entities across environments:

- Blueprints
- Resource blocks
- Cloud APIs

Use the update sets to move the content entities. Update sets, which are XML files, group all the content into named sets and then moves these sets, as a unit, to other systems for testing or deployment. For information on how system update sets work, see System update sets .

Content entity dependencies

Each of the content entities has some direct dependencies that are automatically moved with the content to the new environment.

The direct dependencies for each of these content entities are:

Content entity	Dependencies
Blueprint	• Resource blocks • Cloud APIs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Content entity	Dependencies
	• MID scripts • Resource pools • Policies • Script includes • Cloud scripts
Resource blocks	• Resource pools • Resource operation policies • Script includes • Cloud APIs • Dependant resource blocks • MID scripts
Cloud API	MID scripts

Related topics

Move a blueprint from one environment to another

Move a resource block from one environment to the other

Move a cloud API from one environment to the other

Move a blueprint from one environment to another

Use update sets to move a blueprint and its dependencies from one environment to another. Update sets let you group a blueprint and its dependencies into a named set and then move them as a unit to other systems for testing or deployment. For example, you can move a blueprint from a development environment to a production environment.

Before you begin

Role required: sn_cmp.cloud_service_designer or admin

- Users with the sn_cmp.cloud_service_designer role can only export update sets.
- Users with the admin role can export and import update sets.

About this task

Package the blueprint as an update set. Then export the update set from its current environment and import it to another environment.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Blueprints**.
   All blueprints, in published and draft mode, appear in the Blueprints window.
2. Export the blueprint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Click the Export icon ( ) for the blueprint that you want to export.

Blueprint exporter window

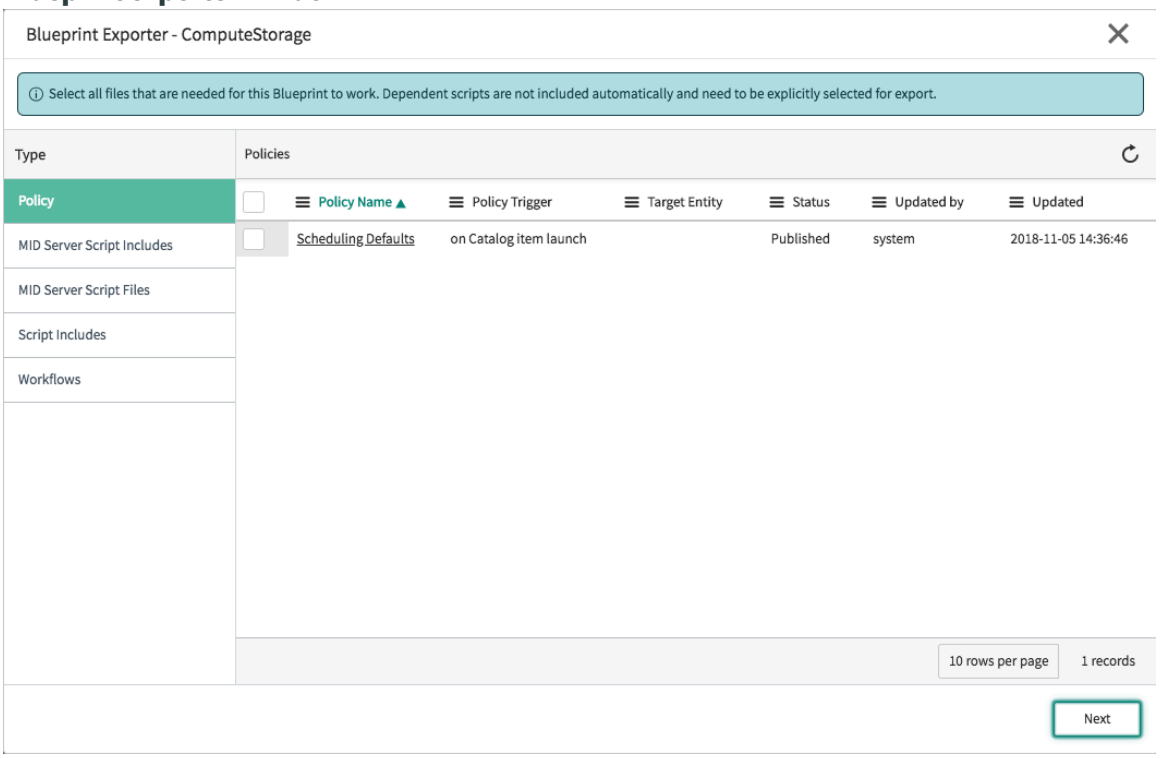

The Blueprint Exporter window has a list of all the indirect dependencies (Policy, MID Server Script Includes, MID Server Script Files, Script Includes, and Workflows) for the blueprint. Objects such as policies, pools, MID scripts, script includes, and workflows are not directly a part of a blueprint. However, a blueprint might depend on these objects to work correctly.

b. In the Blueprint Exporter window, click an object in the Type column.
In the right column, select the object's corresponding entries that you want to export with the blueprint. For example, if you select MID Server Script Includes, all the corresponding script includes appear in the right column. If a workflow contains subworkflows, select the parent workflow and all the subworkflows in order for the workflow to be exported. If there are any custom activities associated with a workflow, export those activities first and then export the workflow.

c. Click **Next**.
A window opens with a summary of the indirect dependencies that you chose to include in the export update set.

d. Click**Export Update Set**.
The Blueprint Exporter window opens with the **Success** check box selected in green indicating that the blueprint has been successfully exported along with all its dependencies. By default, all the files listed under Exported Update Set(s) get downloaded onto your system automatically. However, if the files are not automatically downloaded, follow these steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

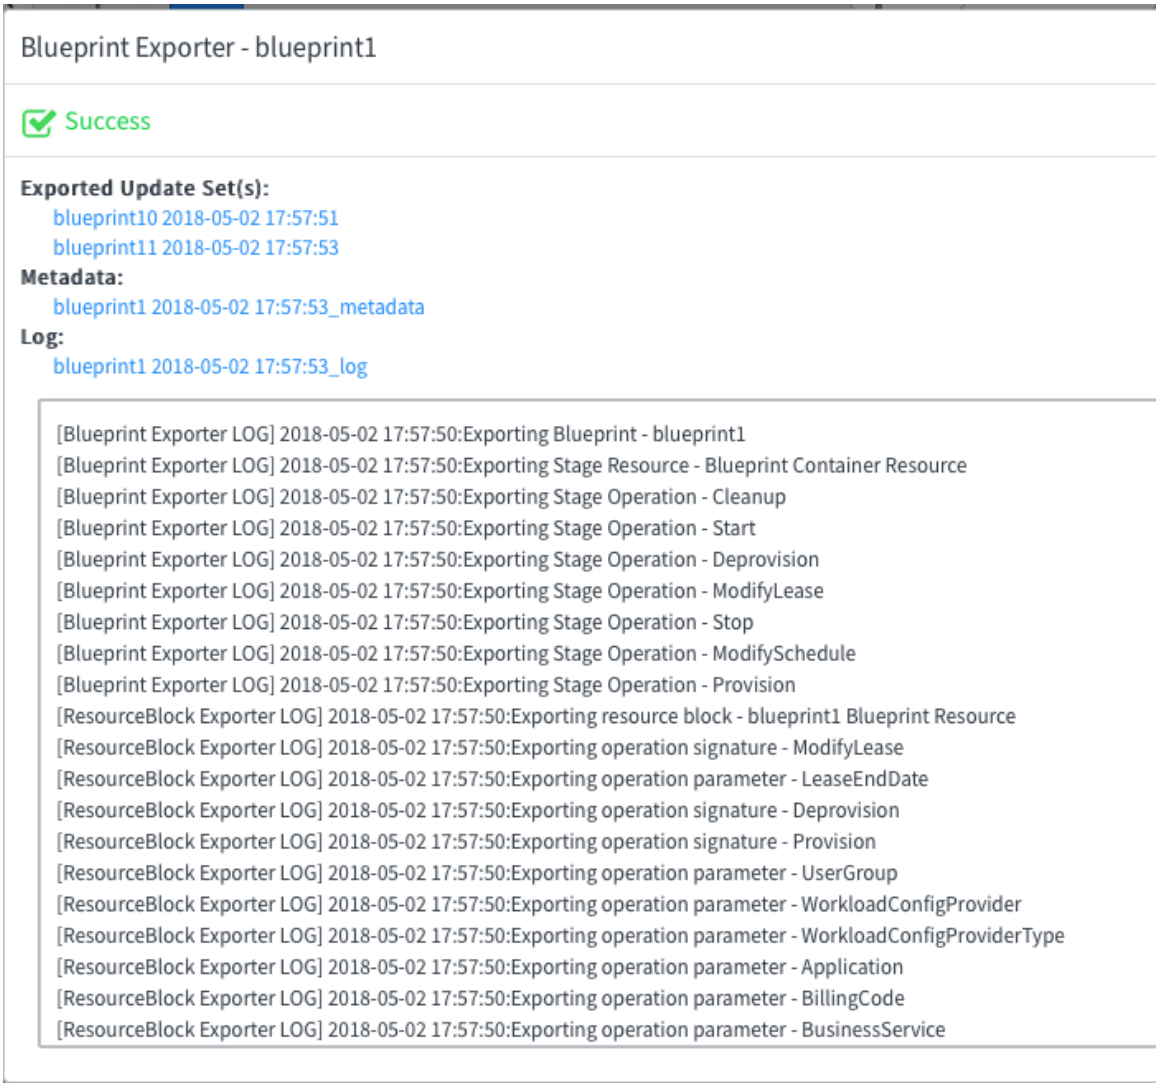

Note: The number of update sets created is based on the scope of records that are being exported. For example, assume the blueprint that you are exporting has 100 records in all: 60 records are in scope one, 20 records are in scope two, and the remaining 20 records are in scope three. In such a scenario, three update sets are created. One update set for each scope.

3. Make sure that the files are downloaded.
   If the files are not downloaded, complete these steps.

   a. Click the metadata file.
   The metadata file mentions the order in which the exported files should be imported. In this example, the file `blueprint10` is the first file to be exported followed by the file `blueprint11`.

   b. Based on the order mentioned in the metadata file, click the appropriate file to open it.
   A window opens with a list of all the files contained in the update set.

   c. To download the first XML file (in this example, the file is `blueprint10`) onto your system, click **Export to XML**.
   Import all the exported files into another environment. Users with the admin role can import the files. For example, you may have created the export update set in a system that runs the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

development environment and want to import it into another system that runs the production environment.

d. Open the other XML file (in this example, the file is `blueprint11`) and download that file onto your system, too.

4. Import the blueprint.
   a. In the environment and the new instance where you want to import the files, enter `Retrieved Update Sets` in the filter navigator and then press the Enter key.
   b. Click the **Import Update Set from XML** related link.
   c. In the Import XML window that appears, click **Choose File**, select the export file, and click **Upload**.

   Note: If a blueprint that you exported from the source environment exists in the target environment, and both the blueprints have a different sys_id, when you import the blueprint, an error appears and the import process stops. Export a different blueprint from the source environment or delete the blueprint with the same name in the target environment. If a blueprint that you exported from the source environment exists in the target environment and both the blueprints have the same sys_id, then the blueprint in the target environment gets updated with the blueprint that you exported from the source environment.

5. To verify that the blueprint is imported into the new environment, go to your instance in the new environment and in the Cloud Admin Portal, navigate to **Design** > **Blueprints**.
   The blueprint you imported should appear in the listed blueprints.

Move a resource block from one environment to the other

Use update sets to effortlessly transfer resource blocks and their dependencies between environments. Group them into a named set, facilitating seamless movement for testing or deployment. Simplify processes such as transferring a resource block from development to production, ensuring efficient and organized transitions across different systems.

Before you begin

- Users with the sn_cmp.cloud_service_designer role can only export update sets.
- Users with the admin role can export and import update sets.

Role required: sn_cmp.cloud_service_designer or admin

About this task

Package the resource block as an update set. Then export the update set from its current environment and import it to the other environment.

Note: If a resource block that you export from the source environment exists in the target environment, when you import this resource block, it results in two resource blocks with the same name: one resource block that existed before the import and another resource block which you just imported. For example, if you export the AWS Datacenter resource block from the development environment and you already have a resource block with the same name in the production environment, then when you import the AWS Datacenter resource block to the production environment, two resource blocks with the same name, AWS Datacenter, reside in the production environment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.
   All resource blocks, in published and draft mode, appear in the Resource Blocks window.
2. Export the resource block.

   a. Click the Export Resource block icon ( ) for the resource block to export.

Resource block exporter window

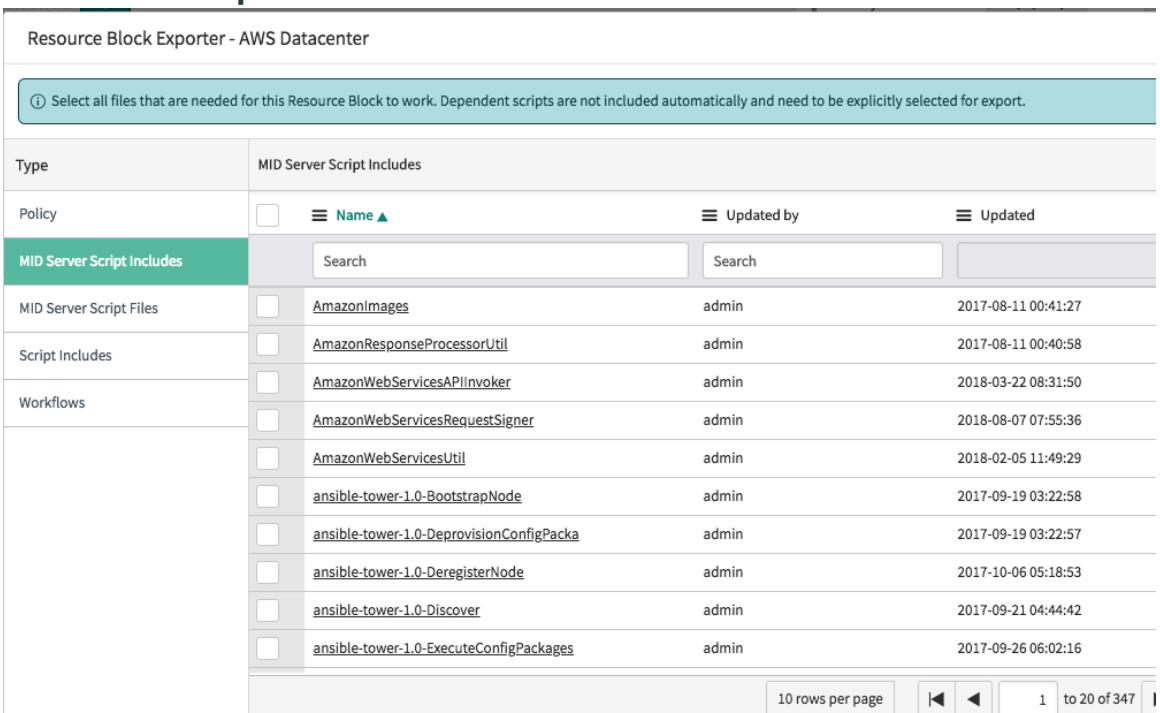

   a. In the Resource Block Exporter window, click an object in the Type column.

   In the right column, select the corresponding entries for the object that you want to export with the resource block. For example, if you select MID Server Script Includes, all the corresponding script includes appear in the right column.

   If a workflow contains subworkflows, select the parent workflow and all the subworkflows in order for the worklow to be exported. If there are any custom activities associated with a workflow, export those activities first and then export the workflow.

   b. Click **Next**.
   A window opens with a summary of the indirect dependencies that you chose to include in the export update set.

   c. Click **Export Update Set**.

   The Resource Block Exporter window opens with the **Success** check box selected in green indicating that the resource block has been successfully exported along with all its dependencies. By default, all the files listed under Exported Update Set(s) get downloaded onto your system automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

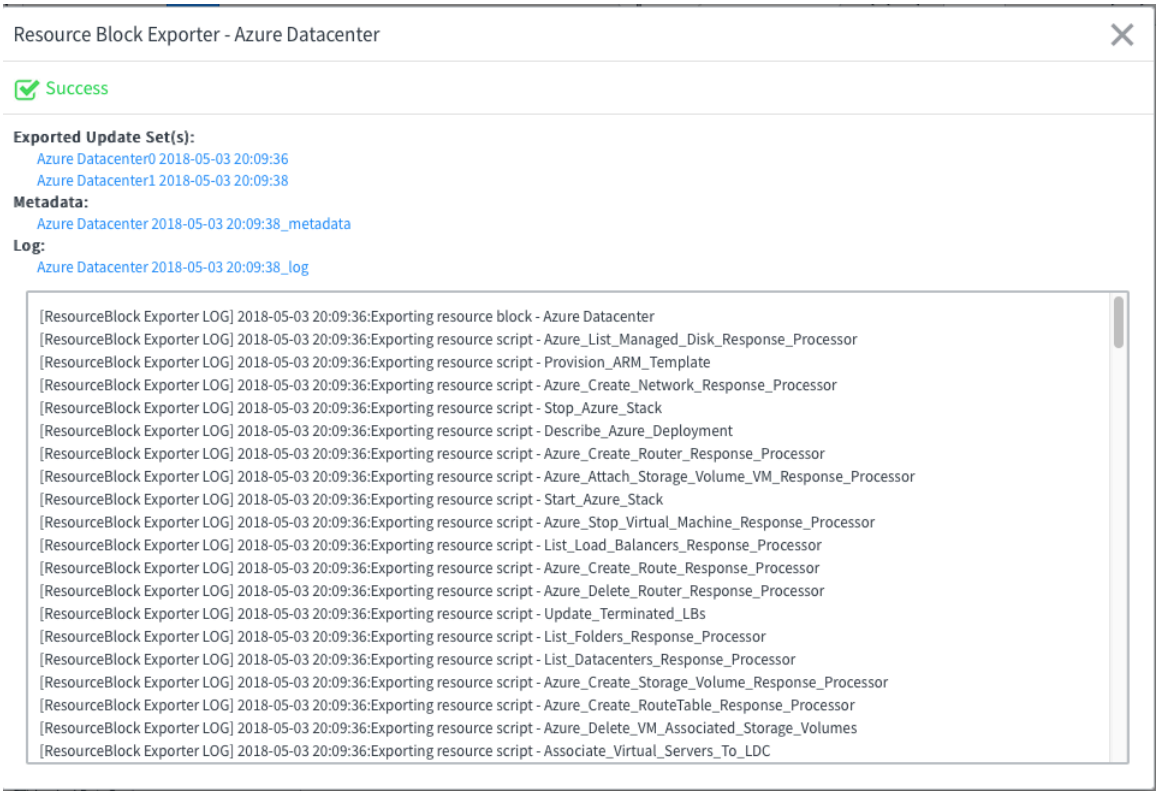

Note: The number of update sets created is based on the scope of records that are being exported. For example, assume the resource block that you are exporting has 100 records in all: 60 records are in scope one, 20 records are in scope two, and the remaining 20 records are in scope three. In such a scenario, three update sets are created. One update set for each scope.

3. Make sure that the files are downloaded.
   If the files are not downloaded, complete these steps.

   a. Click the metadata file.
   The metadata file mentions the order in which the exported files should be imported. In this example, the file `Azure Datacenter0` is the first file to be exported followed by the file `Azure Datacenter1`.

   b. Based on the order mentioned in the metadata file, click the appropriate file to open it.
   A window opens with a list of all the files contained in the update set.

   c. To download the first XML file (in this example, the file is `Azure Datacenter0`) onto your system, click **Export to XML**.

   d. Open the other XML file (in this example, the file is `Azure Datacenter1`), and download that file onto your system.
   You must import all the exported files into another environment. For example, you may have created the export update set in a system that runs the development environment and want to import it into another system that runs the production environment.

4. Import the resource block.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the environment and the new instance where you want to import the files, enter `Retrieved Update Sets` in the filter navigator and then press the Enter key.

b. Click the **Import Update Set from XML** related link.

c. In the Import XML window that appears, click **Choose File**, select the export file, and click **Upload**.

Note: If a resource block that you exported from the source environment exists in the target environment, and both the resource blocks have a different sys_id, when you import the resource block, an error appears and the import process stops. Export a different resource block from the source environment or delete the resource block with the same name in the target environment. If a resource block that you exported from the source environment exists in the target environment and both the resource blocks have the same sys_id, then the resource block in the target environment gets updated with the resource block that you exported from the source environment.

5. To verify that the resource block is imported into the new environment, go to your instance in the new environment and in the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**. The resource block you imported should appear in the listed resource blocks.

Move a cloud API from one environment to the other

Use update sets to move a cloud API from one environment to another. Update sets let you group a cloud API and its dependencies into a named set. You can then move them as a unit to other systems for testing or deployment. For example, you can move a cloud API from a development environment to a production environment.

Before you begin

- Users with the sn_cmp.cloud_service_designer role can only export data sets.
- Users with the admin role can export and import update sets.

Role required: sn_cmp.cloud_service_designer or admin

About this task

Package the cloud API as an update set. Then export the update set from its current environment and import it to the other environment.

Note: If a cloud API that you export from the source environment exists in the target environment, when you import this cloud API, it results in two cloud APIs with the same name: one cloud API that existed before the import and another cloud API which you just imported. For example, if you export the Ansible Configuration API from the development environment and you already have a cloud API with the same name in the production environment, when you import the Ansible Configuration API to the production environment, two cloud APIs with the same name, Ansible Configuration API, reside in the production environment.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud API** > **API**. All cloud APIs appear in the Cloud API window.
2. Export the cloud API.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Click the Export Cloud API icon ( ) for the cloud API to export.

The Cloud API Exporter window

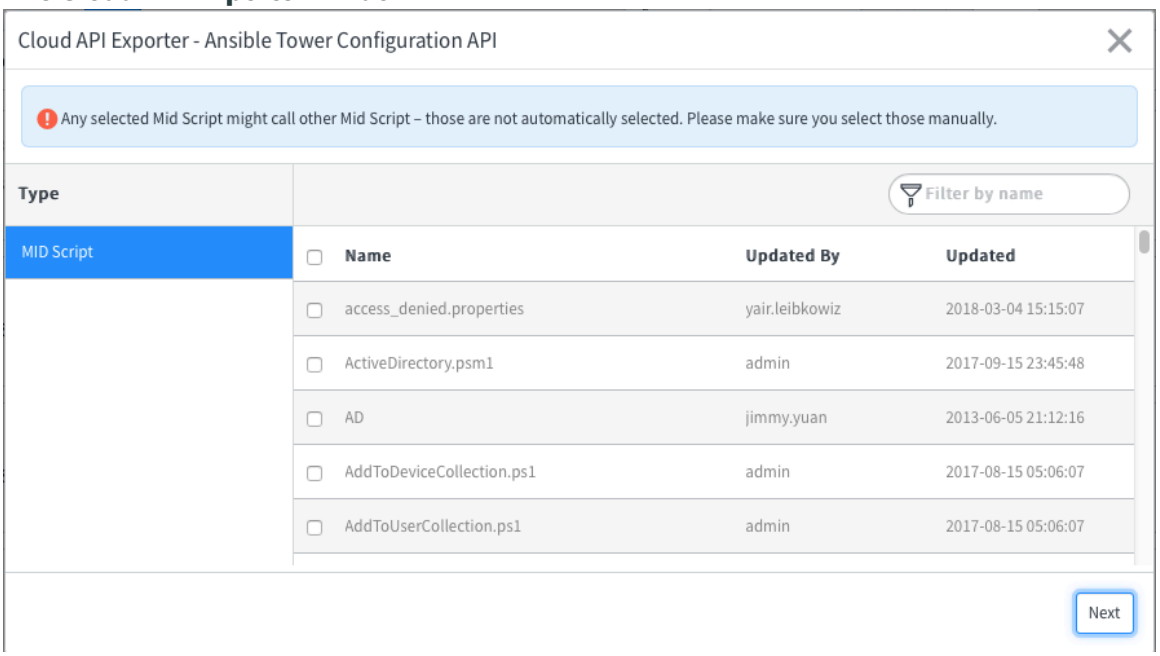

b. In the Cloud API Exporter window, click an object in the Type column.
In the right column, select the corresponding entries for the object that you want to export with the cloud API. For example, if you select MID Script, all the corresponding MID scripts appear in the right column.

c. Click **Next**.
A window opens with a summary of the indirect dependencies of the cloud API that you chose to include in the export update set.

d. Click **Export Update Set**.
The Cloud API Exporter window opens with the **Success** check box selected in green indicating that the cloud API has been successfully exported along with all its dependencies. By default, all the files listed under Exported Update Se(t)s get downloaded onto your system automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

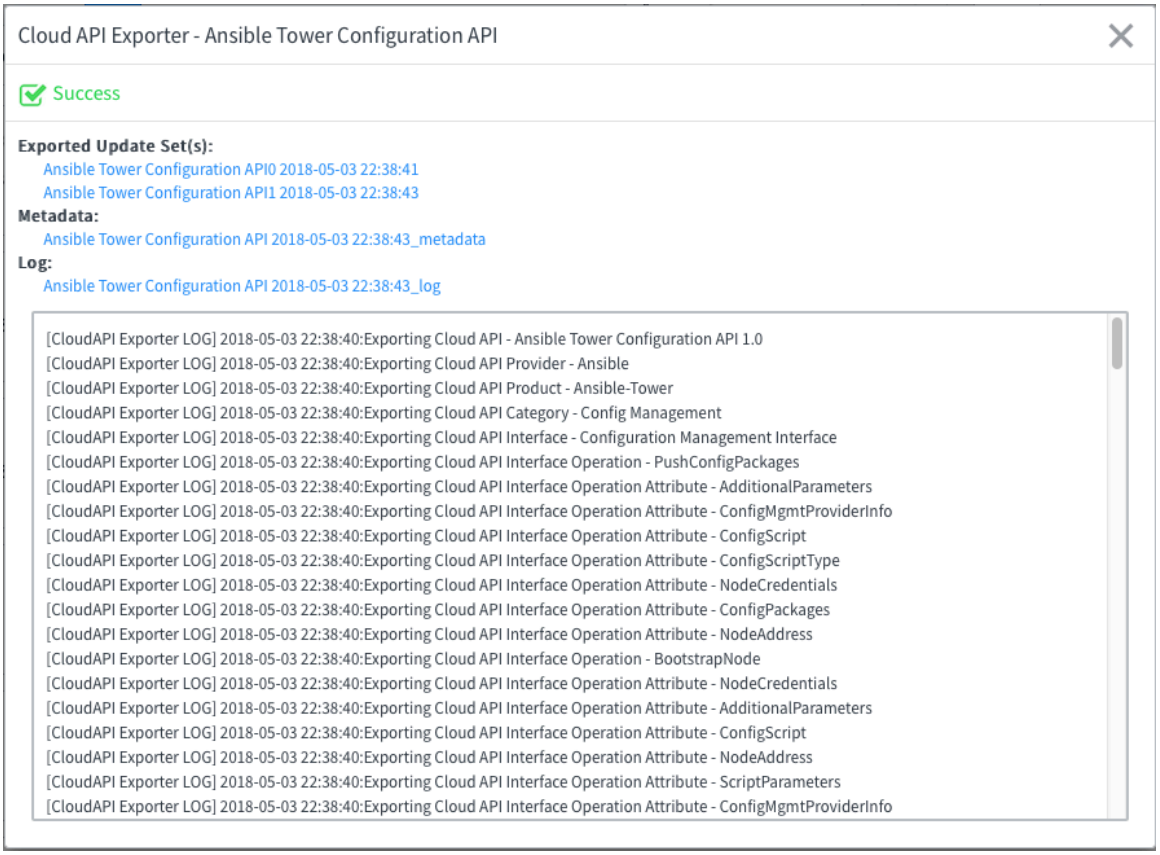

Note: The number of update sets created is based on the scope of records that are being exported. For example, assume the cloud API that you are exporting has 100 records in all: 60 records are in scope one, 20 records are in scope two, and the remaining 20 records are in scope three. In such a scenario, three update sets are created. One update set for each scope.

3. Make sure that the files are downloaded.
   If the files are not downloaded, complete these steps:

   a. Click the metadata file.
   The metadata file mentions the order in which the exported files should be imported. In this example, the file titled `Ansible Tower Configuration API0` is the first file to be exported, followed by the file titled `Ansible Tower Configuration API1`.

   b. Based on the order mentioned in the metadata file, click the appropriate file to open it.
   A window opens with a list of all the files contained in the update set.

   c. To download the first XML file (in this example, the file is `API0 2018-05-03`) onto your system, click **Export to XML**.

   d. Open the other XML file (in this example, the file is `API1 2018-05-03`), and download that file onto your system.
   Import all the exported files into another environment. For example, you may have created the export update set in a development system. You can then import it into a production system.

4. Import the cloud API.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the environment and the new instance where you want to import the files, enter `Retrieved Update Sets` in the filter navigator and then press the Enter key.

b. Click the **Import Update Set from XML** related link.

c. In the Import XML window that appears, click **Choose File**, select the export file, and click **Upload**.
The cloud API gets imported to your new environment.

5. To verify that the blueprint is imported into the new environment, go to your instance in the new environment and in the Cloud Admin Portal, navigate to **Design** > **Cloud API** > **API**.
The cloud API you imported should appear in the list of cloud APIs.

Cloud API (CAPI)

The Cloud API (CAPI) enables you to integrate Cloud Provisioning and Governance with cloud providers using REST APIs.

CAPI components

Integration with cloud providers is performed through REST calls, such as PUT, GET, POST, and DELETE. CAPI provides the framework for you to integrate a cloud vendors REST API so that your instance can communicate with the cloud provider to manage cloud resources.

See Browse APIs by product to learn more.

CAPI contains these components:

CAPI overview

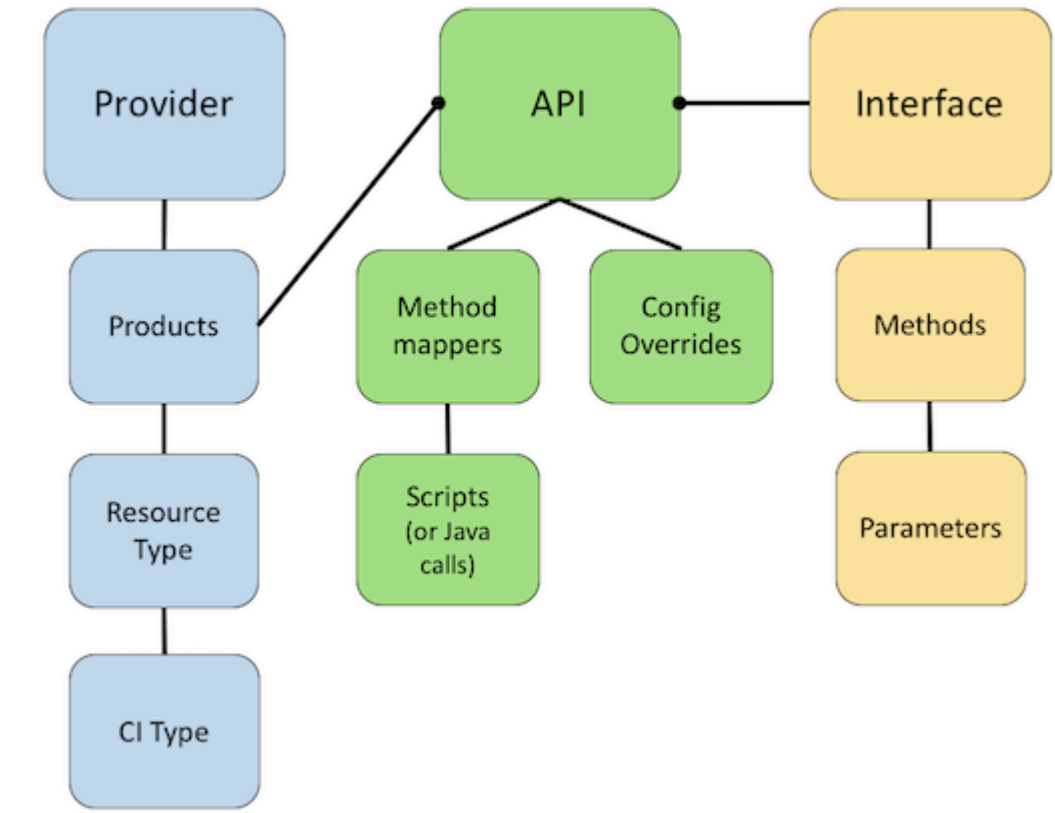

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Providers

Cloud providers are the clouds that you can connect to. By default, Cloud Provisioning and Governance includes the most commonly used providers, such as AWS, Azure, and VMware. Each provider has many products, each of which provide resource types. Each resource type maps to a single CI type. For example, the *`AWS provider`* includes the *`AWS Elastic Compute Cloud`* product, which includes the *`AWS::EC2::Instance`* resource type. This resource type is one of the most common cloud resources that you can create. It maps directly to the Virtual Machine Instance [cmdb_ci_vm_instance] CI type, where virtual machines are saved in the CMDB.

Interfaces

Interfaces define the framework that the system needs to structure the REST calls that the cloud provider APIs expect. Interfaces define operations, also called methods, and the parameters that each method requires.

Interfaces are reusable. If you extend CAPI to include new products and APIs, you can use existing interfaces to make the same REST calls.

APIs

CAPI APIs are the core component of CAPI that ties a product and interface together. The APIs include the actual code that the system executes.

Each CAPI API includes these components:

- CAPI Method Mappers provide the methods that are mapped to the operations defined in the interface. From CAPI method mappers, you create MID Server script includes in JavaScript to tell the cloud provider exactly what to do. It is through the script includes that the connection to the cloud provider occurs.

  Note: If you create custom CAPI APIs, the system provides an empty script include for you to customize. You can also modify existing script includes on method mappers if necessary. However, many of the default APIs that come with the Cloud Provisioning and Governance application do not use modifiable script includes. The connections are hard-coded in Java. You can still use these APIs in new resource blocks that you create, but you cannot modify the APIs.

- API Config Overrides contain the identity, such as a key, and credentials, such as the secret key, and other important parameters required by the cloud provider. These parameters help the cloud provider perform the operations in the CAPI Method Mappers related list. API Config Overrides are necessary because, when the system calls the cloud provider API via REST, credential data is not included. Resource blocks use the parameters and values that you define in the API Config Overrides to query the credential store. When your API executes, the attributes are made available to all method calls in your script includes.

  The overrides are scoped only for this API. The overrides do not override anything in other APIs.

Because you can define multiple versions of a CAPI API with slight variations, you can extend (without overwriting) an existing API while retaining desired functionality.

This image illustrates the components of CAPI that are used when you provision an Azure virtual machine using the default CAPI settings provided with Cloud Provisioning and Governance:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CAPI Azure

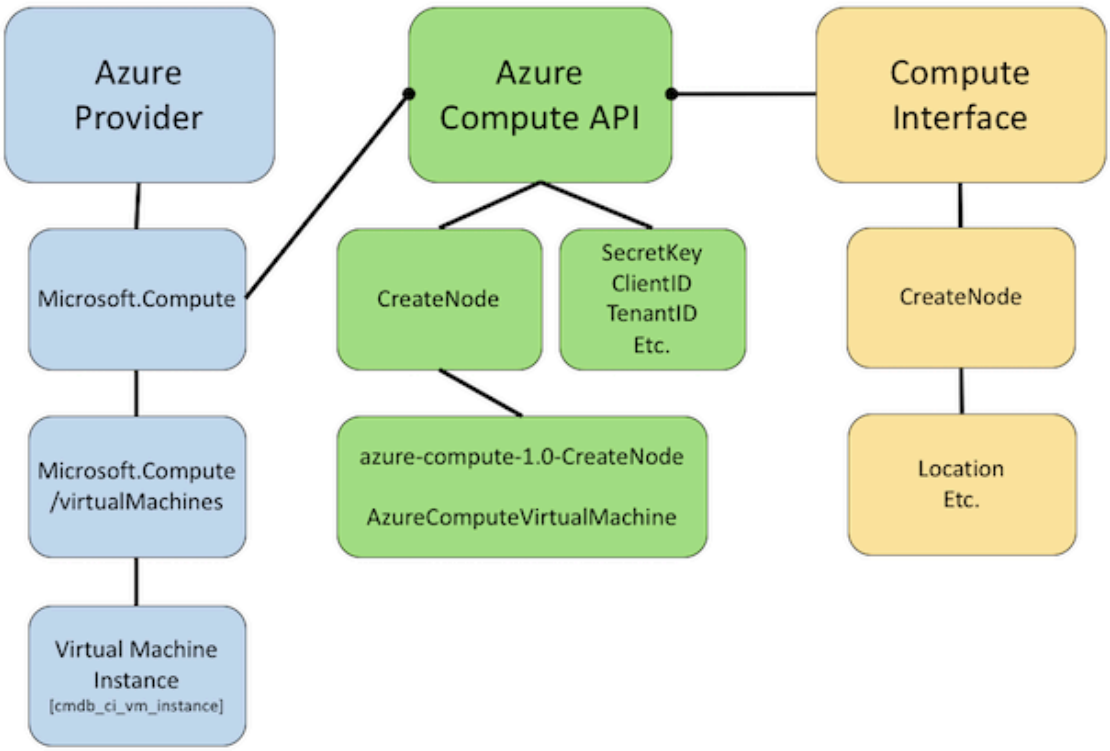

In this example, the *Microsoft.Compute* product is contained in the *Azure* provider. Azure uses the *Microsoft.Compute* product for virtual machines. In your instance, the *Microsoft.Compute* product maps to the *Microsoft.Compute/virtualMachines* resource type, which is associated to the *Virtual Machine Instance* CI type in the CMDB.

The *Compute* interface contains definitions for methods like *CreateNode*, which defines how to create the actual virtual machine. Of the many parameters that *CreateNode* uses, *Location* captures the datacenter where the virtual machine resides.

The *Azure Compute API* pulls together the *Microsoft.Compute* product and the structure defined in the *Compute* interface. The implementation of the *CreateNode* method calls the *azure-compute-1.0-CreateNode* MID Server script include, which calls the *AzureComputeVirtualMachine* MID Server script include. The script includes make the actual calls to the Azure API. To access the Azure account, the *SecretKey*, *ClientID*, *TenantID*, and other methods are passed in Config Overrides.

How CAPI integrates with the instance

CAPI integrates these components in your instance:

Cloud Provisioning and Governance resource blocks

A resource block represents a single cloud resource, such as a virtual server, virtual server storage, or a datacenter. You can also think of it as a CI type in the CMDB. You put many resource blocks together in a blueprint, which appears as a catalog item (also called a stack) to your users in the Cloud Catalog.

In the system, each resource block is like a container that references CAPI, and links responses from the cloud provider to a specific CI. Resource blocks use:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Operational steps that call CAPI for each operation, such as the provision operation, and pass along necessary parameter values that the cloud provider needs to execute the operation.
- Response processors that process and parse the REST response from the provider and update records in the CMDB.

The CMDB

Each resource block is based on a CI type from the CMDB. For Cloud Provisioning and Governance, all cloud-related CI types are based on the *Virtual Machine Object CI* class, which provides all the attributes that you need for all the cloud resources supported by default. If a CI type for a cloud resource does not exist in the base system, you must create a new CI class and add the necessary attributes.

If you do create a new CI class, you must also create:

- A CI class for each of the resources that are available to your users. All CI classes are based on the virtual machine object class.
- An identification rule that specifies an Object ID. Whenever components of Cloud Management refer to a specific cloud resource in the CMDB, they need the Object ID to find the correct cloud resource.
- A relationship rule that specifies how the CI class for the resource is related to other CI classes. For example, a *virtual server* CI must have a *Hosted on::Hosts* relationship with a *datacenter* CI. These relationship rules are necessary for CI uniqueness when processed by the Identification and Reconciliation Engine (IRE). The combination of the service account, the object ID of the resource, and the datacenter (or location) where the resource is located determines the uniqueness.

MID Server script include

Each operation in the CAPI API has a MID Server script include that you configure. The script include calls the JavaScript classes that are already in other script includes in the system, or JavaScript classes that you create. Eventually, the invoker class is called to trigger the REST call. MID Server script includes are configured on your ServiceNow instance, but run on the MID Server.

This image illustrates how the components work together when a user provisions a resource from the Cloud User Portal:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

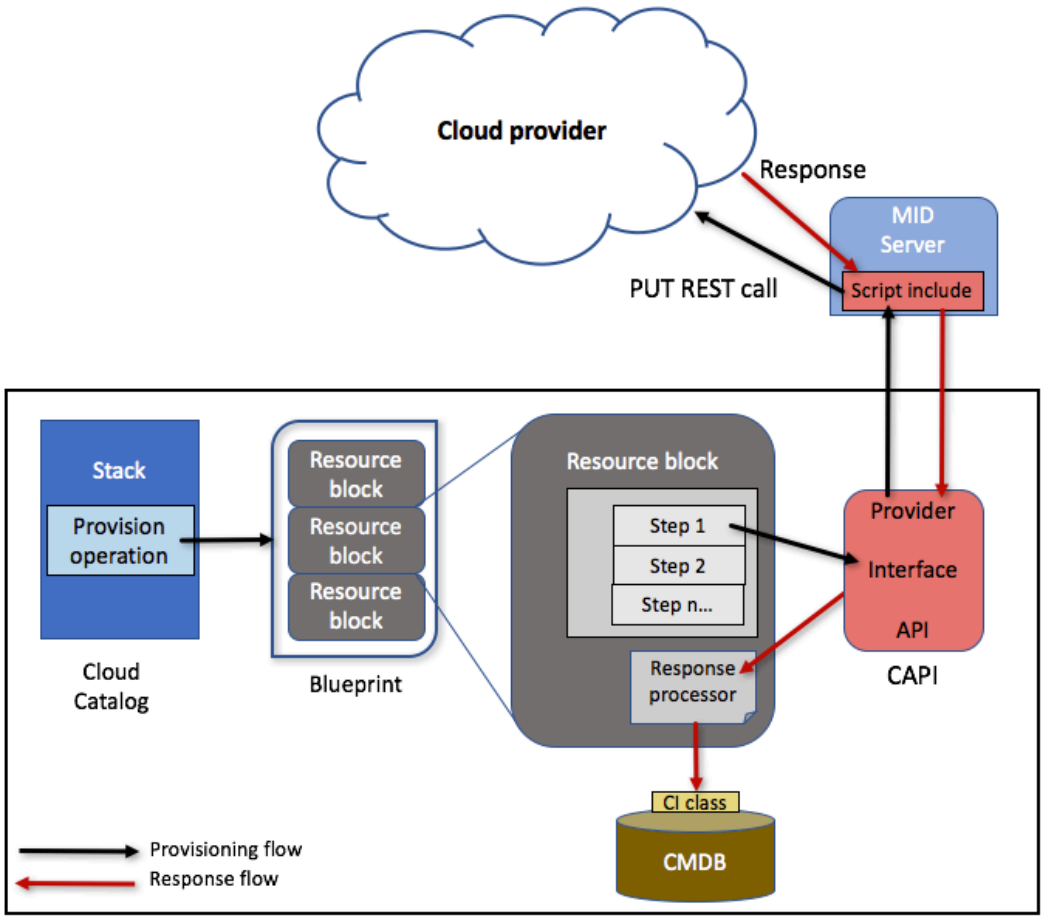

REST calls to the cloud provider

REST calls to the cloud provider are triggered from the MID Server script includes that are referenced from scripted CAPI Method Mappers inside CAPI API records. To create your own CAPI APIs, or to create custom MID Server script includes (which are part of CAPI Method Mappers), you must understand:

- The classes that are available for you to call in MID Server script includes. See CAPI classes in MID Server script includes.
- How the cloud provider implements REST. See:
  - Azure REST API Reference
  - AWS Guides and API References

For example, to learn how to make a REST call to Azure to create a resource group, review this Azure topic: Resource Groups - Create Or Update. You can find the endpoint, parameters, and request body that Azure requires, and the responses it provides. You can see that:

- The endpoint is `management.azure.com`
- The method to call with a PUT operation is `subscriptions/{subscriptionId}/resourcegroups/{resourceGroupName}?api-version=2018-02-01`, where you specify the subscription ID, the resource group name, and the API version.

Remember that the REST API calls take place inside the MID Server script includes that are associated with CAPI API method mappers. Call the methods that CAPI already makes available

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

for you by using the classes extended from *CloudAPIBase* and *CloudRESTAPIInvoker*. You can also create more script includes to extend these base classes and create your own classes. Familiarize yourself with these base classes and the methods available within them.

Get started here

1. Review CAPI APIs that are provided with Cloud Provisioning and Governance by default.
2. Review the CAPI classes that are provided by default. These classes can be called from the MID Server script includes in your CAPI API operations.
3. Walk through the provisioning of an Azure virtual machine and an AWS virtual machine to see how the components work together. The Azure walkthrough uses a MID Server script include so you can see the various CAPI classes used in the provision operation. The AWS walkthrough does not use a MID Server script include.
4. Add a product to an existing provider in CAPI.
5. Create a CI class for a virtual cloud resource.
6. Create or extend a CAPI interface.
7. Create a CAPI API, and a custom MID Server script include that makes the REST calls to the cloud provider. An empty MID Server script include is always generated for new CAPI APIs. Modify it with the calls to other JavaScript classes and methods, such as the methods in the Invoker class.

Default CAPI APIs

Several CAPI APIs are provided by default with the Cloud Provisioning and Governance application.

CAPI APIs refer to the API records that link a CAPI product and a CAPI interface. They do not refer to the CAPI classes that you can call from MID Server script includes in CAPI APIs. For an overview of CAPI classes, see CAPI classes in MID Server script includes.

View Cloud APIs by navigating to **Cloud API** > **Library** in the base system instance interface, or by navigating to **Design** > **Cloud API** in the Cloud Admin Portal.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

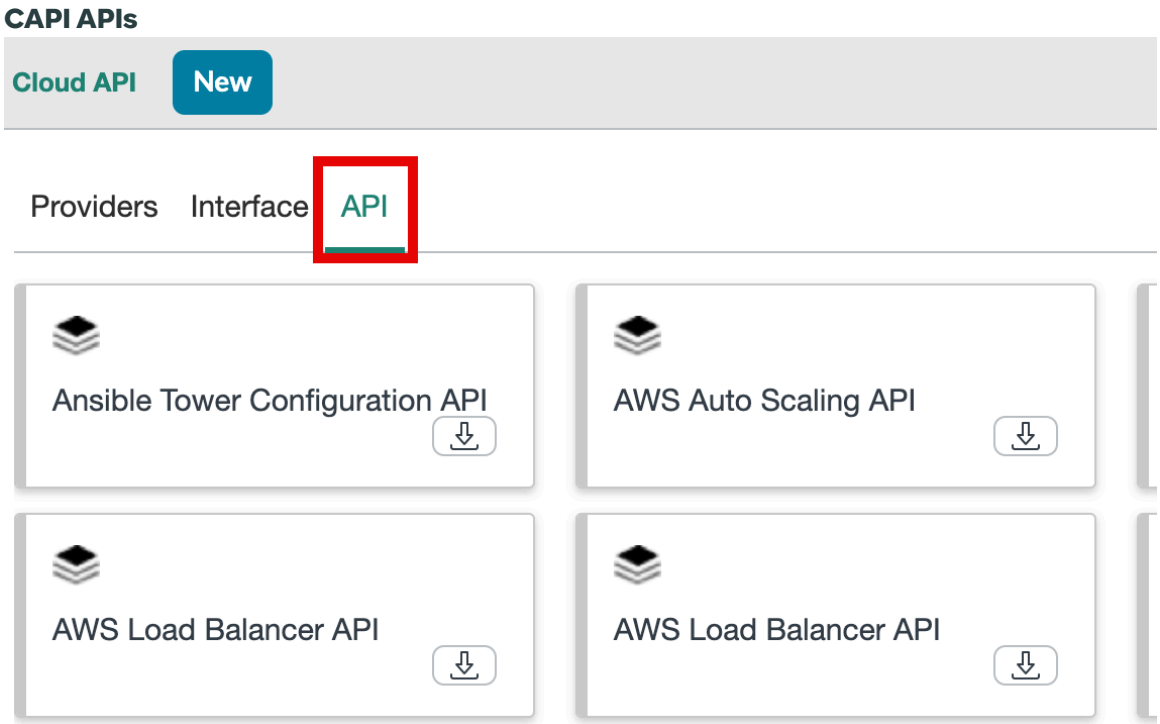

Note: You can customize scripted connectors using MID Server script includes.

API	Interface	Product	Scripted
Ansible Tower Configuration API	Configuration Management Interface	Ansible-Tower	YES*
AWS Auto Scaling API	Auto Scaling Interface	AWS Elastic Compute Cloud	YES*
AWS Blob Storage API	Blob Storage Interface	AWS Simple Storage Service	NO
AWS Block Storage API	Block Storage Interface	AWS Elastic Block Storage	NO
AWS Compute API	Compute Interface	AWS Elastic Compute Cloud	NO
AWS Compute API	Compute Interface	AWS Elastic Compute Cloud	YES*
AWS Load Balancer API	Load Balancer Interface	AWS Elastic Load Balancing	NO
AWS Network API	Network Interface	AWS Virtual Private Cloud	NO
AWS Stack Orchestration API	Stack Orchestration Interface	AWS CloudFormation	YES*
Azure Alert API	Event Interface	Azure Alert	YES*
Azure Blob Storage API	Blob Storage Interface	Microsoft.Storage	YES*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

API	Interface	Product	Scripted
Azure Block Storage API	Block Storage Interface	Microsoft.BlockStorage	YES*
Azure Compute API	Compute Interface	Microsoft.Compute	YES*
Azure Extensions API	Azure Extensions Interface	Microsoft.Compute	YES*
Azure Load Balancer API	Load Balancer Interface	Microsoft.LoadBalancer	YES*
Azure Network API	Network Interface	Microsoft.Network	YES*
Azure Stack Orchestration API	Stack Orchestration Interface	Azure ARM	YES*
Infoblox IPAM API	IPAM Interface	Infoblox	NO
ServiceNow Store IPAM API	IPAM Interface	ServiceNow	YES*
SSH API	SSH Interface	SSH	NO
vSphere Block Storage API	Block Storage Interface	VMware VDisk	NO
vSphere Compute API	Compute Interface	VMware vSphere	NO

*You can customize how scripted APIs make REST calls to the cloud provider.

CAPI classes in MID Server script includes

Cloud Provisioning and Governance provides several JavaScript classes that make REST calls to cloud providers. These classes are called from MID Server script includes (for Azure by default) or from Java calls in the system (for AWS by default). If you create custom providers, interfaces, or CAPI APIs (for any cloud provider), you can use MID Server script includes to call the classes.

Note: In CAPI APIs, there are no MID Server script includes (in Method Mappers) that you can customize to call the AWS classes. By default, the CAPI APIs for AWS make Java calls to then call the AWS classes. However, if you create CAPI APIs, you can also create custom script includes to call the classes.

View CAPI classes by navigating to **MID Servers** > **Script Includes** in the base system instance interface (not the Cloud Admin Portal). You search for the name of a class in the **Go to** search box using the **Name** column:

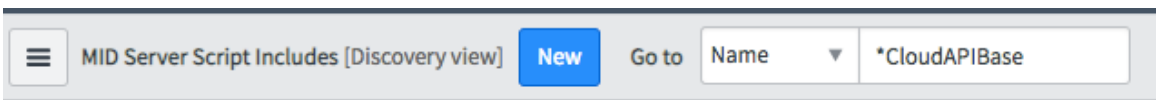

Base classes

These base classes are available by default. They form a basis for the Azure and AWS classes that make the actual REST calls.

CloudAPIBase

CloudAPIBase is a provider-agnostic API that communicates with all types of cloud APIs. It provides functions for basic debugging and logging, and for collecting proxy

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

information on the MID Server. It also defines Headers, Parameters, Endpoints, and `this.parameter`.

AWSCloudAPIBase

This class is extended from CloudAPIBase and provides functions that call the invoker to make GET, PUT, POST, and DELETE requests. The class provides the option of an `executeAction` method, which is an AWS implementation of GET, or the invoke method, which provides REST API calls.

AzureCloudAPIBase

This class is extended from CloudAPIBase and provides functions that do these actions:

- Fetch resource IDs.
- Fetch endpoint properties.
- List resources.
- Call the invoker to make GET, PUT, POST, and DELETE requests.

AzureCloudAPIBase utilizes the `_invokeRequest` method to make REST API calls to Azure using the standard REST API calls.

Base APIs

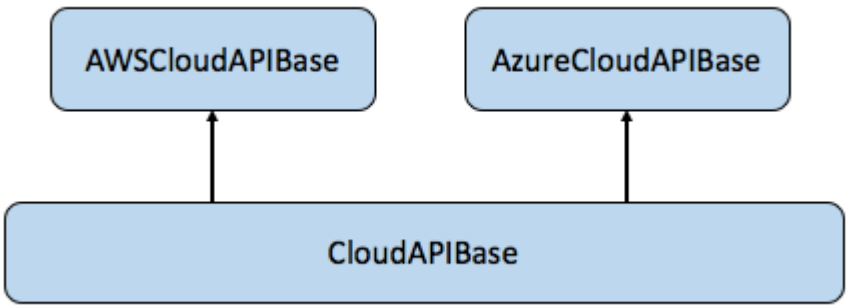

Invoker classes

Invoker classes perform the actual REST calls to the cloud provider. If you create a custom CAPI interface and API, you should extend *`CloudRestAPIInvoker`* to create your REST calls.

CloudRestAPIInvoker

This class is an abstracted class for making REST API calls. If there are headers and/or parameters to send with a request, classes that implement the invoker must extend this object. They also must provide overrides for *`getEndpointUrl`* to return the targeted endpoint, and for *`getHeaders`* and *`getParameters`*. For example, with Azure REST APIs, authentication requires the passing of a token as a header. The extension implementing it should provide a response from *`getHeaders`* that performs the token creation or retrieval and sets it. This base object invokes *`getHeaders`* and sends whatever it gets back with the request. This class also provides a *`getResponse`* method to handle the response from the cloud provider.

If you create a custom CAPI API, you should extend *`CloudRestAPIInvoker`* to make REST calls to the cloud provider.

AmazonWebServicesAPIInvoker

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This class is extended from *CloudRestAPIInvoker* and provides methods for making REST API calls to the cloud provider.

AzureAPIInvoker

This class is extended from *CloudRestAPIInvoker* and provides functions that make these REST API calls to Azure Resource Manager:

- Sets the Service Principal, headers, and parameters for REST calls.
- Gets the endpoint URL, token (client ID, tenant ID, and secret key), and headers/parameters.
- Retrieves and decodes the JSON response body from the cloud provider.

AzureEAAPIInvoker

This class returns the token for access to the Azure Enterprise Agreement (EA) REST API.

AzureStorageAPIInvoker

This class makes REST calls for Azure storage. The functions in this class can get the token, response, headers, and parameters.

This graphic illustrates how these base classes are used to extend AWS-specific and Azure-specific classes:

Invoking REST calls

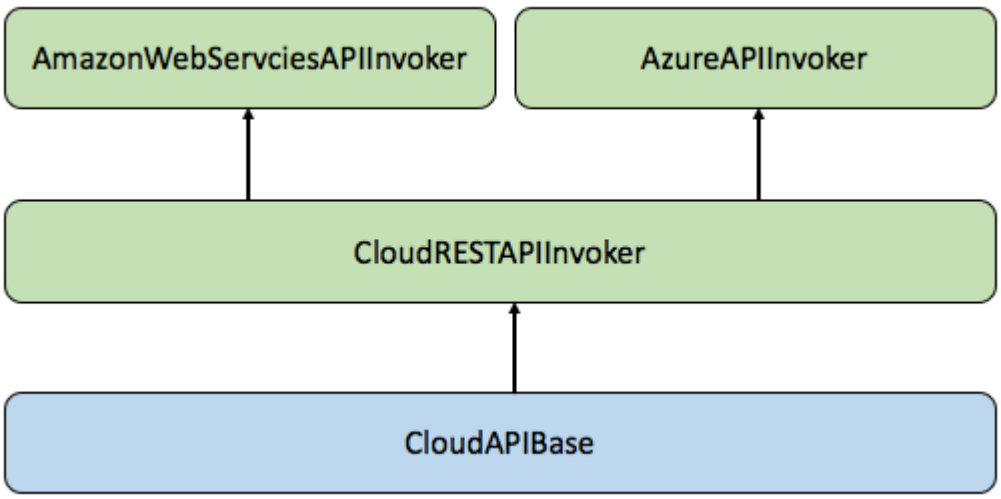

Exceptions

RESTAPIException

This base class handles exceptions that might occur during the REST call. It returns the response code from the cloud provider.

AWSAPIException

This class handles exceptions that might occur during the REST call.

AzureAPIException

This class throws a runtime exception when needed during Azure API invocation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Handling exceptions

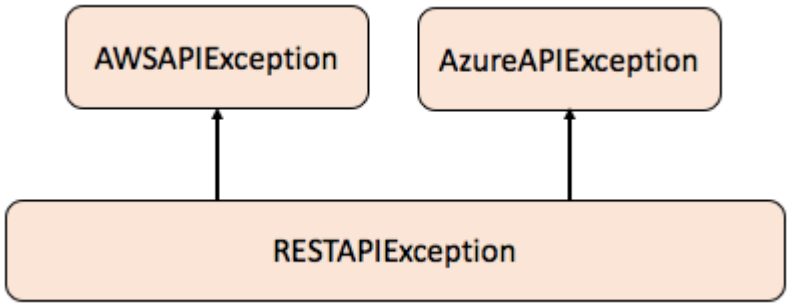

Classes for credentials

CloudRESTRequestSigningUtil

This class uses overrides that are specified in CAPI API Config Overrides. The *AmazonWebServicesRequestSigner* class extends this class. For Azure, many signers are built into the MID Server script includes for Method Mappers in Azure CAPI APIs.

AmazonWebServicesRequestSigner

This class sets up a secure session for the REST API call by taking information about a request, and creating a signature v4 -signed string for the request. This class is extended from CloudRESTRequestSigningUtil.

Other Azure classes

Classes that extend AzureCloudAPIBase

Class	Description
AzureAlertRule	Class to create and delete Azure alert rules.
AzureAuthorizationPolicy	Class to assign tag policies.
AzureComputeAvailabilitySet	Class to fetch Azure availability sets. Refer to: https://docs.microsoft.com/en-us/rest/api/compute/availabilitysets
AzureComputeHardware	Class to fetch Azure sizes. Refer to: https://docs.microsoft.com/en-us/rest/api/compute/virtualmachinesizes
AzureComputeImage	Class to handle Azure API calls against image. Refer to: https://docs.microsoft.com/en-us/rest/api/manageddisks/images/images-rest-api
AzureComputeVirtualMachine	Class to handle Azure API calls against virtual machines. Refer to: https://docs.microsoft.com/en-us/rest/api/compute/virtualmachines VM Size and Capacities: https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-windows-sizes https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-linux-sizes Cloud-Init supported Versions: https://docs.microsoft.com/en-us/azure/virtual-machines/linux/using-cloud-init
AzureDeployment	Class to handle Azure ARM template deployment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classes that extend AzureCloudAPIBase (continued)

Class	Description
AzureDeploymentOperation	Class to handle Azure ARM template deployment operation.
AzureImages	Class to handle Azure Images.
AzureLoadBalancer	Class to handle Azure API calls against load balancers. Refer to: https://msdn.microsoft.com/en-us/library/azure/mt163651.aspx
AzureLocalNetworkGateways	Class to create and retrieve local network gateways.
AzureNetworkNIC	Class to handle Azure API calls against network interface cards. Refer to: https://msdn.microsoft.com/en-us/library/azure/mt163579.aspx
AzureNetworkPublicIP	Class to handle Azure API calls against public IP addresses. Refer to: https://msdn.microsoft.com/en-us/library/azure/mt163638.aspx
AzureNetworkRouter	Class to handle Azure API calls against route table. Refer to: https://docs.microsoft.com/en-us/rest/api/network/route-tables
AzureNetworkSecurityGroup	Class to handle Azure API calls against network interface cards. Refer to: https://msdn.microsoft.com/en-us/library/azure/mt163579.aspx
AzureNetworkVirtualNetwork	Class to handle Azure API calls against virtual networks.
AzureProvider	Class to get Azure provider details, such as the provider URL and version.
AzureRegion	Class to fetch Azure regions.
AzureResource	Deals with the object referenced by ObjectID in the parameters. It also fetches the API version. This class is Intended mostly to get unsupported types or delete resources.
AzureResourceGroupManager	Class to handle Azure ARM resource groups.
AzureResponseUtil	Class to translate Azure API response to DTO object.
AzureStorageAccount	Class to handle Azure API calls against storage account. Refer to: //https://msdn.microsoft.com/en-us/library/azure/mt163638.aspx
AzureStorageBlob	Class to create and manipulate storage containers and blobs.
AzureStorageManagedDisk	Class to handle Azure API calls against managed disk. Refer to: https://docs.microsoft.com/en-us/rest/api/manageddisks/disks/disks-create-or-update
AzureSubscription	Class to fetch Azure subscriptions.
AzureVirtualNetworkGateways	Class to create and manipulate virtual network gateways.
AzureCloudEAAPIBase	Base object for Azure EA API handlers. Azure EA APIs have a base URL of https://consumption.azure.com and require the enrollment number in the path. This

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classes that extend AzureCloudAPIBase (continued)

Class	Description
	class handles marshalling of the EnrollmentNumber and EnrollmentToken parameters.

Other AWS classes

Class	Description
AmazonImages	Class to handle Amazon machine images. http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html
AmazonResponseProcessorUtil	Class to handle AWS REST API response.
AWSAutoScalingGroups	Class to handle AWS API calls to create and manipulate Auto Scaling Groups.
AWSCloudFormation	Handler for AWS CloudFormation API calls.
AWSEC2	Handler for AWS EC2 API calls.
AWSKeyValuePair	Class to import and manipulate key pairs.
AWSLaunchConfigurations	Helper class.
AWSS3	Handler for AWS S3 API calls.
AWSTopics	Class with AWS topics.

Azure VM provisioning walkthrough

This example walks you through the components of Cloud Provisioning and Governance that function during the provisioning of a virtual machine in an Azure datacenter. Topics covered include blueprints, resource blocks, the Cloud API (CAPI), and MID Server script includes.

Before you begin

Role required: admin

About this task

This walkthrough starts with a Linux VM that a user already provisioned in Azure. Next, it walks you through the blueprint with the VM, the resource blocks, and then the CAPI calls specified from the resource blocks. Finally, the walkthrough shows you how a MID Server script include eventually, through several CAPI calls, makes the actual REST calls to the Azure API.

This example uses default resource blocks and script includes that are available in your instance. Therefore, although you might not have a provisioned VM on your instance, you can still follow these steps and view the components used in this example to understand how the components work.

For an example of a VM in AWS, see AWS VM provisioning walkthrough.

Note: The terms virtual machine, VM, and virtual server are used interchangeably in this example.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Look at a provisioned VM in the Cloud User Portal, and focus on some of the details about the VM:

   a. In the base system instance interface, open the Cloud User Portal (**Cloud User Portal**).

   b. On the Cloud User Portal, view a provisioned VM in a stack, such as this VM, by clicking **Stacks**, and then clicking the name of the stack.

      In this example, the stack is named **MyAzureVM**.

   c. Under **Resources**, click the VM in the stack.

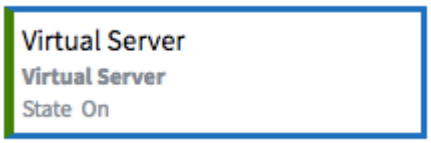

   d. View the properties of the VM, and notice that it is in an **Azure Datacenter** in region **eastus**.

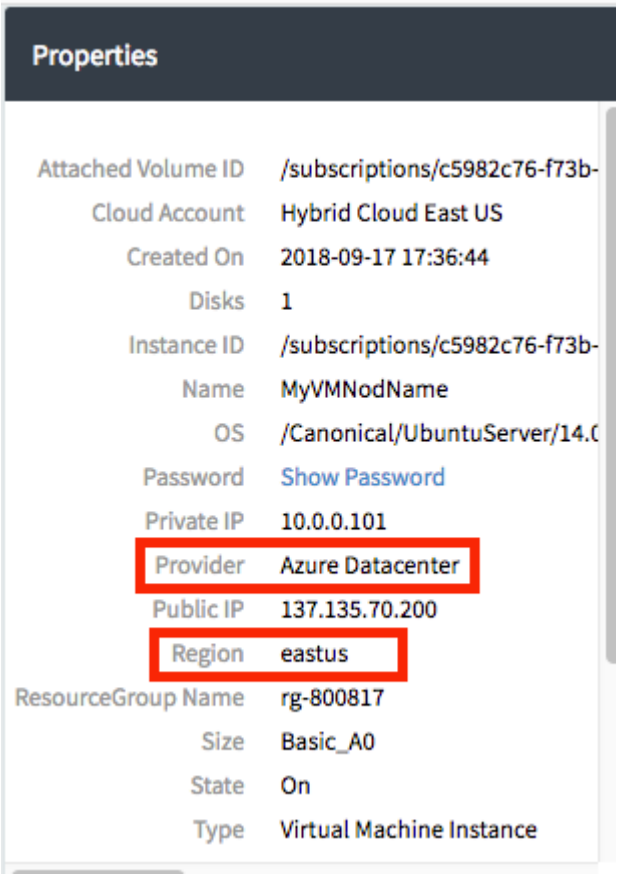

2. Look at the blueprint on which the VM is based:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **Design** > **Blueprints**, and then open a blueprint with a virtual server in the Azure datacenter. The **Deployment Model** tab appears by default, showing you the various components of the blueprint.

This example blueprint has three components: the container, the virtual server, which is the actual VM that is provisioned, and the Azure datacenter.

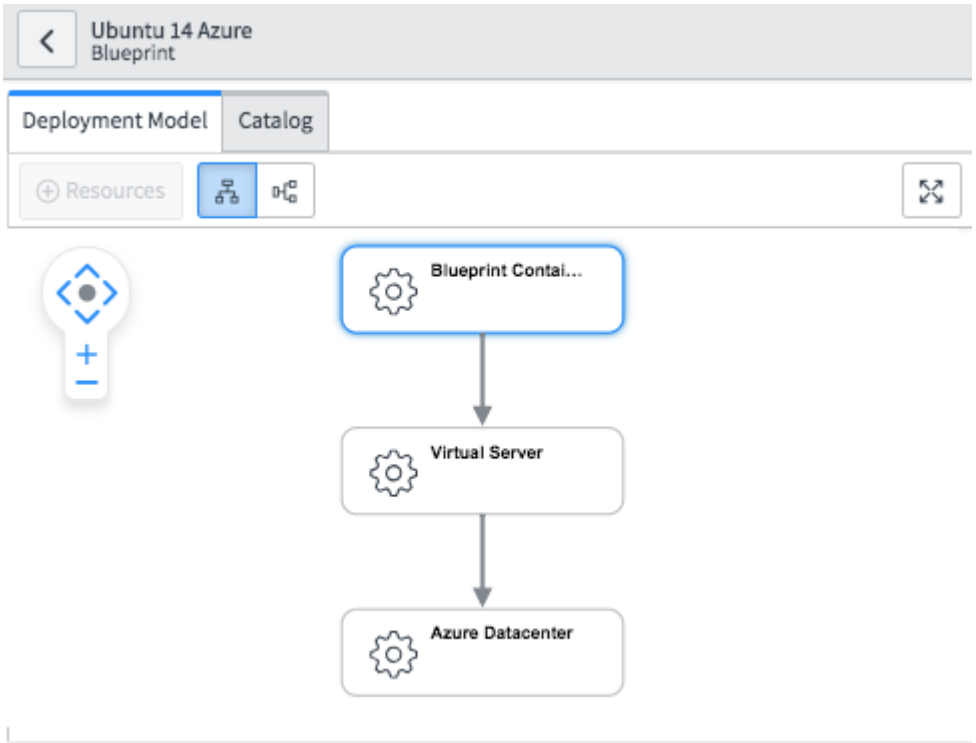

b. Click the **Operations** tab on the bottom, and then click **Provision**.

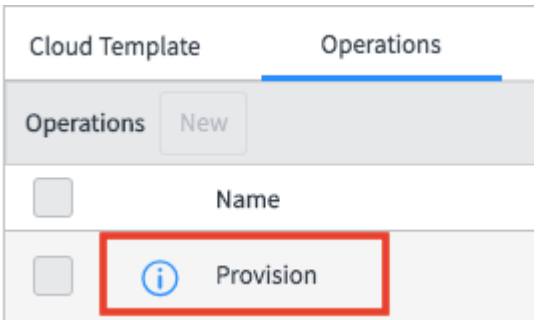

The Provision operation is the operation that the system triggered when it created the VM. Other default operations are available, but this example focuses on the Provision operation.

c. Click the **Provision** block for Blueprint Container Resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

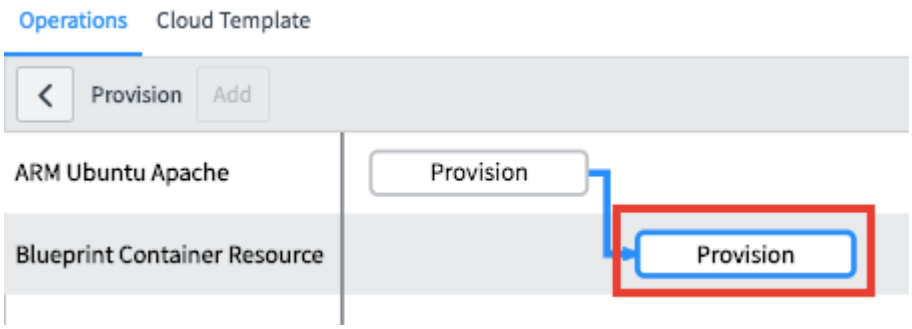

d. On the right, notice that one of the parameters in the Inputs list is Location.
This parameter holds the value eastus, which is where the VM lives in the datacenter.

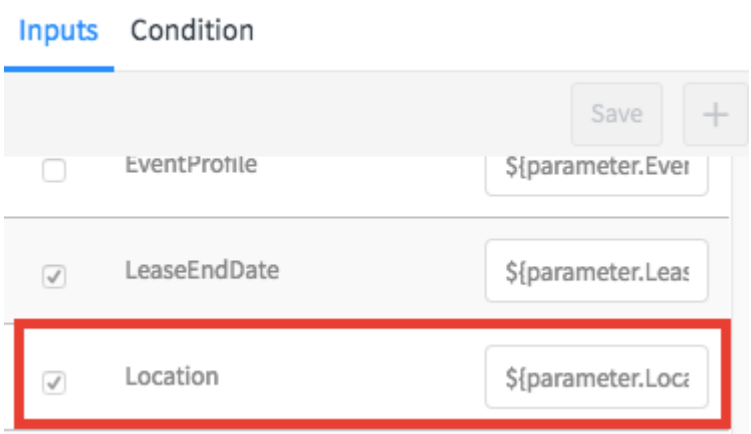

Inputs can be specified on the container, as it is in this example, or on any other resource block. By default, the *Location* parameter is already specified for you in the Blueprint Container resource block, so that you can use it in every blueprint like this one. If you switch the blueprint to **Draft**, you can add more parameters to the Blueprint Container resource block. You cannot add parameters to the Virtual Server resource block. For this walkthrough, no additional parameters are necessary.

3. View the Virtual Server and Datacenter resource blocks used in this blueprint:

a. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.

b. Search for and open the **Virtual Server** resource block, which is provided by default with Cloud Provisioning and Governance.

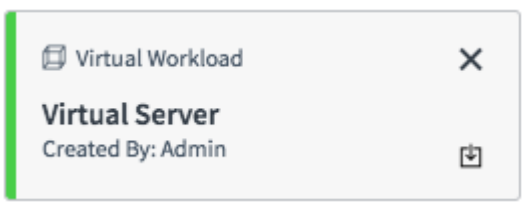

c. On the Resource Block form, click the **Host Resource** related list to view the list of supported host resource blocks.

d. Notice the host interface and host resource that are already specified:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

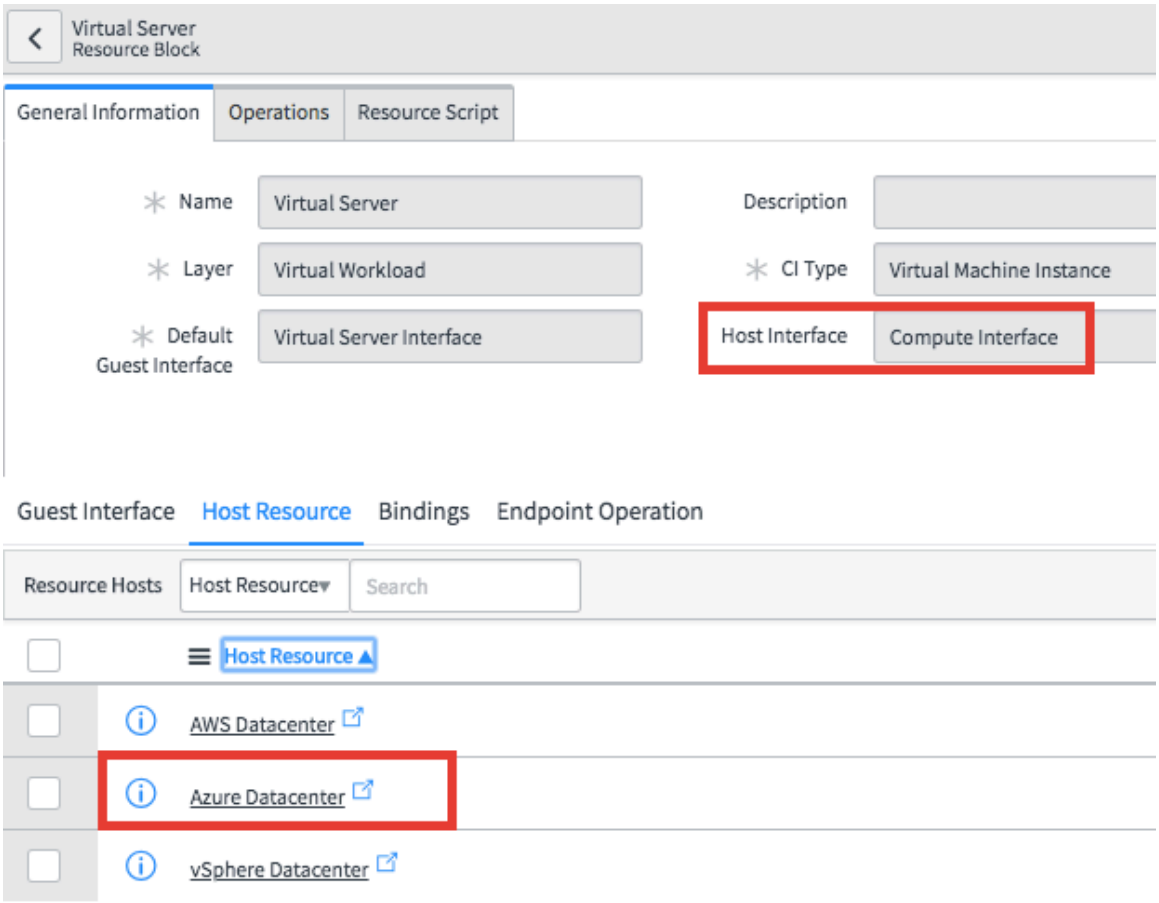

- The **Host Interface** field specifies the interface that must match the guest interface in the resource block that hosts this virtual machine. In this case, the host interface is the Compute Interface, which is also the guest interface on the Azure Datacenter resource block. By default, the datacenter resource blocks provide several guest interfaces that other resource blocks like virtual storage can use to connect to the datacenter.
- The **HostResource** column in the **Host Resource** related list already specifies **Azure Datacenter**, which means that the Azure datacenter resource block is a valid host for this VM.

e. Click the **Operations** tab, and then click the **Steps** subtab.

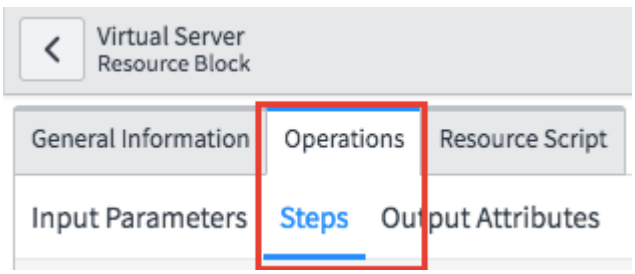

f. Select **Provision** from the **Operation** list.

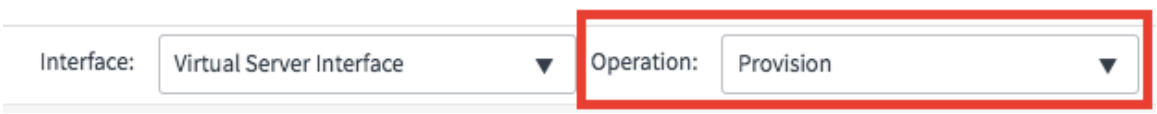

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remember that the Provision operation is the operation that the system used to create the VM. Other default operations are available, but this example focuses on the Provision operation.

g. Notice the step that appears in the list and the full step description that appears above the input parameters:

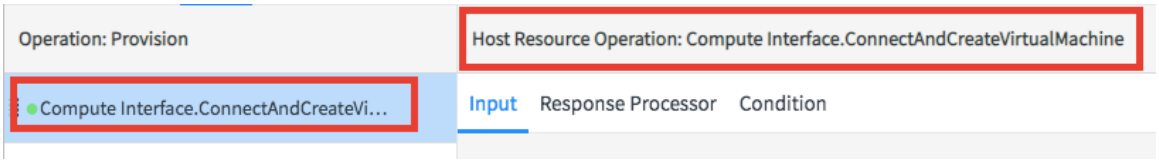

- **Host Resource Operation:** indicates that this step calls an operation on the host resource (the Azure datacenter in this example).
- **Compute Interface** is the guest interface on the Azure datacenter that this step is using.
- **ConnectAndCreateVirtualMachine** is the operation in the Azure datacenter that this step calls.

Note: In this case, the step calls an operation from another resource block: the Azure datacenter. Steps can also call CAPI directly, and then CAPI can execute REST calls to the cloud provider API. You can see that when you look at the Azure datacenter resource block.

h. Navigate back to **Design** > **Resource blocks**.

i. Open the **Azure Datacenter** resource block, which is the host resource block that the virtual server is connected to.

j. Notice the supported guest interfaces in the **Guest Interface** related list:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

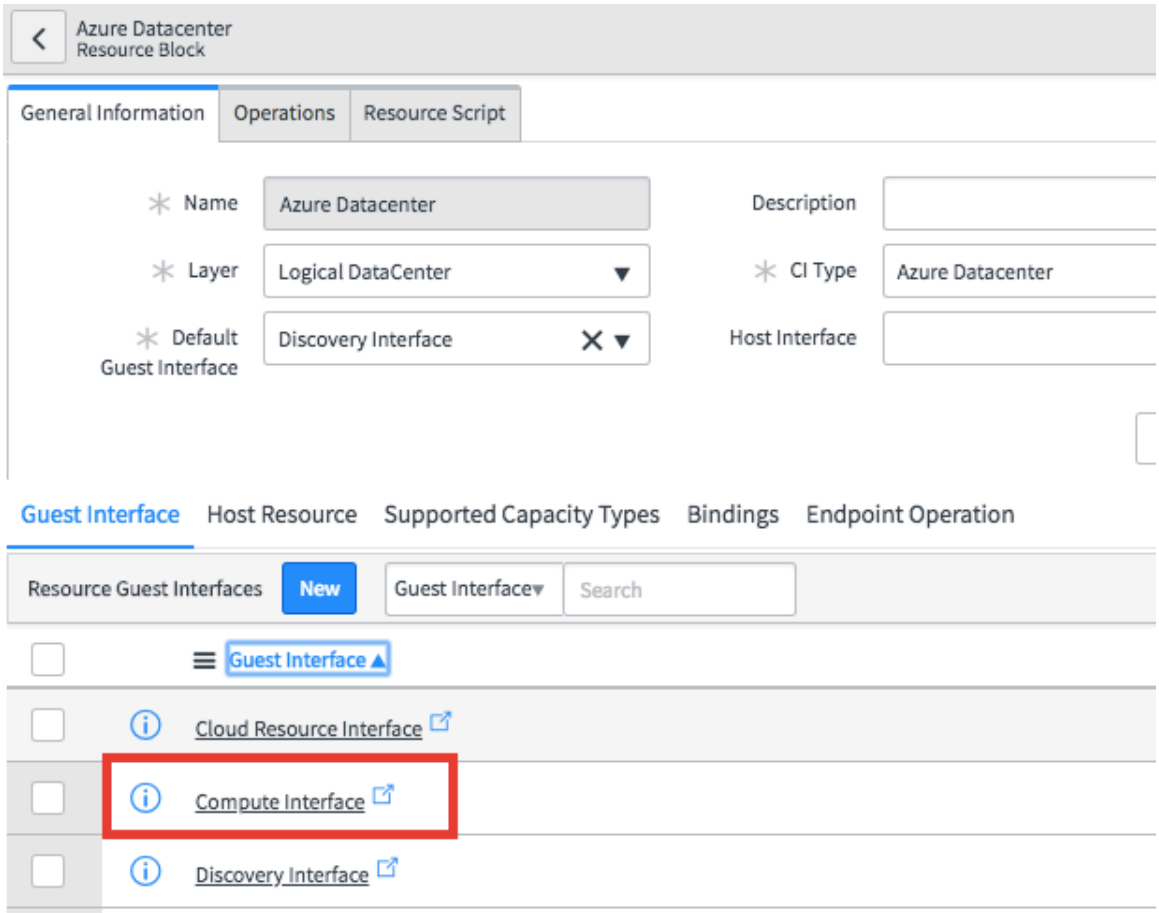

These guest interfaces are the interfaces that the Azure datacenter makes available to other resource blocks. The **Compute Interface** is provided so that the Virtual Server resource block, which specifies the Compute Interface as its host interface, can connect to the datacenter.

k. Click the **Operations** tab, and then click the **Steps** subtab.

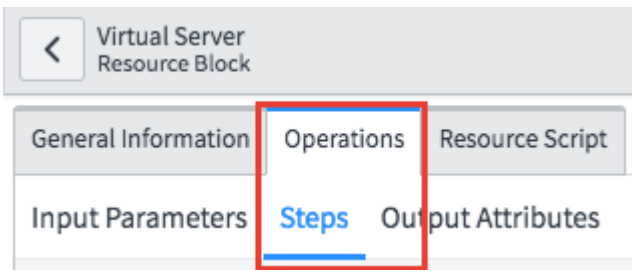

l. In the **Interface** list, select **Compute Interface** if it is not already selected.

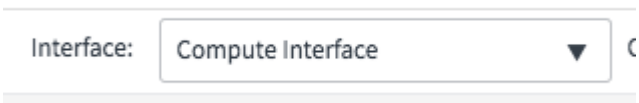

Remember that this interface is specified in the Virtual Server resource block.

m. In the **Operation** list, search for and select **ConnectAndCreateVirtualMachine**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

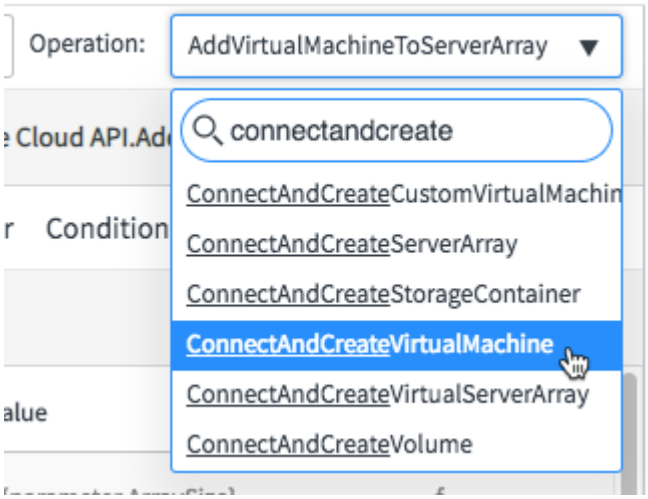

Remember that this operation is specified in the Virtual Server resource block.

n. Notice the CAPI call that is used in the only step for the ConnectAndCreateVirtualMachine operation:

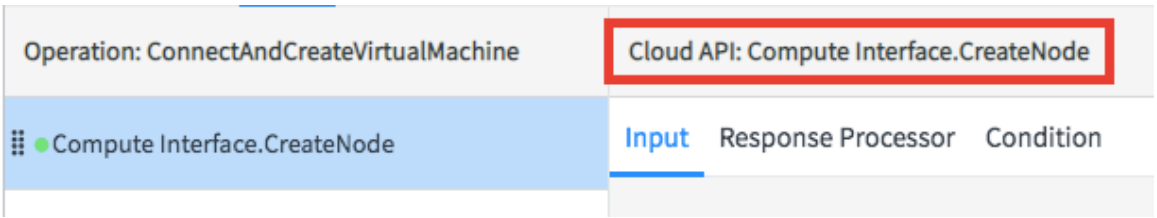

- **Cloud API:** indicates that this step calls CAPI, so that CAPI can execute a REST call to the cloud provider, which in this case is Azure.
- **Compute Interface** specifies the CAPI interface that this step calls.
- **CreateNode** indicates the method that is executed. As the name suggests, this method tells the cloud provider to create the virtual machine.

o. To open the step configuration: With the resource block in the **Draft** state, point to the highlighted (blue) step, and then click the **Edit Step** icon that appears.

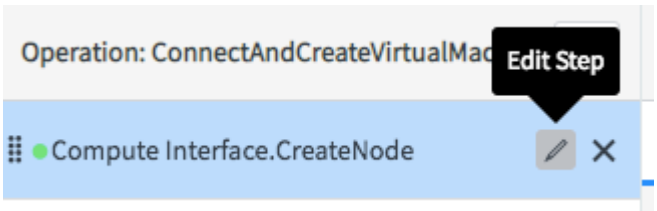

p. Look at the step configuration and notice the settings that integrate with CAPI:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

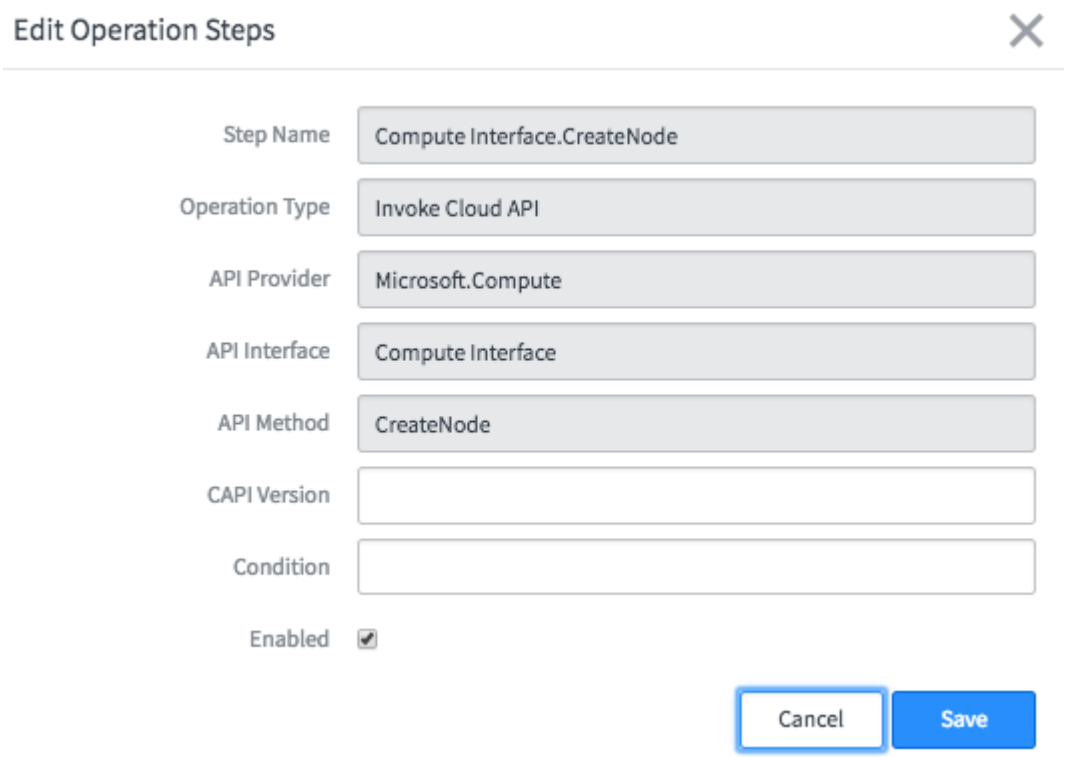

Field	Description
Operation Type	**Invoke Cloud API** specifies that this step should call the CAPI via the specified provider, interface, and method.
API Provider (which is actually the product)	**Microsoft.Compute** is a product (not the actual provider) that belongs to the Azure provider as defined in CAPI.
API Interface	**Compute Interface** is a product that belongs to the Azure provider as defined in CAPI.
API Method	**CreateNode** is the method that calls Azure to create the VM.

Note: The CAPI API definition ties together the provider (Azure), the product (Microsoft.Compute), the interface (Compute Interface), and the method (CreateNode).

q. Close the window.

r. With the **Compute Interface.CreateNode** step selected, click the **Response Processor** tab, and notice the **Create_Virtual_Server_Response_Processor** script.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

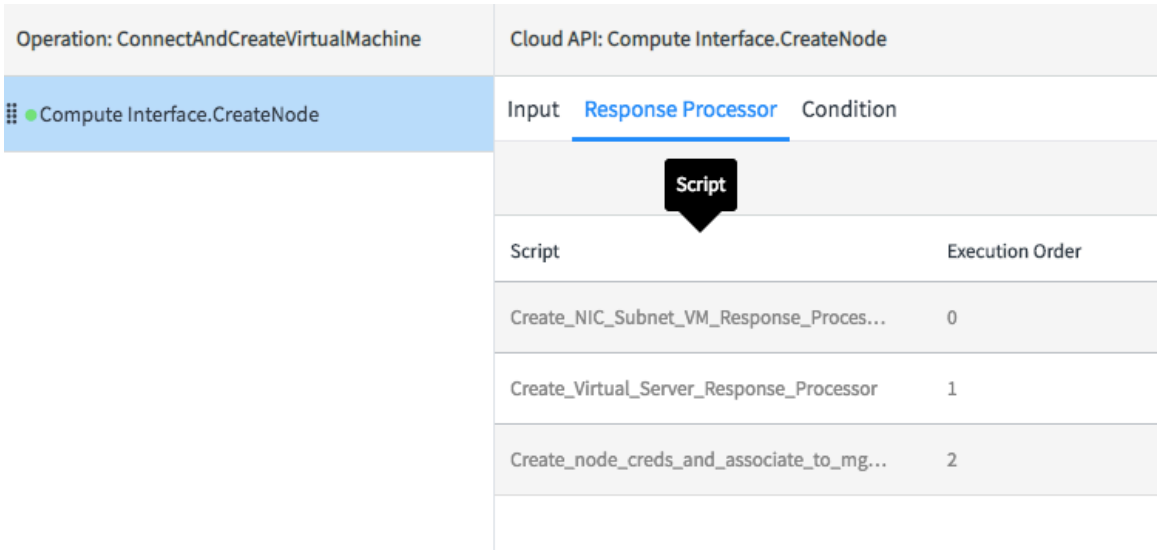

This script is the response processor that updates the CMDB in your instance after the virtual machine is created.

s. View an explanation of the script and the example that is a part of the topic at Create a Response Processor.
When you finish reviewing that topic, return to this topic.

4. Look at the CAPI components that work together to provision the VM:

a. In the Cloud Admin Portal, navigate to **Design** > **Cloud API**.

b. Click the **API** tab and then search for an open **Azure Compute API**.

c. Look at how this CAPI API ties together an interface and a product:

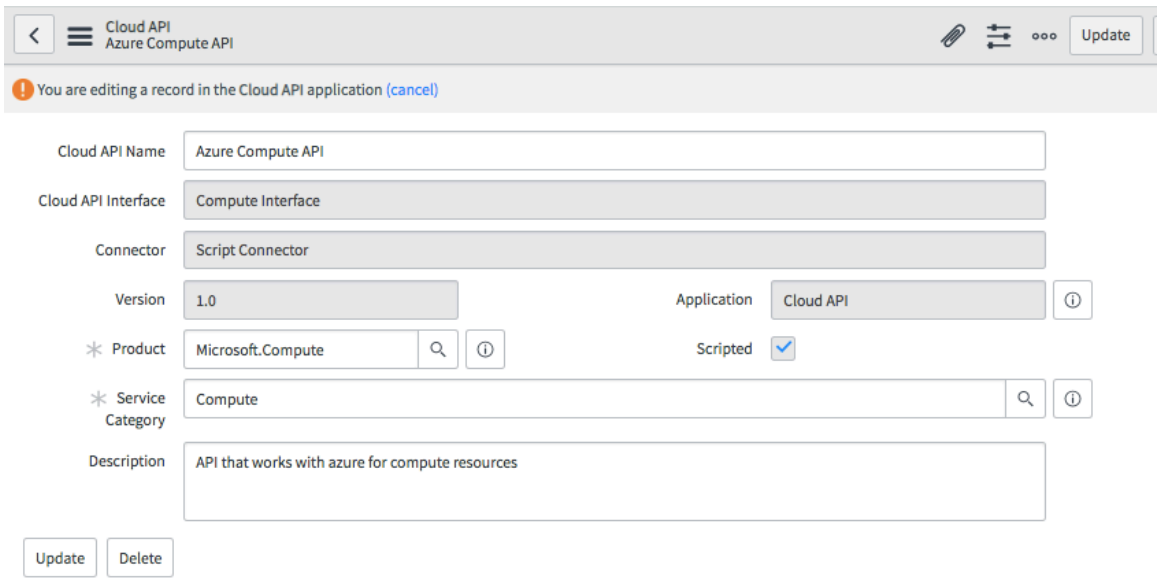

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Cloud API Interface	**Compute Interface** is the same interface that you can see specified in the resource block. The interface contains the definition for methods, including the *`CreateStack`* method.
Connector	**Script Connector** indicates that this CAPI provides a MID Server script include for each operation. The script include calls the API. You can modify the script include or even create a new one.
Version	**1.0** indicates the version of the API. You could have multiple versions of this API with different version numbers. Remember that in the step in the datacenter resource block that creates the VM, a Version field is provided. Although it was blank in this example because there is only one version of this API, you can specify a different version number.
Product	**Microsoft.Compute** is the product that belongs to the provider in CAPI. This product is provided by default with Cloud Provisioning and Governance, and is one of the most commonly used products for creating VMs on Azure.

d. In the CAPI Method Mappers related list, click the record preview icon for the *`CreateNode`* record, and then click **Open Record**.

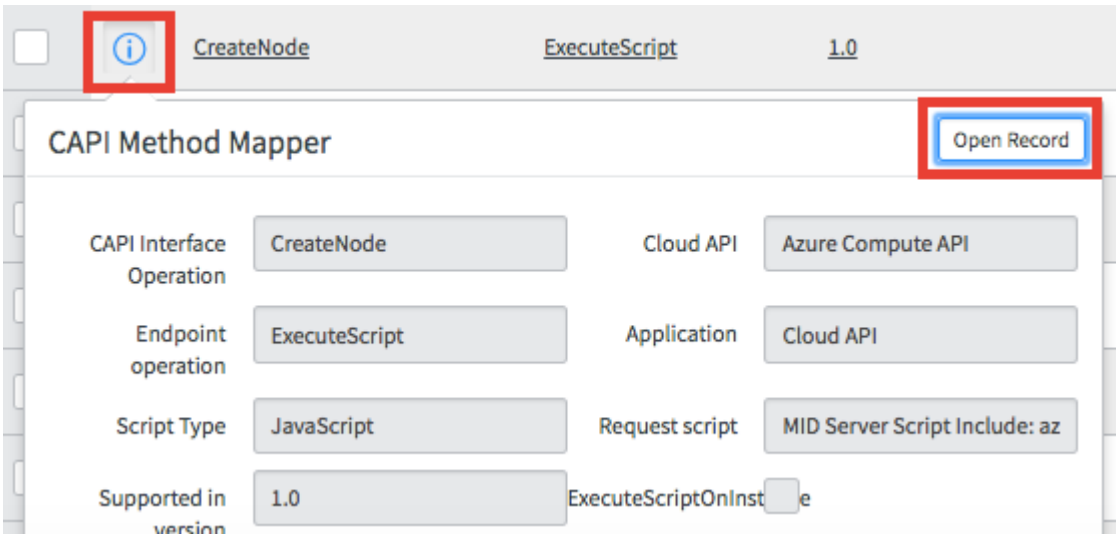

e. Look at the *`CreateNode`* method mapper:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

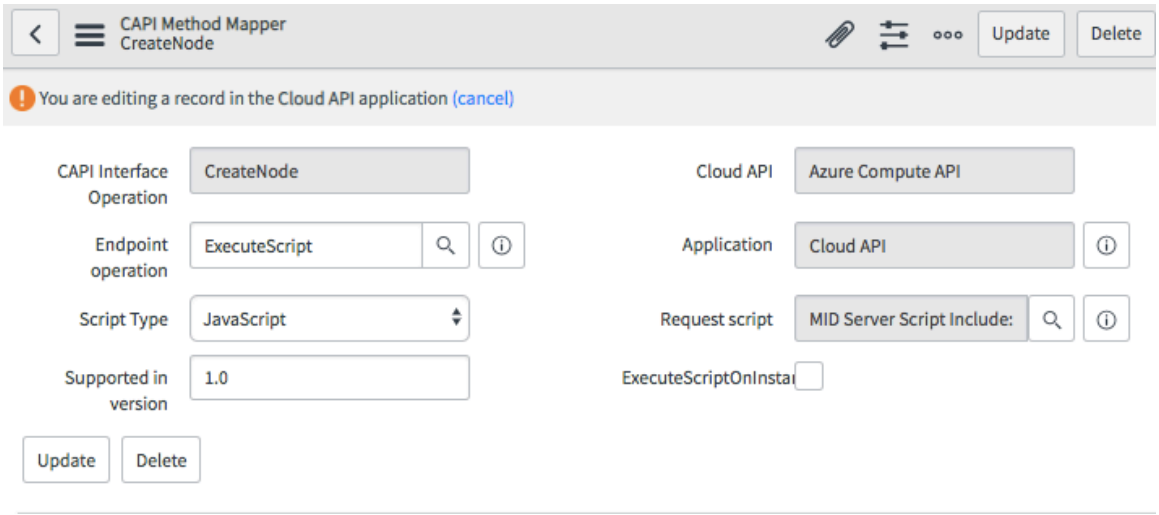

Notice that the **Endpoint operation** is **Execute Script**, which indicates that a customizable script include is used to call functions in CAPI, which eventually makes the REST calls to Azure. The *azure-compute-1.0-CreateNode* MID Server script include is specified in the **Request script** field.

This graphic illustrates how the components — from the resource block to the script include — are connected:

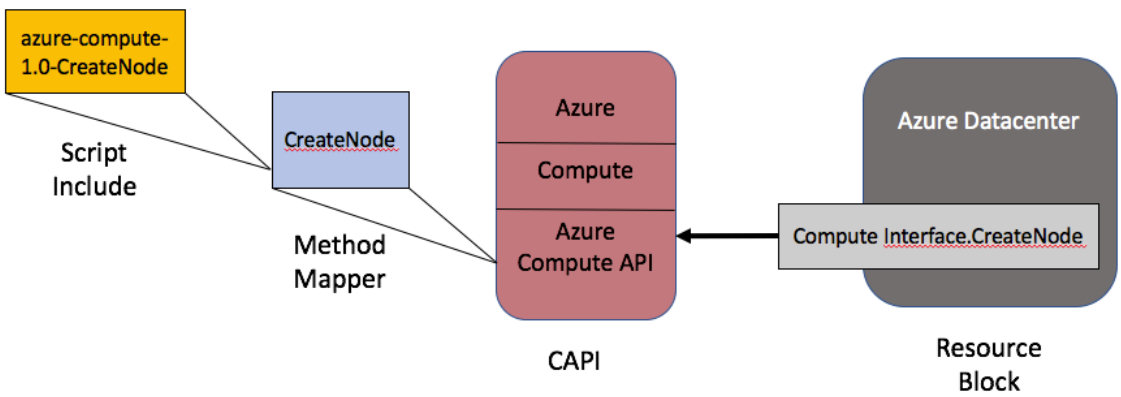

f. Click the record icon (ⓘ) for the **Request script**.

The *azure-compute-1.0-CreateNode* MID Server script include opens.

g. Review the contents of the script.

```
createNode();
function createNode(){
   var acp = new AzureComputeVirtualMachine(this.parameters,
 this.headers);
   var vm = acp.createVirtualMachine();

   return vm;
}
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The script calls the *createVirtualMachine* function in the *AzureComputeVirtualMachine* class. It returns the virtual machine that is created. The *AzureComputeVirtualMachine* class extends *AzureCloudAPIBase*, which is one of the base APIs in Cloud Provisioning and Governance.

h. To see the *AzureComputeVirtualMachine* class, which is a MID Server script include, open the base system instance interface (not the Cloud Admin Portal), and then navigate to **MID Server** > **Script Includes**.

i. Search for and open *AzureComputeVirtualMachine*.

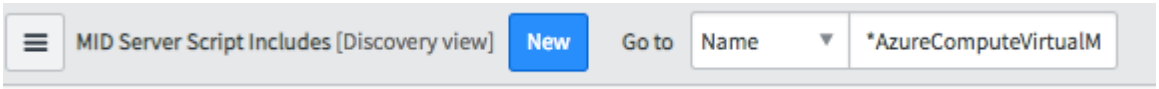

j. Review the `createVirtualMachine` function:

The function defines many variables that are required to create a VM, including the *Location* (datacenter) and the name of the VM, and so on.

```
var location = this.parameters.get('Location');
var vmName = this.parameters.get('NodeName');
...
```

This line of code calls the `_performVMOperation` function in this same script include and passes in the parameters that are necessary. Notice that an empty value `''` is passed for the `action` parameter:

```
var response = this._performVMOperation(subscriptionId,
 resourceGroup, vmName, '', 'PUT', jsonObj.encode(params));
```

The `_performVMOperation` function calls the function `_invokeRequest`, which is located in the *AzureCloudAPIBase* class:

```
return this._invokeRequest(this.pathDefault, param, method,
requestBody);
```

The PUT method for the REST call is specified in the *method* parameter.

Look at the *AzureCloudAPIBase* MID Server script include. The `_invokeRequest` function first calls the `_getEndpoint` function to get the endpoint and the `_getAPIInvoker` function to identify the *AzureAPIInvoker* class as the class that invokes the actual PUT call:

```
var apiInvoker = this._getAPIInvoker(apiEndpoint,
this.apiVersion);
```

AzureAPIInvoker is extended from *CloudRESTAPIInvoker*, which provides the base functions for all invokers. Once the correct invoker class is identified, a *response* variable is created to hold the value returned from `invokePut` function in the *CloudRESTAPIInvoker*.

```
response = apiInvoker.invokePut(requestBody);
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Look at the *CloudRESTAPIInvoker*, which is the base class for all invokers. The *invokePut* function specifies the PUT method and calls `getRequest`, and it returns the response from the *getResponse* function, which gets a status code and the response from Azure. This invoker class is where the connection to Azure is made.

Return to the *AzureComputeVirtualMachine* MID Server script include and find these lines:

```
var response = this._performVMOperation(subscriptionId,
 resourceGroup, vmName, '', 'PUT', jsonObj.encode(params));
this._compareVMState(response.id, 'PowerState/running');
var nodeResponse = this.getVirtualMachine(response.id);
```

Remember that the variable *response* is returned from the call to `_performVMOperation`, which has the response from the API invoker that triggered the PUT method. The *nodeResponse* variable holds the response from the `getVirtualMachine` function, which is passed the ID of the virtual machine that is in the response. The `getVirtualMachine` function calls `_getVMInfo`, which calls the same `_invokeRequest` function that was called earlier. This call is another connection to the invoker, which performs the actual REST calls. In this case, the REST call is a GET REST call to Azure to obtain the identify of the VM.

The `return nodeResponse;` line returns *nodeResponse* in the `createVirtualMachine` function, which was called from the MID Server script include on the *createNode* operation.

This graphic illustrates the classes involved:

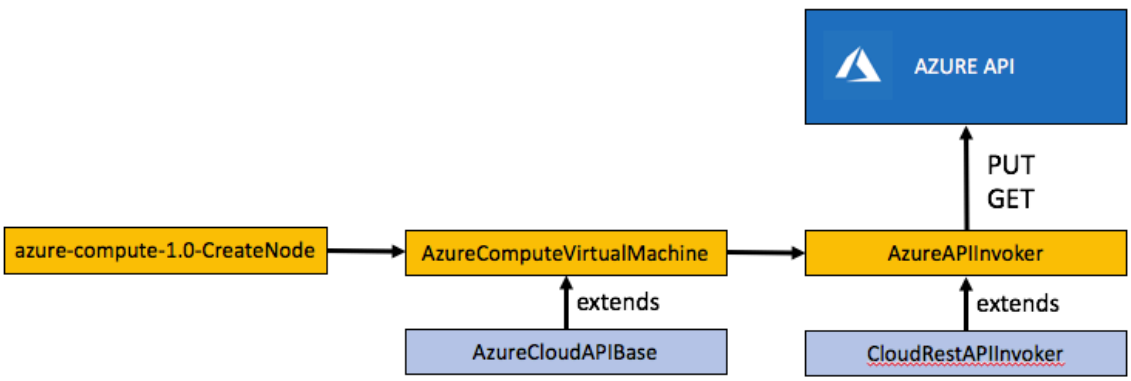

k. Navigate back to the Azure Compute API form.

l. Click the **API Config Overrides** related list, and review the items that are required for authentication.

Config Parameter and Override Value	Description
ClientID $(CloudCredential.client_id)	The client ID is necessary to identify the application in Azure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Config Parameter and Override Value	Description
EndpointURL $(CloudCredential.URL)	The endpoint URL is the authorizing input in Azure.
SecretKey $(CloudCredential.secret_key)	The secret key is used to authorize the client.
TenantID $(CloudCredential.tenant_id)	The tenant ID is the globally unique identifier for your tenant.
EnrollmentNumber $(CloudCredential.sn_cmp_ea_credential.enrollment_number) EnrollmentToken $(CloudCredential.sn_cmp_ea_credential.access_key)	Azure enrollment number and the token to access the enrollme

See Store the Azure service principal credentials in the instance for more information on the Azure information that is used in these overrides.

m. Navigate back to **Design** > **Cloud API**, and then click the **Interface** tab.

n. Search for and open **Compute Interface**.

o. Review the contents of the interface.

Notice that the interface provides REST response structures for methods like *`CreateNode`*. You typically do not need to modify existing interfaces.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

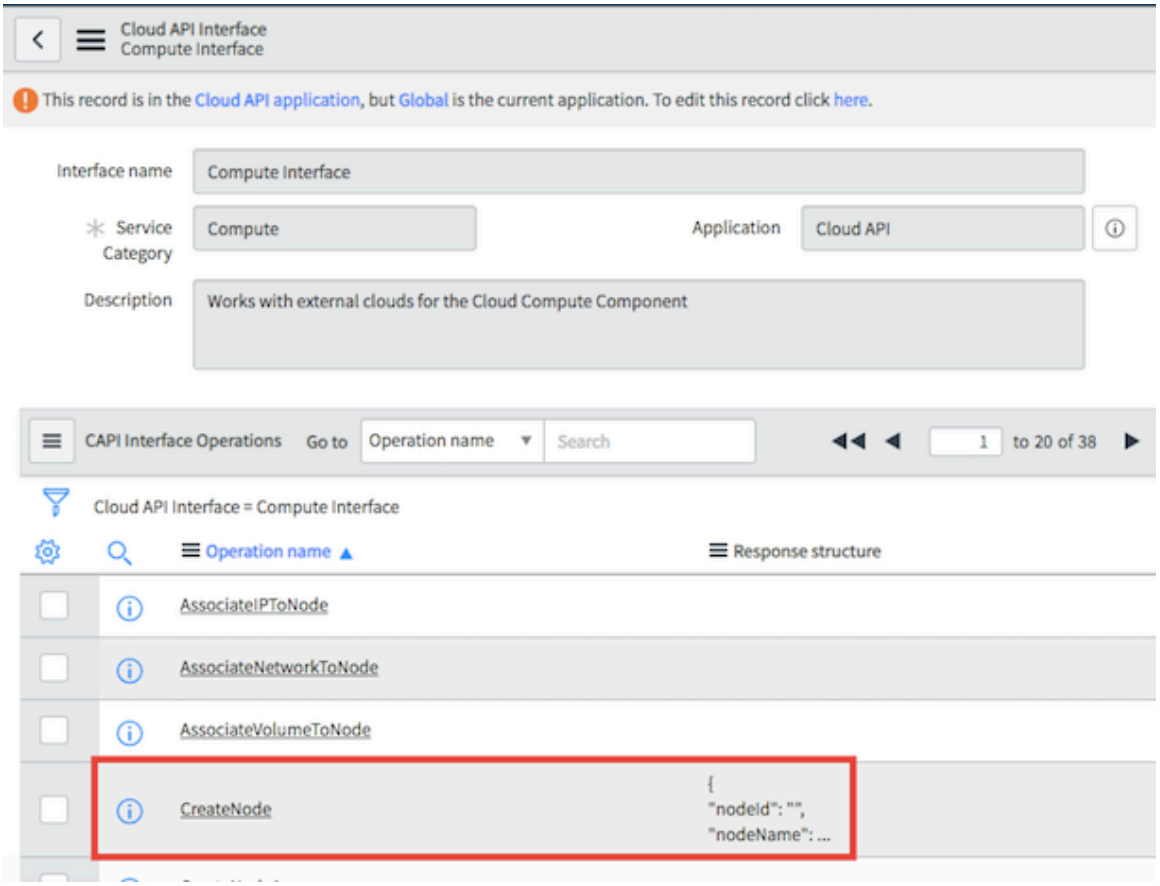

Note the service category and the operations:

Field or related list	Description
Service Category	The service category classifies the interface. The category for the Compute Interface is also **Compute**.
CAPI Interface Operations	The interface operations define the JSON structure for the REST call and the parameters that are required for the interface.

p. Click the **CreateNode** CAPI interface operation.
This operation is the operation that provides the framework for creating the EC2 virtual server.

q. Review the components of the operation:

Field or related list	Description
Response structure	The response structure is the JSON framework for REST API call. It provides a list of attributes that the provider can use to create the virtual resource with empty values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field or related list	Description
Interface Operation Parameters	These parameters are also the parameters that the CAPI interface needs from the system so it can pass it in the REST call to the cloud provider.

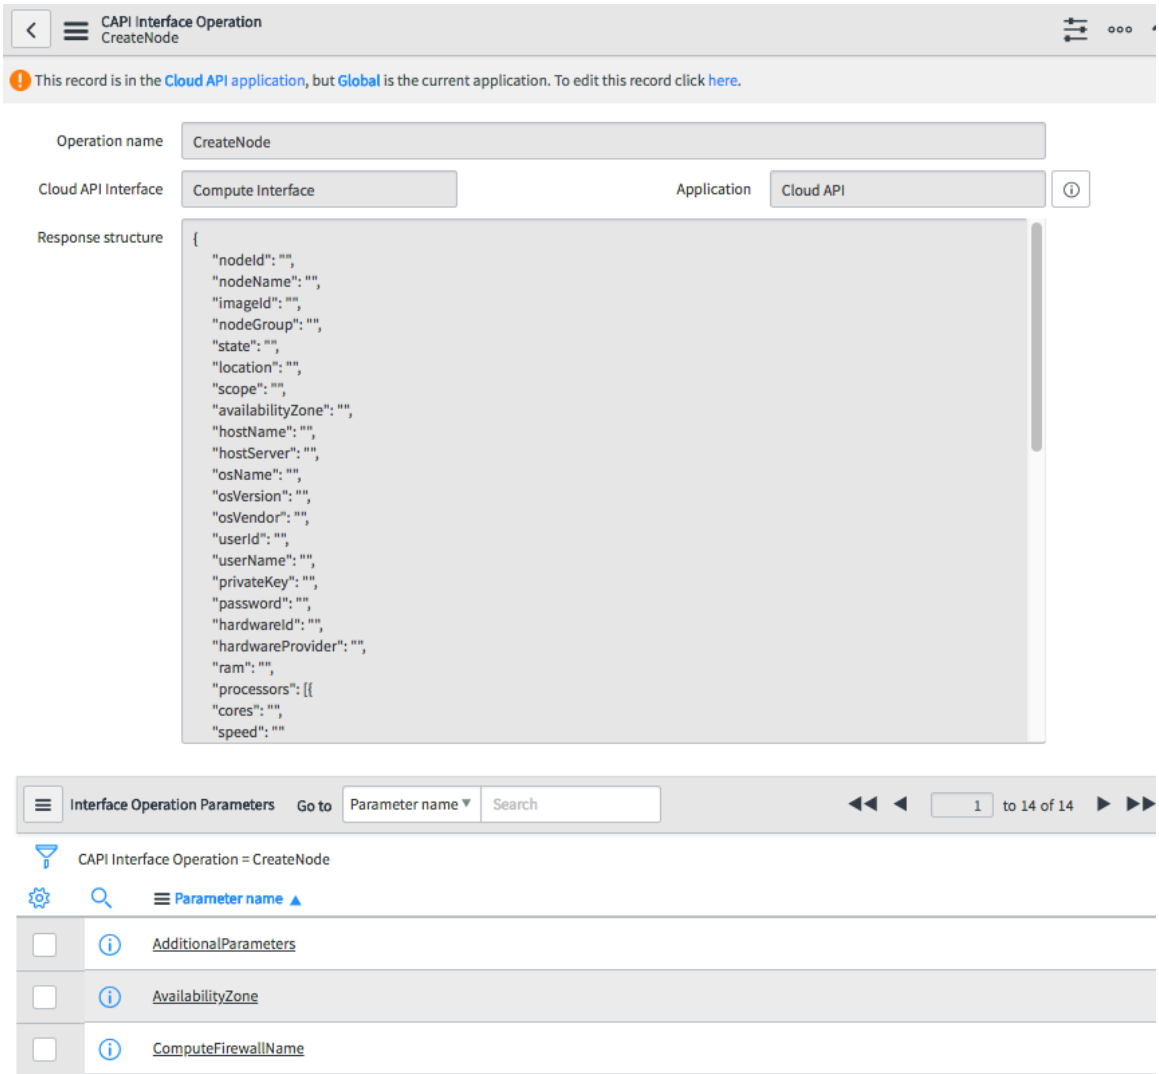

r. Navigate back to **Design** > **Cloud API** and then click the **Provider** tab.

s. Open the **Azure** provider.

The Azure Cloud Provider form opens, showing you that this provider is based on an existing CMDB class: **Azure Datacenter [cmdb_ci_azure_datacenter]**.

t. Click the Cloud Products related list, and sort the list by the **Name** column.

Notice that several Azure products are available by default. One of the most commonly used Azure products created from templates is **Microsoft.Compute**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

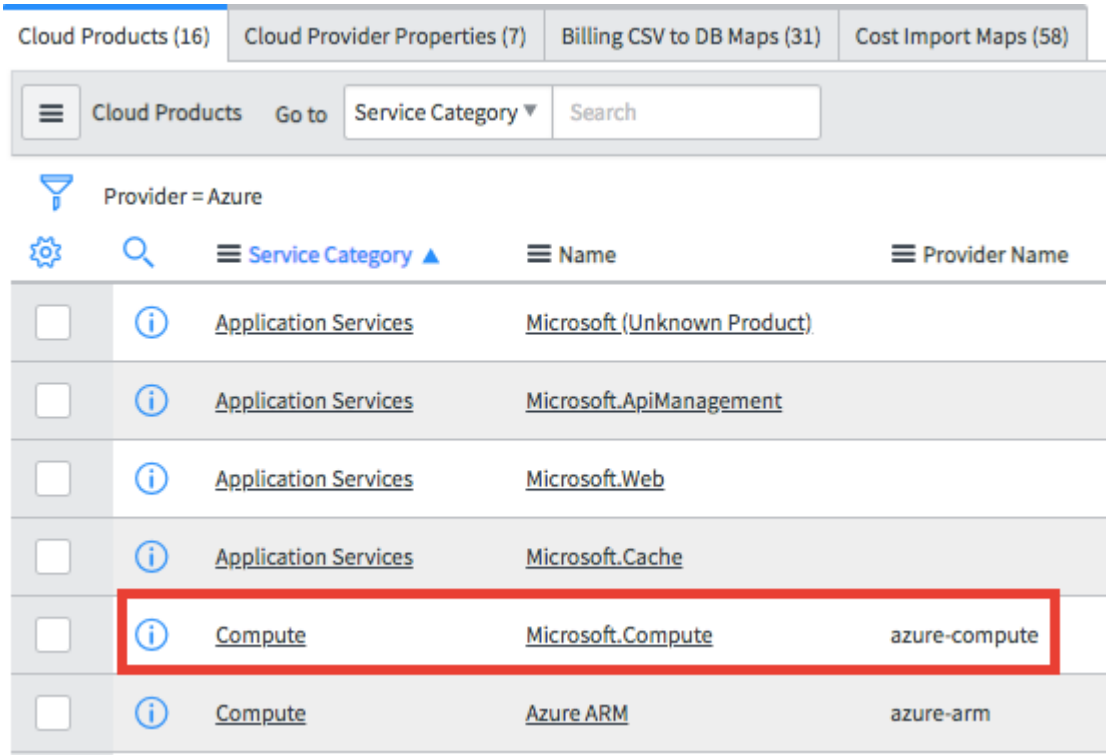

u. Click **Microsoft.Compute** in the **Name** column.
Notice that the product specifies many resource types, each of which is mapped to a CI class.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

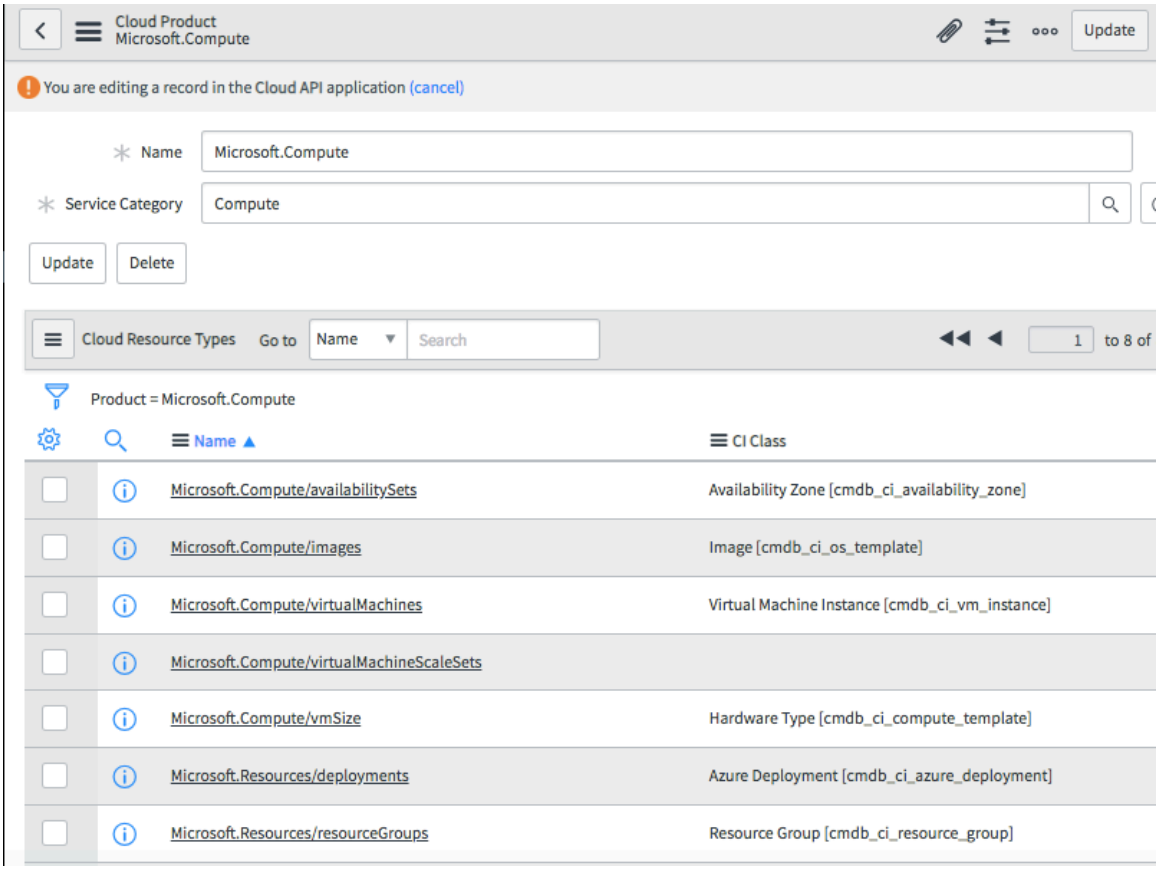

These resource types indicates some of the CIs, but not all, that are related to the virtual machine. The response processor in the resource block populates CIs with data when the cloud provider provisions the VM.

5. To see the important CIs that are related to the VM:

   a. On the Cloud User Portal, click **Stacks**, and then open the stack containing the VM.

   b. Click the **View Dependency** icon.

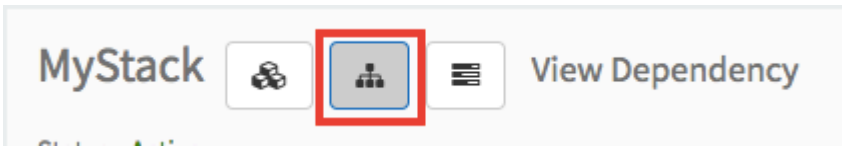

The dependency map displays the stack CI at the top, the VM in the middle, and the various related CIs, such as the image, at the bottom.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

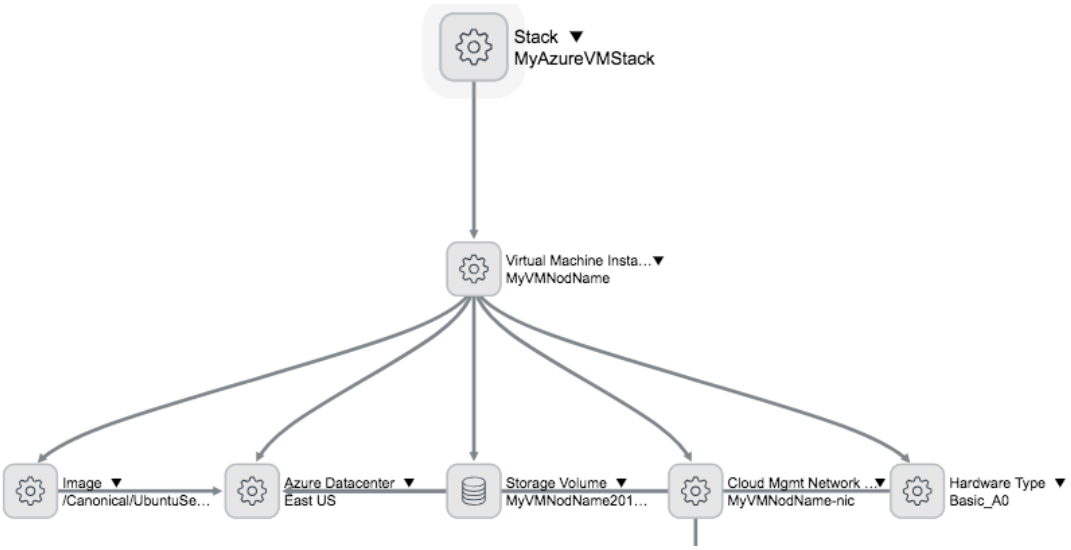

c. To view the form for the VM in the CMDB, right-click the arrow next to any CI.

d. From the menu, select **View Form**.

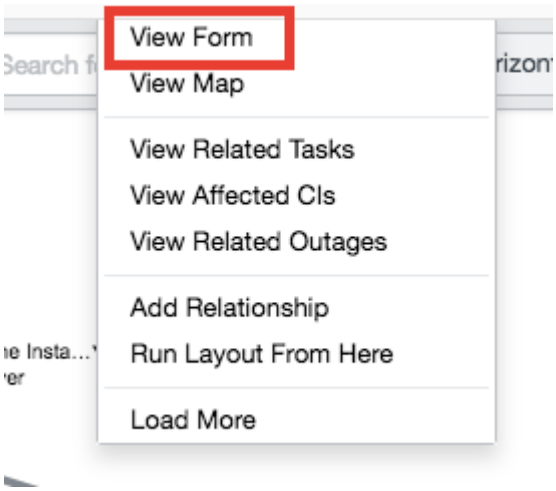

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The CI form opens, showing you much of the information is already available on the Cloud User Portal when you view the properties of the VM.

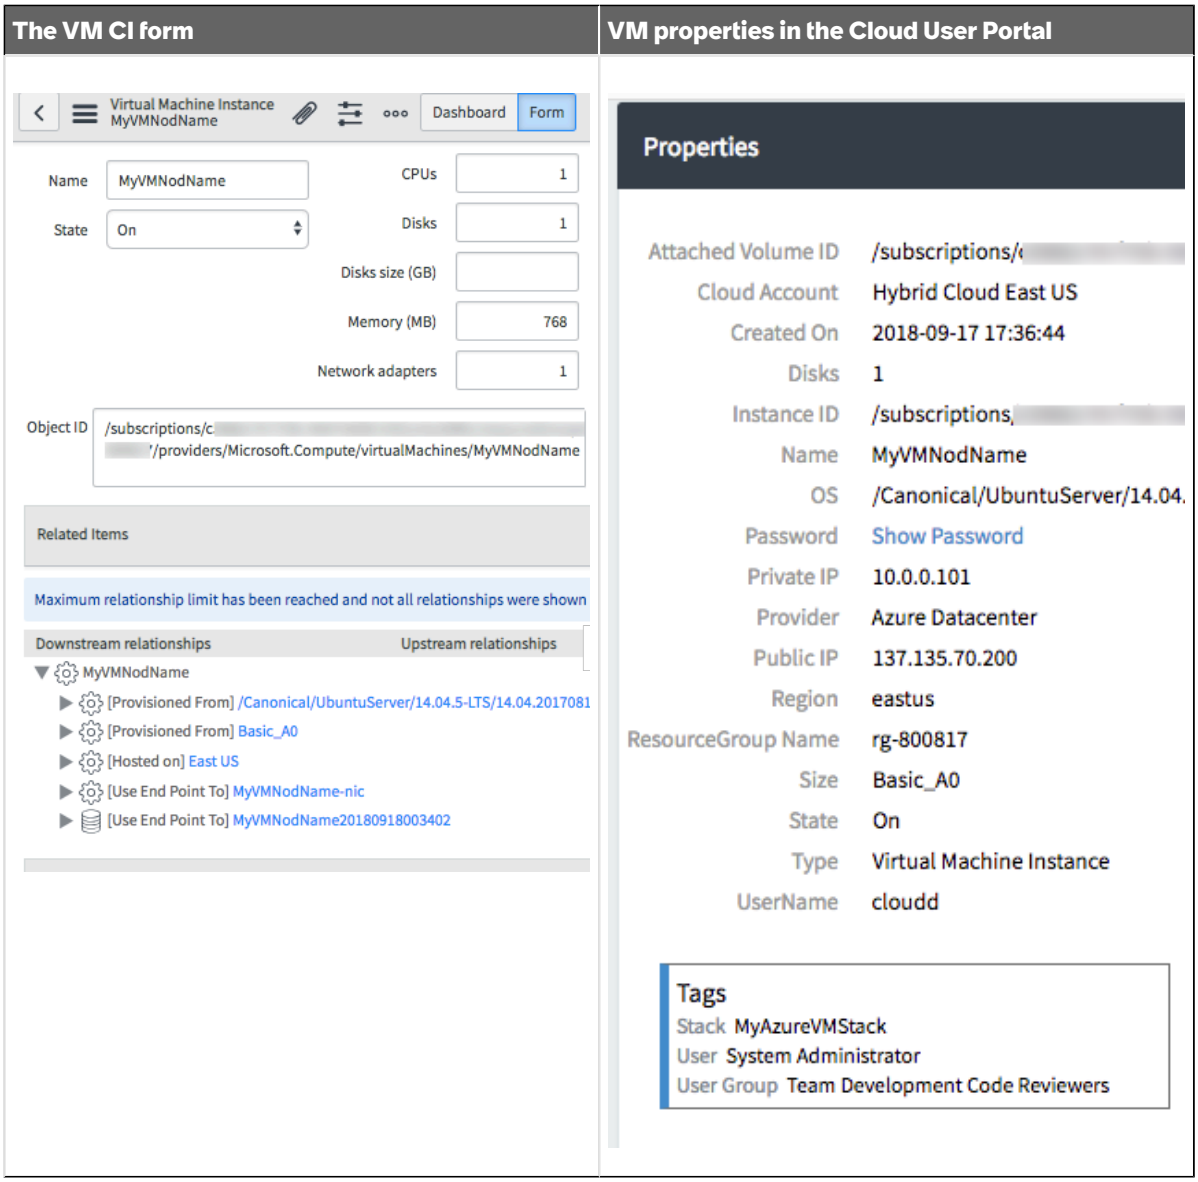

AWS VM provisioning walkthrough

This example walks you through the components of Cloud Provisioning and Governance that function during the provisioning of a virtual machine in an AWS datacenter. Topics covered include blueprints, resource blocks, the Cloud API (CAPI), and MID Server script includes.

Before you begin

Role required: admin

About this task

This walkthrough starts with a Windows VM that a user already provisioned in AWS. Next it walks you through the blueprint with the VM, the resource blocks, and then the CAPI calls specified from the resource blocks. Finally, the walkthrough shows you how a base MID Server script include makes the actual calls to the AWS API.

This example uses default resource blocks and script includes that are available in your instance. Therefore, although you might not have a provisioned VM on your instance, you can still follow

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

these steps and view the components used in this example to understand how the components work.

For an example of a VM in Azure, see Azure VM provisioning walkthrough.

Note: The terms virtual machine, VM, and virtual server are used interchangeably in this example.

Procedure

1. Look at a provisioned VM in the Cloud User Portal, and focus on some of the details about the VM:

   a. On your instance, open the Cloud User Portal (Cloud Provisioning and Governance **Cloud User Portal**).

   b. View a provisioned VM in a stack, such as this Windows VM, by clicking **Stacks**, and then clicking the name of the stack.

   In this example, the stack is named **MyStack**.

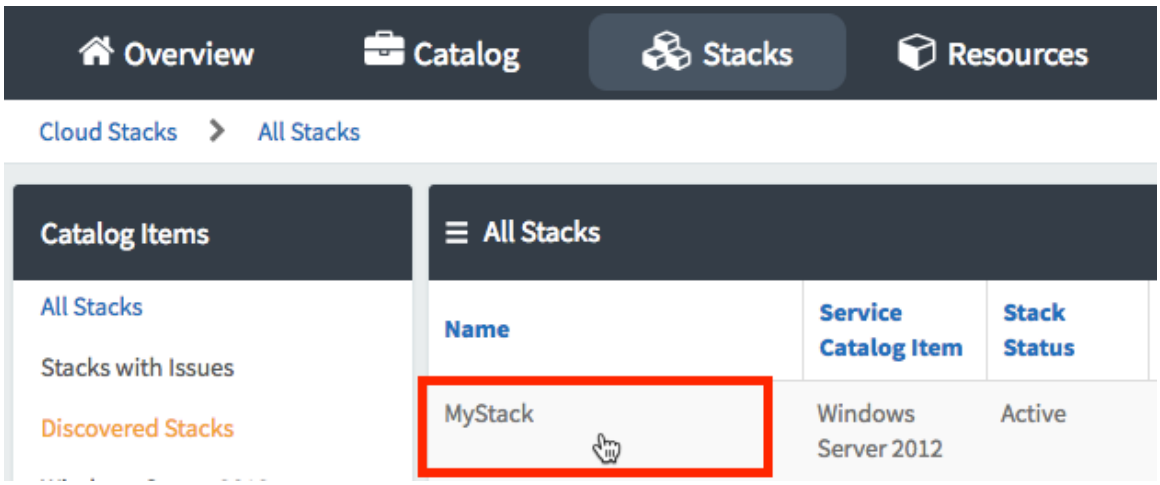

   c. Under **Resources**, click the VM in the stack.

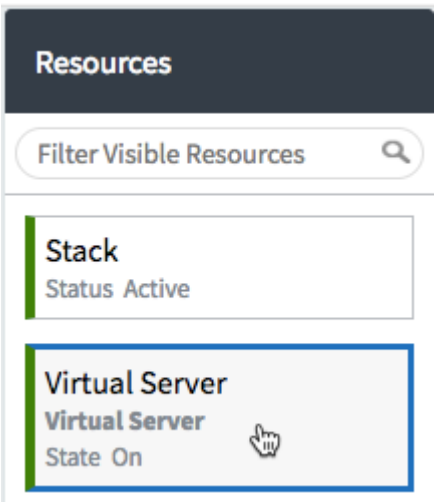

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. View the properties of the VM, and notice that it is in region us-east-1 in an AWS Datacenter.

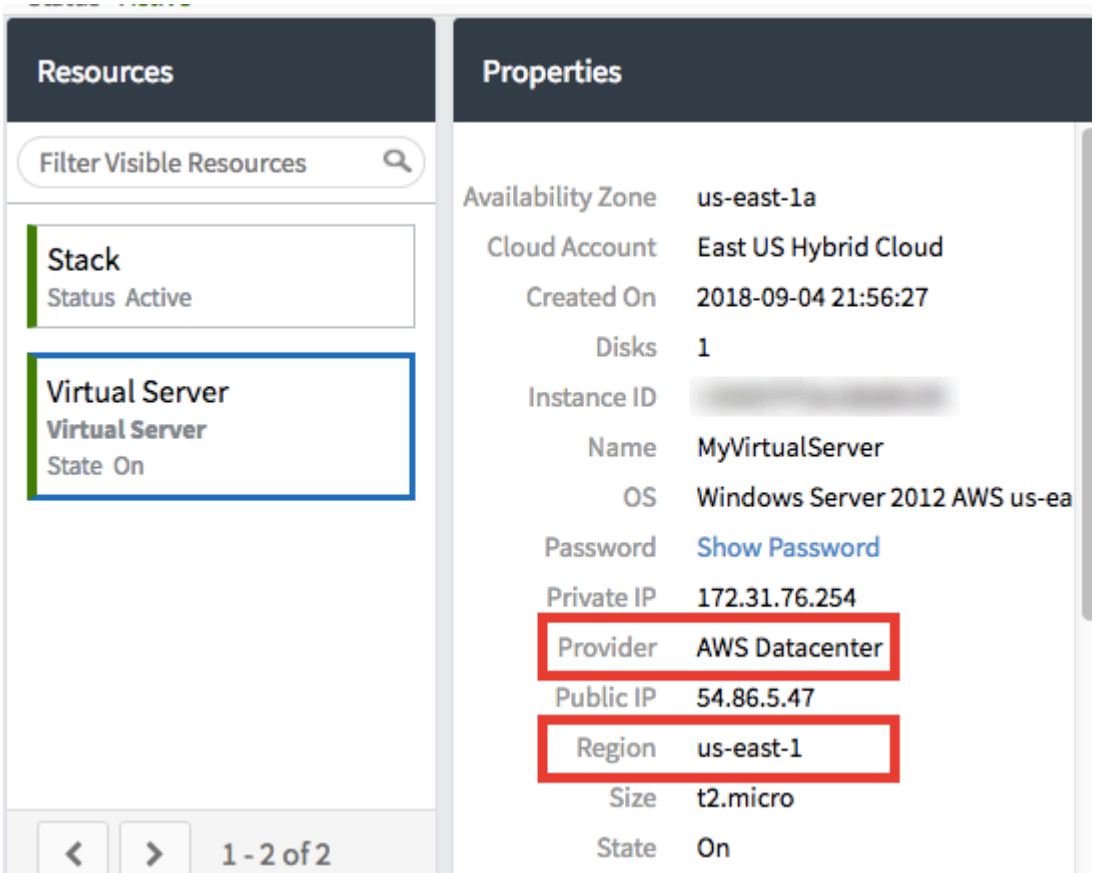

2. Look at the blueprint on which the VM is based:

a. Navigate to **Design** > **Blueprints**, and then open a blueprint with a virtual server in the Azure datacenter. The **Deployment Model** tab appears by default, showing you the various components of the blueprint.

This example blueprint has three components: The container, the virtual server, which is the actual VM that is provisioned, and the AWS datacenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

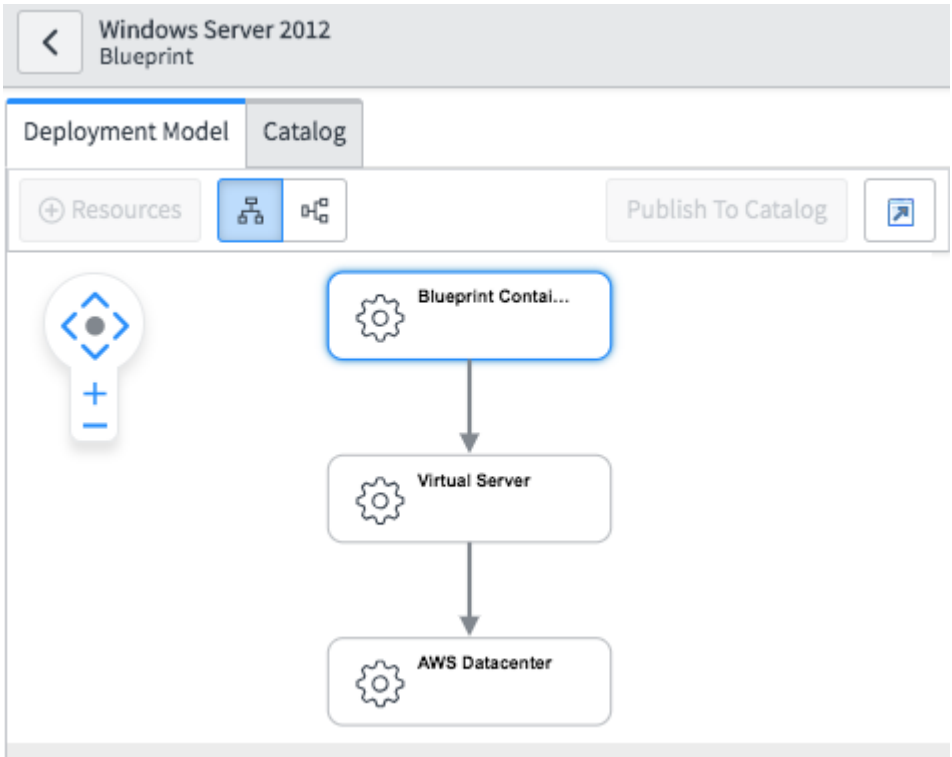

b. Click the **Operations** tab on the bottom, and then click **Provision**.

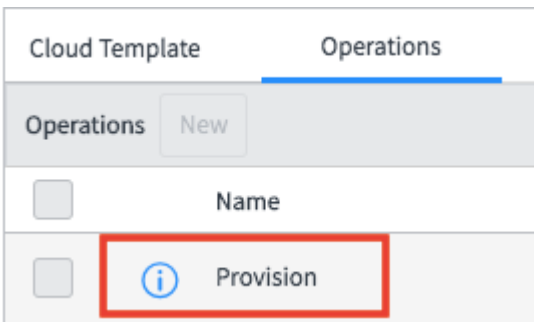

The Provision operation is the operation that the system triggered when it created the VM. Other default operations are available, but this example focuses on the Provision operation.

c. Click the **Provision** block for Blueprint Container Resource.

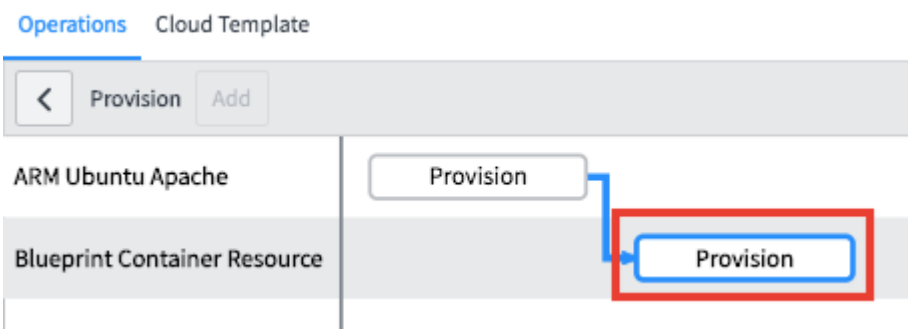

d. On the right, notice that one of the parameters in the Inputs list is Location. This parameter holds the value eastus, which is where the VM lives in the datacenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

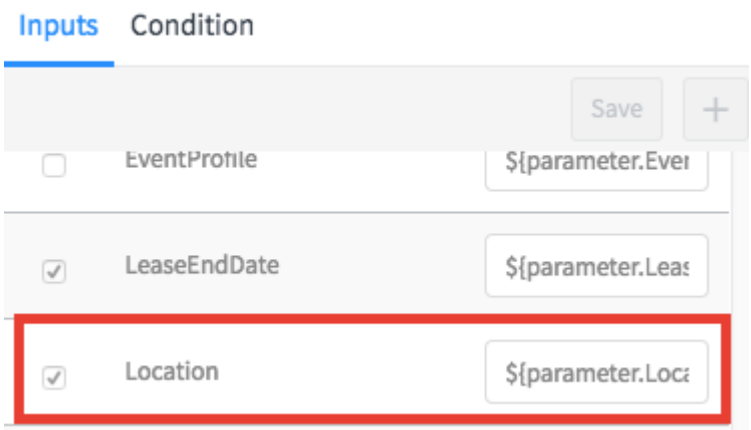

Inputs can be specified on the container, as it is in this example, or on any other resource block. By default, the *Location* parameter is already specified for you in the Blueprint Container resource block, so that you can use it in every blueprint like this one. If you switch the blueprint to **Draft**, you can add more parameters to the Blueprint Container resource block. You cannot add parameters to the Virtual Server resource block. For this walkthrough, no additional parameters are necessary.

3. Look at the Virtual Server and AWS Datacenter resource blocks used in this blueprint:

a. In the Cloud Admin Portal, navigate to **Design** > **Resource Blocks**.

b. Search for and open the **Virtual Server** resource block, which is provided by default with Cloud Provisioning and Governance.

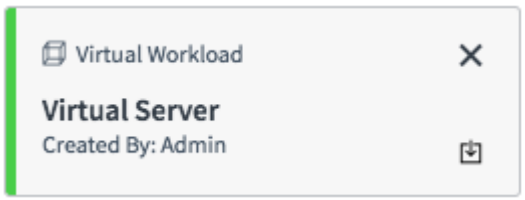

c. On the Resource Block form, click the **Host Resource** related list to view the list of supported host resource blocks.

d. Notice that the host interface and host resource are already specified:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

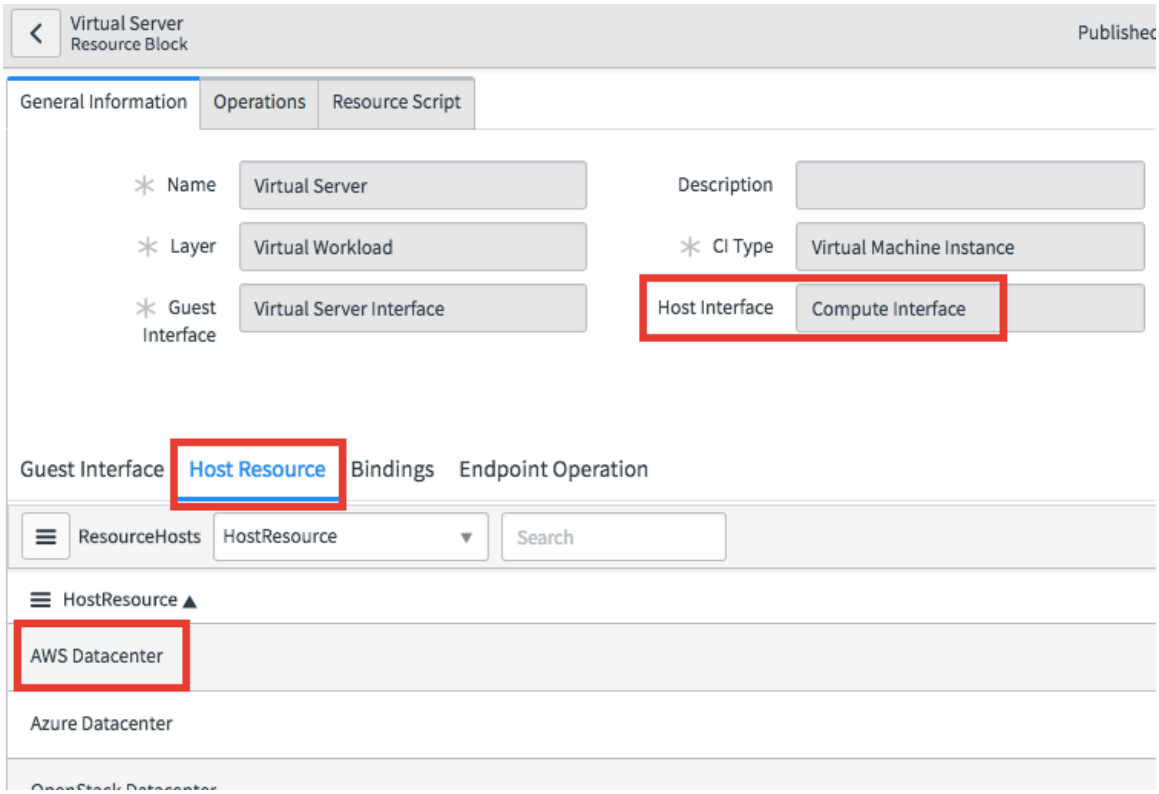

- The **Host Interface** field specifies the interface that must match the guest interface in the resource block that hosts this virtual machine. In this case, the host interface is the Compute Interface, which is also the guest interface on the AWS Datacenter resource block. By default, the datacenter resource blocks provide several guest interfaces that other resource blocks like virtual storage can use to connect to the datacenter.
- The **HostResource** column in the **Host Resource** related list already specifies **AWS Datacenter**, which means that the AWS datacenter resource block is a valid host for this VM.

e. Click the **Operations** tab, and then click the **Steps** subtab.

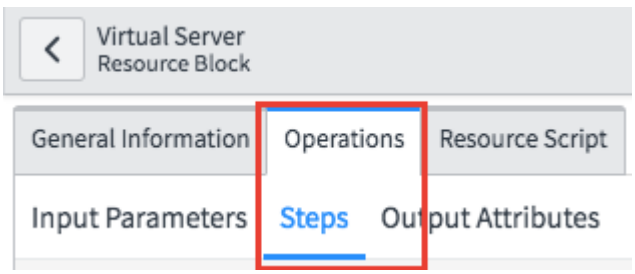

f. Select **Provision** from the **Operation** list.

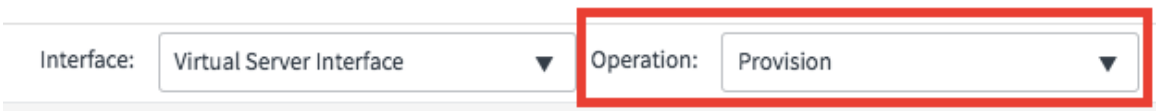

Remember that the Provision operation is the operation that the system used to create the VM. Other default operations are available, but this example focuses on the Provision operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

g. Notice the step that appears in the list and the full step description that appears above the input parameters:

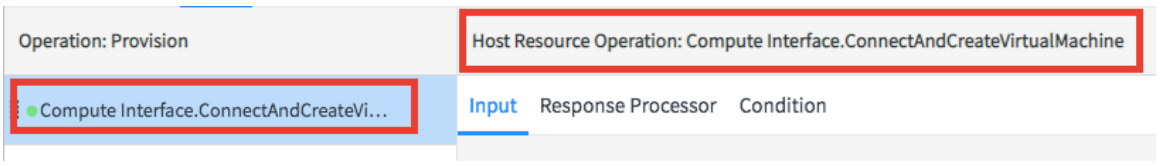

- **Host Resource Operation:** indicates that this step calls an operation on the host resource (the AWS datacenter in this example).
- **Compute Interface** is the guest interface on the AWS datacenter that this step is using.
- **ConnectAndCreateVirtualMachine** is the operation in the AWS datacenter that this step calls.

Note: In this case, the step calls an operation from another resource block: the AWS datacenter. Steps can also call CAPI directly, and then CAPI can execute REST calls to the cloud provider API. You can see that when you look at the AWS datacenter resource block.

h. Navigate back to **Design** > **Resource blocks**.

i. Open the **AWS Datacenter** resource block,which is the host resource block that the virtual server is connected to.

j. Notice the supported guest interfaces in the **Guest Interface** related list:

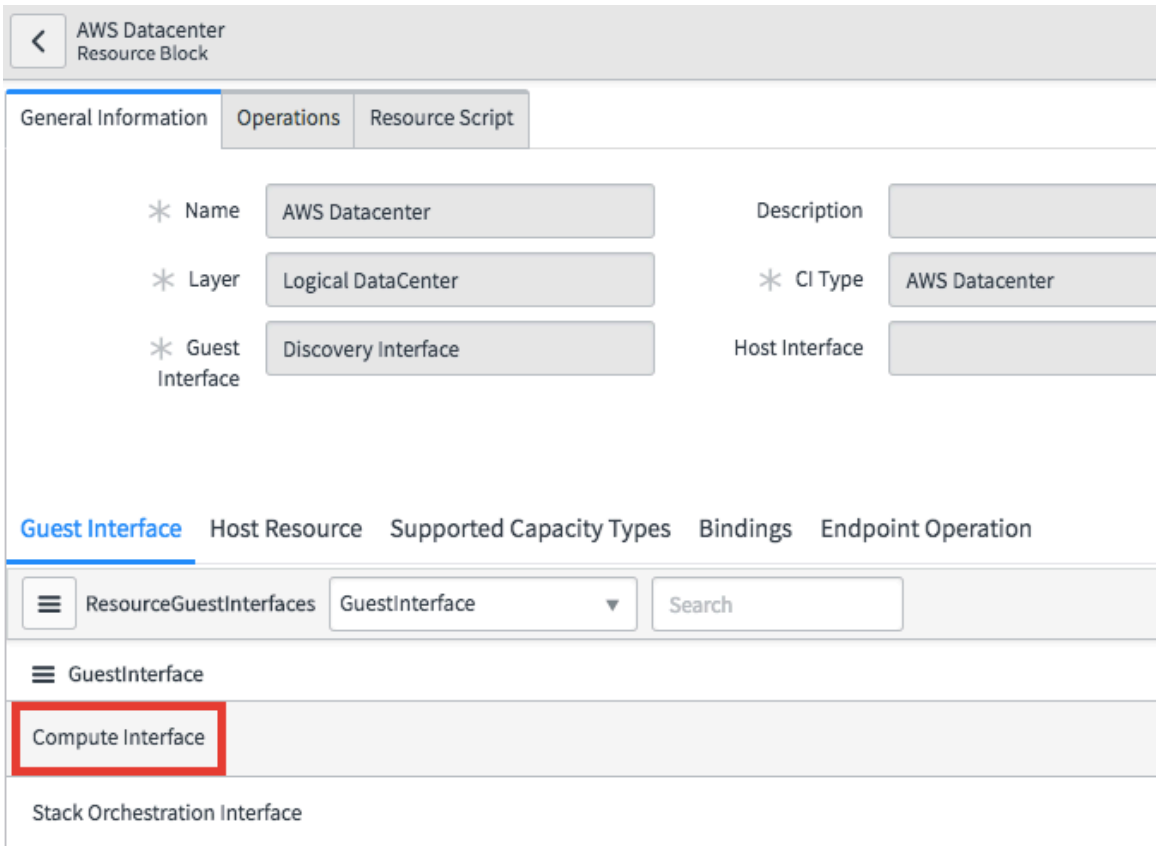

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

These guest interfaces are the interfaces that the AWS datacenter makes available to other resource blocks. The **Compute Interface** is provided so that the Virtual Server resource block, which specifies the Compute Interface as its host interface, can connect to the datacenter.

k. Click the **Operations** tab, and then click the **Steps** subtab.

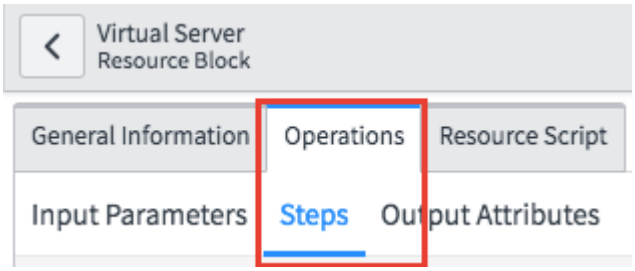

l. In the **Interface** list, select **Compute Interface** if it is not already selected.

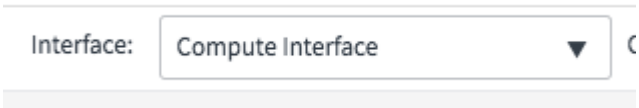

Remember that this interface is specified in the Virtual Server resource block.

m. In the **Operation** list, search for and select **ConnectAndCreateVirtualMachine**.

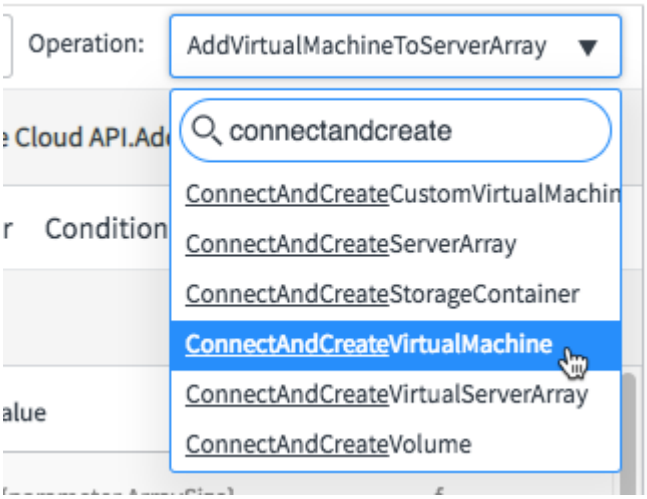

Remember that this operation is specified in the Virtual Server resource block.

n. Notice the CAPI call that is used in the only step for the ConnectAndCreateVirtualMachine operation:

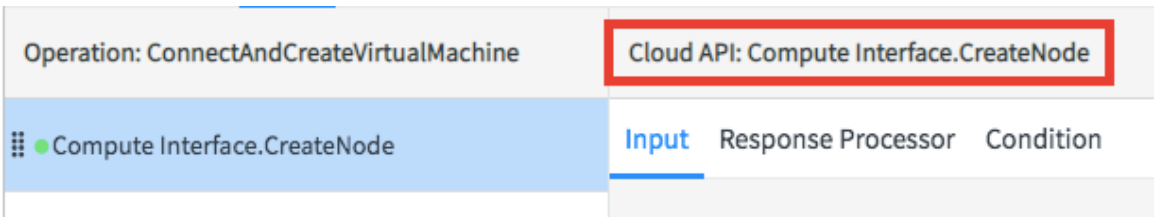

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Cloud API:** indicates that this step calls CAPI, so that CAPI can execute a REST call to the cloud provider, which in this case is AWS.
- **Compute Interface** specifies the CAPI interface that this step calls.
- **CreateNode** indicates the method that is executed. As the name suggests, this method tells AWS to create the virtual machine.

o. With the resource block in the **Draft** state, point to the highlighted (blue) step and then click the **Edit Step** icon to open the step configuration.

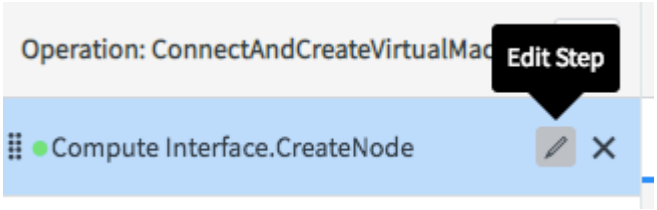

p. Look at the step configuration, and notice the settings that integrate with CAPI:

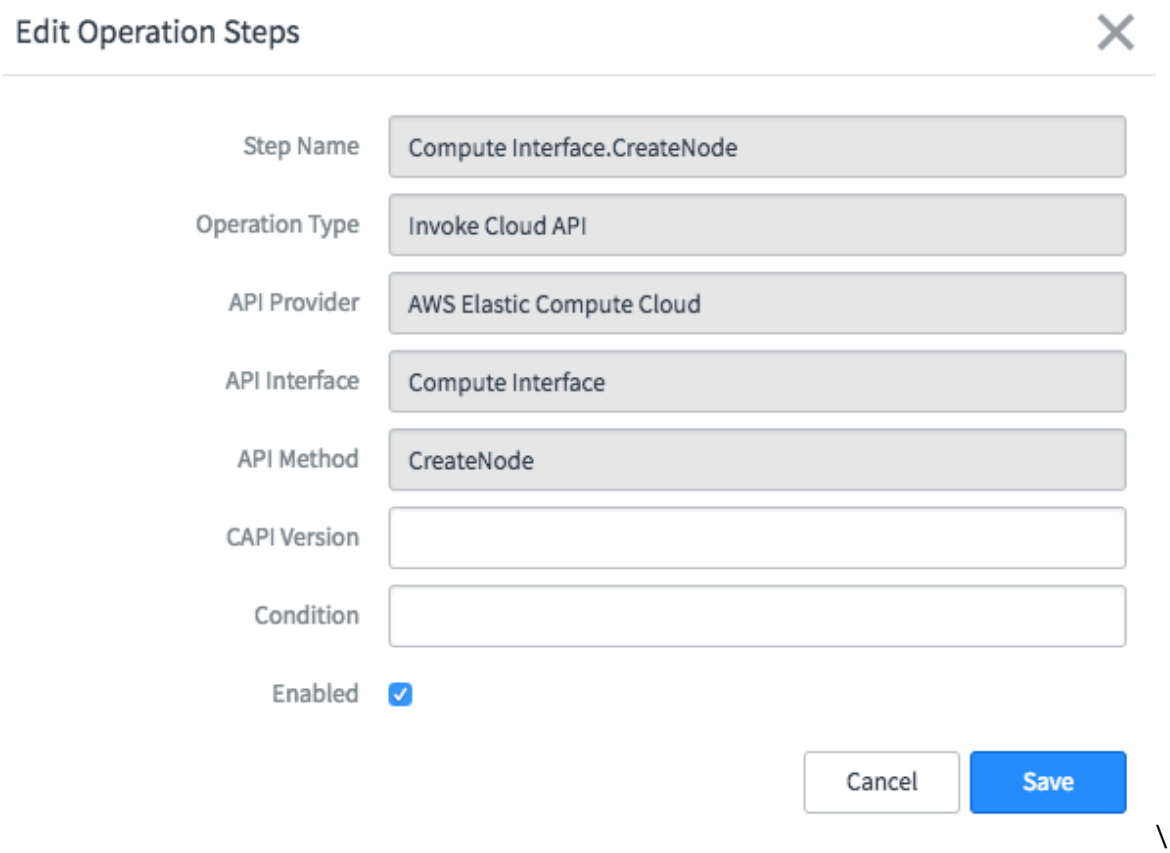

\

Field	Description
Operation Type	**Invoke Cloud API** specifies that this step should call the CAPI via the given provider, interface, and method.
API Provider	**AWS Elastic Compute Cloud** is a product (not that actual provider) that belongs to the AWS provider as defined in CAPI.
API Interface	**Compute Interface** is a product that belongs to the AWS provider as defined in CAPI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
API Method	**CreateNode** is the method that calls AWS to create the VM.

Note: The CAPI API definition ties together the provider (AWS), the product (AWS Elastic Compute Cloud), the interface (Compute Interface), and the method (CreateNode).

q. Close the window.

r. With the **Compute Interface.CreateNode** step selected, click the **Response Processor** tab, and notice the **Create_Virtual_Server_Response_Processor** script.

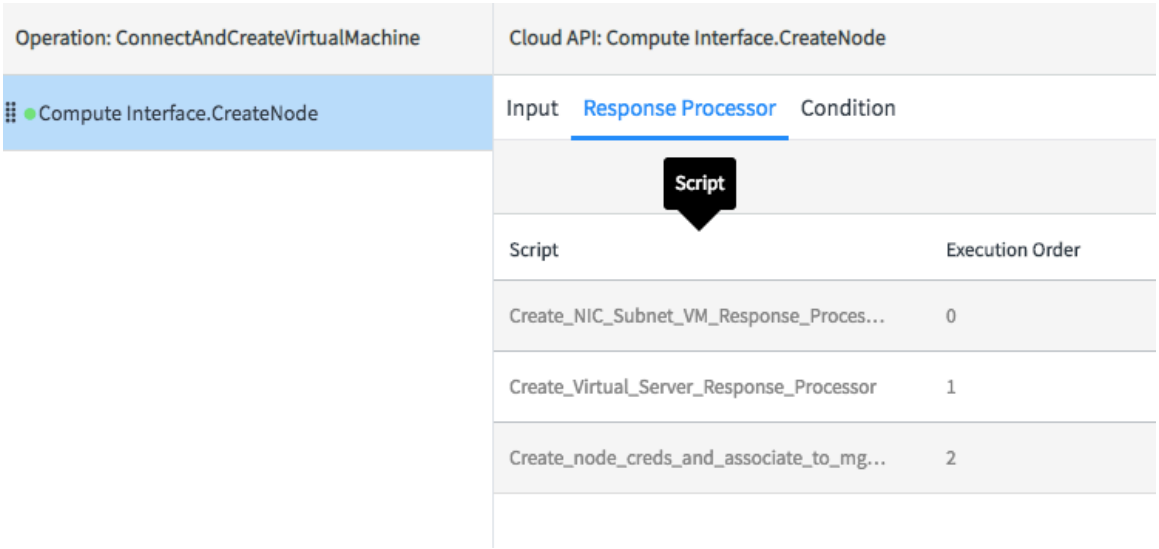

This script is the response processors that updates the CMDB in your instance after the virtual machine is created in AWS.

s. View an explanation of the script and the example that is a part of the topic at Create a Response Processor.
When you finish reviewing that topic, return to this topic.

4. Look at the CAPI components that work together to provision the VM:

a. In the Cloud Admin Portal, navigate to **Design** > **Cloud API**.

b. Click the **API** tab.

c. Search for an open **AWS Compute API**.

There are two AWS Compute API records. Open the first one in the list that matches the image in the next step.

d. Look at how this CAPI API ties together an interface and a product:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

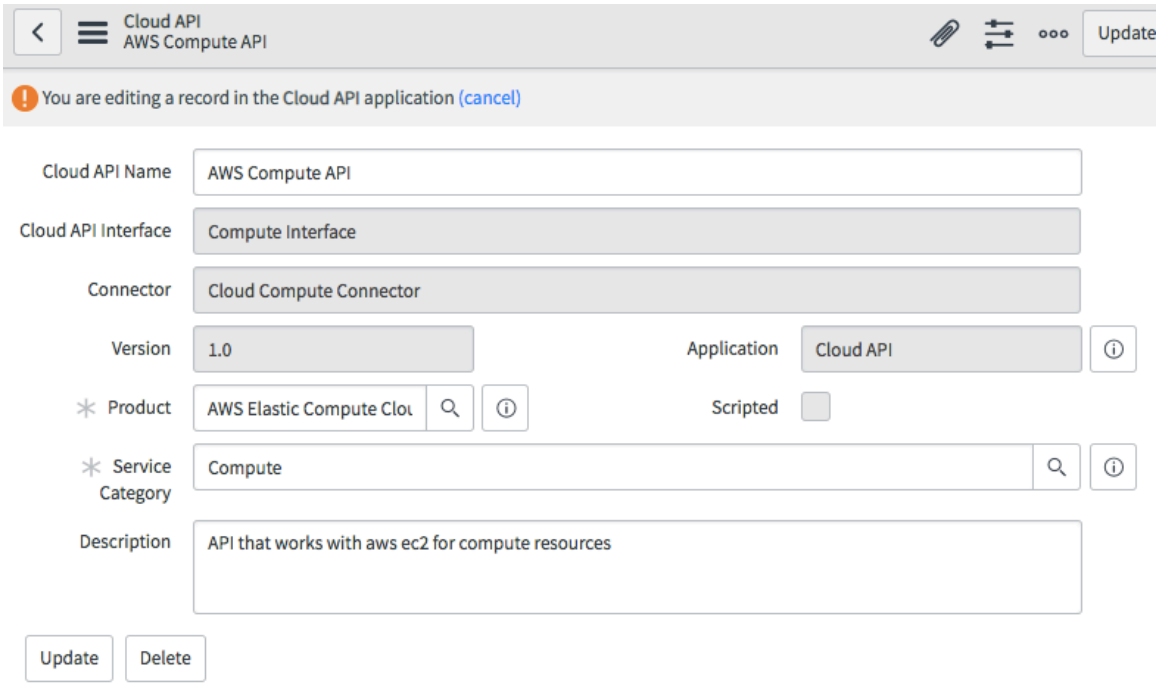

Field	Description
Cloud API Interface	**Compute Interface** is the same interface that you can see specified in the resource block. The interface contains the definition for methods, including the *`CreateNode`* method.
Connector	**Cloud Compute Connector** indicates that this CAPI makes Java calls to the system, which then calls the AWS API. This connector is not scripted. If this were a scripted connector, it was have a MID Server script include that calls the AWS API. **Note:** Most AWS-related APIs that are provided in the Cloud Provisioning and Governance application by default are not scripted, like this one, and cannot be modified. But you can create your own scripted APIs.
Version	**1.0** indicates the version of the API. You could have multiple versions of this API with different version numbers. Remember that in the step in the AWS Datacenter resource block that creates the VM, a Version field is provided. Although it was blank in this example because there is only one version of this API, you could specify a different version number.
Product	**AWS Elastic Compute Cloud** is the product that belongs to the AWS provider in CAPI. This product is provided by default with Cloud Provisioning and Governance, and is one of the most commonly used products for creating VMs on AWS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. In the CAPI Method Mappers related list, click the record preview icon for the *CreateNode* record, and then click **Open Record**.

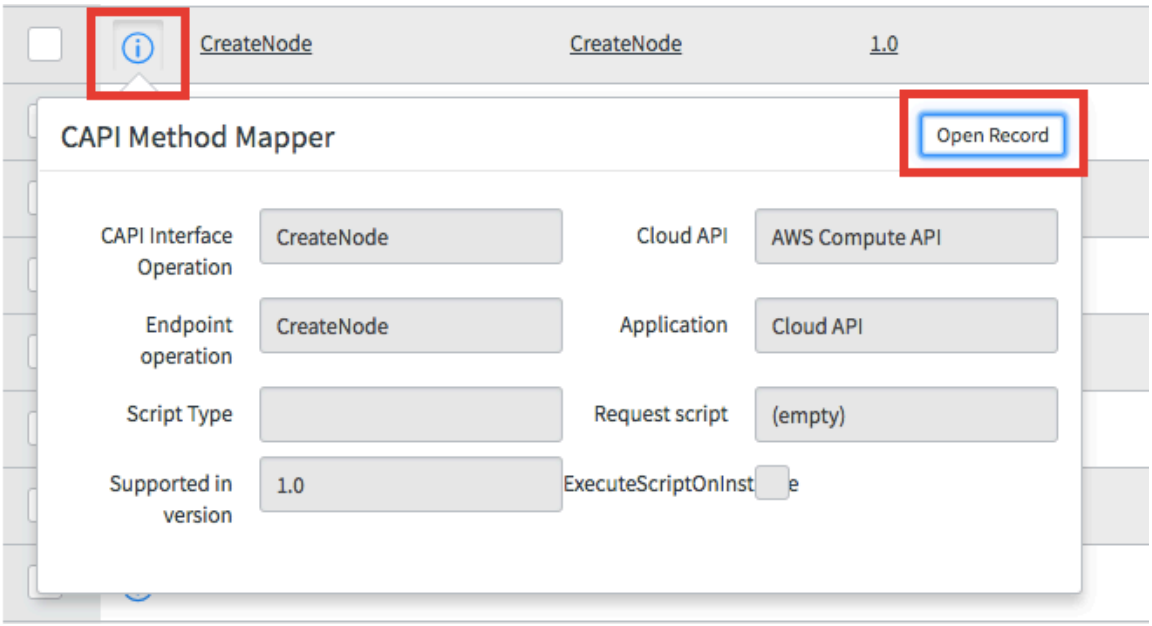

f. Look at the *CreateNode* method mapper:

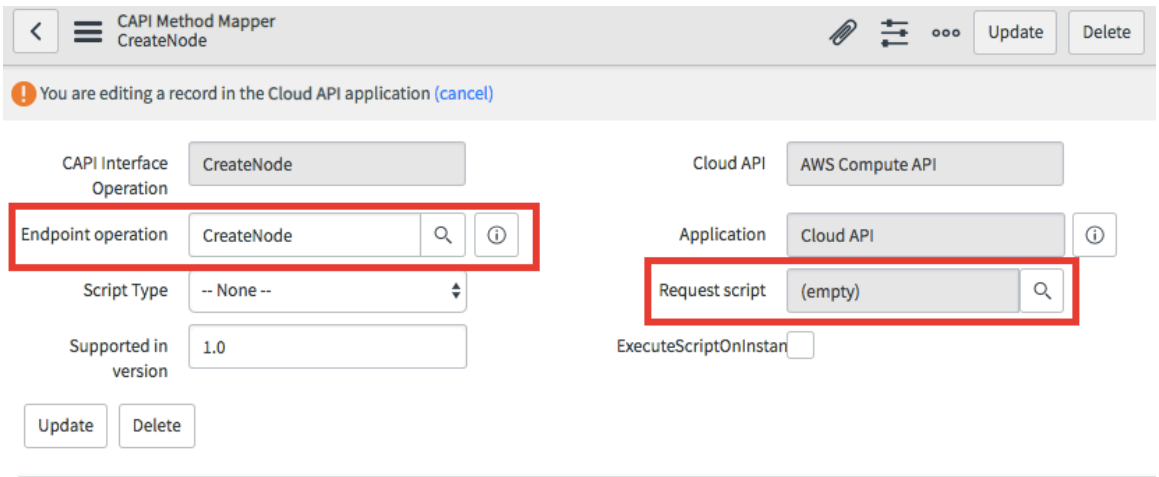

Notice that the Endpoint operation is not **Execute Script**. This value indicates that the CreateNode method is using a Java call within the Cloud Provisioning and Governance application on your instance to make a REST call to the AWS API. Therefore, you cannot modify how the CreateNode method works. If the value was **Execute Script**, you would see a MID Server script include specified in the **Request** script field. You would be able to modify that script include, or specify a new one, to customize the REST calls to the AWS provider.

g. Scroll down and look at parameters in the CAPI Parameter Mappers related list. Notice that important parameters, such as **Location**, are provided.

h. Navigate back to the AWS Compute API form.

i. Click the **API Config Overrides** related list, and review the items that are required for authentication:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Config Parameter and Override Value	Description
AccountAliasName $(CloudCredential.Alias)	The account alias is an optional value that you can create in your AWS account. It is a secondary name for your account ID. See the AWS documentation for more information.
Credentials $(CloudCredential.secret_key)	The secret key is used with the access key for authentication. You configured this value in your AWS credential record during setup. To refer to that procedure, see Configure access to the AWS accounts using permanent AWS credentials.
Endpoint $(CloudCredential.URL)	The endpoint is the URL that your users must log in to and authenticate. It uses the account ID or the account alias. See the AWS documentation for more information.
Identity $(CloudCredential.access_key)	The Identify record holds the AWS access key, which AWS requires for authentication.

j. Navigate back to **Design** > **Cloud API**, and then click the **Interface** tab.

k. Search for and open **Compute Interface**.
This interface is the interface specified in the resource block and the **AWS Compute API** CAPI API.

l. Review the contents of the interface.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Notice that the interface provides REST response structures for methods like *CreateNode*. You typically do not need to modify existing interfaces.

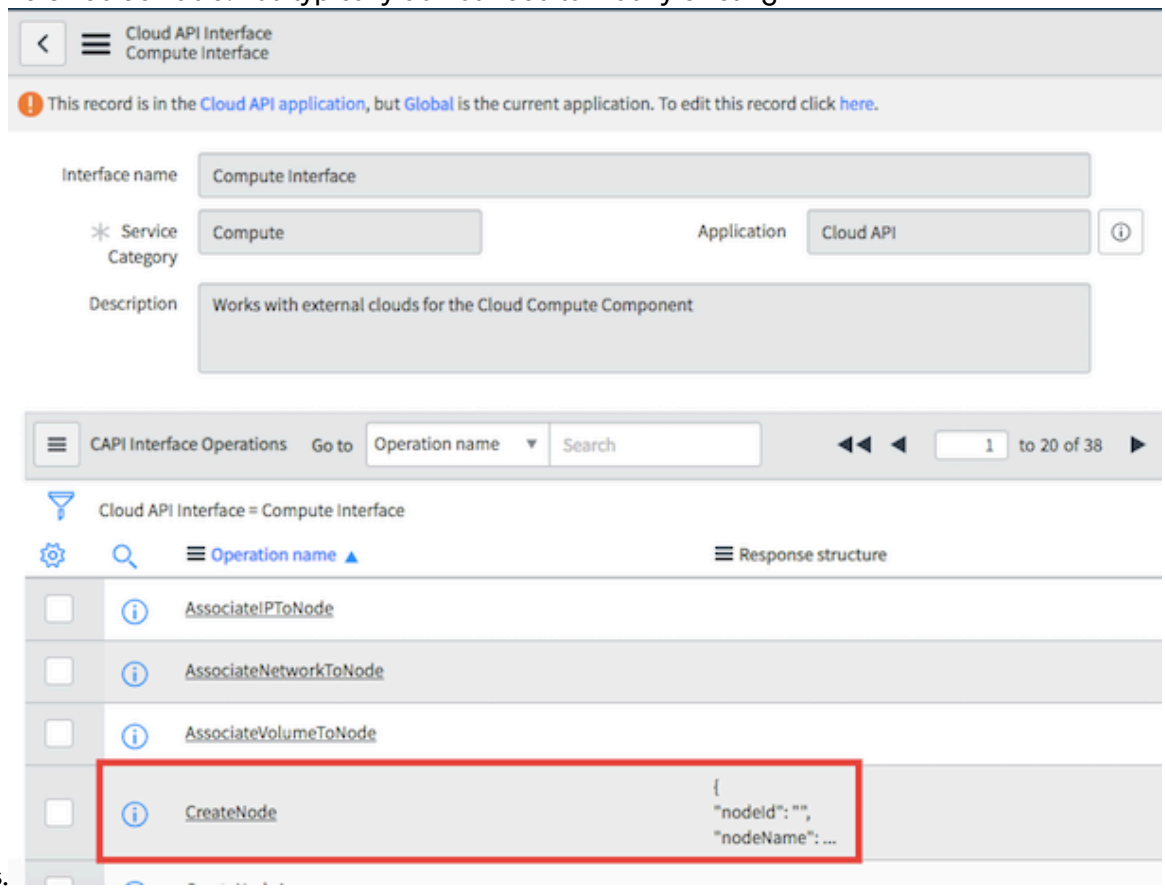

Note the service category and the operations:

Field or related list	Description
Service Category	The service category classifies the interface. The category for the Compute Interface is also **Compute**.
CAPI Interface Operations	The interface operations define the JSON structure for the REST call and the parameters that are required for the interface.

m. Click the **CreateNode** CAPI interface operation.
This operation is the one that provides the framework for creating the EC2 virtual server in AWS.

n. Review the components of the operation:

Field or related list	Description
Response structure	The response structure is the JSON framework for REST API call. It provides a list of attributes that AWS can use to create the virtual resource with empty values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field or related list	Description
Interface Operation Parameters	These parameters are also the parameters that the CAPI interface needs from the system so it can pass it in the REST call to the cloud provider.

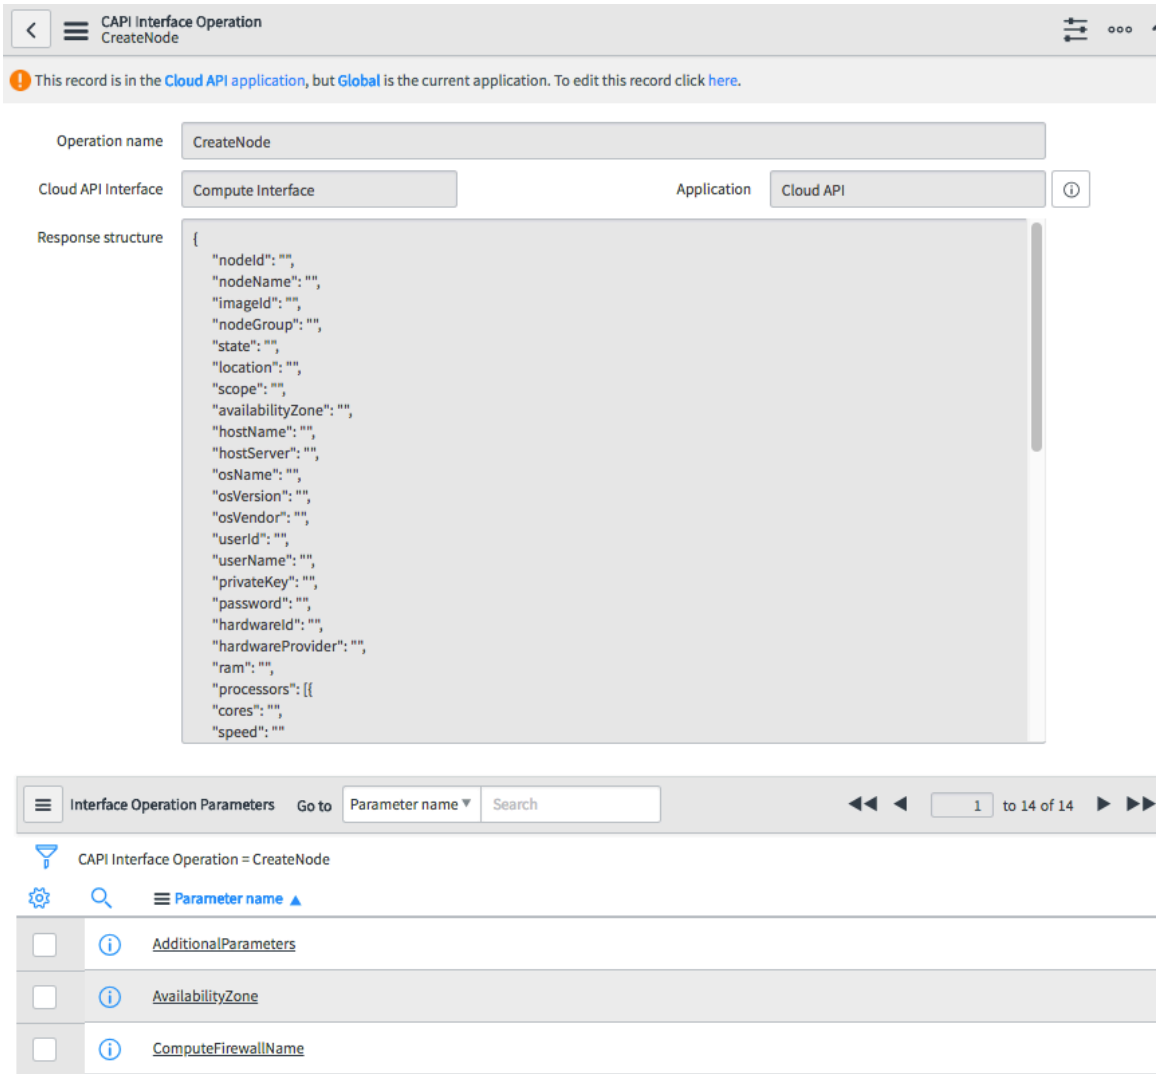

o. Navigate back to **Design** > **Cloud API**, and then click the **Provider** tab.

p. Open the **AWS** provider.

The AWS Provider form opens, showing you that this provider is based on an existing CMDB class: **AWS Datacenter [cmdb_ci_aws_datacenter]**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

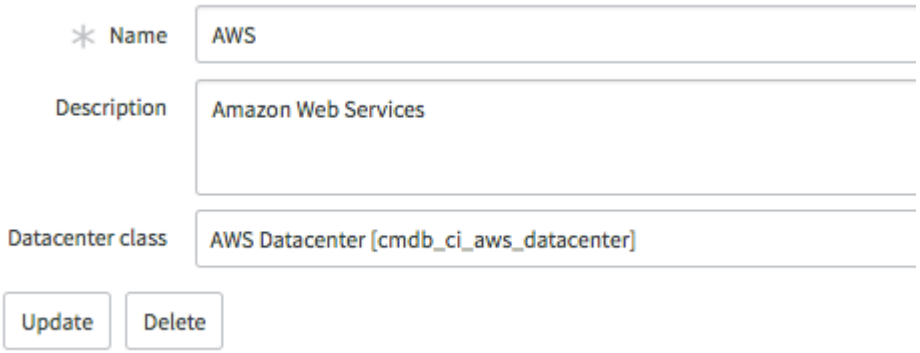

q. Click the Cloud Products related list if it is not already selected, and sort the list by the **Name** column.

Notice that several AWS products are already available by default. One of the most commonly used AWS products is **Elastic Compute Cloud (EC2)**.

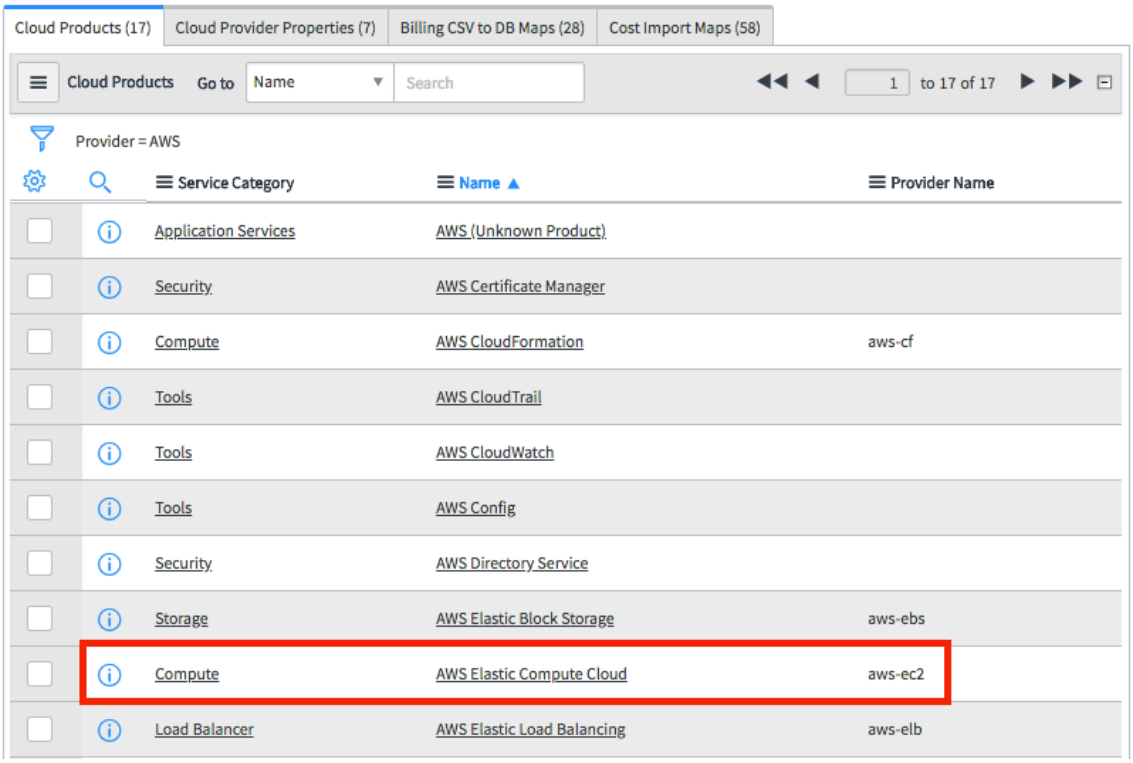

r. Click **AWS Elastic Compute Cloud** in the **Name** column.
Notice that the product specifies many resource types, each of which is mapped to a CI class.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

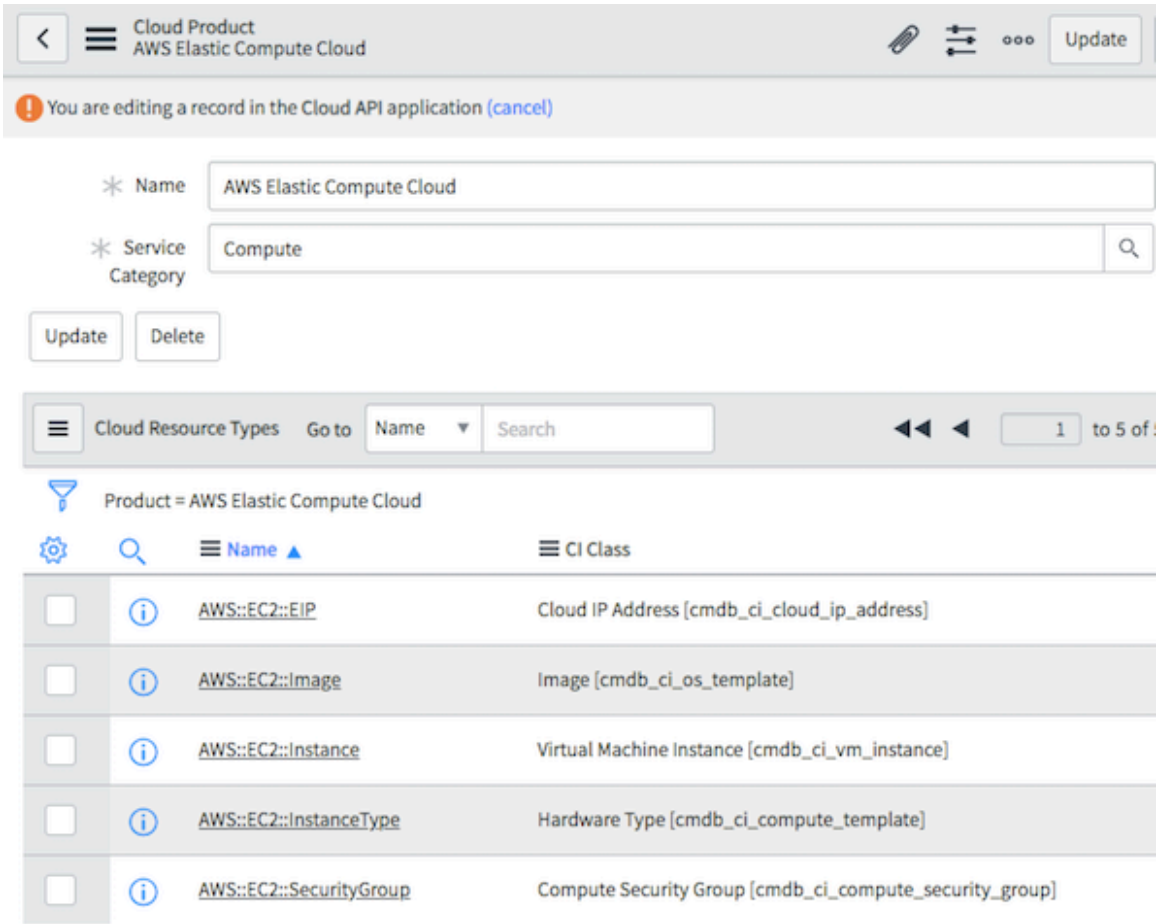

These resource types indicates the some of the CIs, but not all, that are related to the virtual machine. The response processor in the resource block populates CIs with data when AWS provisions the VM.

5. To see the important CIs that are related to the VM:

   a. On the Cloud User Portal, click **Stacks**, and then open the stack containing the VM.

   b. Click the View Dependency icon.

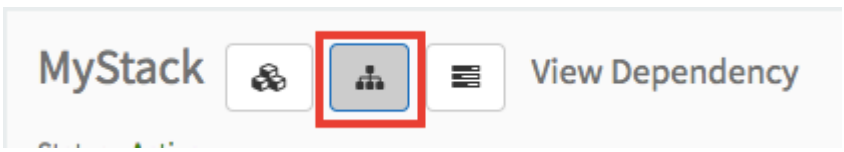

The dependency map displays the stack CI at the top, the VM in the middle, and the various related CIs, such as the image, at the bottom.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

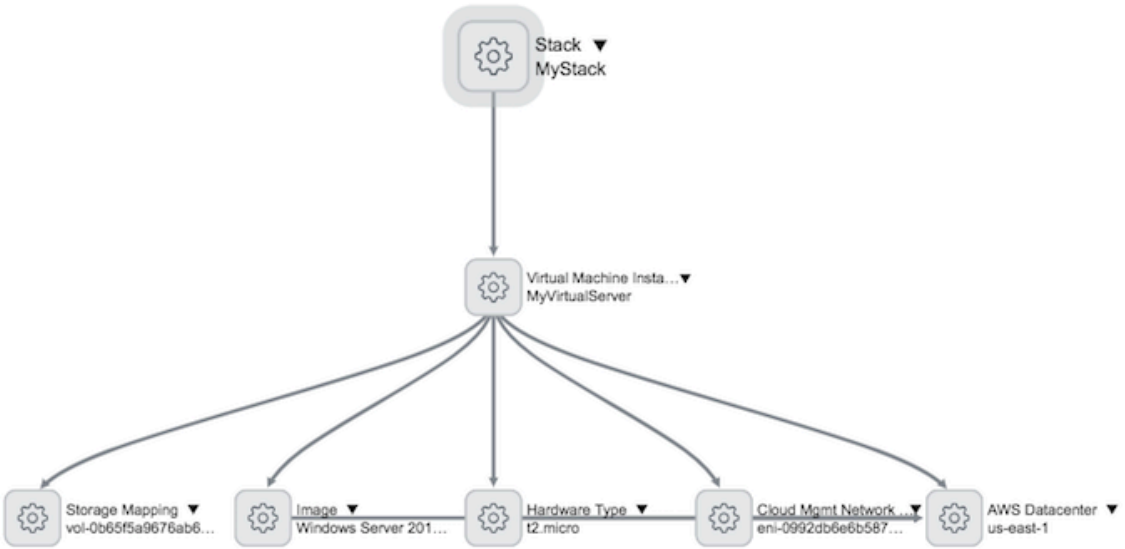

c. To view the form for the VM in the CMDB, right-click the arrow next to any CI, such as the VM.

d. From the menu, select **View Form**.

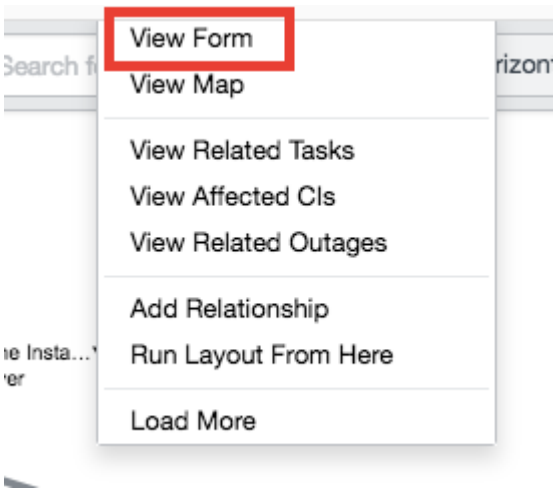

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The CI form opens, showing you that much of the information is already available on the Cloud User Portal when you view the properties of the VM.

	VM properties in the Cloud User Portal
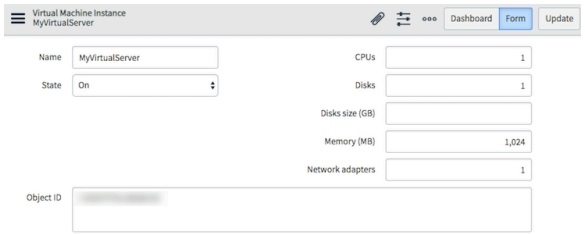	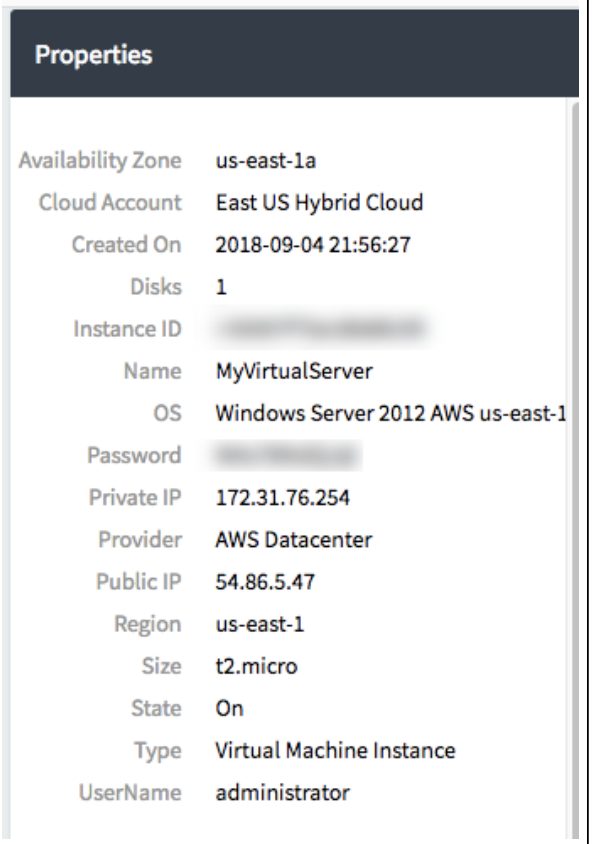

Add a product to an existing provider in CAPI

The Cloud Provisioning and Governance Cloud API (CAPI) offers many providers, such as AWS and Azure, and products, such as EC2 instances. However, you can add products to existing providers for new types of cloud resources.

Before you begin

- Role required: admin
- Review Default CAPI APIs and make sure the provider and product that you want to create are not already provided.
- Familiarize yourself with the cloud provider API so that you know how to create REST calls to the provider. If you are not familiar with the headers, authentication requirements, endpoint, and data that you can pass into the provider API, you cannot build your own CAPI components.

About this task

This topic provides general instructions for extending an existing cloud provider, product, and resource type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In your instance (not in the Cloud Admin Portal), enter the following text into the application filter:

   `sn_capi_product.list`

   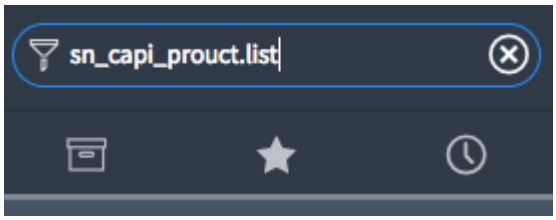

2. Press ENTER.

   The list of products that belong to all providers appears. You can filter the list to verify that the product you want to create is not already provided.

3. Click **New** and then enter a name and choose one of the existing service categories that is most appropriate.
4. Click **Submit**.

   The list of products appears again.

5. Find the product that you created, and then double-click the empty cell in the **Provider** column.

   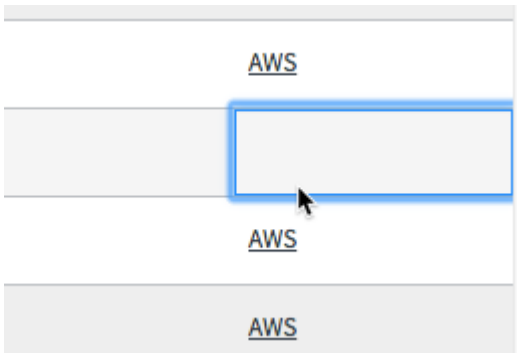

6. Add the correct provider.
7. Add a unique provider name: Double-click the cell in the **Provider Name** field and then change the name.

   The system can auto-generate a provider name for you, but you should modify it to avoid conflicts with existing provider names.

8. In the Cloud Admin Portal, navigate to **Design** > **Cloud API**.
9. Click the tile for the provider to which you added the new product.
10. Verify that your product is in the **Cloud Products** related list.
11. Click the information icon (ⓘ) for the cloud product, and then click **Open Record**.
12. In the Cloud Resource Types related list, click **New** and then fill in the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	A unique and descriptive name.
Product	The Cloud Product for which you're creating the Resource type. Select from the lookup list or create new.
CI Class	A CMDB table you need to associate with the resource type. If you do not see the correct CMDB table, see Create a CI class for a virtual cloud resource for instructions on how to create one.
Pattern	The pattern that is used to discover the resource type. Select from the Discovery patterns lookup list or create a new pattern.
Dependency Sequence	The order in which resource type is provisioned or discovered. Specify the order based on inter-dependencies of resources. For example, if you have a network and a subnet that you want to discover or create, specify a higher order for network than subnet, so that discovery or creation of resources is performed logically and sequentially.

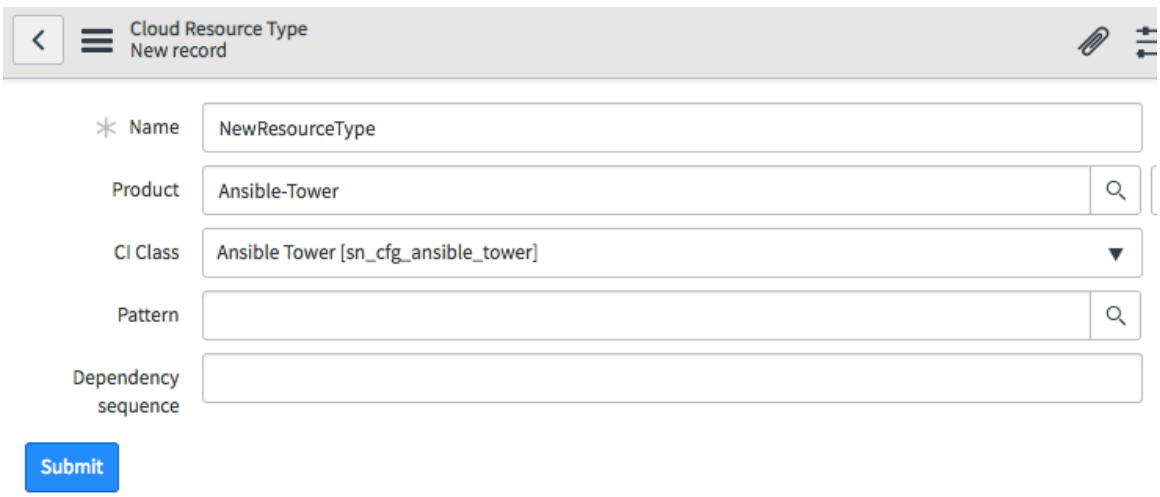

13. Click **Submit**.

What to do next
Create a CI class for a virtual cloud resource if necessary, or Create or extend a CAPI interface

Create a CI class for a virtual cloud resource

If the base system does not provide a CI type for cloud-based virtual resources, you can create a CI class based on the virtual machine object class.

Before you begin
Role required: admin

About this task
This topic provides general instructions for extending an existing cloud provider, product, and resource type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Configuration** > **CI Class Manager**.
2. Click **Open Hierarchy**.

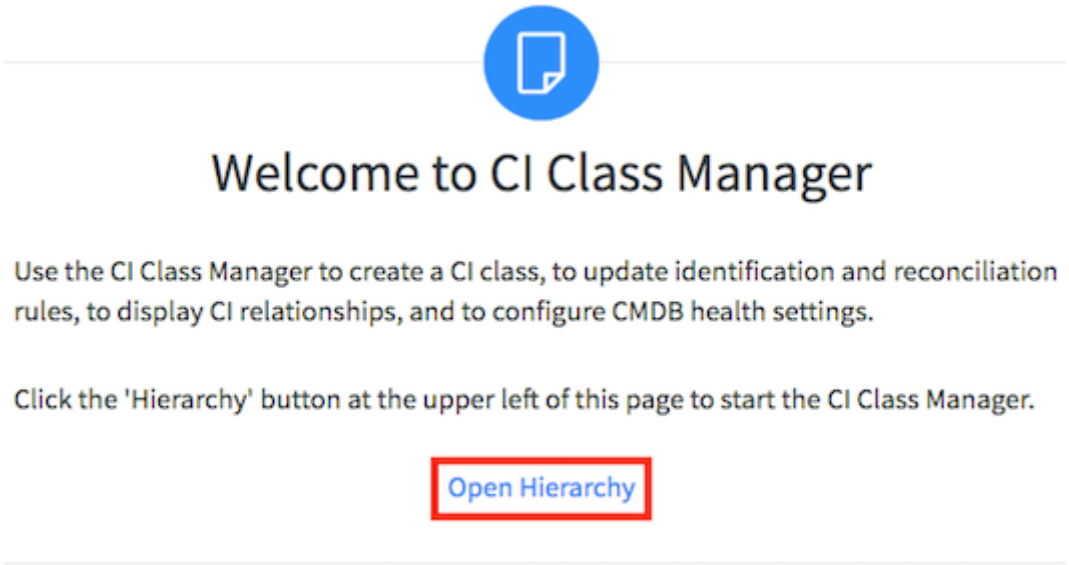

3. On the **CI Classes** pane, enter `Virtual Machine Object` in the search field, and then select **Virtual Machine Object**.

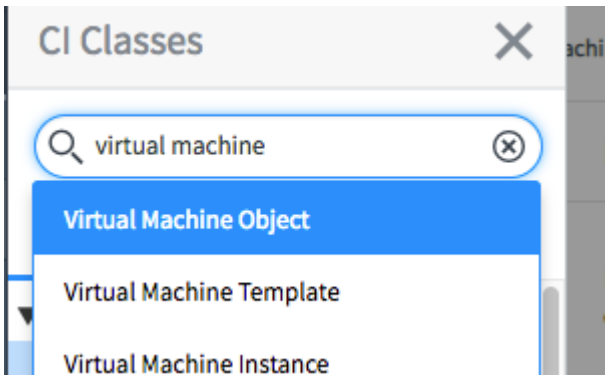

4. Click the **Virtual Machine Object** link.

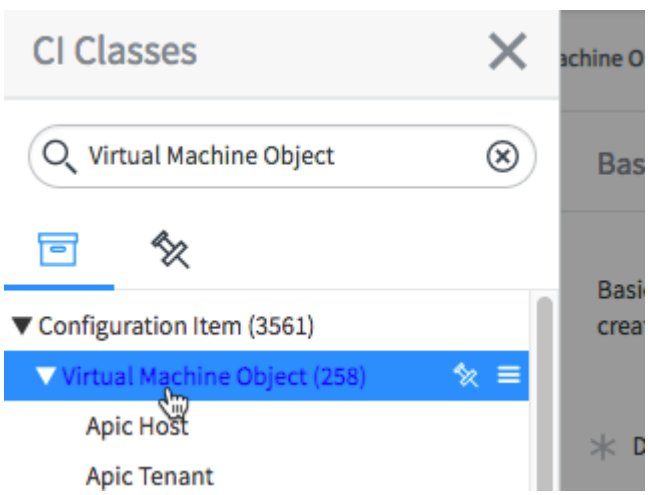

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Click **Add Child Class**.
6. In the **Provide Basic Info** step that appears, enter a name and a description.

   When you move the focus of the cursor away from the **Display Name** field, the **Table Name** field automatically populates with the same name in lower case. For a table in a scoped application, the name is prefixed with a namespace identifier to indicate that it is part of an application. For a table in the global application, the name is prefixed with the string **u_cmdb_ci**. You cannot modify the prefix; however, you can modify the rest of the table name. The name can contain only lowercase alphanumeric ASCII characters and underscores (_). The maximum string length is 80 characters.
7. Click **Next**.
8. On the **Add Attributes** step, click **Derived**, and then review the attributes that are inherited from the Virtual Machine Object class.

   Most attributes that you need are provided.
9. **Optional:** To add an attribute:

   a. Double-click the **Insert a new row** text in the **Column label** field.

   b. Enter a name for the column and then click the green arrow.

   c. Fill out the rest of the fields in the row.
   See Dictionary entry form for a description of all the fields.
10. Click **Next**.
11. On the **Set Identification Rule** step, create an identification rule that specifies a unique Object ID, which is required for identification.

    Important: Cloud Provisioning and Governance depends on the Object ID to uniquely identify an instance of the CI Class Type.

    a. Click **Replace** on the default rule under **Identification Rule**.

    b. Enter a name for the rule and select the **Dependent** option.

    c. Click **Save**.

    d. Under **Identifier Entries**, click **Add**.

    e. Move **Object ID** from the **Available** column to the **Selected** column.

    f. Click **Save**.
12. Click **Next**.
13. On the **Add Dependent Relationships** step, click **Override Dependencies**.
14. Configure the relationship between this virtual resource and the logical datacenter on which it resides:

Field	Value
Rule Type	**Hosting**
Relationship	**Hosted on (Hosted on::Hosts)**
Target Class	**Logical Datacenter**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

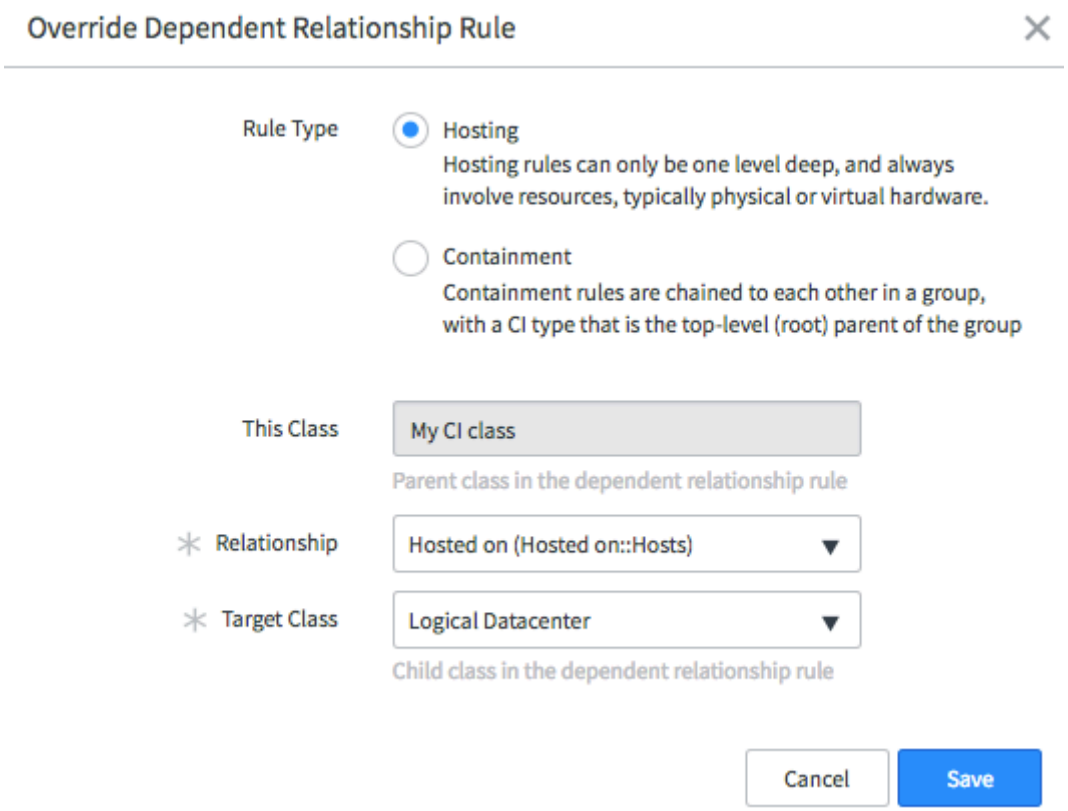

15. Click **Save**.
16. Click **Next**.
17. Click **Next** again to skip the **Add Reconciliation Rules** step, and then click **Done** to skip the **Add Suggested Relationships** step.

What to do next

Use the CI class in a resource type for a cloud provider. See Add a product to an existing provider in CAPI for instructions.

Your response processor, which is part of the resource block, must also specify the necessary CI class types so it can update the CMDB. Usually, the response processor needs the object ID, the specific datacenter, and the cloud service account. See Virtual server response processor example for example response processor code that uses CI classes.

Create or extend a CAPI interface

You can create a new CAPI interface or extend an existing one by adding new operations.

Before you begin

- Role required: admin
- Review Default CAPI APIs and make sure the provider and product that you want to create do not exist.
- Familiarize yourself with the cloud provider API so that you know how to create REST calls to the provider. If you are not familiar with the headers, authentication requirements, endpoint, and data that you can pass into the provider API, you cannot build your own CAPI components.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Usually, you can use the default interfaces that are provided in the Cloud Provisioning and Governance application. To create a custom interface, base it on an existing one that is similar.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud API**, and then click the **Interface** tab.
2. Click **New**.
3. In the Add Cloud interface window, enter a name for the interface, a helpful description, and select the most relevant category to which the API interface should belong.

   This example shows an interface for Azure storage:

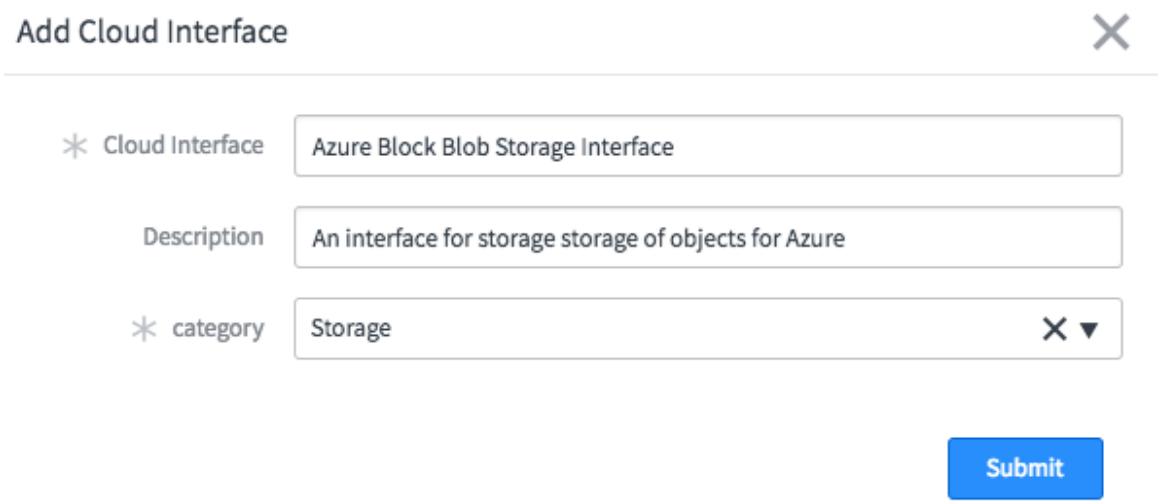

4. In the CAPI Interface Operations related list, click **New**.

   Each operation that you create corresponds to a different REST API call.

5. Provide a name for the operation that corresponds to the API you want to call in the provider.

   For example, you can create an operation called *`CreateContainer`* that corresponds to the Create Container operation in Azure.

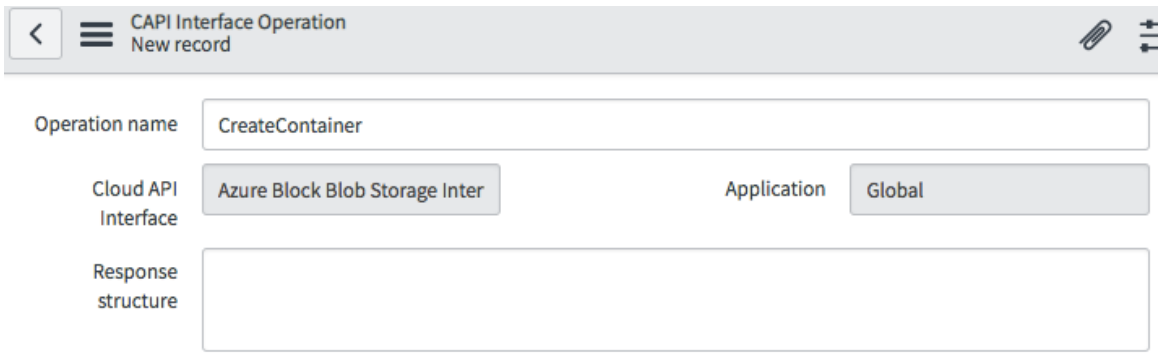

6. Right-click the header and select **Save**.
7. In the Interface Operation Parameters related list, click **New**.
8. Create a parameter that the operation needs as defined by the cloud provider REST API documentation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you do not know the parameters, the operation fails. In the Azure storage example, the *CreateContainer* operation must have the name of the container that is referred to here as *ContainerName*.

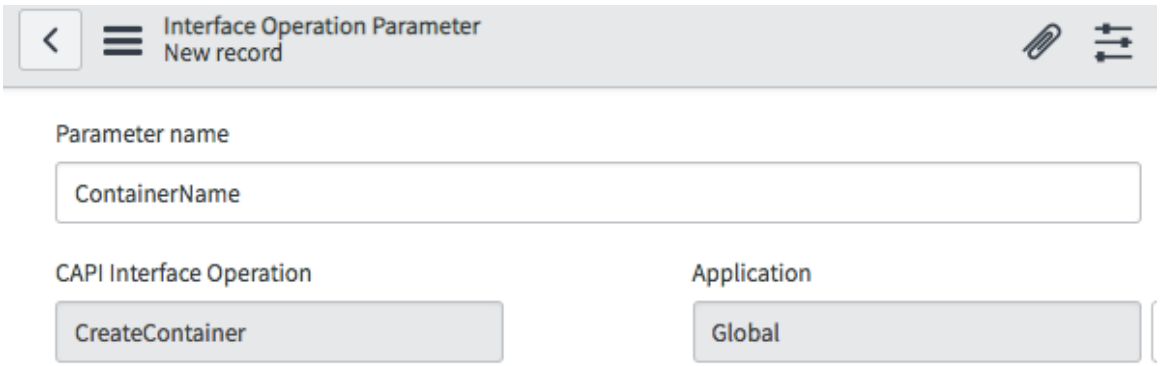

9. Click **Submit**.
10. Add as many parameters as you need for the operation, and then add all the operations you need.

What to do next
Create a CAPI API

Create a CAPI API

If the existing CAPI APIs do not integrate with the CAPI interface for a provider, you can create an API. The steps in resource blocks can call your API methods when operations are executed.

Before you begin

- Role required: admin
- Review Default CAPI APIs and make sure the provider and product that you are creating do not exist.
- Familiarize yourself with the cloud provider API so that you know how to create REST calls to the provider. If you are not familiar with the headers, authentication requirements, endpoint, and data that you can pass into the provider API, you cannot build your own CAPI components.

About this task
This topic provides instructions for creating a CAPI API. It also cover the configuration of CAPI, resource blocks, blueprints, and credential mappings.

Procedure

1. In the Cloud Admin Portal, navigate to **Design** > **Cloud API**.
2. Click the **API** tab.
3. Click **New**.
4. Fill out the form fields (see table):

Field	Description
Cloud API	Provide a descriptive name.
Interface	Select from among the existing interfaces that provide the operations you need.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Connector	**Script Connector** appears in this field if the connector is a Java connector that is available out-of-box and cannot be modified. You do not need to modify these script connectors.
Product	Select the product that is part of the provider. The product includes resource types in the CMDB.
Version	Provide a version number to identify this API. Enter `1.0` if this API is the first version. You can create several versions of the same API and then refer to the specific versions in the operation steps of resource blocks.
Scripted	This field is selected if the **Interface** value is **Script Connector**.
Script Type	Select the language for the script. **Note:** JavaScript is recommended. The method mappers use MID Server script includes written in Javascript.
ExecuteScriptOnInstance	
Description	Provide a description for this API.

5. Verify that all the operations you need appear in the **CAPI Method Mappers** related list.

   The operations in the CAPI Method Mappers list are the same operations in the interface that you selected. They are bound to this API. You cannot add or remove any operations from the list. Some of the operations, especially Azure operations, use MID Server script includes that perform the actual implementation for the API. The script includes a call to the REST API for the provider, such as AWS or Azure. You should be able to find the API in the provider's documentation.

6. To see the script include for a CAPI Method Mapper:

   a. Click the gear icon ( ) on the **CAPI Method Mappers** related list, add the **Request Script** field to the **Selected** column, and then click **OK**.

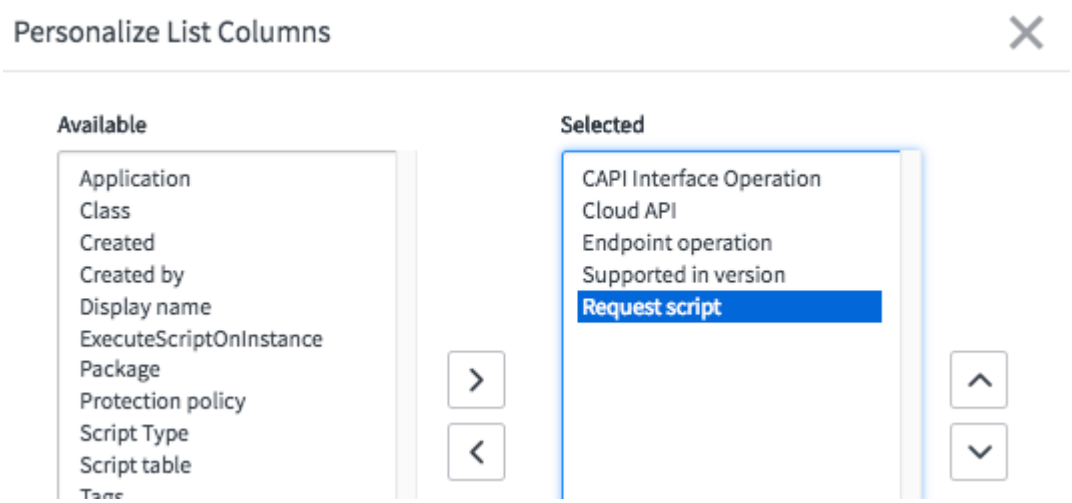

   The MID Server script includes now appear in the list of method mappers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Click the name of a script to view the script include.

In this example, the **CreateNode** operation uses the *azure-compute-1.0-CreateNode* script include.

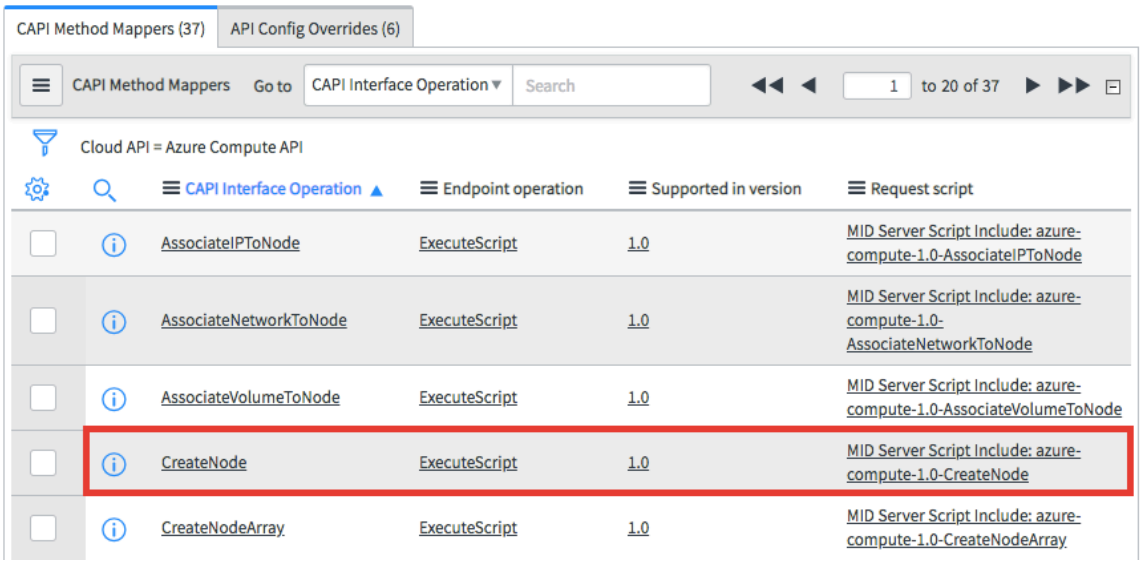

That script include contains the following code:

```
createNode();
function createNode(){
   var acp = new AzureComputeVirtualMachine(this.parameters,
 this.headers);
    //this passes the parameters and headers into the
 AzureComputeVirtualMachine script include.
   var vm = acp.createVirtualMachine();

   return vm;
}
```

Notice that the code calls the *AzureComputeVirtualMachine* script include, which is available by default with the Cloud Provisioning and Governance application. It contains all the API calls to Azure against virtual machines.

7. Create an API config override to pass the necessary credential information to the cloud provider.

You also need to pass header information required by the cloud provider REST API. API config overrides are necessary because credential data is not included when your instance makes the REST call to the cloud provider API.

Tip: If you do not know which API config overrides to create, look at a similar API that exists in CAPI, and use the same overrides. For example, if you are creating an API for something like AWS storage, look at other AWS APIs such as AWS Block Storage or AWS Compute.

a. In the API Config Override related list, click **New**.

b. Click the lookup icon and select an existing config parameter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Many useful parameters are provided by default, such as **Identity**, **Credentials**, and so on.

c. Enter an **Override value**.

You can set a static value, or you can pull a value from the system in the form of `$(CloudCredential.attribute_name)`. The `CloudCredential` portion unlocks it from the credential store. The `attribute_name` portion is an attribute that depends on the particular credential or alias or endpoint URL that you are trying to retrieve.

If you navigate to **Connections and Credentials** > **Credentials** and open your AWS credential, you can right-click the **Access Key ID** field and see the name of the column in the Credentials table where the access key is stored. In this case, the access key is stored in the `access_key` column. So your Override value should be `$(CloudCredential.access_key)`.

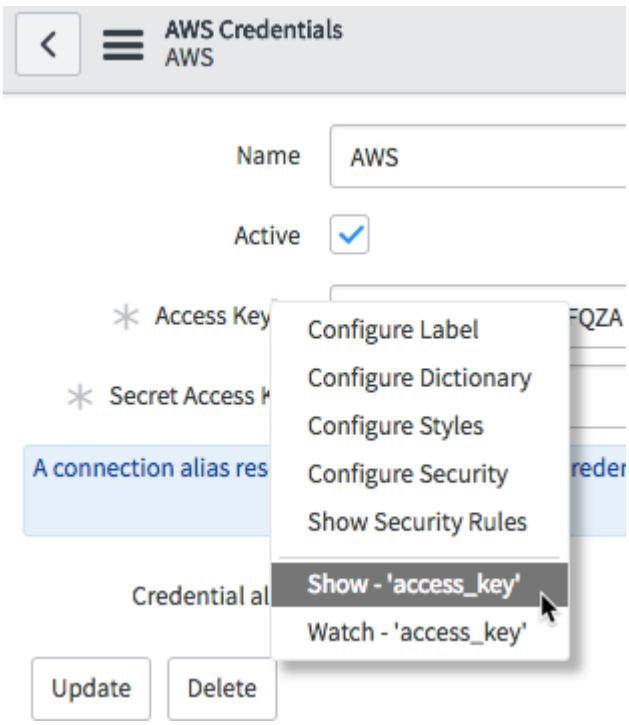

8. If you need to create a configuration parameter to use in the API Config Override, follow this procedure:

a. On the API Config Override form, click the lookup icon in the **Config Parameter** field.

b. Click **New** on the Configuration Parameter window and then fill in the form.

Field	Description
Config parameter name	Enter a descriptive name.
Default value	Set an optional default value.
Endpoint	This field is read only.
Description	Enter a helpful description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Config data type	Select the type of data: ▪ **String** ▪ **DateTime** ▪ **Date** ▪ **Long** ▪ **Password** ▪ **Boolean**
Config type	Select the type of configuration: ▪ Header: Use this value to pass the override values as headers. If your have a MID Server script include as part of your API and it uses `this.headers` as a parameter, the override values are passed to the provider API. ▪ Config ▪ URI Config

9. If necessary, modify the script includes that execute the operations against the cloud provider API:

   a. On the Cloud API form, in the CAPI Method Mappers related list, click the information icon ( ) next to the CAPI interface operation.
   Do not click the CAPI Interface Operation name itself.

   b. Click **Open Record**.

   c. On the operation form, you can modify the following field values:

Field	Description
Endpoint operation	Keep the default **Execute Script** endpoint operation.
Supported in version	You typically do not need to change this value. Change the **Supported in** version to use the method mapper in a different version of the API. By default, the method mappers are set to the same version number as the API they belong to.

   d. In the **Request Script** field, click the information icon ( ) next to the placeholder MID Server script include to modify it.

   e. On the script include form, create your script.

   The following example JavaScript lists Azure load balancers. Note the helpful comments:

```
getResourceInfo();

function getResourceInfo() {
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
var anpi = new AzureLoadBalancer(this.parameters,
this.headers);

      //this calls the 'AzureLoadBalancer' MID Server script
include, which in turn
      //calls the 'AzureCloudAPIBase' script include, both
of which
      //are already in your instance by default.
      //The 'AzureCloudBase' script include runs a query for
resources of a given type in a
//specified region.

var items = anpi.listResourcesByRegionOrIDs();
      //This function is in the 'AzureCloudAPIBase' script
include. It calls the Azure API.

return items;
}
```

The overrides provided for this operation are what Azure needs to find the load balancers: the Client ID, Tenant ID, the Secret Key, and the Endpoint URL:

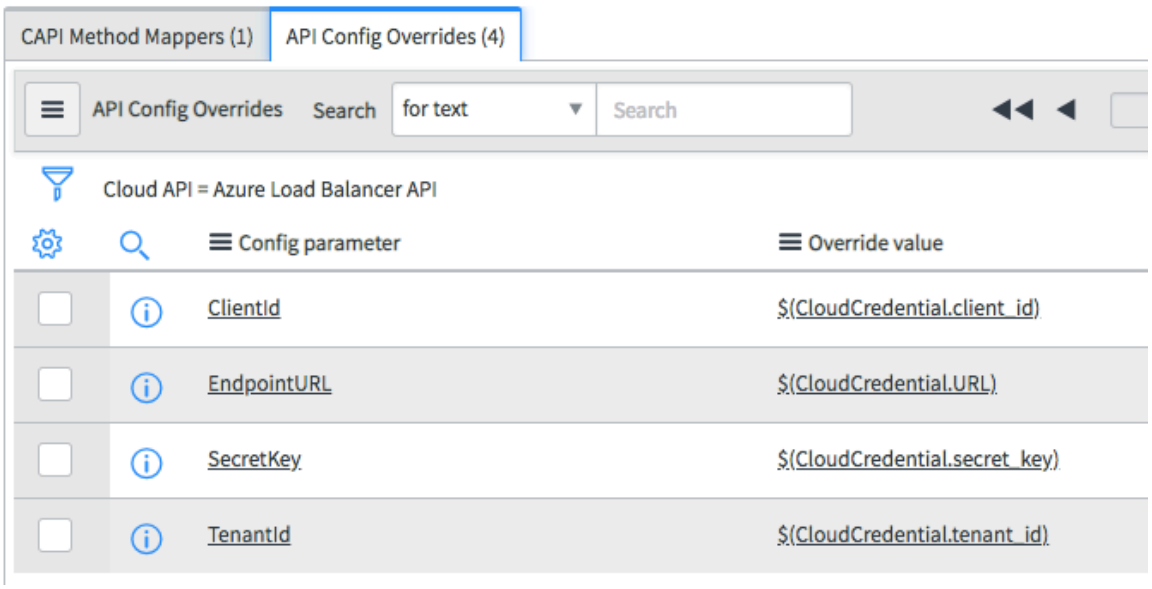

For the Google Cloud integration example, a MID Server JAR file is also needed.

f. Click **Submit**.

g. Modify the script include for each operation.

Cloud User Portal

The Cloud User Portal gives you immediate access to all day-to-day cloud activities.

Cloud User Portal

Use the Cloud User Portal to perform the following actions:

- Monitor your quota, costs, budget, life-cycle events, stack health, and requests.
- Request stacks from the service catalog and track requests.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Request life-cycle operations for stacks and resources (for example, stop, start, or deprovision).
- Create and track incidents.

Using the 'Overview' page

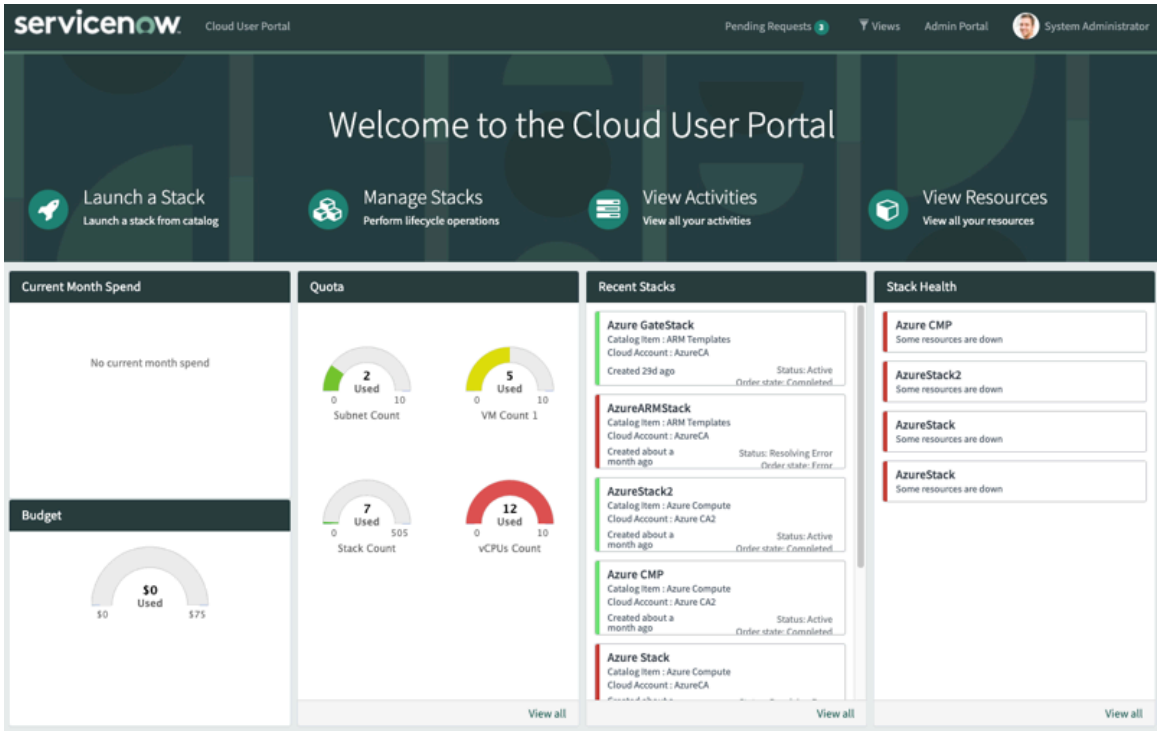

Tasks on the 'Overview' page

Launch a Stack	Open the service catalog of items that you can request. You request an item and, when your request is approved, the system launches the stack. See Launch a stack.
Manage Stacks	View the stacks assigned to you. See Manage a stack.
View Activities	View activities for your stacks and resources. Activities include stack requests, change requests, incidents, cloud events, and SSH keys. See View Activities.
View Resources	Opens the Resource Details page to view details and status of a resource. See Manage a resource.
Pending Requests	When you submit a request for a stack, the **Activities** page displays the progress and status of the stack request, and a **Pending Requests** link and count appears. Click the link to view a summary of pending stack requests. See Track a stack request.

Important: Starting with the Brazil release, the Current Month Spend and Budget widgets are deprecated for non domain-separated instances.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Quick-look status reports on the 'Overview' page	
Views	A filter to view and manage your resources as well as the resources assigned to the group you belong to.
Current Month Spend	A chart that shows costs, by category, for your cloud resources.
Budget	A chart displaying an arc based on your consumption of budget allocated to you.
Quota	A graph shows your quota consumption for quota-defined resource types such as VM Count, vCPUs Count, Storage Volume Size, Network Count, and Stack Count. You can view the quota that you have consumed compared to your maximum allocated quota. See Viewing resource quota limits.
Recent Stacks	Ten stacks that you most recently requested. Each entry shows the **Name**, **Catalog item** requested, **Cloud account** used to request the stack, and the order status of the requested stack. Click a **Name** to view stack details.
Stack Health	Shows the most recent 10 items: Incidents created for provisioned stacks or active stacks with a resource turned off. Also shows the resources which are down. Click an item to view stack details. Color codes for incident type: • Red: Outage. • Orange: Performance issue or Something is broken. • Grey: Question or Request.

Using the task bar

From the task bar, you can perform the following actions:

Overview	Return to the home page of the Cloud User Portal.
Catalog	View the Service Catalog of items that you can request to launch stacks. When you select a **Service Category** (VM, Compute, Network, Storage, and so on), the **All Items** section lists the items in the selected category. Click **Launch** to request an item. See Launch a stack.
Stacks	Lists all stacks that are assigned to you. Select a catalog item to view the stacks that the item launches. Click a stack to view details—all resources, resource properties, activities, and life cycle operations. See Manage a stack.
Resources	Lists individual resources of the provisioned stacks that are assigned to you, regardless of who requested the stack. See Manage a resource.
Activities	Stack requests, change requests, cloud life cycle events (if configured), and incidents that are associated with your stacks and resources. SSH keys that you created or that were generated for you. See View Activities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Search the Cloud User Portal

The Search page returns search results and provides filters that help you to zoom in on the results.

Before you begin

Role required: sn_cmp.cloud_service_user

Procedure

1. On any portal page, enter text in the **Search** text box and press **Return**.
2. To filter the search results on the **Cloud Search** page, click a category.
   In this example, only **Stacks** that include the search text s appear in the **Search results** list:

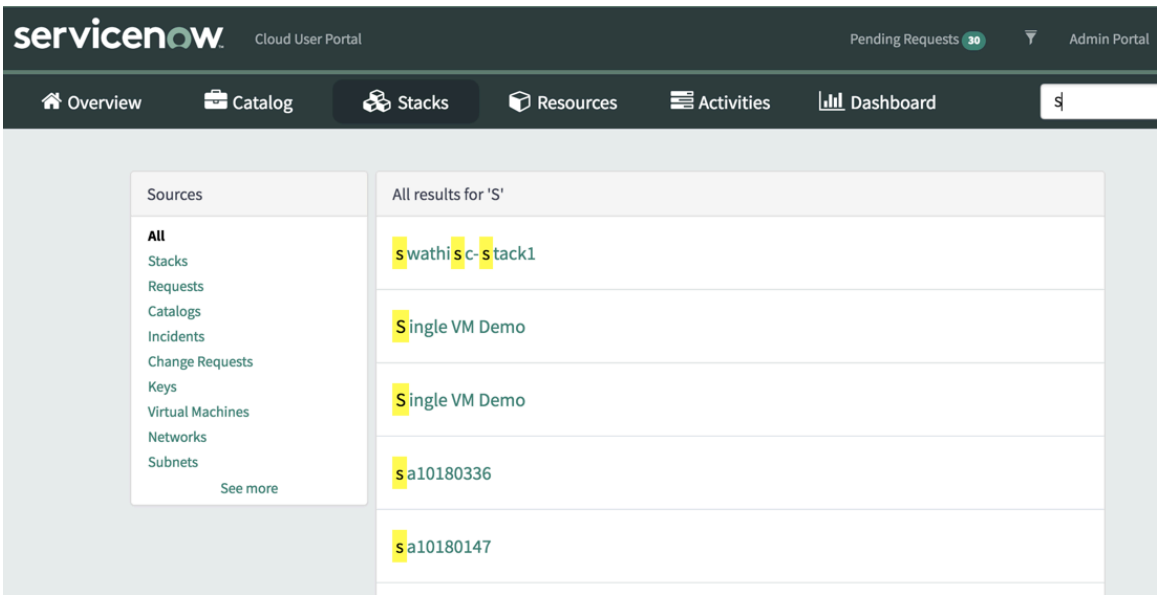

Launch a stack

Launch a stack of cloud resources to submit a request for an item in the service catalog. You can launch a stack from the Cloud User Portal overview page.

Before you begin

Role required: sn_cmp.cloud_service_user
When you submit a request for an item in the service catalog, the system typically provisions the stack immediately. Sometimes, your request must be approved first.

If a resource on a stack requires credentials when being accessed via SSH, then you must create an SSH key pair before you can launch the stack. See Create an SSH key and Download an SSH key.

Procedure

1. Open the request form using either of the following methods:
   - On the overview page, click **Launch a Stack**. By default, all catalog items in all service categories are listed. Select a **Service Category**, and then click a catalog item **Name**.
   - Open the **Catalog** page. By default, all catalog items in all service categories are listed. Select a **Service Category** and then click **Launch** on the catalog item.
2. In each section of the form, enter the values that configure your request.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The form fields differ for each catalog offering. Contact your Cloud Provisioning and Governance administrator if you have any questions about the fields.

Important: When you select a user group on the form, the quota limits for you and the user group (that you are part of) are shown in the Quota section.

The following sub-steps are applicable only if you are using the Terraform connector and Terraform config installables.

a. On the General Info section, select **Yes** for the *UseExistingWorkspace* option, to use existing workspaces that are already discovered.

b. Select the Workspace where your config installable resides, from the **Workspace** list.

Note: If you are configuring a Windows VM, set the **IncludeManagementKey** field to **No**. After the VM is provisioned, you can use the day-2 operation called InfuseKey to add key credentials (**UserName** and **Password** values) to the VM.

3. **Optional:** On the General Info section, select the Agent pool you have created in your Terraform Enterprise organization to communicate with private or on-premise infrastructure.

Note: Agent pool selection is applicable only to VMware via Terraform Enterprise setups for provisioning. You need to have created agent pools in your Terraform Enterprise organization and have VMware-Terraform Enterprise templates in your workspaces for the agent pools option to appear on the GUI.

4. Click the section header to move to the next section.
The remaining quota is updated based on the user group that you select. For required fields, the system selects the first item in the list by default and you can select any value that you need.
5. Click **Submit** to request the stack.
   - The **Activities** page opens to the **Status** tab on the **Track** > **Requests** activity. The tab displays the progress and status of the request. The following color codes show the stage of the request:
   - ▪ Green: Completed or Approved.
     ▪ Orange: Processing.
     ▪ Red: Rejected, Error, or Canceled.
   - If approval is required, then after the administrator approves your request, the system launches the stack and updates the **Request Status**.
   - If the quota limit exceeds for you or your user group, then an error message appears or the system triggers a policy-based approval. For more information about the quota limits, see Viewing resource quota limits.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Select another request to view its **Request Status** and **Stack Status**.
- To view a summary of all your pending requests, click the **Pending Requests** link in the title bar at any time.

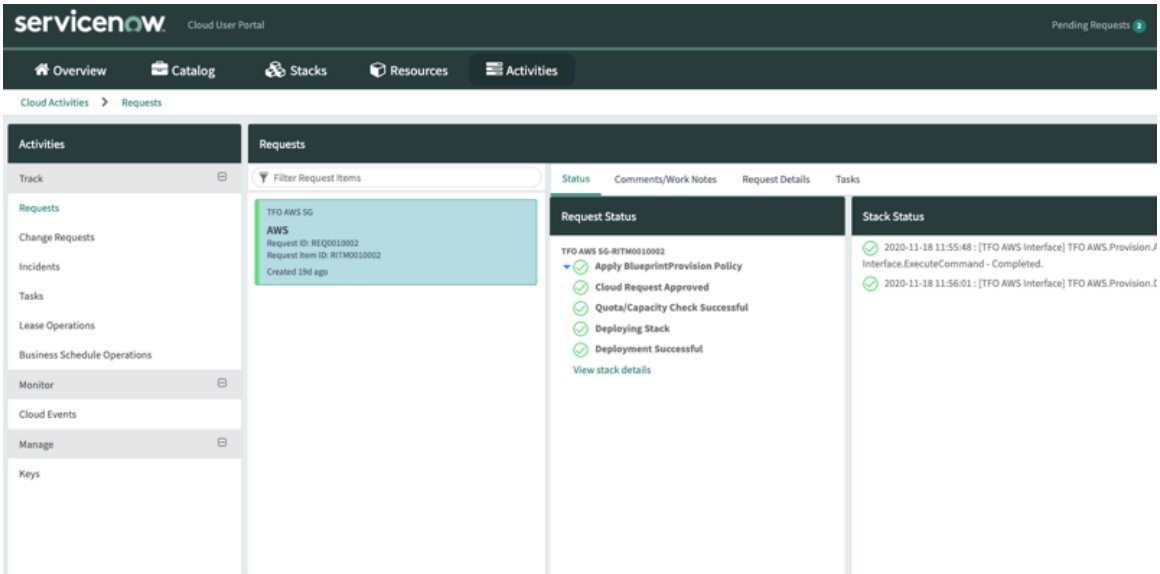

Viewing resource quota limits

View quota limits to see how many resources you have consumed and how many you can provision based on quota limits set for you and your user group.

Allocated and consumed quota limits

To see the quota you have consumed for provisioning resources compared to the maximum quota limit allocated to you, navigate to the Overview page. The following predefined resource types in the system might have quota limits set for you and the user groups to which you belong:

- Stack Count
- VM Count
- vCPUs Count
- Storage Volume Size
- Network Count

Note: The Stack Count is a default base system quota limit available beginning with the Brazil release. It is displayed on the Overview page if your Cloud Provisioning and Governance administrator has set a Stack Count quota for you or your user group.

The Quota section on the Overview page displays the quota limits in a graph.

- Green indicates that sufficient resources are available.
- Red indicates that insufficient resources are available.
- If quota limits have been exceeded, the system might either display an error message or trigger a policy-based approval when you attempt to provision a stack. For more information, see Provisioning stacks based on quota limits.

The quota limit that is displayed on the Quota section is an aggregate from all of the user groups to which you belong. To view quota limits separately for you and each of your user groups, click **View All**. The Quota Details pop-up window is displayed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Quota Details pop-up window

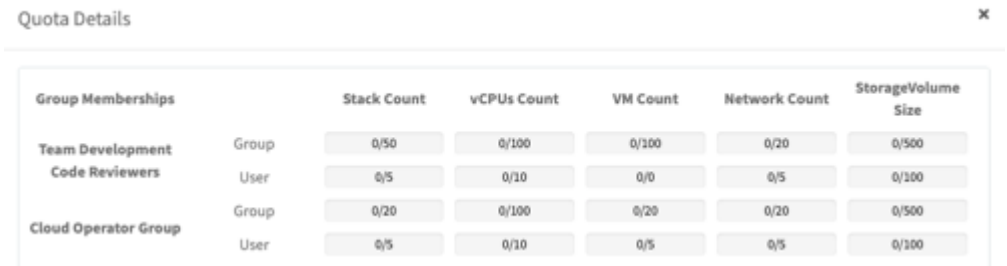

The Quota Details pop-up window indicates your consumed quota compared to your allocated quota limit for each resource type. For example, 0/5 for a resource type means that zero has been consumed and five is the allocated quota limit.

The quota limit for a resource type appears as either a sum or count. For example, the Storage Volume Size is shown as a sum in GB and the VM Count is shown as a number.

You can also view quota limits when launching stacks. For more information, see Launch a stack.

Provisioning stacks based on quota limits

The system calculates the quota allocated to you and the user groups to which you belong when you provision stacks. If quota limits are exceeded, the system either displays an error message or triggers a policy-based approval.

To view quota limits, navigate to the Overview page. For more information, see Viewing resource quota limits.

Quota calculation for template-based stacks

The Stack Count is a default base system quota limit available beginning with the Brazil release. It applies to all template-based stacks in the service catalog.

If your Cloud Provisioning and Governance administrator has set a Stack Count quota for you or your user group, the system calculates the stack number when you provision a template-based stack. After provisioning, the system increases the Stack Count quota by a value of 1 for each provisioned stack.

Stack provisioning after exceeding quota limits

If the quota limit set for you or your user group has been exceeded, the system either displays an error message or triggers an assigned policy when you attempt to provision a stack.

- If the system displays an error message indicating that the quota limit has been exceeded, you cannot provision the stack. If you think the quota limit should be increased, contact your Cloud Provisioning and Governance administrator.
- If a policy is defined in the system for exceeding the quota limit, that policy is triggered. Based on the policy, the system raises an approval request or notification for increasing the quota limit. The request follows a predefined subflow in the system for approval. If the request is approved, the system deploys the stack.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Policy triggered for provisioning a stack after exceeding the quota limit

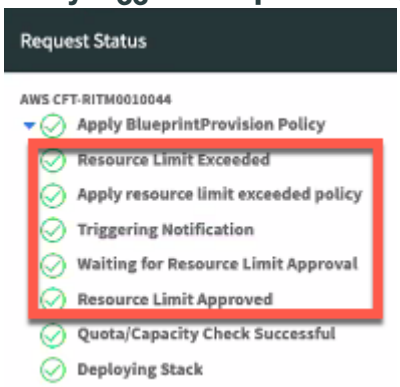

Track a stack request

The system generates a request when you submit the form to launch a stack.

Before you begin

Role required: sn_cmp.cloud_service_user

About this task

If no approval policy is defined, requests are auto-approved. If a policy is defined, then a user on the approver list must approve the request.
If the quota limit exceeds for you or your user group, then an error message appears or the system triggers a policy-based approval. For more information about the quota limits, see Viewing resource quota limits.

Procedure

1. Go to the **Activities** page or navigate to **Overview** > **View Activities**, then select **Track** > **Requests**.
   The following color codes show you the status of the request:
   - Green: Completed or Approved.
   - Orange: Processing.
   - Red: Rejected, Error, or Canceled.
2. In the **Requests** list, click an entry.

Status tab	
Request Status	Actions taken on the request. Click the down-arrow (▼) to toggle between: ○ Verbose mode that describes each action and the status of the action ○ Icon-only mode that indicates only the status (pass/fail) of an action.
Stack Status	Status of each stack request. If deployment fails, errors from cloud orchestration are listed.
Comments / Work Notes tab	
Item requested and request number	Log of comments and notes made for the request. The **itil** role is required to add a comment or work note. To add a comment or work note:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

	a. Enter text. **b.** Select the **Work notes** check box to add a work note. **c.** Click **Send**.
Request Details tab	Values that the requester submitted on the stack request form.
Tasks tab	Lists all the catalog tasks created for stacks that failed to provision. Click a catalog task to view more details about that task. See Resubmit a failed stack request.

Resubmit a failed stack request

The Cloud Provisioning and Governance application creates a remediation or a catalog task when a request for a stack fails to provision. You can resubmit the failed request or assign it to the Cloud Operator group to handle it for you.

Before you begin

Roles required: sn_cmp.cloud_service_user and sn_cmp.cloud_operator.

About this task

Task remediation captures every remediation applied to a failed request. You can view the remediation carried out for a failed request and all the major milestones in the **Comments/Work Notes** sub tab in the Cloud User Portal. Task remediation is available only for new provisioning of stacks and not for life cycle or any other operation.

A stack may fail to provision for many reasons. A stack may fail if you have exceeded your cloud quota, datacenter capacity, or have entered an invalid parameter in the request form (such as a duplicate NIC).

If a stack fails to provision, a catalog task is created and assigned to you. An error message pertaining to the failed request is displayed in the **Status** sub tab in the Cloud User Portal.

Procedure

1. In the Cloud User Portal, navigate to **View Activities** > **Requests** > **Tasks**.
2. Click the catalog task to resubmit.
3. Based on your preference, select an action:
   - **Retry**: Resubmits the failed request. You can try to resubmit the failed request by resolving the error if possible. If the error is due to exceeding the quota limit, you can try to deprovision some resources that are not currently in use and then click **Retry**. When you click **Retry**, the original request is resubmitted. A new request is not created.

     Note: Any policy associated with a resource operation is re-executed.

   - **Ask for Help**: Creates a new catalog task and assigns it to the Cloud Operator group to resubmit the failed request. The Cloud Operator group can look more closely into the error and try to resolve it for you. Every time the comments or work notes in the catalog task gets updated by the Cloud Operator group, the changes are reflected in the **Comments/Work Notes** sub tab in the Cloud User Portal. If necessary, the operator can cancel the failed request.
   - **Cancel Order**: Cancels the failed request and deprovisions all partially-provisioned resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Manage a stack

Use the Stack Details page to view details and status for a stack and to perform life-cycle operations on a stack.

Before you begin

Role required: sn_cmp.cloud_service_user

Procedure

1. Use one of the following methods to open the **Stack Details** page:
   - Click **Stacks** on the toolbar.
   - Use the **Search** text box to search for the stack.
   - On the home page, click a stack in the **Recent Stacks**, or the **Stack Health** section.
   - On the home page, click **Manage Stacks**.

   By default, all stacks in all service categories are listed.
2. In the **Stacks** list, select a stack category, and then click a stack **Name**.
   The **Stack Details** page displays the following information:

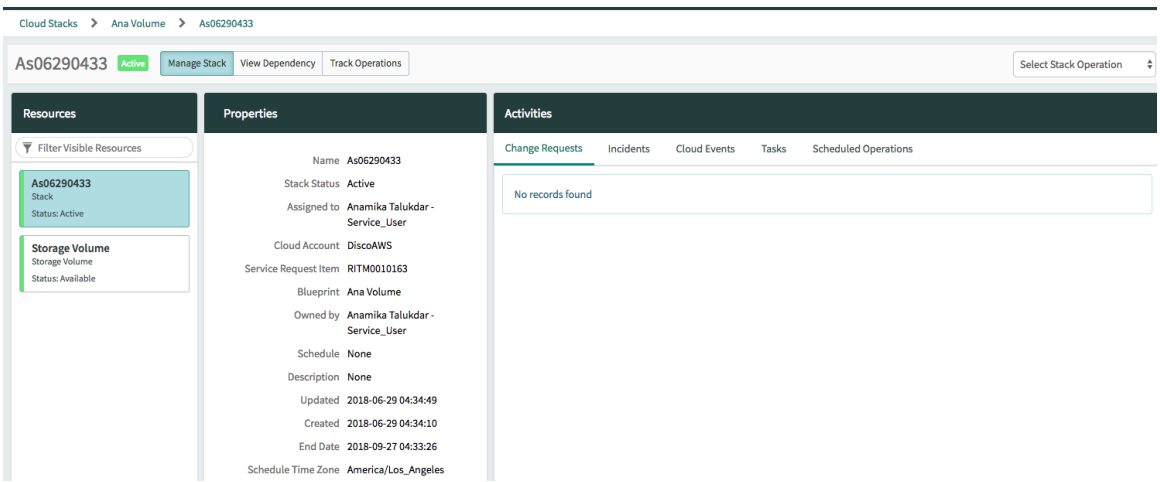

   - Breadcrumb navigation
   - Name of the stack. Click **View Dependency** to view the dependency map for the stack, **Manage Stack** to view the resources, properties, and activities of the stack, and **Track Operations** to view the list of lifecycle operations performed on the stack and its resources.
   - Status information
   - List of all resources in the stack. Click a resource to view details in the **Properties** section. Color codes indicate status of the stack:
     - Green: On/Active.
     - Yellow: Turned off or Processing.
     - Red: Terminated or Error.
   - List of the stack properties. Some properties are set in the stack request form and others are set by policy
   - Activities associated with the stack. Click a tab to view:
     - Change requests that are associated with the stack. See Track a change request.
     - Incidents that were raised for the stack. See Submit an incident for a stack and Track an incident.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Cloud events that are associated with the stack. Cloud Events appears only when a resource is selected. See View cloud events.
- Catalog tasks associated with the stack. See Resubmit a failed stack request.
- Upcoming business hours schedule operations associated with the stack. You can skip an upcoming operation or cancel a schedule associated with the stack. See Modify business schedule operations for a stack.

- Operations you can perform on a stack. See Perform a life-cycle operation on a stack or resource.
- Text box used to search the Cloud User Portal for stacks, resources based on types, stack requests, change requests, incidents, keys, and catalog items.

Stack properties

Name	Name that you specified when requesting the stack.
Stack Status	Current status of the stack: ○ Active (provisioned stacks) ○ Terminated (deprovisioned stacks) ○ Unmanaged (discovered stacks) ○ Error (errored stacks) ○ Processing (stacks being processed)
Assigned To	Owner of the stack configuration. This drives the access controls and views in cloud portal.
Cloud Account	Cloud account that the stack is associated with.
Service Request Item	Record created in Request Item table when you ordered the catalog which resulted in this stack.
Blueprint	Blueprint with which the stack is associated.
Owned by	User name of the requester.
Owner Group	Name of the group the owner of the stack belongs to.
Schedule	The schedule is associated with the stack.
Output	Details specific to ARM/CFT stacks, such as the URL of an application, names, or locations of any resources created.
Updated	Date and time the stack details were last updated.
Created	Date and time the stack was requested.
End Date	Date and time the stack terminates (unprovisioned).

Manage a resource

Use the Resource Details page to view details and status for a resource and to perform life cycle operations on a resource.

Before you begin
Role required: sn_cmp.cloud_service_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Click **Resources** on the toolbar.
   The **Resources** page lists all resource types in the left pane.

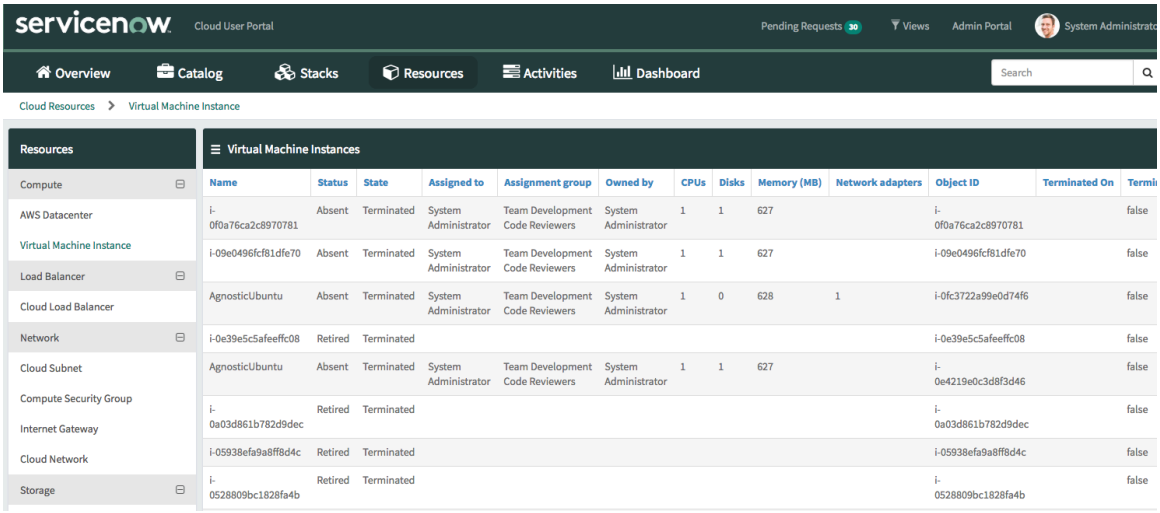

2. Navigate to the appropriate resource type to see the list of resources that you are interested in. Click a resource name to view specific details about that resource.

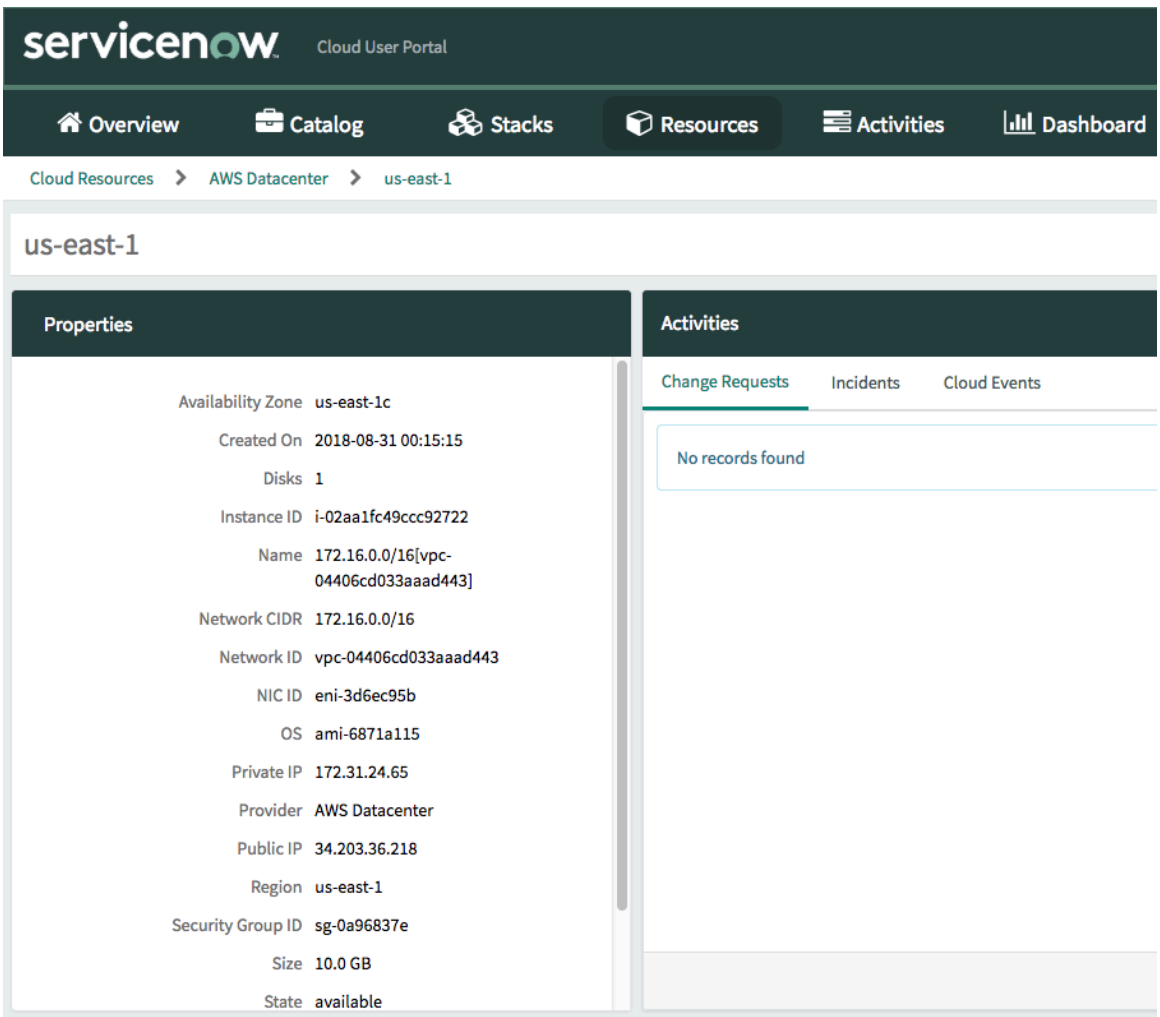

The resource details screen contains the following information:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Breadcrumb navigation
- All properties specific to the resource
- Activities associated with the resource. Click a tab to view:
  - Change requests that are associated with the resource. See Track a change request.
  - Incidents that were raised for the resource. See Submit an incident for a stack and Track an incident.
  - Cloud events that happened to the resource. See View cloud events.
- Text box used to search the Cloud User Portal for stacks, resources based on types, stack requests, change requests, incidents, keys, and catalog items.

Perform a life-cycle operation on a stack or resource

Operations like Start/Stop, Deprovision, and ExecuteScript are called life-cycle operations or day-2 operations. When you request a life-cycle operation on a stack or resource, the system generates a change request. An approval policy specifies either that the change is auto-approved or that a user on the approver list must approve the change.

Before you begin

Role required:

Life-cycle operations require the user with `sn_cmp.cloud_service_user` role who owns the resource. A user with the root_admin, admin, or `sn_cmp.cloud_admin` role can also perform the operation.

About this task

- As noted in the descriptions, not all operations are supported for all providers or for all service categories (resource types).
- All other configuration management providers: If a stack or resource requires login, then you must download the associated SSH key before you can access the resource. See Download an SSH key.

Procedure

1. Select a stack or resource.
2. While viewing the **Stack Details** page, select the operation from the **Select Stack Operation** or **Select Resource Operation** list.
   For for full descriptions of the operations, see Life-cycle operations in Cloud Provisioning and Governance.

   Some operations require settings (for example, the name of a script or new schedule settings).
3. Enter the settings and then click **OK**.
   The system submits a change request and the **Activities** > **Change Requests** page displays the progress and status of the request.

Life-cycle operations in Cloud Provisioning and Governance

Operations like Start/Stop, Deprovision, and ExecuteScript are called life-cycle operations or day-2 operations. When you request a life-cycle operation on a stack or resource, the system generates a change request. An approval policy specifies either that the change is auto-approved or that a user on the approver list must approve the change.

To perform an operation, you select a stack or resource and then click **Select Stack Operation** or **Select Resource Operation**. The list includes only appropriate operations for the selected

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

resource or stack. Some operations are provided by configuration management systems like Ansible or Terraform.

Operations in the base system

Start

The *Start* operation starts selected VM. The resource status changes to On.

No setting is required.

Stop

The *Stop* operation stops the selected VM. The resource status changes to Off.

No setting is required.

Deprovision

The *Deprovision* operation terminates the selected stack and sends a notification to the stack owner.

No setting is required.

Virtual Machine 'Store Extension Interface' operations

Only Google Cloud Platform and IBM Cloud support the *Virtual Machine Store Extension Interface* operations. The Terraform Connector ServiceNow Store app supplies the operations.

Virtual Machine Store Extension Interface.ExecuteScript

The *ExecuteScript* operation runs a script on the selected Windows or Linux VM. For example, a script might install an application on a resource. See KB0756431 for detailed instructions.

> **Note:** To run the *ExecuteScript* operation, a VM must include credentials. Linux VMs inherit management key credentials when they are provisioned. For Windows VMs, use the *InfuseKey* operation to add credentials to the VM.

In the pop-up, specify:

- **ScriptType**: Shell and PowerShell scripts are supported.
- **Script**: Name of the script.

Virtual Machine Store Extension Interface.InfuseKey

For Windows VMs, the *InfuseKey* operation adds key credentials (public and private key pair values) to the VM.

As a result, you can view the credentials on the Properties list of the Windows VM and can use the credentials to perform VM operations. For instructions on generating a key, see Create an SSH key.

To add key credentials to a Windows VM:

1. Select the Windows VM.
2. In the **Select Resource Operation** list, select **Virtual Machine Store Extension Interface.InfuseKey**.
3. In the pop-up, specify the key to use in the **UserKey** field and then click **OK**.

Virtual Machine Store Extension Interface.Execute Job Template

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Ansible supplies the *Execute Job Template* operation. The operation enables you to attach a node to an Ansible server and execute a package on that VME. For example, a script might install an application on a resource. In the pop-up, specify:

WorkloadConfigProviderType	
WorkloadConfigProvider	
Inventory	
HostGroup	The list includes all possible applications (called groups in Ansible), not just the applications that the configuration management provider supports. Therefore, if users select an unsupported application, the provisioning fails.
ApplicationProfile	Select the application profile to execute on the selected resource.

- **ScriptType**: For Google Cloud Platform, **Shell** is supported.
- **Script**: Name of the script.

Virtual Machine Store Extension Interface.Resize

The *Resize* operation changes the **Size** setting for a VM (called **HardwareType** in the CI definition). In the pop-up, specify the new size in the **HardwareType** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resize operation

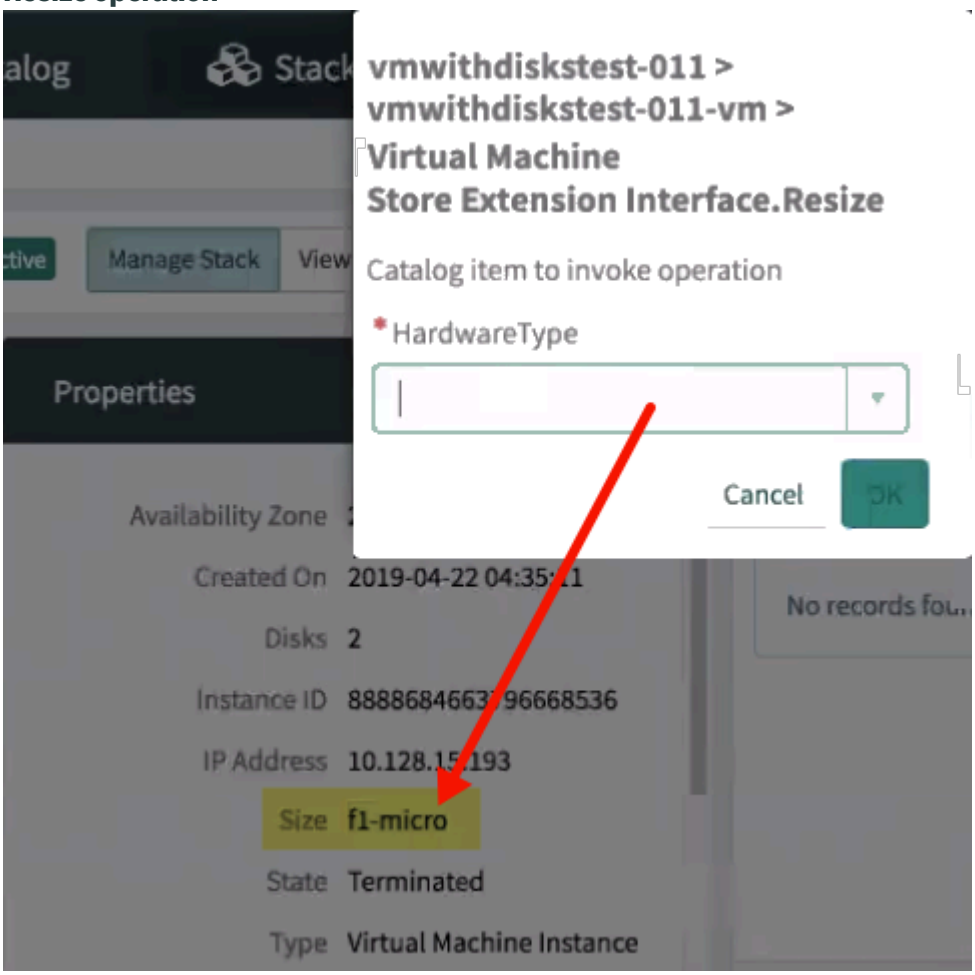

'Compute Security Group Store Extension Interface' operations

The *Compute Security Group Store Extension Interface* operations are supported for IBM Cloud only. The Terraform ServiceNow Store app supplies the operations.

Compute Security Group Store Extension Interface.Deprovision

The *Deprovision* operation terminates the selected compute security group resource and sends a notification to the owner.

No setting is required.

The **Status** of the network interface changes to **Terminated**. In the Stack Health section of the home page, the status of the stack indicates that some resources are down.

'Load Balancer Store Extension Interface' operations

Load Balancer Store Extension Interface.Deprovision

The *Deprovision* operation terminates the selected load balancer and sends a notification to the stack owner.

No setting is required.

The **Status** of the load balancer changes to **Terminated**. In the Stack Health section of the home page, the status of the stack indicates that some resources are down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

'Network Store Extension Interface' operations

'Network Store Extension Interface.Deprovision

The *Deprovision* operation terminates the selected network store and sends a notification to the stack owner.

No setting is required.

The **Status** of the network store changes to **Terminated**. In the Stack Health section of the home page, the status of the stack indicates that some resources are down.

'Storage Volume Store Extension Interface' operations

The *Storage Volume Store Extension Interface* operations are supported for IBM Cloud only. The Terraform ServiceNow Store app supplies the operations.

Storage Volume Store Extension Interface.Create Snapshot

The *Create Snapshot* operation generates a read-only copy of the specified volume.

In the pop-up, specify the **SnapshotName** to use as the label for the snapshot.

The snapshot is saved in the [cmdbci.ci_storage_vol_snapshot_list] table.

If approval is not required, the system immediately takes the snapshot. If approval is required, there may be a delay while an admin approves the change. You receive email when the admin approves. After the admin approves, return to the VM Instance form and click the **Proceed with Change** related link.

Note: If the VM is running, the VM is shut down at the start of the snapshot process and then restarted when the process finishes. The system auto-deletes the oldest snapshot when the **Snapshot eviction policy** limit is exceeded.

Storage Volume Store Extension Interface.Restore From Snapshot

The *Restore From Snapshot* operation restores the image from the specified snapshot.

Select the **SnapshotID** to use to restore the volume.

Storage Volume Store Extension Interface.Delete Snapshot

The *Delete Snapshot* operation deletes the selected snapshot.

Storage Volume Store Extension Interface.Attach

The *Attach* operation attaches a VM to the selected volume.

1. Select the volume to attach. The **Status** of the volume must be **Available**.
2. Select the **Storage Volume Store Extension Interface.Attach** operation.
3. In the **NodeId** field on the pop-up, select the **Instance ID** of the server to attach to the volume.

   The status of the volume changes from Available to In Use.

Storage Volume Store Extension Interface.Detach

The *Detach* operation detaches a VM from the selected volume.

1. Select the volume to attach. The Status of the volume must be In Use.
2. Select the **Storage Volume Store Extension Interface.Detach** operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the **NodeId** field on the pop-up, select the **Instance ID** of the host to detach from the volume.

   If this VM is the last or the only VM that is detached, then the status of the volume changes from In Use to Available.

Storage Volume Store Extension Interface.Deprovision

The *Deprovision* operation terminates the selected volume and sends a notification to the stack owner.

No setting is required.

- The billing item for the storage resource is removed.
- The **Status** of the volume changes to **Terminated**. In the Stack Health section of the home page, the status of the stack indicates that some resources are down.
- For IBM Cloud: The storage volume is available in the IBM Cloud Console for an extended period to enable the user to reclaim the storage volume.

Submit an incident for a stack

Submit an incident for a stack when it is not working as expected or you have questions. When the user responds or solves the issue, the incident is closed.

Before you begin

Role required: sn_cmp.cloud_service_user

About this task

You can submit an incident for a stack or for an individual resource.
You can also Create an incident through the incident application. Make sure that the **Category** field is set to the value Cloud Provisioning and Governance. The configuration item should be the stack name against which the incident should be raised. The caller name should be the name of the person for whom this incident is created.

Procedure

1. View the stack details page using either of the following methods:
   - On the **Stack Details** page, click the **Incidents** tab in the **Activities** section.
   - On the **Activities** page, click **Track** > **Incidents**.

   The **Activities** section lists all submitted incidents for the stack.

Incidents

Activities → Incidents	View and submit incidents of all types, for any stack or resource.
Stacks → Stack Details → Incidents	View and submit incidents for the selected stack.
Resources → Resource Details → Incidents	View and submit incidents for the selected resource.

2. Click **New Incident** and fill in the form:

Incident type	Type that best matches your need.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Stack / Resource [related fields]	Stack and resource that the incident applies to. ○ From the **Activities** page: Select the **Stack** from the list of active stacks. Leave the **Resource** field blank if the issue is not associated with a particular resource. ○ From the **Stack Details** page: The **Stack** value is the current stack. ○ From the **Resource** page: The **Resource** value is the current resource.
Describe the issue	Describe the issue that you are experiencing or enter your question. Give full information on the behavior that is the issue.
What is the business impact	Specify how the incident affects your work and the work of others, including your customers.
Attach screenshots, logs, and so on	Click or drop to add files that clarify the issue or question.

3. Click **Submit**.
   The system generates the incident. The **Incidents** tab on the **Activities** page displays a summary.

What to do next

To view a list of all submitted incidents, navigate to **Activities** > **Incidents**.

Using the Activities page

Use the Activities page to view and manage details of your cloud actions.

Track	
Requests	Your requests for catalog items (stacks). See Track a stack request.
Change Requests	Change requests that have been submitted for stacks or resources that are assigned to you. The system generates change requests when you perform any life cycle operations on a stack or resource. See Track a change request and Perform a life-cycle operation on a stack or resource.
Incidents	Incidents for stacks and their associated resources that are assigned to you. See Track an incident.
Tasks	Catalog tasks created when a request for a stack fails to provision. See Resubmit a failed stack request.
Lease Operations	View a list of all your stacks that are nearing their lease end dates. You can extend the end date of a lease for any of the stacks listed. See Modify the lease for a stack.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Business Schedule Operations	View a list of upcoming business hours schedule operations associated with stacks. You can skip an upcoming operation or cancel a schedule associated with a particular stack. See Modify business schedule operations for a stack.
Monitor	
Cloud Events	List of life cycle events that affect resources assigned to you or to a user in your group. See View cloud events.
Manage	
Keys	View existing SSH keys that are assigned to you or generate a key. See Create an SSH key.

Track a change request

The system generates a change request when you perform a life cycle operation on a stack or resource.

Before you begin
Role required: sn_cmp.cloud_service_user

About this task
Each change request is either approved by default or, if an approval policy is in place, is approved by a user that is on the approver list.

Procedure
1. Go to the **Activities** page or navigate to **Cloud User Portal** > **View Activities**, then select **Track** > **Change Requests**.

Number	Unique auto-generated number.
Configuration item	The affected CI.
Category	Cloud Provisioning and Governance
Short description	Text that was entered by the requester.
Approval	Status.

2. In the **Change Requests** list, click an entry to view full details.

Track an incident

Before an incident is resolved and closed, you can view the state of the incident.

Before you begin
Role required: sn_cmp.cloud_service_user

Procedure
1. View the Incidents page using either of the following methods:
   - On the **Activities** page, click **Track** > **Incidents**.
   - On the **Stack Details** page, click the **Incidents** tab in the **Activities** section.
   The **Activities** section lists all submitted incidents for the stack.
2. In the **Incidents** list, click an entry to view details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Number	Unique auto-generated number.
Category	Cloud Provisioning and Governance.
Subcategory	○ Performance Issue: issue is associated with operation performance. ○ Request: request for a change.
Short description	Text that the submitter of the incident entered.
Configuration item	CI of the resource for the incident.
Impact	○ Urgency, rated as follows: ○ **1 - High** ○ **2 - Medium** ○ **3 - Low**
State	○ **New**: recently created incident. ○ **Closed**: user in the **Assigned to** field responded to or solved the issue and marked the incident as closed. ○ **Canceled**: the 'Assigned to' user reviewed the issue, did not take action, and canceled the incident.
Assigned to	The user who must respond to the incident.

View cloud events

You can view the events that are generated from your cloud resources if your administrator configured Cloud Provisioning and Governance to monitor them. All events are listed on the Cloud Events tab on the stack details page.

Before you begin

Role required: sn_cmp.cloud_service_user

Procedure

1. View the cloud events page using either of the following methods:
   - On the **Activities** page, click **Monitor** > **Cloud Events**.
   - On the **Stack Details** page, click the **Cloud Events** tab in the **Activities** section.

   The **Cloud Events** section lists all events for the stack.

Created	Timestamp of the arrival of the event.
Source	Source of the event: ○ **azure**: Azure Alert ○ **aws**: AWS Config ○ **vmware**: VMware Events
Event name	Name that is provided by Amazon, Azure, or VMware.
Subject	Text that describes the event.
Event time	Timestamp of the event.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource ID	Unique ID of the resource that received the life cycle state change or configuration change event.

2. Click an entry to view full details.

Create an SSH key

A stack might include a virtual machine that requires an SSH key for access. To request such a stack, you can either use a key pair that is assigned to you or you can generate another key.

Before you begin
Role required: sn_cmp.cloud_service_user

Procedure

1. On the Cloud User Portal, click **Activities**.
2. In the **Activities** list, select **Manage** > **Keys**.
3. On the **SSH Key Pair** page, click **Generate Key Pair**, specify a **Name** for the key, and then click **OK**.

Result
The system generates the key pair with the **Name** that you specified.

Download an SSH key

If a resource on a stack requires login, then you must download the associated SSH key before you can access the resource.

Before you begin
Role required: sn_cmp.cloud_service_user

Procedure

1. On the **Cloud User Portal**, click **Activities**.
2. In the **Activities** list, select **Manage** > **Keys**.
3. Click a key in the **SSH Key Pair** list and then click **Download Private Key** in the **SSH private key** section.

Business hour scheduling

You can set up business hour scheduling on all virtual machines in a stack. Business hour scheduling can be set up on stacks at the time of provisioning or on existing stacks.

Schedules can be set for when a stack should be started, stopped, and/or deprovisioned by defining parameters such as business hours, holidays, long weekends.

Note: For Google Deployment Manager (GDM) based template catalog, business hour scheduling and stack level operations are not supported.

The administrator can create schedule profiles in the **Cloud Admin Portal**. See Resource Profiles.

While provisioning a stack in the Cloud User Portal, use the **Schedule Profile** list to select a profile that has already been created by the administrator. Select a time zone in the **Schedule Time Zone** list which defaults to your system's time zone and indicates the time zone in which the schedule's business hour time frames are enforced.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can also define the lease end dates for a stack during provision. The lease end date for a stack can be modified at any time. See Modify the lease for a stack.

Once you provision a stack with a schedule and/or lease dates, a series of operations (start, stop, deprovision) occur at the appropriate time set by you. You can skip a business schedule operation or cancel the business schedule operations associated with a stack. See Modify business schedule operations for a stack.

Modify the lease for a stack

When a stack reaches the end of its lease, the system notifies the stack owner and deprovisions the stack. You can modify the lease for a stack before it approaches its end date.

Before you begin
Role required: sn_cmp.cloud_service_user

About this task
A lease end date is assigned to all provisioned stacks. A built-in policy titled *Lease End ServiceNow* with an *on Lease end* policy trigger has rules defining notifications and operations that are performed before, at, and after the end date of the lease. Lease dates can be modified for active stacks. The system sends a notification to the stack owner one day prior to the end of the lease. On the day the lease ends, the system stops the stack and sends another notification. Seven days after the lease has ended, the stack is terminated and a final notification is sent to the stack owner. You can modify the actions taken and the timings of these actions by modifying the *Lease End ServiceNow* policy. You can modify the default lease duration for all new stacks, going forward, by modifying the associated action script in the *Scheduling Defaults* built-in policy.

Procedure

1. Navigate to **All** > **Overview** > **View Activities**, then select **Lease Operations**.
   A list of stacks with upcoming lease-related operations is displayed.
2. Click the **Select Action** list associated with the stack and then select **Modify Lease**.
   The **Modify Lease** option is also available from the stack details list in **Manage Stacks**.
3. Set a new date and time for the lease to expire and then click **OK**.

Modify business schedule operations for a stack

Business hours scheduling can be performed on stacks at the time of provisioning or on existing stacks.

Before you begin
Role required: sn_cmp.cloud_service_user.

About this task
Once a stack is set up with a business schedule, the system generates a series of tasks for starting and stopping the stack. You can view these tasks and can modify the business schedule for any stack.

Procedure

1. Navigate to **All** > **Overview** > **View Activities**, then select **Business Schedule Operations**.
   A list of stacks with upcoming business schedule operations is displayed.
2. From the **Select Action** list associated with the stack, perform either of the two operations:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Skip Operation**- Click to skip a single business schedule operation on a stack. A confirmation message appears. Click **OK**.
- **Cancel Schedule**- Click to cancel business hour scheduling on a stack. A confirmation message appears. Click **OK**.

For modifying the schedule of an existing stack, navigate to the stacks details list in **Manage Stacks** and select **Modify Schedule** from the **Select Stack Operation** list.

Manage personal and group resources

You can view and manage your resources as well as the resources assigned to the group you belong to. If you belong to more than one group, you can view and manage all resources assigned to those groups.

Before you begin

Role required: Root administrator [sn_cmp.cmp_root_admin] and System administrator [admin].

About this task

Only a Root Administrator and a System Administrator have access to all three filters: **Mine**, **My Groups**, and **All**.

Procedure

1. On the Cloud User Portal, select **Views** filter, towards the right side on the banner.
2. Select an option:
   - **Mine**: View and manage only resources that are assigned to you.
   - **My Groups**: View and manage all resources that are assigned to you and that are assigned to all members of all groups that you belong to.

Schemas of Cloud Provisioning and Governance tables

The tables are cloud-agnostic and can therefore hold data for any cloud provider.

Load balancer table

The load balancer table [cmdb_ci_cloud_load_balancer] extends from [cmdb_ci_vm_object], which extends from [cmdb_ci]. Click to enlarge the image.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Schema load balancer

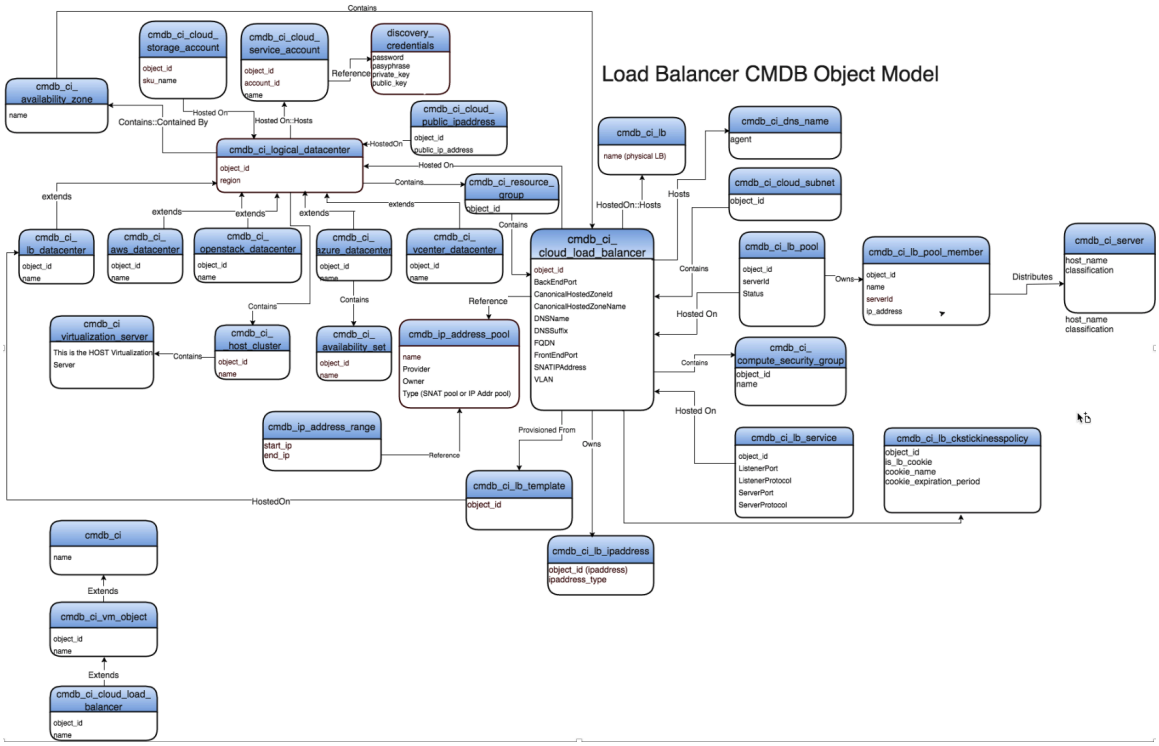

Key native attributes of the cmdb_ci_cloud_load_balancer table

Attribute	Description
object_id	Identifier that typically holds the load balancer name as the value.
canonical_hosted_zone_id	ID of the Amazon Route 53 hosted zone for the load balancer.
canonical_hosted_zone_name	DNS name of the load balancer.
dns_name	Public DNS name of the load balancer.
fqdn	DNS name as fully qualified domain name. Can also be a CNAME record pointed to the DNS name.
state	State of the load balancer: **available** or **terminated**.

Network resource table

The network resource table [cmdb_ci_network] extends from [cmdb_ci_vcenter_object], which extends from [cmdb_ci_vm_object], which extends from [cmdb_ci]. Click to enlarge the image.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network table object model

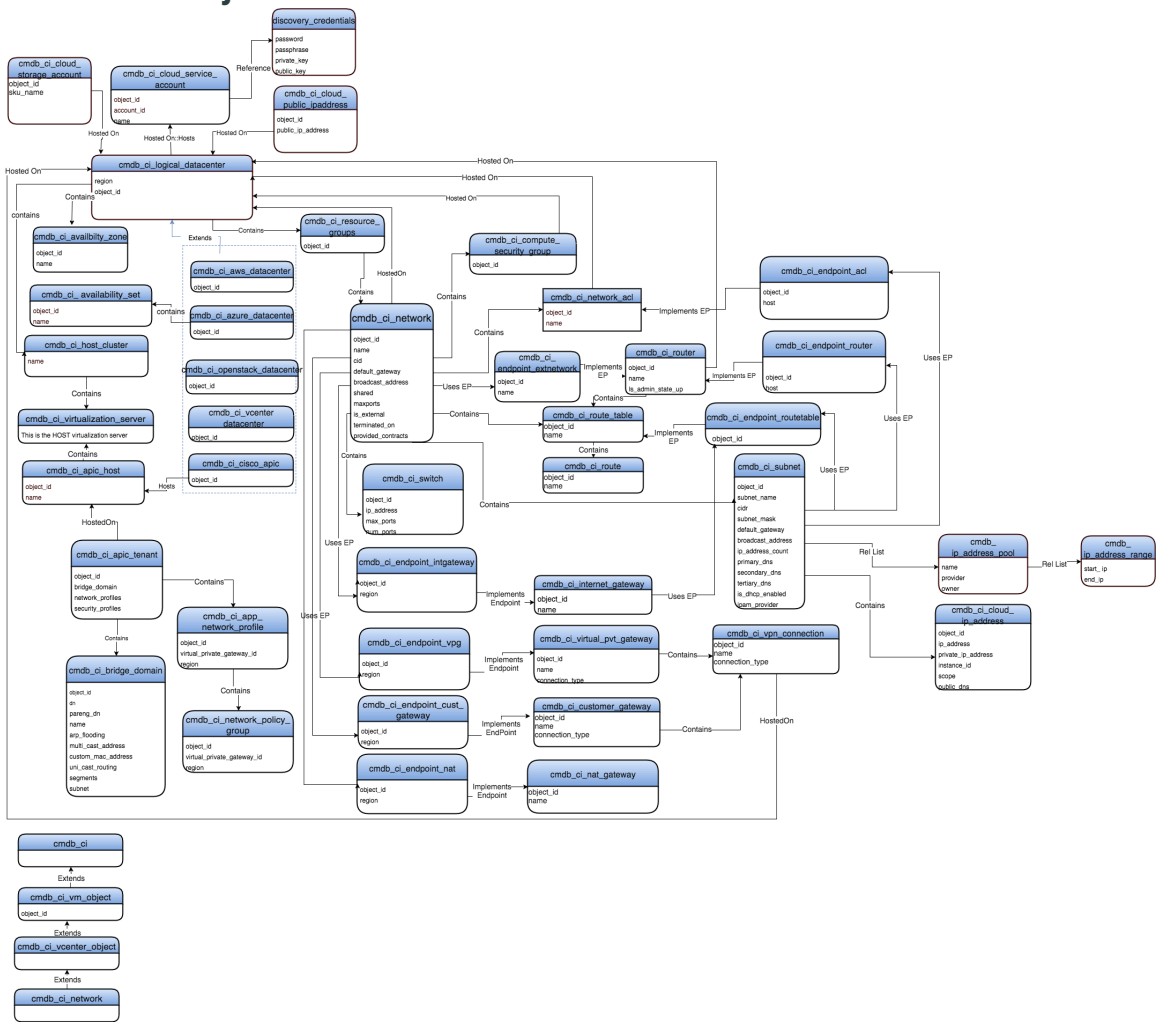

Key native attributes of the cmdb_ci_network table

Attribute	Description
object_id	Identifier that typically holds the network ID as the value. Uniquely identifies an object within a cloud.
name	Name of the network.
cidr	IP Address range. Classless inter-domain routing is a set of internet protocol standards that is used to create unique identifiers for networks.
default_gateway	Holds InternetGatewayID if we attach IntenetGateway to the network. A default gateway serves as an access point or IP router that a networked computer uses to send information to a computer in another network or Internet. The specified gateway is used by default unless an application specifies a different gateway.
Broadcast_address	IP address used to transmit messages and data packets to network systems.
Is_shared	Boolean. true: Network shared across other projects.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key native attributes of the cmdb_ci_network table (continued)

Attribute	Description
	false: Network not shared across other projects.
max_ports	Maximum number of hosts that can be connected to the network
is_external	Boolean. true: Network is external. false: Network is internal.
terminated_on	Time that the network was de-provisioned.
state	State of the network: **available** or **terminated**.
netmask	Type of CIDR. 32-bit mask that divides an IP address into subnets and specify the hosts that are available on the network.
dhcp_enabled	Boolean: true: Dynamic IP address is assigned to host false: Static IP address is assigned to host

Storage volume resource table

The storage volume resource table [cmdb_ci_storage_volume] extends from [cmdb_ci]. Click to enlarge the image.

Schema storage volume

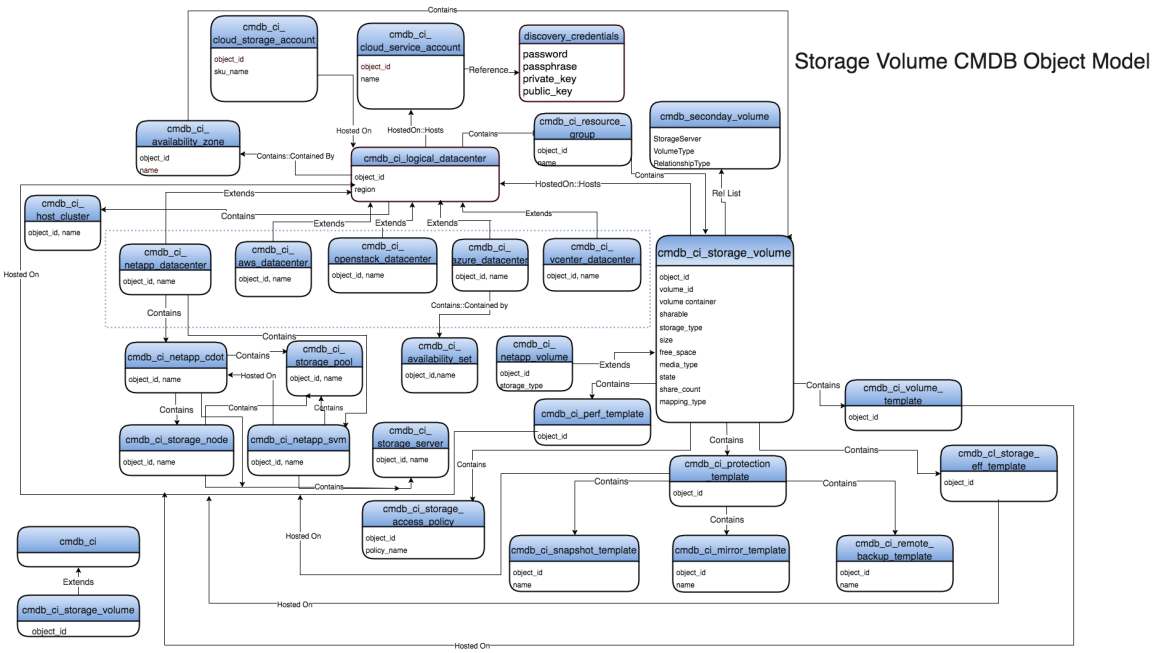

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key native attributes of the cmdb_ci_storage_volume table

Attribute	Description
object_id	Identifier that typically holds the object ID as the value. Uniquely identifies an object within a cloud.
volume_id	Identifier that typically holds the volume ID as the value. Uniquely identifies an object within a cloud.
volume_container	For NetApp only, LUN becomes the volume and NetApp volume becomes the volume container.
sharable	Boolean. true: Volume is shared by multiple VMs. false: Volume is not shared by multiple VMs.
storage_type	Type of storage. • AWS: Block • Azure: PageBlob • vSphere: VMware vdisk
size	Total capacity of the volume.
free_space	Available space of the volume
state	State of the volume: **available** or **in_use**.
share_count	Number of VMs that are shared by the volume.

Virtual server resource table

The virtual server resource table [cmdb_ci_vm_instance] extends from [cmdb_ci_vm_object], which extends from [cmdb_ci]. Click to enlarge the image.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

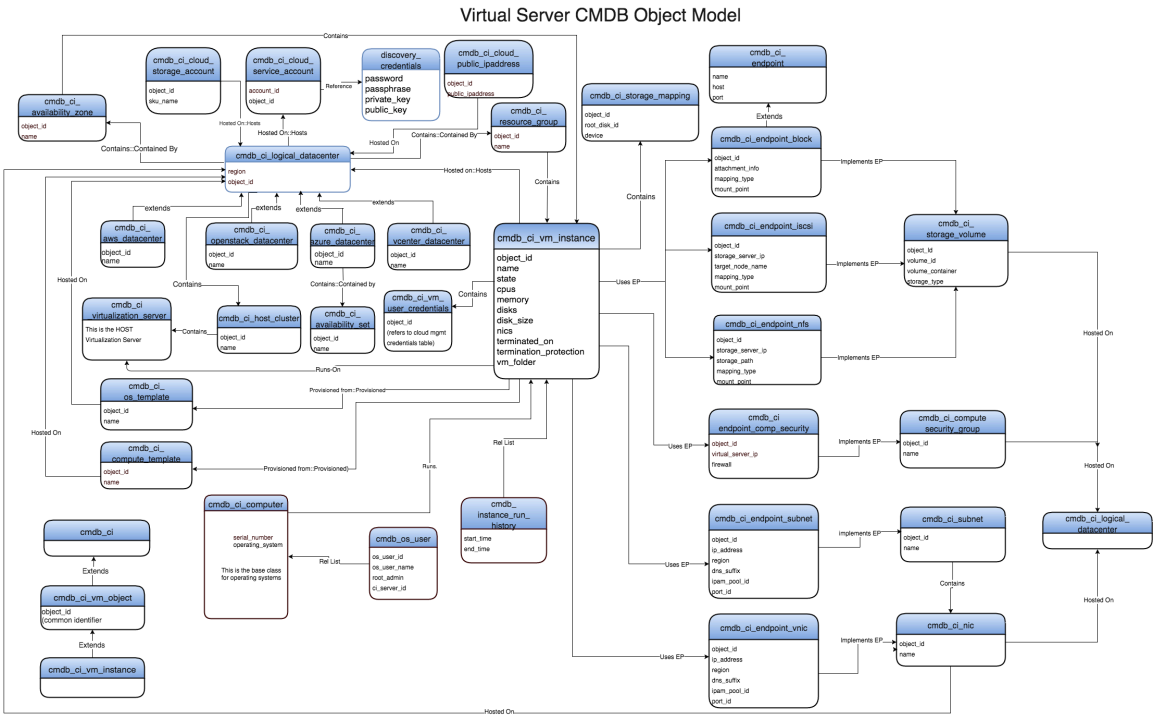

Key native attributes of the cmdb_ci_vm_instance table

Attribute	Description
object_id	Identifier that typically usually holds the VM instance ID. Uniquely identifies an object within a cloud.
name	Name of the VM.
state	State of the VML: **on**, **off**, or **terminated**.
cpus	Number of CPUs.
memory	Memory size in megabytes.
disks	Number of disk drives.
disk_size	Total size of disks in gigabytes.
nics	Number of network interface adapters.
terminated_on	Time that the instance was terminated.
termination_protection	Boolean. Default value is false. true: Can prevent the instance from being accidentally terminated using Amazon EC2. false: Cannot prevent the instance from being accidentally terminated using Amazon EC2.

Quick start tests for Cloud Provisioning and Governance

Validate that Cloud Provisioning and Governance still works after you make any configuration change such as applying an upgrade or developing an application. Copy and customize these quick start tests to pass when using your instance-specific data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Provisioning and Governance

Cloud Provisioning and Governance quick start tests require activating the following Cloud Provisioning and Governance plugins:

- Cloud Provisioning and Governance (com.snc.cloud.mgmt)
- Domain Support- Domain Extensions Installer plugin (com.glide.domain.msp_extensions.installer)
- Service Catalog- Domain Separation plugin (glideapp.servicecatalog.domain_separation)

Cloud Provisioning and Governance: Azure test suite

Test	Description	Release version
CMP: Add Azure Credentials	Add Azure credentials by inserting the fields into the Credentials table.	Orlando
CMP: Azure Service Account	Add Azure service account by inserting the fields into the service account table.	Orlando
CMP: Update datacenter type	Update the Cloud Service account with the datacenter type.	Orlando
CMP: Creating Resource Group	Creating Resource Group to provision Azure stack via ARM Template.	Orlando
CMP: Create Azure Logical datacenter.	Create Azure Logical datacenter for provisioning resources.	Orlando
CMP: Create "Contains: Contained by" relationship.	Create "Contains: Contained by" relationship between Resource Group and Azure datacenter.	Orlando
CMP: Create 'Hosted on' relationship.	Create hosted on relationship between logical datacenter and Cloud Service account.	Orlando
CMP: Create Cloud Account	Create cloud account by inserting the fields into the cloud account table.	Orlando
CMP: Cloud account and Logical datacenter association	Create an association between CMP cloud account and Logical datacenter.	Orlando
CMP: Create Catalog Item	Create Cloud Catalog Item with ARM Template.	Orlando
CMP: Create template version.	Create Cloud Template Version for Cloud Catalog Item.	Orlando
CMP: ARM template body	Update Cloud Template Version with ARM template body.	Orlando

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Provisioning and Governance: Azure test suite (continued)

Test	Description	Release version
CMP: Open the 'ServiceNow Cloud Template Versions'	Open the 'ServiceNow Cloud Template Versions' GUI page.	Orlando
CMP: Activate Cloud Template Version	Activate Cloud Template Version to create resource block, blueprint, and catalog.	Orlando
CMP: Verify the Cloud Template Version.	Validate whether status is Success and state is active in the Cloud Template Version after Activation operation.	Orlando
CMP: Activate Cloud Catalog Item	Activate Cloud Catalog Item to display catalog order form to the cloud user.	Orlando
CMP: Add stack and resource group name in Catalog item	Set the default value for stack name and resource group in the catalog item.	Orlando
CMP: Verify whether stack is present	Validate whether stack is available with name "ATFARMStack".	Orlando
CMP: Open Cloud Catalog Item page	Open Cloud Catalog Item page in the Service Portal.	Orlando
CMP: Submit Catalog form	Submit the Catalog Form to start provisioning ARM stack.	Orlando
CMP: Verify whether stack status is active	Validate the status of a stack after provisioning the stack.	Orlando
CMP: Verify stack status after Day2 - Stop operation	Validate the status of the stack status changed from 'On' to 'Off' after the Stop operation.	Orlando
CMP: Verify stack status after Day2 - Start operation	Validate the status of the stack status changed from 'Off' to 'On' after the Start operation.	Orlando
CMP: Verify stack status after Day2 - Deprovision operation	Validate the status of the stack status changed from 'On' to 'Terminate' after the de-provisioning operation.	Orlando

Cloud Provisioning and Governance: AWS test suite

Test	Description	Release version
CMP: Add AWS Credentials	Add AWS credentials by inserting the fields into the credential table.	Orlando
CMP: AWS Service Account	Add AWS service account by inserting the fields into the service account table.	Orlando

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Provisioning and Governance: AWS test suite (continued)

Test	Description	Release version
CMP: Update datacenter type	Update the cloud service account with the datacenter type.	Orlando
CMP: Create logical datacenter	Create Logical datacenter for provisioning the resources.	Orlando
CMP: Create hosted on relationship	Create hosted on relationship between logical datacenter and Cloud Service account.	Orlando
CMP: Create Cloud Account	Create cloud account by inserting the fields into the cloud account table.	Orlando
CMP: Cloud account and Logical datacenter association	Create an association between CMP cloud account and logical datacenter.	Orlando
CMP: Create Catalog Item	Create Cloud Catalog Item with the CFT Template.	Orlando
CMP: Create template version	Created Cloud Template Version for a Cloud Catalog Item.	Orlando
CMP: CFT template body	Update Cloud Template Version with the CFT template body.	Orlando
CMP: Open the ' ServiceNow Cloud Template Versions'	Open the ' ServiceNow Cloud Template Versions' GUI page.	Orlando
CMP: Activate Cloud Template Version	Activate Cloud Template Version to create resource block, blueprint, and catalog.	Orlando
CMP: Verify the Cloud Template Version	Validate the Cloud Template Version whether status is Success and state is active after Activation operation.	Orlando
CMP: Activate Cloud Catalog Item	Activate Cloud Catalog Item to display the catalog order form to the cloud user.	Orlando
CMP: Add stack name in Catalog item	Set the default value for stack name in the catalog item.	Orlando
CMP: Verify whether stack is present	Validate whether stack is available with name "ATFCFTStack".	Orlando
CMP: Open Cloud Catalog Item page	Open the Cloud Catalog Item page in the Service Portal.	Orlando
CMP: Submit Catalog form	Submit the Catalog Form to start provisioning CFT Stack.	Orlando

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Provisioning and Governance: AWS test suite (continued)

Test	Description	Release version
CMP: Verify whether stack status is active	Validate the status of the stack after provisioning the stack.	Orlando
CMP: Verify stack status after Day2 - Stop operation	Validate the status of the stack status changed from 'On' to 'Off' after the Stop operation.	Orlando
CMP: Verify stack status after Day2 - Start operation	Validate the status of the stack status changed from 'Off' to 'On' after the Start operation.	Orlando
CMP: Verify stack status after Day2 - De-provision operation	Validate the status of the stack status changed from 'On' to 'terminate' after the de-provisioning operation.	Orlando

To learn more about Cloud Provisioning and Governance, see Cloud Provisioning and Governance.

Related topics

Quick start tests

Cloud Configuration Governance

Use the ServiceNow® Cloud Configuration Governance application to check the configuration settings of cloud resources in your organization against a set of policies to identify violations. After identifying the violation, use remediation workflows to mitigate them.

Important: Starting with the Brazil release, Cloud Configuration Governance is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

Explore	Set up	Administer
Exploring Cloud Configuration Governance	• Install Cloud Configuration Governance • Install the CCG Content Pack • Set up Cloud Configuration Governance for AWS • Set up Cloud Configuration Governance for Microsoft Azure	• Policies • Scan configuration • Remediation
Use	**Develop**	**Troubleshoot and get help**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Run the scan configuration
- View dashboard
- Run remediation

- Developer training
- Developer documentation

- Ask or answer questions in the Cloud Configuration Governance forum on the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Cloud Configuration Governance

Understand the typical Cloud Configuration Governance workflow and high-level tasks.

To use Cloud Configuration Governance for managing the configuration of the cloud resources, perform the following tasks:

1. Install Cloud Configuration Governance. For more information, see Install Cloud Configuration Governance.
2. Install the CCG Content Pack to access the base system Cloud Configuration Governance contents. For more information, see Install the CCG Content Pack.
3. Assign Cloud Configuration Governance roles to users. For more information on the Cloud Configuration Governance roles, see Cloud Configuration Governance roles and system properties.
4. To manage the configuration of the Amazon Web Services (AWS) resources through Cloud Configuration Governance, see Set up Cloud Configuration Governance for AWS.
5. To manage the configuration of the Microsoft Azure resources through Cloud Configuration Governance, see Set up Cloud Configuration Governance for Microsoft Azure.
6. Create a scan configuration to scan the cloud resources against one or more policy sets and identify the policy violations. For more information, see Create a scan configuration.
7. Run the appropriate remediation action to fix the non-compliant cloud configuration identified during the scan run. For more information, see Run remediation.

What to do next

See Plugins or applications installed with ITOM Cloud Accelerate to find the list of plugins you activate with Cloud Configuration Governance.

Install Cloud Configuration Governance

You can install the Cloud Configuration Governance application (com.sn.itom.ccg) if you have the admin role. The application installs related ServiceNow® Store applications and plugins if they aren't already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Cloud Configuration Governance uses the ServiceNow® ITOM Cloud Accelerate subscription. When you install this application, it automatically installs several roles and the ServiceNow® Cloud Action Library application (sn.itom.cal).

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Cloud Configuration Governance application (com.sn.itom.ccg) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you can't find the application, you might have to request it from the ServiceNow Store.

   Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.
3. In the application installation dialog box, review the application dependencies.

   Dependent plugins and applications are listed if they'll be installed, are currently installed, or must be installed. If any plugins or applications must be installed, you must install them before you can install Cloud Configuration Governance.
4. Select **Install**.

What to do next

To access the base system contents such as policies, configuration collectors, remediations, and more, install the CCG Content Pack application. For more information, see Install CCG Content Pack.

Install the CCG Content Pack

You can install the CCG Content Pack application (sn_itom_ccg_cp) if you have the admin role.

Important: Starting with the Brazil release, Cloud Configuration Governance Content Pack is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the CCG Content Pack application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

The CCG Content Pack provides the base system contents such as policies, resource collectors, configuration collectors, and remediation actions for the Cloud Configuration Governance app.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the CCG Content Pack application (sn_itom_ccg_cp) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.
3. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
4. Select **Install**.

Set up Cloud Configuration Governance for AWS

Set up access to the Amazon Web Services (AWS) cloud accounts in Cloud Configuration Governance to enable interaction between the application and the cloud. The application requires access to the cloud accounts to scan the cloud resources for non-compliant configurations and remediate them.

Before you begin

Role required: sn_itom_ccg.scheduling_admin

About this task

For the purposes of configuring access for AWS accounts, the following terms are used:

Trusting accounts

The trusting accounts don't have permanent AWS credentials. You configure the trust relationship for IAM roles in these accounts to rely on other accounts for access.

Trusted accounts

The trusted accounts are used by the trusting accounts for access. The ServiceNow UI refers to the trusted accounts as accessor accounts.

Use any one of the following methods to configure access to the AWS accounts:

- Configure the permanent credentials in the ServiceNow AI Platform to connect with the standalone AWS accounts (discreet accounts). The Cloud Service Account [cmdb_ci_cloud_service-account] table contains information on the service account type, such as management or member account and access credentials.
- Configure the member accounts to rely on the management account for access. In this case, configure the permanent credentials of the management account in the ServiceNow AI Platform.
- Configure the accounts to rely on a trusted account for access (lateral access within the same AWS organization or across different AWS organizations). In this case, configure the permanent credentials of the trusted account in the ServiceNow AI Platform.

Note: Cloud Configuration Governance doesn't use a MID Server-based assume role setup to access the trusting accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create credentials for the AWS service accounts.
   a. Navigate to **Connection & Credentials** > **Credentials**.
   b. Select **New**, and then select **AWS Credentials**.
   c. On the form, fill in the fields.

AWS Credentials form

Field	Description
Name	Unique and descriptive name for the AWS credentials.
Active	Option to use the credential.
Access Key ID	The access key ID that you generated on the AWS management console.
Secret Access Key	The secret access key that you generated on the AWS management console.

   d. Select **Save**.
2. Select the **sn_itom_cal.Aws_Creds_Alias** credential alias or create a credential alias.
   a. Unlock the credential alias.
   b. Search for a credential alias.
   c. Select **New**.
   d. On the form, fill in the fields.

Connection & Credential Alias form

Field	Description
Name	Unique name of the alias.
Type	Credential alias type. Select **Credential** from the Type drop-down list.

   e. Select **Submit**.
3. Set the **Authentication Algorithm** field to **AWS Authenticator**.
4. Select **Submit**.
5. Set up an AWS service account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **Cloud Provisioning and Governance** > **Service accounts**.

b. Select **New**.

c. On the form, fill in the fields.

Cloud Service Account form

Field	Description
Name	Unique name of the service account.
Account ID	12-digit user account number. Expand the list under the account name on the AWS Management Console to view the number. **Important:** In the **Account ID** field, remove the hyphen characters (-) from the number.
Discovery credentials	The credentials needed for ServiceNow applications to access the cloud account. You can configure the discovery credentials at a later stage, while configuring access to the AWS accounts. ▪ If you are setting up an independent service account or a management account, select its AWS credential. ▪ To use other AWS accounts to access this account, leave the field empty. For example, you don't have to specify the AWS credentials for accounts assuming Identity and Access Management (IAM) roles or member accounts using their management account for access.
Datacenter URL	URL of the datacenter. Leave this field empty. **Note:** For GovCloud, use the URL https://ec2.us-gov-west-1.amazonaws.com.
Datacenter type	Type of the datacenter where the account is hosted. Select **AWS Datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.
Parent account	Name of the management account that represents the AWS organization to which this member account belongs. This field appears when you select AWS Datacenter. If the account doesn't belong to any AWS organization, leave this field empty.
Is master account	Management account flag. This check box appears when you select AWS Datacenter from the Datacenter Type drop-down. Select the check box to associate the AWS service account with the management account. Select this check box only for accounts that you have previously configured as management accounts

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	with some member accounts belonging to them. For more information on AWS Organization, see AWS documentation.
Accessor account	Name of the trusted account. Configure this field only for accounts that don't use permanent AWS credentials and rely on IAM roles for access.

d. Select **Submit**.

6. Do any one of the following actions.

Option	Description
Create an assume role configuration for the management account	If you want to use a management account to scan the member accounts of the AWS organization, create an assume role configuration for the management account. **a.** If you don't want to use the OrganizationAccountAccessRole to access the member account, configure the trusting account for Cloud Configuration Governance. For more information, see Configure the trusting account for Cloud Configuration Governance and Cloud Action Library. **b.** Repeat step 6.a for all the member accounts that must be scanned through the management account without using the OrganizationAccountAccessRole. **c.** If you want to use the OrganizationAccountAccessRole to access the member account, create an assume role configuration for the management account. For more information, see Create an assume role configuration.
Configure the trusting account for Cloud Configuration Governance	If you want to use a trusted account to scan the trusting account, configure the trusting account for Cloud Configuration Governance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	a. Configure the trusting account for Cloud Configuration Governance. For more information, see Configure the trusting account for Cloud Configuration Governance and Cloud Action Library. **b.** Repeat step 6.a for all the trusting accounts that must be scanned through the trusted account.

7. Install and configure the MID Servers.
   For more information, see Install and configure MID Servers.
8. Run datacenter discovery to identify the datacenters associated with the service accounts.
   For more information, see Run datacenter discovery.

Configure the trusting account for Cloud Configuration Governance and Cloud Action Library

Configure the trusting account whose resources need to be accessed, to rely on the trusted account using the Identity and Access Management (IAM) role.

Before you begin

- Familiarize yourself with the Amazon documentation on Creating a role to delegate permissions to an IAM user.
- Decide which Amazon Web Services (AWS) account is going to be the trusted account. You use the trusted account to configure temporary credentials for Cloud Configuration Governance using IAM roles. The trusted account that you use to access other accounts using IAM roles is referred to as an accessor account.
- Set up the trusted and the trusting account as described in Set up Cloud Configuration Governance for AWS.

Role required: sn_itom_ccg.scheduling_admin

Procedure

1. Create an IAM role for the trusting account and configure the trust relationship between the user assuming this role and the trusted (accessor) account.

   a. Log in to the trusting account on the AWS Management Console.

   b. Create and configure the IAM role specifying the trusted (accessor) account ID in the **Account ID** field.
   For information on creating AWS roles, see the Amazon documentation.

   c. On the Summary page for the IAM role, select the **Trust Relationships** tab.

   d. Select **Edit trust relationship**.
   The Edit Trust Relationship page opens showing the policy document.

   e. Set the *AWS* parameter to the full user ARN of the trusted (accessor) account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

f. Verify that the *`Action`* value is set to *`sts:AssumeRole`*.

g. Select **Update Trust Policy**.

2. Configure the trusted service account for the trusting account in the ServiceNow AI Platform.
   a. Navigate to **All** > **Cloud Provisioning and Governance** > **Service Accounts**
   b. Open the trusting account.
   c. On the Cloud Service Account form, enter the name of the trusted account in the **Accessor account** field.
   d. Select **Update**.
3. Create an assume role configuration for the trusting account.
   For more information, see Create an assume role configuration.

Create an assume role configuration

Create a service account assume role configuration to facilitate cross-account access that is from a management account to a member account or from a trusted account to a trusting account. All the related member accounts or trusted accounts are automatically loaded avoiding the need to add them individually.

Before you begin

Make sure that the accessor account credentials are configured in the ServiceNow AI Platform.

Role required: sn_itom_ccg.scheduling_admin

About this task

If permanent credentials aren't defined for the member account, Cloud Configuration Governance checks the Cloud Service Account Assume Role Config [sn_itom_ccg_service_account_assume_role_config] table to identify the special configuration associated with the management account, and then uses the management account credentials to generate temporary credentials for the member accounts. If the configuration exists in the table, Cloud Configuration Governance uses the temporary credentials acquired by specifying

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a role and its configuration in the Amazon Web Services (AWS) Security Token Services API AssumeRole action.

Procedure

1. Navigate to **All** > **Cloud Configuration Governance** > **Assume Role Config**.
2. Select **New**.
3. Fill the values on the Service Account Assume Role Config form.
   For a description of the form fields, see Cloud Configuration Governance Service Account Assume Role Config form.
4. Select **Submit**.

Set up Cloud Configuration Governance for Microsoft Azure

Set up access to the Microsoft Azure cloud accounts in Cloud Configuration Governance to enable interaction between the application and the cloud. The application requires access to the cloud accounts of your organization to scan the cloud resources for non-compliant configurations and remediate them.

Before you begin

Make sure the follow tasks are complete:

1. You have migrated from CAPI to patterns. For more details, see Migrate from CAPI to Patterns.
2. The Microsoft Azure credentials to access the Microsoft Azure account have been created. For more information, see Create Azure cloud credentials.
3. Your Azure service accounts have been configured. For more information, see Set up Azure service accounts.

Role required: sn_itom_ccg.scheduling_admin

Procedure

1. Create an Integration Hub connection for Microsoft Azure.

   a. Navigate to **Process Automation** > **Flow Designer**, and select the **Integrations** tab.

   b. On the **Azure Connection** tile, select **Add Connection**.

   c. On the form, fill in the fields.

Scan Schedule form

Field	Description
Connection name	Name of the connection.
OAuth client ID	Microsoft Azure application ID.
OAuth client secret	Microsoft Azure application key.
OAuth token URL	Tenant ID of the Microsoft Azure account. Replace `{tenant}` with the Microsoft Azure tenant ID.

   d. Select **Create and Get OAuth Token**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Run datacenter discovery to identify the datacenters attached with the service accounts. For more information, see Run datacenter discovery.

Domain separation and Cloud Configuration Governance

Domain separation is not supported in Cloud Configuration Governance . Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

Related topics

Domain separation for service providers

Using Cloud Configuration Governance

Cloud Configuration Governance provides a set of ready-to-use policies. You can either use these policies to identify the non-compliant cloud resource configurations within your cloud organization, or create custom policies. After identifying the non-compliant cloud resource configurations, you can use the remediation feature of the application to fix the violations.

Note: Starting with Cloud Configuration Governance version 1.3.7, the base system contents are moved to the CCG Content Pack. Install the CCG Content Pack to access the base system Cloud Configuration Governance contents.

Scan configuration

Scan configuration defines the various settings for the Cloud Configuration Governance scan.

When you execute a scan configuration, Cloud Configuration Governance identifies the targeted cloud resource types from the policy sets attached with the scan configuration. Then, it invokes the respective resource collectors and configuration collectors to import the configuration data of the targeted resource types. After completing the configuration data import, Cloud Configuration Governance compares the imported configuration data against the non-compliant configurations defined in the policies. If Cloud Configuration Governance detects a match, it reports a violation in the audit violation report.

You can reduce the amount of configuration data imported into the system by specifying a set of service account-datacenter combinations for the scan configuration. If specified, Cloud Configuration Governance imports the data from the specified service account-datacenter combinations only. In addition, you can also use tag-based cloud resource filters to identify the cloud resources whose configuration data must be imported for the scan.

For more information on creating scan configuration, see Create a scan configuration.

Create a scan configuration

Create a scan configuration in Cloud Configuration Governance to scan the cloud resources against one or more policy sets and identify the policy violations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that a suitable policy set is available for the scan. The policy set enforces a given security or organizational standard. For more information on creating the policy set, see Create a policy set.
- Role required: sn_itom_ccg.scheduling_admin

About this task

Note: Starting with Cloud Configuration Governance version 1.3.7, the base system contents are moved to the CCG Content Pack. Install the CCG Content Pack to access the base system Cloud Configuration Governance contents.

Procedure

1. Navigate to **Cloud Configuration Governance** > **Scan Configurations**.
2. Select **New**.
3. Fill the values on the Scan Configuration form.
   For a description of the form fields, see Cloud Configuration Governance Scan configuration form.
4. Select **Submit**.
5. To create a schedule to execute the scan configuration, select **Setup Schedule**.

   a. **Optional:** Set a criteria to cancel the long running scan runs.
   To import the configuration data of the cloud resources, Cloud Configuration Governance resource collectors and configuration collectors use Workflow Studio subflows. Sometimes, the subflow errors out and it doesn't report the error. In such cases, the scan run can run indefinitely. Setting a cancellation criteria helps to cancel such scans.

What to do next

Run the scan configuration to evaluate the resource configuration of the given cloud against the specified policy sets and identify the policy violations. For more information, see Run the scan configuration.

Run a scan configuration

Run a scan configuration to evaluate the resource configuration of the given cloud against the specified policy sets and identify the policy violations. Cloud Configuration Governance reports the identified policy violations as audit results.

Before you begin

Role required: sn_itom_ccg.scheduling_admin

Procedure

1. Navigate to **Cloud Configuration Governance** > **Scan Configurations**.
2. Open the scan configuration that you want to run.
3. Select **Execute**.

Scan run

Each execution of a scan configuration is called a scan run. The Scan Runs module displays a list of all the scan runs.

Cloud Configuration Governance displays the following information for each scan run in a related list:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Scan Status**: This tab displays the data collection status at the service account level. Open the scan run record to view the configuration data collection status for each datacenter covered in the scan. Use the flow execution context to troubleshoot the execution of the policy.
- **Policy Run Status**: This tab displays information on the policies executed during the scan run and the status of their execution. Use the flow execution context to troubleshoot the execution of the policy.
- **Run Resource Assocs**: This tab displays information about the resources and their configuration data imported during the scan run. You can select any resource to view the following information:
  - Configuration settings of the resource
  - Resource attributes
  - Audit results for the resource

Scan Run Logs tab: It displays the trace of the scan run.

Related topics

Scan configuration

Cloud Configuration Governance policies

Cloud Configuration Governance policy defines the non-compliant configurations for a given cloud resource type.

Each Cloud Configuration Governance policy contains the following information:

- The cloud on which the resource is provisioned.
- The cloud resource type.
- Definition of the non-compliant configuration. For example, unencrypted Amazon Web Services (AWS) S3 buckets or insecure Identity and Access Management (IAM) accounts.
- Definition of the audit violation (policy violation) report.

Note: Starting with Cloud Configuration Governance version 1.3.7, the base system contents are moved to the CCG Content Pack. Install the CCG Content Pack to access the base system Cloud Configuration Governance contents.

Cloud Configuration Governance provides several base system policies. You can either use these policies or create custom policies as per the needs of your organization. Depending on the need and your familiarity with the ServiceNow AI Platform, you can use any one of the following methods to create the policy:

- Create a policy with the condition builder
- Create a policy with Integration Hub flow
- Create a policy through script

To use the policy, add the policy to a policy set. Each policy set can contain one or more policies. For more information on creating policy sets, see Create policy set.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Base system policies

Name	Type	Description
AWS IAM User Activity policy	Condition builder	Policy to check if the password is enabled for the AWS IAM user. To use this policy, ensure that the AWS IAM user account has the following permissions: • Iam:GetCredentialReport • Iam:GenerateCredentialReport
AWS S3 Enforce Bucket encryption	Condition builder	Policy to check if the AWS S3 buckets are encrypted.
AWS Sample flow policy	Integration Hub flow	Policy to illustrate an Integration Hub flow-based policy.
AWS VM HardwareType	Condition builder	Policy to check if the deployed EC2 VMs are using only the approved hardware types.
AWS VM IPAddress	Script	Policy to check if the IP address of the EC2 VM is matching with the Configuration Management Database (CMDB) record.
AWS VM Monitoring State	Condition builder	Policy to check if detailed monitoring is enabled for the EC2 VM.
Azure VM HardwareType	Condition builder	Policy to check if the deployed Azure VMs are using only the approved hardware types.
Azure VM IP Address	Script	Policy to check if the IP address of the Azure VM is matching with the CMDB record.
Azure VM Monitoring State	Condition builder	Policy to check if detailed monitoring is enabled for the Azure VM.

- For more information on creating policy sets, see Create policy set.
- For more information on base system policies and policy sets, see AWS policies, Azure policies, and Policy sets.

Create a policy with the condition builder

Use the condition builder to create a Cloud Configuration Governance policy without writing any code at all. Use it to specify the criteria for identifying and reporting the non-compliant cloud resource configurations.

Before you begin

Role required: sn_itom_ccg.governor

Procedure

1. Navigate to **Cloud Configuration Governance** > **Policies**.
2. Select **New**.
3. On the Policy form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For form field description, see Cloud Configuration Governance Policy form.

Refer to one of the base system condition builder-based policies as an example for creating the custom policy.

4. Select **Submit**.

What to do next
Test the policy to make sure that it is working properly. For more information, see Test the policy.

Create a policy with the Integration Hub subflow

Use the Integration Hub subflow to create a Cloud Configuration Governance policy. An Integration Hub subflow uses graphical coding mechanisms to specify the policy conditions. It offers greater flexibility in audit violation reporting than the condition builder. That is, you can use the same policy to report different types of issues.

Before you begin
Role required: sn_itom_ccg.governor

About this task
Refer to the base system **Sample flow policy** as an example for creating new policies with the Integration Hub subflow.

Procedure
1. Create an Integration Hub subflow for the policy.

   a. Copy the **CCG – Flow Policy Template** subflow.
   For more information on copying the subflow, see Copy an Integration Hub subflow template.

   b. To read the configuration data, add the **CCG – Read Config Setting** action.
   For more information on the **CCG – Read Config Setting** action, see Cloud Configuration Governance actions reference.

   c. Repeat step 1.b for all the configuration keys that you want to use in the policy.

   d. Add the flow logic.

   e. Assign subflow outputs.
   The subflow output settings define how Cloud Configuration Governance reports the audit violations. For more information on assigning the subflow output, see Cloud Configuration Governance actions reference.

   f. **Optional:** If you have not selected the **Report Issue** option in the **Assign Subflow outputs** form, add an action under the **Assign Subflow outputs** action, and then add the **Create Record** action.
   For more information on the **Create Record** action, see Cloud Configuration Governance actions reference.

   g. Save the subflow.

2. Navigate to **Cloud Configuration Governance** > **Policies**.
3. Select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. On the Policy form, fill in the fields.
   For a description of the form fields, see Cloud Configuration Governance Policy form.
5. Select **Submit**.

What to do next
Test the policy to make sure that it is working properly. For more information, see Test the policy.

Create a policy through script

Use the ServiceNow platform native scripting mechanism to define the policy condition. You can externalize the decision making to script includes and reuse the code across different policies.

Before you begin
Role required: sn_itom_ccg.governor

Procedure
1. Navigate to **Cloud Configuration Governance** > **Policies**.
2. Select **New**.
3. On the Policy form, fill in the fields.

   For the description of the form fields, see Cloud Configuration Governance Policy form.

   Refer to one of the base system script-based policies as an example for creating the custom policy.
4. Select **Submit**.

What to do next
Test the policy to make sure that it is working properly. For more information, see Test the policy.

Test the policy

Test the Cloud Configuration Governance policy against the configuration data imported for a previous scan run. Test the policy to make sure that it is working properly.

Before you begin
- Ensure that a scan run is available that contains the configuration data you want to check through the policy.
- Role required: sn_itom_ccg.governor

Procedure
1. Navigate to **Cloud Configuration Governance** > **Policy**.
2. Open the policy.
3. Select **Test Policy**.
4. Select the appropriate scan run.
5. Select **OK**.

Create a configuration key

Create a configuration key in Cloud Configuration Governance to use it in the policies. The configuration key specifies the resource collector and configuration collectors used to import the configuration data from the cloud. It also specifies the data type of the configuration value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that a resource collector is available for the resource type mentioned in the configuration key. If the resource collector is not available, you can create it. For more information, see Create a resource collector.
- Role required: sn_itom_ccg.governor

Procedure

1. Navigate to **Cloud Configuration Governance** > **Configuration Keys**.
2. Select **New**.
3. On the Configuration Key form, fill in the fields.
   For the description of the form fields, see Cloud Configuration Governance Configuration key form.
4. Select **Submit**.

Create a resource collector

Create a resource collector to set up how Cloud Configuration Governance imports all the resources of a given type.

Before you begin

Role required: sn_itom_ccg.governor

About this task

The cloud resources are available under a service account or service account-datacenter combination. The resource collector uses Integration Hub subflows to import the resource list from the cloud. Create only one resource collector per resource type.

Procedure

1. Create an Integration Hub subflow for the configuration collector.

   a. Copy the **CCG – Collector Template** subflow.
   For more information on copying subflow templates, see Copy an Integration Hub subflow template.

   b. Add the appropriate Cloud Action Library (CAL) action and flow logic to read the resource data.

   For more information on the available CAL actions, see Cloud Action Library actions reference.

   If CAL doesn't have the required action, you can refer to an existing CAL action and use the Integration Hub datastream action to implement the custom action. Alternatively, you can also use the REST step. For more information, see Data Stream actions and pagination and REST step.

   c. Add the **CCG – Insert Resource Record** action, and then select the appropriate inputs to insert the configuration data to the Configuration Management Database (CMDB).
   For more information on the **CCG – Insert Resource Record** action, see Cloud Configuration Governance actions reference.

   d. Select **Publish**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Navigate to **Cloud Configuration Governance** > **Resource Collectors**.
3. Select **New**.
4. On the form, fill in the fields.
   For the description of the field values, see Cloud Configuration Governance Resource collector form.
5. Select **Submit**.

What to do next

- If the required configuration data is not available at the cloud resource level, see Create a configuration key.
- If you have created a new resource type in the Cloud Configuration Governance application, you can optionally create a CI finder mapping for the newly created resource type. For more information, see Create a CI finder mapping.

Create a configuration collector

Create a configuration collector to set up how Cloud Configuration Governance imports the configuration data from the cloud. The configuration collector uses Integration Hub subflows to import the configuration data from the cloud.

Before you begin

- Ensure that Cloud Configuration Governance contains suitable configuration keys for all the configuration values that you want to import from the cloud. Create a configuration key if it is not available in the application. For more information, see Create a configuration key.

Role required: sn_itom_ccg.governor

About this task

For some cloud resources, the required configuration data is available at the cloud resource level. Therefore, create a configuration collector only if the required configuration data is not available at the cloud resource level.

For example, most of the configuration data of the **Virtual Machine Instance** resource type are available at the cloud resource level. Therefore, a configuration collector is not implemented to import these configurations.

Refer to a base system configuration collector as an example for creating the configuration collectors.

Procedure

1. Create an Integration Hub subflow for the configuration collector.
   a. Copy the **CCG – Collector Template** subflow.
      For more information on copying templates, see Copy an Integration Hub subflow template.
   b. Add the **CCG - Get Resources** action and fill in the fields.
   c. Add the appropriate Cloud Action Library (CAL) action and flow logic to read the configuration data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

For more information on the available CAL actions, see Cloud Action Library actions reference.

If CAL doesn't have the required action, you can refer to an existing CAL action and use the Integration Hub REST step to implement the custom action. For more information, see REST step.

d. Add the **CCG – Insert Configuration Settings** action to insert the configuration data to the Configuration Management Database (CMDB).

e. Select **Publish**.

2. Navigate to **Cloud Configuration Governance** > **Configuration Collector**.
3. Select **New**.
4. On the form, fill in the fields.

Configuration Collector form

Field	Description
Provider type	Cloud that hosts the resources to be scanned through the policies.
Flow	Configuration collector that can import the configuration data from the cloud.
Resource type	Cloud resource type whose configuration you want to import
Name	Name of the configuration collector.
Description	Brief description of the configuration collector.

5. Select **Submit**.

Create a CI finder mapping

Use the CI finder mapping to link the resources fetched by the ServiceNow Cloud Configuration Governance scan runs with the appropriate Configuration Management Database (CMDB) Configuration Items (CIs), if available. After you create the mapping, you can associate the reported configuration violations with the CIs to which they belong.

Before you begin

Role required: sn_itom_ccg.scheduling_admin or sn_itom_ccg.governor

About this task

After creating a resource type in Cloud Configuration Governance, you can optionally create a CI finder mapping for the new resource type. Cloud Configuration governance calls the CI finder mapping script for all the resources fetched by the scan runs.

Procedure

1. Navigate to **All** > **Cloud Configuration Governance** > **CI Finder Mappings**.
2. Select **New**.
3. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI Finder Mapping form

Field	Description
Provider Type	Cloud service provider that hosts the resource.
Resource Type	Resource type of the cloud resource.
Script	Script to link the resources fetched by the Cloud Configuration Governance application with the existing CI. Use an existing CI finder mapping script as an example to create a script for the new resource type. Cloud Configuration Governance offers several objects and variables for use in the script. For more information, see Cloud Configuration Governance scripting reference.

Create a policy set

Create a policy set in Cloud Configuration Governance to group policies that enforce a given security or organization standard. When the scan configuration runs the policy set, it identifies the cloud resources that don't adhere to the set standards. A policy set can contain one or more policies.

Before you begin

Role required: sn_itom_ccg.governor

Procedure

1. Navigate to **Cloud Configuration Governance** > **Policy Sets**.
2. Select **New**.
3. On the form, fill in the fields.

Policy Sets form

Field	Description
Name	Name of the policy set.
Provider type	Cloud that hosts the resources to be scanned through the policies attached to the policy set.
Description	Brief description of the policy set.
Policies	Policies that you want to group to enforce a given security or organizational standard. The policy set can contain more than one policy.

4. Select **Submit**.

Copy an Integration Hub subflow template

Copy an Integration Hub subflow template to create subflows for the Cloud Configuration Governance application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_itom_ccg.governor

About this task

Procedure
1. Navigate to **Process Automation** > **Flow Designer**.
2. Select **Subflow**.
3. In the Name column, search for the required subflow template, and then open it.
4. Select the more actions icon ( ) icon, and then select **Copy subflow** from the menu.
5. In the **New subflow name** field, enter a unique and descriptive name for the copied subflow.
6. Select **Copy**.
   The new subflow is created and is opened as a tab in the subflow designer.

Remediation

Use the Cloud Configuration Governance remediation actions to fix the non-compliant cloud resource configurations identified during the scan runs.

The Remediations module displays a list of all the available remediation actions for the set of violation types (violation definitions). Cloud Configuration Governance offers several base system remediation actions. If the application doesn't contain the required remediation action, you can create custom actions. For more information, see Create a remediation catalog item and Create a remediation action.

Because Cloud Configuration Governance performs the remediation actions on the cloud resources, you can optionally enforce additional approval requirements for the action. Whenever sn_itom_ccg.ccg_operator submits a remediation request that requires approval, the application creates a change request to track the approval. You can review the list of all the change requests in the Remediation Requests module.

Note: Starting with Cloud Configuration Governance version 1.3.7, the base system contents are moved to the CCG Content Pack. Install the CCG Content Pack to access the base system Cloud Configuration Governance contents.

Create a remediation catalog item

Create a custom remediation catalog item in Cloud Configuration Governance for the remediations that require additional input parameters to execute.

Before you begin
Role required: sn_itom_ccg.governor

About this task
Some remediation actions require additional operational information, such as resource hardware type, to execute. For such cases, create a remediation catalog item with custom input parameters or variables, and associate it with the appropriate remediation action and violation type. After you create the remediation catalog item, a user with the sn_itom_ccg.ccg_operator role can use it to raise a remediation request.

Procedure
1. Navigate to **All** > **Cloud Configuration Governance** > **Remediation Catalog**.
2. Open the CCG Remediation Default Catalog record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Copy** to copy the catalog item.
4. On the form, fill in the fields.

Record Producer form

Field	Description
Name	Name of the Cloud Configuration Governance catalog item.
Table name	The table in which the record producer creates records. Retain the default value.
Owner	Owner of the catalog item who has edit access to it.
Short description	Business purpose of the remediation catalog item. Cloud Configuration Governance displays the short description to the operator.
Activate	Status of the catalog item. The copied catalog item is active by default.

5. Select **Update**.
   The record producer creates the following items:
   - A Workflow Studio action named <Catalog item name>_Action.
   - A Workflow Studio subflow named <Catalog item name>_Subflow.
   - A remediation record for the catalog item. The remediation record is associated with the catalog item and the <Catalog item name>_Subflow.
6. **Optional:** Define catalog item variables to collect additional information through the catalog item form.
   a. From the Variables related list, select **New**.
   b. On the form, fill in the fields.
   For more information on catalog variables, see Service catalog variables.
   c. Save the form.
7. Associate the remediation catalog item with the violation type.
   a. On the Remediations related list, select **New**.
   b. Double-click the violation cell.
   c. Select the violation type.
   d. Select the save icon ( ).
8. Update the <Catalog item name>_Action to interact with the cloud resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance uses Integration Hub action to execute the remediation on the cloud resource. As business needs dictate the remediation logic, the Record Producer only creates a boilerplate for the <Catalog item name>_Action. You must implement the <Catalog item name>_Action per the business needs.

a. Navigate to **All** > **Flow Designer**.

b. Select the **Actions** tab.

c. Open the <Catalog item name>_Action.

d. Refer to the cloud provider documentation and implement the <Catalog item name>_Action per the business needs of your organization.
For more information on building Integration Hub actions, see Building custom actions.

e. Save the updated <Catalog item name>_Action.

f. Test the <Catalog item name>_Action.
For more information, see Test an action.

g. Publish the <Catalog item name>_Action.

9. Associate the <Catalog item name>_Action with the <Catalog item name>_Subflow.
Cloud Configuration Governance uses the Integration Hub subflows to trigger the Integration Hub action and process the returned data.

a. Navigate to the Workflow Studio home page by selecting the flow designer home icon ( ).

b. Select the **Subflows** tab.

c. Open the <Catalog item name>_Subflow.

d. Select **Add an Action, Flow Logic, or Subflow**.

e. Select **Action**.

f. Search for `<Catalog item name>_Action` and select the action.

g. Add appropriate flow logic.

h. Assign the output of the subflow to the output variable.

i. Test the <Catalog item name>_Subflow.

j. Publish the <Catalog item name>_Subflow to activate it.

What to do next

Run the remediation action to fix the identified violation. For more information, see Run remediation.

Create a remediation action

Create a remediation action in Cloud Configuration Governance for the remediations that doesn't require any additional input parameters.

Before you begin

Role required: sn_itom_ccg.ccg_governer

About this task

Cloud Configuration Governance uses Integration Hub subflows to perform the remediations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Cloud Configuration Governance** > **Remediations**.
2. Select **New**.
3. On the form, fill in the fields.

Remediations form

Field	Description
Minimum permission required	Minimum cloud permissions required to execute the remediation. Cloud Configuration Governance doesn't read the permissions information from this field. Use this field to store the permissions information for future reference.
Action	Remediation flow that can fix the selected violation.
Violation	Violation that you want to remediate.
Requires approval	Option to enforce approval for the remediation action. For more information on approvals, see Approvals.
Description	Brief description of the remediation action.
Remediation Catalog Item	Name of the remediation catalog item that can be used to raise a remediation request. Select the CCG Remediation Default Catalog.

4. Select **Submit**.

What to do next

Run the remediation action to fix the identified violation. For more information, see Run remediation.

Run remediation

Run the appropriate remediation action to fix the non-compliant cloud configuration identified during the Cloud Configuration Governance scan run.

Before you begin

- The scan run is complete
- Role required: sn_itom_ccg.ccg_operator

Procedure

1. Navigate to **Cloud Configuration Governance** > **Scan Run**.
2. Open the scan run during which **Cloud Configuration Governance** reported the audit issues.
3. Select **Show audit results**.
4. Select the violation type that you want to fix.
   Cloud Configuration Governance groups the violations per the violation type. During a remediation run, you can fix one or more violations of the same type only.
5. Select one or more violations.
6. Select **Remediate**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can select **Ignore for Now** to ignore the issue. When you ignore an issue, Cloud Configuration Governance doesn't report the issue as a violation and it is not represented in the dashboard reports.

7. Review the list of resources selected for remediation.
8. Select the remediation action.
9. Select **OK**.

Result

What to do next

From the Remediation Runs related list, select the change request and review the status of the remediation. You can also view the list of all the remediation requests in the Remediation requests module.

Cloud Configuration Governance reference

Reference topics provide additional information about the Cloud Configuration Governance application.

Cloud Configuration Governance dashboard

Cloud Configuration Governance is a tool used to manage the configuration of the cloud resources as per your organizational standards and established security standards. Use the dashboard to review the health score of the cloud, policy violation statistics, policy violation trend, remediations overview, and more.

Cloud Configuration Governance dashboard

Required ServiceNow AI Platform roles

sn_itom_ccg.ccg_operator or sn_itom_ccg.report_viewer roles are required to view the dashboard and its widgets.

Access the Cloud Configuration Governance dashboard

To open the dashboard, navigate to **Cloud Configuration Governance** > **Dashboard**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
sn_itom_ccg.ccg_operator or sn_itom_ccg.report_viewer	• Review the cloud heath score. • Time taken to remediate the identified policy violations. • Understand the trend of the policy violation.

Reports

Title	Type	Source table	Description
Cloud health score	Speedo chart	sn_itom_ccg_scan_summary	Percentage of scanned cloud resources that doesn't have any audit issue.
Violating resources	Stacked bar chart	sn_itom_ccg_scan_summary	Severity of the detected audit issues per cloud resource type.
Violations count	Donut chart	sn_itom_ccg_violation_stats	Violation per severity of the audit issue.
Scanned resources	Donut chart	sn_itom_ccg_scan_summary	Number of cloud resources scanned per cloud resource type. Use the **CCG Drilldown Filters** to filter the data and visualize the data of interest.
Issues	Horizontal bar chart	sn_itom_ccg_audit_stats	Number of violations reported per violation definition. The chart can display a maximum of 10 bars, that is one per violation definition type. If the data contains more than 10 violation definition types, all the additional violation definition types are

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
			merged under the **Others** bar. Use the **CCG Drilldown Filters** to filter the data and visualize the data of interest.
Policy compliance	Embedded lists	sn_itom_ccg_policy_compliance_stats	Displays the following information for the given policy: • Count of compliant resources • Count of non-compliant resources
Policy set compliance	Embedded lists	sn_itom_ccg_policy_set_compliance_stats	Displays the following information for the given policy set: • Count of compliant resources • Count of non-compliant resources
Total issues remediated	Half donut chart	sn_itom_ccg_remediation_daily_trend	Remediations performed per violation definition type.
Remediations by type	Pie chart	sn_itom_ccg_remediation_daily_trend	Remediations performed per cloud resource type.
Remediated on	Half donut chart	sn_itom_ccg_remediation_daily_trend	Time taken to remediate the violations.
Trend	Line chart	sn_itom_ccg_daily_trend	Trend of reporting violations over time.

Filters

Name	Type	Description
CCG Drilldown filter	Choice-based filter	Use this filter to select the cloud provider, cloud account, and regions for visualizing their scanned resources and issues data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance roles and system properties

Cloud Configuration Governance users require the appropriate roles and system properties to perform various activities. These roles and system properties are installed with the activation of the application.

Only users with the user_admin or admin role can create users with these roles:

Cloud Configuration Governance roles

Descriptive name and role name	Description and tasks	Nested roles
sn_itom_ccg.scheduling_admin	A scheduling admin can perform the following actions: • Create scan schedules. • Create CI finder mappings.	discovery_admin
sn_itom_ccg.governor	A governor can perform the following actions: • Create policies by using condition builder, Integration Hub flows, or script. • Create a policy set. • Create resource collectors. • Create configuration keys. • Create remediation. • Create CI finder mappings.	• flow_designer • action_designer • catalog_admin
sn_itom_ccg.ccg_operator	An operator can perform the following actions: • View the audit issue reports. • Run remediation. • Run scan configurations. • View the dashboard.	• flow_operator • sn_change_write
sn_itom_ccg.report_viewer	A report viewer can perform the following actions: • View the audit issue reports. • View the dashboard.	

Cloud Configuration Governance system properties

Property name	Description
snc_af.log_level	When you enable this property, Cloud Configuration Governance writes the logs to the syslog. The supported logging levels are as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance system properties (continued)

Property name	Description
	• DEBUG • INFO • WARN • ERROR • OFF Ensure that the application scope is **Global**, when you set or clear this property.

Related topics

Assign a role to a user

Cloud Configuration Governance Scan configuration form

The Cloud Configuration Governance Scan configuration form displays detailed information about the scan configuration.

Field	Description
Name	Name that uniquely identifies the scan configuration.
Description	Brief description of the scan configuration.

Configuration tab

Define the configuration settings, such as cloud provider, policy sets, target service accounts, and datacenters for the scan.

Field	Description
Cloud provider type	Cloud that hosts the resources to be scanned.
Policy sets	Policy sets for the scan. Select one or more policy sets for the scan, and then save the form.
Datacenter filters	Service account or service-account datacenter combination for the scan. **1.** In the empty row of the Datacenter Filters list, double-click the service account cell, and then select a service account. You can select an Amazon Web Services (AWS) management account, and then add all of its member accounts to the scan configuration, subject to the fulfillment of the following conditions: ○ An assume role configuration exists for the management account. ○ Cloud Discovery has completed the datacenter discovery for the member accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	For more information on creating assume role configurations, see Configure temporary credential access for trusted AWS accounts. 2. Save the form. 3. In the empty row of the Datacenter Filters list, double-click the datacenter cell, and then select a datacenter. (Optional) If you do not select a datacenter, Cloud Configuration Governance scans all the datacenters that are associated with the service account.

Filters tab

(Optional) Specify criteria to filter the cloud resources, and then import their configuration data from the cloud. If specified, Cloud Configuration Governance only fetches the configuration of the resources that match the filter criteria. Otherwise, it fetches the configuration of all the cloud resources.

Field	Description
Tag filters	Cloud resource tags to filter the cloud resources. • Enter the tag name in the **Name** field. • Enter the tag value in the **Value** field.
Freeform	Option to specify a comma-separated list of cloud identifiers to filter the cloud resources.
Resources	Comma-separated list of the cloud identifiers. This field appears only when the **Freeform** option is selected.

Cloud Configuration Governance Policy form

The Cloud Configuration Governance Policy form displays detailed information about the policy such as cloud provider, resource type, policy type, and policy violation reporting settings.

Field	Description
Policy name	Name that uniquely identifies the policy.
Description	Brief description of the policy.

Resource type

Define the resource type for which you want to create the policy.

Field	Description
Cloud provider	Cloud that hosts the resources to be scanned.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Resource type	Cloud resource type to be scanned through the policy. If the required resource type is not available, you can create a resource type. For more information, see Create a resource collector.

Policy condition

Define the policy type and the non-compliant resource configuration.

Field	Description
Type	Cloud Configuration Governance supports the following types: • **Condition builder**: The no-code method for creating policies. • **Integration Hub Flow**: The low-code method for creating policies. • **Script**: The code-based method for creating policies. Select the show available keys icon ( ) to view the list of all available configuration keys for the selected resource type. You can use any of the keys in the policy.

Condition	Conditions for reporting the non-compliant cloud resource configuration. Always specify the key and value in a pair. Use the **OR** operator and the **AND** operator to perform logical operations in the policy condition.
	
	Syntax
	
	`Key is <key_name>`
	`<data_type> Value <condition> <value>`
	
	For example,
	
	`Key is AWS:IAM:User:PasswordEnabled`
	`Boolean Value is true`
	
	This field appears only when **Condition Builder** is selected from the **Type** field.

Configuration key	Configuration keys for the policy. This field appears only when **Integration Hub Flow** is selected from the **Type** field.
Integration flow	The appropriate Integration Hub flow. This field appears only when **Integration Hub Flow** is selected from the **Type** field.
Condition script	Script that implements the policy conditions to identify and report the policy violations. Cloud Configuration Governance contains several scriptable objects and variables for use in the policy scripts. For more information see, Scripting reference.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	You can create script includes to externalize the decision making and reuse the code across different scripts. For more information on creating the script includes, see Script includes. **Note:** If you create a custom audit result record through the script, then the **Audit Violation Reporting** configuration defined in the policy doesn't take effect. This field appears only when **Script** is selected from the **Type** field.

Audit violation reporting

Define how Cloud Configuration Governance reports the policy violation.

Field	Description
Report violation as	Violation definition to be included in the audit violation report. Cloud Configuration Governance uses the violation definition to report the policy non-compliances. If an appropriate violation definition is not available, you can create one as follows: **1.** Select the lookup using list icon (). **2.** Select **New**. **3.** Enter a name that uniquely identifies the violation definition. Cloud Configuration Governance includes the name of the violation definition in the audit issue report. **4.** Select the default severity for the violation definition. **5.** (Optional) Enter a brief description of the violation.
Severity	Severity level of the violation. If you do not select the severity level in the policy, Cloud Configuration Governance uses the default severity defined in the violation definition.

Cloud Configuration Governance Configuration key form

The Configuration key form contains detailed information about the configuration key.

Configuration Key form

Name	Description
Name	Name that uniquely identifies the configuration key.
Description	Brief description of the configuration key.
Provider type	Cloud that hosts the resources to be scanned through the policies.
Resource type	Cloud resource type.
Value type	Data type of the configuration value that is associated with the configuration key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration Key form (continued)

Name	Description
Resource collector	Resource collector that can collect the configuration data for the selected resource type.
Configuration collector	(Optional) Configuration collector that can import the configuration keys. If the required configuration data is not available at the cloud resource level, select the configuration collector that can import the configuration keys. If a suitable configuration collector is not available, create a configuration collector and then select it. For more information on creating the configuration collector, see Create a configuration collector.

Cloud Configuration Governance Resource collector form

The Resource collector form contains detailed information about the resource collector.

Resource Collector form

Field	Description
Name	Name that uniquely identifies the configuration key.
Scan mode	Level at which the resource collector imports the resource from the cloud. Supported values are as follows: • **Run per LDC**: Import the resource at the datacenter level. That is, once from each datacenter. • **Run per service account**: Import the resource at the service account level.
Provider type	Cloud that hosts the resources to be imported.
Resource type	Cloud resource type that you want to import through the resource collector. If Cloud Configuration Governance doesn't contain the required resource type, you can create it as follows: **1.** Select the lookup using list icon ( ). **2.** Select **New**. **3.** In the **Name** field, enter the cloud resource type. **4.** Select **Submit**.
Flow	Resource collector flow that can import the given resource type.

Cloud Configuration Governance Service Account Assume Role Config form

The Cloud Configuration Governance Service Account Assume Role Config form displays detailed information about the assume role configuration used to access the trusting account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Account Assume Role Config form

Field	Description
Access Role Name	Name of the Identity and Access Management (IAM) role created for the trusting account that allows the trusted account (accessor) to access the trusting account. To scan the member accounts by using the management account, enter OrganizationAccountAccessRole in this field. If you don't want to use the OrganizationAccountAccessRole to access the member account, enter the name of the IAM role created for the member account. To scan the trusting account by using the trusted account, enter the name of the IAM role created for the trusting account.
Service Account	Name of the service account where assume role exists.
Role Session Name	An identifier for the assumed role session. Do not change the default value.
Credential Alias	Optional: Credential alias for the AWS credential.

Cloud Configuration Governance scripting reference

Cloud Configuration Governance provides several objects and variables that you can use to create script-based policies and CI finder mapping scripts.

Scripting reference for the Cloud Configuration Governance policies

Cloud Configuration Governance objects

Name	Description	Schema
`configSettings`	Represents the configuration data imported from the cloud.	`[ { "config_key":"configuration_key", "type":"data_type", "value":"configuration_value" }, ... ]`
`current`	Represents the current glide record object for the select resource type.	It contains the following fields: • Name: Name of the resource • Type: Resource type

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance objects (continued)

Name	Description	Schema
		• Identifier: Unique identifier of the cloud resource • Details: Cloud resource-specific details
`resourceInformation`	Resource record and its attributes.	`[ { "attribute": "Attribute name", "tableName" : "Referenced table", "documentId": "Referenced record", }, ... ]`
`auditResult`	Audit details that must be written to the Configuration Management Database (CMDB).	`{ "details": "Violation definition "' "severity": "severity of the violation", "auditViolationName": "Violation name" };`

Cloud Configuration Governance variables

Name	Data type	Description
`answer`	Boolean	Variable to specify whether the policy reports the violation or not. • `true`: Report the violation. Create the auditResult object to create a custom audit result record in the CMDB. If you do not create an audit result object, Cloud Configuration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance variables (continued)

Name	Data type	Description
		Governance reports the violation as per the violation definition specified in the policy. • `false`: Do not report the violation.
`violatingConfigSettings`	JSON	Reason of the policy violation. **Syntax** `{ "config_key": "value" };`

Scripting reference for the CI finder mapping script

Cloud Configuration Governance objects

Name	Description	Schema
`attributes`	A map containing the resource attribute key and attribute value from the Resource Attribute table for the given resource.	`{ "LogicalDatacenter": "Referenced record", "ServiceAccount": "Referenced record" }`

Cloud Configuration Governance variables

Name	Data type	Description
`Name`	String	Name of the resource.
`identifier`	Resource record	Identifier of the Cloud Configuration Governance resource record.
`type`	Resource record	Resource type of the Cloud Configuration Governance resource resource.
`answer`	JSON	CI class with which the resource needs to be mapped. The `answer` contains the following information: • `sysId` of the CI class. • `tableName` of the CI class.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance variables (continued)

Name	Data type	Description
		If the CI finder mapping fails to identify the CI class, set this variable to `null`. **Syntax** `{ sysId: "value", tableName: "value" }`

Cloud Configuration Governance actions reference

Cloud Configuration Governance (CCG) uses Integration Hub subflows to interact with the cloud and update the configuration data in the Configuration Management Database (CMDB).

CCG – Read Config Setting

Use this action to read the configuration data of the resource.

To use this action, insert an action and then navigate to **Action** > **Cloud Configuration Governance** > **Utils** > **CCG – Read Config Setting**.

CCG – Read Config Setting action

Field	Description
Resource [Resource]	Resource record that contains the configuration data.
Configuration key [Configuration Key]	Configuration key you want to read.

Assign Subflow Outputs

Assign Subflow Outputs form

Field	Description
Report issue	Option to enable the subflow to report the audit violation. Select the **Report Issue** option in the Data column or clear this check box to set or clear this field. • Selected: Report the issue as per the violation definition selected in the **Audit Violation Reporting** field of the policy. • Cleared: Cloud Configuration Governance doesn't report the violation. Create a custom record for the audit violation. You can specify conditions to control the creation of the audit violation record.
Details	Violation definition that you want to report for the violation. Enter the violation definition in the **Details** field in the Data column. This field is required if you've selected the **Report Issue** option.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create Record

Use this action to create a record in the CMDB.

To use this action, insert an action and then navigate to **Action** > **ServiceNow Core** > **Default** > **Create Record**.

Create Record action

Field	Description
Table	Name of the Configuration Management Database (CMDB) table where the audit result is stored. Set this field to Audit Result [sn_itom_ccg_audit_result].
Fields	Details of the record that you want to create in the Configuration Management Database (CMDB). Add the following fields and configure input for them: • **Scan Run**: Scan run during which Cloud Configuration Governance has identified the audit issue. • **Is Test Run**: Indicates whether Cloud Configuration Governance has reported the audit issue during a test run. • **Details**: Details of the violation. • **Violation Definition**: Violation definition of the audit issue. • **Resource**: Cloud resource for which Cloud Configuration Governance has raised the audit issue. • **Severity**: Severity of the audit issue.

CCG – Insert Resource Record

Use this action to insert a resource record to the Configuration Management Database (CMDB).

To use this action, insert an action and then navigate to **Action** > **Cloud Configuration Governance** > **Utils** > **CCG – Insert Resource Record**.

CCG – Insert Resource Record action

Field	Description
Scan run	Scan run for which the subflow must create the resource record.
Service account	Service account to which the resource is attached.
Logical datacenter	Logical datacenter to which the resource is attached.
Identifier	Identifier of the resource record.
Name	Name of the resource.
Type	Resource type.
Provider	Cloud provider that hosts the resource.
Details	Details of the object that you want to store in the resource record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CCG – Insert Resource Record action (continued)

Field	Description
Attributes	Any additional resource attribute that you want to import to the CMDB.

Cloud Action Library actions reference

Cloud Action Library (CAL) offers various actions that you can use to build the Integration Hub subflows for Cloud Configuration Governance.

Cloud Action Library offers the following actions:

- CAL - AWS generate credential report
- CAL - AWS S3 get encryption
- CAL - Azure VM monitoring metric data
- AWS turn off EC2 instance
- CAL - AWS list VMs
- CAL - AWS list S3 buckets
- AWS put S3 bucket encryption
- AWS enable EC2 instance monitoring
- CAL - Azure VM IP metric
- Azure turn off virtual machine
- AWS get IAM login profile
- Disable AWS IAM user
- AWS list images
- CAL - Azure list VMs

The Cloud Action Library actions are available under the **Default** and **External** categories.

Cloud permissions for collecting Cloud Configuration Governance keys

The Cloud Configuration Governance requires appropriate cloud permissions to collect the base system configuration keys from the cloud. Therefore, you must set the appropriate permissions in the cloud to suit the needs of your organization.

Note: Starting with Cloud Configuration Governance version 1.3.7, the base system contents are moved to the CCG Content Pack. Install the CCG Content Pack to access the base system Cloud Configuration Governance contents.

Amazon Web Services (AWS) datacenter

Cloud Configuration Governance uses the following items to collect the configuration key for the AWS datacenter configuration keys:

- Resource collector: Cloud Service Account
- Cloud API used: Action: `DescribeRegions`
- Cloud permission: `ec2: DescribeRegions`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS datacenter configuration keys

Configuration key	Data type
AWS:EC2:VM:DescribeRegions	*String*

AWS Identity and Access Management (IAM) users

Cloud Configuration Governance uses the following items to collect the configuration key for the AWS IAM user configuration keys:

- Resource collector: AWS IAM User Data Collector
- Cloud API used:
  - Action: *GetCredentialReport* and *GenerateCredentialReport*
  - Service: *AWS IAM service*
- Cloud permission: *Iam:GetCredentialReport* and *Iam:GenerateCredentialReport*

AWS IAM user configuration keys

Configuration key	Data type
AWS:IAM:Policy:ARN	*String*
AWS:IAM:Policy:AttachmentCount	*String*
AWS:IAM:Policy:CreateDate	*String*
AWS:IAM:Policy:PolicyName	*String*
AWS:IAM:Policy:UpdateDate	*String*
AWS:IAM:User:AccessKey1.active	*Boolean*
AWS:IAM:User:AccessKey1.lastRotated	*Date*
AWS:IAM:User:AccessKey1.lastUsedDate	*Date*
AWS:IAM:User:AccessKey1.lastUsedRegion	*String*
AWS:IAM:User:AccessKey1.lastUsedService	*String*
AWS:IAM:User:AccessKey2.active	*Boolean*
AWS:IAM:User:AccessKey2.lastRotated	*Date*
AWS:IAM:User:AccessKey2.lastUsedDate	*Date*
AWS:IAM:User:AccessKey2.lastUsedRegion	*String*
AWS:IAM:User:AccessKey2.lastUsedService	*String*
AWS:IAM:User:Certificate1.active	*Boolean*
AWS:IAM:User:Certificate1.lastRotated	*Date*
AWS:IAM:User:Certificate2.active	*Boolean*
AWS:IAM:User:Certificate2.lastRotated	*Date*
AWS:IAM:User:CreationTime	*Date*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS IAM user configuration keys (continued)

Configuration key	Data type
AWS:IAM:User:LoginProfile.active	*Boolean*
AWS:IAM:User:MfaEnabled	*Boolean*
AWS:IAM:User:PasswordEnabled	*Boolean*
AWS:IAM:User:PasswordLastChanged	*String*
AWS:IAM:User:PasswordLastUsed	*Date*
AWS:IAM:User:PasswordNextRotation	*String*

AWS object storage

Cloud Configuration Governance uses the following items to collect the configuration key for the AWS IAM user configuration keys:

- Configuration collector: AWS S3 Encryption Metric Collector
- Resource collector: AWS S3 Resource Collector
- Cloud API used: Action: *ListBuckets* and *GetBucketEncryption* on S3 service
- Cloud permission: *s3:ListBucket* and *s3:GetEncryptionConfiguration*

AWS object storage configuration keys

Configuration key	Data type
AWS:S3:Encryption:BucketKeyEnabled	*Boolean*
AWS:S3:Encryption:KMSMasterKeyID	*String*
AWS:S3:Encryption:ServerSideEncryptionEnabled	*Boolean*
AWS:S3:Encryption:SSEAlgorithm	*String*

- Configuration collector: AWS S3 ACL Permission Metric Collector
- Resource collector: AWS S3 Resource Collector
- Cloud API used: Action: *GetBucketAcl*
- Cloud permission: *s3:GetBucketAcl*

AWS object storage configuration keys

Configuration key	Data type
AWS:S3:ACL:AuthnUsersListing	*Boolean*
AWS:S3:ACL:AuthnUsersReadACL	*Boolean*
AWS:S3:ACL:AuthnUsersWrite	*Boolean*
AWS:S3:ACL:AuthnUsersWriteACL	*Boolean*
AWS:S3:ACL:OwnerFullControl	*Boolean*
AWS:S3:ACL:OwnerId	*String*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS object storage configuration keys (continued)

Configuration key	Data type
AWS:S3:ACL:OwnerListing	*Boolean*
AWS:S3:ACL:OwnerName	*String*
AWS:S3:ACL:OwnerReadACL	*Boolean*
AWS:S3:ACL:OwnerWrite	*Boolean*
AWS:S3:ACL:OwnerWriteACL	*Boolean*
AWS:S3:ACL:PublicListing	*Boolean*
AWS:S3:ACL:PublicReadACL	*Boolean*
AWS:S3:ACL:PublicWrite	*Boolean*
AWS:S3:ACL:PublicWriteACL	*Boolean*

AWS virtual machine instance

Cloud Configuration Governance uses the following items to collect the configuration key for the AWS virtual machine instance configuration keys:

- Resource collector: AWS VM Data Collector
- Cloud API used: Action: *DescribeInstances* and *AWS EC2 resource*
- Cloud permission: *ec2:DescribeInstances*

AWS virtual machine instance configuration keys

Configuration key	Data type
AWS:EC2:VM:CapacityReservationPreference	*String*
AWS:EC2:VM:CpuOptionsCoreCount	*Numeric*
AWS:EC2:VM:CpuOptionsThreadsPerCore	*Numeric*
AWS:EC2:VM:EbsOptimized	*Boolean*
AWS:EC2:VM:HardwareType	*String*
AWS:EC2:VM:ImageId	*String*
AWS:EC2:VM:InstanceState	*String*
AWS:EC2:VM:KeyName	*String*
AWS:EC2:VM:LaunchTime	*Date*
AWS:EC2:VM:MonitoringState	*String*
AWS:EC2:VM:Platform	*String*
AWS:EC2:VM:PrivateDnsName	*String*
AWS:EC2:VM:PrivateIpAddress	*String*
AWS:EC2:VM:PublicDnsName	*String*
AWS:EC2:VM:PublicIPAddress	*String*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS virtual machine instance configuration keys (continued)

Configuration key	Data type
AWS:EC2:VM:SecurityGroups	*String*
AWS:EC2:VM:SubnetId	*String*
AWS:EC2:VM:Tags	*Map*
AWS:EC2:VM:UsageOperation	*String*
AWS:EC2:VM:VpcId	*String*

AWS profile with minimal permissions

```
{
	"Version": "2012-10-17",
	"Statement": [
		{
			"Sid": "VisualEditor0",
			"Effect": "Allow",
			"Action": [
				"iam:GenerateCredentialReport",
				"s3:GetEncryptionConfiguration",
				"ec2:DescribeInstances",
				"s3:ListBucketVersions",
				"ec2:DescribeRegions",
				"s3:ListBucket",
				"iam:GetCredentialReport"
			],
			"Resource": "*"
		}
	]
}
```

Microsoft Azure virtual machine instance

Cloud Configuration Governance uses the following items to collect the Azure virtual machine instance configuration keys:

- Resource collector: Azure VM Data Collector
- Cloud API used: *Microsoft.ResourceGraph/resources*
- Cloud permission: *Microsoft.ResourceGraph/resources*

Azure virtual machine instance configuration keys

Configuration key	Data type
Azure:VM:HardwareType	*String*
Azure:VM:NICID	*String*
Azure:VM:OSDiskCaching	*String*
Azure:VM:OSDiskCreateoption	*String*
Azure:VM:OSDiskDeleteoption	*String*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure virtual machine instance configuration keys (continued)

Configuration key	Data type
Azure:VM:OSDiskId	*String*
Azure:VM:OSDiskName	*String*
Azure:VM:OSDiskOSType	*String*
Azure:VM:OSDiskSizeGB	*String*
Azure:VM:OSProfileAllowExtensionOperations	*Boolean*
Azure:VM:OSProfileComputerName	*String*
Azure:VM:OSProfileLinuxConfigurationDisablePasswordAuthentication	*Boolean*
Azure:VM:OSProfileLinuxConfigurationPatchSettingsAssessmentMode	*String*
Azure:VM:OSProfileLinuxConfigurationPatchSettingsPatchMode	*String*
Azure:VM:OSProfileLinuxConfigurationProvisionVmAgent	*Boolean*
Azure:VM:OSProfileLinuxConfigurationSSHKeyData	*Map*
Azure:VM:OSProfileLinuxConfigurationSSHPath	*Map*
Azure:VM:OSProfileRequireGuestProvisionSignal	*Boolean*
Azure:VM:OSWindowsConfigurationEnableAutomaticUpdates	*Boolean*
Azure:VM:OSWindowsConfigurationPatchSettingsAssessmentMode	*String*
Azure:VM:OSWindowsConfigurationPatchSettingsEnableHotpatching	*Boolean*
Azure:VM:OSWindowsConfigurationPatchSettingsPatchMode	*String*
Azure:VM:OSWindowsConfigurationProvisionVMAgent	*Boolean*
Azure:VM:PowerState	*String*
Azure:VM:ProvisioningState	*String*
Azure:VM:ResourceGroup	*String*
Azure:VM:StorageProfileDataDisksCaching	*String*
Azure:VM:StorageProfileDataDisksCreateOption	*String*
Azure:VM:StorageProfileDataDisksDeleteOption	*String*
Azure:VM:StorageProfileDataDisksDetachOption	*String*
Azure:VM:StorageProfileDataDisksDiskIopsReadWrite	*String*
Azure:VM:StorageProfileDataDisksDiskMBpsReadWrite	*String*
Azure:VM:StorageProfileDataDisksDiskSizeGb	*Numeric*
Azure:VM:StorageProfileDataDisksImage	*String*
Azure:VM:StorageProfileDataDisksLun	*Numeric*
Azure:VM:StorageProfileDataDisksManagedDiskDiskEncryptionSet	*String*
Azure:VM:StorageProfileDataDisksManagedDiskId	*String*
Azure:VM:StorageProfileDataDisksManagedDiskResourceGroup	*String*
Azure:VM:StorageProfileDataDisksManagedDiskStorageAccountType	*String*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure virtual machine instance configuration keys (continued)

Configuration key	Data type
Azure:VM:StorageProfileDataDisksManagedStorageAccountType	*String*
Azure:VM:StorageProfileDataDisksName	*String*
Azure:VM:StorageProfileDataDisksToBeDetached	*Boolean*
Azure:VM:StorageProfileDataDisksVhd	*String*
Azure:VM:StorageProfileDataDisksWriteAcceleratorEnabled	*Boolean*
Azure:VM:StorageProfileImageReferenceExactVersion	*String*
Azure:VM:StorageProfileImageReferenceId	*String*
Azure:VM:StorageProfileImageReferenceOffer	*String*
Azure:VM:StorageProfileImageReferencePublisher	*String*
Azure:VM:StorageProfileImageReferenceSharedGalleryImageId	*String*
Azure:VM:StorageProfileImageReferenceSku	*String*
Azure:VM:StorageProfileImageReferenceVersion	*String*
Azure:VM:Tags	*Map*
Azure:VM:VMId	*String*

- Resource collector: Azure VM Data Collector
- Configuration collector: Azure VM Metric Collector
- Cloud API used: *Microsoft.ResourceGraph/resources*
- Cloud permission: *Microsoft.ResourceGraph/resources*

Azure virtual machine instance configuration keys

Configuration key	Data type
Azure:VM:PublicIPAddress	*String*
Azure:VM:PublicIPId	*String*

- Resource collector: Azure VM Data Collector
- Configuration collector: Azure VM Monitoring Metric Collector
- Cloud API used: *Microsoft.Compute/virtualMachines/{vmName}/instanceView*
- Cloud permission: *Microsoft.Compute/virtualMachines/{vmName}/instanceView*

Azure virtual machine instance configuration keys

Configuration key	Data type
Azure:VM:MonitoringState	*String*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Use *scope=https://graph.microsoft.com/.default* and *scope=https://management.azure.com/.default* to fetch the oAuth token for the Azure resources.

Azure profile with minimal permissions

```
{
    "properties": {
        "roleName": "CCGAzureMinimalPermission",
        "description": "Grants access to scan compute resources from azure subscription",
        "assignableScopes": [
            "/subscriptions/${subscription_id}"
        ],
        "permissions": [
            {
                "actions": [
                    "Microsoft.ResourceGraph/resources/read",
                    "Microsoft.Compute/virtualMachines/instanceView/read"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
}
```

AWS policies

The Cloud Configuration Governance AWS policies are listed for your reference.

Cloud Configuration Governance AWS policies

Policy Set	Policy Name	Resource Type	Description
	AWS IAM User Activity policy	Condition builder	Policy to check if the password is enabled for the AWS IAM user. To use this policy, ensure that the AWS IAM user account has the following permissions: • Iam:GetCredentialReport • Iam:GenerateCredentialReport
	AWS S3 Enforce Bucket encryption	Condition builder	Policy to check if the AWS S3 buckets are encrypted.
	AWS Sample flow policy	Integration Hub flow	Policy to illustrate an Integration Hub flow-based policy.
	AWS VM HardwareType	Condition builder	Policy to check if the deployed EC2 VMs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			are using only the approved hardware types.
	AWS VM IPAddress	Script	Policy to check if the IP address of the EC2 VM is matching with the Configuration Management Database (CMDB) record.
	AWS VM Monitoring State	Condition builder	Policy to check if detailed monitoring is enabled for the EC2 VM.
AWS: IAM Security	Ensure multi-factor authentication (MFA) is enabled for all IAM users that have a console password (Automated)	IAM User	Multi-Factor Authentication (MFA) adds an extra layer of authentication assurance beyond traditional credentials. With MFA enabled, when a user signs in to the AWS Console, they will be prompted for their user name and password as well as for an authentication code from their physical or virtual MFA token. It is recommended that MFA be enabled for all accounts that have a console password. Enabling MFA provides increased security for console access as it requires the authenticating principal to possess a device that displays a time-sensitive key and have knowledge of a credential.
	Ensure credentials unused for 45 days or greater are disabled (Automated)	IAM User	AWS IAM users can access AWS resources using different types of

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			credentials, such as passwords or access keys. It is recommended that all credentials that have been unused in 45 or greater days be deactivated or removed Disabling or removing unnecessary credentials will reduce the window of opportunity for credentials associated with a compromised or abandoned account to be used.
	Ensure there is only one active access key available for any single IAM user (automated)	IAM User	Access keys are long-term credentials for an IAM user or the AWS account 'root' user. You can use access keys to sign programmatic requests to the AWS CLI or AWS API (directly or using the AWS SDK) Access keys are long-term credentials for an IAM user or the AWS account 'root' user. You can use access keys to sign programmatic requests to the AWS CLI or AWS API. One of the best ways to protect your account is to not allow users to have multiple access keys.
	Ensure access keys are rotated every 90 days or less (automated)	IAM User	Access keys consist of an access key ID and secret access key, which are used

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			to sign programmatic requests that you make to AWS. AWS users need their own access keys to make programmatic calls to AWS from the AWS Command Line Interface (AWS CLI), Tools for Windows PowerShell, the AWS SDKs, or direct HTTP calls using the APIs for individual AWS services.It is recommended that all access keys be regularly rotated. Rotating access keys will reduce the window of opportunity for an access key that is associated with a compromised or terminated account to be used. Access keys should be rotated to ensure that data cannot be accessed with an old key which might have been lost, cracked, or stolen.
	Ensure IAM Users Receive Permissions Only Through Groups (Automated)	IAM User	IAM users are granted access to services, functions, and data through IAM policies. There are three ways to define policies for a user:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			1. Edit the user policy directly, that is an inline, user, policy. 2. Attach a policy directly to a user. 3. Add the user to an IAM group that has an attached policy. Only the third implementation is recommended. Assigning IAM policy only through groups unifies permissions management to a single, flexible layer consistent with organizational functional roles. By unifying permissions management, the likelihood of excessive permissions is reduced.
	Ensure IAM policies that allow full "*:*" administrative privileges are not attached (Automated)	IAM User	IAM policies are the means by which privileges are granted to users, groups, or roles. It is recommended and considered a standard security advice to grant least privilege-that is, granting only the permissions required to perform a task. Determine what users need to do and then craft policies for them that let the users perform only those tasks, instead of allowing full administrative privileges. It's more secure to start with a minimum

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			set of permissions and grant additional permissions as necessary, rather than starting with permissions that are too lenient and then trying to tighten them later. Providing full administrative privileges instead of restricting to the minimum set of permissions that the user is required to do exposes the resources to potentially unwanted actions. IAM policies that have a statement with "Effect": "Allow" with "Action": "" over "Resource": "" should be removed.
	Ensure a support role has been created to manage incidents with AWS Support (Automated)	IAM User	AWS provides a support center that can be used for incident notification and response, as well as technical support and customer services. Create an IAM Role to allow authorized users to manage incidents with AWS Support. By implementing least privilege for access control, an IAM Role will require an appropriate IAM Policy to allow Support Center Access in order to manage Incidents with AWS Support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
	Ensure that all the expired SSL/TLS certificates stored in AWS IAM are removed (Automated)	IAM User	To enable HTTPS connections to your website or application in AWS, you need an SSL/TLS server certificate. You can use ACM or IAM to store and deploy server certificates. Use IAM as a certificate manager only when you must support HTTPS connections in a region that is not supported by ACM. IAM securely encrypts your private keys and stores the encrypted version in IAM SSL certificate storage. IAM supports deploying server certificates in all regions, but you must obtain your certificate from an external provider for use with AWS. You cannot upload an ACM certificate to IAM. Additionally, you cannot manage your certificates from the IAM Console. Removing expired SSL/TLS certificates eliminates the risk that an invalid certificate will be deployed accidentally to a resource such as AWS Elastic Load Balancer (ELB), which can damage the credibility of the application/website behind the ELB. As a best practice, it is recommended

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			to delete expired certificates.
	Ensure that IAM Access analyzer is enabled for all regions (Automated)	AWS::AccessAnalyzer::Analyzer	Enable IAM Access analyzer for IAM policies about all resources in each region. IAM Access Analyzer is a technology introduced at AWS reinvent 2019. After the Analyzer is enabled in IAM, scan results are displayed on the console showing the accessible resources. Scans show resources that other accounts and federated users can access, such as KMS keys and IAM roles. So the results allow you to determine if an unintended user is allowed, making it easier for administrators to monitor least privileges access. Access Analyzer analyzes only policies that are applied to resources in the same AWS Region. AWS IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. This lets you identify unintended access to your resources and data. Access Analyzer identifies resources

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			that are shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment. IAM Access Analyzer continuously monitors all policies for S3 bucket, IAM roles, KMS(Key Management Service) keys, AWS Lambda functions, and Amazon SQS(Simple Queue Service) queues.
	Ensure no 'root' user account access key exists (Automated)	IAM User	The 'root' user account is the most privileged user in an AWS account. AWS Access Keys provide programmatic access to a given AWS account. It is recommended that all access keys associated with the 'root' user account be removed. Removing access keys associated with the 'root' user account limits vectors by which the account can be compromised. Additionally, removing the 'root' access keys encourages the creation and use of role based accounts that are least privileged.
	Ensure MFA is enabled for the 'root' user account (Automated)	IAM User	The 'root' user account is the most privileged user in an AWS account. Multi-

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			factor Authentication (MFA) adds an extra layer of protection on top of a username and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for their username and password as well as for an authentication code from their AWS MFA device. Enabling MFA provides increased security for console access as it requires the authenticating principal to possess a device that emits a time-sensitive key and have knowledge of a credential.
	Ensure hardware MFA is enabled for the 'root' user account (Automated)	IAM User	The 'root' user account is the most privileged user in an AWS account. MFA adds an extra layer of protection on top of a user name and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for their user name and password as well as for an authentication code from their AWS MFA device. For Level 2, it is recommended that the 'root' user account be protected with a hardware MFA. A hardware MFA has a smaller attack surface than a virtual MFA. For example, a hardware

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			MFA does not suffer the attack surface introduced by the mobile smartphone on which a virtual MFA resides.
	Eliminate use of the 'root' user for administrative and daily tasks (Automated)	IAM User	With the creation of an AWS account, a 'root user' is created that cannot be disabled or deleted. That user has unrestricted access to and control over all resources in the AWS account. It is highly recommended that the use of this account be avoided for everyday tasks. The 'root user' has unrestricted access to and control over all account resources. Use of it is inconsistent with the principles of least privilege and separation of duties, and can lead to unnecessary harm due to error or account compromise.
	Ensure IAM password policy requires minimum length of 14 or greater (Automated)	IAM User	Password policies are, in part, used to enforce password complexity requirements. IAM password policies can be used to ensure password are at least a given length. It is recommended that the password policy require a minimum password length 14. Setting a password complexity policy increases account

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			resiliency against brute force login attempts.
	Ensure IAM password policy prevents password reuse (Automated)	IAM User	IAM password policies can prevent the reuse of a given password by the same user. It is recommended that the password policy prevent the reuse of passwords. Preventing password reuse increases account resiliency against brute force login attempts.
AWS: Data Security	Ensure that S3 Buckets are configured with 'Block public access (bucket settings)	AWS::S3::Bucket	Amazon S3 provides Block public access (bucket settings)and Block public access (account settings)to help you manage public access to Amazon S3 resources. By default, S3 buckets and objects arecreated with public access disabled. However, an IAM principal with sufficient S3 permissions can enable public access at the bucket and/or object level. While enabled, Block public access (bucket settings)prevents an individual bucket, and its containedobjects, from becoming publicly accessible. Similarly, Block public access (account settings)prevents all buckets, and contained objects,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			from becoming publicly accessible across the entire account. Amazon S3 Block public access (bucket settings) prevents the accidental or malicious public exposure of data contained within the respective bucket(s). Amazon S3 Block public access (account settings) prevents the accidental or malicious public exposure of data contained within all buckets of the respective AWS account. Whether blocking public access to all or some buckets is an organizational decision that should be based on data sensitivity, least privilege, and use case.
AWS: CloudTrail Logging	Ensure that Object-level logging for write events is enabled for S3 bucket (Automated)	AWS::S3::Bucket	S3 object-level API operations such as GetObject, DeleteObject, and PutObject are called data events. By default, CloudTrail trails don't log data events and so it is recommended to enable Object-level logging for S3 buckets. Rationale:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			Enabling object-level logging will help you meet data compliance requirements within your organization, perform comprehensive security analysis, monitor specific patterns of user behavior in your AWS account or take immediate actions on any object-level API activity within your S3 Buckets using Amazon CloudWatch Events.
AWS: CloudTrail Logging	Ensure that Object-level logging for read events is enabled for S3 bucket (Automated)	AWS::S3::Bucket	S3 object-level API operations such as GetObject, DeleteObject, and PutObject are called data events. By default, CloudTrail trails don't log data events and so it is recommended to enable Object-level logging for S3 buckets. Rationale: Enabling object-level logging will help you meet data compliance requirements within your organization, perform comprehensive security analysis, monitor specific patterns of user behavior in your AWS account or take immediate actions on any object-level API

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			activity using Amazon CloudWatch Events.
AWS: Network Security	Ensure no Network ACLs allow ingress from 0.0.0.0/0 to remote server administration ports (Automated)	AWS::EC2::NetworkAcl	The Network Access Control List (NACL) function provide stateless filtering of ingress and egress network traffic to AWS resources. It is recommended that no NACL allows unrestricted ingress access to remote server administration ports, such as SSH to port 22and RDP to port 3389. Public access to remote server administration ports, such as 22 and 3389, increases resource attack surface and unnecessarily raises the risk of resource compromise.
AWS: Network Security	Ensure no security groups allow ingress from 0.0.0.0/0 to remote server administration ports (Automated)	AWS::EC2::SecurityGroup	Security groups provide stateful filtering of ingress and egress network traffic to AWS resources. It is recommended that no security group allows unrestricted ingress access to remote server administration ports, such as SSH to port 22 and RDP to port 3389 Public access to remote server administration ports, such as 22 and 3389, increases resource attack surface and unnecessarily raises

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance AWS policies (continued)

Policy Set	Policy Name	Resource Type	Description
			the risk of resource compromise.

Azure policies

The Cloud Configuration Governance AWS policies are listed for your reference.

Cloud Configuration Governance Azure policies

Policy Set	Policy Name	Resource Type	Description
	Azure VM HardwareType	Condition builder	Policy to check if the deployed Azure VMs are using only the approved hardware types.
	Azure VM IP Address	Script	Policy to check if the IP address of the Azure VM is matching with the CMDB record.
	Azure VM Monitoring State	Condition builder	Policy to check if detailed monitoring is enabled for the Azure VM.
Azure: Security Center Configuration	Ensure That Auto provisioning of 'Log Analytics agent for Azure VMs' is Set to 'On' (Automated)	Azure Subscription	Enable automatic provisioning of the monitoring agent to collect security data. When 'Log Analytics agent for Azure VMs' is turned on, Microsoft Defender for Cloud provisions the Microsoft Monitoring Agent on all existing supported Azure virtual machines and any new ones that are created. The Microsoft Monitoring Agent scans for various security-related configurations and events such as system updates, OS vulnerabilities, endpoint protection, and provides alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
Azure: Security Center Configuration	Ensure 'Additional email addresses' is Configured with a Security Contact Email (Automated)	Azure Subscription	Microsoft Defender for Cloud emails the subscription owners whenever a high-severity alert is triggered for their subscription. You should provide a security contact email address as an additional email address. Microsoft Defender for Cloud emails the Subscription Owner to notify them about security alerts. Adding your Security Contact's email address to the 'Additional email addresses' field ensures that your organization's Security Team is included in these alerts. This ensures that the proper people are aware of any potential compromise in order to mitigate the risk in a timely fashion.
Azure: Security Center Configuration	Ensure That 'Notify about alerts with the following severity' is Set to 'High' (Automated)	Azure Subscription	Enables emailing security alerts to the subscription owner or other designated security contact. Enabling security alert emails ensures that security alert emails are received from Microsoft. This ensures that the right people are aware of any potential security issues and are able to mitigate the risk.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
Azure: Security Center Configuration	Ensure That 'All users with the following roles' is set to 'Owner' (Automated)	Azure Subscription	Enable security alert emails to subscription owners. Enabling security alert emails to subscription owners ensures that they receive security alert emails from Microsoft. This ensures that they are aware of any potential security issues and can mitigate the risk in a timely fashion.
Azure: Storage Security	Ensure Storage logging is Enabled for Blob Service for 'Read', 'Write', and 'Delete' requests (Automated)	Microsoft.Storage/ StorageAccounts	The Storage Blob service provides scalable, cost-efficient objective storage in the cloud. Storage Logging happens server-side and allows details for both successful and failed requests to be recorded in the storage account. These logs allow users to see the details of read, write, and delete operations against the blobs. Storage Logging log entries contain the following information about individual requests: Timing information such as start time, end-to-end latency, and server latency, authentication details, concurrency information and the sizes of the request and response messages. Storage Analytics logs contain detailed

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			information about successful and failed requests to a storage service. This information can be used to monitor each individual request to a storage service for increased security or diagnostics. Requests are logged on a best-effort basis. Storage Analytics logging is not enabled by default for your storage account.
Azure: Storage Security	Ensure Storage Logging is Enabled for Table Service for 'Read', 'Write', and 'Delete' Requests (Automated)	Microsoft.Storage/ StorageAccounts	The Storage Table storage is a service that stores structure NoSQL data in the cloud, providing a key/ attribute store with a schema less design. Storage Logging happens server-side and allows details for both successful and failed requests to be recorded in the storage account. These logs allow users to see the details of read, write, and delete operations against the tables. Storage Logging log entries contain the following information about individual requests: Timing information such as start time, end-to-end latency, and server latency, authentication details, concurrency information and the sizes of the request and response messages.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			Storage Analytics logs contain detailed information about successful and failed requests to a storage service. This information can be used to monitor each individual request to a storage service for increased security or diagnostics. Requests are logged on a best-effort basis. Storage Analytics logging is not enabled by default for your storage account.
Azure: Storage Security	Ensure Storage Logging is Enabled for Queue Service for 'Read', 'Write', and 'Delete' requests (Automated)	Microsoft.Storage/StorageAccounts	The Storage Queue service stores messages that may be read by any client who has access to the storage account. A queue can contain an unlimited number of messages, each of which can be up to 64KB in size using version 2011-08-18 or newer. Storage Logging happens server-side and allows details for both successful and failed requests to be recorded in the storage account. These logs allow users to see the details of read, write, and delete operations against the queues. Storage Logging log entries contain the following information about individual requests: Timing information such as start time, end-to-end latency,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			and server latency, authentication details, concurrency information and the sizes of the request and response messages. Storage Analytics logs contain detailed information about successful and failed requests to a storage service. This information can be used to monitor individual request and to diagnose issues with a storage service. Requests are logged on a best-effort basis.Storage Analytics logging is not enabled by default for your storage account.
Azure: SQL Security	Ensure that Vulnerability Assessment (VA) is enabled on a SQL server by setting a Storage Account (Automated)	Microsoft.Sql/servers	Enable Vulnerability Assessment (VA) service scans for critical SQL servers and corresponding SQL databases.
Azure: SQL Security	Ensure that VA setting 'Periodic recurring scans' to 'on' for each SQL server (Automated)	Microsoft.Sql/servers	Enable Vulnerability Assessment (VA) Periodic recurring scans for critical SQL servers and corresponding SQL databases.
Azure: SQL Security	Ensure that VA setting 'Send scan reports to' is configured for a SQL server (Automated)	Microsoft.Sql/servers	Configure 'Send scan reports to' with email IDs of concerned data owners/stakeholders for a critical SQL servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
Azure: SQL Security	Ensure that Vulnerability Assessment Setting 'Also send email notifications to admins and subscription owners' is Set for	Microsoft.Sql/servers	Enable Vulnerability Assessment (VA) setting 'Also send email notifications to admins and subscription owners'.
Azure: SQL Security	Ensure that Azure Active Directory Admin is configured (Automated)	Microsoft.Sql/servers	Use Azure Active Directory Authentication for authentication with SQL Database to manage credentials in a single place.
Azure: SQL Security	Ensure SQL server's TDE protector is encrypted with Customermanaged key (Automated)	Microsoft.Sql/servers	TDE with Customer-managed key support provides increased transparency and control over the TDE Protector, increased security with an HSM-backed external service, and promotion of separation of duties. With TDE, data is encrypted at rest with a symmetric key (called the database encryption key) stored in the database or data warehouse distribution. To protect this data encryption key (DEK) in the past, only a certificate that the Azure SQL Service managed could be used. Now, with Customer-managed key support for TDE, the DEK can be protected with an asymmetric key that is stored in the Key Vault. Key Vault is a highly available and scalable cloud-based key store

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			which offers central key management, leverages FIPS 140-2 Level 2 validated hardware security modules (HSMs), and allows separation of management of keys and data, for additional security. Based on business needs or criticality of data/databases hosted a SQL server, it is recommended that the TDE protector is encrypted by a key that is managed by the data owner (Customer-managed key).
Azure: Logging and Monitoring	Ensure Diagnostic Setting captures appropriate categories (Automated)	Azure Subscription	The diagnostic setting should be configured to log the appropriate activities from the control/management plane. A diagnostic setting controls how the diagnostic log is exported. Capturing the diagnostic setting categories for appropriate control/ management plane activities allows proper alerting A diagnostic setting controls how the diagnostic log is exported. Capturing the diagnostic setting categories for appropriate control/ management plane activities allows proper alerting.
Azure: Logging and Monitoring	Ensure that logging for Azure KeyVault is 'Enabled' (Automated)	Microsoft.KeyVault	Enable AuditEvent logging for key vault

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			instances to ensure interactions with key vaults are logged and available. Monitoring how and when key vaults are accessed, and by whom enables an audit trail of interactions with confidential information, keys and certificates managed by Azure Keyvault. Enabling logging for Key Vault saves information in an Azure storage account that the user provides. This creates a new container named insights-logs-auditevent automatically for the specified storage account, and this same storage account can be used for collecting logs for multiple key vaults.
Azure: Data Security	Ensure that 'Unattached disks' are encrypted with CMK (Automated)	Microsoft.Compute/disks	Ensure that unattached disks in a subscription are encrypted with a Customer Managed Key (CMK). Managed disks are encrypted by default with Platform-managed keys. Using Customermanaged keys may provide an additional level of security or meet an organization's regulatory requirements.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			Encrypting managed disks ensures that its entire content is fully unrecoverable without a key and thus protects the volume from unwarranted reads. Even if the disk is not attached to any of the VMs, there is always a risk where a compromised user account with administrative access to VM service can mount/attach these data disks which may lead to sensitive information disclosure and tampering.
Azure: Access Control	Ensure the key vault is recoverable (Automated)	Microsoft.KeyVault	The key vault contains object keys, secrets and certificates. Accidental unavailability of a key vault can cause immediate data loss or loss of security functions (authentication, validation, verification, non-repudiation, etc.) supported by the key vault objects. It is recommended the key vault be made recoverable by enabling the "Do Not Purge" and "Soft Delete" functions. This is in order to prevent loss of encrypted data including storage accounts, SQL databases, and/or dependent services provided by key vault objects (Keys, Secrets,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			Certificates) etc., as may happen in the case of accidental deletion by a user or from disruptive activity by a malicious user. There could be scenarios where users accidently run delete/ purge commandson key vault or attacker/ malicious user does it deliberately to cause disruption. Deleting or purging a key vault leads to immediate data loss as keys encrypting data and secrets/certificates allowing access/ services will become non-accessible. There are 2 key vault properties that plays role in permanent unavailability of a key vault.1.enableSoftDelete:Setting this parameter to true for a key vault ensures that even if key vault is deleted, Key vault itself or its objects remain recoverable for next 90days. In this span of 90 days either key vault/objects can be recovered or purged (permanent deletion). If no action is taken, after 90 days key vault and its objects will be purged.2.enablePurgeProtection:enable only ensures that key vault is not deleted permanently and will be recoverable for 90 days from date of deletion. However, there are chances that the key vault

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
			and/or its objects are accidentally purged and hence will not be recoverable. Setting enablePurgeProtection to "true" ensures that the key vault and its objects cannot be purged.Enabling both the parameters on key vaults ensures that key vaults and their objects cannot be deleted/purged permanently.
Azure: App Service Security	Ensure Web App Redirects All HTTP traffic to HTTPS in Azure App Service (Automated)	Microsoft.Web/sites	Azure Web Apps allows sites to run under both HTTP and HTTPS by default. Web apps can be accessed by anyone using non-secure HTTP links by default. Non-secure HTTP requests can be restricted and all HTTP requests redirected to the secure HTTPS port. It is recommended to enforce HTTPS-only traffic. Enabling HTTPS-only traffic will redirect all non-secure HTTP request toHTTPS ports. HTTPS uses the TLS/SSL protocol to provide a secure connection, which is both encrypted and authenticated. So it is important to support HTTPS for the security benefits.
Azure: App Service Security	Ensure Web App is using the latest version of	Microsoft.Web/sites	The TLS(Transport Layer Security) protocol secures

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance Azure policies (continued)

Policy Set	Policy Name	Resource Type	Description
	TLS encryption (Automated)		transmission of data over the internet using standard encryption technology. Encryption should be set with the latest version of TLS. App service allows TLS 1.2 by default, which is the recommended TLS level by industry standards, such as PCI DSS. App service currently allows the web app to set TLS versions 1.0, 1.1 and 1.2. It is highly recommended to use the latest TLS 1.2 version for web app secure connections.

Policy sets

The Cloud Configuration Governance policy sets and its policies are listed for your reference.

Cloud Configuration Governance policy sets

Policy Set Name	Policies	Description
AWS 1.4.0	Ensure IAM Users Receive Permissions Only Through Groups (Automated) Ensure no Network ACLs allow ingress from 0.0.0.0/0 to remote server administration ports (Automated) Ensure IAM password policy requires minimum length of 14 or greater (Automated) Ensure a support role has been created to manage incidents with AWS Support (Automated) Eliminate use of the 'root' user for administrative and daily tasks (Automated) Ensure hardware MFA is enabled for the 'root' user account (Automated) Ensure that Object-level logging for write events is enabled for S3 bucket (Automated) Ensure access keys are rotated every 90 days or less (automated)	Amazon Web Services Foundations Benchmark (Automated) v1.4.0 - 05-28-2021

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance policy sets (continued)

Policy Set Name	Policies	Description
	Ensure there is only one active access key available for any single IAM user (automated) Ensure IAM policies that allow full "*:*" administrative privileges are not attached (Automated) Ensure that all the expired SSL/TLS certificates stored in AWS IAM are removed (Automated) Ensure IAM password policy prevents password reuse (Automated) Ensure that IAM Access analyzer is enabled for all regions (Automated) Ensure multi-factor authentication (MFA) is enabled for all IAM users that have a console password (Automated) Ensure credentials unused for 45 days or greater are disabled (Automated) Ensure that S3 Buckets are configured with 'Block public access (bucket settings)' Ensure no security groups allow ingress from 0.0.0.0/0 to remote server administration ports (Automated) Ensure no 'root' user account access key exists (Automated) Ensure MFA is enabled for the 'root' user account (Automated) Ensure that Object-level logging for read events is enabled for S3 bucket (Automated)	
Azure 1.4.0	Ensure that logging for Azure KeyVault is 'Enabled' (Automated) Ensure Web App is using the latest version of TLS encryption (Automated) Ensure that Vulnerability Assessment (VA) is enabled on a SQL server by setting a Storage Account (Automated) Ensure Storage logging is Enabled for Blob Service for 'Read', 'Write', and 'Delete' requests (Automated) Ensure Diagnostic Setting captures appropriate categories (Automated) Ensure that VA setting 'Send scan reports to' is configured for a SQL server (Automated) Ensure That 'Notify about alerts with the following severity' is Set to 'High' (Automated)	Microsoft Azure Foundations Benchmark v1.4.0 - 11-26-2021

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Configuration Governance policy sets (continued)

Policy Set Name	Policies	Description
	Ensure 'Additional email addresses' is Configured with a Security Contact Email (Automated) Ensure Web App Redirects All HTTP traffic to HTTPS in Azure App Service (Automated) Ensure That 'All users with the following roles' is set to 'Owner' (Automated) Ensure Storage Logging is Enabled for Table Service for 'Read', 'Write', and 'Delete' Requests (Automated) Ensure Storage Logging is Enabled for Queue Service for 'Read', 'Write', and 'Delete' requests (Automated) Ensure that 'Unattached disks' are encrypted with CMK (Automated) Ensure that VA setting 'Periodic recurring scans' to 'on' for each SQL server (Automated) Ensure SQL server's TDE protector is encrypted with Customermanaged key (Automated) Ensure the key vault is recoverable (Automated)	

Cloud Action Library

Use the ready-to-use actions and subflows of the ServiceNow® Cloud Action Library application to interact with the cloud resources of the organization. ServiceNow® ITOM Cloud Accelerate features, such as Cloud Configuration Governance, use these actions to operate.

Important: Starting with the Brazil release, Cloud Action Library is being prepared for future deprecation. It will be hidden and no longer installed on new instances but will continue to be supported. For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

Explore	Set up	Develop
Explore Cloud Action Library	• Install Cloud Action Library • Set up Cloud Action Library for Amazon Web Services (AWS) • Set up Cloud Action Library for Microsoft Azure	• Developer training • Developer documentation • Cloud Action Library reference

Troubleshoot and get help

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Ask or answer questions in the Cloud Action Library forum on the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Cloud Action Library

Understand the typical Cloud Action Library workflows and high-level tasks.

To use the Cloud Action Library actions and workflows to interact with the cloud resources, perform the following tasks:

1. Install Cloud Action Library. For more information, see Install Cloud Action Library.
2. To use Cloud Action Library for interactions with Amazon Web Services (AWS) resources, see Set up Cloud Action Library for AWS.
3. To use Cloud Action Library for interactions with Microsoft Azure resources, see Set up Cloud Action Library for Microsoft Azure.

See Plugins or applications installed with ITOM Cloud Accelerate to find the list of plugins you activate with Cloud Action Library.

Install Cloud Action Library

You can install the Cloud Action Library application (sn.itom.cal) if you have the admin role. The application installs related ServiceNow® Store applications and plugins if they are not already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Cloud Action Library application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

Several plugins are installed with the Cloud Action Library application. For more information, see Cloud Action Library.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Cloud Action Library application (sn.itom.cal) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.

4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

Set up Cloud Action Library for AWS

Set up access to the Amazon Web Services (AWS) cloud accounts in Cloud Action Library to enable interaction between the application and the cloud.

Before you begin

Role required: admin

About this task

For the purposes of configuring access for AWS accounts, the following terms are used:

Trusting accounts

The trusting accounts don't have permanent AWS credentials. You configure the trust relationship for IAM roles in these accounts to rely on other accounts for access.

Trusted accounts

The trusted accounts are used by the trusting accounts for access. The ServiceNow UI refers to the trusted accounts as accessor accounts.

Use any one of the following methods to configure access to the AWS accounts:

- Configure the permanent credentials in the ServiceNow AI Platform to connect with the standalone AWS accounts (discreet accounts). The Cloud Service Account [cmdb_ci_cloud_service-account] table contains information on the service account type, such as management or member account and access credentials.
- Configure the member accounts to rely on the management account for access. In this case, configure the permanent credentials of the management account in the ServiceNow AI Platform.
- Configure the accounts to rely on a trusted account for access (lateral access within the same AWS organization or across different AWS organizations). In this case, configure the permanent credentials of the trusted account in the ServiceNow AI Platform.

Note: Cloud Action Library doesn't use a MID Server-based assume role setup to access the trusting accounts.

Procedure

1. Create credentials for the AWS service accounts.

   a. Navigate to **Connection & Credentials** > **Credentials**.

   b. Select **New**, and then select **AWS Credentials**.

   c. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Credentials form

Field	Description
Name	Unique and descriptive name for the AWS credentials.
Active	Option to use the credential.
Access Key ID	The access key ID that you generated on the AWS management console.
Secret Access Key	The secret access key that you generated on the AWS management console.

d. Select **Save**.

2. Select the **sn_itom_cal.Aws_Creds_Alias** credential alias or create a credential alias.

a. Unlock the credential alias.

b. Search for a credential alias.

c. Select **New**.

d. On the form, fill in the fields.

Connection & Credential Alias form

Field	Description
Name	Unique name of the alias.
Type	Credential alias type. Select **Credential** from the Type drop-down list.

e. Select **Submit**.

3. Set the **Authentication Algorithm** field to **AWS Authenticator**.
4. Select **Submit**.
5. Set up an AWS service account.

a. Navigate to **Cloud Provisioning and Governance** > **Service accounts**.

b. Select **New**.

c. On the form, fill in the fields.

Cloud Service Account form

Field	Description
Name	Unique name of the service account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Account ID	12-digit user account number. Expand the list under the account name on the AWS Management Console to view the number. **Important:** In the **Account ID** field, remove the hyphen characters (-) from the number.
Discovery credentials	The credentials needed for ServiceNow applications to access the cloud account. You can configure the discovery credentials at a later stage, while configuring access to the AWS accounts. ▪ If you are setting up an independent service account or a management account, select its AWS credential. ▪ To use other AWS accounts to access this account, leave the field empty. For example, you don't have to specify the AWS credentials for accounts assuming Identity and Access Management (IAM) roles or member accounts using their management account for access.
Datacenter URL	URL of the datacenter. Leave this field empty. **Note:** For GovCloud, use the URL https://ec2.us-gov-west-1.amazonaws.com.
Datacenter type	Type of the datacenter where the account is hosted. Select **AWS Datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.
Parent account	Name of the management account that represents the AWS organization to which this member account belongs. This field appears when you select AWS Datacenter. If the account doesn't belong to any AWS organization, leave this field empty.
Is master account	Management account flag. This check box appears when you select AWS Datacenter from the Datacenter Type drop-down. Select the check box to associate the AWS service account with the management account. Select this check box only for accounts that you have previously configured as management accounts with some member accounts belonging to them. For more information on AWS Organization, see AWS documentation.
Accessor account	Name of the trusted account. Configure this field only for accounts that don't use permanent AWS credentials and rely on IAM roles for access.

d. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. **Optional:** To use a management account to scan the member accounts of the AWS organization, create an assume role configuration for the management account.

   a. **Optional:** If you don't want to use the OrganizationAccountAccessRole to access the member account, Configure trusting account for Cloud Action Library.
   For more information, see Configure the trusting account for Cloud Configuration Governance and Cloud Action Library.

   b. **Optional:** Repeat step 6.a for each member account that needs to be scanned through the management account without using the OrganizationAccountAccessRole.

   c. **Optional:** If you want to use the OrganizationAccountAccessRole to access the member account, create an assume role configuration for the management account.
   For more information, see Create an assume role configuration.

7. **Optional:** To use a trusted account to scan the trusting account, configure trusting account for Cloud Action Library.
   The trusted account and trusting account can be in different AWS organizations or they can be in the same AWS organization.

   a. Configure trusting account for Cloud Configuration Governance.
   For more information, see Configure the trusting account for Cloud Configuration Governance and Cloud Action Library.

   b. Repeat step 7.a for each trusting account that needs to be scanned through the trusted account.

Set up Cloud Action Library for Microsoft Azure

Set up access to the Microsoft Azure cloud accounts in Cloud Action Library to enable interaction between the application and the cloud.

Before you begin

Role required: admin

Procedure

1. Create Microsoft Azure credentials that can access the Microsoft Azure account.
   For more information, see Create Azure cloud credentials.
2. Create an Integration Hub connection for Microsoft Azure.

   a. Navigate to **Process Automation** > **Flow Designer**, and select the **Integrations** tab.

   b. On the **Azure Connection** tile, select **Add Connection**.

   c. On the form, fill in the fields.

Scan Schedule form

Field	Description
Connection name	Name of the connection.
OAuth client ID	Microsoft Azure application ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
OAuth client secret	Microsoft Azure application key.
OAuth token URL	Tenant ID of the Microsoft Azure account. Replace `{tenant}` with the Microsoft Azure tenant ID.

d. Select **Create and Get OAuth Token**.

Domain separation and Cloud Action Library

Domain separation is not supported in Cloud Action Library . Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

Related topics

Domain separation for service providers

Cloud Action Library reference

Reference topics provide additional information about the Cloud Action Library application.

AWS Enable EC2 Instance Monitoring action

Action that turns on monitoring for the selected Amazon Web Services (AWS) Virtual Machine (VM) instance.

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *`ec2:MonitorInstances`* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Instances

Data type: *`String`*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS instance ID. Enter the comma-separated list of AWS instance IDs.

Example: i-1234567890abcdef0, i-1234567761abcdek1

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Location

Data type: *String*

Name of the datacenter.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Gov Cloud

Data type: *True/False*

Not supported.

Outputs

You can use these outputs as inputs for other actions.

Is Access Denied

Data type: *True/False*

Server authorization status of the call.

- *True*: The server has authorized the call
- *False*: The server didn't authorize the call

Error Message

Data type: *String*

Error message returned for the failed call.

Result

Data type: *True/False*

Status of the action.

- *True*: Successful
- *False*: Failed

Error Code

Data type: *String*

Client error code returned for the failed call.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Get IAM Login Profile action

Action that retrieves the user name for the specified Amazon Web Services (AWS) Identity Access Management (IAM) user.

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *iam:GetLoginProfile* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Location

Data type: *String*

Name of the datacenter.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Gov Cloud

Data type: *True/False*

Not supported.

IAM Account

Data type: *String*

AWS IAM user name.

Outputs

You can use these outputs as inputs for other actions.

Result

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data type: *True/False*

Status of the action.

- *True*: Successful
- *False*: Failed

AWS List Images action

Action that describes the specified images (AMIs, AKIs, and ARIs) available to you or all of the images available to you.

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *ec2:DescribeImages* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Location

Data type: *String*

Name of the datacenter.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Version

Data type: *String*

Version of the AWS API.

Limit

Data type: *Integer*

Number of pages to return per call.

Next Token

Data type: *String*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Starting page token.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Filters

Data type: *Object*

A filter name and value pair that is used to return a specific list of results from the describe operation. Filters can be used to match a set of resources based on criteria such as tags, attributes, or IDs.

Outputs

You can use these outputs as inputs for other actions.

Result

Data type: *Object*

A complex object containing information about the images.

AWS Put S3 Bucket Encryption action

Action that configures the default encryption and Amazon Web Services (AWS) S3 Bucket Key for an existing bucket.

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *s3:PutEncryptionConfiguration* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Bucket Name

Data type: *String*

Name of the AWS S3 bucket.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Location

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data type: *String*

Name of the datacenter.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Gov Cloud

Data type: *True/False*

Not supported.

Outputs

You can use these outputs as inputs for other actions.

Is Access Denied

Data type: *String*

Server authorization status of the call.

- *True*: The server has authorized the call
- *False*: The server didn't authorize the call

Bucket Region

Data type: *String*

Datacenter where the specified AWS S3 bucket is hosted.

Error Code

Data type: *String*

Client error code returned for the failed call.

Result

Data type: *True/False*

Status of the action.

- *True*: Successful
- *False*: Failed

Redirected

Data type: *String*

If the AWS S3 bucket is not available in the specified datacenter, the *Bucket Region* variable specifies the datacenter of the AWS S3 bucket.

Error Message

Data type: *String*

Error message returned for the failed call.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Turn Off EC2 Instance action

Action that stops the specified Amazon Web Services (AWS) Virtual Machine (VM).

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *ec2:StopInstances* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Instances

Data type: *String*

AWS instance ID. Enter the comma-separated list of AWS instance IDs.

Example: i-1234567890abcdef0, i-1234567761abcdek1

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Location

Data type: *String*

Name of the datacenter.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Gov Cloud

Data type: *True/False* Not supported.

Outputs

You can use these outputs as inputs for other actions.

Is Access Denied

Data type: *True/False*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Server authorization status of the call.

- *True*: The server has authorized the call
- *False*: The server didn't authorize the call

Error Message

Data type: *String*

Error message returned for the failed call.

Result

Data type: *True/False*

Status of the action.

- *True*: Successful
- *False*: Failed

Error code

Data type: *String*

Client error code returned for the failed call.

Azure Turn Off Virtual Machine action

Action that stops the specified Microsoft Azure Virtual Machine (VM).

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *Microsoft.Compute/virtualMachines/*/powerOff* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Location

Data type: *String*

Name of the datacenter.

Use MID

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Gov Cloud

Data type: *True/False*

Not supported.

VM Identifier

Data type: *String*

URI of the VM resource.

Outputs

You can use these outputs as inputs for other actions.

Is Access Denied

Data type: *True/False*

Server authorization status of the call.

- *True*: The server has authorized the call
- *False*: The server didn't authorize the call

Error Message

Data type: *String*

Error message returned for the failed call.

Result

Data type: *True/False*

Status of the action.

- *True*: Successful
- *False*: Failed

Error Code

Data type: *String*

Client error code returned for the failed call.

CAL - AWS List S3 Buckets action

Action that returns a list of all the Amazon Web Services (AWS) buckets owned by the authenticated sender of the request.

Roles and availability

Subscription requirements

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *s3:ListBucket* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Next Token

Data type: *String*

Starting page token.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Limit

Data type: *Integer*

Number of pages to return per call.

Outputs

You can use these outputs as inputs for other actions.

Result

Data type: *Object*

A complex object containing information about the images owned by the authenticated sender of the request.

CAL - AWS List VMs action

Action that returns a list of all Amazon Web Services (AWS) Virtual Machines (VM) that belong to the specified logical datacenter.

Roles and availability

Subscription requirements

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *ec2:DescribeInstances* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Location

Data type: *String*

Name of the datacenter.

Use Mid

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Version

Data type: *String*

Version of the AWS API.

Limit

Data type: *Integer*

Number of pages to return per call.

Next Token

Data type: *String*

Starting page token.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Filters

Data type: *Object*

A filter name and value pair that is used to return a specific list of results from the describe operation. Filters can be used to match a set of resources based on criteria such as tags, attributes, or IDs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Outputs

You can use these outputs as inputs for other actions.

Result

Data type: *Object*

A complex object containing information about the VMs hosted on the specified logical datacenter.

CAL - Azure List VMs action

Action that returns the list of Microsoft Azure Virtual Machines (VM) that belong to the specified logical datacenter.

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *Microsoft.ResourceGraph/resources* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Location

Data type: *String*

Name of the datacenter.

Version

Data type: *String*

Version of the Microsoft Azure API.

Credential Alias

Data type: *Record*

Credential alias for the Microsoft Azure credential.

Limit

Data type: *Integer*

Number of pages to return per call.

Next Token

Data type: *String*

Starting page token.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Filters

Data type: *Object*

A filter name and value pair that is used to return a specific list of results from the describe operation. Filters can be used to match a set of resources based on criteria such as tags, attributes, or IDs.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

Account Id

Data type: *String*

Microsoft Azure account ID.

Outputs

You can use these outputs as inputs for other actions.

Result

Data type: *Object*

A complex object containing information about the VMs hosted on the specified logical datacenter.

CAL - Azure VM IP Metric action

Action that returns the IP details of the specified Microsoft Azure Virtual Machine (VM).

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *Microsoft.ResourceGraph/resources* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Credential Alias

Data type: *Record*

Credential alias for the Microsoft Azure credential.

Version

Data type: *String*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Version of the Microsoft Azure API.

VM

Data type: *Array.Object*

URI of the VM resource.

Outputs

You can use these outputs as inputs for other actions.

Result

Data type: *Object*

A complex object containing the IP details of the specified Microsoft Azure VM.

CAL - Azure VM Monitoring Metric Data action

Action that returns monitoring information for the specified Microsoft Azure Virtual Machine (VM).

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *Microsoft.Compute/virtualMachines/{vmName}/instanceView* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Credential Alias

Data type: *Record*

Credential alias for the Microsoft Azure credential.

Version

Data type: *String*

Version of the Microsoft Azure API.

VM

Data type: *Object.Array*

URI of the VM resource.

Outputs

You can use these outputs as inputs for other actions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

Data type: *String*

The result variable contains the monitoring information for the specified Microsoft Azure VM.

CAL - Integer Stream action

Action that returns an array of integers in the specified range.

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

You don't need any cloud permission to execute this action.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

start

Data type: *Integer*

An integer specifying at which position to start.

end

Data type: *Integer*

An integer specifying at which position to stop. The returned integer stream doesn't include the end integer.

step

Data type: *Integer*

An integer specifying the increment interval for the integer stream.

Outputs

You can use these outputs as inputs for other actions.

Integer Stream

Data type: *Array.Integer*

An array of integers.

Disable AWS IAM User action

Action that deletes the password for the specified IAM user, which terminates the user's ability to access Amazon Web Services (AWS) services through the AWS Management Console.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles and availability

Subscription requirements

To use this action in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This action requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this action, the caller must have the *iam:DeleteLoginProfile* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Location

Data type: *String*

Name of the datacenter.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Gov Cloud

Data type: *True/False*

Not supported.

IAM Account

Data type: *String*

AWS IAM user name.

Outputs

You can use these outputs as inputs for other actions.

Is Access Denied

Data type: *String*

Server authorization status of the call.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- *True*: The server has authorized the call
- *False*: The server didn't authorize the call

error_code

Data type: *String*

Client error code returned for the failed call.

result

Data type: *True/False*

Status of the action.

- *True*: Successful
- *False*: Failed

error_message

Data type: *String*

Error message returned for the failed call.

CAL - AWS Generate Credential Report subflow

Subflow that lists all users in your account and the status of their various credentials, including passwords, access keys, and MFA devices.

Roles and availability

Subscription requirements

To use this subflow in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This subflow requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this subflow, the caller must have the *Iam:GetCredentialReport* and *Iam:GenerateCredentialReport* cloud permission.

Inputs

Provide a value for each input that your subflow needs. To add dynamic values, you can also select data pills using the pill picker.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Use Mid

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

Mid Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data type: *Record*

MID Server for making the outbound calls.

Api Version

Data type: *String*

Version of the AWS API.

Outputs

You can use these outputs as inputs for other subflows.

Error

Data type: *String*

Description of the error returned by the AWS endpoint.

credential_report_list

Data type: *Array.Object*

Complex object containing all fields of the credential report.

CAL - AWS S3 Get Encryption subflow

Subflow that returns the default encryption configuration for an Amazon Web Services (AWS) S3 bucket.

Roles and availability

Subscription requirements

To use this subflow in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This subflow requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this subflow, the caller must have the *s3:GetEncryptionConfiguration* cloud permission.

Inputs

Provide a value for each input that your subflow needs. To add dynamic values, you can also select data pills using the pill picker.

Bucket Name

Data type: *String*

Name of the AWS S3 bucket.

Credential Alias

Data type: *String*

Credential alias for the AWS credential.

Use Mid

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

Data type: *Record*

MID Server for making the outbound calls.

Outputs

You can use these outputs as inputs for other subflows.

Is Access Denied

Data type: *True/False*

Server authorization status of the call.

- *True*: The server has authorized the call
- *False*: The server didn't authorize the call

Is Encrypted

Data type: *True/False*

Encryption status of the AWS resource.

Bucket Region

Data type: *String*

Datacenter of the AWS S3 bucket.

KMS Masterkey Id

Data type: *String*

Customer AWS Key Management Service (KMS) key ID used for the default encryption.

Encryption Algorithm

Data type: *String*

Server-side encryption algorithm used for the default encryption.

Bucket Key Enabled

Data type: *String*

Specifies whether Amazon S3 uses an S3 Bucket Key with server-side encryption using KMS (SSE-KMS) for new objects in the bucket.

Error Code

Data type: *String*

Client error code returned for the failed call.

CAL - Resolve Credential Alias subflow

Subflow that returns the credentials alias for the mentioned credential.

Roles and availability

Subscription requirements

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To use this subflow in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This subflow requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

You don't need any cloud permission to execute this action.

Inputs

Provide a value for each input that your subflow needs. To add dynamic values, you can also select data pills using the pill picker.

Service Account

Data type: *Record*

Service account for which a suitable credential alias is sought.

Outputs

You can use these outputs as inputs for other subflows.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Error

Data type: *String*

If the underlying Cloud Action Library (CAL) action fails to return a credential alias, it returns an error message that explains the reason behind the failure. The Error variable contains the error message returned by the underlying CAL action.

CAL - Resolve MID subflow

Subflow that returns a suitable MID Server based on the specified MID Server selection criteria.

Roles and availability

Subscription requirements

To use this subflow in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This subflow requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

You don't need any cloud permission to execute this action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Inputs

Provide a value for each input that your subflow needs. To add dynamic values, you can also select data pills using the pill picker.

Service Account

Data type: *Record*

Service account for which a suitable MID Server is sought.

Outputs

You can use these outputs as inputs for other subflows.

Mid Server

Data type: *Record*

Resolved MID Server.

Error

Data type: *String*

Error encountered during identification of the MID Server.

CAL - AWS S3 Get ACL subflow

Subflow that retrieves the Access Control List (ACL) details for the specified Amazon Web Services (AWS) S3 bucket.

Roles and availability

Subscription requirements

To use this subflow in custom flows, you must obtain an Integration Hub Enterprise subscription or an App Engine subscription. For more information, see Request Integration Hub.

Role requirements

This subflow requires roles granted by delegated development or assigned to the user. For more information, see User access to Flow Designer.

Cloud permission

To execute this subflow, the caller must have the *s3:GetBucketAcl* cloud permission.

Inputs

Provide a value for each input that your action needs. To add dynamic values, you can also select data pills using the pill picker.

Bucket Name

Data type: *String*

Name of the AWS S3 bucket.

Use MID

Data type: *True/False*

Selection to indicate whether to use a MID Server to make the outbound calls.

MID Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data type: *Record*

MID Server for making the outbound calls.

Credential Alias

Data type: *Record*

Credential alias for the AWS credential.

Outputs

You can use these outputs as inputs for other actions.

Is Access Denied

Data type: *True/False*

Server authorization status of the call.

- *True*: The server has authorized the call
- *False*: The server didn't authorize the call

Bucket Region

Data type: *String*

Datacenter where the specified AWS S3 bucket is hosted.

Error Code

Data type: *String*

Client error code returned for the failed call.

Owner write

Data type: *True/False*

The bucket owner can write into the bucket.

Owner full control

Data type: *True/False*

The owner has all the permissions for the bucket.

Owner listing

Data type: *True/False*

The bucket owner can list the contents of the bucket.

Public write

Data type: *True/False*

Anyone (public access) can write into the bucket.

Public listing

Data type: *True/False*

Anyone (public access) can list the contents of the bucket.

Public full control

Data type: *True/False*

Everyone (public access) has all the permissions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Auth Users Write

Data type: *True/False*

Authenticated user groups (users with an AWS account) can write into the bucket.

Auth Users Full Control

Data type: *True/False*

Authenticated user groups (users with an AWS account) have all the permissions for the bucket.

Owner Read ACL

Data type: *True/False*

Owner of the bucket can read the bucket ACL.

Owner Write ACL

Data type: *True/False*

Owner of the bucket can write or update the bucket ACL.

Public Read ACL

Data type: *True/False*

Anyone (public access) can read the bucket ACL.

Public Write ACL

Data type: *True/False*

Anyone (public access) can write or update the bucket ACL.

Auth Users Read ACL

Data type: *True/False*

Authenticated user groups (users with an AWS account) can read the bucket ACL.

Auth Users Write ACL

Data type: *True/False*

Authenticated user groups (users with an AWS account) can write or update the bucket ACL.

Auth Users Listing

Data type: *True/False*

Authenticated user group (users with an AWS account) can list the contents of the bucket.

Owner ID

Data type: *String*

Canonical user ID of the AWS account.

Owner Display Name

Data type: *String*

Display name of the resource owner.

For federated login, the subflow doesn't return any Owner Display Name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Minimal cloud permission profiles for Cloud Action Library actions and subflows

You need appropriate cloud permissions to execute the Cloud Action Library actions and subflows. Edit the cloud permissions profile JSON to suit the needs of your organization.

Amazon Web Services (AWS) profile with minimal permissions

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Sid": "VisualEditor0",
  "Effect": "Allow",
  "Action": [
   "iam:GenerateCredentialReport",
   "s3:GetEncryptionConfiguration",
   "ec2:DescribeInstances",
   "s3:ListBucketVersions",
   "ec2:DescribeRegions",
   "s3:ListBucket",
   "iam:GetCredentialReport",
   "iam:DeleteLoginProfile",
   "ec2:MonitorInstances",
   "iam:GetLoginProfile",
   "ec2:DescribeImages",
   "s3:PutEncryptionConfiguration",
   "ec2:StopInstances",
   "s3:GetBucketAcl"
  ],
  "Resource": "*"
 }]
}
```

Microsoft Azure profile with minimal permissions

```
{
 "properties": {
  "roleName": "CCGAzureMinimalPermission",
  "description": "Grants access to scan compute resources from the Azure subscription",
  "assignableScopes": [
   "/subscriptions/${subscription_id}"
  ],
  "permissions": [{
   "actions": [
    "Microsoft.ResourceGraph/resources/read",
    "Microsoft.Compute/virtualMachines/instanceView/read",
    "Microsoft.Compute/virtualMachines/*/powerOff"
   ],
   "notActions": [],
   "dataActions": [],
   "notDataActions": []
  }]
 }
}
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM Cloud Accelerate reference

Reference topics provide additional information for configuring and using ITOM Cloud Accelerate.

Plugins or applications installed with ITOM Cloud Accelerate

Tables that list the plugins or applications that are installed with ITOM Cloud Accelerate applications. When you update your application, any newly required application dependencies are installed.

Important: Starting with the Brazil release, the following applications are being prepared for future deprecation. They will be hidden and no longer installed on new instances but will continue to be supported:

- Cloud Services Catalog
- Cloud Services Catalog Content Pack
- Cloud Configuration Governance
- Cloud Configuration Governance Content Pack
- Cloud Action Library
- Cloud Services Catalog Terraform Connector
- Cloud Config Management
- Cloud Provisioning and Governance

For details, see the Deprecation Process [KB0867184] article in the Now Support Knowledge Base.

Cloud Services Catalog

The CSC Content Pack plugin (sn_itom_csc_cp) provides the required predefined content for the Cloud Services Catalog application. This content includes governance policies and actions to deliver an end-to-end governed deployment workflow.

Plugins or applications installed with Cloud Services Catalog (sn_itom_csc)

Plugin or application name	Description
Cloud Services Catalog Terraform Connector (sn_cmp_terraform)	Integrates Terraform Open Source, Terraform Enterprise, and Terraform Cloud with ITOM Cloud Provisioning and Governance (CPG), enabling you to provision and manage services in multiple clouds like AWS, Azure, GCP, and VMware. The Terraform templates are codified as cloud template catalogs that you can use to provision cloud resources.
Cloud Discovery Workspace (sn_cloud_ops_ws)	Offers a comprehensive solution to manage the cloud operations of your organization.
Cloud Provisioning and Governance (com.snc.cloud.mgmt)	Provides a single interface to access cloud resources, publish cloud offerings to a catalog, and manage the use of those resources.
Guided Setup (sn_ads_setup)	Enables you to step through the initial configuration of any product or application that provides a Guided Setup module. Guided Setup

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Cloud Services Catalog (sn_itom_csc) (continued)

Plugin or application name	Description
	assists you with planning the roll-out of a product and performing the basic configuration necessary to go live.
ITOM/OT SU Licensing (sn_itom_licensing)	Calculates and displays how many subscriptions for ITOM applications your organization purchases and consumes.

Plugins or applications installed with CSC Content Pack (sn_itom_csc_cp)

Plugin or application name	Description
Cloud Services Catalog (sn_itom_csc)	Provides a simplified interface to automate cloud management and access cloud resources. The application enables you to publish cloud offerings to a catalog and manage the usage and life cycle of those resources.
ServiceNow IntegrationHub Action Step - SSH (com.glide.hub.action_step.ssh)	The SSH step executes SSH commands on an external *nix system through a ServiceNow® MID Server. The step also stores scripts and commands for the *nix systems.
Agent Client Collector Framework (sn_agent)	Manages Agent Client Collector on the ServiceNow instance and the MID Server.
CCG Content Pack (sn_itom_ccg_cp)	Provides optional predefined content for the Cloud Configuration Governance application such as policies, resource collectors, configuration collectors, and remediation actions.
Flow Designer Action Step - Email (com.glide.hub.action_step.email)	Provides the Send Email step to send an email to specified users or groups as an action in a flow.

Cloud Migration Assessment

Plugins or applications installed with Cloud Migration Assessment (sn_cloud_migration)

Plugin or application name	Description
ITOM/OT SU Licensing (sn_itom_licensing)	Calculates and displays how many subscriptions for ITOM applications your organization purchases and consumes.
Discovery (com.snc.discovery)	Finds applications and devices on your network, and then updates the CMDB with the information it finds.
Cloud Migration Base (com.sn_cloud_mig_base)	A base plugin to support the Cloud Migration Assessment application.

Cloud Configuration Governance

The CCG Content Pack plugin (sn_itom_ccg_cp) provides optional predefined content for the Cloud Configuration Governance application such as policies, resource collectors, configuration collectors, and remediation actions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Cloud Configuration Governance (sn_itom_ccg)

Plugin or application name	Description
Cloud Action Library (sn_itom_cal)	Provides predefined actions and sublows to enable interaction with the cloud resources of your organization. ServiceNow® ITOM Cloud Accelerate features such as Cloud Configuration Governance use these actions to operate.

Plugins or applications installed with CCG Content Pack (sn_itom_ccg_cp)

Plugin or application name	Description
Cloud Configuration Governance (sn_itom_ccg)	Checks the configuration settings of cloud resources in your organization against a set of policies to identify violations. After the violation is identified, you can use remediation workflows to mitigate them.

Cloud Action Library

Plugins or applications installed with Cloud Action Library (sn_itom_cal)

Plugin or application name	Description
ITOM/OT SU Licensing (sn_itom_licensing)	Calculates and displays how many subscriptions for ITOM applications your organization purchases and consumes.
Flow Designer Action Step - Payload Builder (com.glide.hub.action_step.payload)	Enable action designers to easily create name-value pairs in JSON and XML payloads using dynamic data.
ServiceNow IntegrationHub Action Step - Get Connection Info (com.glide.hub.action_step.get_connection_info)	Provides the connection and credential details from another step, such as a representational state transfer (REST) step, to other steps in your action.
ServiceNow IntegrationHub Action Step - XML Parser (com.glide.hub.action_step.xmlparser)	Enables the XML parser action step in Workflow Studio.
ServiceNow IntegrationHub Runtime (com.glide.hub.integration.runtime)	Enables execution of IntegrationHub actions and flows.
Cloud Provisioning and Governance Core (com.snc.cloud.core)	Contains resource blocks and cloud Application Programming Interface (cloud API) mappings related to Cloud API-based discovery and events processed by either response mapping or calling patterns.
Flow Designer Action Step - Script on MID (com.glide.hub.action_step.script_mid)	Provides an action step to execute custom JavaScript within a reusable action on a MID Server.
ServiceNow IntegrationHub Action Step - REST (com.glide.hub.action_step.rest)	Enables the REST action step in Workflow Studio.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Cloud Action Library (sn_itom_cal) (continued)

Plugin or application name	Description
ServiceNow IntegrationHub Action Template - Data Stream (com.glide.hub.action_type.datastream)	Send REST, SOAP, or JDBC requests from Workflow Studio to APIs that return a stream of response data larger than 10 MB, or that return paginated results. Parse stream data into a series of complex object outputs and use the data pills in other actions in a flow.

Cloud Account Management

Plugins or applications installed with Cloud Account Management (sn_cloud_acct_mgt)

Plugin or application name	Description
Discovery and Service Mapping Patterns (sn_itom_pattern)	Provides the latest versions of the Discovery and Service Mapping Patterns released on the ServiceNow Store. ServiceNow releases new patterns on the ServiceNow® Store on a monthly basis, enabling your organization to discover the latest industry-standard devices and applications.
ServiceNow IntegrationHub Runtime (com.glide.hub.integration.runtime)	Enables execution of IntegrationHub actions and flows.
ServiceNow IntegrationHub Action Step - Get Connection Info (com.glide.hub.action_step.get_connection_info)	Provides the connection and credential details from another step, such as a representational state transfer (REST) step, to other steps in your action.
ServiceNow IntegrationHub Action Template - Data Stream (com.glide.hub.action_type.datastream)	Send REST, SOAP, or JDBC requests from Workflow Studio to APIs that return a stream of response data larger than 10 MB, or that return paginated results. Parse stream data into a series of complex object outputs and use the data pills in other actions in a flow.
CMDB CI Class Models (sn_cmdb_ci_class)	Adds class models that extend the CMDB class hierarchy, including class descriptions, identification rules, identifier entries, and dependent relationships if applicable. Applications such as Discovery and Service Mapping can use these class extensions to populate configuration items (CIs) and discover various technologies and software.
Flow Designer Action Step - Payload Builder (com.glide.hub.action_step.payload)	Enable action designers to easily create name-value pairs in JSON and XML payloads using dynamic data.
ServiceNow IntegrationHub Action Step - REST (com.glide.hub.action_step.rest)	Enables the REST action step in Workflow Studio.
ServiceNow IntegrationHub Action Step - JSON Parser (com.glide.hub.action_step.jsonparser)	Identify structured data from a JSON payload without having to write a script.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Cloud Account Management (sn_cloud_acct_mgt) (continued)

Plugin or application name	Description
	Map incoming JSON content to a complex object output that you can use in other steps or actions. At runtime, values from a JSON payload populate the complex object output.

Related topics

- Plugins or applications installed with ITOM Visibility
- Plugins or applications installed with ITOM AIOps
- Plugins activated with ITOM Optimization
- Applications installed with ServiceNow Otto for ITOM

ITOM Optimization

The IT Operations Management (ITOM) Optimization application provides automation for the cloud workflows used to manage the cloud resources throughout their life cycle. It enables certified and enterprise-compliant cloud deployment, cost visibility, and other cloud management processes.

Features of ITOM Optimization

ITOM Optimization features provide an integrated, service-centric approach that maximizes business service quality, drives efficient processes, and promotes strong governance.

Cloud Provisioning and Governance

The ServiceNow® Cloud Provisioning and Governance (CPG) application serves as a unified interface for accessing cloud resources, delivering cloud offerings to a catalog, and overseeing resource usage. This application is transformed as Cloud Services Catalog application, offering refined and streamlined management of usage and life cycle of cloud resources.

ITOM Optimization licensing

The ServiceNow AI Platform® uses a licensing method where your organization is billed for using ITOM Optimization applications. The ServiceNow Product Documentation doesn't provide information on prices, packaging, or other details determined by your organization customer contract. For general information about licensing and subscriptions, see ITOM/OT SU Licensing and subscriptions.

Using guided setup to implement IT Operations Management applications

The IT Operations Management guided setup application provides a sequence of tasks that help you configure IT Operations Management applications on your ServiceNow instance. To open IT Operations Management guided setup, navigate to **Guided Setup** > **ITOM Guided Setup**. For more information about using the guided setup interface, see Using guided setup.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

- To find the list of plugins you activate with IT Operations Management, see Plugins activated with ITOM Optimization .
- For more details on Cloud Provisioning and Governance, see Cloud Provisioning and Governance.

ITOM Optimization reference

Reference topics provide additional information for configuring and using ITOM Optimization .

Plugins activated with ITOM Optimization

Table that lists the plugins that are activated with the ITOM Optimization application Cloud Provisioning and Governance. When you update your application, any newly required application dependencies are installed.

Plugins activated with Cloud Provisioning and Governance (com.snc.cloud.mgmt)

Plugin name	Description
Cloud Provisioning and Governance Core (com.snc.cloud.core)	Contains resource blocks and cloud Application Programming Interface (cloud API) mappings related to Cloud API-based discovery and events processed by either response mapping or calling patterns.
Cloud Config Management (com.snc.config.mgmt)	A Configuration Management Framework that provides API implementations to integrate with configuration management vendors.
Orchestration (com.snc.runbook_automation)	Automates often-repeated manual operations such as deployments, configurations, and information gathering. Orchestration uses workflow and MID Server technologies to enable activities to execute commands on systems inside your enterprise network.
Service Portal - Core (com.glide.service-portal)	A portal framework that enables administrators to build a mobile-friendly self-service experience for users.
ITOM Optimization Licensing (com.sn_itom_opt_licensing)	A proxy application that enables you to install all applications related to Optimization. Once you purchase the Optimization license, you can install this proxy application, which enables you to install any application that is part of Optimization license.

Related topics

Plugins or applications installed with ITOM Cloud Accelerate

Plugins or applications installed with ITOM Visibility

Plugins or applications installed with ITOM AIOps

Applications installed with ServiceNow Otto for ITOM

MID Server configuration and exception handling

You can configure a MID Server with additional configuration parameters or create a new Mid Server property to alter any default behavior.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A system property 'mid.property.probe.api.timeout.minutes' enables you to decide the amount of time mid-server should wait before throwing a timeout error while provisioning.

This Cloud Provisioning and Governance MID Server property with name and value parameters handles the timeout exception.

See MID Server properties for more information on creating a MID Server property.

ITOM Visibility

The ServiceNow® ITOM Visibility product consists of ServiceNow® Discovery, Firewall Audits and Reporting, Certificate Inventory and Management, Service Mapping, Cryptographic Asset Compliance, AI Agent Topology Mapping, ITOM Content Service, Tag Governance, Cloud Discovery Workspace, and Cloud License Estimator applications. Discovery and Service Mapping give you a unified, connected view of your entire IT network and the services that it supports. The Service Graph Connectors and CMDB 360 applications provide data to ITOM Visibility.

https://player.vimeo.com/video/1226069277?h=82e1a4c4f6&badge=0&autopause=0&player_id=0&app_id=58479%22

Who uses ITOM Visibility

ITOM Visibility enables the IT departments of enterprises and cloud companies providing platform as a service to discover their IT resources.

ServiceNow® Configuration Management Database (CMDB) is not just an operational tool, it is a strategic necessity in today's IT landscape. Maintaining an accurate and complete CMDB provides the foundation for maintaining critical services and drives multiple outcomes important to IT departments.

Data collected by ITOM Visibility provides a foundation for operation of the following business units and products of ServiceNow AI Platform:

ITOM AIOps

Use ITOM AIOps to track and maintain the health of services in your organization. ITOM AIOps gathers alerts from infrastructure events captured by third-party monitoring tools. It then uses IT-related information gathered by Discovery to map alerts to configuration items. Based on the collected information, ITOM AIOps provides dashboards showing a consolidated view of all service-impact events. You can also use ITOM AIOps to proactively analyze your IT infrastructure to spot issues and prevent service outages. Using advanced machine learning to analyze information about your IT infrastructure, the application automatically determines dynamic thresholds and identifies anomalies that may indicate potential service outages.

ITOM Optimization

ITOM Optimization gives you tools to provision private and public cloud infrastructure and services and to achieve consistent management and cost visibility. The Cloud Insights application, available in the ServiceNow Store , helps you analyze cloud asset costs, identify savings opportunities, and optimize operations.

Software Asset Management

Understand the software running in your IT environment. Software Asset Management works together with the CMDB powered up by ITOM Visibility. Use

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Software Asset Management to track configurations that impact software license consumption across your IT environments and datacenter.

Customer Service Management

Efficiently diagnose and resolve issues related to the IT infrastructure by using near real-time data supplied by ITOM Visibility.

IT Service Management

Rely on the IT infrastructure discovered by ITOM Visibility to manage and deliver services to your users. See changes and incidents created and managed by IT Service Management applications in ITOM Visibility service maps.

Strategic Portfolio Management

Use data collected by ITOM Visibility to gain a comprehensive understanding of the applications used in your organization.

Security Operations

View security incidents in the context of ITOM Visibility to find out which application services are at risk. Use this information to prioritize and resolve threats based on the impact they pose to your organization.

How you use ITOM Visibility

The Discovery feature offers a replicable and reliable method of identifying the enterprise IT infrastructure. Discovery can find computers, servers, software, printers, routers, and switches. It can also find IP-enabled devices and applications that run on them, resources from various cloud providers, and TLS certificates. This method is referred to as horizontal discovery. Connections between the devices and applications are not included in horizontal discovery.

The Service Mapping feature maps dependencies, based on a connection between devices and applications. This method is referred to as top-down mapping. The top-down mapping helps you immediately see the impact of a problematic object on the rest of the application service operation. Application service maps show infrastructure objects and connections between them. Service Mapping regenerates application service maps regularly, to keep them updated and relevant. Any faulty objects are shown along with the devices and applications they affect, providing a visual clue of the state of the application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison of horizontal discovery and top-down mapping results

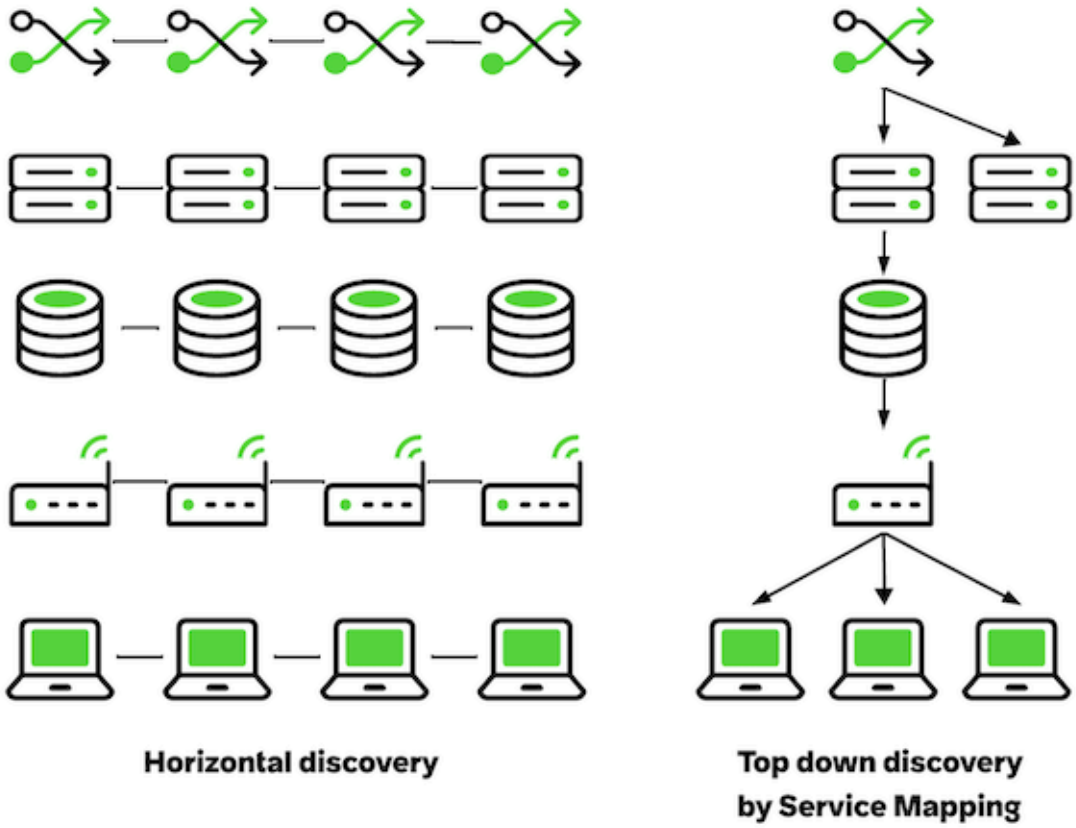

How ITOM Visibility works

Discovery can use scripts to collect and process data on a host and then update the CMDB. Scripts that explore or investigate CIs on your network are called probes. Sensors are the scripts that parse the data returned from the probes. In addition, Discovery uses discovery patterns. A pattern is a sequence of operations whose purpose is to detect attributes of devices and applications and, when used by Service Mapping, their outbound connections.

Service Mapping can deploy different methods for creating application services. Discovery patterns are the main method used by Service Mapping. However, you can also map application services using tags, and traffic connections between devices and applications. For more information, see Choose the right method for mapping application services.

ITOM Visibility licensing

The ServiceNow AI Platform® uses a licensing method where your organization is billed for using ITOM Visibility applications. Licensing options can vary based on your organization's agreement. The ServiceNow Product Documentation doesn't provide information on prices, packaging, or other details determined by your organization customer contract. For general information about licensing and subscriptions, see ITOM/OT SU Licensing and subscriptions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to know before you begin

ITOM Visibility is available with activation of the Discovery (com.snc.discovery) plugin and the Service Mapping (com.snc.service-mapping) plugin, which require an ITOM Visibility subscription. For details, see Request Discovery and Request Service Mapping. For full ITOM Visibility functionality, install the latest ITOM Visibility applications from the ServiceNow Store. ITOM Visibility installation is also available in the ITOM Product Hub for ServiceNow Otto users.For more information, see Install ITOM Visibility using ServiceNow Otto for Setup

Define users and configure credentials to enable ITOM Visibility access to applications and devices inside your organization network. For details, see Prerequisites for performing top-down discovery using Service Mapping.

Install ITOM Visibility using ServiceNow Otto for Setup

Install all required ITOM Visibility applications and plugins from the IT Operations Management Product Hub.

Before you begin
Verify the following:

- You're using the Australia Patch 3, Zurich Patch 10, or later version of the ServiceNow AI Platform.
- You have installed the ServiceNow Otto for Platform plugin (sn_genai_platform).
- You have installed the Generative AI Controller plugin (sn_generative_ai).
- You have installed the AI Skill Kit plugin (sn_skill_builder).
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.
- You have activated the ServiceNow Otto panel. For more information, see Activate the ServiceNow Otto panel standard chat.
- You have activated the following assistants:
  - ServiceNow Otto - Platform assistant
  - ServiceNow Otto for Virtual Agent
    For more information, see Assistants overview.
- You have activated Now LLM Service as a provider. For more information, see Manage AI models.
- You have set up AI Search. For more information, see AI Search readiness for AI on the ServiceNow AI Platform.
- You have installed the following ServiceNow Otto for Setup applications from the ServiceNow Store:
  - ServiceNow Otto for Setup (sn_ia)
  - ServiceNow Otto for Setup Common (sn_ia_common)
  - ServiceNow Otto for Setup Content (sn_ia_content)
  - ServiceNow Otto for Setup Core (sn_ia_now_assist)
  - ServiceNow Otto for Setup Configuration (sn_ia_config)

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Note: Applications are installed only if you have the license. If you don't have the required license for an application, it isn't installed.

Procedure

1. Navigate to **Admin** > **Admin Home** on your instance.
2. Select the IT Operations Management tile.
3. In the IT Operations Management Product Hub, select **Start setup**.

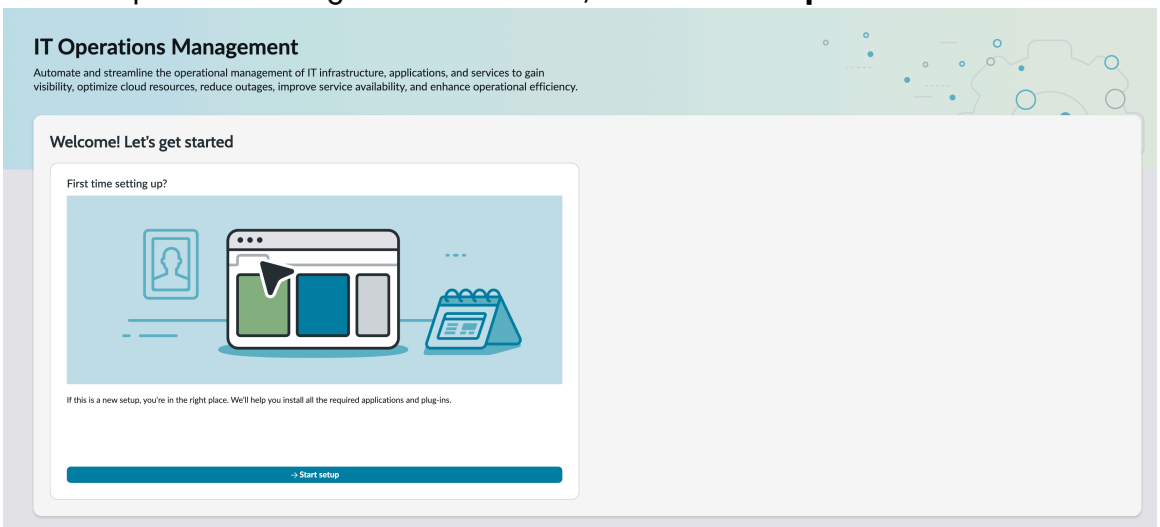

4. Select the install icon ( ) on the ITOM Visibility card.

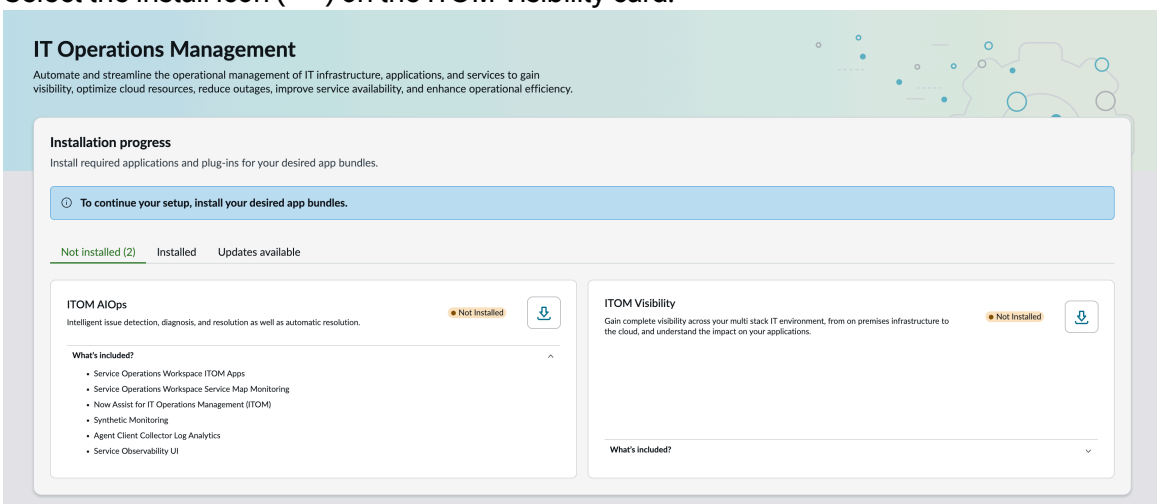

Note: Expand the What's included section to view the applications installed with ITOM Visibility.

AI in ITOM Visibility

AI features in ITOM Visibility help administrators automate service maps, diagnose discovery issues, manage certificates, and resolve incidents using generative AI and agentic workflows.

AI features in ITOM Visibility span several products. The following tables summarize all available AI features by type and product.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

Agentic workflows

Agentic workflows in ITOM Visibility use AI agents to complete investigation and management tasks autonomously or with minimal administrator input.

Agentic workflows in ITOM Visibility

Product	AI feature	Description	Use case	Resources
Discovery	Pattern diagnostic agentic workflow	Investigates missing CI attributes in pattern-based discovery. The workflow identifies the attribute gap, parses discovery logs, determines the root cause, and suggests a remediation — all from the ServiceNow Otto panel. Uses two agents: the Pattern diagnostic agent and the EF Remediation Agent.	A discovery administrator notices that a CI attribute is missing and wants to identify the root cause and a fix without manually navigating discovery logs and tables.	Pattern Diagnostic agentic workflow
Firewall Audits and Reporting	Firewall Management Task Creation agentic workflow	Uses AI to accept natural-language firewall rule requests, extract required parameters, run a risk analysis, and create a rule task. The risk level determines whether the task is created automatically or requires confirmation.	A firewall administrator wants to request a new firewall rule by describing it in plain language rather than filling out a form manually.	• Firewall rule requests using agentic workflows • Request firewall rules using agentic workflow
Service Mapping	Analyze potential impact agentic workflow	Analyzes how a change request might affect relevant servers and services. The agent verifies prerequisites, retrieves the change request, selects up to 10 affected servers, identifies matches with suggested services, and displays an impact analysis in the ServiceNow Otto panel. The analysis is also saved to the change request work notes.	An operator or change manager wants to assess the risk of a change request before approving it, without manually reviewing all affected CIs.	• Analyze potential impact agentic workflow • Assess a change request with the Analyze potential impact workflow • Activate the Service Mapping Candidate skill

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agentic workflows in ITOM Visibility (continued)

Product	AI feature	Description	Use case	Resources
				• Activate the Service Mapping Candidates Impact skill
LEAP	Ansible discovery agent	Analyzes automation opportunities and uses AI-powered semantic matching to identify relevant Ansible job templates from Ansible Automation Platform. Presents discovered templates to the LEAP administrator for step-to-job mapping.	A LEAP administrator wants to find the most relevant Ansible job templates for an automation opportunity without manually searching the Ansible catalog.	• Ansible automation integration • Ansible discovery agent
LEAP	Ansible execution agent	Automatically launches mapped Ansible job templates during incident remediation in the Service Operations Workspace. Reports execution status back to the incident responder.	An incident responder wants to trigger an Ansible automation directly from the Service Operations Workspace during active incident resolution, without switching to the Ansible platform.	• Ansible automation integration • Execute Ansible automations for incidents

AI agents

Standalone AI agents in ITOM Visibility run autonomously on a schedule or in response to a trigger, without requiring user input for each operation.

AI agents in ITOM Visibility

Product	AI feature	Description	Use case	Resources
Service Mapping	Service Mapping AI Agent	Autonomously creates service maps from ML-powered Application Service Candidates (ASCs). The agent evaluates candidates, filters noise such as monitoring clients and operating system processes, and creates the service topology in the CMDB. Created service maps are set to non-operational	A Service Mapping administrator wants to process a large volume of ML-powered candidates and create	• AI Agents for Service Mapping • Activate AI Agents for Service Mapping

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI agents in ITOM Visibility (continued)

Product	AI feature	Description	Use case	Resources
		by default for administrator review. Runs automatically every 15 minutes after activation.	service maps automatically, without reviewing each candidate individually.	
Service Mapping	Business App Mapping AI Agent	Automatically creates CSDM "Consumes::Consumed by" relationships between Business Applications and discovered Application Services by invoking the Business App Identification skill. High-confidence matches are connected automatically. Medium-confidence matches are saved to a staging table for administrator review. **Note:** The Business App Mapping AI Agent is available through the Feature Preview Program. For more information, see Feature Preview Program.	A Service Mapping administrator wants to connect discovered application services to their business application context in the CSDM without manually mapping each relationship.	• AI Agents for Service Mapping • Activate AI Agents for Service Mapping
Certificate Inventory and Management	Certificate renewal AI agent	Uses AI to automatically renew a single TLS certificate through a conversational prompt. The agent guides the user through the renewal process and provides a link to the task record when complete.	A PKI administrator wants to renew a certificate without manually navigating the renewal workflow.	Renew certificates using the AI agent
LEAP	LEAP AI agent	Uses automation opportunities created by LEAP analysis to generate artifacts — problem records, AI-enhanced knowledge base articles, or playbooks — based on user requests. The agent is turned on by default.	A LEAP administrator wants to convert a recurring incident pattern into a problem record, a structured knowledge base article, or a playbook without manually authoring each artifact.	• AI agent • Create LEAP problem records • Generate LEAP knowledge base articles • Generate LEAP playbooks

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generative AI skills

Generative AI skills in ITOM Visibility are discrete, on-demand capabilities that administrators invoke. They generate analyses, summaries, or insights for a specific discovery error set, service map, or incident pattern.

Generative AI skills in ITOM Visibility

Product	AI feature	Description	Use case	Resources
Service Mapping	Service Mapping Candidate skill	Classifies application service candidates and generates names and descriptions for the processes within them. Retrieves publisher, product, and description details for each process, then generates a name and description for the candidate. Active by default. Used by the Analyze potential impact agentic workflow.	A Service Mapping administrator wants AI to automatically name and describe application service candidates during a change impact analysis, without manually reviewing each process.	Activate the Service Mapping Candidate skill
Service Mapping	Service Mapping Candidates Impact skill	Analyzes connections and effects on servers following a change request, and generates an impact summary of affected components. Active by default. Used by the Analyze potential impact agentic workflow.	A Service Mapping administrator wants AI to automatically summarize which components are affected by a change request, without manually tracing server dependencies.	Activate the Service Mapping Candidates Impact skill
Discovery	AI Insights for discovery errors	Analyzes current Discovery errors and surfaces the highest-impact and highest-leverage errors in the Discovery Admin Workspace Diagnostics page. Each insight includes a summary of affected schedules and devices, and recommended remediation steps. Powered by ServiceNow Otto for Error Framework.	A discovery administrator wants to quickly identify which errors to resolve first to have the greatest effect on discovery health, without reviewing every error individually.	Discovery Admin Workspace Diagnostics
LEAP	AI-generated resolution steps	Uses generative AI to analyze grouped incidents and generate structured resolution steps for each automation opportunity. Resolution steps can be reviewed and modified before use.	A LEAP administrator wants to produce actionable resolution guidance for a recurring incident pattern without manually authoring steps from scratch.	Generate and modify resolution steps in LEAP

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery

ServiceNow Discovery is an automated process that continuously scans and identifies all the components within the IT infrastructure. It plays a crucial role in maintaining an accurate and up-to-date CMDB 360 with the information it finds.

https://player.vimeo.com/video/1092921849?h=0dcb800150&badge=0&autopause=0&player_id=0&app_id=58479

Discovery is available as a separate subscription from the rest of the ServiceNow AI Platform. See Request Discovery for details.

Discovery supports database encryption and full-disk encryption. See Database Encryption for details.

Get started

Introduction

- Exploring Discovery

Set up

- Request Discovery
- Discovery setup
- MID Server installation
- Create and test your credentials
- Migrate from CAPI to Patterns

Configure

- Configuring Discovery
- Discovery properties
- Create a Discovery behavior
- Create a Discovery CI classification
- Create a Discovery process classification

Use

- Schedule a horizontal discovery
- Run a Quick Discovery
- Discovery status
- Validate discovery results
- Discovery Admin Workspace
- Discovery for cloud environment
- Certificate Inventory and Management

Integrate and develop

- Activate SCCM Asset Intelligence scheduled imports
- Discovery APIs on the ServiceNow® Developer Site

Troubleshoot and get help

- Discovery monitoring and issue resolution
- Search the Known Error Portal for known error articles
- Ask or answer questions in the ITOM forum
- Contact Customer Service and Support

Exploring Discovery

Discovery finds computers, servers, printers, a variety of IP-enabled devices, and the applications that run on them. It can then update the configuration items (CIs) in your Configuration Management Database (CMDB) with the data it collects.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Horizontal discovery and top-down discovery

There are two types of discovery:

Horizontal discovery

Horizontal discovery is a technique that Discovery uses to scan your network, find computers and devices, and then populate the CMDB with the CIs it finds. Horizontal discovery creates direct relationships between CIs, such as a `runs on` relationship between an application CI and the actual computer CI that it runs on. Horizontal discovery is not aware of business services and does not create relationships between CIs based on the business service they are in.

Top-down discovery

Top-down discovery is a technique that Service Mapping uses to find and maps CIs that are part of business services, such as an email service. For example, top-down discovery can map a website business service by showing the relationships between an Apache Tomcat web server service, a Windows server, and the MSSQL database that stores the data for the business service.

Typically, Service Mapping and Discovery work together to run horizontal discovery first to find CIs, and then top-down discovery to establish the relationships between business services that you need to know.

Probes, sensors, and patterns

Discovery uses these components to find CIs:

Probes and sensors

Probes and sensors are scripts that collect and process data on a host and then update the CMDB. More specifically, probes explore or investigate CIs on your network, and sensors parse the data returned from the probes. Several probes and sensors are provided by default, but you can customize them to find different information, or you can create ones. You can also configure several parameters to control the behavior of a particular probe every time it is triggered.

Patterns

Patterns, like probes and sensors, are a series of operations that also collect data on a host, process it, and update the CMDB. Patterns differ from probes and sensors in that they are written in Neebula Discovery Language (NDL) rather than JavaScript, and they are called into action in the later stages of the horizontal discovery process. Default patterns are provided, but you can also customize or create patterns using the Pattern Designer.

Horizontal discovery phases

The phases of horizontal discovery are:

Discovery follows these phases:

Scanning

Discovery sends a probe called Shazzam to the network to see if commonly used ports are open and if these ports can respond to queries. For example, if Shazzam finds a device that responds on port 135, Discovery knows that it's a Windows server.

Classification

If Discovery finds devices or computers, it sends additional probes to find the type of device or the operating system on the device. For example, Discovery sends the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

WMI probe to a Windows machine to detect the Windows 2012 operating system. Then Discovery uses records called classifiers, which specify the trigger probe or probes that run during the next two phases. If you're using patterns, the classifier specifies a trigger probe that in turn launches a pattern.

Identification

Discovery tries to gather more information about the device and then tries to determine if a CI for the device exists in the CMDB. Discovery then uses additional probes, sensors, and identifiers to update existing CIs in the CMDB or create new ones. Identifiers, also known as identification rules, specify the attributes that the probes look at when reconciling data with the CIs in the CMDB. If you're using patterns, Discovery uses the appropriate identification rule for the CI type specified in the pattern.

Exploration

The identifier launches additional probes configured in the classifier. These probes gather additional information about the device, including the applications running on it and attributes such as memory, network cards, and drivers. Discovery then creates relationships between applications and devices and between applications. If you're using patterns, the operations in the pattern perform the exploration of the CI.

Discovery communication through MID Servers

A MID Server, which constantly queries the instance for probes to run, executes the instructions in the probe or in the pattern that the probe specifies. The MID Server then returns the results to the instance, where sensors process it. The MID Server does not retain any discovery information.

The MID Server starts all communications, using SOAP on HTTPS, which means that all communications are secure, and all communications are initiated inside the enterprise's firewall. No special firewall rules or VPNs are required.

Because Discovery is agentless, meaning that it does not require any permanent software to be installed on any computer or device to be discovered, the MID Server uses several techniques to probe devices without using agents. For example, the MID Server uses SSH to connect to a Unix or Linux computer, and then it can run a standard command, as specified in the probe, to gather information. Similarly, it uses the Simple Network Management Protocol (SNMP) to gather information from a network switch or a printer.

Types of discovery

The types of horizontal discovery that the Discovery application can perform are explained in the following table:

Type	Description
Network discovery	Run this type of discovery to find the internal IP networks within your organization. If you already know the IP address ranges in your network, it is not necessary to run network discovery.
CI discovery	Run this type of discovery to find the devices, computers, and applications on your network. This is essentially the standard type of discovery that you run most often.
Cloud discovery	Run this type of discovery to find AWS and Azure resources in your organization's cloud.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Description
Serverless discovery	Run this type of discovery to find applications on host machines without the need to discover the host first. Serverless discovery relies on patterns to explore CIs on a host.

IP service affinity

IP Service affinity saves the IP service information that is used to successfully find a device and associates it with the IP address of the device. Using this information, Discovery can target the device in subsequent runs with the accurate protocol. Discovery records the IP Service along with the IP address. Discovery can store the successful IP service information in the IP Service Affinity table [ip_service_affinity].

For example: A network device has both an SSH port and an SNMP port open. By its agentless design, Discovery tries SSH first. However, network devices should be discovered through SNMP. Discovery tries the SSH probe and it fails. This triggers the SNMP probe, which succeeds. With the association between the IP address and the IP service, subsequent discovery runs that target this IP address use SNMP first, because that is the probe that succeeded.

Related topics

- Discovery setup
- Discovery monitoring and issue resolution

Horizontal discovery process flow with probes and sensors

The horizontal discovery process passes through the four phases of discovery using probes, which gather information on the target machine. Sensors then help Discovery determine what to do with that information.

Kicking off Discovery

A user triggers horizontal discovery by configuring a discovery schedule or by launching an on-demand discovery with Discover now or Quick Discovery. The schedule specifies one or more IP addresses or range of IP addresses.

Scanning phase

1. Discovery first takes the Shazzam probe (and then port probes) and places it in a request in the External Communication Channel (ECC) queue.
2. The MID Server checks the ECC queue, retrieves the discovery request, and runs the probes against the host and discovers open ports.
3. The port probes scan common ports using several protocols, such as WMI, HTTP, SSH, and SNMP.
4. If one or more ports respond, the Shazzam probe sends information about the port back to the ECC queue through the MID Server.
5. Discovery checks the ECC queue to find out which ports responded, which identifies the type of machine. For example, if Shazzam detects that the machine is listening on port 22, Discovery treats the machine as a UNIX or Linux machine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classification phase

1. The Discovery application determines which classification probe to send to the newly discovered device by using information in the record of the port probe that successfully responded.
2. Discovery puts the classification probe into the ECC queue.
3. The MID Server checks the ECC queue, retrieves the discovery request, and runs the classification probe.
4. The classification probe retrieves additional information, such as which version of the operating system is running on a machine. This information determines the class of the CI that Discovery found. There is only one classification probe per discovered device.
5. The classification probe sends information back to the instance ECC queue through the MID Server.

The Identification phase

1. Discovery determines which classifier to use based on the class of the CI and the criteria specified in all CI classifier records. The classifier specifies which probes to use for the next two phases.
2. Discovery puts the identification trigger probe for the CI classifier into the ECC queue. For example, a UNIX machine running HP-UX would require the HP-UX classifier, which specifies that the Multi Probe-HP-UX Identity identification trigger probe. These probes use identification rules to determine whether or not to insert or update a CI in the CMDB.

   Note: The trigger probe could also be the Horizontal Pattern probe, which tells Discovery to follow the operations in the specified pattern, rather than sending out additional probes. The operations in the pattern cover both the identification and exploration phases. Discovery knows which identification rules to use based on the CI type, and Discovery makes inserts or updates to the CMDB based on these rules. Probes and sensors are not used.

3. The MID Server checks the ECC queue, retrieves the discovery request, and runs the identification trigger probe.
4. The identification probe accumulates identification data for each device and sends that data back to the instance via the MID Server.
5. Discovery uses sensors for the identifier probe to process the information.
6. Discovery performs the analysis on the CMDB using CI identifiers. Discovery can update existing CIs in the CMDB or create new ones.

The Exploration phase

1. Discovery looks at the Triggers Probes related list in the classifier to find exploration probes to run.
2. Discovery puts the exploration trigger probe into the ECC queue.
3. The MID Server checks the ECC queue, retrieves the discovery request, and runs the exploration trigger probes.
4. The probes send data back to the instance via the MID Server and sensors make updates to the CMDB, just as in the identification phase.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Patterns and horizontal discovery

Horizontal discovery process flow with patterns

Horizontal discovery with patterns has four phases, just as horizontal discovery with probes does. However, for the last two phases, Discovery triggers operations from a pattern, rather than additional sets of probes.

Scanning phase

1. Discovery first takes the Shazzam probe (and then port probes) and places it in a request in the External Communication Channel (ECC) queue.
2. The MID Server checks the ECC queue, retrieves the discovery request, and runs the probes against the host and discovers open ports.
3. The port probes scan common ports using several protocols, such as WMI, HTTP, SSH, and SNMP.
4. If one or more ports respond, the Shazzam probe sends information about the port back to the ECC queue through the MID Server.
5. Discovery checks the ECC queue to find out which ports responded, which identifies the type of machine. For example, if Shazzam detects that the machine is listening on port 22, Discovery treats the machine as a UNIX or Linux machine.

Classification phase

1. The Discovery application determines which classification probe to send to the newly discovered device by using information in the record of the port probe that successfully responded.
2. Discovery puts the classification probe into the ECC queue.
3. The MID Server checks the ECC queue, retrieves the discovery request, and runs the classification probe.
4. The classification probe retrieves additional information, such as which version of the operating system is running on a machine. This information determines the class of the CI that Discovery found. There is only one classification probe per discovered device.
5. The classification probe sends information back to the instance ECC queue through the MID Server.

The Identification and Exploration phases

Patterns unify the Identification and Exploration phases of discovery.

1. Discovery determines which classifier to use based on the class of the CI and the criteria specified in all CI classifier records. The classifier specifies the Horizontal Pattern probe, which in turn specifies which pattern to launch. The Horizontal pattern probe also contains a sensor which does the actual work of updating the CMDB.

   Note: Patterns need applicative credentials to find applications running on host machines. Make sure you have applicative credentials configured along with the credentials required to access the host machine itself.

2. The operations in the pattern specify the actions that Discovery needs to take for both the identification and exploration phases. Discovery knows which identification rules to use

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

based on the CI type in the pattern, and Discovery makes inserts or updates to the CMDB based on these rules. Only the Horizontal Discovery Sensor is used. Other probes and sensors are not used.

Discovery Admin Workspace

The Discovery Admin Workspace serves as a central location for monitoring, tracking, and completing discovery-related tasks. Experience a streamlined discovery process and greater efficiency with the integration of schedules, diagnostics, tuning, anomaly detection, and more within this single workspace.

Prerequisites

To access Discovery Admin Workspace, make sure the Discovery plugin is installed and active.

You must also have the latest ServiceNow Store version of the following applications:

- Discovery Admin Workspace
- Visibility Content
- CMDB Workspace
- (optional) Discovery and Service Mapping Patterns

Roles

The following roles have access to the Discovery Admin Workspace:

- discovery_admin
- admin

Dependencies

For a list of Discovery Admin Workspace application dependencies, see Plugins or applications installed with ITOM Visibility.

Discovery Admin Workspace Home

The Discovery Admin Workspace Home page features tools to help you identify and address the most critical discovery errors.

To access the Discovery Admin Workspace, navigate to **Workspaces** > **Discovery Admin Workspace**.

Note: The capabilities described here are available starting with Discovery Admin Workspace v1.17.0. If the *sn_disco_workspace.enable_error_framework* system property is turned off, the Home page displays the legacy experience. Specific version requirements are noted for individual features where applicable.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

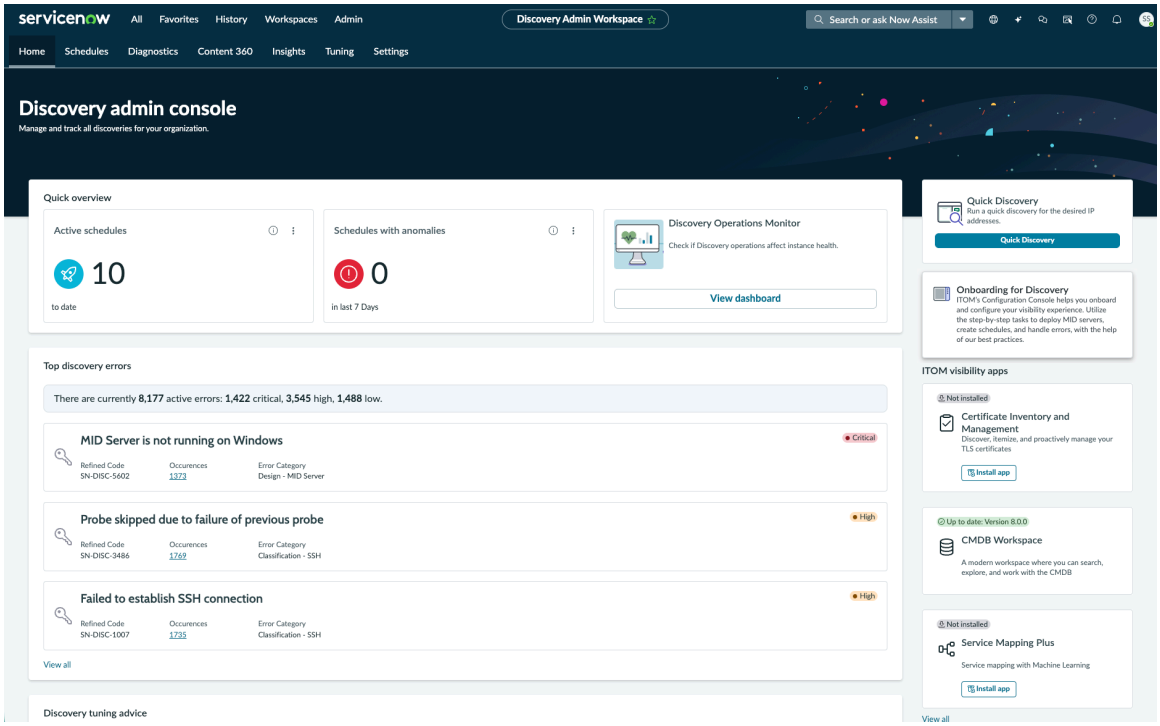

Onboarding for Discovery

The onboarding experience on the Home page changes based on your configuration progress. Before you create an IP or cloud-based Discovery schedule, the Home page displays an Onboarding for Discovery section. Select **Get started** to access the ITOM Configuration Console.

Note: Note: The ITOM Configuration Console is available in Zurich Patch 8 or later versions of the ServiceNow AI Platform. For more information, see ITOM Configuration Console for Discovery.

After you create a schedule, the Onboarding for Discovery section is replaced by an Onboard Discovery progress indicator. This indicator displays the number of completed configuration tasks. Select **Continue** to resume configuration in the ITOM Configuration Console.

After you complete configuration, the Home page displays an Onboarding for Discovery link as a persistent entry point to the ITOM Configuration Console.

Quick overview

View the status of discovery using data counts and identify any irregularities that might impact discovery. For detailed information, select the appropriate tile.

Select the **More Options** icon (⋮), then select **Refresh** to refresh the data for each visualization in this section.

Note: You can configure the time scale reflected in the displayed counts on the Settings page.

The following data counts display in the Quick overview section:

Active schedules

Displays the number of active IP-based and Cloud Discovery schedules run to date.

Schedules with anomalies

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Displays the number of Discovery schedules with anomalies detected over a certain time period. This count only displays when anomaly detection is enabled. If anomaly detection is inactive, select **Turn on** to access the Settings page.

Selecting this number redirects you to the Anomaly detection tab on the Diagnostics page, where you can view details about the type and severity of anomalies.

Discovery Operations Monitor

Monitor discovery performance across your IT environment. Select **View dashboard** to open the dashboard and view performance metrics for transactions, sensor jobs, MID Server queue load, and probe processing times. For more information, see Discovery Operations Monitor dashboard.

Top discovery errors

View the most critical Discovery errors currently active on your instance, including a summary count of active errors by severity. Each error card displays the error title, severity, refined code, occurrence count, and error category. Selecting an error card or the **Occurrences** link opens the Error Details page, where you can view the root cause, remediation steps, and individual error instances. Select **View all** to access the Diagnostics page.

Discovery tuning advice

Fine-tune Discovery and MID Server settings with automated suggestions derived from scans of your instance. These findings identify potential issues, arranged in order of their criticality. Selecting a finding redirects you to the **Tuning check** page for more details. To view all tuning checks and latest reports, select **View all** to access the Tuning page.

Quick Discovery

Discover devices in your network outside of a scheduled discovery. Enter a single IP address, multiple comma-separated IP addresses, or an IP network as shown in this example:

- IP network address - 10.0.1.0/24
- IP address range - 10.0.2.1-10.0.2.15
- IP address list - 10.0.3.176,10.0.3.222,2001::100,2600:1bdc::fffe

Select a MID Server, a MID cluster, or enable Discovery to select the MID Server automatically.

ITOM Visibility apps

Enhance the functionality of the Discovery Admin Workspace by integrating additional ITOM Visibility applications. When you select **View all**, you can view both the applications installed on your instance and others that are available for installation.

Note: Updating applications requires you to have the admin role.

While you can access details about the apps installed on your instance, information regarding pricing and packages isn't provided, as it varies based on each contract. For a general overview of licensing and subscription details, see ITOM/OT SU Licensing and subscriptions.

Learnings

Receive guidance on managing the Discovery process and resolving Discovery errors with the ITOM Visibility Knowledge Base. Access relevant articles from the ServiceNow Knowledge Base, explore blog posts, or watch instructional video tutorials for further information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

- To view the resources in this section, confirm that you have installed ITOM Content Service version 1.2.8.
- If you encounter an **Error 153: Video player configuration error** message when attempting to play a video on the **Videos** tab, update the system property *com.glide.security.referrerpolicy* to **origin-when-cross-origin**. For more information, see Enforce secure referrer policy.

Discovery Admin Workspace Schedules

The Schedules page provides a single place to monitor Discovery performance, efficiently manage schedules and statuses, and set up new IP-based or Cloud Discovery schedules.

To access the Discovery Admin Workspace Schedules page, navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.

Note: The capabilities described here are available starting with Discovery Admin Workspace v1.19.0. Specific version requirements are noted for individual features where applicable.

Key features

Quick Discovery option

Enables you to perform an IP-based discovery in just a few steps. You can select an IP range, choose the MID selection method, pick the MID Server, and run the discovery all within one interface. This option is available from any tab on the **Schedules** page.

New Discovery option

Enables you to choose how you want to discover your resources, either through IP-based or Cloud-based discovery. This option is available from any tab on the **Schedules** page.

Enables you to choose how you want to discover your resources, through IP-based, Cloud, or Certificate discovery. This option is available from any tab on the **Schedules** page.

Note: Certificate discovery is available in Discovery Admin Workspace starting with v1.20.0 and requires the Brazil, Australia, or Zurich release starting with patch 8.

Overview tab

The **Schedules overview** section enables you to make data-driven decisions through powerful visualizations.

Note: You can configure the time scale reflected in the displayed counts on the Settings page.

Select the **More Options** icon (⋮), then select **Refresh** to refresh the data for each visualization in this section.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data visualizations

Report title	Type	Description
Active schedules	Indicator	Displays the total number of IP-based and Cloud Discovery schedules that are configured aP-based and Cloud Discovery schedulesecurring basis.
Schedules completed without anomalies	Indicator	Displays the number of IP-based and Cloud Discovery schedules that were completed without anomalies in a certain time period. **Important:** This count only displays when anomaly detection is turned on. If anomaly detection is turned off, select **Turn on** to be redirected to the Settings page.
X day schedule trends	Line chart	Displays Discovery schedule trends over a certain time period. Schedules that failed to run are marked as critical anomalies. Schedules that have a low CI discovery rate are marked as major anomalies. Schedules that have a long run time or have a high error count are marked as minor anomalies. The timeline displays events that occurred during the schedule. Use the magnifying glass icon ( ) or scroll to zoom in on the timeline and hover over the indicator to expose event details. You can toggle event indicators on and off using the **Show legend** feature. Select **View discovery status** to view the status of all scheduled IP-based and Cloud discoveries. Select **View anomalies** to view anomalous schedules.
Schedules by location	Map	Displays Discovery schedules by location. Select an area of the map to view Discovery schedules in the region selected.
Schedules with anomalies	Bar chart	Displays Discovery schedules with anomalies detected in a certain time period, indicating the schedules that failed, had a low CI discovery rate, had increased error counts, or had long run times. Select the number or **View all** to view the anomalous schedules. **Important:** This count only displays when anomaly detection is turned on. For more information, see Discovery Admin Workspace Settings.
Errors by category	Donut chart	Displays the count of all Discovery errors by category. Select a category to add or remove it from the chart. Select **View details** to view a list of all errors from all discoveries.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP-based discovery tab

Provides a centralized view for monitoring and managing IP-based Discovery schedules and coverage.

Depending on your IPv6 IP Address Management (IPAM) integration, items that require attention are displayed. These items include IPs not covered by Discovery schedules and schedules that were auto-created from IPAM that require activation. Select **Review missing coverage** to access the CMDB Coverage page where you can view coverage analysis results and create Discovery schedules from IP Ranges. If you haven't enabled auto-created Discovery schedules via IPAM, select **Review in Settings** to access the Discovery Admin Workspace Settings page. When auto-created schedules are enabled, a notification displays the number of Discovery schedules that require activation. Select **View and activate schedules** to open the Auto-created schedules page. For details on activating these schedules, see Activate auto-created Discovery schedules. For more information about IPv6 IPAM integration, see IPAM Discovery integration.

Select **View Inventory** to open the IP inventory page. On the IP inventory page, you can manage the IP ranges, networks, and IPAM data that Discovery relies on. For more information, see Discovery Admin Workspace IP inventory.

Select a Discovery schedule from the table to view extensive details and update schedule field parameters. For more information, see Discovery Admin Workspace schedule details.

Cloud discovery tab

Gives you quick access to all of your cloud-based Discovery schedules.

Select a Discovery schedule from the table to view extensive details and update schedule field parameters. For more information, see Discovery Admin Workspace schedule details.

Certificate discovery

Provides a centralized view for monitoring and managing certificate Discovery schedules. Depending on the state of your schedules, items that require attention are displayed. Select **Review that list of IP-based schedules** to open the IP-based discovery tab. To resolve certificate schedule errors, select **Review schedules** to open the Diagnostics page. To review discovered certificates, select **Go to Certificate Management Workspace** to access the Certificate Inventory and Management Workspace.

Discovery Admin Workspace offers two methods for certificate discovery:

- Port scans: Collects certificates from devices that have listening ports found during your existing IP-based Discovery schedules. For more information, see Run Certificate Discovery via port scans in Discovery Admin Workspace.
- URL scans: Collects certificates from a list of URLs that you specify in a new certificate Discovery schedule. For more information, see Run Certificate Discovery via URL scans in Discovery Admin Workspace.

Select a Discovery schedule from the table to view extensive details and update schedule field parameters. For more information, see Discovery Admin Workspace schedule details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Certificate discovery is available in Discovery Admin Workspace starting with v1.20.0 and requires the Brazil, Australia, or Zurich release starting with Patch 8.

Discovery status tab

Shows the status of scheduled discoveries. You can select a Discovery status number to view detailed information. For more details, see Discovery Admin Workspace status details.

Related topics

- Create an IP-based Discovery schedule in Discovery Admin Workspace
- Create an Alibaba Cloud Discovery schedule in Discovery Admin Workspace
- Create an AWS Discovery schedule in Discovery Admin Workspace
- Create an Azure Discovery schedule in Discovery Admin Workspace
- Create a GCP Discovery schedule in Discovery Admin Workspace
- Create an IBM Discovery schedule in Discovery Admin Workspace
- Create an OCI Discovery schedule in Discovery Admin Workspace
- Create an OpenStack Discovery schedule in Discovery Admin Workspace
- Create an oVirt Discovery schedule in Discovery Admin Workspace
- Create a VMware Discovery schedule in Discovery Admin Workspace

Discovery Admin Workspace IP inventory

The IP inventory page centralizes the IP addresses, ranges, networks, and range sets that Discovery uses to scan your environment. Use it to review existing IP data and add entries to Discovery range sets.

To access the IP inventory page, navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules** > **IP-based Discovery** > **View inventory**.

You can also access the IP inventory page from the Discovery Admin Workspace Settings page.

Note: The IP inventory page is available starting with Discovery Admin Workspace v1.19.0 and requires the Brazil, Australia, Zurich Patch 8, or later release of the ServiceNow AI Platform. Specific version requirements are noted for individual features where applicable.

Key features

Each tab displays a list of IP data with the following controls:

- **New**: Create IP data. For more information, see Add IP data.
- **Add to Range Set**: Map IP data to a range set. For more information, see Add IP data to a range set.
- **Refresh** icon: Update the data
- **Export**: Download the list as Excel, CSV, JSON, or PDF, or send it by email.
- **More Actions** icon: Add or remove columns or copy the URL.

The following sections describe the IP data on each tab.

IP Address Lists

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **IP Address Lists** tab displays your IP address lists, grouped by list. A list can contain both IPv4 and IPv6 addresses, but not IPv6 subnets or networks. An IP address can belong to more than one list. Discovery removes duplicate addresses at the start of a discovery when the lists are on the same schedule.

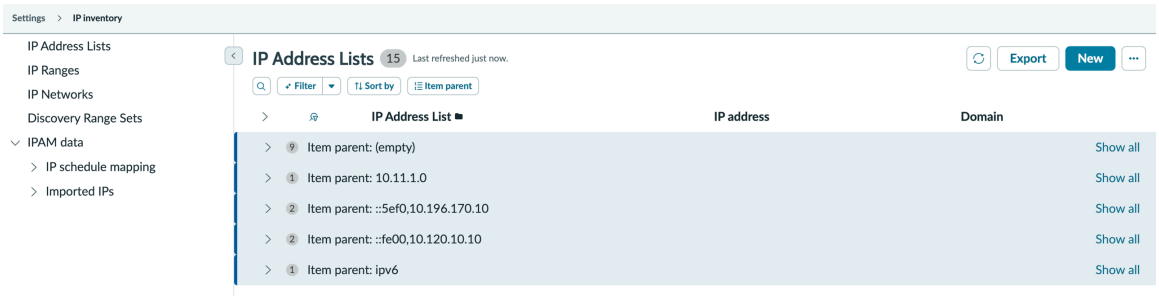

An IP address list requires a name. The Name column is hidden by default because existing records may not have a name. To add the Name column to the table, select the **More Actions** icon ( ) and select **Personalize fields**.

IP Ranges

The **IP Ranges** tab displays your IP ranges.

If an IP range belongs to a range set used in multiple Discovery schedules, the **Discovery Schedule** column displays a Multiple label. Select the **Multiple** link to view the associated records in the Discovery Schedule Range [discovery_schedule_range] table, filtered for the selected discovery range.

Note: The **Multiple** link is available starting with Discovery Admin Workspace v1.20.0.

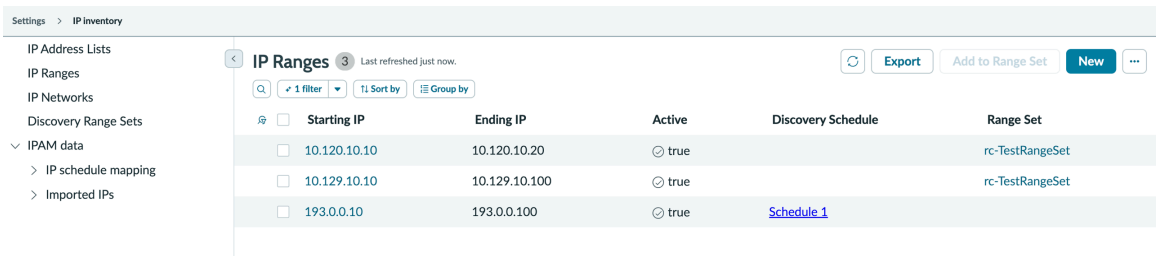

Select one or more IP ranges from the table to enable the **Add to Range Set** option. For more information, see Add IP data to a range set.

Note: An IP range can belong to only one range set.

IP Networks

The **IP Networks** tab displays your IP networks.

If a network belongs to a range set used in multiple Discovery schedules, the **Discovery Schedule** column displays a Multiple label. Select the **Multiple** link to view the associated records in the Discovery Schedule Range [discovery_schedule_range] table, filtered for the selected network range.

Note: The **Multiple** link is available starting with Discovery Admin Workspace v1.20.0.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

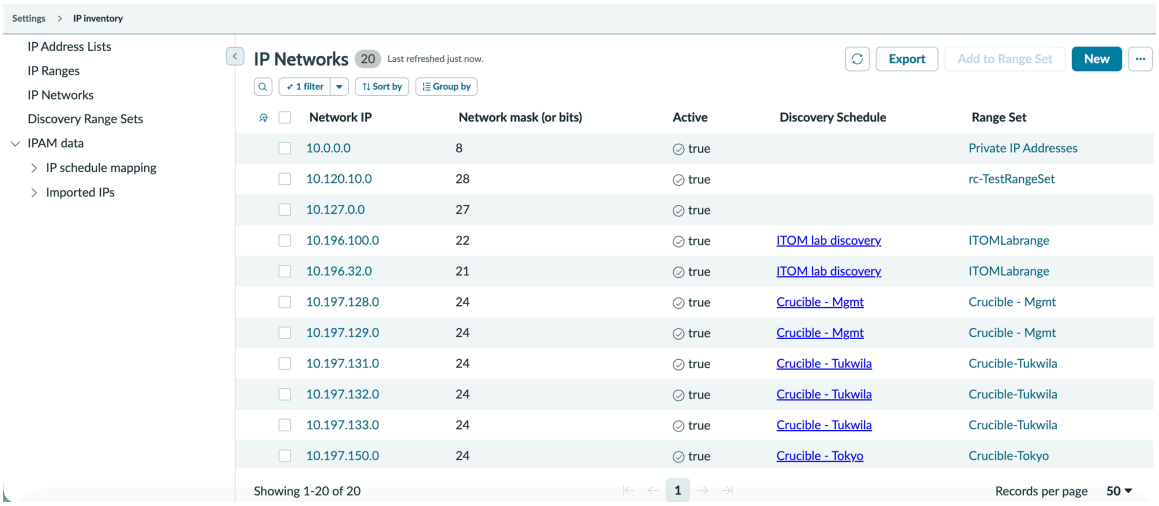

Select one or more IP networks from the table to enable the **Add to Range Set** option. For more information, see Add IP data to a range set.

Note: An IP range can belong to only one range set.

Discovery Range Sets

The **Discovery Range Sets** tab displays your discovery range sets. A range set can be associated with one or more Discovery schedules.

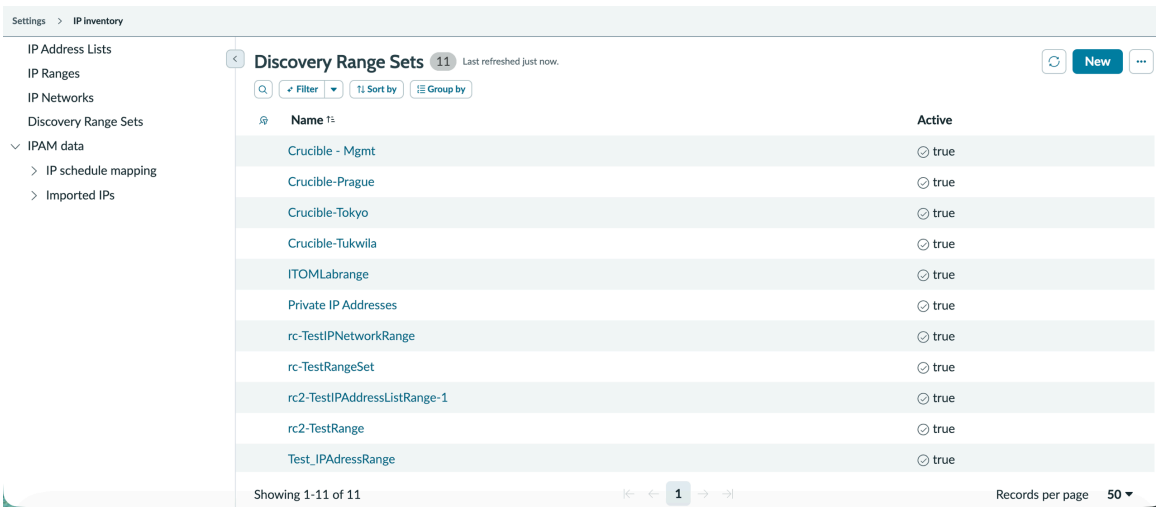

This tab doesn't include the **Add to Range Set** action, because the records are already range sets. Select **New** to create a range set.

Shazzam

The **Shazzam** tab contains the **Shazzam Summary** and **Shazzam Status** tabs, which display the results of the Shazzam probe for the IPs in your inventory. When Discovery runs, the Shazzam probe scans your network to identify responding IP addresses and open ports.

Unlike the other IP inventory tabs, which focus on IP data such as addresses, ranges, and networks, the Shazzam tabs provide scan results and discovery status information.

The **Shazzam Summary** tab provides aggregated scan results for each network range in your inventory. The **Shazzam Status** tab provides detailed results for

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

individual IP addresses that Shazzam scanned. Use these tabs to verify scan coverage and investigate IPs that didn't classify as expected. For more information about scan data, see Shazzam Insights dashboard.

The Shazzam tables support **Export** and other read-only platform actions, but don't support create, edit, or delete actions for Shazzam data.

Note:

- Shazzam data is available starting with Discovery Admin Workspace v1.20.0.
- A completed Discovery schedule run is required before Shazzam data is available.

IPAM data

The **IPAM data** tab contains the **IP schedule mapping** and **Imported IPs** sections.

The IP schedule mapping section contains the **IPv6 addresses** and **IPv4 subnets** tabs, which display the IPv6 addresses and IPv4 subnets that are mapped to auto-created Discovery schedules.

The **Imported IPs** section contains the **IP addresses** and **IP subnets** tabs. These tabs display the IP addresses and IP subnets imported from your IPAM solutions.

Warning: To avoid data corruption, don't edit or add IP data in the Imported IPs table.

Related topics

Discovery IP address configuration

Add IP data to a range set

Add IP data

Discovery Admin Workspace schedule details

Discovery Admin Workspace enables you to conveniently view, edit, and run Discovery schedules conveniently within a single interface.

To access Discovery schedule details in Discovery Admin Workspace, navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules** > **Discovery schedules**.

Note: The capabilities described here are available starting with Discovery Admin Workspace v1.8.0. Specific version requirements are noted for individual features where applicable.

After selecting a schedule name from the table, the schedule header displays key information such as Discovery details, MID Server details, and anomaly severity when anomaly detection is enabled. For more information, see Discovery Admin Workspace Settings.

The actions available in the schedule header depend on the discovery type:

- For an IP-based Discovery schedule, select **Save** to save your changes, or select **Quick ranges** to create a range of IP addresses to discover.
- For a cloud or certificate Discovery schedule, select **Edit** to update the schedule. To run any Discovery schedule on demand, select **Discover now**.

Note: Certificate discovery is available in Discovery Admin Workspace starting with v1.20.0 and requires the Brazil, Australia, or Zurich release starting with Patch 8.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key features

Overview

The **Overview** tab includes visualizations that provide detailed information about the Discovery schedule. These visualizations offer a comprehensive view of the schedule's performance and status, showing key metrics like the number of discoveries completed, success rate, and any errors encountered. Additionally, the visualizations highlight trends over time, enabling you to quickly identify patterns and potential issues. For a full list of the visualizations available on this tab, see Schedule details data visualizations.

Note: The time scale reflected on this page can be configured on the Settings page. For more information, see Discovery Admin Workspace Settings.

Select the **More Options** icon (⋮), then select **Refresh** to refresh the data for each visualization in this section.

Schedule Details

The **Schedule Details** tab provides in-depth information about the Discovery schedule and enables you to update the information directly within the interface.

Select the **More Options** icon (⋮), to access additional actions for customizing and managing the form interface.

Run History

The **Run History** tab displays the status of the last 10 scheduled discoveries, including icons for anomaly detection, when enabled. Select a Discovery status number to view detailed information. For more details, see Discovery Admin Workspace status details.

Discovery Admin Workspace status details

The Discovery Status Details page offers a summary of a discovery initiated from a schedule, detailing the devices identified, any errors encountered, and any anomalies found.

To access Discovery status details in Discovery Admin Workspace, navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules** > **Discovery status**.

Note: The capabilities described here are available starting with Discovery Admin Workspace v1.8.0. Specific version requirements are noted for individual features where applicable.

After selecting a discovery status from the table, the schedule header displays key information such as Discovery details, MID Server details, and anomaly severity.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery schedule and status headers

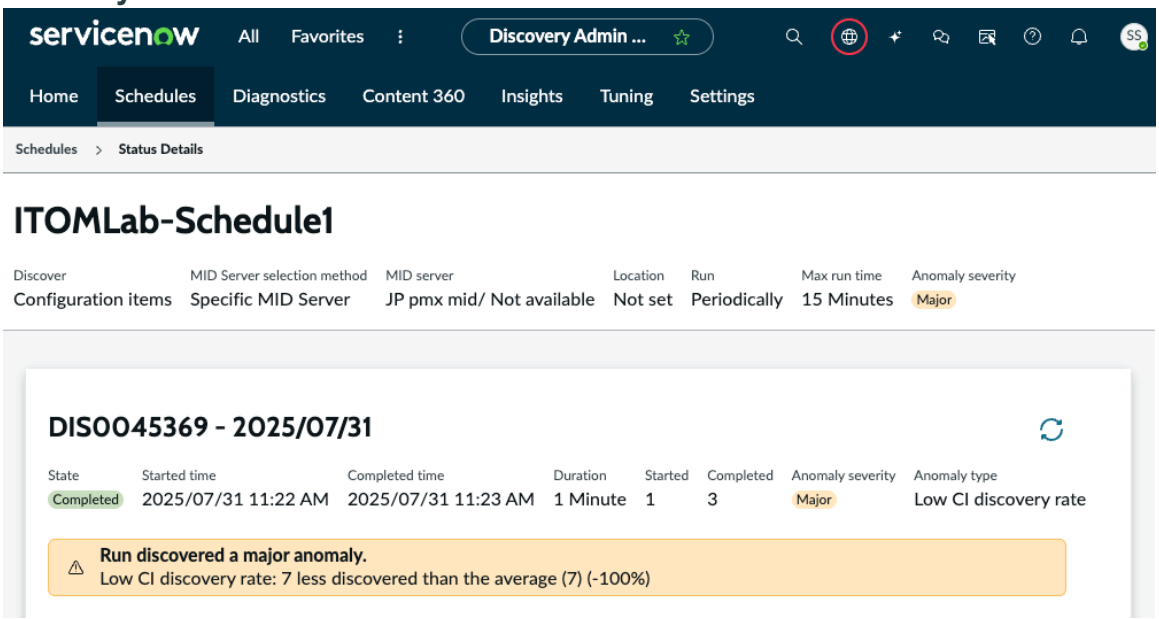

Note: Starting with v1.10.0, the schedule header displays 'Quick Discovery' when the schedule is created using the Quick Discovery feature. Additionally, if the schedule associated with a run is deleted, the header no longer displays the schedule name.

The status header shows run-related details, including start and end times, the number of probes triggered and completed, and any anomalies detected.

Important: Anomaly information only displays when anomaly detection is enabled. For more information, see Discovery Admin Workspace Settings.

If the status is Active or Starting, selecting the **Refresh** icon ( ) updates the Started and Completed values in the header in real time.

Key features

Details

The **Details** tab includes visualizations that provide detailed information about the Discovery status. The visualizations provide an overview of the schedule's performance and current status for IP-based, cloud-based, and certificate discoveries. Key metrics include devices and IPs discovered, cloud resources identified, certificates collected, and errors encountered during the run.

Note: Certificate discovery is available in Discovery Admin Workspace starting with v1.20.0 and requires the Brazil, Australia, or Zurich release starting with Patch 8.

Select the **More Options** icon ( ), then select **Refresh** to refresh the data for each visualization in this section, or export the data as a CSV or Excel file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Details tab data visualizations

Report title	Discovery type	Description
Errors	IP-based, Cloud-based, and Certificate	Displays the number of errors that were detected during the run.
Total Devices	IP-based	Displays the number of devices that were discovered during the run.
New Devices		Displays the number of new devices that were discovered during the run.
Total IPs		Displays the number of IP addresses that were discovered during the run.
Duplicate IPs		Displays the number of duplicate IP addresses that were discovered during the run.
Total Cloud Resources	Cloud-based	Displays the number of cloud resources that were discovered during the run.
Total Certificates	Certificate	Displays the number of certificates that were collected during the run.

Selecting an indicator reveals related information in a table. By default, detected errors display, sorted by priority. Each error card includes details such as the error title, severity, refined code, occurrence count, and error category. Selecting an error card or the **Occurrences** link opens the Error Details page, where you can view the root cause, remediation steps, and individual error instances. For more information, see Discovery Admin Workspace Error Details.

For IP-based schedules, the **Total Devices**, **New Devices**, and **Duplicate IPs** tables provide additional details such as the Source, Classification probe, and Scan status. Selecting the **Source** hyperlink opens a page where you can view more information about the device, apply tags, and access the Discovery Log and ECC Queue details. Selecting the **Total IPs** indicator opens the Shazzam Summary table, where you can access details such as IP addresses, IP Range, and Network Range. To learn more about a specific item, select its hyperlink in the table.

For cloud-based schedules, selecting the **Total Cloud Resources** indicator reveals a bar chart that categorizes each discovered cloud resource by its CI type. Select the **More Options** icon (⋮) to refresh the data or export the chart as a CSV, JPEG, PNG, or Excel file.

For certificate discovery schedules, selecting the **Total Certificates** indicator reveals a table of certificates collected during the run. Select a hyperlink in the table to learn more about a specific certificate. Select the **More Options** icon (⋮) to refresh the data or export it as a CSV or Excel file.

Debugging

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Debugging** tab provides information about the Discovery Log and ECC Queue.

Select the **More Options** icon (⋮), then select **Refresh** to refresh the data for each visualization in this section, or export the data as a CSV or Excel file.

By default, the **Discovery Log** table displays information such as classification failures, CMDB updates, and authentication failures. A Discovery Log record is created for each action associated with a discovery status.

Select the *ECC Queue* indicator to display entries in the **ECC Queue**. The entries show a connected flow of probe and sensor activity and the actual XML payload sent to or from an instance.

Discovery Admin Workspace Diagnostics

The Diagnostics page helps you prioritize and address errors and anomalies in IP-based and Cloud Discovery schedules.

To access the Discovery Admin Workspace Diagnostics page, navigate to **Workspaces** > **Discovery Admin Workspace** > **Diagnostics**.

Note: The capabilities described here are available starting with Discovery Admin Workspace v1.17.0. If the *sn_disco_workspace.enable_error_framework* system property is turned off, the Diagnostics page displays the legacy experience. Specific version requirements are noted for individual features where applicable.

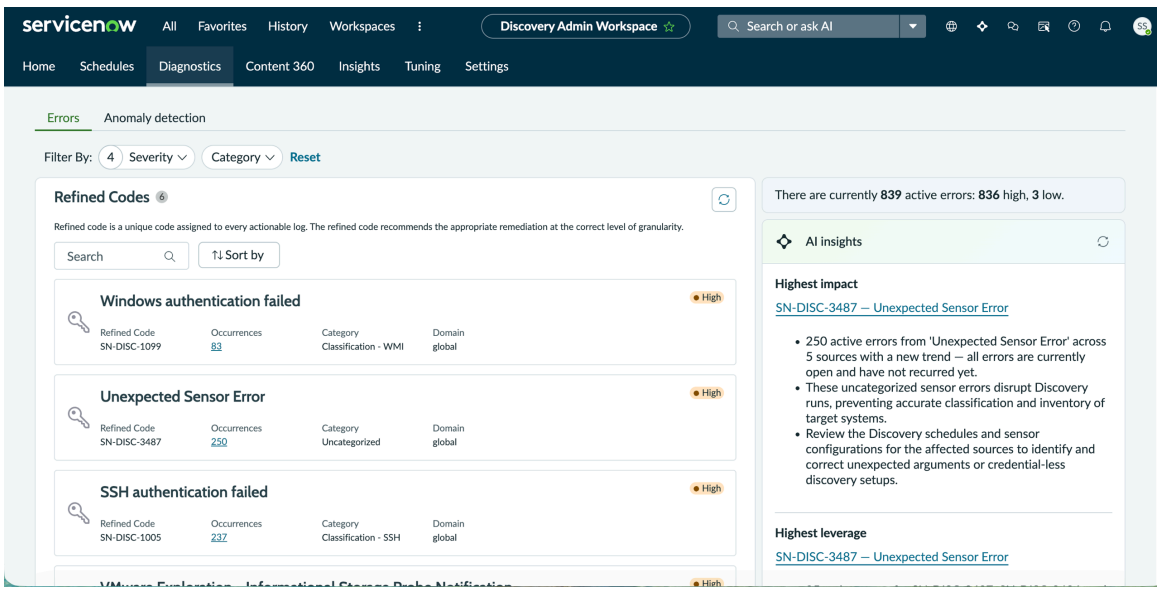

Errors

The Errors tab displays all Discovery errors grouped by refined code. Each error card shows the refined code, occurrence count, category, and domain. Use this tab to identify high-impact errors and trigger resolution actions.

Use the **Filter by** options to prioritize the error list by Severity and Category. Selecting the **Domain** drop-down enables you to filter errors based on a specific domain.

You can filter the list of errors using the **Search** (🔍) icon. Use the **Sort by** drop-down menu to organize the results by Refined Code, Category, Severity, and Occurrences in either ascending or descending order.

Refined Codes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A refined code is a unique code assigned to every actionable error generated by Discovery. The refined code identifies the error type and recommends the appropriate remediation. Refined codes use the SN-DISC format.

Select an error card or the **Occurrences** link within an error card to open the Error Details page. The page displays the root cause, remediation steps, and available actions for that refined code. For more information, see Discovery Admin Workspace Error Details.

AI Insights

The Errors tab includes an AI Insights panel powered by ServiceNow Otto for Error Framework. The panel analyzes your current Discovery errors and surfaces the following:

- **Highest-impact errors**: Errors that are affecting the most devices or schedules.
- **Highest-leverage errors**: Errors where resolving a single instance clears the largest number of occurrences.

Each insight includes a summary of affected schedules and devices, and recommended remediation steps. Select an error code link in the panel to navigate directly to that error.

Note: ServiceNow Otto for Error Framework (com.sn_ef_gen_ai) is available starting with Australia Patch 3 and requires the ServiceNow Otto for Platform (sn_genai_platform) plugin. If the plugin is not enabled, the panel displays a prompt to enable it.

For more information about how the Error Framework works, see Error Framework in Discovery Admin Workspace.

Anomaly detection

The Anomaly detection tab displays all of the Discovery schedules with anomalies detected in a certain time period. You can configure this time scale on the Settings page.

Important: This tab only displays anomalies when anomaly detection is enabled. Starting with v1.11.0, anomaly detection requires a minimum of five CIs to trigger. For more information, see the Skip Anomaly Detection for small schedules article in the Now Support Knowledge Base.

Select the **More Options** icon (⋮), then select **Refresh** to refresh the data for each visualization in this section.

Data visualizations

Report title	Type	Description
Schedules with anomalies	Indicator	Number of Discovery schedules with anomalies detected over a certain time period. Schedules are categorized as the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data visualizations (continued)

Report title	Type	Description
		• Critical anomalies: Failed to run. • Major anomalies: Have a low CI discovery rate. • Minor anomalies: Have a long run time or a high error count.
Failed schedules	Indicator	Number of Discovery schedules that failed to run.
Low CI discovery rate	Indicator	Number of Discovery schedules that have a low CI discovery rate. **Note:** The threshold for the CI discovery rate can be configured on the Settings page.
Increased error count	Indicator	Number of Discovery schedules that have an increased error count. **Note:** The threshold for the error count can be configured on the Settings page.
Long run time	Indicator	Number of Discovery schedules that had a long run time. **Note:** The threshold for the run time can be configured on the Settings page.

Selecting any of the indicators filters the table to reflect the schedules represented in the data count. For each entry in the table, you can view details around the type and severity of anomalies.

The **Actions** drop-down menu provides a direct shortcut to the Settings page for configuring anomaly detection. Additionally, the menu includes a **Clear all anomalies** option so that administrators can quickly remove records in the Discovery Schedule Anomaly [discovery_schedule_anomaly] and Discovery Schedule Anomaly Metrics [discovery_schedule_anomaly_metrics] tables.

Use the **Filter by** options to prioritize and address anomalies. Filter Discovery schedules by name using the **Schedules** drop-down menu. Additionally, you can prioritize the results with the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Severity drop-down menu, which includes options such as Critical, Major, and Minor. Use the **Anomaly type** drop-down menu to quickly view schedules with the same anomaly. Available options include Failed schedule, Low CI discovery rate, Increased error count, and Long run time.

Discovery Admin Workspace Error Details

The Error Details page displays the root cause and remediation steps for a specific Discovery error, along with the list of individual error instances associated with that error.

To access Discovery error details in Discovery Admin Workspace, navigate to **Workspaces** > **Discovery Admin Workspace** > **Diagnostics** > **Errors**.

Note: The capabilities described here are available starting with Discovery Admin Workspace v1.17.0. This page is only accessible when the *sn_disco_workspace.enable_error_framework* system property is enabled. Specific version requirements are noted for individual features where applicable.

After selecting an error card, the header displays key information including the error title, severity, refined code, and category.

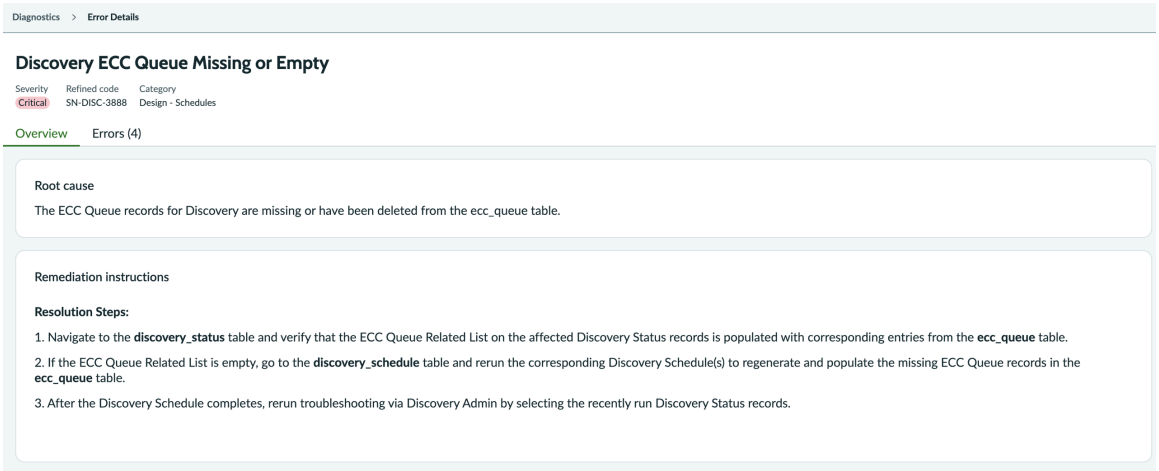

Key features

Overview

The Overview tab is the starting point for resolving an error. The **Root cause** section explains what triggered the error, and the **Remediation instructions** section provides the steps to resolve it.

Errors

The Errors tab displays the list of individual error instances associated with the refined code. Each instance represents a specific device or schedule where the error occurred.

Use the filter, **Sort by**, and **Group by** options to apply advanced filtering to the table. The table includes Fixed filters, which are pre-applied and can't be removed or modified. Select **Export** to download the list of error instances in Excel, CSV, JSON, or PDF formats.

The Errors table includes the following information:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
IP address / Error key	The IP address or unique key that identifies the affected device or target. Select the hyperlink to view additional details and perform actions on this error in a side panel.
State	The current state of the error instance. Possible values are Open, Closed, and Ignored.
Error re-opened	The number of times this error instance has re-opened after being closed.
Impacted schedule	The Discovery schedule associated with the error instance. Select the schedule name to navigate to its schedule details page.
Last occurrence	The date and time when the error was most recently detected.
Consecutive #	The number of consecutive Discovery runs in which this error has occurred.
Refined on	The date and time when the error classification was last refined.

Selecting an IP address / Error key hyperlink opens a side panel, as shown in the following image.

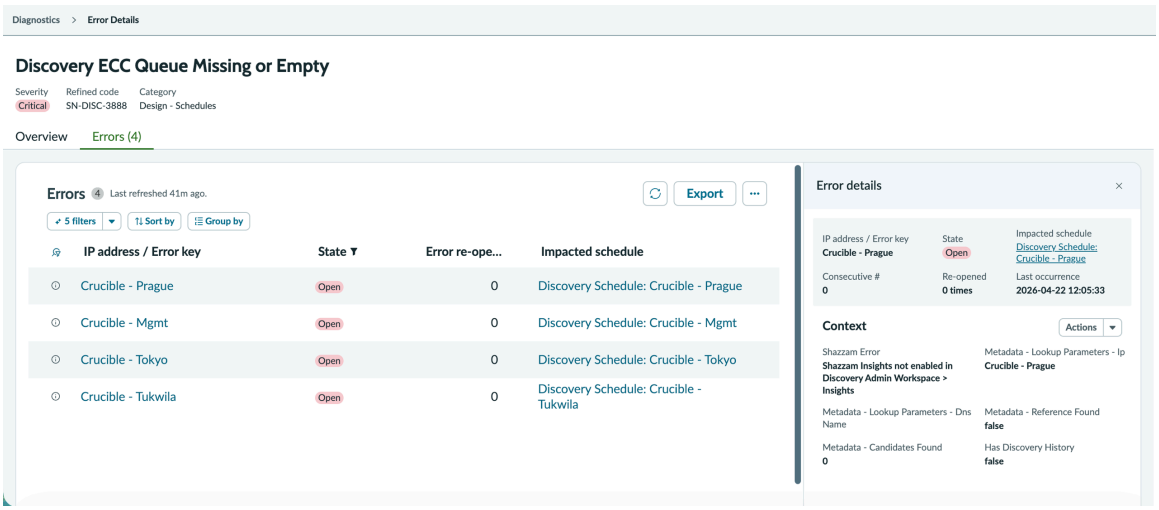

The **Context** section of the panel displays metadata collected during the Discovery run for the affected device, such as probe results, lookup parameters, and CMDB match data. Select the **Actions** drop-down to perform an action on the error instance. The fields and actions in this section vary by error type.

You can also select the ServiceNow Otto icon ( ) to open the panel chat and ask Otto to analyze the error code. Otto generates a summary and remediation steps based on the error context.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- ServiceNow Otto for Error Framework (com.sn_ef_gen_ai) is available starting with Australia Patch 3 and requires the ServiceNow Otto for Platform (sn_genai_platform) plugin.
- The Error Analysis and Remediation Workflow and Error Analysis and Remediation Agent are enabled by default with the ServiceNow Otto for Error Framework plugin.
- To interact with an agent, you must first Activate the ServiceNow Otto panel standard chat.

Error Framework in Discovery Admin Workspace

Discovery Admin Workspace uses the Error Framework to surface actionable Discovery errors. You can review errors, understand their causes, and take steps to resolve them directly from the Diagnostics page.

Benefits

The Error Framework in Discovery Admin Workspace provides the following benefits:

- A centralized view of Discovery errors grouped by category and severity.
- Contextual details for each error, including affected devices and Discovery schedule.
- Solution-focused error records that guide you to the root cause and remediation steps first.
- Actions you can trigger directly from an error record to help resolve the issue.
- AI-powered insights that prioritize your errors and surface high-impact items and quick wins.
- Aggregated error statistics to identify patterns across discovery runs.
- A domain filter to focus on errors relevant to your domain.

How it works

When Discovery runs, it logs actionable errors to the Error Framework. Each error includes a unique error code, context about where and why it occurred, and one or more resolution actions. The Diagnostics page in Discovery Admin Workspace displays these errors on the Errors tab.

After Discovery completes, the framework refines error classifications using enrichment data from the discovery run. For example, if a device initially logged an SNMP connection error but was successfully discovered through another protocol, refinement updates the error to an informational state. Errors downgraded to informational severity are filtered from the default view so you can focus on errors that need attention.

Error Framework is enabled by default. Administrators can disable it by setting the *sn_disco_workspace.enable_error_framework* system property to `false`. When Error Framework is inactive, Discovery reverts to the legacy Diagnostics experience.

Note: The Error Framework only logs actionable errors that have documented resolution steps. It doesn't replace system-level debug logging.

Error Classification

The framework classifies Discovery errors using two levels of rules:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Base rules**: ServiceNow-provided rules that map known Discovery log patterns to standardized error codes in the SN-DISC format.
- **Refinement rules**: Rules that further classify errors based on additional context gathered after the discovery run completes. Refinement can change an error's code or downgrade its severity. Default refinement rules are provided for known Discovery scenarios.

The framework processes errors through base classification first, then applies refinement. As part of this process, each error record is enriched with data from related Discovery sources, including CMDB records, credential affinities, IP address details, and Discovery schedule configuration. This context is available on the error record and reduces the need to navigate between multiple tables during investigation. The refined error code is what appears on the Errors tab of the Diagnostics page. Errors downgraded to informational severity are filtered from the default view, so you can focus on errors that need attention.

Error severity levels

Each error is assigned a severity level based on its impact on discovery. The following table describes the severity levels.

Severity	Description
Critical	The error prevents discovery from completing for the affected device or schedule.
High	The error significantly affects discovery results but doesn't stop the run.
Medium	The error may affect some discovery results. Review and resolve when possible.
Low	The error has minimal impact. Address when convenient.
Info	Informational only. No action is required.

Error states

Errors move through the following states during their life cycle:

- **Open**: The error is active and hasn't been resolved.
- **Closed**: The error has been resolved.
- **Ignored**: The error has been acknowledged and suppressed from active views.

Note: Closed and ignored error records are retained for seven days. Error instance records are retained for 90 days.

Error actions

Each error code can have one or more actions associated with it. Actions are available from the error record and can include the following types:

- Changing the state of an error instance to Open, Ignored, or Resolved.
- Adding the affected IP address to the global exclusion list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Retrying discovery for the affected device or schedule.
- Displaying a message with remediation steps to resolve the error.

Note: Available actions vary by error type.

For more information, see Discovery Admin Workspace Error Details.

AI insights

The Diagnostics page includes an AI Insights panel powered by ServiceNow Otto for Error Framework. The panel analyzes your current Discovery errors and generates a prioritized summary to help you focus on the most important issues first. Insights are scoped to your active domain and update automatically when you change filters on the Errors tab.

The panel surfaces the following:

- Highest-impact errors that are affecting the most devices or schedules.
- Quick wins, which are errors that can be resolved with minimal effort.

You can also use ServiceNow Otto agents to analyze a specific error code and receive a summary with remediation steps on demand.

Note:

- ServiceNow Otto for Error Framework (com.sn_ef_gen_ai) is available in Australia Patch 3 and requires the ServiceNow Otto for Platform (sn_genai_platform) plugin. If the plugin is not enabled, the panel displays a prompt to enable it.
- The Error Analysis and Remediation Workflow and Error Analysis and Remediation Agent are enabled by default with the ServiceNow Otto for Error Framework plugin.
- To interact with an agent, you must first Activate the ServiceNow Otto panel standard chat.

Roles

The discovery_admin role is required to access Discovery Admin Workspace and interact with errors.

Requirements

The following are required to use Error Framework features in Discovery Admin Workspace:

- The ServiceNow AI Platform must be running the Brazil release, or the Australia release starting with Patch 3.
- Discovery Admin Workspace must be installed, starting with v1.17.0.

Discovery Admin Workspace Content 360

The Content 360 page enables you to discover applications and evaluate application suggestions based on machine learning or crowd-sourced resources. Then, create configuration items with a single click.

Content 360 uses data visualizations to offer extensive visibility of your infrastructure's applications and the resources available to create application configuration items (CIs). It also

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

provides access to a list of classifiers to add to discovery. Refresh the page to view the most up-to-date visualizations.

To access the Discovery Admin Workspace Content 360 page, navigate to **Workspaces** > *Discovery Admin Workspace* > **Content 360**.

Key features

AI-suggested application fingerprints

View a list of suggested applications with fingerprints that you can add to the discovery process through the AI-suggested application fingerprints feature. You can then select a group name to view extensive details about an application.

- Modify certain details such as the discovery pattern name or suggested CMDB class.
- View additional information about process samples and related servers.
- Choose to ignore viewed suggestions the next time Discovery runs.
- Add suggestions to the Configuration Management Database (CMDB) the next time Discovery runs.

The AI-suggested application fingerprints widget doesn't refresh automatically. To verify you're viewing the most up-to-date information, refresh the page manually.

For more information about discovery based on application fingerprints, see Discover applications based on fingerprints and Discover applications using application fingerprints or ITOM Content Service.

Application CIs discovered

Monitor information about application CIs discovered based on application fingerprints or ITOM Content Service suggestions through the Discovery Admin Workspace. A pie report shows the total number of user-defined application CIs discovered. A bar report presents the top five application CIs discovered, created within the last two days, and grouped by class.

ITOM Content Service

ITOM Content Service collects unrecognized application fingerprints from different environments. AI librarians analyze these fingerprints, creating dedicated classifiers that are updated and shipped to you weekly. The latest version of Discovery Admin Workspace enables you to refresh the content service suggestions manually before reviewing them. As new IT products are rapidly released to the market, this feature grants you rapidly updating visibility and provides the following benefits:

- It reviews the products that you’re currently using every week and delivers new application fingerprints.
- It identifies irrelevant processes that aren’t suitable candidates for CIs.
- Based on AI, it helps you identify running processes fingerprints and Simple Network Management Protocol (SNMP) Object Identifiers (OIDs) and create the CIs you need for monitoring your infrastructure.
- Discovers a higher number of products by using AI capabilities that cluster and classify running application processes.

On the latest version of Discovery Admin Workspace, starting from Xanadu instances, the Content service suggestions widget refreshes automatically and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

displays the most up-to-date information. This is indicated by the real time icon ( ) that appears next to the widget name.

To learn more about ITOM Content Service, see ITOM Content Service.

Sharing data on *ITOM Content Service*

The Continuance Delivery System (CDS) is a ServiceNow® technology. It synchronizes data between the Data Services instance—which stores and distributes data—and customer instances, using only processes and SNMP OIDs data. Using the ITOM Content Service feature through the CDS, you can opt in to share application fingerprint process data. This data can be tagged and classified by AI librarians, enabling ITOM Content Service to deliver new CIs every week. Your data is shared with your ServiceNow® instance in a safe environment. No personal information is revealed, and all information is kept anonymous. Participation is optional and you can opt out at any time.

Starting with ITOM Content Service version 1.6.2 and Discovery Admin Workspace version 1.11.0, you can also opt in to share Service Info data of mapped candidates

For more information, see Share data on ITOM Content Service.

Prerequisites

Plugins required for ITOM Content Service:

- Discovery
- Normalization Data Services

Discovery Admin Workspace Insights

The Insights page provides widgets, dashboards, and reports for IT operations management and infrastructure monitoring. Use these tools to improve operational visibility and infrastructure health.

To access the Discovery Admin Workspace Insights page, navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights**.

Information accessed through the Insights page

You can access the following information through the Insights page:

- Discovery Operations Monitor dashboard
- Oracle GLAS data collection using patterns
- Kubernetes Explorer
- Nutanix Dashboard
- MID Server dashboard
- Agent Health dashboard
- Missing compute resources from Discovery
- Shazzam Insights
- Tag Governance Insights
- Virtual Machine Explorer
- URL Discovery insights for SAM

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Operations Monitor dashboard

The Discovery Operations Monitor dashboard displays performance metrics for your Discovery environment. Use this dashboard to monitor discovery health and identify issues across your IT landscape.

To access the dashboard, navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Discovery Operations Monitor**. You can also access this dashboard from the Discovery Admin Workspace Home page.

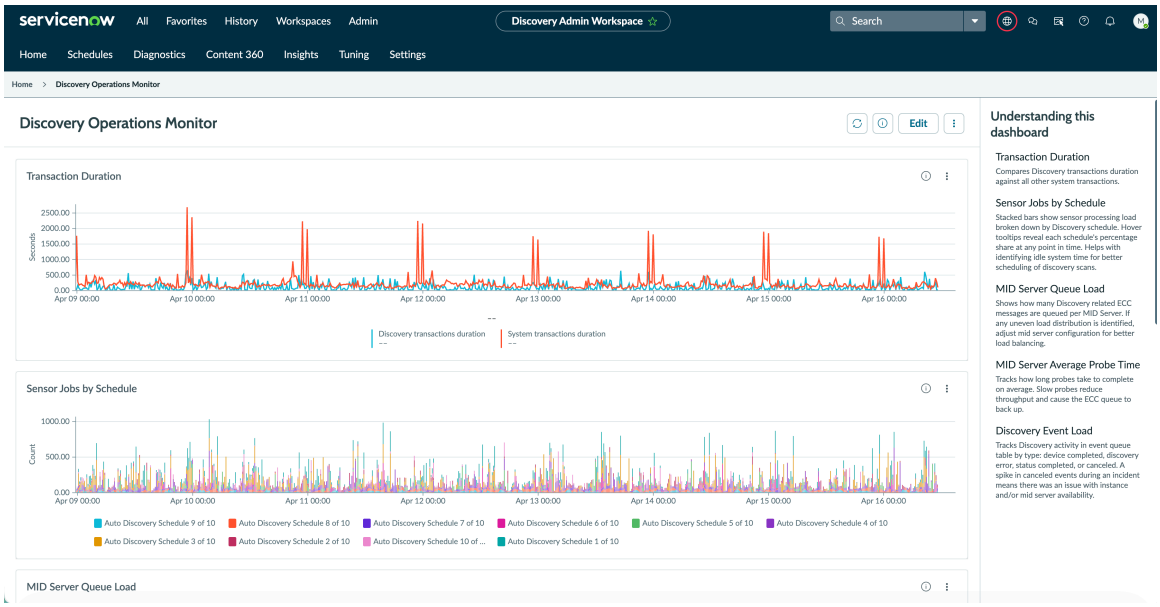

Required ServiceNow AI Platform roles

- discovery_admin
- admin

Prerequisites

Verify that you have the required setup

- The ServiceNow AI Platform must be running the Brazil, Australia, Zurich, or Yokohama release.
- Discovery Admin Workspace, starting with v1.15.0, must be installed.

Key features

The Discovery Operations Monitor dashboard displays visuals for discovery performance data.

Note: Metrics are collected for the most recently completed 15-minute window. Data may not reflect activity from the current interval until collection has run. For more information, see Configure the time interval for a visualization.

Select the info icon ( ) to view dashboard details, such as the owner and usage. Select the **More Actions** icon ( ) to duplicate, share, export, schedule, or print the dashboard. Select **Edit** to access editing mode, where you can update the data source and visualization settings. For more information about editing dashboards, see Common dashboard tasks in the in-line editor.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select the **More Options** icon (⋮) on a visualization to export it in CSV, JPEG, PNG, or Excel format.

Data visualizations

Report title	Type	Description
Transaction Duration	Line chart	Displays discovery transaction duration compared to overall system transaction duration over time. Use this chart to identify discovery-specific spikes or anomalies relative to system activity. Select a record to view the related transactions.
Sensor Jobs by Schedule	Bar chart	Displays the count of discovery sensor jobs over time, stacked by schedule. Use this chart to identify which schedules contribute most to sensor job load.
MID Server Queue Load	Line chart	Displays ECC queue volume over time for MID Servers associated with discovery. Metrics are tracked separately for each MID Server. Select a record to view the ECC queue records.
MID Server Average Probe Time	Line chart	Displays the average ECC probe processing time in seconds over time. Select a record to view the related ECC queue records.
Discovery Event Load	Bar chart	Displays the volume of discovery events over time. The chart tracks four event types: discovery.error, discovery.device.complete, discovery.complete, and discovery.canceled. Select a record to view related events.

Note: The dashboard displays data from the last 7 days by default. Data is retained for 30 days. For more information, see Configure the date range for a visualization.

Configure the time interval for a visualization

Update the data collection time interval for individual visualizations on the Discovery Operations Monitor dashboard. By default, each visualization displays data in 15-minute intervals.

Before you begin

Verify you're using Discovery Admin Workspace v1.15.0 or later.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Discovery Operations Monitor**.
2. Select **Edit**.
3. Select the visualization that you want to update.
   For example, to update the Transaction Duration visualization, select that chart while in editing mode.
4. In the Configuration panel, under Trendy by, select the edit icon ( ) on the Minute by interval end time card.
5. Select a time interval from the **Interval in minutes** drop-down list.

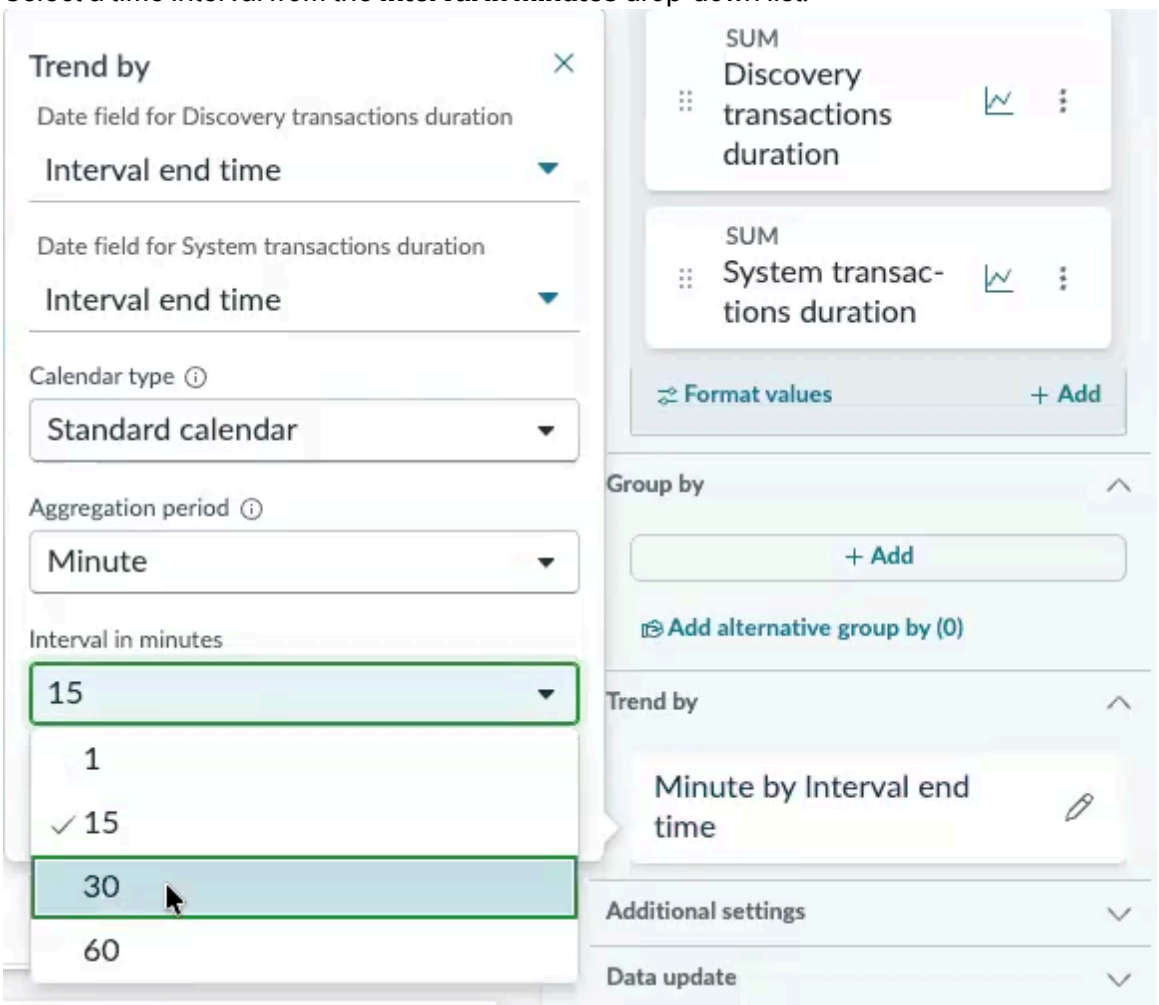

> **Warning:** Available intervals include 1, 15, 30, and 60 minutes. Although 1 minute appears in the list, it isn't a supported interval and may produce unreliable results.

6. Select **Apply**.
   The visualization updates to display data at the new interval.
7. Select **Save**.
8. Then, select **Exit editing mode**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Discovery Operations Monitor dashboard
- Configure the date range for a visualization
- Common dashboard tasks in the in-line editor

Configure the date range for a visualization

Update the data time frame for individual visualizations on the Discovery Operations Monitor dashboard. By default, each visualization displays data from the last 7 days. You can extend the time frame up to 30 days, which is the maximum data retention period for this dashboard.

Before you begin

Verify you're using Discovery Admin Workspace v1.15.0 or later.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Discovery Operations Monitor**.
2. Select **Edit**.
3. Select the visualization you want to update.
4. In the Configuration panel, under Data Sources, select the data source card for the visualization you want to update.
   For example, for the Transaction Duration visualization, select the **Discovery Transaction Load** card.

   The Edit data source page displays.
5. In the Filters section, select **Edit**.
6. Select the **Last 7 Days** drop-down list to choose a new time frame for the filter condition.

> ⚠ **Warning:** The list includes options beyond 30 days, but the dashboard retains data for a maximum of 30 days. Selecting a longer time frame doesn't display additional data.

7. Select **Apply**.
   You're returned to the Discovery Operations Monitor dashboard. The visualization updates to display data for the selected time frame.
8. Select **Exit editing mode**.

Related topics

- Discovery Operations Monitor dashboard
- Configure the time interval for a visualization
- Common dashboard tasks in the in-line editor

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes explorer dashboard

The Discovery Admin Workspace Kubernetes explorer dashboard displays a consolidated view of the Kubernetes resources discovered in your environment.

To access the dashboard, navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Kubernetes Explorer**.

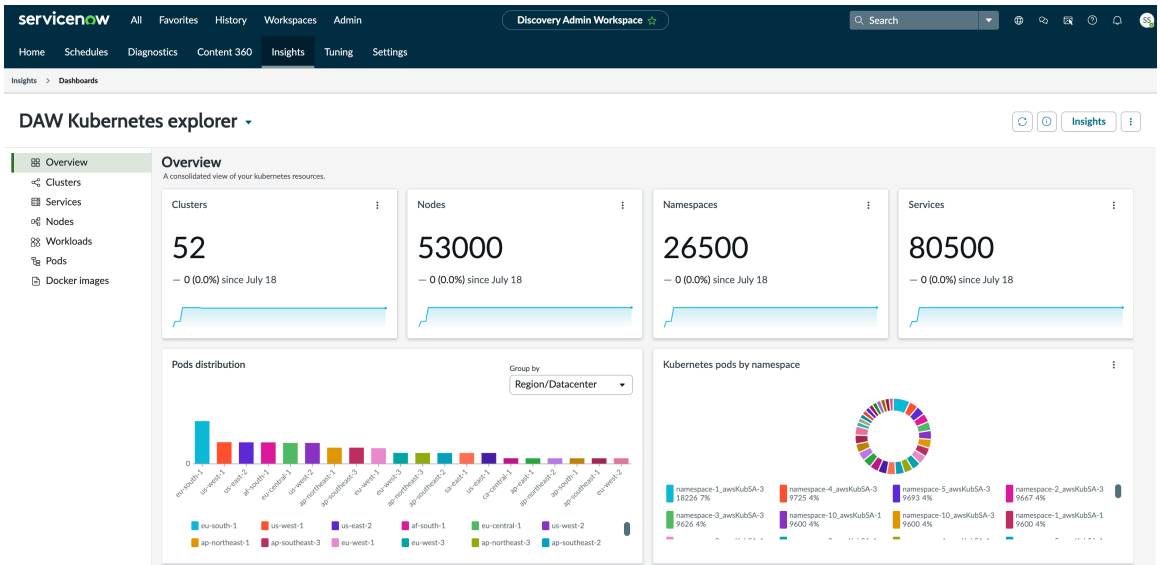

Prerequisites

Verify that you have the required setup

- The ServiceNow AI Platform must be running the Brazil, Australia, or the Zurich release starting with Patch 8.
- Discovery Admin Workspace must be installed, starting with v1.20.0.
- You must have a completed Kubernetes discovery, so that resource data is available to display. For more information, see Kubernetes discovery using patterns.

Required ServiceNow AI Platform roles

The discovery_admin role is required to view the Kubernetes explorer dashboard.

Key features

The Kubernetes Explorer dashboard enables you to make data-driven decisions through visualizations. Each tab lists the discovered resources in a table, and most tabs also include visualizations such as data counts, bar charts, and donut charts. From the dashboard, you can do the following:

- Interact with a visualization to filter the resources in the table on the same tab.
- Select a resource in a table to open its details page. The details page shows the resource properties and a Dependency View that maps the resource to its related configuration items.
- Refresh or export the data for a visualization.

Overview tab

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Overview** tab displays a consolidated view of your Kubernetes resources. To refresh or export the data for a visualization, select the **More Options** icon (⋮).

The **Overview** tab includes these data visualizations:

Overview data visualizations

Report title	Type	Source table	Descrip
Clusters	Data count and trend line	Kubernetes Cluster Analytics [sn_cow_k8s_cluster_analytics]	Display total nu and tre Kubern cluster discov since a date. Select report open th Cluste
Nodes	Data count and trend line	Kubernetes Node Analytics [sn_cow_k8s_node_analytics]	Display total nu and tre Kubern nodes a given Select report open th Nodes
Namespaces	Data count and trend line	Kubernetes Namespace Analytics [sn_cow_k8s_namespace_analytics]	Display total nu and tre Kubern names discov since a date. Select report open th Pods ta
Services	Data count and trend line	Kubernetes Service Analytics [sn_cow_k8s_service_analytics]	Display total nu and tre Kubern service discov since a date.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview data visualizations (continued)

Report title	Type	Source table	Descrip
			Select report open t Servic
Pods distribution	Bar chart	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	Displa distrib Kuberr pods, separa on-pre and cl infrast Use th Group list to g the po region data ce geo loc or cou Select in the to ope Pods t
Kubernetes pods by namespace	Donut chart	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	Displa distrib Kuberr pods groupe names Select in the to ope Pods t

Clusters tab

The **Clusters** tab lists the discovered Kubernetes clusters. The table displays each cluster with its port, provider, and datacenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

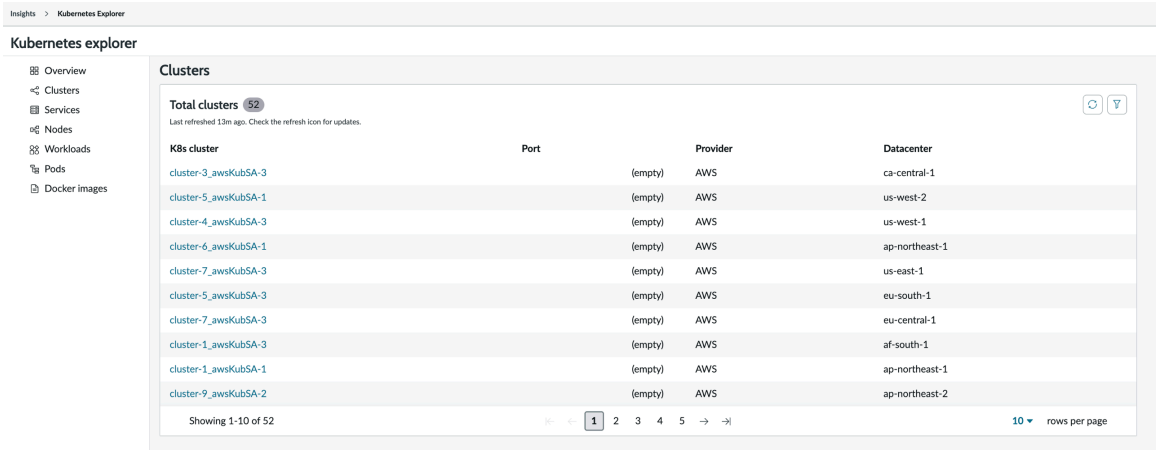

Select a cluster from the table to open its details page. The details page includes a properties panel, Dependency View, and the Key values, Nodes, and Pods tables.

Services tab

The **Services** tab lists the discovered Kubernetes services. A Total services visualization shows the number and trend of services. The Kubernetes services by platform chart groups services by platform or type. Use the Group by list to change the grouping.

The table displays each service with its service type, Kubernetes cluster, namespace, and IP address. Select a platform or type in the chart to filter the table.

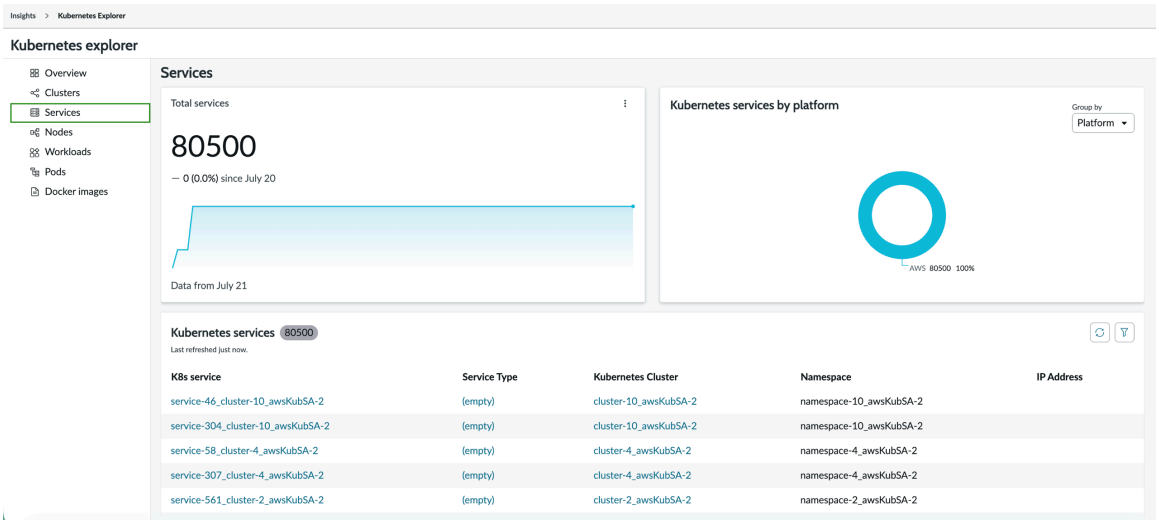

Select a service in the table to open its details page. The details page includes a properties panel, Dependency View, and the Key values table.

Nodes tab

The **Nodes** tab displays the discovered Kubernetes nodes. The Nodes by cluster and Nodes by region charts show how nodes are distributed. The table displays each node with its Kubernetes cluster and IP address.

Select a region in the Nodes by region chart, or a cluster in the Nodes by cluster chart, to filter the Total nodes table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

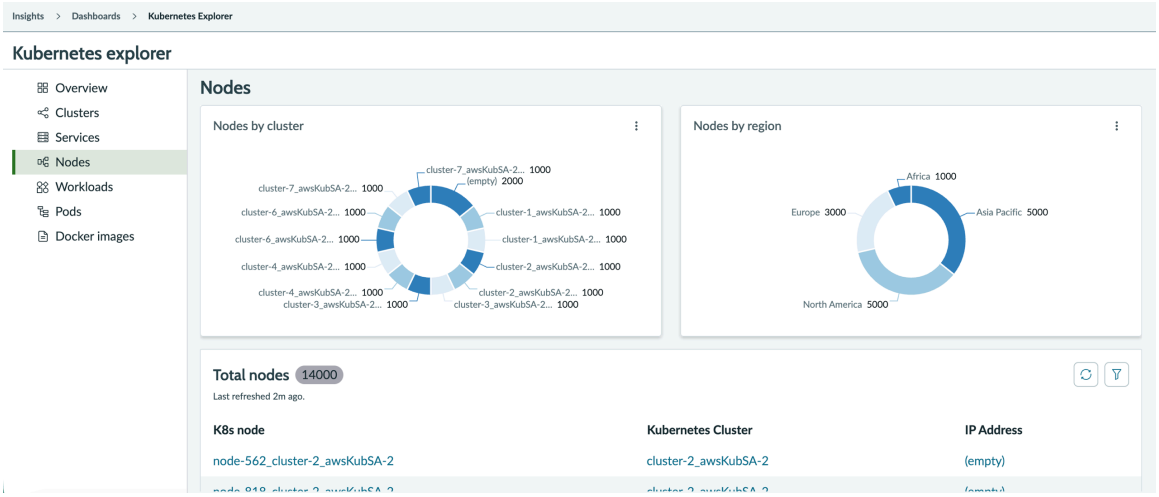

Select a node in the table to open its details page. The details page includes a properties panel, Dependency View, and the Key values and Pods tables.

Workloads tab

The **Workload** tab lists the discovered Kubernetes workloads. The Total workloads count and the Desired/available replicas mismatch count show workload totals. The Workload by class chart groups workloads by class, such as Kubernetes Deployment, DaemonSet, ReplicaSet, StatefulSet, Cronjob, and Job.

The table displays each workload with its Kubernetes cluster, namespace, and class. Select a class in the Workload by class chart to filter the table by that class.

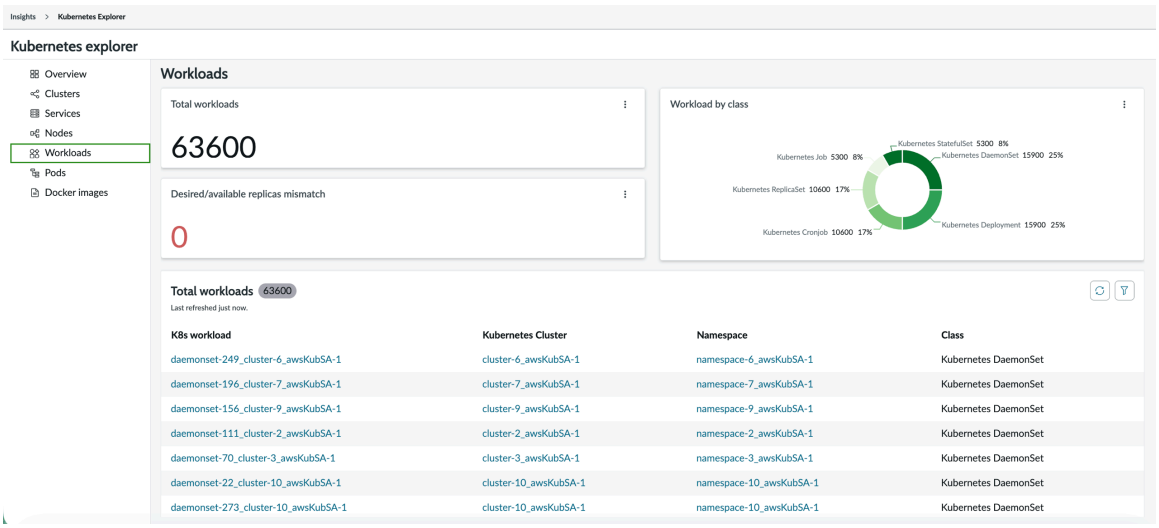

Select a workload in the table to open its details page. The details page includes a properties panel, Dependency View, and the Key values table. The properties in the panel reflect the workload class, such as Kubernetes DaemonSet or Kubernetes StatefulSet.

Pods tab

The **Pods** tab lists the discovered Kubernetes pods. Data counts show the total number of pods, and the number of running, pending, creating, and crashloop backoff pods. The table displays each pod with its state, Kubernetes cluster, namespace, owner kind, and IP address.

Select a count, such as Running pods, to filter the Total pods table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

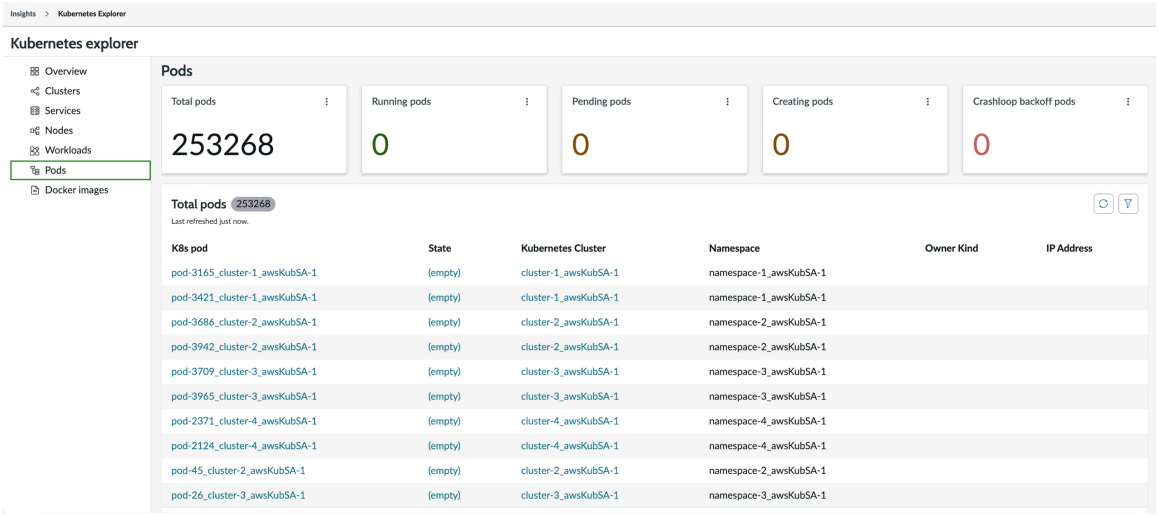

Select a pod in the table to open its details page. The details page includes a properties panel, Dependency View, and the Key values table.

Docker images tab

The **Docker images** tab lists the discovered Docker images. The table displays each image with its image ID, operating system, OS family, and OS version.

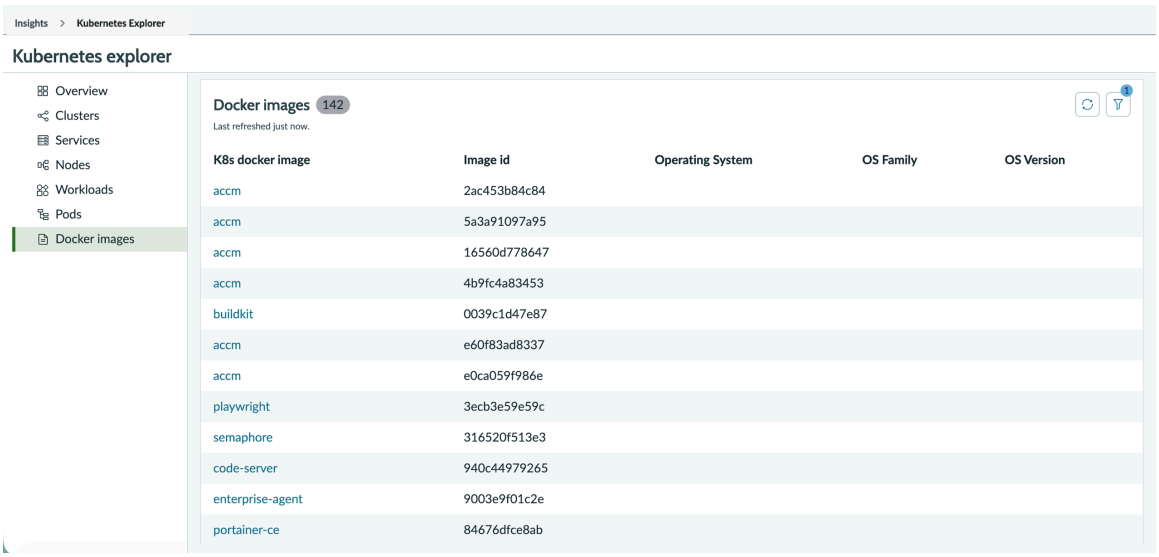

Select a Docker image in the table to open its details page. The details page includes a properties panel, Dependency View, and the Key values, Container Image OS Packages, Docker Containers, and Pods tables.

Shazzam Insights dashboard

The Discovery Admin Workspace displays your port scanning information, IP address utilization, and Discovery schedules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

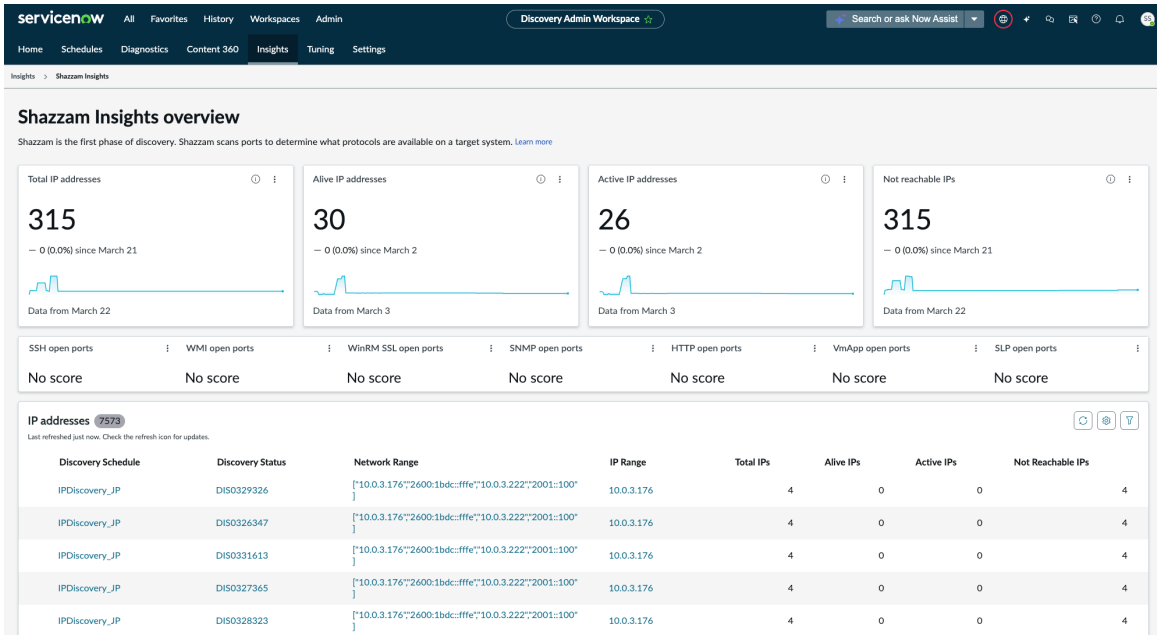

Required ServiceNow AI Platform roles

- discovery_admin
- admin

Prerequisites

Verify that you have the required setup

- Discovery.
- Discovery Admin Workspace 1.3.1. (August 2024 Store release).
- CMDB Workspace
- Visibility Content

Enable the data collection

Role required: system_administrator.

Set the *glide.discovery.shazzaminsights* property to **True**.

Run a Discovery schedule. For more information, see: Running discoveries in your network.

Access the Shazzam Insights dashboard

To open the dashboard, navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Shazzam Insights**.

Indicators

Total IP addresses

The total number and the trend of IP address data that Shazzam scanned since a given date. The data is extracted from the Shazzam Summary [shazzam_summary] table.

Alive IP addresses

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The number and the trend of responsive and pingable IP addresses that Shazzam scanned since a given date. The data is extracted from the Shazzam Status [shazzam_status] table.

Active IP addresses

The number and the trend of connected IP addresses that Shazzam scanned since a given date. The data is extracted from the Shazzam Status [shazzam_status] table.

Not reachable IPs

The number and the trend of unresponsive IP addresses that Shazzam scanned since a given date. The data is extracted from the Shazzam Summary [shazzam_summary] table.

Open ports

- The score of SSH open ports
- The score of WMI open ports
- The score of WinRM SSL open ports
- The score of SNMP open ports
- The score of HTTP open ports
- The score of VmApp SSL open ports
- The score of SLP open ports

Filters

Name	Type	Description
IP addresses	Boolean	The filter enables you to select all Discovery schedules you want to review by category. The available categories are: • IP Range • Alive • Active

Tag Governance Insights dashboard

View metrics like tag policy coverage, compliance status, and usage trends directly from the Discovery Admin Workspace.

Before you begin

Verify the following setup:

- Tag Governance (app version 1.9.0 or later) is installed.
- Discovery Admin Workspace (app version 1.13.0 or later) is installed.
- The Discovery role is discovery_admin or admin.
- Tag Governance role is sn_itom_tag.tag_governance_admin.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Access the dashboard

To access the dashboard, navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Tag Governance Insights**.

For details about the dashboard, see the "Dashboard details" section in View Tag Governance metrics.

Virtual Machine Explorer dashboard

The Discovery Admin Workspace Virtual Machine Explorer dashboard displays information on the activity and capacity of your virtual machine (VM) instances as revealed by the discovery process.

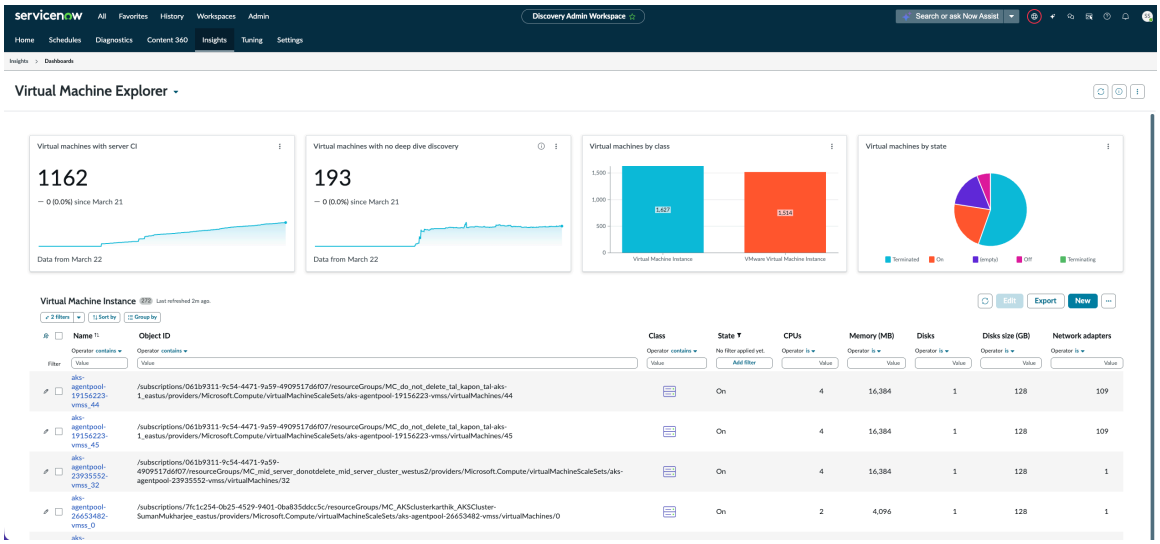

To access the dashboard, navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Virtual Machine Explorer**.

Key features

The VM Explorer dashboard enables you to make data-driven decisions through powerful visualizations.

Select the **More options** icon (⋮), then select **Refresh** to refresh the data for each visualization in this section.

Data visualizations

Report title	Type	Description
Virtual machines with server CI	Indicator and line chart	Displays the total number and trend of VMs that have a server configuration item (CI) discovered since a given date.
Virtual machines with no deep dive discovery	Indicator and line chart	Displays the total number and trend of the VMs that have no created CI relationships discovered since a given date.
Virtual machine by class	Bar chart	Displays the number of VMs in your infrastructure sorted by class.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data visualizations (continued)

Report title	Type	Description
		Virtual Machine Instances [cmdb_ci_vm_instance] table is the source table for this chart.
Virtual machines by state	Pie chart	Displays the percent of VMs grouped by the following states: on, off, or other. Virtual Machine Instances [cmdb_ci_vm_instance] table is the source table for this chart.

Note: When you first access the VM Explorer dashboard after installing Discovery Admin Workspace, the Virtual machines with server CI and Virtual Machines with no deep dive discovery indicators might not display scores. For instructions on how to collect data for these scores, see Run historical data collection for Virtual Machine Explorer.

Virtual Machine Instance

This section offers an overview of VM instances discovered, featuring key information such as the Object ID, Class, and Network adapters. You can select an instance from the table to view detailed information about it, or select the **Edit form** icon ( ) to edit details and view activity.

The filter icon ( ) enables you to apply fixed filters or create custom filters that can be saved and reused. To export the table, choose the **Export** option and select your preferred format: Excel, CSV, JSON, or PDF.

Run historical data collection for Virtual Machine Explorer

After installing Discovery Admin Workspace, run historical data collection jobs to populate Virtual Machine Explorer reports.

Before you begin

The application scope must be set to **Global**.

Role required: pa_data_collector or admin

Procedure

1. Navigate to **All** > **Platform Analytics** > **Data Collector** > **Jobs**.
2. Select the **filter** icon.
3. **[Name][contains]**`historic`**AND [Name][contains]**`VM`.
4. Run the filter by selecting **Run**.
5. Select **[PA VMs with links] Historic Data Collection** from the table.
6. Select **Execute Now**.
   The historical data collection job collects data for the last 60 days.
7. Select the **back** icon to return to the previous table.
8. Select **[PA VMs with NO links] Historic Data Collection** from the table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. Select **Execute Now**.
   The historical data collection job collects data for the last 60 days.

What to do next

Manually refresh the indicators on the Virtual Machine Explorer dashboard to populate the Virtual machines with server CI and Virtual machines with no deep dive discovery reports. This enables you to view relevant data and trends. For more information, see Virtual Machine Explorer dashboard.

URL Discovery insights dashboard

The Discovery Admin Workspace URL Discovery insights dashboard displays captured data and analytics for web domains accessed on managed Windows and macOS devices.

To access the dashboard, navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **URL Discovery insights for SAM**.

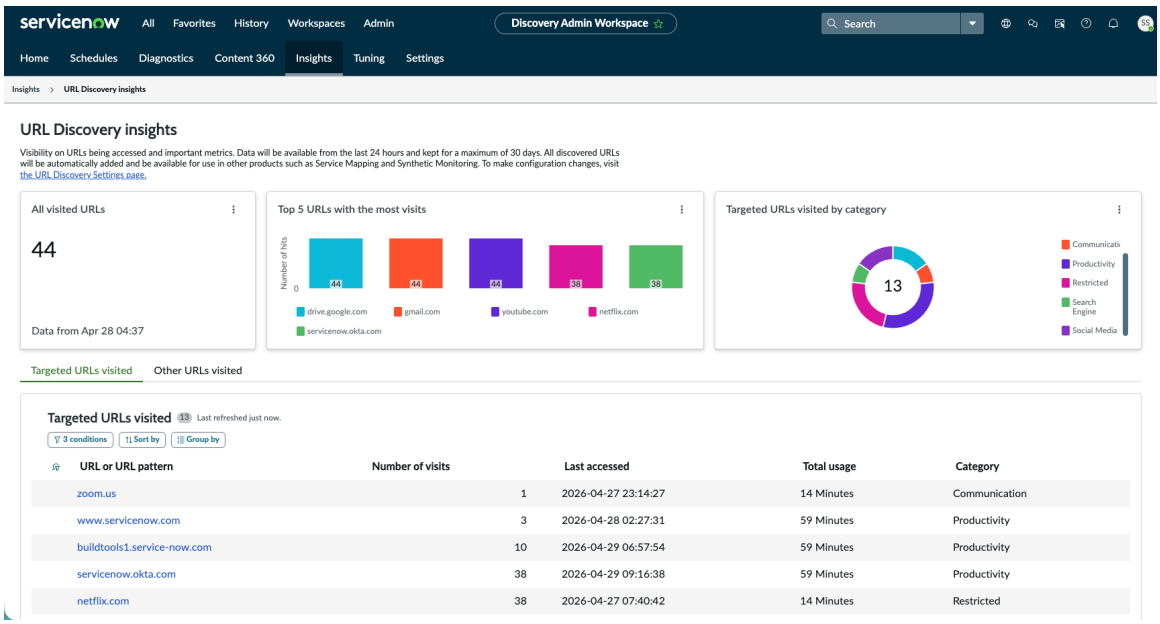

Prerequisites

Verify that you have the required setup

- ServiceNow AI Platform: Brazil, Australia, Zurich, or the Yokohama release starting with Patch 6
- Software Asset Management Professional plugin (com.snc.samp)
- Agent Client Collector for Visibility - Content (ACC-VC), starting with v1.8.0
- DEX Browser Extension 4.1.0 enabled on the host.
- Google Chrome 114 or later, or Microsoft Edge 96 or later.
- ITOM URL Discovery, starting with v1.1.0
- Discovery Admin Workspace, starting with v1.14.0

Enable the data collection

Role required: discovery_admin or admin

Set the *sn_itom_url_disc.enable_full_monitoring* property to `True`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Collect web usage data using Agent Client Collector for Visibility Content.

Required ServiceNow AI Platform roles

- discovery_admin
- admin

Key features

The URL Discovery insights dashboard enables you to make data-driven decisions through visualizations. The data is extracted from the Application Domain [sn_acc_vis_content_application_domain] table.

Select the **More options** icon (⋮), then select **Refresh** to refresh the data for each visualization in this section.

Data visualizations

Report title	Type	Description
All visited URLs	Data count	Displays the total number of URLs visited across managed devices since a given date. **Note:** This data count displays only when Broad discovery is enabled.
Top 5 URLs with the most visits	Bar chart	Displays the five URLs with the highest visit counts across managed devices since a given date.
Targeted URLs visited by category	Donut chart	Displays the distribution of targeted URL visits grouped by URL category since a given date.

Note: Daily aggregation jobs update domain-level metrics, including total page views and last accessed time. Data may not reflect activity from the current day until the job has run.

Targeted URLs visited

This section lists the targeted URLs visited across managed devices since a given date. Targeted URLs are those configured for monitoring in **Workspaces** > **Discovery Admin Workspace** > **Settings** > **URL Discovery**. For more information, see Discovery Admin Workspace Settings.

Select a URL from the list to open its domain details page, which displays domain-level usage metrics including total page views, total accessed time, last accessed time, and the number of distinct users.

Note: These metrics are aggregated per domain by the daily aggregation job and may not reflect current-day activity until the job runs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select the filter icon ( ) to apply filters to the table.

Other URLs visited

The Broad URLs visited list displays URLs accessed on managed devices that aren't in the targeted URLs list and is populated only when broad URL discovery is enabled. The list includes the URL and the number of hits recorded for each domain. To add a URL to your targeted discovery list, select the check box next to the URL and select **Add to targeted discovery**. In the dialog, enter the following information:

Targeted discovery

Field	Description
Internal or External	Whether the domain is internal or external. The default value is External.
Category	The URL category to assign to the domain. The Entity Categories [sn_acc_vis_content_entity_category] table populates the available categories.
Product	The product associated with the domain. The Software Products [samp_sw_product] table populates the available products. **Note:** Beginning with URL Discovery v1.2.1, this field is optional. However, if you select a product, **Application Service** is set to a required field.
Publisher	The publisher of the product associated with the domain.
Application Service	The application service associated with the domain. The Service Instances [cmdb_ci_service_auto] table populates the available services. **Note:** Beginning with URL Discovery v1.2.1, this field is optional unless a Product is selected.

Select **Save** to add the URL to the Targeted URL discovery. The URL is then monitored and appears in the **Targeted URLs visited** tab.

Select the filter icon ( ) to apply filters to the table.

Note: Mapping a SAM product to a URL requires the sam_user role.

Discovery Admin Workspace Tuning

The Tuning page enables you to validate your existing configurations and resolve configuration issues.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Admin Workspace Tuning Home page

The Discovery Admin Workspace Tuning Home page provides tools to assess instance health through configured checks and displays failed checks by priority, category, and suite.

Tuning check results

Use the tuning tools to interrogate your instance for configurations that indicate health issues and identify opportunities to address best practices. Tuning checks your existing configurations and helps you resolving issues. For more information on how to run checks, see Run Tuning checks.

Discovery Admin Workspace Settings

The Settings page enables you to customize and manage high-level Discovery properties so they're tailored to meet your specific needs.

To access the Discovery Admin Workspace Settings page, navigate to **Workspaces** > **Discovery Admin Workspace** > **Settings**.

Note: The capabilities described here are available starting with Discovery Admin Workspace v1.19.0. Specific version requirements are noted for individual features where applicable.

General

The General settings control how Discovery visualizations are displayed in the workspace.

Dashboard time scale

Enables you to see the health of your schedules over your preferred time span.

Use the **Time scale** drop-down list to choose between 7, 14, and 30-day time periods. The default time period is 7 days. After you configure it, all relevant data visualizations across the Home, Schedules, and Diagnostics pages update to measure over this time period.

Important: After these settings are saved, you must refresh your browser to update the data visualizations with the new time scale.

Note: No changes are applied until you select **Save**.

Anomaly detection

The Anomaly Detection settings control how Discovery identifies and surfaces irregular schedule behavior within the Discovery Admin Workspace.

Anomaly detection

Identifies unusual behavior, like failed runs and significant deviations from the thresholds for high error counts, longer discovery status duration, and fewer discovered configuration items (CIs) or cloud resources.

Anomalies are categorized into the following severity levels:

- **Critical**: Schedules that have failed to run
- **Major**: Schedules that have a low CI discovery rate
- **Minor**: Schedules that have a long run time or a high error count

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Enable anomaly detection** toggle enables you to turn the feature on and off. This feature is on by default.

> ⚠ **Warning:** Turning off this feature significantly impacts the Discovery experience across multiple pages, charts, and tables.

Detection method

Identifies anomalies in Discovery schedules with the following approaches:

- **ML-based**: Uses machine learning (ML) algorithms to analyze historical schedule data, recognize complex patterns, and detect anomalies without predefined rules.
- **Stats-based**: Relies on traditional mathematical techniques, such as calculating standard deviations to identify outliers based on fixed numerical thresholds.

Select the **Detection method** drop-down list to choose an approach. The ML-based approach is the default detection method.

You can also adjust the sensitivity of anomaly detection. Lower sensitivity reports more anomalies, which might lead to more false positives, which means that normal scans are flagged as anomalies. Higher sensitivity reports fewer anomalies, which might result in more false negatives, which means that actual anomalies might not be detected.

Use the **Sensitivity** drop-down list to choose your preferred detection sensitivity. Choices include Lowest, Low, High, Highest. By default, anomaly detection is set to Low sensitivity.

Adjust thresholds

Adjusting thresholds helps prevent false positives, making sure that normal scans aren't mistakenly flagged as anomalies. When the anomaly detection method flags a Discovery status as an anomaly, the status metrics are analyzed for deviations from the average measurements based on the specified thresholds.

Scans are marked as anomalous when the run time, error count, or CI measurements deviate from the average by a specified threshold. The default threshold is 20%, but you can change the number that defines this threshold in the relevant text box. Selecting **Reset** reverts the threshold back to the default 20%.

> ℹ **Note:** No changes are applied until you select **Save**.

IP Address Management (IPAM)

The IPAM settings manage how Discovery integrates with your IPAM sources, controlling schedule creation and connection behavior within the workspace. For more information, see IPAM Discovery integration.

> ℹ **Important:** This feature requires the ServiceNow AI Platform to be running on the Brazil release, the Australia release, the Zurich release starting with Patch 8, or the Yokohama release starting with Patch 13. You must also install and configure Service Graph Connector Central (SGC Central), starting with v2.4.0, and Service Graph Connector for Infoblox, starting with v1.5.0. For more information, see Configuring SGC Central and Configure Service Graph Connector for Infoblox using SGC Central.

Auto-create Discovery schedules

Automatically create and manage Discovery schedules based on your IPv6 network infrastructure data stored in IPAM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After you install and configure the connectors through SGC Central, use the toggle to enable auto-created schedules. When enabled, Discovery creates schedules from incoming IPAM data and keeps them updated based on your auto-creation criteria and IP data refresh settings. Schedules created automatically by IPAM are inactive by default. To activate a schedule created by IPAM, see Activate auto-created Discovery schedules.

Use the Auto-creation criteria setting to define how Discovery organizes IPs from your IPAM connections into new Discovery schedules. Discovery evaluates the key that you provide, retrieves the corresponding value from your IPAM data by looking it up in the Key Values [cmdb_key_value] table, where it searches only subnet records. IPs that share values are assigned to the same schedule, enabling the schedule structure to follow meaningful attributes such as location or environment.

Note: If you don't want to provide a key, you can enter a single space in the **Tag key** field.

Use the data refresh settings to keep auto-created schedules aligned with the latest IPAM information. Discovery checks for new, updated, or removed IPs during each refresh and updates schedule mappings accordingly. This helps keep schedules accurate as your network changes.

IPAM connectors and connections

Provides a centralized view of your installed IPAM connectors and connections, along with access to their configuration details and drafts.

The **Installed connections** tab displays the details of your existing IPAM connections. Selecting a connection from this list opens the connection record, where you can view information such as connection properties, status, and data sources.

The **Installed connectors** tab shows the IPAM connectors that are installed on your instance. Selecting a connector from this list opens its connector record, where you can review its configuration and confirm that it's ready to send imported IP data to Discovery.

The **Drafts connections** tab lists any IPAM connections that were created but not fully configured. Select **Resume setup** next to the draft connection to complete the required settings before it is set to active.

IP data from IPAM connections

Enables you to view the IP address information imported from your IPAM solutions and understand how that data is organized for Discovery. This includes lists that show the IPs retrieved through your IPAM connections and how those IPs are mapped to auto-created schedules.

Select the **IP schedule mapping** link to access the IP inventory page, where you can manage the IP ranges, networks, and IPAM data that Discovery relies on. For more information, see Discovery Admin Workspace IP inventory.

Note: This link displays only when auto-created schedules are enabled.

Select the **Imported IPs** link to access the IPAM Data page, displaying a list of IP data imported from your IPAM solutions through SGC connections. To avoid data corruption, don't edit or add IP data to the Imported IPs list.

Note: No changes are applied until you select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To synchronize your IP data immediately instead of waiting for the next refresh time, select **Refresh IP data**.

Discovery notifications

Discovery notifications enable administrators to receive real-time alerts or daily summaries of critical Discovery errors and schedule failures through Microsoft Teams and email, directly from the Discovery Admin Workspace.

Important: This feature requires the ServiceNow AI Platform to be running on the Brazil, Australia, Zurich, or the Yokohama release starting with Patch 6. Before you can set up notifications, you must configure the Microsoft Teams Graph spoke. For more information, see Set up the .

Notification destinations

Enables you to configure where your team receives Discovery notifications to respond quickly to urgent anomalies, errors, and failures.

To integrate your Microsoft Teams channel or email for notifications, select **Configure** the first time you set up the integration, or use the **Edit** icon ( ) to update it at any time. For Microsoft Teams, enter a channel name and URL in the corresponding fields.

Note: The channel name is unique to Discovery Admin Workspace, but the channel URL must be obtained directly from Microsoft Teams.

Select **Send test notification** to validate the integration. A failed test displays an error message, while a successful test sends a notification to your Teams channel. After you send the test notification and select **Save**, the connection status updates. A successfully validated channel URL displays a **Connected** status. If validation fails, the status changes to **Disconnected** and includes a timestamp for the last attempted connection.

For email, enter a valid email address in the text box and select **Add**. There's no limit to the number of recipients that you can add. To remove a recipient, select **X**. Recipients can opt out of email notifications by selecting the **Unsubscribe** link in the email and sending the auto-generated response.

Note: The **Unsubscribe** option is available only if the recipient's email address is associated with a sys_user record.

Notification settings

After configuring your notification destinations, you can use the settings here to create and update notification rules, view connected Microsoft Teams channels, and adjust integration details as needed.

Use the toggle to enable Discovery notifications. This feature is off by default.

When enabled, the notification table displays four notification categories:

- **Auto created IPAM schedules**: Notifies you when an auto-created IPAM schedule requires manual activation.
- **Critical anomalies**: Notifies you when anomaly detection identifies a schedule that has failed to run.
- **MID cluster down**: Notifies you when a Discovery-specific MID cluster goes down.
- **MID server down**: Notifies you when a Discovery-specific MID Server goes down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To configure a notification, select the **Edit** icon ( ) or the **Name** hyperlink. You can also use inline editing within the table.

By default, all Microsoft Teams and email notifications are enabled. The notification frequency is set to **Immediately** for all notifications, except critical anomalies, which is set to **Daily**.

URL discovery

The URL discovery settings control how Agent Client Collector for Visibility (ACC-VC) collects web usage data from managed Windows and macOS devices. You can enable broad URL monitoring and manage targeted URLs to monitor from this page.

Important: This feature requires the ServiceNow AI Platform to be running on the Brazil, Australia, Zurich, or the Yokohama release starting with Patch 6. Other app dependencies include:

- Software Asset Management Professional plugin (com.snc.samp)
- ITOM URL Discovery, starting with v1.1.0
- Discovery Admin Workspace, starting with v1.14.0
- Agent Client Collector for Visibility (ACC-VC), starting with v1.8.0

The browser extension must be enabled on the host.

Broad URL Discovery

Enables full monitoring of web usage data from managed devices. Use the **Discover all URLs** toggle to enable data to be captured for all URLs that are visited.

Important: To enable or disable the **Discover all URLs** toggle, the application scope must be set to **ITOM URL Discovery**. If another scope is selected, the toggle doesn't work.

This feature is inactive by default. Data is captured only for your targeted URLs. When enabled, data is captured for all URLs accessed on Chrome-based browsers across managed devices and any targeted URLs. Data is kept for a maximum of 30 days. Data is deleted if a URL is removed or if the discovery type is changed.

Warning: Enabling this feature captures data for all URLs accessed on managed devices. This method may not be compliant with EU General Data Protection Regulation (GDPR) and other privacy regulations. Review your company policies before enabling.

Targeted URL discovery

The Targeted URL Discovery list displays the URL or URL domain configured for monitoring. When broad URL discovery is turned off, the system captures data only for domains in this list.

Select **Add new** to add URLs to the monitored list. In the Add new URL dialog, select **Individual** or **Bulk** to choose how to add URLs.

- Individual: Complete the required fields in the Add new URL form and select **Save**.

Note: For a description of the fields in this form, see URL Discovery insights dashboard.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Bulk: Add multiple URLs at once using a file upload. The bulk upload follows a three-step process:

  1. Download template: Select a template to download.
  2. Add file: Select a `.csv` or `.xlsx` file to upload. Select **+ Add another file** to include additional files. Select **Next** when at least one file is attached.
  3. Review: Review the parsed URL entries in the snapshot preview table before submitting. The table displays the URL or URL Domain, Internal or External classification, Category, Product, Publisher, and Application Service for each entry, along with a total URL count. Select **Add** to submit all URLs. Or, select **Cancel** at any step to close the dialog without saving.

  Important: Beginning with URL Discovery v1.2.1, when you add a URL, the following rules apply:

  - A URL with no application and no product is added as a standalone domain.
  - A URL with an application but no product creates the application without a SAM product.
  - Selecting a **Product** makes the **Application Service** field required.
  - A product without an application is rejected.
  - A duplicate URL (same normalized host and source) is rejected.

Edit an individual URL by selecting the URL or URL domain hyperlink.

Note:

- Mapping a SAM product to a URL requires the sam_user role.
- No changes are applied until you select **Save**.

To remove a URL from the targeted monitoring list, select the corresponding check box and then select **Remove**. Confirm the removal when prompted. The URL is removed from the list immediately and is no longer actively monitored. Historical monitoring data for the removed URL is retained for 30 days before being deleted.

Select the filter icon ( ) to apply filters to the table.

URL Discovery Insights

Select the link icon to access the URL Discovery insights dashboard, where you can view captured data and analytics for accessed URLs. For more information, see URL Discovery insights dashboard.

X devices with ACC agent installed

Select the link icon to access the Agent Client Collectors [sn_agent_cmdb_ci_agent] table.

AI in Discovery

AI-powered features in Discovery help administrators investigate missing CI attributes, diagnose discovery errors, manage certificates, and request firewall rules using automation and generative AI.

Discovery includes several AI-powered features. All features require ServiceNow Otto for IT Operations Management (ITOM) to be installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

AI features in Discovery

AI feature	Description	Use case	Resources
Pattern diagnostic agentic workflow	Investigates missing CI attributes in pattern-based discovery. The workflow identifies the attribute gap, parses discovery logs, determines the root cause, and suggests a remediation — all from the ServiceNow Otto panel. Uses two agents: the Pattern diagnostic agent and the EF Remediation Agent.	A discovery administrator notices that a CI attribute is missing and wants to identify the root cause and a fix without manually navigating discovery logs and tables.	Pattern Diagnostic agentic workflow
AI Insights for discovery errors	Analyzes current Discovery errors and surfaces the highest-impact and highest-leverage errors in the Discovery Admin Workspace Diagnostics page. Each insight includes a summary of affected schedules and devices, and recommended remediation steps. Powered by ServiceNow Otto for Error Framework.	A discovery administrator wants to quickly identify which errors to resolve first to have the greatest effect on discovery health, without reviewing every error individually.	Discovery Admin Workspace Diagnostics

Discovery for cloud environment

ITOM Visibility cloud discovery solutions enable you to collect detailed information about your cloud-based infrastructure and your resources in major cloud service providers: Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), IBM Cloud Platform, Oracle Cloud Infrastructure (OCI), and Alibaba Cloud.

https://player.vimeo.com/video/1188194804?h=11d049501f&badge=0&autopause=0&player_id=0&app_id=58479

ITOM Visibility cloud discovery can be understood through the layers at which discovery operates and the methods used to perform discovery. Different methods or combination of methods at the different layers support various business outcomes, giving you the flexibility to customize your discovery approach to your organization's specific needs.

Business outcomes

Effective cloud discovery is crucial for maintaining an accurate and up-to-date inventory of cloud resources. This comprehensive visibility into the cloud infrastructure enables organizations to drive several key business outcomes. It empowers informed decision-making, accelerates innovation by optimizing resource allocation, and helps cloud investments align directly with strategic business objectives.

Business outcomes supported by cloud discovery

Business outcome	Cloud discovery role
Regulatory compliance	Provides the necessary visibility and data to support regulatory compliance mandates.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Business outcomes supported by cloud discovery (continued)

Business outcome	Cloud discovery role
SAM (software asset management)	Enables more effective software asset management by providing comprehensive visibility into cloud software deployments.
FinOps (financial operations)	Provides comprehensive visibility into cloud resources and their usage, enabling FinOps teams to optimize cloud spend across cloud providers, organizations, and cost centers.
SecOps (security operations)	Offers continuous visibility into cloud resources and their configurations, supporting the security and compliance of cloud environments.
Certificate management	Discovers and tracks certificates, their expiry, and usage, enabling teams to reduce the risk of certificate-related outages and assess potential business impact.
AIOps (artificial intelligence for IT operations)	Identifies and maps cloud resources, their configurations, and interdependencies. Cloud discovery provides the foundational data for AIOps platforms to monitor proactively, analyze, and predict potential issues, ultimately reducing outage risks and supporting business continuity.

For more information on which cloud discovery methods best support each business outcome, see the **Matching methods to business outcomes** section.

Cloud discovery layers

To achieve these business outcomes, you need to understand how cloud discovery works. Cloud discovery can be understood in two distinct layers: the cloud metadata (or topology) layer and the operating system (OS) layer. Each layer provides a different level of visibility into cloud infrastructure, from structural inventory to a detailed operational state. Think of it like mapping a neighborhood versus inspecting individual buildings.

The cloud metadata layer provides a high-level view of cloud infrastructure by outlining its structural components. This outline includes identifying resources like virtual machines, networks, and subnets, effectively creating an inventory of what exists. It's similar to mapping all the buildings in a neighborhood, showing how everything is organized and interconnected. This layer helps to understand the overall architecture and resource relationships within a cloud environment.

The OS layer reveals the internal state of cloud resources, including installed software, active services, running processes, and system configurations. It's like stepping inside each building in a neighborhood to inspect its internal setup. This deeper level of discovery helps monitor operations, verify configurations, and maintain compliance and security across environments.

These layers can be discovered through the different cloud discovery methods and together, the two layers enable comprehensive visibility from structural inventory to detailed operational insights.

Cloud discovery methods

ITOM Visibility supports several methods for cloud discovery, depending on the layer. Different discovery methods address different organizational needs and constraints.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Cloud metadata layer: Pattern-based Cloud Discovery and Service Graph Connectors
- OS layer: IP-based Horizontal Discovery, Agent Client Collector for Visibility Content, and Service Graph Connectors

Note: The pattern-based Cloud Discovery and IP-based Horizontal Discovery methods operate at different layers but both use the Discovery and Service Mapping Patterns application.

Pattern-based Cloud Discovery/IP-based Horizontal Discovery

The Discovery and Service Mapping Patterns application supports both cloud metadata and OS layer discovery. Patterns execute commands in a pull-based, scheduled approach by targeting devices within specified IP ranges. Key characteristics include:

- Requires OS and application credentials
- Needs firewall access from a MID Server
- Works with AWS, Azure, GCP, OCI, Alibaba Cloud, and IBM clouds
- Supports continuous visibility for AWS, Azure, and GCP

For more information, see Discovery patterns used by ITOM Visibility.

Agent Client Collector for Visibility Content

An agent-based, push-based cloud discovery that provides near real-time discovery and supports OS layer discovery. Key characteristics include:

- Supports deployment with or without a MID Server
- Doesn't require OS-level credentials
- Works on Windows, Linux, and macOS
- Uses system accounts, local service accounts, or Microsoft gMSA
- Supports OS layer discovery only
- Works with AWS, Azure, and GCP

For more information, see Agent Client Collector Discovery.

Service Graph Connectors

Third-party connectors that perform scheduled imports by pulling data directly from cloud provider APIs. Service Graph Connectors support both cloud metadata and limited OS layer discovery. Key characteristics include:

- Requires read access to the provider's system
- Doesn't require OS credentials or MID Server firewall access
- Works with AWS, Azure, and GCP
- Supports continuous visibility for AWS

For more information, see Integrating third-party data into CMDB.

Each method has its own requirements and supports different use cases, depending on the cloud provider. For more information, see:

- AWS: AWS cloud discovery methods and use cases
- Azure: Azure cloud discovery methods and use cases
- GCP: GCP cloud discovery methods and use cases

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a list of AWS, Azure, and GCP resources discovered by Discovery and Service Mapping Patterns and Service Graph Connectors, see Cloud discovery solutions comparison.

Continuous visibility

Beyond choosing a cloud discovery method, you must also consider when and how often discovery occurs. Cloud environments are highly dynamic, with resources constantly being created, modified, or deleted. Traditional scheduled discovery often fails to keep up with these rapid changes, leading to a stale and inaccurate CMDB. Continuous visibility addresses this challenge by capturing changes as they happen in the cloud. There are two main types of continuous visibility: pull-based and push-based.

In pull-based discovery, the ServiceNow instance connects to cloud providers at set intervals to retrieve and update CMDB data. This approach promotes regular synchronization and offers more frequent updates than traditional scheduled scans. The pattern-based Cloud Discovery method supports this approach for both Azure and GCP.

In push-based discovery, the cloud environment itself is configured to send change events directly to the ServiceNow instance. This approach enables near real-time updates to the CMDB. Both Service Graph Connectors and the pattern-based Cloud Discovery method support push-based discovery for AWS.

Note: Continuous visibility for Amazon AWS Cloud, Microsoft Azure Cloud, and Google Cloud Platform doesn't require Cloud Provisioning and Governance.

For more information, see:

- Azure: Microsoft Azure Alert driven discovery
- GCP: Configure the Google Cloud Logging service to auto-update the CMDB
- AWS: AWS events-driven discovery

Choosing methods for desired business outcomes

Choosing the right discovery methods at the cloud metadata and OS layers for your desired business outcomes is essential for optimizing your configuration management strategy and achieving operational goals.

Cloud metadata layer discovery methods support the following business outcomes:

Cloud metadata layer discovery

Business outcomes	Pattern-based Cloud Discovery	Service Graph Connectors
Regulatory compliance	Yes	Yes
Basic AIOps	Yes	Yes
Continuous visibility	AWS, Azure, and GCP	AWS only

OS layer discovery methods support the following business outcomes:

OS layer discovery

Business outcomes	IP-based Horizontal Discovery	Agent Client Collector for Visibility Content	Service Graph Connectors
Regulatory compliance	Full	Full	Limited

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OS layer discovery (continued)

Business outcomes	IP-based Horizontal Discovery	Agent Client Collector for Visibility Content	Service Graph Connectors
SAM	Enterprise	Enterprise	Basic
FinOps	Yes	Yes	No
SecOps	Yes, including file-based discovery	Yes, including file-based discovery	No
Certificate management	Yes	Yes	No
AIOps	Full	Full	Limited
Oracle GLAS (Global Licensing Advisory Services) audits	Yes	Yes	No
Hibernating or autoscaling resource visibility	No	Yes	No

Managing your cloud resources

Discovery Admin Workspace provides a centralized location for managing and monitoring discovery operations, including performance monitoring, issue diagnosis, and maintaining configuration accuracy. From tracking schedules and diagnosing errors to tuning configurations and accessing insights, Discovery Admin Workspace integrates key tools to improve efficiency and visibility. You can identify critical tasks, manage discovery processes, and resolve issues with enhanced diagnostics and reporting. For more information, see Discovery Admin Workspace.

If you have Cloud Account Management (CAM) within Cloud Workspace, you can use CAM to access your cloud account and resource inventory through the Cloud Asset Explorer, Cloud Compliance, and Cloud Accounts dashboards. Cloud Workspace provides comprehensive visibility into cloud accounts and assets across providers, enabling you to track ownership, monitor account trends, and promote regulatory compliance. For more information, see Cloud Account Management.

Cloud Discovery setup using Discovery Home

Get started with Cloud Discovery by configuring credentials, service accounts, schedules, and MID Servers.

Setting up Cloud Discovery is the first stage in performing cloud discovery and using Cloud Provisioning and Governance for managing discovered cloud resources.

Cloud Discovery is part of the ServiceNow AI Platform and deploys some of its platform-wide mechanisms and features. At the same time, there are some configurations that are specific to Cloud Discovery only.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Perform the following tasks in the exact order they are listed below:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Evaluate your cloud deployment and plan how you are going to discover it.
2. Request Discovery.
3. Install and configure the MID Servers.
4. Set up service accounts and configure access to them.
   - Setup AWS service accounts
   - Create Azure service accounts
   - Create GCP service accounts
   - Create VMware service accounts
   - Set up Alibaba Cloud service accounts
5. For Amazon AWS Cloud and Microsoft Azure Cloud, configure notifications and alerts from the Amazon AWS Cloud and Microsoft Azure Cloud to make the necessary updates to your CMDB without additional scanning.
   - AWS events-driven discovery
   - Configure the Microsoft Azure Alert service to auto-update the CMDB
   - Configure the Google Cloud Logging service to auto-update the CMDB
   - Configure the VMware Events service to auto-update the CMDB

   Note: Event-driven discovery does not require Cloud Provisioning and Governance.
6. Ensure that your ServiceNow instance has the relevant patterns and CI classes (types) for cloud resources. You may need to deploy the latest version of the following applications from the ServiceNow Store.
   - Discovery and Service Mapping Patterns
   - CMDB CI Class Models
7. Create Discovery schedules for cloud resources.

Cloud discovery setup using Cloud Discovery Workspace

Perform the necessary procedures to prepare for discovering resource CIs on the cloud.

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

Setting up Cloud Discovery is the first stage in performing cloud discovery and using Cloud Provisioning and Governance for managing discovered cloud resources.

Cloud Discovery is part of the ServiceNow AI Platform and deploys some of its platform-wide mechanisms and features. At the same time, there are some configurations that are specific to Cloud Discovery only.

Cloud Discovery Workspace present an enhanced way of creating discovery schedules. You can also check prerequisites required for discovering cloud resources, like credentials and service accounts.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Perform the following tasks in the exact order they are listed below:

1. Evaluate your cloud deployment and plan how you're going to discover it.
2. Request Discovery.
3. Install Cloud Discovery Workspace.
4. Install and configure the MID Servers.
5. (For Amazon AWS Cloud and Microsoft Azure Cloud) Configure notifications and alerts from the Amazon AWS Cloud and Microsoft Azure Cloud to make the necessary updates to your CMDB without additional scanning.
   - AWS events-driven discovery
   - Configure the Microsoft Azure Alert service to auto-update the CMDB
   - Configure the Google Cloud Logging service to auto-update the CMDB
   - Configure the VMware Events service to auto-update the CMDB

   Note: Event-driven discovery does not require Cloud Provisioning and Governance.
6. Ensure that your ServiceNow instance has the relevant patterns and CI classes (types) for cloud resources. You may need to deploy the latest version of the following applications from the ServiceNow Store.
   - Discovery and Service Mapping Patterns
   - CMDB CI Class Models
7. Create a discovery schedule in Cloud Discovery Workspace.

Related topics

Cloud Discovery Workspace

Create Discovery schedules for cloud resources

Create schedules for discovering cloud resources using service accounts or IP ranges.

Discovery Manager interface

Use the Cloud Discovery Schedule Configuration interface to input the information that you need to access your service accounts and sub-accounts. You can specify the credentials to validate the selected accounts. Test the connections to your accounts and then create schedules to discover the cloud resources in your selected datacenters.

The following private and public cloud platforms support VMware Discovery:

- Alibaba Cloud
- Amazon Web Services (AWS)
- Google Cloud Platform (GCP)
- IBM Cloud
- Microsoft Azure
- OpenStack
- Oracle Cloud Infrastructure (OCI)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- oVirt
- VMware

Note: The term "sub-accounts" is used generically in the Cloud Discovery Schedule Configuration interface. Its meaning varies by cloud provider:

- In AWS, sub-accounts refer to member accounts under a management account.
- In Microsoft Azure, the equivalent is a subscription.
- In Google Cloud Platform (GCP), Cloud Discovery targets projects.

Progress indicator

As you complete each configuration phase, the tabs at the top of the Discovery Manager track your progress. A check mark indicates a completed phase. A dark underline marks the current phase. The tabs are interactive, you can move to a phase by clicking that tab.

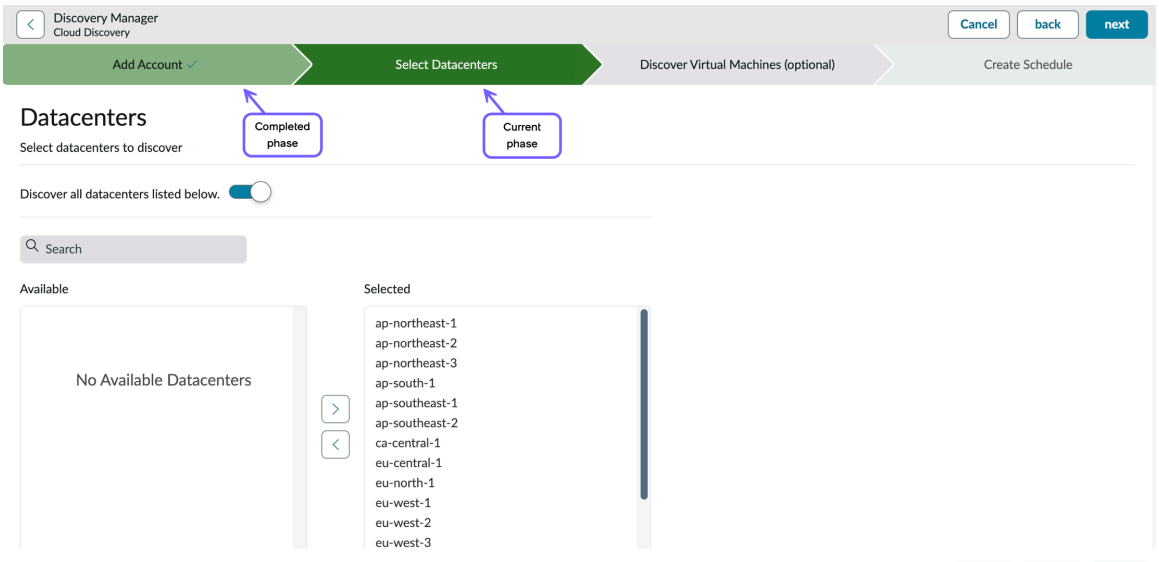

Discovery Manager label customization

The Service Account Field Options [service_account_field_options] table contains fields to modify the labels in Discovery Manager. These fields enable you to customize the Discovery Manager to handle new service providers or to alter existing labels to meet internal policies. For example, the Microsoft Azure account ID label is Microsoft Azure Subscription GUID by default. You can change it to Subscription ID, Account ID, or anything else based on your needs.

Select a service account and test credentials

Specify your cloud provider and select the service account to discover. You can include one or all available sub-accounts on a schedule. Provide the credentials for the cloud service account and test the connection.

Before you begin

Before you attempt to configure Cloud Discovery, you must have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows. Ensure to configure the MID Server used for Cloud Discovery with the Cloud Management capability.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin

About this task

The ServiceNow® Cloud Discovery wizard begins with the account selection and testing phase. To prevent long waits or unexpected timeouts, use the *cdu.record.watcher.timeout* system property to define the maximum waiting period for these Cloud Discovery configuration processes:

- Service account validation
- Member account discovery
- Discovery of logical datacenters

The default timeout value is 300000 milliseconds or 5 minutes.

Procedure

1. Navigate to **All** > **Discovery** > **Home**.
2. Select **Discover Cloud Resources** in the Welcome screen.

   Note: The Welcome screen includes a link to the Discovery Quick Start. Quick Start sets up a basic Discovery for evaluating the product or exploring new network devices. It is not intended for discovering cloud resources.

   If the system can't detect an appropriate MID Server, you're prompted to install and configure a MID Server by using the MID Server Guided Setup. This notification is also displayed if an existing MID Server is not available because it is upgrading.

3. In the Basic Info pane of the Discovery Manager, select the name of your cloud provider in the **Provider** field, such as **AWS** or **Azure**.
   The wizard displays a count of the validated and available MID Servers in the MID Servers pane.
4. If you're using AWS, Azure, GCP, or Alibaba Cloud you can specify the MID Server selection type in the **MID Selection Type** field.

   Note: AWS and Azure selection is only available after migration to pattern-based Discovery. For more information, see Migrate from CAPI to Patterns.

   The following options are available:

   - **Auto-Select MID Server**: The wizard automatically uses an appropriate MID Server.
   - **Specific MID Server**: The wizard uses the MID Server specified in the **MID Server** field. To filter and view MID Servers that are configured to assume an IAM role for discovering cloud resources, select the **Credential-Less Discovery** option.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. If no MID Server has the Cloud Management capability, select **Configure MID Servers** and configure one appropriately.

   a. Navigate to **MID Servers** > **Capabilities**.

   b. From the MID Server Capabilities list, select **Cloud Management**.
   You can also select **ALL**, which includes Cloud Management by default.

   c. In the MID Servers related list, select **Edit** and then add your MID Servers to the capability.

   d. Select **Save**.
   Your MID Servers are ready to communicate with your cloud service accounts.

6. In the **Schedule Name** field, enter a unique and descriptive name for the schedule.
7. In the Service Account pane, choose whether you want to add an account or select an existing account.
   If you select an existing account, the system auto-completes the form with the available information. If you add an account, you must provide the account ID, credentials, and URL, if applicable. For VMware, you must specify the URL.
   - The AWS account ID is the account ID.
   - The Azure account ID is the Azure Subscription GUID.
   - The GCP account ID is the project ID.
   - The VMware account ID is the virtual machine instance UUID.
   - The Alibaba Cloud account ID is the Compartment ID.

   Note: VMware discovery can run without a service account. To discover VMware resources, configure the schedule to run on the IP address of the vCenter server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Select the credentials for this cloud service account or add new credentials.

   a. To add a credential, select the plus icon (**+**).

   b. Select the appropriate cloud credential type from the list.

   c. Complete the form and select **Save**.

   > **Note:** You can configure AWS to generate temporary credentials for service accounts. This capability eliminates the need to define some credentials on the instance, and it can increase security on your member accounts (sub-accounts). For details, see Access setup for AWS service accounts.

9. To validate the account with the provided details, select **Test Account**.

   Messages in the header bar notify you of the validation results. If you specified a management account and the connection test was successful, then the wizard displays a selection screen for any members accounts (sub-accounts) that are associated with the management account. By default, all current member accounts are selected for the schedule and can't be moved to the Available list.

   > **Note:**
   >
   > If the validation process does not return any sub-accounts or any sub-accounts that you added recently, select **Refresh Members** or rerun the validation process. If validation times out repeatedly, increase the timeout interval in the *`cdu.record.watcher.timeout`* MID Server property.

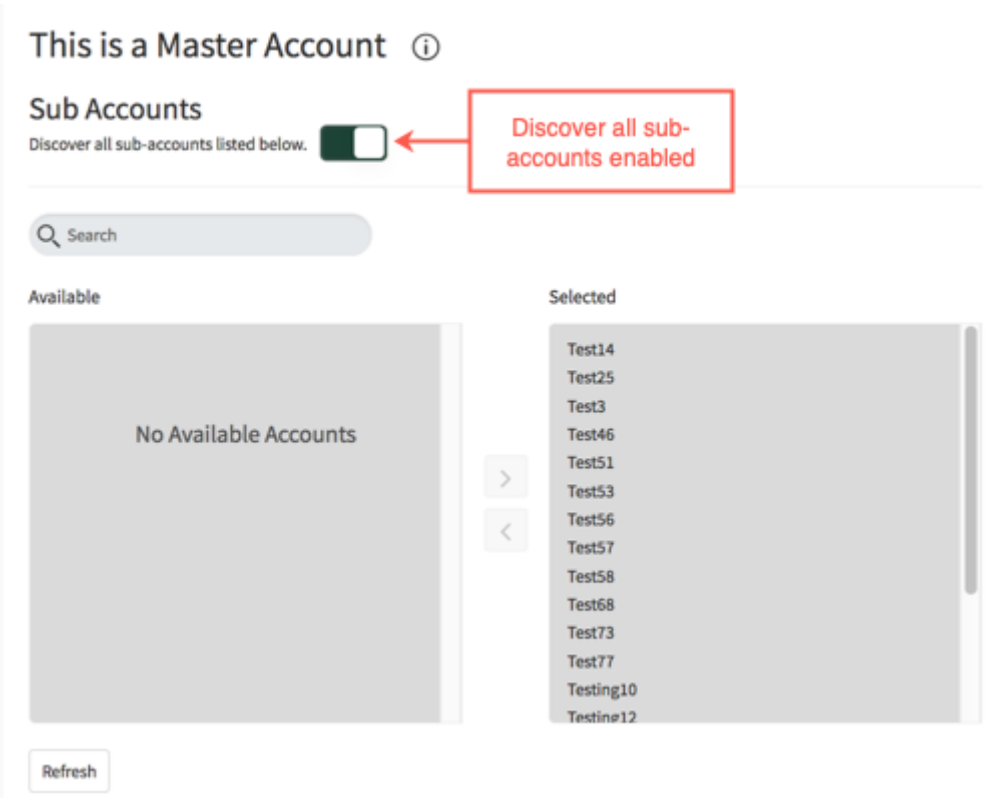

10. To select specific sub-accounts for the schedule, deactivate the auto-select feature.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, all sub-accounts are selected. To disable the feature and to permit the selection of specific datacenters to discover, slide the toggle switch to the left.

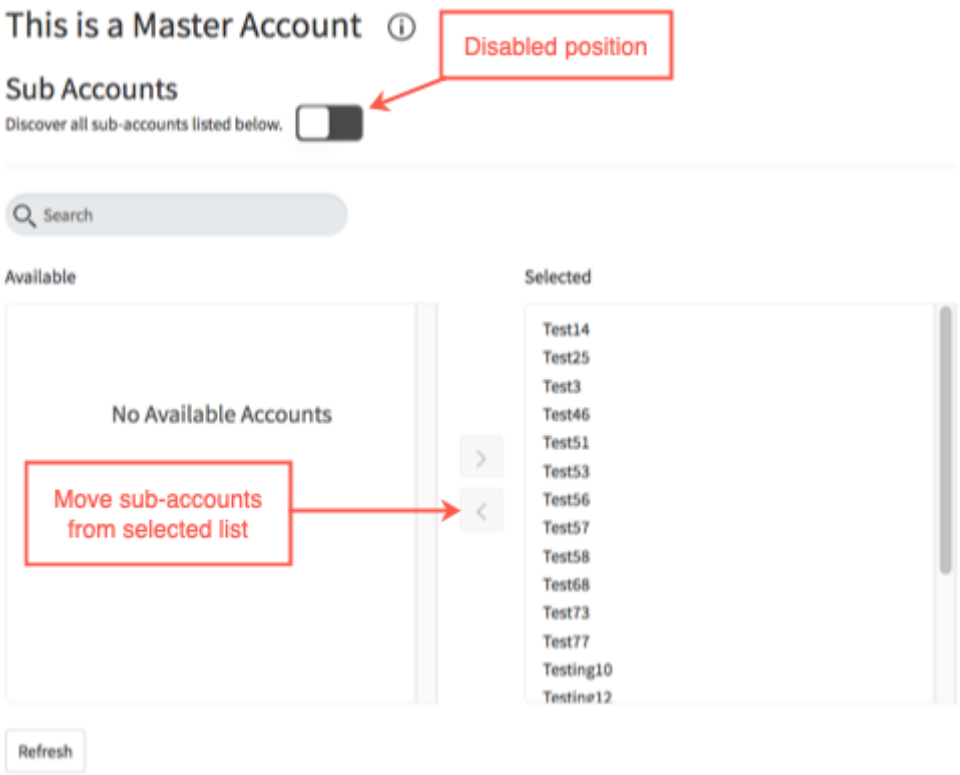

11. To exclude a sub-account, drag the sub-account from the Selected list to the Available list.
12. Select **Next**.

Select a cloud provider datacenter to discover

Cloud Discovery displays the list of available datacenters when you validate the service account. Specify which datacenters you want to discover through the schedule.

Before you begin

Role required: discovery_admin

Procedure

1. To select the datacenters to discover, click the **Select Datacenters** phase in the Discovery Manager tab bar.
   By default, all datacenters associated with the service account are selected for the Cloud Discovery schedule.

Note:

- If no datacenters appear, select **Refresh Datacenters** or re-run the account validation. If validation times-out repeatedly, increase the timeout interval in the *cdu.record.watcher.timeout* MID Server property.
- You can configure Cloud Discovery to auto-refresh the list of sub-accounts and datacenters before each scheduled discovery. For more information, see Enable auto-refreshing of sub-accounts and datacenters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. To select specific datacenters for the schedule, deactivate the auto-select feature by moving the **Discover all datacenters listed below** toggle switch to the left.
3. Move the datacenters you want to discover from the **Available** list to the **Selected** list.
4. Select the next phase in the tab bar.
   If required, you can configure IP ranges to discover virtual machines (VM), or move to the schedule phase and complete Cloud Discovery setup.

Enable discovery of virtual machines

You can discover virtual machines (VMs) during discovery of the cloud resources. Cloud Discovery performs a deep discovery of the virtual machines identified by the cloud schedule, using the IP addresses of the VMs.

Before you begin

All MID Servers you intend to use for Cloud Discovery must be validated, have a status of **Up**, and they can access the IP ranges in the cloud provider account.

Important: The upgrade process from Madrid to New York deactivates the existing Cloud Discovery schedules. Ensure to re-activate any schedules in the Discovery Manager that might have been deactivated during an upgrade from Madrid . These schedules have a **Run** status of **undefined** in the Discovery Schedule [discovery_schedule] table.

Role required: discovery_admin

About this task

You do not need to configure IP ranges for the MID Servers used with Cloud Discovery. Cloud Discovery automatically selects the IP addresses of the virtual machines in the datacenters you selected in the wizard.

Procedure

1. To discover the VMs that are part of the current schedule, click the **Discover Virtual Machines (optional)** phase in the Discovery Manager tab bar.

   By default, VM discover is active.

   Note: To learn more about the working of **Discover Virtual Machines (optional)** in Cloud Discovery, see the KB article How does the "Discover Virtual Machine" (optional) work in Cloud Discovery.

2. Select the MID Server to use.
   The following options are available:
   - **Auto-Select**: Discovery selects an appropriate MID Server automatically.
   - **Specific MID Cluster**: Discovery uses the named MID Server cluster.
   - **Specific MID Server**: Discovery uses the named MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

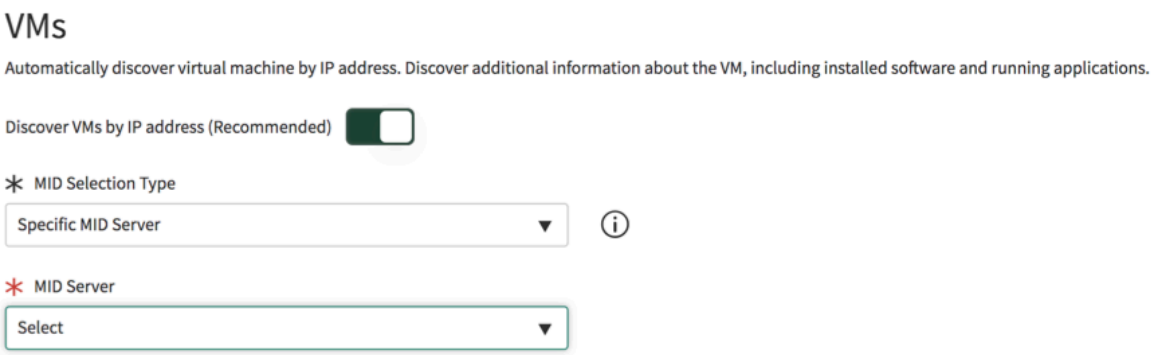

3. To disable VM discovery, move the **Discover VMs by IP address** toggle switch to the left.
   The **MID Selection Type** field disappears.
4. Select the **Create Schedule** phase in the tab bar.

Set the run time for a Cloud Discovery schedule

The final phase in the Discovery Manager wizard creates the schedule that triggers discovery of the cloud resources.

Before you begin

Role required: discovery_admin

About this task

You can also access the schedule phase of the Discovery Manager from the Discovery Schedule form. Navigate to **Discovery** > **Discovery Schedules** and select **Cloud Discovery** in the schedules list.

Procedure

1. To configure the Cloud Discovery schedule, select **Create Schedule** in the Discovery Manager tab bar.
   By default, the schedule is active.
2. Configure when and how often the schedule should run and define a timeout threshold if you want to restrict Cloud Discovery to a specific duration.

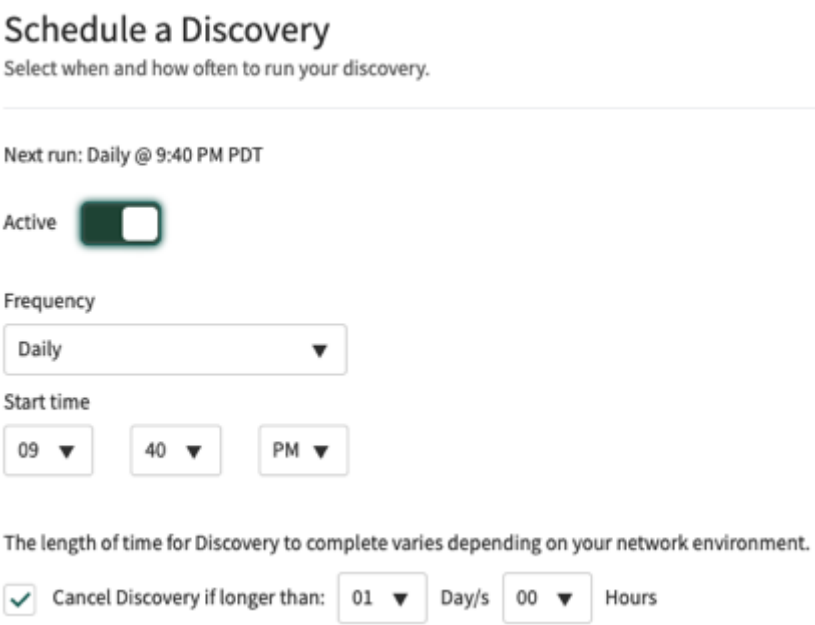

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. To deactivate the schedule, toggle the **Active** switch to the left.
4. Complete the Cloud Discovery scheduling process by selecting one of the following options:
   - **Finish and Run**: Runs Discovery immediately, using the current configuration. All subsequent cloud discoveries run on the schedule you provided.
   - **Finish**: Exits the wizard and runs Cloud Discovery on the schedule you provided.

What to do next
Open the Discovery Home page or Discovery Admin Workspace Schedules page and view the results of your Discovery.

Enable auto-refreshing of sub-accounts and datacenters

Configure the Cloud Discovery to auto-refresh the list of sub-accounts and datacenters before each scheduled discovery. If Cloud Discovery identifies any new sub-accounts or datacenters, it adds them to the Configuration Management Database (CMDB) and includes them in the discovery schedule.

Before you begin
Role required: discovery_admin

About this task
Cloud Discovery can automatically refresh the list of sub-accounts and datacenters covered in a discovery schedule, subject to the fulfilment of the following criteria:

- The cloud provider is Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform (GCP)
- The *glide.discovery.cdu.auto_refresh_sub_accounts_and_ldcs* property to `true`.
- The Cloud Discovery schedule is created for a management account.
- The Cloud Discovery schedule is configured to discover the resources of all the sub-accounts and datacenters of the management account.

Note: The term "sub-accounts" is used generically in the Cloud Discovery Schedule Configuration interface. Its meaning varies by cloud provider:

- In AWS, sub-accounts refer to member accounts under a management account.
- In Azure, the equivalent is a subscription.
- In Google Cloud Platform (GCP), Cloud Discovery targets projects.

Procedure
1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for the following system property record, and then open the record.

   glide.discovery.cdu.auto_refresh_sub_accounts_and_ldcs

3. On the form, change the Value field to `true`.
4. Select **Update**.

Configure Cloud Discovery for code-signed instances

Configure a system property to disable unsupported script tracking for Cloud Discovery operations on instances with code signing enabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Confirm the following:

- You're using the Brazil version of the ServiceNow AI Platform.
- You're using the following ServiceNow Store applications and versions:
  - Discovery Admin Workspace v1.13.0.
  - Visibility Content v6.13.10.
  - Discovery and Service Mapping Patterns v1.30.2.
- You have Code Signing enabled with the Circle of Trust established. For more information, see Configuring Code Signing.

Role required: security_admin

About this task

The *com.glide.codesigning.tracking.unsupported_script_tracking.enabled* property controls script tracking for code signing on Cloud Discovery. When this property is set to **true** and code signing is enabled, ECC queue records inserted through unsupported scripts aren't notarized, which causes Cloud Discovery operations to fail. Setting this property to **false** enables Cloud Discovery on code-signed instances.

Procedure

1. Navigate to **All** > **System Properties** > **All properties**.
2. Search for the property *com.glide.codesigning.tracking.unsupported_script_tracking.enabled*
3. Set the **Value** field to **true**.
4. Select **Update**.
   Cloud Discovery is now configured to work with code signing. You can run Cloud Discovery on your code-signed instances.

Related topics

Discovery on Code Signing instances

Code Signing

Create a discovery schedule in Cloud Discovery Workspace

Create schedules for discovering cloud resources based on the discovery method that you choose: service accounts or IP ranges. The discovery schedule defines the various settings for the cloud discovery.

Before you begin

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

- Verify that version 1.7.1 or later of the Cloud Discovery Workspace application is installed on the ServiceNow AI Platform. For more information, see Install Cloud Discovery Workspace.
- Before you attempt to configure Cloud Discovery, you must have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

instructions for Linux and Windows. Ensure to configure the MID Server used for Cloud Discovery with the Cloud Management capability.

- Ensure that suitable credentials are available in the ServiceNow AI Platform to connect with the service accounts on which you want to perform the discovery.
- Pattern-based discovery runs by default. In scenarios where the latest patterns application is not installed, discovery fails and you are prompted to install or migrate from CAPI-based discovery. For more information, see Migrate from CAPI to Patterns.

Role required: discovery_admin

About this task

The ServiceNow® Cloud Discovery wizard begins with the account selection and testing phase. To prevent long waits or unexpected timeouts, use the *cdu.record.watcher.timeout* system property to define the maximum waiting period for these Cloud Discovery configuration processes:

- Service account validation
- Member account discovery
- Discovery of logical datacenters

The default timeout value is 300000 milliseconds or 5 minutes.

You can configure Cloud Discovery to auto-refresh the list of subaccounts and datacenters before each scheduled discovery. For more information, see Create Discovery schedules for cloud resources.

Procedure

1. Navigate to **Workspaces** > *Cloud Discovery Workspace* > *Cloud Discovery*.
2. Select **New discovery schedule**.
3. Select the cloud provider.

   a. Fill the values on the Choose the provider form.
   For description of the form fields, see Create a new discovery schedule form reference.

   b. Select **Next**.

4. Select service account for the Cloud Discovery schedule or create a service account.

   For description of the service account form fields, see Cloud Discovery service account form reference. If you select an existing account, the system auto-completes the form with the available information.

   Note: VMware discovery can run without a service account. To discover VMware resources, configure the schedule to run on the IP address of the vCenter server.

   a. **Optional:** To check the connection with the service account, select **Test account**.

   b. **Optional:** If you're creating the discovery schedule for a management account, specify whether you want to discover all the member accounts or only specific member accounts. By default, all the existing member accounts are selected for the discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. **Optional:** If you're creating the discovery schedule for a management account and the validation process doesn't return the subaccounts, select **Refresh** or rerun the account validation.

d. Select **Next**.

5. Select datacenters for the discovery schedule.
   You can either select all the datacenters attached with the service account or select specific datacenters.
   a. **Optional:** If you want to discover specific datacenters, select the datacenters from the **Select the data centers** drop-down list.
   b. **Optional:** If the datacenters don't appear, select **Refresh Datacenters** or rerun the account validation.
   c. Select **Next**.
6. **Optional:** Specify whether you want to discover the Virtual Machines (VM) and their attributes.
   a. **Optional:** Specify whether you want to perform VM discovery using AWS Systems Manager.

      Note: This option displays only when the system property *glide.discovery.enable_ssm* is enabled and **AWS** is selected as the cloud provider in Step 3.

   b. Select a MID Server or MID cluster that can discover the VMs.
      If you're not sure about the MID selection, select **Auto-Select** *MID Server*.
   c. Select **Next**.
7. Create a schedule to execute the Cloud Discovery.
   a. Select an execution frequency for the Cloud Discovery schedule.
   b. Specify a start time for executing the discovery.
   c. **Optional:** Specify criteria to cancel the discovery.
   d. Activate the discovery schedule.
8. Do any of the following actions.

Option	Description
Finish and run	Cloud Discovery saves the schedule and exe cutes it.
Finish	Cloud Discovery saves the schedule and exe cutes it as per the defined frequency.

Create an OCI Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create an Oracle Cloud Infrastructure (OCI) Discovery schedule.

Before you begin

Verify the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have configured an OCI user in the Oracle Cloud console with an API signing key and the required Identity and Access Management (IAM) policies so that the MID Server can access the cloud resources you plan to discover.
- You have created a suitable OCI API credentials in the ServiceNow AI Platform to connect with the OCI user you want to use for discovery. For more information, see Create Oracle API credentials.
- You're using Discovery Admin Workspace v1.13.0.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **OCI** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters pro vide failover protection and load balancing be tween MID Servers. See Configure a MID Serv for more information.
Automatically select a MID server	An available MID Server is automatically se lected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validat ed, up, and have the Cloud Management ca pability are displayed.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your in stance. Once an account is selected, the **Ac count name** and **Account ID** fields auto-popu late.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
New cloud account	Add a new cloud account to your instance by entering an account name, compartment ID, and Datacenter URL. **Note:** ○ You can access the compartment ID, also referred to as the Oracle Cloud Identifier (OCID) of the compartment, in the Oracle Cloud console via **Identity & Security** > **Compartments**. ○ Example URLs: ▪ For commercial cloud: `https://{service}.ap-mumbai-1.oraclecloud.com` ▪ For GovCloud: `https://{service}.us-gov-ashburn-1.oraclegovcloud.com`

Even if the cloud account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

9. Select **Discover all Projects** or **Discover specific Projects**.

   Note: This option displays only when the target cloud account is an Oracle compartment account.

   a. If discovering specific projects, select **Add projects**.
   The Discoverable Projects dialog displays.

   b. Select the check box next to one or more corresponding names and then select **Add**.

10. Select **Next**.
11. Select cloud regions for the discovery schedule.

    You can choose to include all regions linked to the cloud account or select specific regions. There's no restriction on the number of cloud regions that you can add.

    a. **Optional:** To discover specific cloud regions, select the check box for the cloud region you want to scan from the **Available cloud regions** list.

    b. **Optional:** If the cloud regions don't appear, select **Refresh cloud regions**.

12. Select **Next**.
13. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.

14. Select **Next**.
15. Configure the Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the schedule runs.
Run on demand	The schedule only runs when triggered manually. The schedule can be triggered by selecting **Finish and run** at the end of this setup, or by navigating to the schedule in the Discovery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Discovery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Finish and run	After being selected, all the information provided is validated. Then, a Discovery schedule is created in the background, a Discovery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information provided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_schedule] table. You can edit the schedule information or run it by selecting **Discover Now**.

Create an OpenStack Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create an OpenStack Discovery schedule.

Before you begin

Verify the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have configured an OpenStack user with the required roles for the project so that the MID Server can access the selected cloud resources you plan to discover. You can create this user through the Horizon dashboard, the OpenStack CLI, or directly through the Keystone API. All three methods write the user to the Keystone identity service.
- You have created suitable credentials in the ServiceNow AI Platform to connect with the OpenStack user that you want to use for discovery. For more information, see OpenStack resource discovery.
- You're using Discovery Admin Workspace v1.13.0 or later.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **OpenStack** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters pro vide failover protection and load balancing be tween MID Servers. See Configure a MID Serv for more information.
Automatically select a MID server	An available MID Server is automatically se lected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validat ed, up, and have the Cloud Management ca pability are displayed.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your in stance. Once an account is selected, the **Ac count name**, **OpenStack Account ID**, and **Cloud URL** fields auto-populate.
New cloud account	Add a new cloud account to your instance by entering an account name, ID, and Cloud URL. **Note:** ○ You can access the Cloud URL in Hori zon via **Project** > **API Access** ○ Example Cloud URL (OpenStack iden tity service endpoint): `https://<openstack dns>:<port>`

Even if the cloud account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

9. Select **Next**.
10. Select cloud regions for the discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can choose to include all regions linked to the cloud account or select specific regions. There's no restriction on the number of cloud regions that you can add.

a. **Optional:** To discover specific cloud regions, select the check box for the cloud region you want to scan from the **Available cloud regions** list.

b. **Optional:** If the cloud regions don't appear, select **Refresh cloud regions**.

11. Select **Next**.
12. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.
    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.
13. Select **Next**.
14. Configure the Discovery schedule.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the schedule runs.
Run on demand	The schedule only runs when triggered manually. The schedule can be triggered by selecting **Finish and run** at the end of this setup, or by navigating to the schedule in the Discovery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Discovery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Finish and run	After being selected, all the information provided is validated. Then, a Discovery schedule is created in the background, a Discovery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information provided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_schedule] table. You can edit the schedule information or run it by selecting **Discover Now**.

Create an oVirt Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create an oVirt Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Verify the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have configured an oVirt user in the oVirt Administration Portal with the required roles and permissions. This configuration enables the MID Server to access the cloud resources you plan to discover.
- You have created suitable credentials in the ServiceNow AI Platform to connect with the oVirt user that you want to use for discovery. For more information, see Oracle Linux Virtualization Manager and Red Hat Virtualization discovery.
- You're using Discovery Admin Workspace v1.13.0 or later.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **oVirt** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters provide failover protection and load balancing between MID Servers. See Configure a MID Server cluster for more information.
Automatically select a MID server	An available MID Server is automatically selected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validated, up, and have the Cloud Management capability are displayed.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your instance. Once an account is selected, the **Account name** and **Manager URL** fields auto-populate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
New cloud account	Add a cloud account to your instance by en tering an account name and Manager URL. **Note:** ○ The Manager URL is the URL you use to access the oVirt Administration Por tal. ○ Example: `https://my-manager-instance.com:433`

Even if the cloud account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

9. Select **Next**.
10. Select datacenters for the discovery schedule.

    You can choose to include all datacenters linked to the cloud account or select specific datacenters. There's no restriction on the number of datacenters that you can add.

    a. **Optional:** To discover specific datacenters, select the datacenters from the **Select the data centers** drop-down list.

    b. **Optional:** If the datacenters don't appear, select **Refresh data centers**.
11. Select **Next**.
12. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.
13. Select **Next**.
14. Configure the Discovery schedule.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the sched ule runs.
Run on demand	The schedule only runs when triggered manu ally. The schedule can be triggered by select ing **Finish and run** at the end of this setup, or by navigating to the schedule in the Discov ery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Dis covery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Finish and run	After being selected, all the information provided is validated. Then, a Discovery schedule is created in the background, a Discovery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information provided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_schedule] table. You can edit the schedule information or run it by selecting **Discover Now**.

Discovery for AWS

Amazon Web Services (AWS) cloud discovery enables visibility to your AWS cloud resources, to populate and update the Configuration Management Database (CMDB). Visibility into AWS supports business outcomes such as cloud transformation and optimizing efficiency for operations (ITOM/ITSM/AIOps).

What is AWS cloud discovery

AWS cloud discovery is an automated process that continuously identifies and maps AWS resources and populates the data in the Configuration Management Database (CMDB).

AWS discovery can be performed by a combination of approaches such as agent-based or agentless, cloud APIs for metadata discovery, or IPs for OS-level discovery. The visibility can be provided near real time (using event-based discovery for example) or by timed discovery schedules.

The Discovery and Service Mapping apps perform discovery by methods refereed to as horizontal discovery and top-down mapping. Horizontal discovery identifies configuration items (CIs) without dependency mapping. Top-down service mapping identifies application dependencies, connection paths, and service impact.

Key outcomes and business value

AWS discovery facilitates several vital business outcomes by populating the CMDB with essential cloud data:

- Regulatory compliance enabled by the data support. Visibility can promote alignment with compliance frameworks such as the Mutual Recognition Agreement (MRA) or Digital Operational Resilience Act (DORA).
- Software asset management (SAM) enabled by visibility into cloud software deployments.
- Financial operations(FinOps) enabled by comprehensive visibility into AWS resources and their usage,
- Security operations(SecOps) enabled by continuous visibility into cloud resources and their configurations.
- Certificate management enabled by the discovery of certificates, their expiry, and usage.
- Artificial intelligence for IT operations (AIOps) enabled by Identifying and mapping all cloud resources, and their configurations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS discovery approaches

There are several approaches for discovering AWS environments.

1. Cloud metadata discovery: Provides a high-level view of the AWS cloud infrastructure.
2. Cloud OS-Level discovery: Provides a deeper level of discovery that indicates the state of the AWS cloud resources, such as installed software, active services, running processes, and system configurations.
3. Event-driven cloud discovery: Tracks changes in the life-cycle state or the configuration of AWS cloud resources. For more information, see AWS events-driven discovery
4. Collecting data with AWS Systems Manager: Provides a streamlined, agent-based approach to discovering Amazon Elastic Compute Cloud (EC2). For more information, see AWS SSM discovery
5. Collecting data with Agent Client Collector (ACC-VC): Performs horizontal IP-based discovery for OS-related attributes such as system configurations, network interfaces, and running process. For more information, see Agent Client Collector Discovery.
6. Collecting data with Service Graph Connectors: Imports and integrates AWS data into CMDB and non-CMDB tables. Specializes in collecting the data for AWS Organizations. For more information, see AWS discovery solutions comparison and Service Graph Connector for AWS .

For comparison of AWS cloud discovery methods and requirements, see AWS cloud discovery methods and use cases

How to perform AWS cloud discovery

Multiple ITOM Visibility apps can collect (or discover) your data, visualize it, and help you monitor your AWS resources.

- Discovery Admin Workspace
- Discovery for Amazon Elastic Kubernetes Service (EKS)
- Service Mapping
- Discovery and Service Mapping Patterns
- Certificate Inventory and Management

Enabling Discovery or other Visibility solutions to access your AWS infrastructure depends on roles and permissions configured both in AWS and in ServiceNow AI Platform. The discovery process requires configuration within AWS, like setting up Identity and Access Management roles.

AWS user discovery permissions

AWS Users	Discovery permissions
AWS Organizations with master and member accounts	Access is based on the IAM roles defined for the master and member accounts.
AWS account root user	Has complete access to all AWS services and resources in the account.
IAM users/IAM user group	Has access to specific resources and services based on IAM roles or temporary access based on assumed roles.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Access to cloud environments for ITOM products

In the ServiceNow AI Platform side, there are user configurations needed if you choose to use Discovery.

- You must configure the discovery_admin role for a user, to be able to run the discovery. For more information, see Managing roles
- Discovery runs commands and API queries to access and discover your AWS infrastructure. Before starting to configure Discovery roles and permissions, review the Cloud discovery spreadsheet and verify the REST API permissions.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

AWS cloud discovery methods and use cases

Comparison of use cases and requirements for cloud discovery methods in AWS.

ITOM Visibility offers several methods for AWS cloud discovery. Each method supports different use cases and has its own requirements.

AWS cloud discovery use cases

Comparison of AWS cloud discovery methods and use cases

Use case	Cloud Discovery (Pattern-based)	Horizontal discovery with AWS Systems Manager (AWS SSM)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
IT Service Management (ITSM) - basic and generic virtual CI classes	Yes	Yes	Yes	No	No	No
ITSM - extended virtual CI classes	Yes	Yes	No	No	No	No
ITSM - hardware CI classes	No	Yes	Yes	Yes	Yes	Yes
Pattern discovery of applications	No	Yes	No	No	Most	Yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison of AWS cloud discovery methods and use cases (continued)

Use case	Cloud Discovery (Pattern-based)	Horizontal discovery with AWS Systems Manager (AWS SSM)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
Tag-based service mapping	Yes	Yes	Yes	No	No	No
Machine Learning (ML)-based service mapping	No	Some	Some	Yes	Yes	Yes
Top-down service mapping	No	No	No	No	Most	Yes
Software Asset Management (SAM) - package manager	No	Yes	Yes	Yes	Yes	Yes
SAM - file-based	No	No	No	No	Yes	Yes
Oracle Global License Advisory Services (GLAS)	No	No	No	No	Yes	Yes
Certificate discovery	No	No	No	No	Yes	Yes, without host relationships

AWS cloud discovery requirements

Comparison of AWS cloud discovery methods and requirements

Requirement	Cloud Discovery (Pattern-based)	Horizontal discovery with AWS SSM	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
MID Server	Yes	Yes	No	No	Yes	Yes
Cloud credential	No	No	Yes	No	No	No
OS credential	No	No	No	No	No	Yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison of AWS cloud discovery methods and requirements (continued)

Requirement	Cloud Discovery (Pattern-based)	Horizontal discovery with AWS SSM	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
ServiceNow agent	No	No	No	Yes	Yes	No
AWS SSM agent - on all virtual machines (VMs)	No	Yes	Yes	No	No	No
AWS S3 bucket	No	Yes	Yes	No	No	No
Network Connectivity Target VM to MID Server	No	No	No	No	Yes	No
Network Connectivity MID Server to Target VM	No	No	No	No	No	Yes

AWS events-driven discovery

The Amazon Web Services (AWS) Config service can raise events for any changes in the life-cycle state or the configuration of a cloud resource. The ServiceNow® event-driven discovery uses the events to auto-update the latest resource information in the Configuration Management Database (CMDB).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview of the AWS events-driven discovery

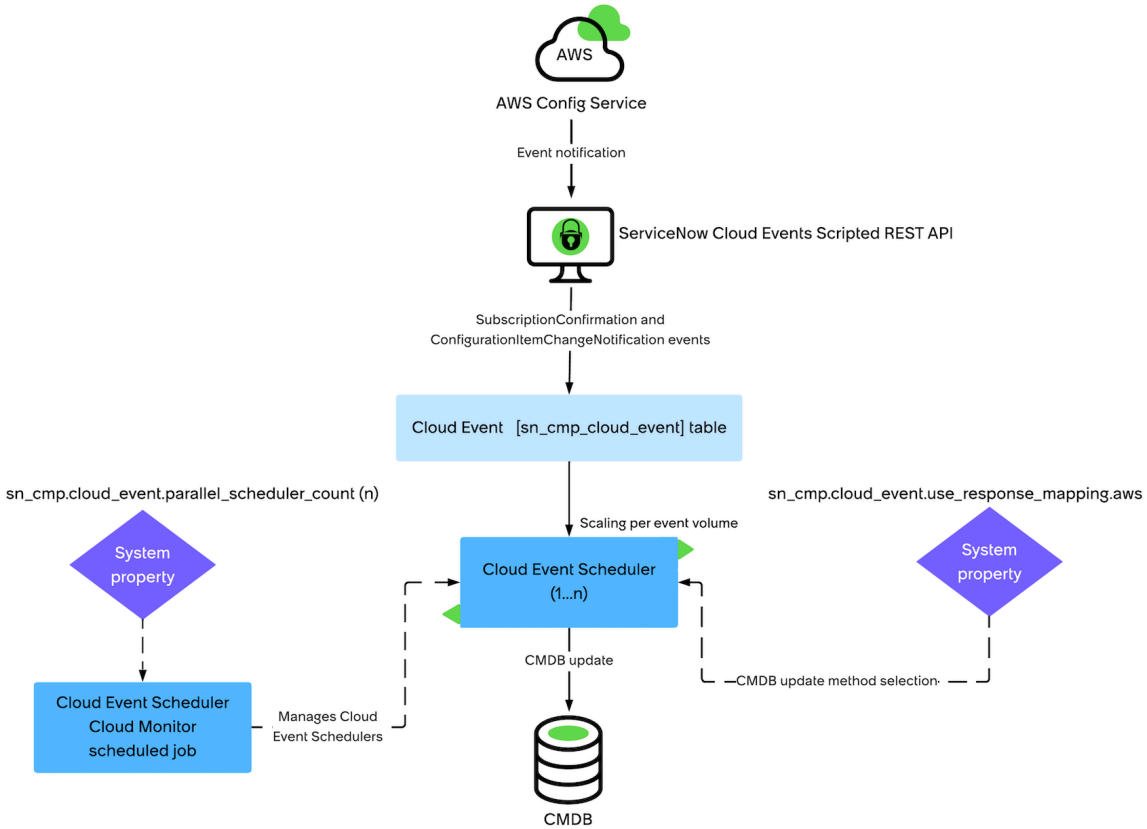

Configure the AWS Config service to send Simple Notification Service (SNS) notifications to the ServiceNow instance. For more information, see Configure AWS Config event notifications. After you configure the service, it starts sending event notifications to the ServiceNow Cloud Events REST API. When an event meets any one of the following criteria, the API writes it to the Cloud Events [sn_cmp_cloud_event] table:

- The *Type* of the event is *SubscriptionConfirmation*.
- The *Type* of the event is *Notification* and *messageType* is *ConfigurationItemChangeNotification*.
- Amazon CloudWatch has raised the event for a change in the tag associated with the Configuration Item (CI).

The Cloud Event Scheduler then picks the events in the Ready state for batch processing. During event processing, the event-driven discovery uses response mappings or patterns to update the details of the affected resource in the CMDB. The *sn_cmp.cloud_event.use_response_mapping_aws* property determines the CMDB update method. To understand the status of an event, review its state in the Cloud Events [sn_cmp_cloud_event] table.

Starting with the Brazil release, the *sn_cmp.cloud_event.use_response_mapping_aws* property is set to True by default. When this property is set to True, and suitable response mappings are available, the event-driven discovery uses the response mappings to create or update the CI in the CMDB. Otherwise, the event-driven discovery triggers the appropriate patterns to discover the affected resource and create or update the CI in the CMDB.

Starting with the Brazil release, use the *sn_cmp.cloud_event.parallel_scheduler_count* property to scale the Cloud Event Scheduler per the event inflow rate. Running multiple cloud event schedulers in parallel

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

helps to improve the event processing rate of the instance. For more information on scaling the Cloud Event Schedulers, see Scale the AWS cloud event schedulers.

During event processing, the Cloud Event Scheduler identifies the domain of the associated service account and assigns it to the event. If domain identification fails before processing, the event can remain unassigned and be visible across all domains. To limit the visibility of failed events to the service-provider administrator, set the *sn_cmp.error_events.default_domain* property to the sys_id of the service-provider domain.

Configure AWS Config event notifications

Configure the Amazon Web Services (AWS) Config service to send event notifications to the ServiceNow instance for any changes in the lifecycle state of a resource.

Before you begin

- Verify that the Discovery (com.snc.discovery) plugin is activated in the instance.
- Verify that you have valid AWS subscriptions (service accounts) and its associated logical datacenters are discovered.
- Verify that the user account password used to subscribe the instance to the Simple Notification Service (SNS) does not contain the @ or # characters.
- Verify that the AWS Config recorder is properly configured with continuous recording enabled to prevent events from being nested with a "detail" JSON node.

Roles required:

- Verify that an AWS role is available that can access the following services and resources:
  - SNS
  - AWS Config service
  - Resource types for which you want to track the configuration change
- ServiceNow roles:
  - discovery_admin
  - sn_cmp.cloud_event_integration: The access credentials of a ServiceNow user with the sn_cmp.cloud_event_integration role is required to subscribe the instance to the SNS notifications. For more information, see Create a user and Assign a role to a user.

About this task

ServiceNow® event-driven discovery uses the events to update the latest resource information in the Configuration Management Database (CMDB). For more information, see AWS events-driven discovery.

Many of the steps in the topic are performed in the AWS portal. For more information, see the following AWS documents:

- Set Up Amazon SNS Notifications
- Viewing the AWS Config dashboard

Note:

If you're using domain separation for Cloud Discovery, the events are also domain-separated. Therefore, you can view the details of a processed event only if it belongs to your domain. If an event isn't associated with any service account, then it's associated with the global domain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

During event processing, the Cloud Event Scheduler identifies the domain of the associated service account and assigns it to the event. If domain identification fails before processing, the event can remain unassigned and be visible across all domains. To limit the visibility of failed events to the service-provider administrator, set the *sn_cmp.error_events.default_domain* property to the sys_id of the service-provider domain.

Procedure

1. Log in to the AWS account.
2. On the Services page, navigate to **AWS Services** > **Find Services** > **Simple Notification Service**.
3. Create an Amazon SNS topic.
   The AWS Config service uses the SNS topic to publish the event notifications.

   a. On the SNS dashboard, navigate to **Topics** > **Create topic**.

   b. On the form, fill in the fields.

   Create topic form

Field	Description
Type	Notification topic type. Select the **Standard** topic type. **Important:** Do not select FIFO. FIFO topics only support the SQS subscription protocol. Standard is required because the ServiceNow subscription uses HTTPS.
Name	Name of the SNS topic.
Display name - optional	Display name of the SNS topic.

   c. Select **Create topic**.

   For more information, see Creating an Amazon SNS topic.

4. Subscribe the instance to the SNS topic.
   After subscribing to the SNS topic, the instance can receive the event notifications from the AWS cloud.

   a. On the Topic Details page, select **Create subscription**.

   b. On the form, fill in the fields.

   Create subscriptions form

Field	Description
Protocol	Communication protocols to use. For the subscription that auto-updates the CMDB, select HTTPS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: After creating the HTTPS subscription, you can create a separate subscription that sends emails to a specified person.

Endpoint	URL of the ServiceNow Cloud Events REST API.
	
	Ensure that the URL adheres to the following syntax:
	`https://<username>:<user_password>@<instance_URL>/api/now/cloud_event`
	
	Use the alternate endpoint if there is any expectation of high load from AWS. High load on default endpoint can significantly slow down the instance as default endpoint is also used for other transactions.
	
	Ensure that the URL adheres to the following syntax:
	`https://<username>:<user_password>@<instance_URL>/api/now/cloud_event?sysparm_rest_integration_pool=true`

c. Select **Create subscription**.

For more information, see Subscribing to an Amazon SNS topic.

5. Configure AWS Config data and delivery channel settings.

a. Navigate to **AWS Config** > **Settings**.

b. In the Data and delivery section, select **Edit**.

c. On the Edit data and delivery channel settings form, fill in the fields.

Edit data and delivery channel settings form

Field	Description
Data retention period	Length of time to retain AWS Config data. Select **Retain AWS Config data for 7 years** or set a custom retention period.
Amazon S3 bucket	S3 bucket for storing configuration history and snapshots. Select one of the following options: ▪ **Create a bucket** ▪ **Choose a bucket from your account** ▪ **Choose a bucket from another account**
Amazon SNS topic	SNS topic for streaming configuration changes and notifications. Select the **Stream configuration changes and notifications to an Amazon SNS topic** check box. Select **Choose a topic from your account** and select the SNS topic created in 3.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Select **Save**.

For more information, see Updating the delivery channel.

6. Configure AWS Config recorder settings.

a. On the Settings page, in the Recorder section, select **Edit**.

b. On the Edit customer managed recorder settings form, fill in the fields.

Edit customer managed recorder settings form

Field	Description
Enable recording	Turns the recorder on or off. Select the **Enable recording** check box.
Recording strategy	Determines which resource types to record. Select one of the following options: ▪ **All resource types with customizable overrides**: Records all current and future supported resource types in the Region. You can override the recording frequency for specific resource types or exclude specific resource types entirely. ▪ **Specific resource types**: Records only the resource types you specify.
Recording frequency	How often to record configuration changes. Select one of the following options: ▪ **Continuous recording**: Records configuration changes whenever a change occurs. **This option is required for proper event processing in ServiceNow.** ▪ **Daily recording**: Records configuration data once per day if a change occurred.
Override settings	Optional overrides for specific resource types. You can override the recording frequency for specific resource types or exclude specific resource types from recording. You can add up to 100 frequency overrides and 596 exclusions.
IAM role for AWS Config	IAM role that AWS Config uses to access other AWS services. Select one of the following options: ▪ **Use an existing AWS Config service-linked role**: Uses a predefined role that includes the permissions AWS Config requires. ▪ **Choose a role from your account**: Uses a pre-existing IAM role from your account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Select **Save**.

For more information, see Managing the configuration recorder.

Warning:

Ensure that the recorder settings are configured correctly with continuous recording enabled. If the AWS Config recorder is not properly configured, AWS Config events may be nested with a "detail" JSON node, which prevents ServiceNow from processing the events.

Trouble?

Symptom: Events from AWS Config are not being inserted into the ServiceNow instance.

Cause: The AWS Config recorder is not properly configured. Without correct recorder settings, event information from AWS is nested with a "detail" JSON node, and the ServiceNow event processing logic cannot detect the events.

Resolution: Verify the recorder settings on the AWS Config Settings page. Ensure that:

- The **Enable recording** check box is selected.
- The recording frequency is set to **Continuous recording**.
- The SNS topic is correctly configured in the data and delivery channel settings.

What to do next
After some events are generated, navigate to the Cloud User Portal to view the events.

Scale the AWS cloud event schedulers

Scale the cloud event schedulers to improve the Amazon Web Services (AWS) event processing rate of the ServiceNow instance.

Before you begin
Role required: cloud_admin

About this task
If the instance receives a huge number of events and it doesn't process them within a reasonable time, create additional Cloud Event Schedulers. Running multiple cloud event schedulers in parallel helps to improve the event processing rate of the instance. When the event inflow rate reduces, you can deactivate some of the Cloud Event Schedulers. For more information, see AWS events-driven discovery.

Procedure
1. Configure the number of Cloud Event Schedulers that you want to use.

   a. Navigate to **All** > **System Properties** > **All Properties**.

   b. In the Name column, search for *sn_cmp.cloud_event.parallel_scheduler_count*.

   c. On the *sn_cmp.cloud_event.parallel_scheduler_count* property form, set the **Value** field to an appropriate value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

- Make sure that the value of the *sn_cmp.cloud_event.parallel_scheduler_count* property doesn't exceed the count of available worker nodes.
- If you have multiple domains where cloud events are configured, use the *sn_cmp.cloud_event.parallel_scheduler_count* property with caution because each cloud event scheduler has multiple iterations for each domain in which cloud events are active.

d. Select **Update**.

2. Execute the Cloud Event Scheduler Count Monitor scheduled job.
   The Cloud Event Scheduler Count Monitor scheduled job manages the count of Cloud Event Schedulers per the value specified in the *sn_cmp.cloud_event.parallel_scheduler_count* property.
   a. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
   b. Open the Cloud Event Scheduler Count Monitor scheduled job.
   c. Select **Execute Now**.

Create a database index for the Cloud Events table

Create a database index from the specified columns of the Cloud Events [sn_cmp_cloud_event] table to improve the AWS cloud event processing performance.

Before you begin

Ensure that the scope is set to Cloud Provisioning and Governance.

Role required: admin

About this task

After upgrading Cloud Provisioning and Governance to the Utah release, create a database index from the specified columns of the Cloud Events [sn_cmp_cloud_event] table. The database index helps to improve the Amazon Web Services (AWS) event processing performance of the instance. If you've started using Cloud Provisioning and Governance with the Utah release, the database index is automatically created on application installation.

Instead of manually creating the database index, you can use a scheduled job to create it. For more information, see ServiceNow® events-driven discovery: Create a database index for the Cloud Events table (KB1207069).

You can also Scale the cloud event schedulers to improve the Amazon Web Services (AWS) event processing rate of the ServiceNow instance. For more information, see Scale the AWS cloud event schedulers.

Procedure

1. Navigate to `https://<instance name>.service-now.com/sn_cmp_cloud_event_list.do`
2. Select **Column options** > **Configure** > **Table**.
3. Select the **Database Indexes** related list.
4. Select **New**.
5. From the available columns list, select the following columns in the stated order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. **State**

b. **Process Id**

c. **Resource ID**

6. Select **Create Index**.
7. **Optional:** Select an option to specify whether you want to receive a notification on index creation.
8. Select **OK**.
9. Rename the new index for easy identification.

AWS SSM discovery

AWS Systems Manager (SSM) Agent discovery introduces a streamlined, agent-based approach to discovering Amazon Elastic Compute Cloud (EC2) using AWS SSM. This integration enhances Discovery by leveraging SSM agents to reduce dependency on traditional MID Server configurations, simplify credential management, and improve scalability across multi-region environments.

Workflow

The overall process from a high level is as follows:

1. The ServiceNow AI Platform® sends discovery commands to the MID Server.
2. The MID interacts with AWS services (SSM, Simple Storage Service (S3), Parameter Store) to execute commands on target devices.
3. The SSM agents run the commands and return results to S3.
4. The MID retrieves and processes results.
5. The MID sends the results back to the ServiceNow AI Platform® via the ECC queue, which updates the CMDB.

Benefits and usage

The following examples highlight the primary advantages and practical uses of AWS SSM Agent discovery:

- Execute discovery without needing additional credentials local to the operating system.
- Simplify deployment without the need of Agent Client Collector (ACC) or Virtual Private Cloud (VPC) access.
- Minimize the need for multiple MID Servers and direct network access to target devices.
- Securely manage credentials and command execution using AWS services.

Requirements

Confirm that you have the required versions of the following applications and plugins:

- Discovery
- Cloud Discovery Workspace version 1.7.1 or later.
- CMDB CI Class Models version 1.74.0 or later.
- Discovery and Service Mapping Patterns version 1.27.0 or later.

Install the MID Server. For more information, see Install and configure the MID Servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify that you have an AWS user account with administrative access.

Verify that you have a ServiceNow AI Platform® user account with the discovery_admin role.

Unsupported features

Currently, the following features don't support AWS SSM discovery:

- File-based discovery
- Certificate-based discovery
- Top-down discovery
- Enhanced ADM
- Change/Unchange user step in patterns

Note: SSM supports only sudo for privileged command execution and defaults to the sh shell, with no support for alternate command or shell types.

AWS environment configuration

Create IAM roles and permissions

Define Identity and Access Management (IAM) roles to support AWS SSM operations, specifying the required permissions for executing Read, Write, and List commands. For more information, see IAM Policies.

Configure EC2 roles and instances

Create EC2 roles and assign the necessary permissions for these roles to interact with SSM. For more information about setting up an EC2 instance, see EC2 Instances.

Create S3 Buckets

Create S3 buckets to support large data transfers and configure appropriate bucket policies and life cycle rules. To overcome SSM's 24,000-character output limit, command output is redirected to S3, enabling full payload capture. Additionally, S3 facilitates file transfers to EC2 instances. For more information about creating S3 buckets, see Setting up an S3 Bucket.

(Optional) Configure KMS keys

Create a custom AWS Key Management Service (KMS) key to encrypt sensitive credentials stored as SecureString parameters in the AWS Systems Manager Parameter Store. Using a dedicated key enhances security by verifying that credentials are encrypted and securely retrieved at runtime, without exposing them in plain text. For more information about creating KMS keys, see Create a Customer Managed AWS KMS Key.

Important: If you're using applicative credentials in Discovery, you must create a custom KMS key.

Import custom SSM documents

SSM command execution depends on the availability of required documents. Before initiating SSM-based discovery, verify that all YAML files provided by ServiceNow have been successfully deployed to each AWS region where the discovery process will run. For more information on this process, see Import Custom AWS Documents. You can download the YAML files directly from this article.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information on AWS Management Console configuration, see the Amazon SSM Discovery - AWS Environment Setup Instructions article in the Now Support Knowledge Base.

ServiceNow AI Platform® instance configuration

Configure System Properties

Enable the following Discovery system properties:

- *glide.discovery.enable_ssm*
- *glide.discovery.ssm.enable_windows*

The AWS SSM Agent runs with root (Linux) or SYSTEM (Windows) privileges, meaning any command sent through it can execute with full system access. Due to this high level of access, SSM is turned off by default for security reasons. For more information, see Enable AWS SSM-based discovery.

Define root and non-root credentials

The MID Server property *mid.discovery.aws_ssm.linux.fallback_root_user* enables EC2 instances to default to the root user. By default, this setting is **false**, but if you don't want to set up alternate user credentials, you can enable it to run commands as root. If you leave this setting turned off, you must create a credential record in the AWS SSM Instance Users [aws_ssm_instance_user_credentials] table, where you simply provide the user name that should be used to run commands on the instance. For more information, Enable root fallback and Configure custom user credentials.

Configure MID Server properties

After you set up KMS keys or S3 buckets in the AWS Management Console, configure the following MID Server system properties on the ServiceNow AI Platform®:

- *mid.discovery.aws_ssm.kms_key_name*
- *mid.discovery.aws_ssm.kms_key_region*
- *mid.discovery.aws_ssm.s3_bucket_name*
- *mid.discovery.aws_ssm.s3_bucket_region*

For more information, see Configure MID Server for AWS S3 access and Configure MID Server for AWS KMS access.

Enable MID Server capability

You must enable the new MID Server capability for AWS SSM so the MID Server can support running SSM discovery. By default, this capability is included if you're using the **ALL** capability. However, if you're using individual capabilities, you must manually add the new **AWS SSM** capability to your MID Servers. For information about adding a capability, see Configure MID Server capabilities.

Cloud Discovery schedule configuration

Create a Cloud Discovery schedule

SSM discovery can be deployed within an AWS-based Cloud Discovery schedule in the Discovery Admin Workspace. To do this, go to the Deep discovery step, enable the **Create an IP-based discovery schedule** toggle, and select **Discover servers though AWS Systems Manager (SSM) Agent** as the Discovery method.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Create an AWS Discovery schedule in Discovery Admin Workspace.

Enable AWS SSM-based discovery

Configure Discovery to use AWS Systems Manager (SSM) Agent.

Before you begin

Role required: discovery_admin

About this task

The AWS SSM Agent runs with root (Linux) or SYSTEM (Windows) privileges, meaning any command sent through it can execute with full system access. Due to this high level of access, SSM is turned off by default for security reasons. On Linux, you can configure non-root credentials to run Discovery commands more securely. On Windows, Discovery commands require SYSTEM-level access, so you must decide whether to enable this capability based on your security requirements. Enabling *glide.discovery.enable_ssm* doesn't automatically activate SSM discovery on Windows. Unlike Linux, Windows lacks support for using a non-SYSTEM user, so for security reasons, you must explicitly enable the *glide.discovery.ssm.enable_windows* property to use AWS SSM discovery on Windows devices.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. In the **Name** field, enter `glide.discovery.enable_ssm`.
3. Double-click the **Value** field to perform inline editing.
4. Enter `true`.
5. Select the check mark ( ) icon.
6. **Optional:** To enable AWS SSM-based discovery on Windows machines, repeat steps 2-5 for the *glide.discovery.ssm.enable_windows* property.

What to do next

Define fallback root behavior using the *mid.discovery.aws_ssm.linux.fallback_root_user* property or create non-root credentials to execute Discovery commands using the AWS SSM Instance Users [aws_ssm_instance_user_credentials] table. For more information, see Enable root fallback and Configure custom user credentials.

Configure MID Server for AWS S3 access

Configure MID Server properties to enable Simple Storage Service (S3) access from the ServiceNow AI Platform® when performing Systems Manager (SSM) Agent discovery.

Before you begin

Confirm the following:

- You have at least one MID Server installed and in the **Up** state. For more information, see Install and configure MID Servers to access cloud environments.
- You have an AWS credential that allows the MID Server to access and make API calls into the AWS environment. For more information, see AWS Access for MID Server.
- You created a custom S3 bucket with necessary access enabled for the IAM roles. For more information, see Setting up an S3 Bucket and Setting up a Policy for accessing the S3 Bucket.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin

About this task

Because the SSM Run Command console has a 24,000-character limit, results that exceed this limit must be stored in and retrieved from an S3 bucket. For Put File commands, the S3 bucket also serves as a transfer point between the MID Server and the AWS environment, enabling file delivery to target devices via SSM. To configure the MID Server to support these S3-based operations, set the *mid.discovery.aws_ssm.s3_bucket_name* property to the name of the S3 bucket used for SSM Agent discovery. Then, specify the bucket's region in the *mid.discovery.aws_ssm.s3_bucket_region* property.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** field, enter `mid.discovery.aws_ssm.s3_bucket_name`.
3. Double-click the **Value** field to perform inline editing.
4. Enter the name of the S3 bucket.
5. Select the check mark (✓) icon.
6. In the **Name** field, enter `mid.discovery.aws_ssm.s3_bucket_region`.
7. Double-click the **Value** field to perform inline editing.
8. Enter the name of the region where the S3 bucket is located.
9. Select the check mark (✓) icon.

Configure MID Server for AWS KMS access

Configure MID Server properties to enable Key Management Service (KMS) key access from the ServiceNow AI Platform® when performing Systems Manager (SSM) Agent discovery.

Before you begin

Confirm the following:

- You have at least one MID Server installed and in the **Up** state. For more information, see Install and configure MID Servers to access cloud environments.
- You have an AWS credential that allows the MID Server to access and make API calls into the AWS environment. For more information, see AWS Access for MID Server.
- You created a custom AWS KMS key to encrypt sensitive credentials. For more information, see Create a Customer Managed AWS KMS Key.

Role required: discovery_admin

About this task

Using a dedicated key enhances security by verifying that applicative credentials are encrypted and securely retrieved at runtime, without being exposed in plain text. This step is optional and only necessary if you plan to use applicative credentials during discovery. To configure the MID Server to support secure credential handling during SSM Agent discovery, set the *mid.discovery.aws_ssm.kms_key_name* property to the name of the AWS KMS key. Then, set the *mid.discovery.aws_ssm.kms_key_region* property to the region where the key is located.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** field, enter `mid.discovery.aws_ssm.kms_key_name`.
3. Double-click the **Value** field to perform inline editing.
4. Enter the name of the KMS key.
5. Select the check mark (✓) icon.
6. In the **Name** field, enter `mid.discovery.aws_ssm.kms_key_region`.
7. Double-click the **Value** field to perform inline editing.
8. Enter the name of the region where the KMS key is located.
9. Select the check mark (✓) icon.

Enable root fallback

Enable the fallback root user property to allow AWS SSM discovery to run commands as root on Linux systems when alternate user credentials aren't configured.

Before you begin

Confirm that AWS SSM Agent discovery is enabled. For more information, see Enable AWS SSM-based discovery.

Role required: discovery_admin

About this task

When the *glide.discovery.enable_ssm* property is enabled, AWS SSM Discovery on Linux defaults to using the root user for command execution. If you prefer not to configure alternate user credentials on your EC2 instances, you must enable the *mid.discovery.aws_ssm.linux.fallback_root_user* MID Server property to allow fallback to root execution. When enabled, it automatically creates a record in the AWS SSM Instance Users [aws_ssm_instance_user_credentials] table with the user name root. However, if you leave this property set to **false**, you must configure non-root credentials to run Discovery commands. For more information, see Configure custom user credentials.

Procedure

1. Navigate to **All** > **MID Server** > **All Properties**.
2. In the **Name** field, enter `mid.discovery.aws_ssm.linux.fallback_root_user`.
3. Double-click the **Value** field to perform inline editing.
4. Enter `true`.
5. Select the check mark (✓) icon.

Result

A new record is added to the AWS SSM Instance Users [aws_ssm_instance_user_credentials] table with the user name "root".

Configure custom user credentials

Define a non-root user for AWS SSM discovery by creating a credential record that specifies a user name with sufficient privileges to execute discovery commands on Linux targets.

Before you begin

Confirm the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The system property *glide.discovery.enable_ssm* is turned on. For more information, see Enable AWS SSM-based discovery.
- The user you're assigning as the non-root user has access to run all the necessary commands on the target server. This user requires the same access as credentials used for regular Linux discovery, including the necessary root-level access for privileged commands. For more information, see Credentials required for host discovery and SSH credentials

Role required: discovery_admin

About this task

Instead of relying on root, you can define a custom user with sufficient privileges to execute the required Discovery commands. Only a user name is needed—no password or key—provided that the user has the same level of access as traditional Linux Discovery credentials. This approach promotes restricted access and better alignment with enterprise security policies.

Note: Currently, SSM supports only sudo for privileged command execution and defaults to the sh shell, with no support for alternate command or shell types.

Procedure

1. Navigate to **All** > **Discovery** > **Credentials**.
2. Select **New**.
3. Select **AWS SSM instance User**.
4. Enter a unique name for the user.
5. Enter the user name from the alternate user credential that you created on your EC2 instance in the AWS Management Console.
   For more information, see the Amazon SSM Discovery - AWS Environment Setup Instructions article in the Now Support Knowledge Base.
6. Select **Submit**.

Result

A new record is added to the AWS SSM Instance Users [aws_ssm_instance_user_credentials] table.

Create an AWS Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create an Amazon Web Services (AWS) Discovery schedule.

Before you begin

Verify the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have created suitable credentials in the ServiceNow AI Platform to connect with the service accounts on which you want to perform the discovery.
- Cloud Discovery has adequate permission to discover the resources of the target cloud. For more information on AWS cloud permissions, see Control AWS access and permissions using policies.
- You're using Discovery Admin Workspace v1.11.0 or later.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **AWS** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters pro vide failover protection and load balancing be tween MID Servers. See Configure a MID Serv for more information.
Automatically select a MID server	An available MID Server with an Identity and Access Management (IAM) role is automati cally selected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validat ed, up, and have the Cloud Management ca pability are displayed. **Note:** If you select this option, you can restrict the search to MID Servers with credential-less Discovery enabled by checking the corresponding box.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your in stance. Once an account is selected, the **Ac count name** and **Account ID** fields auto-popu late.
New cloud account	Add a new cloud account to your instance by entering an account name and ID.

9. Choose whether you want to **Assume an IAM role** or **Use permanent credentials**.

Option	Description
Assume an IAM role	The MID Server temporarily assumes an IAM role using a trust relationship. After selecting, provide the Access Role Name and account for access.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	Note: Before entering the Access Role Name and account, verify the following: ○ The IAM role exists on the target ac count, with the required permissions and policies for Cloud Discovery. ○ The IAM role must trust the MID Serv er's account so it can assume the role. If these prerequisites aren't met, verifi cation fails. It's best to confirm with your cloud admin that the role and trust rela tionship are properly configured before proceeding.
Use permanent credentials	The MID Server uses static access keys tied to an IAM user.

Even if the service account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

10. Select **Discover all member accounts** or **Discover specific member accounts**.

    Note: This option displays only when the target cloud account is an AWS management account.

    a. If discovering specific subscriptions, select **Add member accounts**.
    The Discoverable Accounts dialog displays.

    b. Select the check box next to one or more corresponding names and then select **Add**.

11. Select **Next**.
12. Select cloud regions for the discovery schedule.

    You can choose to include all regions linked to the cloud account or select specific regions. There’s no restriction on the number of cloud regions that you can add.

    a. **Optional:** To discover specific cloud regions, select the check box for the cloud region you want to scan from the **Available cloud regions** list.

    b. **Optional:** If the cloud regions don't appear, select **Refresh cloud regions**.

13. Select **Next**.
14. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a method from the **Discovery method** drop-down list.
    Discover servers with credentials uses operating system (OS)-level credentials to perform deep discovery on EC2 instances. **Discover servers with AWS Systems Manager (SSM)** uses AWS SSM Agent to perform deep discovery on EC2 instances. The system property

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glide.discovery.enable_ssm must be enabled to use SSM discovery. For more information, see Enable AWS SSM-based discovery.

b. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.

15. Select **Next**.
16. Configure the Discovery schedule.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the schedule runs.
Run on demand	The schedule only runs when triggered manually. Trigger the schedule by selecting **Finish and run** at the end of this setup, or by navigating to the schedule in the Discovery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Discovery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Finish and run	After being selected, all the information provided is validated. Then, a Discovery schedule is created in the background, a Discovery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information provided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_schedule] table. You can edit the schedule information or run it by selecting **Discover Now**.

Discovery for Microsoft Azure

If your cloud resources are in an Azure cloud, you must create a user identity called a service principal that grants permissions to the MID Server to access selected resources.

Azure management groups and subscriptions

An Azure management group contains other management groups and subscriptions. The management groups in an Azure Cloud environment form a hierarchy, but don’t contain volumes or virtual machines. Subscriptions contain cloud resources, such as virtual machines. The subscriptions that belong to management groups are called sub-accounts.

The advantages of using management groups are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Easy population of sub-accounts

After you configure the management group and supply the necessary credentials, you can test the account. If the test succeeds, Discovery returns a list of subscriptions in that management group. From this list, you can choose one or more subscription sub-accounts to include in the Discovery schedule using the management group. For more information on the hierarchy of management groups and subscriptions, see Organize your resources with Azure management groups

Discovery of sub-account resources using dynamically acquired credentials

When you run Discovery on your subscriptions, you do not need separate credentials for each sub-account. Discovery finds the credentials for the management group and maps them to all of the subscription sub-accounts. The Cloud Discovery process handles credentials automatically by acquiring a temporary credential for each sub-account via an Azure API. You can elect to use the default configuration or customize the MID Server to assume other roles for additional controls and security. In addition, Discovery can automatically refresh the list of sub-accounts and datacenters covered in a discovery schedule. For more information, see the KB article Retrieve newer accounts/sub-accounts automatically via Cloud Discovery.

A service principal for Azure cloud services is similar to a Microsoft Windows service account that enables Windows processes to communicate with each other within an Active Directory domain.

To create the Azure service principal in your ServiceNow instance, copy the service principal credential values from the Azure portal into a text editor, and then transfer those values into the instance.

The text file that you generate during this procedure

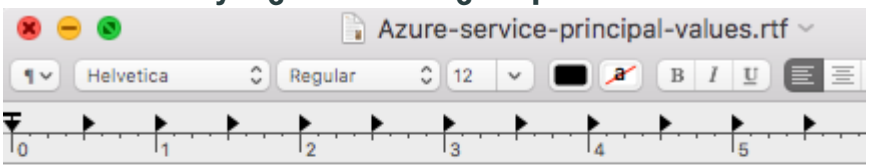

Directory ID: 8bcdd120-9979-491e-8683-d9a4550bdd

Application ID: e6821ea9-e5c5-428a-b50c-7ea62993b2d

Application key: d348050-f893ccb7-hg84yg8-rg200hd

Subscription ID: 4803be48c3ee89-96d2-4fed-bedb-22e49013ee89

This table shows you the Azure Service Principal value and the location in Azure where you can find the values you need for the credentials.

Cloud Provisioning and Governance setting	Azure Service Principal value	Location of the Azure value
Tenant ID	Azure **Directory ID** value from the text file.	**Azure Active Directory** > **Properties** > **Directory ID**
Client ID	Azure **Application ID** value from the text file.	**Azure Active Directory** > **App registrations** > **Registered App.Application ID**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Provisioning and Governance setting	Azure Service Principal value	Location of the Azure value
Secret Key	Azure **Application key** value from the text file.	**Azure Active Directory** > **App registrations** > **Registered App** > **Settings** > **Keys (hidden)**
Account ID	Azure **Subscription ID** associated with the Tenant ID.	**Azure Active Directory** > **Subscriptions** > **Subscription ID**

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected by Service Mapping during top-down discovery

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Service Mapping uses tag-based discovery to create service instance maps including the Cloud components. The Service Mapping application comes with the following preconfigured CI relationships used for tag-based discovery. These CI relationships are available from the 1.0.68 release on the ServiceNow Store.

CI	Relationship	CI
Configuration Item [cmdb_ci]	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]
Logical Datacenter [cmdb_ci_logical_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Azure Cloud Discovery API list

Logical Datacenter (cmdb_ci_logical_datacenter)

CI Attributes	Azure Attributes
object_id	response.name
name	response.name
region	response.displayName
status	Installed

Availability Zone (cmdb_ci_availability_zone)

CI Attributes	Azure Attributes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Availability Zone (cmdb_ci_availability_zone) (continued)

object_id	response.id
name	response.name
status	Installed/ Retired
state	

Resource Group (cmdb_ci_resource_group)

CI Attributes	Azure Attributes
object_id	id
name	name
state	available
status	Installed/ Retired

Network (cmdb_ci_network)

CI Attributes	Azure Attributes
object_id	id
state	
name	name
cidr	properties.addressSpace.addressPrefixes

Subnet (cmdb_ci_cloud_subnet)

CI Attributes	Azure Attributes
subnetName	response.name
subnetId	response.id
resourceGroup	response.properties.resourceGuid
networkId	response.id.split('/subnets/')
networkName	getNetwork(networkId)
cidrBlock	response.properties.addressSpace.addressPrefixes

Storage Volume (cmdb_ci_storage_volume)

CI Attributes	Azure Attributes
state	properties.provisioningState
storage_type	properties.BlobType
volume_id	id

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume (cmdb_ci_storage_volume) (continued)

name	name
size_bytes	properties.diskSizeGB * 1024 * 1024 * 1024
object_id	id
size	response.properties.diskSizeGB
"volume_container"	containerName
status	Installed/ Retired

Security Groups (cmdb_ci_compute_security_group)

CI Attributes	Azure Attributes
object_id	id
name	name
state	properties.provisioningState

Virtual Server (cmdb_ci_vm_instance)

CI Attributes	Azure Attributes
memory	properties.hardwareProfile.vmSize
state	The instance statuses: • succeeded: on • running: succeeded/stopping • deallocating: stopping/stopped • deallocated: off • terminated: error
object_id	id
cpus	properties.hardwareProfile.vmSize
disks	properties.storageProfile.dataDisks
nics	properties.networkProfile.networkInterfaces[].size
vm_inst_id	properties.vmId
name	name
status	Installed/ Retired

Hardware Template (cmdb_ci_compute_template)

CI Attributes	Azure Attributes
name	name

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hardware Template (cmdb_ci_compute_template) (continued)

object_id	name
vcpus	numberOfCores
memory_mb	memoryInMB
local_storage_gb	resourceDiskSizeInMB
cores	numberOfCores

Cloud Public IP Address (cmdb_ci_cloud_public_ipaddress)

CI Attributes	Azure Attributes
object_id	response.id
name	response.name
public_dns	properties.dnsSettings.fqdn
public_ip_address	properties.ipAddress

Cloud LB IP Address (cmdb_ci_cloud_lb_ipaddress)

CI Attributes	Azure Attributes
object_id	"properties.frontendIPConfigurations.properties.privateIPAddress OR properties.frontendIPConfigurations.properties.publicIPAddress, then call Public IP Address API"
name	"properties.frontendIPConfigurations.properties.privateIPAddress OR properties.frontendIPConfigurations.properties.publicIPAddress, then call Public IP Address API"
ipaddress_type	"properties.frontendIPConfigurations.properties.privateIPAddress ==> Private IP Address OR properties.frontendIPConfigurations.properties.publicIPAddress ==> Public IP Address"
status	Installed

Cloud Network Interfaces [cmdb_ci_nic]

CI Attributes	Azure Attributes
object_id	id
name	name
private_ip	properties.ipConfigurations
public_dns	call public ip address api - properties.dnsSettings.fqdn
state	properties.provisioningState
is_static	properties.ipConfigurations
mac_address	properties.macAddress

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Network Interfaces [cmdb_ci_nic] (continued)

public_ip	call public ip address api -

Image (cmdb_ci_os_template)

CI Attributes	Azure Attributes
name	response.id
object_id	response.name
guest_os	properties.storageProfile.osDisk.osType
image_source	id
status	Installed/ Retired

Cloud Storage Account cmdb_ci_cloud_storage_account

CI Attributes	Azure Attributes
name	name
object_id	id
sku_name	sku.name
state	properties.provisioningState

Load Balancer (cmdb_ci_cloud_load_balancer)

CI Attributes	Azure Attributes
object_id	response.id
name	response.name
state	
dns_name	properties.
fqdn	
canonical_hosted_zone_name	

Useful information

- Microsoft Azure discovery solutions comparison
- Set up Azure service accounts
- Install and configure MID Servers to access cloud environments
- Microsoft Azure Cloud discovery using patterns

Azure cloud discovery methods and use cases

Comparison of use cases and requirements for cloud discovery methods in Azure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM Visibility offers several methods for Azure cloud discovery. Each method supports different use cases and has its own requirements.

Azure cloud discovery use cases

Comparison of Azure cloud discovery methods and use cases

Use case	Cloud Discovery (Pattern-based)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
IT Service Management (ITSM) - basic and generic virtual CI classes	Yes	Yes	No	No	No
ITSM - extended virtual CI classes	Yes	No	No	No	No
ITSM - hardware CI classes	No	Yes	Yes	Yes	Yes
Pattern discovery of applications	No	No	No	Most	Yes
Tag-based service mapping	Yes	Yes	No	No	No
Machine Learning (ML)-based service mapping	No	Some	Yes	Yes	Yes
Top-down service mapping	No	No	No	Most	Yes
Software Asset Management (SAM) - package manager	No	Yes	Yes	Yes	Yes
SAM - file-based	No	No	No	Yes	Yes
Oracle Global License Advisory Services (GLAS)	No	No	No	Yes	Yes
Certificate discovery	No	No	No	Yes	Yes, without host relationships

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure cloud discovery requirements

Comparison of Azure cloud discovery methods and requirements

Requirement	Cloud Discovery (Pattern-based)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
MID Server	Yes	No	No	Yes	Yes
Cloud credential	Yes	Yes	No	No	No
OS credential	No	No	No	No	Yes
ServiceNow agent	No	No	Yes	Yes	No
Azure Monitor Agent/change tracking - on all virtual machines (VMs)	No	Yes	No	No	No
Azure Log Analytics workspace	No	Yes	No	No	No
Azure storage account	No	Yes	No	No	No
Network Connectivity Target VM to MID Server	No	No	No	Yes	No
Network Connectivity MID Server to Target VM	No	No	No	No	Yes

Create Azure cloud credentials

If your cloud resources are in an Azure cloud, create credentials that can access the Azure account. This procedure requires configuration in your Azure account.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Before you begin

Roles required:

- discovery_admin, service_mapping_admin, sn_cmp.cloud_admin roles in Cloud Provisioning and Governance or sn_cloud_ops_ws.cloud_ops_admin role in Cloud Discovery Workspace.
- Operations on the Microsoft Azure portal require one of the following roles:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Azure or Azure AD (Active Directory) Administrator
- Application Administrator
- Application Developer
- Cloud Application Administrator

and the Resource Policy Contributor role to create or modify resource policies.

- Enable internal network connection between the MID Servers and the Azure Cloud API endpoints: `management.azure.com`

Procedure

1. Log in to the Azure portal and navigate to **Azure Active Directory**.
2. Navigate to the **App registrations** section and select **New application registration**. Enter the following information for your application:

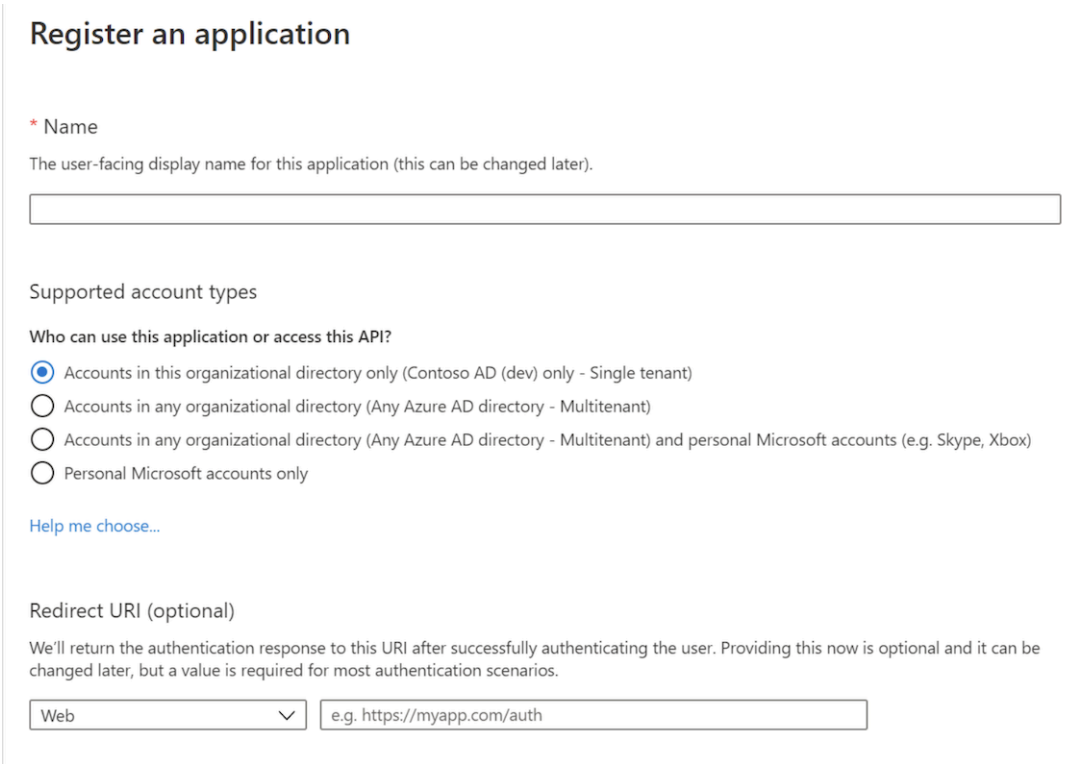

Field	Description
Name	Unique name for the application and its integration credentials. For example, `ServiceNow Integration`.
Supported account types	Specify who can use the application.
Redirect URI (Optional)	URL that will access Azure. Typically the URL of the ServiceNow instance.

3. Select **Register** to complete the app registration.
4. When registration completes, copy the Application (client) ID and Directory (tenant) ID values, and paste them in the text editor.
5. Label the values **Application ID** and **Directory ID** respectively.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. In the Azure portal, navigate to the **Certificates & secrets** section and **New client secret** then specify the following values:

Field	Description
Key description	Description for the key.
Duration	Expiration for the key. **Note:** Your organization may apply policies to restrict key durability. Select the appropriate duration.

7. Select **Add**.
8. Copy and paste the key value into the text editor and label the value **Application key**.
9. Enable the service principal to access Azure subscriptions based on your environment.

   Choose the option that matches your setup:

   - Management group: Use this option to grant Reader access to all subscriptions under the management group
   - Individual subscription: Use this option to grant Reader access to a specific subscription only

Option	Steps
Management group	**a.** In the Azure portal, navigate to **Management groups** and select the required management group. **b.** Select from the menu **Access Control (IAM)**. **c.** Select **+ Add** and then **Add role assignment**. **d.** In the **Role** field, select the **Reader** value. **Note:** The **Resource Policy Contributor** role is only required for provisioning. **e.** In the **Assign access to** field, keep the default **User, group, or service principal** value. **f.** In the **Select** field, select the name you created in the registering the application step. **g.** Select **Save**.
Individual subscription	**a.** In the Azure portal, navigate to **Subscriptions** and select the required subscription. **b.** Enter the subscription ID into the text editor and label it `Subscription ID`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Steps
	c. Select from the menu **Access Control (IAM)**. **d.** Select **+ Add** and then **Add role assignment**. **e.** In the **Role** field, select the **Reader** value. **Note:** The **Resource Policy Contributor** role is only required for provisioning. **f.** In the **Assign access to** field, keep the default **User, group, or service principal** value. **g.** In the **Select** field, select the name you created in the registering the application step. **h.** Select **Save**.

10. Select the **Azure Service Principal** type credential.
    a. Navigate to **All** > **Connections & Credentials** > **Credentials**.
    b. Select **New**.
    c. Select **Azure Service Principal**.
11. Specify the following values on the Azure Service Principal form:

Field	Value
Name	Name of the service principal to register with the instance. For example, `Azure service principal credentials`.
Authentication Method	Select **Client secret**. The **Secret key** field appears when you select **Client secret**. **Note:** **Client assertion** is not supported.

12. Copy and paste values from the temporary text file into the remaining fields.

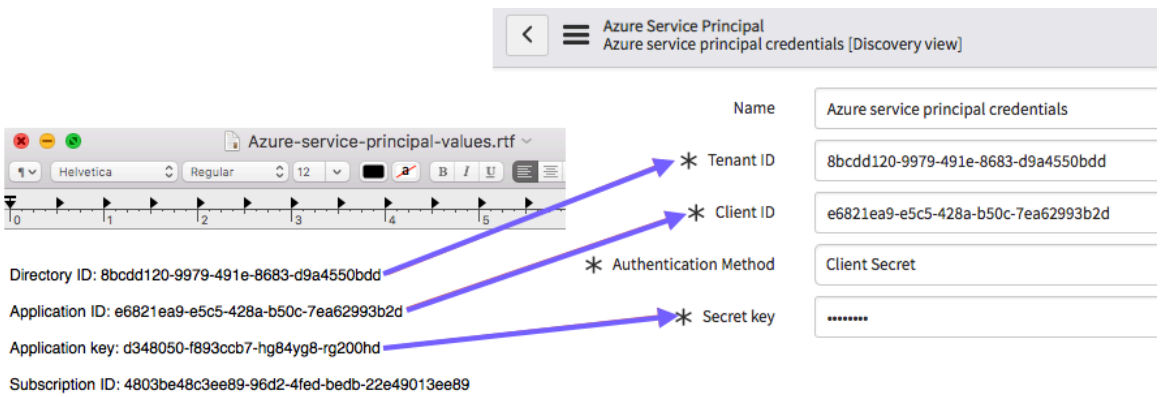

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credentials form field	Azure Service Principal value
Tenant ID	Azure **Directory ID** value from the text file.
Client ID	Azure **Application ID** value from the text file.
Secret key	Azure **Application key** value from the text file.

13. Select **Save** to create the Azure service principal.
14. Select the **Discover Subscriptions** related link to find all subscriptions for the Azure service principal.
    The instance creates a service account for each discovered subscription. The **Azure Subscriptions** related list displays all subscriptions for the Azure service principal.
15. Select a subscription to view the service account created for the subscription.
16. Select a Discovery status entry in the **Credential Discovery Status** list to view the Discovery log.
    Each time you select **Discover Subscription**, the instance generates a new Discovery status and displays it in the **Credential Discovery Status** list.

Run Azure Cloud Discovery using Service Principal with SSH Certificates

Discover Linux virtual machines on Azure using Service Principal (SP) with short-lived SSH certificates. Using these certificates circumvents the need for passwords or public and private key-pairs.

Before you begin

Before creating a Cloud Discovery credentials, Service Principles must be set up on the Azure Portal. See the Microsoft Azure documentation site for more information. After creating a Linux VM with Azure AD login enabled, be sure to verify the requirements for login with Azure AD using OpenSSH certificate-based authentication for Linux VMs. Configure suitable role assignments for the Service Principle and Resource Group.

Before creating credentials, the **External Credential Storage** plugin is required to connect Azure VM using OpenSSH certificates.

Role required: admin

Procedure

1. Navigate to **Discovery** > **Credentials**.
2. Select **New**.
3. For the type of credential, select **Azure Service Principal**.
4. Fill in the form with the required information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

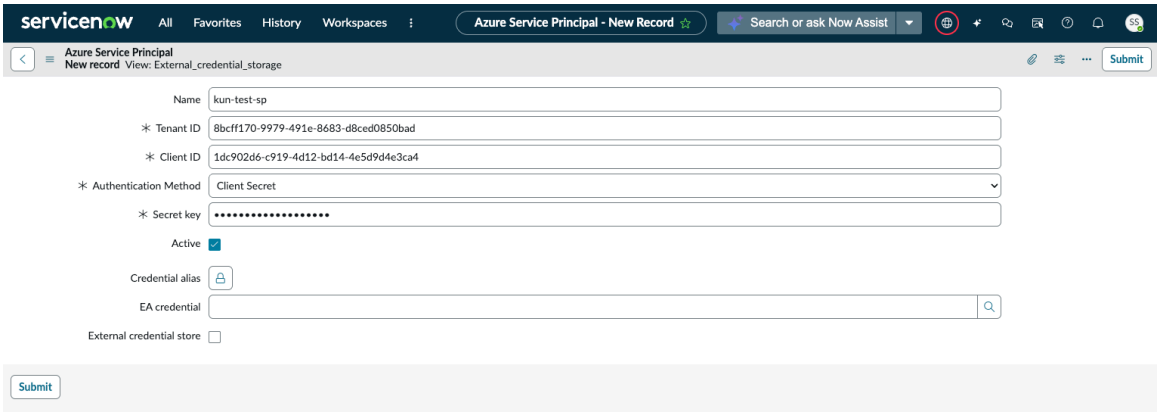

5. Select **Submit**.
6. Navigate to **Discovery** > **Credentials**.
7. Select **New**.
8. For this type of credential, select **Azure SSH Certificate Credential**.
9. Fill in the form with the necessary information, including linking the Service Principle credential you created.
   The Azure Service Provider and SSH Certificate credentials have been created and linked. Continue the procedure to create the Cloud Discovery schedule.
10. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
11. Select **New Discovery** from the header of any tab on the Schedules page.
12. Select **Cloud-based discovery** and select **Continue**.
13. Provide a name for the schedule and select **Azure** as the cloud provider.
14. Choose a MID Server.
15. Select **next**.
16. Create a cloud account using your Azure Service Principal credential.
    For more information, see Set up Azure service accounts.
17. Select **next**.
18. Select cloud regions for the Discovery schedule.

    You can choose to include all regions linked to the cloud account or select specific regions. There's no restriction on the number of cloud regions that you can add.

    a. **Optional:** To discover specific cloud regions, select the check box for the cloud region you want to scan from the **Available cloud regions** list.

    b. **Optional:** If the cloud regions don't appear, select **Refresh cloud regions**.
19. Select **next**.
20. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.
21. Select **next**.
22. Complete the schedule creation by selecting **Finish and run**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

The Discovery schedule should start, and the Discovery Admin Workspace should show the running status for the newly created schedule. After some time, the scheduled discovery should be completed and a new schedule for the VM discovery is then created and run. The new VM discovery schedule utilizes the SP we created for the generation of SSH certs to authenticate with the VMs. You can observe this in the Discovery IP Affinity section for the credential.

Create an Azure Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create an Azure Discovery schedule.

Before you begin

Verify the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have configured a service principal in the Azure Active Directory that has the necessary permissions so that the MID Server can access the cloud resources you plan to discover. For more information, see Discovery for Microsoft Azure.
- You have created suitable credentials in the ServiceNow AI Platform to connect with the service accounts on which you want to perform the discovery. For more information, see Create Azure cloud credentials.
- You're using Discovery Admin Workspace v1.11.0 or later.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **Azure** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters pro vide failover protection and load balancing be tween MID Servers. See Configure a MID Serv for more information.
Automatically select a MID server	An available MID Server is automatically se lected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validat ed, up, and have the Cloud Management ca pability are displayed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your in stance. Once an account is selected, the **Ac count name** and **Subscription ID/Manage ment Group ID** fields auto-populate.
New cloud account	Add a new cloud account to your instance by entering a name and subscription ID or man agement group ID. **Note:** You can access the Subscrip tion or Group ID in the Azure portal, un der the **Subscriptions** or **Management Groups** section, respectively.

Even if the service account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

9. Select **Discover all Subscriptions** or **Discover specific Subscriptions**.

Note: This option displays only when the target cloud account is an Azure management account.

a. If discovering specific subscriptions, select **Add subscriptions**.
The Discoverable Subscriptions dialog displays.

b. Select the check box next to one or more corresponding names and then select **Add**.

10. Select **Next**.
11. Select cloud regions for the discovery schedule.

You can choose to include all regions linked to the cloud account or select specific regions. There's no restriction on the number of cloud regions that you can add.

a. **Optional:** To discover specific cloud regions, select the check box for the cloud region you want to scan from the **Available cloud regions** list.

b. **Optional:** If the cloud regions don't appear, select **Refresh cloud regions**.

12. Select **Next**.
13. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.

14. Select **Next**.
15. Configure the Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the sched ule runs.
Run on demand	The schedule only runs when triggered manu ally. The schedule can be triggered by select ing **Finish and run** at the end of this setup, or by navigating to the schedule in the Discov ery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Dis covery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Enable event-driven discovery	If Azure sends resource change events, the Configuration Management Database (CMDB) is updated without running the full discovery schedule. Only the affected re sources are processed. For more information, see Azure change processing.
Finish and run	After being selected, all the information pro vided is validated. Then, a Discovery sched ule is created in the background, a Discov ery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information pro vided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_sched ule] table. You can edit the schedule informa tion or run it by selecting **Discover Now**.

Microsoft Azure Alert driven discovery

The Microsoft Azure Alert service can raise alerts for any changes in the life-cycle state or the configuration of a cloud resource. You can configure the service to auto-update the latest resource information in the Configuration Management Database (CMDB) without waiting for the next scheduled Cloud Discovery to run.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: The Azure change processing offers an optimal performance when compared to the Microsoft Azure alert-driven discovery that uses a legacy form of authentication and should no longer be used. Starting with the Discovery and Service Mapping Patterns app version 1.1.0, the Azure change processing has replaced the Microsoft Azure alert-driven discovery. Use the new secure webhook or migrate to Azure change processing for additional performance gains. For more information, see on the secure webhook, see the knowledge base article Azure Alert service to auto-update the CMDB has webhook url without credentials in the azure portal.

During event processing, the Cloud Event Scheduler identifies the domain of the associated service account and assigns it to the event. If domain identification fails before processing, the event can remain unassigned and be visible across all domains. To limit the visibility of failed events to the service-provider administrator, set the *sn_cmp.error_events.default_domain* property to the sys_id of the service-provider domain.

You can configure the Cloud Event Scripted REST API to receive alerts from the Microsoft Azure Alert service. For each alert—also called an event—the API creates a record in the Cloud Events [sn_cmp_cloud_event] table. The Cloud Event Scheduler reads the events from the [sn_cmp_cloud_event] table for batch processing. The scheduler invokes the pattern to process the events that are in ready state. The patterns discover the affected resource and add or update its details in the CMDB. This method helps to update the CMDB Configuration Item (CI) data and view the latest data in the Cloud User portal.

Important: If you're using domain separation for Cloud Discovery, the events are also domain-separated. Therefore, you can view the details of a processed event only if it belongs to your domain. If an event is not associated with any service account, then it is associated with the global domain.

For each event, the Cloud Event Scripted REST API stores the following information [sn_cmp_cloud_event] table:

- **Source**: This field is always set to `azure activity log`.
- **Event Name**: Name of the event received from the Microsoft Azure console.
- **Resource ID**: ID of the affected resource. If there are VM events, Cloud Discovery uses this information to identify the VM.
- **Resource Type**: Resource type of the affected resource.

Configure the Microsoft Azure Alert service to auto-update the CMDB

Configure the Microsoft Azure Alert service to auto-update the Configuration Management Database (CMDB) without waiting for the next scheduled Cloud Discovery to run.

Before you begin

- Ensure that either Discovery or Cloud Provisioning and Governance is activated in the instance.
- ○ An active MID Server with Microsoft Azure discovery capability is available.

  Note: The alert configuration continues to ping alerts even when the MID Server is down. Which in turn leads to a high number of errors in the ECC queue. If no MID Server is available, you must manually deactivate the Microsoft Azure alert rule. The Azure Alert Rules [azure_alert_rule] table contains all the Microsoft Azure alert rules.

Role required:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- ServiceNow role: sn_cmp.cloud_event_integration
- Microsoft Azure roles:
  - Microsoft Azure credentials with the Reader role
  - Microsoft Azure custom role with the following permissions:

```
"permissions": [
            {
                "actions": [
                    "Microsoft.Insights/ActionGroups/Write",
                    "Microsoft.Insights/ActionGroups/Delete",
                    "Microsoft.Insights/ActionGroups/Read",

 "Microsoft.Insights/ActivityLogAlerts/Write",

 "Microsoft.Insights/ActivityLogAlerts/Delete",

 "Microsoft.Insights/ActivityLogAlerts/Read",

 "Microsoft.Insights/ActivityLogAlerts/Activated/Action",
                    "Microsoft.Insights/AlertRules/Write",
                    "Microsoft.Insights/AlertRules/Delete",
                    "Microsoft.Insights/AlertRules/Read",
                    "Microsoft.Resources/deployments/write",
                    "Microsoft.Resources/deployments/delete",
                    "Microsoft.Resources/deployments/read"
                ],
```

Note: To set up the alert-driven discovery, you can use a Microsoft Azure service principal with the reader role only. For more information, see the community article Azure alert driven discovery with reader privileges.

About this task

After you configure the alert service, it facilitates tracking of the life-cycle state and configuration of the Microsoft Azure cloud resources.

Important: You can use a secure webhook to forward the Microsoft Azure alerts to the ServiceNow AI Platform. For more information, see Use a secure webhook to forward the Microsoft Azure alerts to the ServiceNow AI Platform [KB1208357].

Procedure

1. Disable authentication for the Cloud Event Scripted REST API, because the Microsoft Azure Alert doesn't support basic authentication.

   a. As an administrator, navigate to **Scripted Web Services** > **Scripted REST APIs**.

   b. Open the Cloud Event record.

   c. From the **Resources** tab, open the Cloud Config Event Post record.

   d. Set the application scope to **Global**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. On the **Security** tab, clear the **Requires Authentication** check box.

> **Important:** The ServiceNow instance uses a webhook URL for the alert-driven discovery. However, the Microsoft Azure alert configuration of the ServiceNow instance doesn't support authentication for creating the webhook. You can enable authentication by providing the user name and password in the webhook URL. For more information, see KB0963526.

2. Navigate to **Discovery** > **Cloud Alerts** > **Azure Alert Configurations**.

> **Note:** If you have Discovery activated in the instance, you can navigate directly to **Alert Configurations** using your instance URL. For example,
>
> ```
> https://{instancename}.service-now.com/sn_cmp_azure_alert_configuration_list.do
> ```

3. On the **Azure** tab, select **New**.
4. On the form, fill in the fields.

Azure Alert Configuration form

Field	Description
Service Account	Microsoft Azure service account for which you want to create the Azure alert configuration. Make sure that Cloud Discovery has discovered all the Logical Datacenters (LDC) associated with the cloud account.
Resource Group	Resource group that contains the alert rules. Select the Microsoft Azure service account. For more information on adding a Microsoft Azure service account, see Add an Azure service account. The instance creates the alert rules for all the resources that are available in the service account.
Activate	Option to indicate the status of the alert configuration. This option is selected by default.
User	User name of a ServiceNow user having the sn_cmp.cloud_event_integration role.
Password	ServiceNow user account password.

5. Select **Submit**.
   On successfully configuring the Microsoft Azure Alert service, the status of the alert configuration changes from **None** to **Activated**.

Add a Microsoft Azure alert rule

Enable alert-driven discovery for a new Microsoft Azure alert type by adding it to the Azure Alert Rules [azure_alert_rule] table.

Before you begin
Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The Azure Alert Rules [azure_alert_rule] table contains all the supported alert types. For more information, see Supported Microsoft Azure alert types. The Microsoft Azure Alert service uses the information stored in the Azure Alert Rules [azure_alert_rule] table to raise the alerts. Cloud Discovery also supports all the alerts generated for the members of the resource group specified in the Microsoft Azure alert configuration.

Procedure

1. Navigate to the Azure Alert Rules [azure_alert_rule] table by entering the following URL in the navigation bar of the browser:
   `https://<instance name>.service-now.com/azure_alert_rule_list.do`
2. Select **New**.
3. On the form, fill in the fields.

Azure Alert Rule form

Field	Description
Rule name	Name of the Microsoft Azure alert type.
Action	Microsoft Azure resource operation API path. For more information on the Microsoft Azure resources for which you can raise alerts, go to the Microsoft Azure documentation and see "Azure resource provider operations".
Active	Option to indicate the status of the Microsoft Azure alert type. This option is selected by default.

4. Select **Submit**.

Supported Microsoft Azure alert types

The base system Cloud Discovery supports alert driven discovery for several Microsoft Azure alert types. The Azure Alert Rules [azure_alert_rule] table contains the supported alert types.

The base system Cloud Discovery supports the following Microsoft Azure alert types:

```
Microsoft.Compute/virtualMachines/deallocate/action
Microsoft.Network/privateDnsZones/delete
Microsoft.Compute/availabilitySets/write
Microsoft.Network/loadBalancers/write
Microsoft.Network/connections/delete
Microsoft.Network/publicIPAddresses/delete
Microsoft.Storage/storageAccounts/write
Microsoft.Network/virtualNetworks/delete
Microsoft.Sql/servers/databases/delete
Microsoft.Network/expressRouteCircuits/write
Microsoft.Network/localnetworkgateways/delete
Microsoft.Network/networkInterfaces/delete
Microsoft.Compute/virtualmachines/write
Microsoft.Network/expressRouteCircuits/delete
Microsoft.Network/natGateways/write
Microsoft.Network/loadBalancers/delete
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Microsoft.Compute/virtualMachineScaleSets/delete
Microsoft.Network/virtualNetworks/virtualNetworkPeerings/delete
Microsoft.Storage/storageAccounts/delete
Microsoft.Compute/availabilitySets/delete
Microsoft.Network/connections/write
Microsoft.Compute/virtualMachines/start/action
Microsoft.Compute/virtualMachines/restart/action
Microsoft.Sql/servers/databases/write
Microsoft.Compute/virtualMachines/delete
Microsoft.Network/networkInterfaces/write
Microsoft.Network/privateDnsZones/write
Microsoft.Network/localnetworkgateways/write
Microsoft.Network/natGateways/delete
Microsoft.Compute/virtualMachines/stop/action
Microsoft.Network/publicIPAddresses/write
Microsoft.Compute/virtualMachineScaleSets/write
Microsoft.Network/virtualNetworks/virtualNetworkPeerings/write
Microsoft.Network/virtualNetworks/write
Microsoft.Sql/servers/delete
Microsoft.Sql/servers/write
```

Azure change processing

The Azure change processing feature collects information about Microsoft Azure resources that have undergone a life-cycle state change or configuration change near real time. Then, it uses the collected information to update the Configuration Management Database (CMDB).

The Azure Change Processing scheduled jobs query the Azure API and collect the resource change information. During each execution cycle, the scheduled job collects information about all the resources that have changed after the last Azure Change Processing cycle. After collecting the change information, the Azure Change Processing uses response mappings to update the change information in the CMDB. During the next discovery, Cloud Discovery triggers appropriate patterns, if available, and populates detailed resource information in the CMDB.

By default, the Azure change processing scheduled jobs run in every 5 minutes. If necessary, update the execution frequency of the schedule according to your needs. Set the execution frequency within the minimum value of one minute.

When you execute the Azure change processing for the first time, it may run a few hours. This is due to processing resource changes that occurred in the last four hours. If you want to increase the default schedule execution period, ensure that enough worker nodes are available to execute the schedule.

During event processing, the Cloud Event Scheduler identifies the domain of the associated service account and assigns it to the event. If domain identification fails before processing, the event can remain unassigned and be visible across all domains. To limit the visibility of failed events to the service-provider administrator, set the *sn_cmp.error_events.default_domain* property to the sys_id of the service-provider domain.

Note:

Azure change processing can fetch resource change information only from the service accounts created by the discovery_admin or sn_cmp.cloud_admin.

Update your Discovery and Service Mapping Patterns plugin on ServiceNow Store to have at least 1.21.0 version of Azure Change Processing. See the upgrade scenario in KB1705862

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Starting in the Discovery and Service Mapping Patterns 1.21.0 version, the scheduled jobs were enhanced and renamed:

- The *Azure Process Changes* scheduled job was renamed as *ACP Pull Changes*.
- The *CPG Build Azure Changes IRE Payload* scheduled job was renamed as *ACP Process Changes*.

ACP tables

Starting with Discovery and Service Mapping Patterns 1.21.0 version, Azure change processing enhancement populates the following tables:

ACP Resource Type [sn_itom_pattern_acp_resource_types]

> **Note:** In Discovery and Service Mapping Patterns versions 1.21.0-1.30.2, resource type information was stored in the sn_itom_pattern_acp_supported_resource_types table.

You can use this table to review all supported resource types, add a resource type for change processing, and disable or enable the processing for a certain resource.

For more information on adding support for an Azure resource type, see Add change processing support for an Azure resource type.

Key fields:

- Resource Type.
- Query: The resource graph query that retrieves changes from Azure.
- Active: Indicates if the resource type is actively being queried.

ACP Order Status [sn_itom_pattern_acp_order_status]

This table tracks the status of the Azure Resource Graph query for each resource type, based on the associated credentials that are used for querying Azure. You can use this table to review the processing status. For more information about the query, see KB1705862.

Key fields:

- Resource Type.
- Credential: the specific credential used for pulling changes.
- Last Processed time.
- Status: Tracks the success or failure of the event pull process. After the next ACP Pull Changes scheduled job, this status will be updated based on the CMP Order status.
- CMP Order: Reference to the CMP Order [sn_cmp_order] field, which provides details about the current order progress and outcome.

ACP Resource Change [sn_itom_pattern_acp_resource_change]

This table records changes detected in Azure resources. The [sn_cmp_resource_changes_payload_info] table that used to store this info before the 1.21.0 version is no longer available.

Key fields:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Resource ID: the unique identifier for the changed resource in Azure.
- Change Payload: a snapshot of the attributes of interest for the resource.
- Order: A reference to the CMP Order [sn_cmp_order] field that pulled the change.
- Resource Status: A reference to the ACP Resource Status table to track the processing status of the change.
- State: the current processing state of the change. If the state is marked as 'Error', review the associated Resource Status record for further details.

ACP Resource Status [sn_itom_pattern_acp_resource_status]

This table tracks the overall status of resource processing.

Key fields:

- Resource ID: the identifier for the resource being processed.
- Status: the processing state. The states are: Ready, In Progress, Failure, and Success.
- Error Message: If the processing failed, this field contains details about the error.

Supported Azure resource types and changes

Azure Change Processing tracks a set of resource properties for each supported resource type. When a change is detected in one of these properties, Azure Change Processing retrieves an updated snapshot of the resource from Azure. It then processes the snapshot in the CMDB to reflect the change.

Resource type	Resource Property
Microsoft.Compute/ virtualMachine	• *CreateOrDeleteVm : properties.changeType in ('Create','Delete')* • *ChangeStatus: (properties.extended.instanceView.powerState.code) in ('PowerState/running', 'PowerState/stopped', 'PowerState/deallocated'))* • *ChangeDisk: properties.storageProfile* • *ChangeTag: tags.** • *ChangeNetworkInterface : properties.networkProfile.networkInterfaces*
Microsoft.Compute/ disks	• *CreateOrDeleteDisk : properties.changeType in ('Create','Delete')* • *ChangeState : properties.diskState* • *ChangeSize : properties.storageProfile* • *ChangeTag : tags.**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource type	Resource Property
Microsoft.Network/networkSecurityGroups	• *CreateOrDeleteNsg : properties.changeType in ('Create','Delete')* • *ChangeTag : tags.**
Microsoft.Network/networkinterfaces	• *CreateOrDeleteNic : properties.changeType in ('Create','Delete'),* • *changePrimary : properties. properties.primary* • *ChangeIp : properties.ipConfigurations[0]'* • *ChangeTag : tags.**
Microsoft.Network/publicIPAddresses	• *CreateOrDeleteIp : properties.changeType in ('Create','Delete'),* • *ChangeIpAddress : 'properties.ipaddress',* • *ChangeFqdn : properties.dnsSettings.fqdn* • *ChangeIpConfig : properties.ipConfiguration* • *ChangeTag : tags.**

Advantages of the Azure change processing

The Azure change processing offers the following advantages over the Microsoft Azure alert-driven discovery:

- Improved performance and reduced probability of Azure API throttling
- Simple setup

Improved performance and reduced probability of Azure API throttling

The Microsoft Azure alert-driven discovery triggers a targeted discovery for each affected resource. Therefore, when the ServiceNow AI Platform receives a huge number of alerts, the targeted discovery may cause the Azure APIs to throttle. As a result, the alert processing performance of the ServiceNow AI Platform may decline. In contrast, the Azure change processing doesn't trigger targeted discovery for each affected resource. Instead, it uses response mappings to update the CMDB per the available change information. During the next discovery, Cloud Discovery triggers appropriate patterns, if available, and populates detailed resource information in the CMDB. Therefore, Azure change processing improves the change handling performance of the ServiceNow AI Platform and reduces the probability of Azure API throttling.

Simple setup

The Microsoft Azure alert-driven discovery uses a webhook to send the alerts to the ServiceNow AI Platform. As the Azure cloud generates alerts at the subscription level, the Microsoft Azure alert-driven discovery needs a webhook for each subscription you want to monitor. In contrast, Azure change processing uses CAPI and MID Server to interact with the Azure resource changes API. The API can provide change information at the management group level. Therefore, Azure change processing removes the need for webhooks and simplifies the setup.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can configure Azure change processing to get resource change information from the Microsoft Azure cloud and use it to update the CMDB.

If you're using the Microsoft Azure alert-driven discovery, then you can migrate to Azure change processing to improve the change processing performance of the ServiceNow AI Platform and take advantage of the simplified setup.

Related topics

https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB1705862

Microsoft Azure Cloud discovery using patterns

Enable Azure change processing

Configure the Azure change processing to collect resource change information from the Microsoft Azure cloud and use it to update the Configuration Management Database (CMDB).

Before you begin

- Ensure that the latest version of Discovery and Service Mapping Patterns store app is installed on the ServiceNow AI Platform.
- Ensure that Microsoft Azure alert-driven discovery isn't configured on the ServiceNow AI Platform.

Role required: discovery_admin or sn_cmp.cloud_admin

About this task

For information about the resources which are supported by Azure change processing, see Azure change processing.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Configure** *Azure* **Change Processing**.
2. Select the accounts which you want to sync Azure change processing for.
   a. Select the accounts from the **Service Accounts without 'pull-based' change processing** list.
   b. Add the selected accounts to the **Service Accounts with 'pull-based' change processing** list.

   Note: When you enable Azure change processing for an Azure management group, it is also enabled for all the subscriptions in the management group.
3. Select **OK**.

Related topics

https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB1705862

Migrate to Azure change processing

Migrate to Azure change processing to improve the change processing performance of the ServiceNow AI Platform and take advantage of the simplified setup.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that the latest version of Discovery and Service Mapping Patterns store app is installed on the ServiceNow AI Platform.
- Ensure that Microsoft Azure alert-driven discovery is configured on the ServiceNow AI Platform.
- Ensure that the Cloud Events [sn_cmp_cloud_event] table doesn't contain any event in ready state.
- Ensure that the scope is set to Global.
- Ensure that at least one active MID Server is available.

Role required: discovery_admin or sn_cmp.cloud_admin

About this task

If you are using the Microsoft Azure alert-driven discovery and the ServiceNow AI Platform receives many alerts from the Azure cloud, processing those alerts and updating the Configuration Management Database (CMDB) can take a significant amount of time. Migrate to the Azure change processing to overcome the performance issues. For more information on the Azure change processing, see Azure change processing.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Configure** *Azure* **Change Processing**.
2. Select Retrieve Azure alert configurations to retrieve the list of all the active Azure alert configurations.
3. Select **Disable** *Azure* **alert configurations**.
   The ServiceNow AI Platform deactivates the Azure alert configurations and removes the webhooks from the Azure portal.
4. Select the accounts for which you want to configure Azure change processing.

   a. From the **Service Accounts without 'pull-based' change processing** list, select the accounts for which you want to configure Azure change processing.

   b. Add the selected accounts to the **Service Accounts with 'pull-based' change processing** list.

   Note: When you enable Azure change processing for an Azure management group, it is also enabled for all the subscriptions in the management group.

5. Select **OK**.

Result

The migration script marks the Azure configurations module as deprecated. In the Cloud Discovery. The Azure Alert Configurations page is hidden in the Cloud Discovery Workspace.

Related topics

https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB1705862

Add change processing support for an Azure resource type

Add Azure change processing support for an Azure resource type per the needs of your organization.

Before you begin

Role required: discovery_admin or cloud_root_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

This procedure is applicable only for Discovery and Service Mapping Patterns store versions that are older than 1.21.0. Starting the 1.21.0 version, all supported resource types are listed in the ACP Resource Type table. All the resource types are activated by default. The MID Server property *mid.cmp.azure.event.supported_resource_types* is no longer available. You can use the ACP Resource Type table to add an entry for a resource type and to add a query. For information on adding a query see the "Example VM Query" section in KB1705862.

To add change processing supports for an Azure resource type in legacy versions, add the resource types to the *mid.cmp.azure.event.supported_resource_types* property. As a comma-separated value. Then, create response mappings between the change fields and the Configuration Item (CI) fields. Response mappings define a one-on-one mapping between the source data fields and the CI fields. Azure change processing uses the mapping information to update resource change data in the appropriate CI class fields.

Procedure

1. Navigate to **All** > *MID Server* > **Properties**.
2. Open the *mid.cmp.azure.event.supported_resource_types* property.
3. In the **Value** field, add the *Azure* resource type as a comma-separated value.
4. Select **Update**.
5. Create response mappings for the change fields that you want to update in the Configuration Management Database (CMDB).

   a. Navigate to the Response Mapping [sn_cmp_response_mapping] table.
   `https://<instance name>.service-now.com/sn_cmp_response_mapping_list.do`

   b. Select **New**.

   c. On the form, fill in the fields.

   Response Mapping form

Field	Description
CI class	Name of the CI class that represents the Azure resource in the CMDB.
Datasource	Name of the input datasource. Set the Datasource field to Azure Resource Changes.
CI class fields	Name of the CI class field that is the logical equivalent of the source field.
Source field	Name of the source field that you want to update in the specified CI class field.

   For more information on response mapping, see Using response mapping topic in the ServiceNow Community.

   d. Select **Submit**.

   e. Repeat steps 5.b to 5.d for all the change fields that you want to update in the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery for Google Cloud Platform

Discovery finds Google Cloud Platform and its components. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Data discovered for the Google Cloud Platform components offers visibility for multiple ServiceNow products, including Cloud Provisioning and Governance.

Discovery uses multiple patterns to find GCP components.

You can discover global and regional resources.

Discoverable global and regional GCP resources

Global	Regional
Global static external IP addresses	Regional static external IP addresses
Global Load Balancers (Global Http Load Balancers)	Regional Load Balancers (Regional HTTP Load Balancers, TCP/UDP Load Balancers)
Images	Subnets
Snapshots	Regional managed instance groups
Instance Templates	Regional persistent disks
VPC Network	
Firewalls	

Requirements for GCP discovery

Verify that the applications are up to date.

- Discovery and Service Mapping Patterns starting version 1.8.0 (September 2023)
- CMDB CI Class Models starting version 1.49.0 (September 2023)
- Visibility Content starting version 6.13.0 (September 2023)

Verify the MID Server requirements

- Use a dedicated MID Server for cloud discovery. The MID Server must have 16 Cores, 8 GB with 4 GB Java Virtual Machine (JVM) memory allocation. For information on verifying MID Server specifications, see KB article KB2156327.
- Deploy and connect the dedicated MID Server with internet access.
- For more information on the MID Server configurations, see Install and configure MID Servers to access cloud environments.

Verify Google Cloud BigQuery and Google Cloud Bigtable requirements

To discover Google Cloud BigQuery and Google Cloud Bigtable, navigate to **All** > **Discovery** > **MID Server** > **Properties** and ensure that the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

sn_itom_pattern.google_dbservice_table_collection property is set to **True**.

Verify user role settings

Set the user permission for the Google Cloud Platform member to **Viewer**:

1. In the Google Cloud Platform console, navigate to **IAM**.
2. Select the relevant member from the list and click the **Edit** icon.
3. In the Edit permissions window, select `Viewer` from the **Role** list.
4. Click **Save**.

Configure Google API Credentials

For more information, see Configure the Google API credentials for Cloud Discovery.

Create a Google Cloud Platform service account

For more information, see Create GCP service accounts.

Enable collecting cloud logging events

The Google Cloud Platform patterns can collect cloud logging events generated by the native Google Cloud APIs. Configure the Google Cloud Logging service to auto-update the CMDB.

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Verify that a Discovery schedule is configured

For more information, see Create a GCP Discovery schedule in Discovery Admin Workspace.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected by Service Mapping during tag-based discovery

Service Mapping uses tag-based discovery to create service instance maps including the Cloud components. The Service Mapping application comes with the following preconfigured CI relationships used for tag-based discovery. These CI relationships are available from the 1.0.68 release on the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Configuration Item [cmdb_ci]	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]
Logical Datacenter [cmdb_ci_logical_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Useful information

- GCP discovery solutions comparison
- Set up Google Cloud Platform service accounts
- Google Cloud Platform (GCP) Cloud discovery using Patterns

GCP cloud discovery methods and use cases

Comparison of use cases and requirements for cloud discovery methods in GCP.

ITOM Visibility offers several methods for GCP cloud discovery. Each method supports different use cases and has its own requirements.

GCP cloud discovery use cases

Comparison of GCP cloud discovery methods and use cases

Use case	Cloud Discovery (Pattern-based)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
IT Service Management (ITSM) - basic and generic virtual CI classes	Yes	Yes	No	No	No
ITSM - extended virtual CI classes	No	No	No	No	No
ITSM - hardware CI classes	No	Yes	Yes	Yes	Yes
Pattern discovery of applications	No	No	No	Most	Yes
Tag-based service mapping	Yes	Yes	No	No	No
Machine Learning (ML)-based service mapping	No	Some	Yes	Yes	Yes
Top-down service mapping	No	No	No	Most	Yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison of GCP cloud discovery methods and use cases (continued)

Use case	Cloud Discovery (Pattern-based)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
Software Asset Management (SAM) - package manager	No	Yes	Yes	Yes	Yes
SAM - file-based	No	No	No	Yes	Yes
Oracle Global License Advisory Services (GLAS)	No	No	No	Yes	Yes
Certificate discovery	No	No	No	Yes	Yes, without host relationships

GCP cloud discovery requirements

Comparison of GCP cloud discovery methods and requirements

Requirement	Cloud Discovery (Pattern-based)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
MID Server	Yes	No	No	Yes	Yes
Cloud credential	Yes	Yes	No	No	No
OS credential	No	No	No	No	Yes
ServiceNow agent	No	No	Yes	Yes	No
Google Cloud Software Development Kit (SDK) - all virtual machines (VMs)	No	Yes	No	No	No
Google OS Config agent - all VMs	No	Yes	No	No	No
Google Cloud Storage buckets	No	Yes	No	No	No
Network Connectivity Target VM to MID Server	No	No	No	Yes	No

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison of GCP cloud discovery methods and requirements (continued)

Requirement	Cloud Discovery (Pattern-based)	Service Graph Connectors	Agent Client Collector for Visibility Content (without MID Server)	Agent Client Collector for Visibility Content (with MID Server)	IP-based Horizontal Discovery (Pattern-based)
Network Connectivity MID Server to Target VM	No	No	No	No	Yes

Configure the Google Cloud Logging service to auto-update the CMDB

You can activate the Google Cloud Logging (formerly Stackdriver Logging) service to auto-update Configuration Management Database (CMDB) configuration items (CI) data whenever Google Cloud Connector or your Google account makes a life-cycle state or configuration change to a Google Cloud Platform (GCP) resource. As a result, the CI data in the CMDB is updated without having to wait for Discovery to run.

Before you begin

- Activate Discovery and Service Mapping Patterns. For more information, see Google Cloud Platform Discovery.
- Verify that Discovery has successfully discovered logical datacenters of the Google Cloud project. For more information, see Set up a cloud account and service account for Google Cloud Platform.

> **Note:** Verify that you have selected the **Should pull events** check box in the Google Service Account page before you execute datacenter discovery.

Role required:

- discovery_admin for Cloud Discovery and `sn_cmp.cloud_event_integration` if authentication to cloud event endpoint is required. This role is required to set up an Alert Rule in the instance.
- Google Cloud Platform administrator for configuration steps in the GCP console.
- Logging Role for the Google Cloud Platform integration user that would allow OAuth scopes for the following
  - `https://www.googleapis.com/auth/logging.admin`
  - `https://www.googleapis.com/auth/cloud-platform.read-only`
  - `https://www.googleapis.com/auth/cloud-platform`
  - `https://www.googleapis.com/auth/logging.read`

About this task

Discovery integrates with Google Cloud's Cloud Logging API to collect log data from resources in your Google Cloud projects. This integration enables Discovery to automatically update the CMDB based on log entries that reflect changes in your environment. The Cloud Logging service provides:

- Centralized log collection for Google Cloud resources.
- APIs to query and export logs for operational workflows.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When configured, Discovery pulls the Cloud Logging API at regular intervals to detect updates and synchronize them with the CMDB.

Procedure

1. Log in to your Google Cloud Platform account.
   For your project, navigate to **APIs & Services** > **Library**.
2. Enable the **Cloud Logging API**.
   a. In the library, search for `Cloud Logging`.
   b. Select the **Cloud Logging API** card.
   c. On the Cloud Logging API page, select **Enable**.
3. To modify the schedule, navigate to **System Scheduler** > **Scheduled Jobs** and then select **GCP-Events-job** and then update the **Repeat** setting as needed.

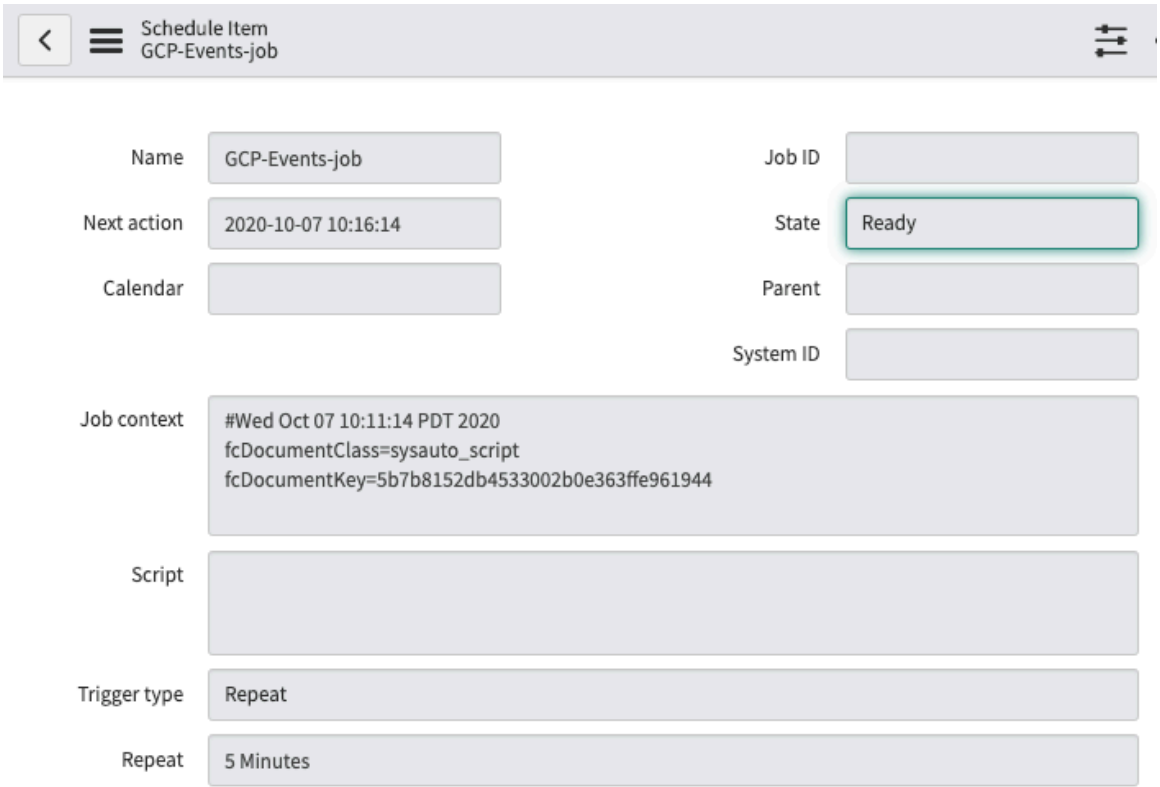

4. Select **Update**.

Result

- Once the Google Cloud Platform generates an event and Discovery receives the event information, the 'GCP-Events-job' scheduled jobs execute in the background.
- The scheduled job passes the event payload to the Identification and Reconciliation Engine (IRE). The IRE then takes the necessary actions to modify the state of the resources based on the event.

Create a GCP Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create a Google Cloud Platform (GCP) Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Verify the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have configured a service account in the Google Cloud console that has the necessary permissions so that the MID Server can access the cloud resources you plan to discover.
- You have created suitable credentials in the ServiceNow AI Platform to connect with the service accounts on which you want to perform the discovery. For more information, see Set up a cloud account and service account for Google Cloud Platform.
- You're using Discovery Admin Workspace v1.11.0 or later.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **Google** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters provide failover protection and load balancing between MID Servers. See Configure a MID Server cluster for more information.
Automatically select a MID server	An available MID Server with an Identity and Access Management (IAM) role is automatically selected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validated, up, and have the Cloud Management capability are displayed.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your instance. Once an account is selected, the **Account name** and **Project ID/Folder ID** fields auto-populate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
New cloud account	Add a new cloud account to your instance by entering a name and subscription ID or man agement group ID. **Note:** You can access the Project or Folder ID in the Google portal via **IAM & Admin** > **Settings** or **IAM & Admin** > **Folders**, respectively.

Even if the cloud account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

9. Select **Discover all Projects** or **Discover specific Projects**.

   Note: This option displays only when the target cloud account is a Google management account.

   a. If discovering specific projects, select **Add projects**.
   The Discoverable Projects dialog displays.

   b. Select the check box next to one or more corresponding names and then select **Add**.

10. Select **Next**.
11. Select cloud regions for the discovery schedule.

    You can choose to include all regions linked to the cloud account or select specific regions. There's no restriction on the number of cloud regions that you can add.

    a. **Optional:** To discover specific cloud regions, select the check box for the cloud region you want to scan from the **Available cloud regions** list.

    b. **Optional:** If the cloud regions don't appear, select **Refresh cloud regions**.

12. Select **Next**.
13. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.

14. Select **Next**.
15. Configure the Discovery schedule.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the sched ule runs.
Run on demand	The schedule only runs when triggered manu ally. The schedule can be triggered by select ing **Finish and run** at the end of this setup, or by navigating to the schedule in the Discov ery Schedules [discovery_schedule] table and selecting **Discover Now**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Run after series	The schedule runs after another existing Dis covery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Enable event-driven discovery	If Google Cloud Connector or your Google account makes a life-cycle state or configu ration change to a GCP resource, the Con figuration Management Database (CMDB) is updated without running the full discovery schedule. Only the affected resources are processed. For more information, see Config ure the Google Cloud Logging service to au to-update the CMDB.
Finish and run	After being selected, all the information pro vided is validated. Then, a Discovery sched ule is created in the background, a Discov ery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information pro vided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_sched ule] table. You can edit the schedule informa tion or run it by selecting **Discover Now**.

Discovery for IBM Cloud Platform

If your cloud resources are in an IBM cloud, create credentials that can access your IBM account.

IBM credentials form

Field	Value
Name	Unique name of the IBM credentials to register with the instance.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Softlayer API User	Softlayer cloud user name.
Softlayer API Key	Softlayer API key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM credentials form (continued)

Field	Value
Bluemix API Key	Bluemix API key.

Data collected by Service Mapping during top-down discovery

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Service Mapping uses tag-based discovery to create service instance maps including the Cloud components. The Service Mapping application comes with the following preconfigured CI relationships used for tag-based discovery. These CI relationships are available from the 1.0.68 release on the ServiceNow Store.

CI	Relationship	CI
Configuration Item [cmdb_ci]	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]
Logical Datacenter [cmdb_ci_logical_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

IBM Cloud Discovery API list

Region

CI Attributes	IBM Constraints
object_id	name
dns_domain	domain
name	alias
short_description	display_name
region	name

cmdb_ci_availability_zone

CI Attributes	IBM Constraints
object_id	id
short_description	longName
name	name
install_status	statusId

Location Group (cmdb_ci_cloud_location_group)

CI Attributes	IBM Constraints

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Location Group (cmdb_ci_cloud_location_group) (continued)

object_id	id
short_description	description
location_group_type_id	locationGroupTypeId
name	name
location_group_type	locationGroupType.name

Virtual Server (cmdb_ci_vm_instance)

CI Attributes	IBM Constraints
object_id	id
cpus	maxCpu/startCpus
memory	maxMemory
disks	<not mapped>
disks_size	<not mapped>
nics	<not mapped>
terminated_on	<not mapped>
state	statusId
termination_protection	<not mapped>
short_description	notes
name	hostname
dns_domain	domain
fqdn	fullyQualifiedDomainName

Image (cmdb_ci_os_template)

CI Attributes	IBM Constraints
cmdb_ci_os_template	
image_type	
guest_os	
version	
offer	
OS Status	
root_device_type	
image_source	
credentials	
infuse_key	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Image (cmdb_ci_os_template) (continued)

CI Attributes	IBM Constraints
update_host_name	
name	name
object_id	id
short_description	description

Add public key (cmdb_ci_cloud_key_pair)

CI Attributes	IBM Constraints
object_id	id
finger_print	fingerprint
created_date	createDate
key	key
name	label
modify_date	modifyDate
short_description	notes

Storage Volume (cmdb_ci_storage_volume)

CI Attributes	IBM Constraints
volume_id	id
size	capacityGb
size_bytes	capacityGb
free_space	<not mapped>
free_space_bytes	<not mapped>
media_type	<not mapped>
file_system	<not mapped>
label	serviceResourceName
lun (Logical Unit Number)	<not mapped>
volume_container	<not mapped>
state	?
storage_type	storageTypeId
protected_by	<not mapped>
share_count	<not mapped>
sharable	FALSE
device	guestId
delete_on_termination	<not mapped>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume (cmdb_ci_storage_volume) (continued)

CI Attributes	IBM Constraints
mapping_type	<not mapped>
object_id	id
vdisk_id	<not mapped>
storage_type	nasType
name	serviceResourceName

Storage Volume Snapshot (cmdb_ci_storage_vol_snapshot)

CI Attributes	IBM Constraints
object_id	id
capacity	capacityGb
is_encrypted	
volume_name	serviceResourceName
storage_server	
start_time	
state	
snapshotSizeBytes	snapshotSizeBytes

Network (cmdb_ci_network)

CI Attributes	IBM Constraints
state	
cidr	<not_mapped>
is_external	true/false
netmask	<not_mapped>
broadcast_address	<not_mapped>
default_gateway	<not_mapped>
is_shared	<not_mapped>
is_dhcp_enabled	<not_mapped>
scope	<not_mapped>
max_ports	<not_mapped>
terminated_on	<not_mapped>
provided_contracts	<not_mapped>
object_id	id
access_port_vlan	vlanNumber
<not_mapped>	primarySubnetId

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network (cmdb_ci_network) (continued)

CI Attributes	IBM Constraints
name	vlanNumber
<not_mapped>	accountId

Subnet (cmdb_ci_cloud_subnet)

CI Attributes	IBM Constraints
subnet_mask	netmask
cidr	networkIdentifier +'/' + cidr
available_ip_count	usableIpAddressCount
broadcast_address	broadcastAddress
primary_dns	<not mapped>
secondary_dns	<not mapped>
tertiary_dns	<not mapped>
is_dhcp_enabled	<not mapped>
ipam_provider	<not mapped>
domain_name	<not mapped>
parent_domain_name	<not mapped>
gateway	gateway
route_table_association_id	<not mapped>
state	<not mapped>
<not mapped>	modifyDate
<not mapped>	version
<not mapped>	networkVlanId
<not mapped>	isCustomerRoutable
object_id	id
<not mapped>	broadcastAddress
<not mapped>	isCustomerOwned
<not mapped>	sortOrder
<not mapped>	networkIdentifier
<not mapped>	totalIpAddresses
<not mapped>	addressSpace
category	subnetType
name	networkIdentifier +'/' + cidr

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Compute Security Group (cmdb_ci_compute_security_group)

CI Attributes	IBM Constraints
region	<not_mapped>
object_id	id
intId	<not_mapped>
tenant_name	<not_mapped>
state	<not_mapped>
<not_mapped>	port
<not_mapped>	description
<not_mapped>	createDate
<not_mapped>	modifyDate
name	name
short_description	description

Security Group Rule (cmdb_ci_compute_security_group_rule)

CI Attributes	IBM Constraints
created	createDate
direction	direction
ethertype	ethertype
object_id	id
updated	modifyDate
portRangeMax	portRangeMax
portRangeMin	portRangeMin
protocol	protocol
remoteGroupId	remoteGroupId
securityGroupId	securityGroupId

Load Balancer (cmdb_ci_cloud_load_balancer)

CI Attributes	IBM Constraints
state	"operatingStatus"
canonical_hosted_zone_id	
canonical_hosted_zone_name	
dns_name	address
snat_ip_address	
vlan	
app_code	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer (cmdb_ci_cloud_load_balancer) (continued)

CI Attributes	IBM Constraints
component_code	
env_code	
ip_address_pool	
	"accountId"
	"address"
	"createDate"
	"description"
object_id	"id"
	isPublic
	"locationId"
	"modifyDate"
name	"name"
	"provisioningStatus"
	"useSystemPublicIpPool"
	"uuid"

Load Balancer Service (cmdb_ci_lb_service)

CI Attributes	IBM Constraints
cmdb_ci_lb_service	
pool	
port	protocolPort
last_port	
input_url	
host_name	
back_end_port	
front_end_port	protocolPort
listener_protocol	protocol
server_port	
server_protocol	
traffic_domain	
ssl_certificate_id	
object_id	id
hit_count	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

LB Health Service (cmdb_ci_lb_health_service exntends cmdb_ci_vm_object)

CI Attributes	IBM Constraints
check_interval_sec	interval
max_retries	maxRetries
monitor_type	monitorType
provisioning_status	provisioningStatus
timeout_sec	timeout
object_id	id

LB Pool (cmdb_ci_lb_pool)

CI Attributes	IBM Constraints
	createDate
load_balancing_method	loadBalancingAlgorithm
	modifyDate
	protocol
	protocolPort
pool_status	provisioningStatus
object_id	uuid

LB Pool Member (cmdb_ci_lb_pool_member)

CI Attributes	IBM Constraints
server_id	address
object_id	id
	createDate
	modifyDate
	provisioningStatus
	uuid
	weight

Resource Group (cmdb_ci_resource_group)

CI Attributes	IBM Constraints
name	name
need to see what is this?	quota_id
	tags
object_id	id

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Org (cmdb_ci_cloud_org)

CI Attributes	IBM Constraints
name	name
url	url
billing_enabled	billing_enabled
quota_definition_url	quota_definition_url
spaces_url	spaces_url
guid	guid
status	status
object_id	metadata.guid

Org (cmdb_ci_cloud_space)

CI Attributes	IBM Constraints
name	name
url	url
object_id	metadata.guid
allow_ssh	allow_ssh
guid	metadata.guid

Create an IBM Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create an IBM Discovery schedule.

Before you begin

Verify the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have configured a Service ID in the IBM Cloud Console and created an associated API key that has the required Identity and Access Management (IAM) roles, so that the MID Server can access the cloud resources you plan to discover.
- You have created suitable IBM Cloud credentials in the ServiceNow AI Platform to connect with the IBM Cloud Service ID you created. For more information, see Discovery for IBM Cloud Platform.
- You're using Discovery Admin Workspace v1.13.0.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **IBM** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters pro vide failover protection and load balancing be tween MID Servers. See Configure a MID Serv for more information.
Automatically select a MID server	An available MID Server is automatically se lected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validat ed, up, and have the Cloud Management ca pability are displayed.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your in stance. Once an account is selected, the **Ac count name** and **Account ID** fields auto-popu late.
New cloud account	Add a new cloud account to your instance by entering an account name and ID.

   Even if the cloud account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

9. Select **Next**.
10. Select datacenters for the discovery schedule.

    You can choose to include all datacenters linked to the cloud account or select specific datacenters. There's no restriction on the number of datacenters that you can add.

    a. **Optional:** To discover specific datacenters, select the datacenters from the **Select the data centers** drop-down list.

    b. **Optional:** If the datacenters don't appear, select **Refresh data centers**.

11. Select **Next**.
12. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.

13. Select **Next**.
14. Configure the Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the schedule runs.
Run on demand	The schedule only runs when triggered manually. The schedule can be triggered by selecting **Finish and run** at the end of this setup, or by navigating to the schedule in the Discovery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Discovery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Finish and run	After being selected, all the information provided is validated. Then, a Discovery schedule is created in the background, a Discovery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information provided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_schedule] table. You can edit the schedule information or run it by selecting **Discover Now**.

Discovery for VMware

Applications that access VMware cloud resources need access to VMware credentials.

For example, the VMware credential type enables Discovery to explore VMware's vCenter running on a Windows machine to discover ESX machines, virtual machines, and resource pools. The VMware Discovery and automation API (vCenter API) now provides the globally unique serial number for computer CIs. CIM credentials aren’t needed to enable access to each VMware host.

Note: Windows credentials aren’t necessary for vCenter Discovery, when valid VMware credentials are used.

VMWare credentials form

Field	Description
Name	Enter a unique and descriptive name for the VMware credentials.
Active	Enable or disable these credentials for use.
User name	Enter the user name that you use for your VMware account. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMWare credentials form (continued)

Field	Description
	The VMware credentials must have the read-only role in vCenter.
Password	Enter the password for the VMware account.
Applies to	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

Related topics

Configure the VMware Events service to auto-update the CMDB

Create a VMware Discovery schedule in Discovery Admin Workspace

Configure the VMware Events service to auto-update the CMDB

Configure events from the cloud environment to make the necessary updates to your CMDB without additional scanning. The VMware Events service can auto-update CI data in the CMDB whenever Cloud Provisioning and Governance makes a life-cycle state or configuration change to a VMware resource. As a result, the CI data in the CMDB is updated without having to wait for Discovery to run.

Before you begin

Important: You can configure the VMware event service to auto-update the CMDB for event-based discovery or configuration changes even if you haven't subscribed to and activated the Cloud Provisioning and Governance plugins.

- You must have a cloud account with VMware, service accounts and associated logical datacenters.
- The following Cloud Provisioning and Governance plugins must be activated (Not mandatory for Cloud Discovery):
  - Cloud Provisioning and Governance (com.snc.cloud.mgmt)
  - Cloud Provisioning and Governance Core (com.snc.cloud.core)
  - Cloud API (com.snc.cloud.api)
  - Cloud Config Management (com.snc.config.mgmt)
- A MID Server must be available with **All** or **VMware** capabilities.
- The vCenter datacenters are successfully discovery and associated data is populated.

Role required: sn_cmp.cloud_event_integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

When VMware Events sends an update, the instance processes the event and creates or updates the CI entry in the CMDB and the CI information in the Cloud User Portal. Each event is saved as a record in the Events table on your instance. Configure all event/alert rules from the vCenter.

To connect to a different MID Server or vCenter: Update the settings and then click the **Start** related link. The **Status** value changes to **Updating** and then to **Started**.

Important:

If you are on a domain separated instance, only those events that are updated to the CMDB and belong to your domain are visible. Events create configuration items (CI) in the same domain as the cloud service account they are mapped to. Events that are not associated to a service account are visible to all domains.

During event processing, the Cloud Event Scheduler identifies the domain of the associated service account and assigns it to the event. If domain identification fails before processing, the event can remain unassigned and be visible across all domains. To limit the visibility of failed events to the service-provider administrator, set the *sn_cmp.error_events.default_domain* property to the sys_id of the service-provider domain.

Procedure

1. Navigate to the VMware Alert Configuration in one of the following ways:
   - Use the `https://baseURL/ecc_agent_ext_context_vcenter` to navigate directly using the URL.
   - On the **Cloud Admin Portal**, navigate to **Manage** > **Alert Configurations**.
2. On the **VMware** tab, click **New**.
3. On the form, fill in the fields.

vCenter Event Collector Context

Field	Description
Name	A unique name for the Event Collector record.
Short description	A brief description of the record.
Execute on	Select **Specific MID Server**
MID Server	Select a MID Server.
vCenter	Select the vCenter to monitor.
Extension [read only]	The MID Server extension of this context. Updated when data collection starts.
Status [read only]	The status of the data collection process, for example, Starting, Started, and so on.
Executing on [read only]	The MID Server that is executing the event collector

4. Right-click in the header and select **Save**.
5. Test the connection between the vCenter and the MID Server: Click **Test parameters**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A pop-up displays **Parameters verified** for success. The **Status** field does not change. To update a parameter (for example, to change the MID Server or vCenter), change the setting and then click **Update parameters**.

6. Click the **Start** related link to start collecting events for vCenter.
   - To confirm that the instance is receiving events, verify that the **Status** field displays **Started**.
   - Click **Stop** to stop collecting events for a particular vCenter.
   - Click **Restart** to restart event collection.

   vCenter events are being populated in *sn_cmp_cloud_event* table.

Result

- If the vCenter has events in the Event Manager and the same events are configured, Cloud Provisioning and Governance gathers the events and take action on the CIs accordingly.
- Identification and Reconciliation Engine (IRE) changes the CI state based on event payload received and the IRE cannot be customized.

Create a VMware Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create a VMware Discovery schedule.

Before you begin

Verify the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- Configure a vSphere user in the VMware vSphere Client with the required roles and permissions. This grants the MID Server access to the cloud resources you plan to discover.
- You have created suitable credentials in the ServiceNow AI Platform to connect with the vSphere user that you want to use for discovery. For more information, see Discovery for VMware.
- You're using Discovery Admin Workspace v1.13.0 or later.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **VMware** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters pro vide failover protection and load balancing be tween MID Servers. See Configure a MID Serv [illegible] for more information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Automatically select a MID server	An available MID Server is automatically selected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validated, up, and have the Cloud Management capability are displayed.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your instance. Once an account is selected, the **Account name**, **UUID**, and **URL** fields auto-populate.
New cloud account	Add a cloud account to your instance by entering an account name, UUID, and URL. **Note:** ○ You can access the UUID in the VMware vSphere Client via **VMs** > **Manage Columns** > **UUID** ○ Example URL: `https://<openstack dns>:<port>`

   Even if the cloud account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**
9. Select **Next**.
10. Select datacenters for the discovery schedule.

    You can choose to include all datacenters linked to the cloud account or select specific datacenters. There's no restriction on the number of datacenters that you can add.

    a. **Optional:** To discover specific datacenters, select the datacenters from the **Select the data centers** drop-down list.

    b. **Optional:** If the datacenters don't appear, select **Refresh data centers**.
11. Select **Next**.
12. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.
13. Select **Next**.
14. Configure the Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the schedule runs.
Run on demand	The schedule only runs when triggered manually. The schedule can be triggered by selecting **Finish and run** at the end of this setup, or by navigating to the schedule in the Discovery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Discovery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Enable event-driven discovery	If Cloud Provisioning and Governance makes a life-cycle state or configuration change to a VMware resource, the Configuration Management Database (CMDB) is updated without running the full discovery schedule. Only the affected resources are processed. For more information, see Configure the VMware Events service to auto-update the CMDB.
Finish and run	After being selected, all the information provided is validated. Then, a Discovery schedule is created in the background, a Discovery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information provided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_schedule] table. You can edit the schedule information or run it by selecting **Discover Now**.

Discovery for Alibaba Cloud

Alibaba Cloud discovery is one of the overall Cloud discovery offerings within the IT Operations Management (ITOM) Visibility framework. It’s an automated process used to scan and identify Alibaba Cloud resources within your organization's cloud infrastructure. This discovery process is critical for maintaining an accurate and trustworthy data foundation—the Configuration Management Database (CMDB).

Pattern-based cloud discovery

Using Discovery and Service Mapping Patterns to perform horizontal discovery enables you to find and map your organization's Alibaba Cloud resources. You can discover Alibaba Cloud metadata including:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Cloud service accounts.
- Datacenters.
- Availability zones.
- Hardware types.

Patterns also support OS level discovery, for example OS images.

Discovery and Service Mapping Patterns create configuration items (CIs) for your Alibaba Cloud resources. Additionally, patterns discover the relationships between your organization's Alibaba Cloud resources, such as Hosted On :: Hosts.

See Alibaba Cloud discovery using patterns to learn about all Alibaba Cloud resources you can discover using Patterns.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Alibaba Cloud discovery configuration

The basic steps to configure pattern-based discovery for Alibaba Cloud involve preparation on the ServiceNow AI Platform side like installing necessary plugins and setting up credentials. The discovery_admin role in ServiceNow AI Platform is required for creating Alibaba Cloud API credentials and service accounts.

The discovery process requires configuration within Alibaba Cloud, like setting up Identity and Access Management roles. The discovery permissions of Alibaba Cloud users are determined by their access levels within Alibaba Cloud.

To promote proper discovery, the Alibaba Cloud user must have at least read-only access to the necessary Alibaba Cloud services.

Alibaba Cloud user discovery permissions

Alibaba Cloud user	Discovery permissions
Root Account (Master Account)	Full access to all Alibaba Cloud resources and services, including Elastic Compute Service (ECS), Object Storage Service (OSS), Relational Database Service (RDS), and Resource Access Management (RAM). Can create and manage RAM users, assign permissions, and perform billing operations.
RAM user	Access to specific Alibaba Cloud resources and services based on assigned policies. Can be granted read-only access for discovery purposes.
RAM Role (AssumedRoleUser)	Temporary access to Alibaba Cloud resources and services based on assumed role policies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alibaba Cloud user discovery permissions (continued)

Alibaba Cloud user	Discovery permissions
	Useful for cross-account access, temporary access, or access by ECS instances.

Alibaba Cloud discovery users and tasks

Typical persona	Roles and permissions	Responsibility	Link to detailed documentation
ServiceNow administrator or IT Implementation Specialist	admin	Install the store applications and update them on every store release: • Discovery • Discovery and Service Mapping Patterns • Visibility content. • CMDB CI Class Models • Discovery Admin Workspace	IT Operations Management
ServiceNow administrator	admin	• Create a MID Server user. • Assign the MID Server role to the user.	Create the MID Server user and grant the role
ServiceNow administrator	admin, mid_server	Install a MID Server.	• Install a MID Server on Linux • Install a MID Server on Windows • Install and configure MID Servers to access cloud environments
ServiceNow administrator	admin	Validate that the MID Server is installed correctly.	Validate the MID Server
ServiceNow administrator	admin	Assigning users with discovery_admin roles and giving them permission for discovery.	Managing roles

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alibaba Cloud discovery users and tasks (continued)

Typical persona	Roles and permissions	Responsibility	Link to detailed documentation
Cloud administrator or Discovery administrator	discovery_admin	Creating Alibaba Cloud service accounts	Set up Alibaba Cloud service accounts
Cloud administrator or Discovery administrator	The person configuring the API credentials must have the discovery_admin role in ServiceNow and must have access to the Alibaba Cloud *Access Key ID* and *Access Key Secret*.	Configuring Alibaba Cloud API credentials	Create Alibaba Cloud API Credentials
Discovery administrator	discovery_admin	Use Discovery and Service Mapping Patterns	Alibaba Cloud discovery using patterns
Discovery administrator	discovery_admin	Set up a discovery schedule for Alibaba Cloud	• Create a Discovery schedule for Alibaba Cloud • Create an Alibaba Cloud Discovery schedule in Discovery Admin Workspace

Create Alibaba Cloud API Credentials

Create Alibaba Cloud API credentials on the ServiceNow AI Platform to enable access to your Alibaba Cloud resources during Alibaba Cloud discovery.

Before you begin

Verify that your Alibaba Cloud user has at least read-only access to all Alibaba Cloud services. For information on access permissions, consult your system administrator.

Retrieve the following properties from the Alibaba Cloud Console:

- *AccessKey ID*
- *AccessKey secret*

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Credentials**.
2. Select **New**.
3. Select **Alibaba Cloud API Credentials**.
4. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alibaba Cloud API Credentials form

Field	Value
Name	Unique and descriptive name for the Alibaba Cloud credentials.
Active	Option that makes the credential active.
AccessKey ID	The AccessKey ID for your Alibaba Cloud account.
AccessKey secret	The AccessKey secret used to authenticate and securely sign API requests for your Alibaba Cloud account.

5. Select **Submit**.

What to do next

Set up Alibaba Cloud service accounts

Set up Alibaba Cloud service accounts

Create Alibaba Cloud infrastructure service accounts on the ServiceNow AI Platform to access your Alibaba Cloud account during Alibaba Cloud discovery.

Before you begin

Verify that Alibaba Cloud API credentials have been created. For more information, see Create Alibaba Cloud API Credentials.

Retrieve the *Account ID*, which can be found under **Root Account** or **RAM User** in the Alibaba Cloud Console. For more information, see the Alibaba Cloud documentation.

Role required: discovery_admin

Procedure

1. In the navigation filter, enter `cmdb_ci_cloud_service_account.list`.
2. Select **New**.
3. On the form, fill in the fields.

Cloud Service Account New record form

Field	Description
Name	A unique name for the Alibaba Cloud service account.
Account Id	Compartment ID, which is the Alibaba Cloud Account Identifier of the compartment associated with this service account.
Discovery Credentials	Name of credentials created for an Alibaba Cloud.
Datacenter URL	URL of the datacenter.
Datacenter Type	The Alibaba Cloud datacenter [cmdb_ci_Alibaba_datacenter].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Accessor account	Account name associated to the discovery credentials.

4. Select **Submit**.

What to do next
Create a schedule for Alibaba Cloud in Discovery or Discovery Admin Workspace.

Create a Discovery schedule for Alibaba Cloud

Create a Discovery schedule to run Alibaba Cloud discovery.

Before you begin
Role required: discovery_admin

Procedure

1. Navigate to **Discovery** > **Discovery Schedules**.

   Note: Starting with v1.13.0, you can also create an Alibaba Cloud in Discovery Admin Workspace. For more information, see Create an Alibaba Cloud Discovery schedule in Discovery Admin Workspace.

2. Select the **Cloud Discovery** button.
3. In the **Provider** field, select **Alibaba**.
4. Provide a name in the **Schedule Name** field.
5. Choose the method of selecting the MID Server from the **MID Selection Type** drop-down list.
   - **Auto-Select MID Server**: Automatically chooses the most appropriate MID Server.
   - **Specific MID Server**: A specific MID Server that you choose from the **MID Server** drop-down list.
   - **Specific MID Cluster**: A MID Server cluster that you choose from the **MID Cluster** drop-down list.
6. Either create a new account or select an existing one.

Existing or new account	Actions
Add account	**a.** Select **Add Account**. **b.** Provide a name in the **Name** field. **c.** Enter the ID in the **Account ID** field. **d.** In the **Credentials** field, select existing cre dentials or create new credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Existing or new account	Actions
	▪ To use existing credentials, select the cre dential from the drop-down list. ▪ To create new credentials, select the plus button and fill in the fields. For more information, see Create Alibaba Cloud API Credentials. **e.** Enter the URL of a datacenter in the **Data center URL** field.
Select Account	**a.** Select **Select Account**. **b.** Select the account from the drop-down list. The form fields are populated with the infor mation from the selected account.

7. **Optional:** Validate the account by selecting **Test Account**.
8. Select **Next**.
9. On the Select Datacenters page, determine whether to select specific datacenters to be discovered or to discover all datacenters.
   - Manually select the datacenters to be discovered by moving them into the **Selected** column.
   - Discover all available datacenters by turning on the **Discover all datacenters listed below** option.
10. Select **Next**.
11. **Optional:** On the **Discover Virtual Machines** page, set up the running of VM OS Discovery.
    a. Retain the **Discover VMs by IP address** option.
    b. Choose the method of selecting the MID Server from the **MID Selection Type** drop-down list.
12. Select **Next**.
13. On the Create Schedule page, determine when and how often to run your discovery.
    a. Activate the Discovery schedule by turning the **Active** option on.
    b. Choose the discovery frequency from the **Frequency** drop-down list.
    c. Set the **Start time** for the discovery schedule.
    d. **Optional:** Automatically cancel discovery if it exceeds a specified time limit by selecting the **Cancel discovery if longer than:** check box and setting the time limit.
14. Determine whether to create the schedule or to create the schedule and run discovery.
    - To create the schedule but not run discovery, select **Finish** from the Finish and Run drop-down list.
    - To create the schedule and run discovery, select **Finish and Run.**

Create an Alibaba Cloud Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create an Alibaba Cloud Discovery schedule.

Before you begin
Verify the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- The MID Server used for Cloud Discovery is configured with the Cloud Management capability.
- You have configured a Alibaba Cloud Random Access Memory (RAM) user or RAM role in the Alibaba Cloud console. The RAM user or role must have the permissions required for the MID Server to access the cloud resources you plan to discover.
- You have created suitable Alibaba Cloud API credentials in the ServiceNow AI Platform. For more information, see Create Alibaba Cloud API Credentials.
- You're using Discovery Admin Workspace v1.13.0 or later.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Cloud-based discovery** and select **Continue**.
4. Enter a unique and descriptive name for the Discovery schedule.
5. Select **Alibaba** as the cloud provider.
6. Choose a MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server. Clusters provide failover protection and load balancing between MID Servers. See Configure a MID Server for more information.
Automatically select a MID server	An available MID Server is automatically selected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Choose a MID Server to run the Discovery schedule. Only MID Servers that are validated, up, and have the Cloud Management capability are displayed.

7. Select **Next**.
8. Select a cloud account.

Option	Description
Existing cloud account	Select an existing cloud account on your instance. Once an account is selected, the **Account name**, **Account ID**, **Discover Gov Cloud regions**, and **Datacenter URL** fields auto-populate.
New cloud account	Add a cloud account to your instance by entering an account name and ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Even if the cloud account was previously verified, you can reverify if policies or permissions may have changed by selecting **Verify information.**

9. Select **Next**.
10. Select cloud regions for the discovery schedule.

    You can choose to include all regions linked to the cloud account or select specific regions. There's no restriction on the number of cloud regions that you can add.

    a. **Optional:** To discover specific cloud regions, select the check box for the cloud region you want to scan from the **Available cloud regions** list.

    b. **Optional:** If the cloud regions don't appear, select **Refresh cloud regions**.
11. Select **Next**.
12. **Optional:** To complement your cloud schedule with an IP-based Discovery schedule, enable the **Create a server discovery schedule** toggle.

    a. **Optional:** Select a MID Server or cluster that can discover the Virtual Machines (VM) and their attributes.
13. Select **Next**.
14. Configure the Discovery schedule.

Option	Description
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the sched ule runs.
Run on demand	The schedule only runs when triggered manu ally. The schedule can be triggered by select ing **Finish and run** at the end of this setup, or by navigating to the schedule in the Discov ery Schedules [discovery_schedule] table and selecting **Discover Now**.
Run after series	The schedule runs after another existing Dis covery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Set discovery time out	If the schedule exceeds the maximum run time, it's canceled. After toggling, configure the runtime threshold.
Finish and run	After being selected, all the information pro vided is validated. Then, a Discovery sched ule is created in the background, a Discov ery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information pro vided is validated. A Discovery schedule is created, and you're redirected to its entry in the Discovery Schedules [discovery_sched ule] table. You can edit the schedule informa tion or run it by selecting **Discover Now**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud discovery reporting

Cloud discovery reporting provides a user-friendly interface for navigating and analyzing your cloud resource inventory. It offers drill-down and slice-and-dice views to gain granular insights into your cloud resource inventories, with details obtained from the discovery. It provides an easier and faster way to locate and identify the state and location of cloud resources, streamlining your management tasks.

Learn about Cloud Discovery reporting:

- Analytics and Reporting for Cloud Resources
- Cloud Resources dashboard
- Cloud Resources Explorer

Note: For better results, use the Cloud Resources Explorer instead of the Platform Analytics Solutions featured on the Cloud Resources dashboard.

Analytics and Reporting for Cloud Resources

Platform Analytics Solutions contain prepackaged Performance Analytics and Reporting content for use with other ServiceNow AI Platform products.

Performance Analytics Content Pack for Cloud Resources overview

The Performance Analytics Content Pack for Cloud Resources shows the aggregated view of all your Cloud Resources. You can filter on specific information for more details. There are also analytics tables and views that associate tags with resources so that you can build your own custom dashboards and reports.

The Performance Analytics Content Pack for Cloud Resources works with Discovery and Cloud Provisioning and Governance and it supports the following cloud providers:

- Amazon Web Services
- Microsoft Azure public clouds
- VMware vCenter
- Google Cloud Platform (GCP) (Supported with Performance Analytics Content Pack for Cloud Resources version 1.2.1 and higher).

The aggregate count of each cloud resource and the latest Cloud Discovery error logs are displayed on the dashboard. You can filter the cloud resources by the service account, datacenter, or resource group (only for Microsoft Azure) under a specific provider and view the resource count and the resources.

The base-system Performance Analytics Content Pack for Cloud Resources dashboard doesn't contain any widget for the GCP resources. You can create the reports and widgets by consuming data from the appropriate tables. For more information on the tables that contain the GCP resources data, see Cloud Resources dashboard. Alternatively, you can install the Cloud Discovery Workspace application to view the GCP reports and widgets. For more information, see Cloud Discovery Workspace.

For each cloud provider, a separate scheduled job is triggered every 30 minutes. The latest CI data is updated in the cloud resource analytic tables, which allows the dashboard to be more realistic and have the updated count. The information about each scheduled job is available in the Cloud Dashboard Job Status [sn_disco_cd_cloud_dashboard_job_status] table. The job details including the provider type, status, start time, end time, and duration are stored in this table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The reports for the cloud dashboard widgets are generated from the cloud resource analytical tables. The analytical tables are populated through the scheduled job. The scheduled job queries the new database views, and inserts and updates the records into the analytical tables.

The scheduled job can run for a maximum of 360 minutes. The maximum job timeout value is set using the system property (sn_disco_cd.scheduled_job_time_out_in_mins). If the scheduled job exceeds 360 minutes, or the time set in this property, then the job is canceled.

Required roles

The following roles and tasks are associated with this solution:

- ServiceNow AI Platform administrator (admin): Install and activate this Analytics and Reporting Solution and can make any necessary changes to system properties.
- Performance Analytics administrator (pa_admin): Review the indicators, breakdowns, widgets, and dashboards. Set up and start data collection and share the dashboards with appropriate stakeholders.
- Cloud Dashboard user (sn_disco_cd.user): Crew the cloud resource widgets in the Cloud Dashboard and has read-only access on the Cloud Resource analytics tables and the dashboard reports.
- Cloud Dashboard administrator (sn_disco_cd.admin): View, add, or edit the dashboard widgets and reports and has read-only access on the Cloud Resource analytics tables. By default (sn_disco_cd.admin) contains the (sn_disco_cd.user) role and (discovery_admin) contains the (sn_disco_cd.admin) role.

Activate Analytics and Reporting for Cloud Resources

Activate the Performance Analytics Content Pack for the Cloud Resources ServiceNow Store application for the Analytics and Reporting Solution.

Before you begin

Make sure the Cloud Provisioning and Governance Core (com.snc.cloud.core) plugin is installed and activated and that you have upgraded to Orlando or later release. Cloud Provisioning and Governance Core is available with Discovery as well as the Cloud Provisioning and Governance Core plugin.

Role required: Admin

About this task

Procedure

1. Navigate to **System Applications** > **All Applications** > **ServiceNow Store**.
2. Browse for Performance Analytics Content Pack for Cloud Resources and verify that you have the necessary valid entitlements.

   You need a HI account and permission to request applications for the instances under consideration. The HI account and the instance must be part of the same company. For more information, see Install a ServiceNow Product on the ServiceNow Store.
3. Install the `Performance Analytics Content Pack for Cloud Resources`application.

Cloud Resources dashboard

The Cloud Resources dashboard shows the aggregated view of all your cloud resources, their combined analytical views, and a visual summary and detail of all your cloud resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: The Cloud Resources dashboard has been replaced by the Cloud Resources explorer in the Cloud Discovery Workspace store application. For more information, see Cloud Resources Explorer.

Cloud Resources dashboard tabs

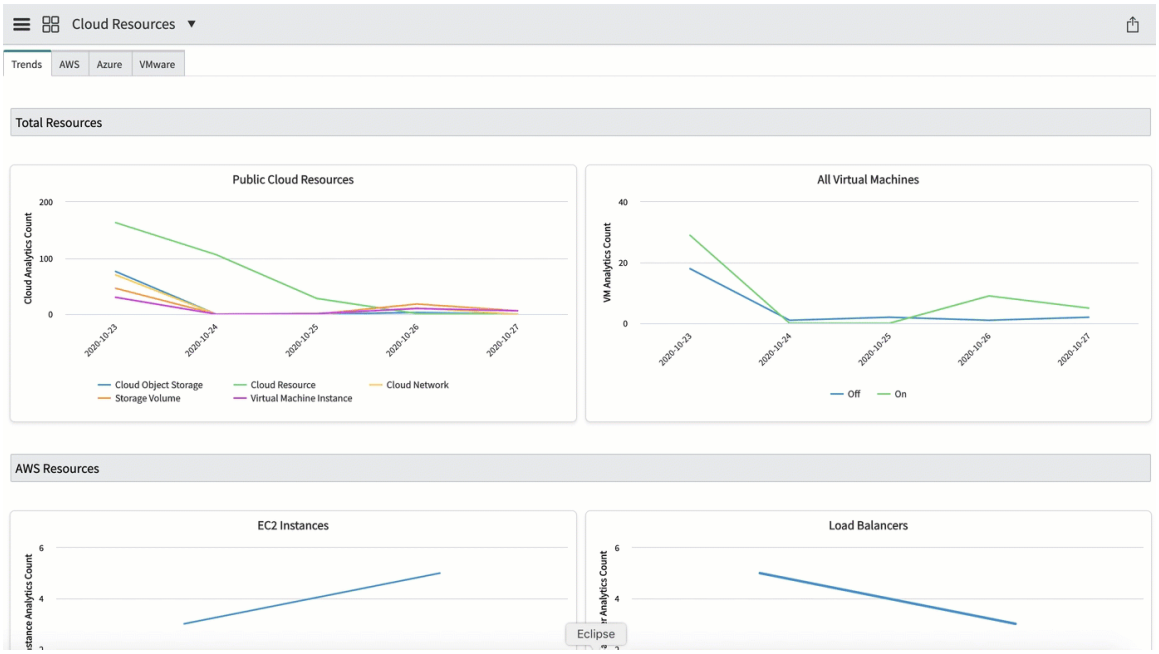

Use cases

For examples of how someone in your organization could use this dashboard, see the use cases.

User	Dashboard use
Cloud user	Discovery customers have an enormous number of CIs obtained from Cloud Discovery which can be viewed in one place.
Cloud user with sn_disco_cd.admin role	Use the analytics tables that are available to build custom dashboards. These tables are specific baseline database views with resource tags. See Customize the Cloud Resources dashboard for more information.

Data visualizations

Title	Type	Source table
Cloud Networks	Line	VMware Network Analytics [sn_disco_cd_vmware_network_analytics]
Cloud Networks	Single Score	VMware Network Analytics [sn_disco_cd_vmware_network_analytics]
Public Cloud Resources	Line	Cloud Analytics [sn_disco_cd_analytics]
Clusters	Single Score	VMware Cluster Analytics [sn_disco_cd_vmware_cluster_analytics]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table
Clusters	Line	VMware Cluster Analytics [sn_disco_cd_vmware_cluster_analytics]
Compute Security Groups	Single Score	AWS Security Group Analytics [sn_disco_cd_aws_sec_group_analytics]
Databases	Single Score	Azure Database Analytics [sn_disco_cd_azure_database_analytics]
Databases	Line	Azure Database Analytics [sn_disco_cd_azure_database_analytics]
Datastore Disk	Single Score	VMWare Datastore Disk Analytics [sn_disco_cd_vmware_ds_disk_analytics]
Datastores	Single Score	VMware Datastore Analytics [sn_disco_cd_vmware_datastore_analytics]
Disks	Line	Azure Disk Analytics [sn_disco_cd_azure_disk_analytics]
Disks	Single Score	Azure Disk Analytics [sn_disco_cd_azure_disk_analytics]
Distributed Virtual Port Groups	Single Score	VMware Distributed Virtual Port Group Analytics [sn_disco_cd_vmware_dv_port_grp_analytics]
Distributed Virtual Switches	Single Score	VMWare Distributed Virtual Switch Analytics [sn_disco_cd_vmware_dvs_analytics]
EC2 Instances	Single Score	AWS EC2 Instance Analytics [sn_disco_cd_aws_ec2_instance_analytics]
EC2 Instances	Line	AWS EC2 Instance Analytics [sn_disco_cd_aws_ec2_instance_analytics]
EC2 Instances	Pie	AWS EC2 Instance Analytics [sn_disco_cd_aws_ec2_instance_analytics]
ESX Servers	Single Score	ESX Server Analytics [sn_disco_cd_vmware_esx_analytics]
ESX Servers	Line	ESX Server Analytics [sn_disco_cd_vmware_esx_analytics]
Folders	Single Score	vCenter Folder Analytics [sn_disco_cd_vmware_folder_analytics]
Functions	Single Score	Azure Function Analytics [sn_disco_cd_azure_function_analytics]
Generic Cloud Resources	Single Score	AWS Generic Cloud Resource Analytics [sn_disco_cd_aws_gen_resource_analytics]
Generic Cloud Resources	Line	AWS Generic Cloud Resource Analytics [sn_disco_cd_aws_gen_resource_analytics]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table
Generic Cloud Resources	Single Score	Azure Generic Cloud Resource Analytics [sn_disco_cd_azure_gen_resource_analytics]
Images	Single Score	AWS Image Analytics [sn_disco_cd_aws_image_analytics]
Images	Single Score	Azure Image Analytics [sn_disco_cd_azure_image_analytics]
Lambdas	Single Score	AWS Lambda Analytics [sn_disco_cd_aws_lambda_analytics]
Load Balancers	Line	AWS Load Balancer Analytics [sn_disco_cd_aws_loadbalancer_analytics]
Load Balancers	Single Score	AWS Load Balancer Analytics [sn_disco_cd_aws_loadbalancer_analytics]
Load Balancers	Single Score	Azure Load Balancer Analytics [sn_disco_cd_azure_loadbalancer_analytics]
Load Balancers	Line	Azure Load Balancer Analytics [sn_disco_cd_azure_loadbalancer_analytics]
Network Interfaces	Single Score	AWS Network Interface Analytics [sn_disco_cd_aws_nw_interface_analytics]
Network Interfaces	Single Score	Azure Network Interface Analytics [sn_disco_cd_azure_nw_interface_analytics]
Network Security Groups	Single Score	Azure Network Security Groups Analytics [sn_disco_cd_azure_nw_sec_group_analytics]
Relational Database Services (RDS)	Line	AWS RDS Analytics [sn_disco_cd_aws_rds_analytics]
Relational Database Services (RDS)	Single Score	AWS RDS Analytics [sn_disco_cd_aws_rds_analytics]
S3 Buckets	Single Score	AWS S3 bucket Analytics [sn_disco_cd_aws_s3_bucket_analytics]
Storage Volumes	Single Score	AWS Storage Volume Analytics [sn_disco_cd_aws_storage_volume_analytics]
Storage Volumes	Line	AWS Storage Volume Analytics [sn_disco_cd_aws_storage_volume_analytics]
Virtual Machines	Line	Azure VM Analytics [sn_disco_cd_azure_vm_analytics]
Virtual Machines	Single Score	Azure VM Analytics [sn_disco_cd_azure_vm_analytics]
Virtual Machines	Pie	Azure VM Analytics [sn_disco_cd_azure_vm_analytics]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table
Virtual Machines	Line	VMware VM Analytics [sn_disco_cd_vmware_vm_analytics]
Virtual Machines	Pie	VMware VM Analytics [sn_disco_cd_vmware_vm_analytics]
Virtual Machines	Single Score	VMware VM Analytics [sn_disco_cd_vmware_vm_analytics]
All Virtual Machines	Line	VM Analytics [sn_disco_cd_vm_analytics]
Virtual Networks	Single Score	Azure Virtual Network Analytics [sn_disco_cd_azure_virtual_nw_analytics]
Virtual Private Clouds (VPC)	Line	AWS VPC Analytics [sn_disco_cd_aws_vpc_analytics]
Virtual Private Clouds (VPC)	Single Score	AWS VPC Analytics [sn_disco_cd_aws_vpc_analytics]
VM Templates	Single Score	VMWare VM Template Analytics [sn_disco_cd_vmware_vm_template_analytics]

Access the Cloud Resources dashboard

To access the Cloud Resources dashboard, navigate to **Discovery** > **Cloud Resource Dashboard** or **Cloud Provisioning and Governance** > **Cloud Resource Dashboard**.

Customize the Cloud Resources dashboard

Customize your dashboard In each cloud provider tab by filtering these attributes:

- Service Account: All the service accounts that are specific to the cloud provider tab selected are listed.
- Datacenters: The datacenters are populated based on the cloud provider.
- Resource Group: This filter is only applicable for Microsoft Azure and restricts objects and resources to the chosen resource group.

Use these database views with resource tags to create your own dashboard.

AWS Database Views with Resource Tags

Database View Table Name	Label
sn_disco_cd_aws_api_gateway_analytics_tags_db_view	AWS API Gateway Analytics with Tags
sn_disco_cd_aws_ec2_instance_analytics_tags_db_view	AWS EC2 Instance Analytics with Tags
sn_disco_cd_aws_gen_resource_analytics_tags_db_view	AWS Cloud Resource Analytics with tags
sn_disco_cd_aws_image_analytics_tags_db_view	AWS Image Analytics with Tags
sn_disco_cd_aws_lambda_analytics_tags_db_view	AWS Lambda Analytics with Tags

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Database Views with Resource Tags (continued)

Database View Table Name	Label
sn_disco_cd_aws_loadbalancer_analytics_tags_db_view	AWS Load Balancer Analytics with Tags
sn_disco_cd_aws_nw_interface_analytics_tags_db_view	AWS NW Interface Analytics with Tags
sn_disco_cd_aws_rds_analytics_tags_db_view	AWS RDS Analytics with Tags
sn_disco_cd_aws_s3_bucket_analytics_tags_db_view	AWS S3 bucket Analytics with Tags
sn_disco_cd_aws_sec_group_analytics_tags_db_view	AWS Security Group Analytics with Tags
sn_disco_cd_aws_storage_volume_analytics_tags_db_view	AWS Storage Volume Analytics with Tags
sn_disco_cd_aws_vpc_analytics_tags_db_view	AWS VPC Analytics with Tags
sn_disco_cd_aws_zone_analytics_tags_db_view	AWS Zones Analytics with tags

Microsoft Azure Database Views with Resource Tags

Database View Table Name	Label
sn_disco_cd_azure_database_analytics_tags_db_view	Azure Database Analytics with Tags
sn_disco_cd_azure_disk_analytics_tags_db_view	Azure Disk Analytics with Tags
sn_disco_cd_azure_function_analytics_tags_db_view	Azure Function Analytics with Tags
sn_disco_cd_azure_gen_resource_analytics_tags_db_view	Azure Cloud Resource Analytics with Tags
sn_disco_cd_azure_image_analytics_tags_db_view	Azure Image Analytics with Tags
sn_disco_cd_azure_loadbalancer_analytics_tags_db_view	Azure Load Balancer Analytics with Tags
sn_disco_cd_azure_nw_interface_analytics_tags_db_view	Azure NW Interface Analytics with Tags
sn_disco_cd_azure_nw_sec_group_analytics_tags_db_view	Azure Security Group Analytics with Tags
sn_disco_cd_azure_virtual_nw_analytics_tags_db_view	Azure Virtual NW Analytics with Tags
sn_disco_cd_azure_vm_analytics_tags_db_view	Azure VM Analytics with Tags
sn_disco_cd_azure_zone_analytics_tags_db_view	Azure Zones Analytics with Tags

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware Database Views with Resource Tags

Database View Table Name	Label
sn_disco_cd_vmware_cluster_analytics_tags_db_view	VMware Cluster Analytics with Tags
sn_disco_cd_vmware_datastore_analytics_tags_db_view	VMware Datastore Analytics with Tags
sn_disco_cd_vmware_ds_disk_analytics_tags_db_view	VMware DatastoreDisk Analytics with Tags
sn_disco_cd_vmware_dvs_analytics_tags_db_view	VMware DVS Analytics with Tags
sn_disco_cd_vmware_dv_port_grp_analytics_tags_db_view	VMware DV Port Group Analytics with Tags
sn_disco_cd_vmware_esx_analytics_tags_db_view	VMware ESX Analytics with Tags
sn_disco_cd_vmware_folder_analytics_tags_db_view	VMware Folder Analytics with Tags
sn_disco_cd_vmware_network_analytics_tags_db_view	VMware Network Analytics with Tags
sn_disco_cd_vmware_vm_analytics_tags_db_view	VMware VM Analytics with Tags
sn_disco_cd_vmware_vm_template_analytics_tags_db_view	VMware VM Template Analytics with Tags

The below resources that are supported for GCP are populated by GCPResourcePopulateScheduledJob.

Google Cloud Platform (GCP) Analytics tables and their corresponding Configuration Management Database (CMDB) Configuration Item (CI) resource tables

Analytic Table	CMDB CI resources
GCP VM Analytics [sn_disco_cd_gcp_database_analytics]	Virtual Machine Instance [cmdb_ci_vm_instance]
GCP Storage Volume Analytics [sn_disco_cd_gcp_storage_volume_analytics]	Storage Volume [cmdb_ci_storage_volume]
GCP Database Analytics [sn_disco_cd_gcp_database_analytics]	Cloud Database [cmdb_ci_cloud_database]
GCP Virtual Network Analytics [sn_disco_cd_gcp_virtual_nw_analytics]	Cloud Network [cmdb_ci_network]
GCP Load Balancer Analytics [sn_disco_cd_gcp_loadbalancer_analytics]	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
GCP Image Analytics [sn_disco_cd_gcp_image_analytics]	Cloud OS Images [cmdb_ci_cloud_os_images]
GCP Generic Cloud Resource Analytics [sn_disco_cd_gcp_gen_resource_analytics]	Cloud Resource [cmdb_ci_cmp_resource]
GCP Network Interface Analytics [sn_disco_cd_gcp_nw_interface_analytics]	Cloud Mgmt Network Interface [cmdb_ci_nic]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scheduled jobs

Scheduled jobs installed

Job	Description
GCPResourcePopulateScheduledJob	Scheduled job that fetches and populates the GCP CI data in the analytics table.
AWSResourcePopulateScheduledJob	Scheduled job that fetches and populates the AWS CI data in the analytics table.
AzureResourcePopulateScheduledJob	Scheduled job that fetches and populates the Microsoft Azure CI data in the analytics table.
VMwareResourcePopulateScheduledJob	Scheduled job that fetches and populates the VMware CI data in the analytics table.

Since a separate scheduled job is triggered every 30 minutes for each cloud provider, you can verify if the last job was successful. If the scheduled job was not successful, you can manually trigger the job by going to the **System Definition** > **Scheduled Jobs**. Select the job for the affected cloud provide and click **Execute Now**.

The refresh jobs run every 30 minutes to populate and update the date in the analytics tables. You can increase this frequency to get more updated data.

Synchronization is needed if:

- The Resource Status was updated from **installed** to another status (for example, retired or canceled).
- Any new resources were recently discovered.

Synchronization is not needed if:

- No resources were updated.
- There are new resources which were not recently discovered.

Important to note:

- Only one job can run at a time. For example, if an AWS job is already running, you cannot trigger another AWS job.
- If the LDC and service account fields of a CI resource is empty, the resource may not be associated with any LDC in the CI Relationship [cmdb_rel_ci] table or the resource is a generic cloud resource of IAM type and the fields are empty.
- Only the resources with installed status set to installed are counted and represented in the cloud resources dashboard widgets.
- If a new resource is discovered after the scheduled job run, then the next scheduled job adds the resource details in the analytical tables.
- If a job is marked as failed in the Cloud Dashboard Job Status [sn_disco_cd_cloud_dashboard_job_status] table, the reason for the failure is logged in the system logs. Filter the system log table by the job sys_id to fetch the error message.
- If IP-based discovery is used to discover the vCenter, the service account field of the VMware resource is left blank.
- The System Administrator should not delete the records from the analytics table since records deleted from the analytics table are not populated again by the scheduled job. Deletion of records leads to data inconsistency in the analytics table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To ensure complete refresh of the VMware records (including the service accounts) after changing the VMware resource discovery mechanism, delete the VMware records from the Cloud Dashboard Job Status [sn_disco_cd_cloud_dashboard_job_status] table and run the schedule job. In the Cloud Dashboard Job Status table, the provider field of all the VMware resources is set to *vmware*. You can use either vCenter IP discovery or service account discovery to discover the VMware resources.

Cloud Resources Explorer

Use the ServiceNow® Cloud Resources Explorer to filter and visualize the distribution of resources across your multi cloud estate. Gain insight into operational information such as stale resources and cloud event inflow rates.

Important: Starting with the Zurich release, Cloud Discovery Workspace is being prepared for future deprecation. It will be hidden and no longer activated on new instances, but will continue to be supported. For similar reporting, refer to the Cloud asset explorer in the Cloud Discovery store application.

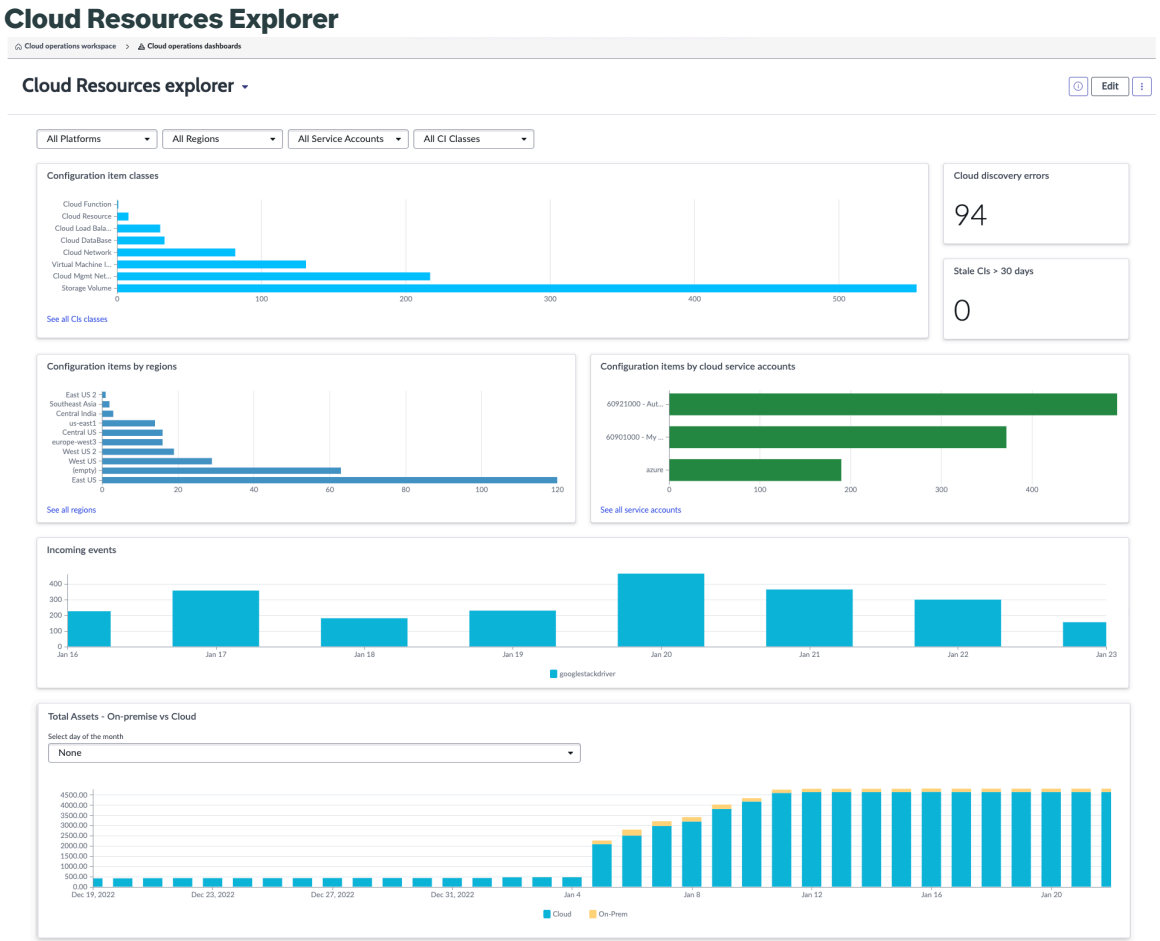

The Cloud Resources Explorer supports the following cloud providers:

- Amazon Web Services.
- Microsoft Azure public clouds.
- VMware vCenter.
- Google Cloud Platform (GCP) (Supported with Performance Analytics Content Pack for Cloud Resources version 1.2.1 and higher).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Required ServiceNow AI Platform roles

The sn_cloud_ops_ws.cloud_ops_admin role is required to view the Cloud Resources Explorer.

Access the Cloud Resources Explorer

To open the Cloud Resources Explorer, navigate to **All** > **Cloud Discovery Workspace** > **Cloud Resources Explorer**.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
sn_cloud_ops_ws.cloud_ops_admin	• Gain insight into the cloud resources inventory count and the current size of the Configuration Management Database (CMDB), and how many objects are under the purview. • Gain insight into fault-related information such as Cloud Discovery errors. • Gain insight into operational information such as the inflow and processing rate of the event. • Gain insight into the health of your CMDB. • Gain insight into the distribution of resources across the on-prem cloud and the public clouds.

Indicators

Cloud Discovery errors

Total Cloud Discovery errors encountered over time.

Stale CIs > 30 days

If a CI isn't updated in the last 30 days, it's considered as a stale CI.

A stale CI indicates that the CI exists in the Configuration Management Database (CMDB), but it may or may not exist in the datacenter. Identifying the stale CIs is important for maintaining a healthy CMDB.

Here are some of the scenarios that leads to the creation of stale CIs:

- The cloud resource doesn't exist, but its deletion isn't updated in the CMDB.
- The device may still be on the network, but the automated discovery process can't connect to the device to rediscover it.
- The device exists on the network and it's accessible, but the identification process isn't recognizing it uniquely. Therefore, Cloud Discovery has created a duplicate CI in the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data visualizations

Title	Type	Source table	Description
Configuration item classes	Horizontal bar report 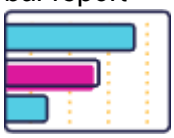	Cloud Analytics [sn_disco_cd_analytics]	This report displays the count of the top 10 CI classes discovered over time. You can select a report entry from the Configuration item classes report to view the analytics table on which the report entry is based. If you've selected a specific cloud provider, then the analytics table of that specific CI appears. Otherwise, the parent analytics table for that CI type appears. For example, selecting Google from the Cloud resources filter and the Storage Volume bar opens the dedicated Google Storage Volume CI analytics table. If you select All Platforms from the Cloud resources filter and select the Storage Volume bar from the report, the parent analytics table of the storage volume CI appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
Configuration items by regions	Horizontal bar report	Cloud Analytics [sn_disco_cd_analytics]	This report displays the distribution of CIs per region (datacenter).
Configuration items by cloud service accounts	Horizontal bar report	Cloud Analytics [sn_disco_cd_analytics]	This report displays the top 10-service accounts with the maximum number of discovered CIs.
Incoming events	Bar report	Cloud Event [sn_cmp_cloud_event]	This report displays the trend of cloud resource events received over time. To view the list of all the cloud events received over a given day, select the corresponding report entry.
Total cloud resources in on-prem vs public clouds	Stacked bar report	Daily resource count by providers [sn_disco_cd_daily_resource_cnt_by_provider]	This report displays the daily count of the resources of your organization that are spread across the on-prem cloud and the public clouds. Use this report to understand the public cloud adoption and Virtual Machine (VM) count over time. In this report, the VMware Cloud is

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
			considered on-prem.

To populate the Cloud Analytics [sn_disco_cd_analytics] table, the Cloud Resources Explorer uses the same set of scheduled jobs as the Performance Analytics Content Pack for Cloud Resources. For more information on the scheduled jobs, see Cloud Resources dashboard.

To view the details of the OS images discovered in the AWS, Microsoft Azure, and GCP clouds, Discovery and Service Mapping Patterns store app version 1.0.87 or later must be installed in the instance. Also, the *sn_cmdb_ci_class.use_single_cloud_os_image* property must be set to `true`.

Filters

Name	Type	Description
Cloud resources filter	Choice-based filter	Use the cloud resources filter to select and visualize the cloud resources data of interest. By default, the Cloud Resources Explorer displays data for all the cloud resources across the cloud providers of your organization. Only the cloud providers that you've selected in the Cloud provider selection page of the Cloud Discovery Workspace app appear in the cloud resources filter. You have the choice of the following cascading filters: • Platform: Select the cloud provider with the resources data that you want to visualize. • Region: Select the region or datacenter with the resources data that you want to visualize. • Service account: Select the service account with the resources data that you want to visualize. • CI class: Select the CI class with the resources data that you want to visualize. • Resource group: Select the resource group with the resources data that you want to visualize. This field is available only when `Azure` is selected from the Platform drop-down list. You can filter the Incoming events report based on the Platform filter only. The data displayed in the Cloud Discovery errors indicator isn't affected by the criteria selected in the Cloud resources filter.
Day navigator	Choice-based filter	Use the Day navigator to slice the Total cloud resources in on-prem vs public clouds report across the following dimensions: • None: The report displays the data collected over the last one year if available. • First day of the month: The report displays the historical data collected for the first day of the month over the last one year if available.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Type	Description
		• Mid of the month: The report displays the historical data collected for the mid day of the month over the last one year if available. • Last day of the month: The report displays the historical data collected for the last day of the month over the last one year if available. • Pick a date of the month: The report displays the historical data collected for the selected day of the month over the last one year if available.

Cloud discovery solutions comparison

Comparing cloud discovery solutions provides insights on the relative strengths of each solution. The comparison showcases the number of resource types supported by the solutions across AWS, Microsoft Azure, and GCP. The comparison can help you understand the capabilities of each solution and better manage your CMDB when using one or multiple methods.

Discovery and Service Mapping Patterns vs. Service Graph Connectors

ITOM Visibility offers two primary solutions for discovering cloud resources and mapping them into the Configuration Management Database (CMDB): Cloud Discovery and Service Mapping Patterns and Cloud Service Graph Connectors.

Patterns are the core of ITOM Visibility's native discovery capabilities. They provide a deeper, more dynamic, and holistic view of your infrastructure. Patterns use a variety of discovery methods: Agentless, agent-based, and cloud-native, to create a unified, business-aware view of your services. By leveraging machine learning, patterns automatically map technical components to their business context, creating rich relationships, and a comprehensive knowledge graph.

Service Graph Connectors, on the other hand, provides fast, point-in-time data ingesting from existing systems. They assist you quickly ingest data from specific sources like public cloud providers (AWS, Microsoft Azure, GCP), endpoint management systems, and security or network tools.

While Service Graph Connector provides rapid data import, Patterns (regularly updated on the ServiceNow Store) provide the foundation for ongoing, automated service mapping and deeper operational insights. This difference is crucial for solving complex challenges like certificate management, firewall auditing, and a wide range of other service and operational use cases.

The choice between patterns and connectors affects coverage, level of detail, how data is modeled in the CMDB, and frequency of update. While some resource types are supported by both solutions, the data might be populated in different CMDB CI classes. Understanding these differences helps you plan a consistent and effective discovery strategy.

By reviewing the coverage tables for AWS, Azure, and GCP, you can

- Identify which resources are supported by patterns, by connectors, or both.
- Understand how the same resource type may be handled differently (for example, patterns might bring in all zone data, while connectors might bring in only zones tied to an instance).
- Decide on a solution per provider, considering its CI coverage or data model. For example: Use connectors to discover your GCP resources, but use patterns for AWS and Microsoft Azure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison of Discovery and Service Mapping Patters and Service Graph Connectors

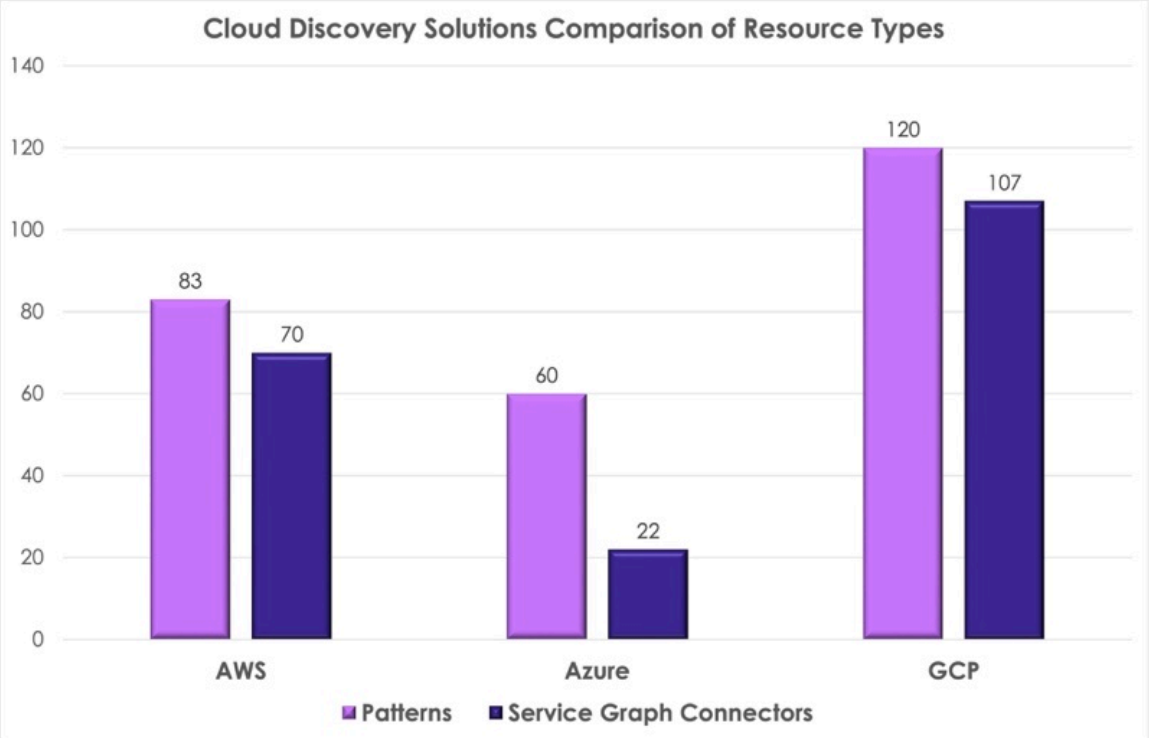

AWS discovery solutions comparison

ITOM Visibility applications discover a variety of AWS resources and populate the relevant configuration item (CI) classes in the Configuration Management Database (CMDB) with their attributes.

ITOM Visibility solutions

ITOM Visibility can offer you comprehensive visibility into your organization's cloud infrastructure. Each application offers unique strengths and capabilities, and by comparing them you can gain holistic insights on their capabilities and benefits.

ITOM Visibility solutions for discovering AWS

ITOM Visibility solution	Overview	Configuration
Discovery and Service Mapping Patterns and Visibility Content	Discovery patterns used by ITOM Visibility	• You must update your store apps to the latest release version. • You must provide ITOM Visibility with access permissions. For more information, see Setup AWS service accounts
Service Graph Connector for AWS	Service Graph Connector for AWS	Configure Service Graph Connector for AWS using the guided setup

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

The CI class, in most cases, is applicable to both solutions. If it's not applicable to one of the solutions, the solution is marked "N/A".

Discovery and Service Mapping Patterns vs Service Graph Connector for AWS

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
AWS::AppSync::API	Yes	No	cmdb_aws_appsync_api
AWS::Athena::WorkGroup	Yes	No	cmdb_aws_athena_workgroup
AWS::Backup::BackupPlan	Yes	No	cmdb_aws_backup_backup_plan
AWS::Backup::BackupVault	Yes	No	cmdb_aws_backup_backup_vault
AWS::CloudFront::Distribution	Yes	N/A	cmdb_aws_cloudfront_distribution
AWS::CloudFront::Distribution	N/A	Yes	cmdb_ci_cmp_resource
AWS::CloudWatch::Log	Yes	No	cmdb_aws_cloudwatch_log
AWS::CodeDeployment::Deployment	Yes	No	cmdb_aws_codedeploy_deployment
AWS::CodePipeline::Pipeline	Yes	N/A	cmdb_aws_codepipeline_pipeline
AWS::CodePipeline::Pipeline	N/A	Yes	cmdb_ci_cmp_resource
AWS::DMS::Endpoint	Yes	No	cmdb_aws_dms_endpoint
AWS::ElasticBeanstalk::Application	Yes	N/A	cmdb_aws_elastic_beanstalk_applicati
AWS::ElasticBeanstalk::Application	N/A	Yes	cmdb_ci_cmp_resource
AWS::EMR::Cluster	Yes	No	cmdb_aws_emr_cluster
AWS::EventBridge::EventBus	Yes	No	cmdb_aws_eventbridge_event_bus
AWS::FSx::Backup	Yes	No	cmdb_aws_fsx_backup
AWS::FSx::FileSystem	Yes	No	cmdb_aws_fsx_file_system
AWS::Glue::Database	Yes	No	cmdb_aws_glue_database
AWS::IAM::Policy	Yes	N/A	cmdb_aws_iam_policy
AWS::IAM::Policy	N/A	Yes	cmdb_ci_cmp_resource
AWS::IAM::Role	Yes	N/A	cmdb_aws_iam_role
AWS::IAM::Role	N/A	Yes	cmdb_ci_cmp_resource
AWS::IAM::User	Yes	N/A	cmdb_aws_iam_user
AWS::IAM::User	N/A	Yes	cmdb_ci_cmp_resource
AWS::KMS::Key	Yes	No	cmdb_aws_kms_key
AWS::MQ::Broker	Yes	No	cmdb_aws_mq_broker
AWS::MQ::Configuration	Yes	No	cmdb_aws_mq_configuration
AWS::S3::GlacierVault	Yes	No	cmdb_aws_s3_glacier_vault
AWS::SageMaker::NotebookInstance	Yes	No	cmdb_aws_sagemaker_notebook_insta

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for AWS (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
AWS::SecretsManager::Secret	Yes	No	cmdb_aws_secrets_manager_secret
AWS::SimpleEmailService::Identity	Yes	No	cmdb_aws_simple_email_service_iden
AWS::SQS::Queue	Yes	No	cmdb_aws_sqs_queue
AWS::StepFunctions::StateMachine	Yes	No	cmdb_aws_step_functions_state_mach
AWS::StorageGateway::Gateway	Yes	No	cmdb_aws_storage_gateway_gateway
AWS::SystemsManager::Document	Yes	No	cmdb_aws_systems_manager_docume
AWS::SystemsManager::ParameterStore	Yes	No	cmdb_aws_systems_manager_parame
AWS::TransferFamily::Server	Yes	No	cmdb_aws_transfer_family_server
AWS::VPC::FlowLog	Yes	No	cmdb_aws_vpc_flow_log
AWS::VPC::ManagedPrefixList	Yes	No	cmdb_aws_vpc_managed_prefix_list
AWS::XRay::SamplingRule	Yes	No	cmdb_aws_xray_sampling_rule
ec2:DescribeAvailabilityZones	Yes	Yes	cmdb_ci_availability_zone
ec2:DescribeRegions	Yes	Yes	cmdb_ci_aws_datacenter
AWS::Kinesis::Stream	Yes	No	cmdb_ci_aws_kinesis_data
AWS::KinesisAnalytics	Yes	No	cmdb_ci_aws_kinesis_data_analytics
AWS::KinesisFirehose	Yes	No	cmdb_ci_aws_kinesis_firehose
AWS::KinesisVideo	Yes	No	cmdb_ci_aws_kinesis_video
AWS::Redshift::Cluster	Yes	N/A	cmdb_ci_aws_redshift
AWS::Redshift::Cluster	N/A	Yes	cmdb_ci_cmp_resource
AWS::CertificateManager::Certificate	Yes	No	cmdb_ci_certificate
AWS::Cognito::UserPool	Yes	No	cmdb_ci_cloud_authentication
AWS::RDS::DBInstance	Yes	Yes	cmdb_ci_cloud_database
AWS::ECS::Cluster	Yes	No	cmdb_ci_cloud_ecs_cluster
AWS::ECS::Service	Yes	No	cmdb_ci_cloud_ecs_service
AWS::ECS::TaskDefinition	Yes	No	cmdb_ci_cloud_ecs_task_definition
AWS::Lambda::Function	Yes	Yes	cmdb_ci_cloud_function
AWS::ApiGateway::RestApi	Yes	Yes	cmdb_ci_cloud_gateway
ec2:DescribeInstanceTypes	Yes	Yes	cmdb_ci_cloud_hardware_type
AWS::EC2::Host	Yes	No	cmdb_ci_cloud_host
AWS::EC2::KeyPair	Yes	No	cmdb_ci_cloud_key_pair
AWS::ElasticLoadBalancingV2::LoadBalancer	Yes	Yes	cmdb_ci_cloud_load_balancer
Amazon::S3::Bucket	Yes	Yes	cmdb_ci_cloud_object_storage

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for AWS (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
Organizations	Yes	Yes	cmdb_ci_cloud_org or cmdb_ci_cloud_service_account
Organizations::AWS accounts	Yes	Yes	cmdb_ci_cloud_service_account
AWS::EC2::Subnet	Yes	Yes	cmdb_ci_cloud_subnet
AWS SSM agents	Yes	No	cmdb_ci_cloud_system_management_
AWS::EC2::Subnet	Yes	No	cmdb_ci_cloud_subnet
AWS::WAF::WebACL	N/A	Yes	cmdb_ci_cmp_resource
AWS::WAFRegional::WebACL	N/A	Yes	cmdb_ci_cmp_resource
AWS::ACM::Certificate	N/A	Yes	cmdb_ci_cmp_resource
AWS::Redshift::EventSubscription	N/A	Yes	cmdb_ci_cmp_resource
AWS::RDS::DBSnapshot	N/A	Yes	cmdb_ci_cmp_resource
AWS::Redshift::ClusterParameterGroup	N/A	Yes	cmdb_ci_cmp_resource
AWS::AutoScaling::LaunchConfiguration	N/A	Yes	cmdb_ci_cmp_resource
AWS::Redshift::ClusterSnapshot	N/A	Yes	cmdb_ci_cmp_resource
AWS::CloudFormation::Stack	N/A	Yes	cmdb_ci_cmp_resource
AWS::CodeBuild::Project	N/A	Yes	cmdb_ci_cmp_resource
AWS::ElasticBeanstalk::ApplicationVersion	N/A	Yes	cmdb_ci_cmp_resource
AWS::Redshift::ClusterSecurityGroup	N/A	Yes	cmdb_ci_cmp_resource
AWS::CloudTrail::Trail	N/A	Yes	cmdb_ci_cmp_resource
AWS::CloudFront::StreamingDistribution	N/A	Yes	cmdb_ci_cmp_resource
AWS::WAF::RuleGroup	N/A	Yes	cmdb_ci_cmp_resource
AWS::RDS::DBSecurityGroup	N/A	Yes	cmdb_ci_cmp_resource
AWS::RDS::EventSubscription	N/A	Yes	cmdb_ci_cmp_resource
AWS::WAF::Rule	N/A	Yes	cmdb_ci_cmp_resource
AWS::AutoScaling::ScheduledAction	N/A	Yes	cmdb_ci_cmp_resource
AWS::ElasticBeanstalk::Environment	N/A	Yes	cmdb_ci_cmp_resource
AWS::AutoScaling::ScalingPolicy	N/A	Yes	cmdb_ci_cmp_resource
AWS::Shield::Protection	N/A	Yes	cmdb_ci_cmp_resource
AWS::CloudWatch::Alarm	N/A	Yes	cmdb_ci_cmp_resource
AWS::XRay::EncryptionConfig	N/A	Yes	cmdb_ci_cmp_resource
AWS::ShieldRegional::Protection	N/A	Yes	cmdb_ci_cmp_resource
AWS::Redshift::ClusterSubnetGroup	N/A	Yes	cmdb_ci_cmp_resource

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for AWS (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
AWS::IAM::Group	N/A	Yes	cmdb_ci_cmp_resource
AWS::WAFRegional::RateBasedRule	N/A	Yes	cmdb_ci_cmp_resource
AWS::WAFRegional::RuleGroup	N/A	Yes	cmdb_ci_cmp_resource
AWS::WAF::RateBasedRule	N/A	Yes	cmdb_ci_cmp_resource
AWS::RDS::DBSubnetGroup	N/A	Yes	cmdb_ci_cmp_resource
AWS::SSM::AssociationCompliance	N/A	Yes	cmdb_ci_cmp_resource
AWS::EC2::SecurityGroup	Yes	Yes	cmdb_ci_compute_security_group
ec2:DescribeInstanceTypes	Yes	Yes	cmdb_ci_compute_template
AWS::ECR::Repository	Yes	No	cmdb_ci_container_repository
AWS::EC2::CustomerGateway	Yes	No	cmdb_ci_customer_gateway
Docker container	Yes	Yes	cmdb_ci_docker_container
ecr:DescribeImages	Yes	No	cmdb_ci_docker_image
AWS::DynamoDB::GlobalTable	Yes	Yes	cmdb_ci_dynamodb_table
AWS::AutoScaling::AutoScalingGroup	Yes	N/A	cmdb_ci_instance_scale_set
AWS::AutoScaling::AutoScalingGroup	N/A	Yes	cmdb_ci_cmp_resource
AWS::EC2::InternetGateway	Yes	No	cmdb_ci_internet_gateway
AWS::EKS::Cluster	Yes	Yes	cmdb_ci_kubernetes_cluster
WS::ElasticLoadBalancingV2::TargetGroup	Yes	No	cmdb_ci_lb_pool
AWS::ElasticLoadBalancingV2::TargetGroup TargetDescription	Yes	No	cmdb_ci_lb_pool_member
AWS::ElasticLoadBalancingV2::Listener	Yes	No	cmdb_ci_lb_service
AWS::EC2::NatGateway	Yes	No	cmdb_ci_nat_gateway
AWS::EC2::VPC	Yes	Yes	cmdb_ci_network
AWS::EC2::NetworkAcl	Yes	No	cmdb_ci_network_acl
AWS::EC2::NetworkInterface	Yes	Yes	cmdb_ci_nic
AWS::EC2::RouteTable	Yes	No	cmdb_ci_route_table
AWS::EC2::VolumeAttachment	Yes	Yes	cmdb_ci_storage_mapping
AWS::EC2::Volume	Yes	Yes	cmdb_ci_storage_volume
AWS::EC2::VPNGateway	Yes	No	cmdb_ci_virtual_pvt_gateway
AWS::EC2::VPNConnection	Yes	No	cmdb_ci_vpn_connection
AWS::WAFv2::WebACL	Yes	N/A	cmdb_ci_web_acl
AWS::WAFv2::WebACL	N/A	Yes	cmdb_ci_cmp_resource
Tags	Yes	Yes	cmdb_key_value

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for AWS (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
Kubernetes DaemonSet	No	Yes	cmdb_ci_kubernetes_daemonset
Kubernetes deployment	No	Yes	cmdb_ci_kubernetes_deployment
Kubernetes namespace	No	Yes	cmdb_ci_kubernetes_namespace
Amazon EKS node	No	Yes	cmdb_ci_kubernetes_node
Amazon EKS pod	No	Yes	cmdb_ci_kubernetes_pod
Amazon EKS service	No	Yes	cmdb_ci_kubernetes_service
Kubernetes Volume	No	Yes	cmdb_ci_kubernetes_volume
AWS::SSM::ManagedInstanceInventory	No	Yes	cmdb_ci_spkg
AWS::EC2::VPCPeeringConnection	Yes	No	cmdb_aws_elastic_compute_cloud_vpc
AWS::Elastic Compute Cloud::VPC Endpoint Service	Yes	No	cmdb_aws_elastic_cloud_compute_vpc
AWS::Elastic Compute Cloud::EBS Snapshot	Yes	No	cmdb_aws_elastic_cloud_compute_ebs
AWS::Elastic Compute Cloud::Reserved Instance	Yes	No	cmdb_aws_elastic_cloud_compute_res
AWS::RDS::DBSnapshot	Yes	No	cmdb_aws_relational_database_servic
AWS::SageMaker::Training Job	Yes	No	cmdb_aws_sagemaker_training_job
AWS::Elastic File System::File System	Yes	No	cmdb_aws_elastic_file_system_file_sys
AWS::Redshift Serverless::Namespace	Yes	No	cmdb_aws_redshift_serverless_names
AWS::Redshift Serverless::Workgroup	Yes	No	cmdb_aws_redshift_serverless_workgr
AWS::Redshift Serverless::Snapshot	Yes	No	cmdb_aws_redshift_serverless_snapsh
AWS::ElastiCache::Snapshot	Yes	No	cmdb_aws_elasticache_snapshot
AWS::CloudTrail::Trail	Yes	No	cmdb_aws_cloudtrail_trail
AWS::Batch::ComputeEnvironment	Yes	No	cmdb_aws_batch_compute_environme
AWS::StorageGateway::FileShare	Yes	No	cmdb_aws_storage_gateway_file_share
AWS::DataSync::Task	Yes	No	cmdb_aws_datasync_task

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for AWS (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
AWS::API Gateway::Domain Name	Yes	No	cmdb_aws_api_gateway_domain_name
AWS DynamoDB Cluster	Yes	No	cmdb_ci_cluster_node
AWS::GlobalAccelerator::Accelerator	Yes	No	cmdb_aws_global_accelerator_accelera
AWS::Timestream::InfluxDBInstance	Yes	No	cmdb_aws_timestream_for_influxdb_da
AWS::CloudHSM::HSM	Yes	No	cmdb_aws_cloudhsm_hsm
AWS::MWAA::Environments	Yes	No	cmdb_aws_managed_workflows_for_ap
AWS::NetworkFirewall::Firewall	Yes	No	cmdb_aws_network_firewall_firewall

Microsoft Azure discovery solutions comparison

ITOM Visibility applications discover a variety of Microsoft Azure resources and populate the relevant configuration item (CI) classes in the Configuration Management Database (CMDB) with their attributes.

ITOM Visibility solutions

ITOM Visibility can offer you comprehensive visibility into your organization's cloud infrastructure. Each application offers unique strengths and capabilities, and by comparing them you can gain holistic insights on their capabilities and benefits.

ITOM Visibility solutions for discovering Azure

ITOM Visibility solution	Overview	Configuration
Discovery and Service Mapping Patterns and Visibility Content	Discovery patterns used by ITOM Visibility	• You must update your store apps to the latest release version. • You must provide ITOM Visibility with access permissions. For more information, see Set up Azure service accounts
Service Graph Connector for Microsoft Azure	Service Graph Connector for Microsoft Azure	Configure Service Graph Connector for Microsoft Azure using the guided setup

Note:

The CI class, in most cases, is applicable to both solutions. If it's not applicable to one of the solutions, the solution is marked "N/A".

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for Microsoft Azure

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
microsoft.compute/availabilitysets	Yes	N/A	cmdb_ci_availability_set
microsoft.compute/availabilitysets	Yes	Yes	cmdb_ci_availability_zone
Microsoft.Management/managementGroups	Yes	No	cmdb_ci_azure_management_group
microsoft.KeyVault/vaults	Yes	No	cmdb_ci_certificate
• Microsoft.Sql/servers/databases • Microsoft.DBforMySQL/servers • Microsoft.DocumentDB/databaseAccounts • Microsoft.Cache/Redis	Yes	No	cmdb_ci_cloud_database
• Microsoft.Sql/servers • Microsoft.DBforPostgreSQL/servers	Yes	Yes	cmdb_ci_cloud_database
microsoft.network/expressroutecircuits	Yes	No	cmdb_ci_cloud_direct_connect
Microsoft.Storage/storageAccounts/fileServices	Yes	No	cmdb_ci_cloud_file_service
Microsoft.Storage/storageAccounts/fileServices/shares	Yes	No	cmdb_ci_cloud_file_share
microsoft.web/sites (kind matches regex functionapp.*)	Yes	Yes	cmdb_ci_cloud_function
VM Size	Yes	No	cmdb_ci_cloud_hardware_type
microsoft.compute/hostgroups/hosts	Yes	No	cmdb_ci_cloud_host
microsoft.network/applicationgateways	Yes	No	cmdb_ci_cloud_load_balancer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for Microsoft Azure (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
microsoft.network/loadbalancers	Yes	Yes	cmdb_ci_cloud_load_balancer
Microsoft.Storage/storageAccounts/blobService	Yes	No	cmdb_ci_cloud_object_service
Tenant	Yes	No	cmdb_ci_cloud_org
microsoft.network/publicipaddresses	Yes	Yes	cmdb_ci_cloud_public_ipaddress
Microsoft.Storage/storageAccounts/queueServices	Yes	No	cmdb_ci_cloud_queue_service
Microsoft.Resources/subscriptions	Yes	Yes	cmdb_ci_cloud_service_account
Microsoft.Management/managementGroup	Yes	No	cmdb_ci_cloud_service_account
microsoft.storage/storageaccounts	Yes	Yes	cmdb_ci_cloud_storage_account
Microsoft.Network/virtualNetworks/subnets	Yes	No	cmdb_ci_cloud_subnet
Microsoft.Storage/storageAccounts/tableServices	Yes	No	cmdb_ci_cloud_table_service
microsoft.web/sites (kind doesn't contain functionapp)	Yes	Yes	cmdb_ci_cloud_webserver
microsoft.network/networksecuritygroups	Yes	Yes	cmdb_ci_compute_security_group
microsoft.network/privatednszones	Yes	No	cmdb_ci_dns_zone
microsoft.compute/virtualmachinescalesets	Yes	Yes	cmdb_ci_instance_scale_set
microsoft.containerservice/managedclusters	Yes	Yes	cmdb_ci_kubernetes_cluster
Microsoft.Network/loadBalancers/backendAddressPools	Yes	No	cmdb_ci_lb_pool
Microsoft.Network/loadBalancers/inboundNatRules	Yes	No	cmdb_ci_lb_service

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for Microsoft Azure (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
microsoft.network/natgateways	Yes	No	cmdb_ci_nat_gateway
microsoft.network/virtualnetworks	Yes	Yes	cmdb_ci_network
microsoft.network/networkinterfaces	Yes	Yes	cmdb_ci_nic
microsoft.compute/images	Yes	No	cmdb_ci_os_template
Microsoft.Resources/resourceGroups	Yes	Yes	cmdb_ci_resource_group
Microsoft.Network/routeTables/routes	Yes	No	cmdb_ci_route_table
microsoft.network/routetables	Yes	No	cmdb_ci_route_table
microsoft.compute/disks	Yes	Yes	cmdb_ci_storage_volume
microsoft.network/localnetworkgateways	Yes	No	cmdb_ci_virtual_pvt_gateway
microsoft.network/virtualnetworkgateways	Yes	No	cmdb_ci_virtual_pvt_gateway
microsoft.compute/virtualmachines	Yes	Yes	cmdb_ci_vm_instance
microsoft.sqlvirtualmachines/sqlvirtualmachines	Yes	No	cmdb_ci_vm_instance
Microsoft.Network/virtualNetworks/virtualNetworkPeerings	Yes	No	cmdb_ci_vnet_peering
microsoft.network/connections	Yes	No	cmdb_ci_vpc_gateway_connection
Microsoft.Resources/tags	Yes	Yes	cmdb_key_value
Microsoft.Compute/virtualMachines	No	Yes	cmdb_ci_server
Microsoft.Compute/virtualMachines	No	Yes	cmdb_ci_win_server
Microsoft.Network/networkInterfaces	No	Yes	cmdb_ci_hardware
Microsoft.Compute/virtualMachines	No	Yes	cmdb_ci_linux_server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for Microsoft Azure (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
microsoft.containerservices/managedclusters	Yes	Yes	cmdb_ci_kubernetes_cluster
Microsoft.Compute/virtualmachinescalesets	No	Yes	cmdb_ci_vm_instance
Microsoft.ServiceBus/namespaces/queues	Yes	No	cmdb_azure_service_bus_queue
Microsoft.ServiceBus/namespaces/topics	Yes	No	cmdb_azure_service_bus_topic
microsoft.appconfiguration/configurationstores	Yes	No	cmdb_azure_app_configuration_configu
microsoft.automation/automationaccounts	Yes	No	cmdb_azure_automation_automation_a
microsoft.compute/diskencryptionsets	Yes	No	cmdb_azure_disk_encryption_set_disk
microsoft.compute/galleries/images	Yes	No	cmdb_azure_compute_gallery_image_c
microsoft.compute/snapshots	Yes	No	cmdb_azure_compute_snapshot
microsoft.containerregistry/registries	Yes	No	cmdb_azure_container_registry_registry
microsoft.databricks/workspaces	Yes	No	cmdb_azure_databricks_workspace
microsoft.datafactory/factories	Yes	No	cmdb_azure_data_factory_data_factory
microsoft.dataprotection/backupvaults	Yes	No	cmdb_azure_data_protection_backup_
microsoft.eventgrid/systemtopics	Yes	No	cmdb_azure_event_grid_system_topic
microsoft.eventhub/namespaces	Yes	No	cmdb_azure_event_hub_namespace
microsoft.insights/components	Yes	No	cmdb_azure_application_insight_comp
microsoft.network/applicationgatewaywebapplicationfirewallpolicies	Yes	No	cmdb_azure_web_application_firewall_
microsoft.operationalinsights/workspaces	Yes	No	cmdb_azure_log_analytics_workspace
microsoft.network/azurefirewalls	Yes	No	cmdb_azure_network_security_azure_fi

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for Microsoft Azure (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
microsoft.devcenter/devcenters	Yes	No	cmdb_azure_dev_centers_dev_center
microsoft.kusto/clusters	Yes	No	cmdb_azure_data_explorer_clusters_da
microsoft.recoveryservices/vaults	Yes	No	cmdb_azure_recovery_services_vault
microsoft.recoveryservices/vaults/backupfabrics/protectioncontainers/protecteditems	Yes	No	cmdb_azure_recovery_services_vault_b
microsoft.dbforpostgresql/servergroupsv2	Yes	No	cmdb_azure_cosmos_db_for_postgresc
microsoft.insights/datacollectionrules	Yes	No	cmdb_azure_application_insight_data_
microsoft.network/serviceendpointpolicies	Yes	No	cmdb_azure_service_endpoint_policy_

GCP discovery solutions comparison

ITOM Visibility applications discover a variety of Google Cloud Platform (GCP) resources and populate the relevant configuration item (CI) classes in the Configuration Management Database (CMDB) with their attributes.

ITOM Visibility solutions

ITOM Visibility can offer you comprehensive visibility into your organization's cloud infrastructure. Each application offers unique strengths and capabilities, and by comparing them you can gain holistic insights on their capabilities and benefits.

ITOM Visibility solutions for discovering GCP

ITOM Visibility solution	Overview	Configuration
Discovery and Service Mapping Patterns and Visibility Content	Discovery patterns used by ITOM Visibility	• You must update your store apps to the latest release version. • You must provide ITOM Visibility with access permissions. For more information, see Set up Google Cloud Platform service accounts

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM Visibility solutions for discovering GCP (continued)

ITOM Visibility solution	Overview	Configuration
Service Graph Connector for GCP	Service Graph Connector for GCP	Configure Service Graph Connector for GCP using the guided setup

Note:

The CI class, in most cases, is applicable to both solutions. If it's not applicable to one of the solutions, the solution is marked "N/A".

If an entry is marked with one asterisk, the resource type is supported by Patterns only by utilizing the **Google Asset inventory** pattern's inclusion list. For more information, see Google Cloud Platform (GCP) resource inventory discovery with Patterns.

If an entry is marked with two asterisks, the resource type is supported by Service Graph Connecters only by utilizing the [sn_cmdb_int_util_service_graph_resource_inclusion_whitelist] table.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
admissionregistration.k8s.io/mutatingwebhookconfiguration	Yes*	Yes**	cmdb_ci_cmp_resource
admissionregistration.k8s.io/validatingwebhookconfiguration	Yes*	Yes**	cmdb_ci_cmp_resource
alloydb.googleapis.com/Cluster	Yes	No	cmdb_ci_cloud_db_cluster
apigateway.googleapis.com/api	Yes*	Yes**	cmdb_ci_cmp_resource
apigateway.googleapis.com/apiconfig	Yes*	Yes**	cmdb_ci_cmp_resource
apigee.googleapis.com/instance	Yes*	Yes**	cmdb_ci_cmp_resource
apigee.googleapis.com/organization	Yes*	Yes**	cmdb_ci_cmp_resource
apikeys.googleapis.com/key	Yes*	Yes**	cmdb_ci_cmp_resource
appengine.googleapis.com/application	Yes*	Yes**	cmdb_ci_cmp_resource
apps.k8s.io/daemonset	Yes*	Yes**	cmdb_ci_cmp_resource
apps.k8s.io/Deployment	Yes*	Yes	cmdb_ci_kubernetes_deployment

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
apps.k8s.io/ReplicaSet	Yes*	Yes	cmdb_ci_kubernetes_replicaset
apps.k8s.io/statefulset	Yes*	Yes**	cmdb_ci_cmp_resource
artifactregistry.googleapis.com/dockerimage	Yes*	Yes**	cmdb_ci_cmp_resource
artifactregistry.googleapis.com/repository	Yes*	Yes**	cmdb_ci_cmp_resource
bigquery.googleapis.com/bigquery-public-data/datasets	Yes	No	cmdb_ci_cloud_public_database
bigquery.googleapis.com/bigquery-public-data/datasets/tables	Yes	No	cmdb_ci_database
bigquery.googleapis.com/Dataset	Yes	N/A	cmdb_ci_cloud_database
bigquery.googleapis.com/Dataset	N/A	Yes**	cmdb_ci_cmp_resource
bigquery.googleapis.com/Table	Yes	No	cmdb_ci_database
bigtableadmin.googleapis.com/Instance	Yes	Yes	cmdb_ci_cloud_database
bigtableadmin.googleapis.com/Table	Yes	No	cmdb_ci_database
cloudasset.googleapis.com/feed	Yes*	Yes**	cmdb_ci_cmp_resource
cloudbilling.googleapis.com/billingaccount	Yes*	Yes**	cmdb_ci_cmp_resource
cloudbilling.googleapis.com/projectbillinginfo	Yes*	Yes**	cmdb_ci_cmp_resource
cloudfunctions.googleapis.com/CloudFunction	Yes*	Yes	cmdb_ci_cloud_function
cloudresourcemanager.googleapis.com/Folder	Yes	Yes	cmdb_ci_gcp_folder
cloudresourcemanager.googleapis.com/Organization	Yes	Yes	cmdb_ci_cloud_org
cloudresourcemanager.googleapis.com/Project	Yes	Yes	cmdb_ci_gcp_project
cloudresourcemanager.googleapis.com/tagkey	Yes*	Yes	cmdb_ci_cmp_resource

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
cloudresourcemanager.googleapis.com/tagvalue	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/Address	Yes	No	cmdb_ci_cloud_public_ipaddress
compute.googleapis.com/autoscaler	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/BackendService	N/A	Yes	cmdb_ci_cloud_load_balancer
compute.googleapis.com/BackendService	Yes	N/A	cmdb_ci_lb_service
compute.googleapis.com/Disk	Yes	Yes	cmdb_ci_storage_volume
compute.googleapis.com/DiskTypes	Yes	No	cmdb_ci_disk_type
compute.googleapis.com/Firewall	N/A	Yes	cmdb_ci_compute_security_group
compute.googleapis.com/Firewall	Yes	Yes	cmdb_ci_network_acl_rule
compute.googleapis.com/ForwardingRule	Yes	Yes	cmdb_ci_lb_service
compute.googleapis.com/HttpHealthCheck	Yes	Yes	cmdb_ci_lb_health_service
compute.googleapis.com/Image	Yes	No	cmdb_ci_os_template
compute.googleapis.com/Instance	Yes	N,A	cmdb_ci_lb_pool_member
compute.googleapis.com/Instance	N/A	Yes	cmdb_ci_server
compute.googleapis.com/Instance	Yes	Yes	cmdb_ci_vm_instance
compute.googleapis.com/Instance/networkInterfaces	Yes	No	cmdb_ci_endpoint_vnic
compute.googleapis.com/Instance/networkInterfaces	Yes	Yes	cmdb_ci_nic
compute.googleapis.com/InstanceGroup	Yes	N/A	cmdb_ci_lb_pool
compute.googleapis.com/InstanceGroup	N/A	Yes**	cmdb_ci_cmp_resource

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
compute.googleapis.com/Instancegroupmanager	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/Instancesettings	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/InstanceTemplate	Yes	No	cmdb_ci_instance_template
compute.googleapis.com/Instancetemplate	No	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/MachineImage	No	Yes	cmdb_ci_os_template
compute.googleapis.com/MachineTypes	Yes	No	cmdb_ci_cloud_hardware_type
compute.googleapis.com/MachineTypes	Yes	No	cmdb_ci_compute_template
compute.googleapis.com/Network	Yes	Yes	cmdb_ci_network
compute.googleapis.com/Networkendpointgroup	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/Regions	Yes	No	cmdb_ci_google_datacenter
compute.googleapis.com/resourcepolicy	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/router	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/Snapshot	Yes	Yes	cmdb_ci_storage_vol_snapshot
compute.googleapis.com/ssh	Yes	No	cmdb_ci_cloud_key_pair
compute.googleapis.com/sslcertificate	Yes*	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/Subnetwork	Yes	Yes	cmdb_ci_cloud_subnet
compute.googleapis.com/targetHttpProxy	Yes	N/A	cmdb_ci_lb_service
compute.googleapis.com/targetHttpProxy	Yes	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/TargetPool	Yes	N/A	cmdb_ci_cloud_load_balancer
compute.googleapis.com/TargetPool	N/A	Yes	cmdb_ci_lb_pool

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
compute.googleapis.com/targetSslProxy	Yes	No	cmdb_ci_lb_service
compute.googleapis.com/targetTcpProxy	Yes	No	cmdb_ci_lb_service
compute.googleapis.com/UrlMap	Yes	N/A	cmdb_ci_cloud_load_balancer
compute.googleapis.com/UrlMap	N/A	Yes**	cmdb_ci_cmp_resource
compute.googleapis.com/zones	Yes	No	cmdb_ci_availability_zone
container.googleapis.com/Cluster	Yes	Yes	cmdb_ci_kubernetes_cluster
container.googleapis.com/NodePool	No	Yes	cmdb_ci_kubernetes_node_pool
containerregistry.googleapis.com/image	Yes*	Yes**	cmdb_ci_cmp_resource
dataplex.googleapis.com/entrygroup	Yes*	Yes**	cmdb_ci_cmp_resource
firebase.googleapis.com/firebaseproject	Yes*	Yes**	cmdb_ci_cmp_resource
firebasedatabase.googleapis.com/instances	Yes	No	cmdb_ci_cloud_database
firebaserules.googleapis.com/release	Yes*	Yes**	cmdb_ci_cmp_resource
firebaserules.googleapis.com/ruleset	Yes*	Yes**	cmdb_ci_cmp_resource
firestore.googleapis.com/Database	Yes	Yes	cmdb_ci_cloud_database
iam.googleapis.com/role	Yes*	Yes**	cmdb_ci_cmp_resource
iam.googleapis.com/serviceaccount	Yes*	Yes**	cmdb_ci_cmp_resource
iam.googleapis.com/serviceaccountkey	Yes*	Yes**	cmdb_ci_cmp_resource
integrations.googleapis.com/integration	Yes*	Yes**	cmdb_ci_cmp_resource
integrations.googleapis.com/integrationversion	Yes*	Yes**	cmdb_ci_cmp_resource
K8s.io/endpoints	Yes*	Yes**	cmdb_ci_cmp_resource

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
K8s.io/Namespace	Yes*	Yes	cmdb_ci_kubernetes_namespace
k8s.io/Node	Yes*	Yes	cmdb_ci_kubernetes_node
K8s.io/ persistentvolume	Yes*	Yes**	cmdb_ci_cmp_resource
k8s.io/ persistentvolumeclaim	Yes*	Yes**	cmdb_ci_cmp_resource
K8s.io/Pod	Yes*	Yes	cmdb_ci_docker_container
K8s.io/Pod	N/A	Yes	cmdb_ci_docker_image
K8s.io/Pod	N/A	Yes	cmdb_ci_kubernetes_pod
K8s.io/Pod	N/A	Yes	cmdb_ci_kubernetes_volume
K8s.io/resourcequota	Yes*	Yes**	cmdb_ci_cmp_resource
k8s.io/Service	Yes*	Yes	cmdb_ci_kubernetes_service
K8s.io/serviceaccount	Yes*	Yes**	cmdb_ci_cmp_resource
labels	Yes	No	cmdb_key_value
logging.googleapis.com/ logbucket	Yes*	Yes**	cmdb_ci_cmp_resource
logging.googleapis.com/ logmetric	Yes*	Yes**	cmdb_ci_cmp_resource
logging.googleapis.com/ logsink	Yes*	Yes**	cmdb_ci_cmp_resource
logging.googleapis.com/ recentquery	Yes*	Yes**	cmdb_ci_cmp_resource
memcache.googleapis.com/ Instance	Yes*	Yes	cmdb_ci_cloud_database
monitoring.googleapis.com/ alertpolicy	Yes*	Yes**	cmdb_ci_cmp_resource
monitoring.googleapis.com/ dashboard	Yes*	Yes**	cmdb_ci_cmp_resource
monitoring.googleapis.com/ notificationchannel	Yes*	Yes**	cmdb_ci_cmp_resource
networkmanagement.googleapis.com/ connectivitytest	Yes*	Yes**	cmdb_ci_cmp_resource
orgpolicy.googleapis.com/ policy	Yes*	Yes**	cmdb_ci_cmp_resource
policy.k8s.io/ poddisruptionbudget	Yes*	Yes**	cmdb_ci_cmp_resource

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
pubsub.googleapis.com/subscription	Yes*	Yes**	cmdb_ci_cmp_resource
pubsub.googleapis.com/topic	Yes*	Yes**	cmdb_ci_cmp_resource
rbac.authorization.k8s.io/ClusterRole	No	Yes	cmdb_ci_kubernetes_cluster_role
rbac.authorization.k8s.io/ClusterRoleBinding	No	Yes	cmdb_ci_kubernetes_cluster_role_binc
rbac.authorization.k8s.io/role	Yes*	Yes**	cmdb_ci_cmp_resource
rbac.authorization.k8s.io/rolebinding	Yes*	Yes**	cmdb_ci_cmp_resource
redis.googleapis.com/clusters	Yes	No	cmdb_ci_cloud_db_cluster
redis.googleapis.com/Instance	No	Yes	cmdb_ci_cloud_database
run.googleapis.com/revision	Yes*	Yes**	cmdb_ci_cmp_resource
run.googleapis.com/service	Yes*	Yes**	cmdb_ci_cmp_resource
servicedirectory.googleapis.com/endpoint	Yes*	Yes**	cmdb_ci_cmp_resource
servicedirectory.googleapis.com/namespace	Yes*	Yes**	cmdb_ci_cmp_resource
servicedirectory.googleapis.com/service	Yes*	Yes**	cmdb_ci_cmp_resource
servicemanagement.googleapis.com/managedservice	Yes*	Yes**	cmdb_ci_cmp_resource
servicenetworking.googleapis.com/connection	Yes*	Yes**	cmdb_ci_cmp_resource
serviceusage.googleapis.com/service	Yes*	Yes**	cmdb_ci_cmp_resource
spanner.googleapis.com/Database	Yes	No	cmdb_ci_database
spanner.googleapis.com/Instance	Yes	Yes	cmdb_ci_cloud_database
sqladmin.googleapis.com/backup	Yes*	Yes**	cmdb_ci_cmp_resource
sqladmin.googleapis.com/backuprun	Yes*	Yes**	cmdb_ci_cmp_resource

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns vs Service Graph Connector for GCP (continued)

Resource type	Supported by Cloud Discovery and Service Mapping Patterns	Supported by Cloud Service Graph Connectors	CI Class
sqladmin.googleapis.com/Instance	Yes*	Yes	cmdb_ci_cloud_database
sqladmin.googleapis.com/Instances/databases	Yes*	No	cmdb_ci_database
storage.googleapis.com/Bucket	N/A	Yes	cmdb_ci_cloud_object_storage
storage.googleapis.com/Bucket	Yes	N/A	cmdb_ci_cmp_resource
storage.k8s.io/storageclass	Yes*	Yes**	cmdb_ci_cmp_resource

* These resource types are supported in Patterns only by utilizing the **Google Asset inventory** pattern's inclusion list. For more information, see Google Cloud Platform (GCP) resource inventory discovery with Patterns.

** These resource types are supported in Service Graph Connecters only by utilizing the [sn_cmdb_int_util_service_graph_resource_inclusion_whitelist] table.

Discovery for containerized resources

Container discovery enables IT departments of companies to collect detailed information about their container management tools in the infrastructure. Discovery finds resources in major container management services providers.

Related topics

Amazon ECS resource discovery with Patterns

Pivotal Cloud Foundry discovery

Docker virtualization

Kubernetes discovery using patterns

The ServiceNow ITOM Visibility finds Kubernetes and OpenShift components using patterns and creates application services containing them. Discovery also finds Kubernetes events and frequently updates the CMDB to reflect the dynamic Kubernetes environment.

Discovery uses the Kubernetes pattern and its extension sections to discover Kubernetes components:

- The *`Collect OpenShift info`* extension section of the Kubernetes pattern discovers the OpenShift components of the Kubernetes deployment. The OpenShift Build Config extension section is available from Store version 1.0.53.
- The *`Service Mesh`* extension discovers service mesh details. This information enables the pattern to create service-to-service relations, shown as Connects to::Connected. Service mesh discovery requires deploying Istio on your K8s (Kubernetes) cluster. The Service Mesh extension section is available from Kubernetes extension classes. It's supported on the ServiceNow AI Platform using the Madrid release or later.
- The *`Collect Container Repository`* and extension section finds container registries and images in these registries.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Starting from Discovery and Service Mapping Patterns version 1.35.0, the Kubernetes Cluster - Per-Namespace LP pattern is available for large cluster discovery. For more information, see the **Large-payload Kubernetes discovery** section.

In addition, Discovery uses the Kubernetes Event pattern to discover events for Kubernetes components.

Starting from the 1.0.68 release on ServiceNow Store, Service Mapping can use CI relationships to add the Kubernetes components to application services during tag-based discovery.

Discovery uses the following patterns to discover the entire Kubernetes infrastructure deployed on GCP, AWS, and Azure:

- Google Cloud Platform (GCP) – Get Kubernetes Clusters.
- Amazon AWS Cloud - Get Kubernetes Clusters.
- Microsoft Azure - Get Kubernetes Clusters.

These patterns query the Cloud, collect data on all Kubernetes clusters, and create a serverless schedule for each cluster. When the cluster is deleted, the schedule is marked as inactive. This feature eliminates the overhead of creating and managing multiple credentials and serverless discovery schedules per cluster. The Cloud infrastructure patterns are triggered through standard Cloud discovery.

Starting with Discovery and Service Mapping Patterns version 1.31.0, you can choose to discover Docker image CIs only, without discovering Docker container CIs. Check your entitlements to determine whether you have access to 2026 Packaging SKU. For more information, see Disable Docker container CI discovery.

Supported versions

For the list of Kubernetes distribution versions validated against the Kubernetes and Kubernetes Event patterns, see Kubernetes Versions Tested with ServiceNow Discovery [KB3145280].

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Large-payload Kubernetes discovery

By default, Discovery uses the Kubernetes pattern to discover an entire cluster in a single execution. On large clusters, this can result in incomplete discovery data in the CMDB. The Kubernetes Cluster - Per-Namespace LP pattern is available as an alternative for large cluster environments, starting from Discovery and Service Mapping Patterns version 1.35.0.

The Kubernetes Cluster - Per-Namespace LP pattern identifies the cluster and then discovers each namespace separately, improving reliability on large clusters. The data collected and the CIs populated in the CMDB are identical between the two patterns. For configuration on on-premises clusters, see the **Prerequisites for on-premises Kubernetes discovery** section. For cloud clusters, see the prerequisites section for your cloud provider.

General prerequisites for Kubernetes discovery

Note: For additional prerequisites for Kubernetes Cloud infrastructure discovery, see below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Running automatic serverless Kubernetes schedules fetches the Bearer token. Adding credentials is unnecessary.

Perform the following steps so that Discovery can use the pattern to successfully find Kubernetes.

1. Deploy the latest Discovery and Service Mapping Patterns application from ServiceNow Store.
2. On the Kubernetes platform, find the parameters to set up Kubernetes discovery:
   - Find the URL of the kubeapi server:
     a. On the Kubernetes platform, run the following command:

        `kubectl cluster-info`

     b. In the output, find the line that states the URL of the kubeapi server. For example, Kubernetes control plane is running at

        `https://10.154.144.146:443`

   - Find the namespaces of the kubeapi server:
     a. On the Kubernetes platform, run this command:

        `kubectl get namespaces`

     b. In the output, find the line that states the namespaces. For example, `kube-system`.
   - Find the Kubernetes username and password:
     a. On the Kubernetes platform, run this command:

        `kubectl config view`

     b. In the output, find the username and

password.

Note: If in a certain environment, `kubectl config view` command is not showing the expected details, use the `supported` command from the Kubernetes admin to fetch the user name and password details.

   - Find the valid Bearer token with the proper permissions:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you know the default token name, use the command in the following format: `kubectl describe secret <default-token-token name>`.

  For example: `kubectl describe secret default-token-g6pwc`.

- If you don't know the default token name, use the command: `kubectl describe secret`.

3. Verify that the API Server is reachable from the MID Server for successful Kubernetes discovery.
4. Verify that the user configured on the Kubernetes platform has **GET** permissions to run the following /api/v1 elements:
   - https://<url>/api/v1/namespaces/
   - https://<url>/api/v1/namespaces/<namespace>
   - https://<url>/api/v1/namespaces/kube-system/endpoints/kube-controller-manager
   - https://<url>/api/v1/services
   - https://<url>/api/v1/pods
   - https://<url>/api/v1/nodes
   - https://<url>/api/v1/replicationcontrollers
   - https://<url>/apis/networking.k8s.io/v1/ingresses
   - https://<url>/apis/apps/v1/deployments
   - https://<url>/apis/apps/v1/statefulsets
   - https://<url>/apis/apps/v1/daemonsets
   - https://<url>/apis/apps/v1/replicasets
   - https://<url>/apis/batch/v1/cronjobs
   - https://<url>/apis/batch/v1/jobs

   Note: When using the **Kubernetes Cluster - Per-Namespace LP** pattern, verify your RBAC setup before switching. The LP pattern requires the same permissions listed above, but lists resources per namespace instead of cluster-wide. A ClusterRole with a ClusterRoleBinding covers all namespaces automatically and requires no changes. If you use namespace-scoped Role and RoleBinding objects, verify that a binding exists in each namespace that discovery will enumerate, because missing bindings could cause incomplete discovery for that namespace.

5. To discover the OpenShift components of the Kubernetes deployment, verify that the user configured on the Kubernetes platform has **GET** permissions to run the following /api/v1 elements:
   - /apis/apps.openshift.io/v1/deploymentconfigs
   - /apis/build.openshift.io/v1/buildconfigs
   - /apis/route.openshift.io/v1/routes
   - /apis/user.openshift.io/v1/groups
   - /apis/user.openshift.io/v1/users
   - /apis/project.openshift.io/v1/projects

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- /apis/image.openshift.io/v1/images
- /apis/image.openshift.io/v1/imagestreams

To discover service mesh information:

- Deploy Istio on your K8s cluster.
- Provide the Prometheus URL.
- Configure Prometheus to scrape metrics from Istio.

6. Activate **Get Kubernetes Config Files** extension to:
   - Discover configuration files.
   - Create tracked configuration files.
   - Map the configuration files workloads and services with a relationship.

   Note: Tracked files content is in the JSON format from version 1.0.92. Tracked files content is in YAML format in version 1.0.91 and earlier.

7. Create the Kubernetes credentials on the ServiceNow platform:
   a. On the ServiceNow AI Platform, navigate to **All** > **Discovery** > **Credentials**.
   b. Select **New**.
   c. Select **Kubernetes Credentials**.
   d. On the form, fill in the fields.

Field	Description
Name	Unique and descriptive name for this credential.
User name	User name associated with this credential. Leading or trailing spaces should be avoided; if any are detected, a warning will appear. Only one authentication method should be used: either a user name and password or a Bearer token. Don't use both.
Password	Password associated with this credential. Only one authentication method should be used: either a user name and password or a Bearer token. Don't use both.
Bearer Token Authentication	This option enables advanced authentication using a Bearer token. When the check box is selected, the Bearer Token field is displayed.
Bearer Token	Discovery uses the Bearer token for advanced authentication when accessing Kubernetes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	The Bearer token should be in BASE64 encoded format, using the character sequence as the token. For example: 31ada4fd-adec-460c-809a-9e56ceb75269. Only one authentication method should be used: either a user name and password or a Bearer token. Don't use both.
Credential alias	An alias is configured to use the Kubernetes credential for devices and applications other than Kubernetes. This alias is also used when defining a serverless discovery schedule for discovering the Kubernetes deployment. **i.** Select the padlock icon, and then select the search icon. **ii.** On the Connection & Credential Aliases form, select **New**. **iii.** Specify a name for the credential alias record. **iv.** Define attributes for the alias. Set the *Type* to `Credential`. **v.** Select and hold (or right-click) the form header and select **Save**, then select **Update**. **vi.** On the Connection & Credential Aliases form, select the newly added alias. The alias appears in the Credential alias field.

e. On the Kubernetes credentials form, select **Update**.

8. Create a serverless discovery schedule for the Kubernetes pattern.

a. Navigate to **All** > **Discovery** > **Discovery Schedules** and select **New**.

b. On the form, fill in the fields.

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type, which should be **Serverless**.
MID Server selection method	Select the method that Discovery uses to select a MID Server: ▪ `Specific MID Cluster`: Use a preconfigured cluster of MID Servers. The MID Server can't be part of multiple clusters

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	▪ `Specific MID Server`: Use only one MID Server. If that MID Server is part of a cluster, only that MID Server is used. The cluster isn't used.
MID server	Name of the MID Server to use for this schedule. This field is available if **MID Server selection method** is set to `Specific MID Server`.
MID Server cluster	Name of the MID Server cluster to use for this schedule. This field is available if **MID Server selection method** is set to `Specific MID Cluster`.

c. Select **Submit**.

d. In the **Execution Patterns** related list, select **New**.

e. On the form, fill in the fields.

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern to use for this schedule: ▪ **Kubernetes** ▪ **Kubernetes Cluster - Per-Namespace LP** (starting from Discovery and Service Mapping Patterns version 1.35.0)
Run Child Patterns	Enable this option if you selected **Kubernetes Cluster - Per-Namespace LP** as the pattern.

f. Select **Submit**.

g. In the **Pattern Launcher Parameters** related list, configure the parameters.

Configuring execution pattern attributes

Field	Description
url	The identifier for the hostname, IP, or FQDN and the port of the Kubernetes apiserver. Use the following format: example_hostname:example_port or example_ip:example_port. Provide the correct protocol (HTTP or HTTPS) in the URL.
namespace	The namespaces that the system passes in the Kubernetes Discovery Configuration. Enter one of the following values:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	▪ Individual namespace: enter the namespace and then "kube-system". For example: `dev,kube-system` ▪ The default value: enter `default, kube-system` ▪ Multiple namespaces: enter the namespaces, use a comma (,) to separate the values, and then enter "kube-system". For example: `automation,application,test,kube-system` ▪ All namespaces: Use an asterisk (*) to enter all namespaces
credentials alias	The alias associated with the previously created Kubernetes credentials.
cluster name	The name of the Kubernetes cluster, in the following format: <serviceaccountid><space><clustername>.
provider	The cloud provider or deployment type: ▪ For cloud hosted clusters: GCP, AWS, or Azure ▪ For on-premises clusters: `OnPrem`
cluster_resource_id	Cluster resource ID example: ▪ Azure Kubernetes clusters - Resource ID. ▪ AWS - cluster ARN. ▪ GCP - cluster global name.
cluster_uid_cache	Internal parameter used by the **Kubernetes Cluster - Per-Namespace LP** pattern. Leave this value empty.

9. Create a serverless discovery schedule for the Kubernetes Event pattern. Configure the schedule to run every 5 or 10 minutes.

   Note: When the pattern is run for the first time, it stores an event_timestamp. Later on it collects only the delta events based on the timestamp. The more often the pattern is run, the fewer updates to the CMDB IRE are needed.

   Create a serverless execution pattern for the discovery schedule and assign it to the Kubernetes Events pattern. Configure the parameters required by the Kubernetes pattern as described in Configuring execution pattern attributes.

10. To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Additional prerequisites for Kubernetes Cloud infrastructure discovery

For the Google Cloud Platform (GCP) – Get Kubernetes Clusters pattern, perform the following:

1. In the ServiceNow instance, set up a Google Cloud Platform (GCP) service account with valid credentials and permissions.
2. On the GCP infrastructure, set up the MID Server with full access to all Cloud APIs: Set Cloud API access scopes to "Allow full access to all Cloud APIs". The MID Server instance can access only the Clusters specific to the project.
3. Navigate to `sys_properties.list` and, using the admin role, configure the following properties:
   - *`sn_itom_pattern.k8s_midserver`*: Specify a valid MID Server or MID Server cluster name (MID Server cluster support starting with Discovery and Service Mapping Patterns version 1.35.0).

     Note: To target a specific service account or cluster, you can configure the *`sn_itom_pattern.k8s_<service_account_id>_midserver`* and the *`sn_itom_pattern.k8s_<service_account_id>_<clustername>_midserver`* properties. The most specific property takes precedence.

   - *`sn_itom_pattern.k8s_create_schedule_enabled`*: Set the value to **true**.

     Note: Enabling the *`sn_itom_pattern.k8s_create_schedule_enabled`* property automatically creates a serverless schedule for your cloud clusters, eliminating the need for manual scheduling. If you have an existing manual schedule and want to convert it to an automatic one, enable the property. Your manual schedule will be updated; no additional schedule will be created. An automatically created schedule has no "Max run time" defined.

   - *`sn_itom_pattern.k8s_entry_pattern`*: To use the per-namespace large-payload (LP) discovery pattern, set the value to **Kubernetes Cluster - Per-Namespace LP** (starting from Discovery and Service Mapping Patterns version 1.35.0).
4. Create and run Google Cloud Discovery

Note:

To fetch the Bearer token, while running GKE Kubernetes schedule, use the gcloud command:

```
gcloud config config-helper --
format="value(credential.access_token)"
```

Configuring gcloud in the MID Server instance grants access to the GKE cluster to fetch the token.

For the Amazon Elastic Kubernetes Service (EKS) cluster discovery, perform the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. In the ServiceNow instance, set an AWS service account with valid management account credentials and permissions.

   a. Verify that the Amazon Elastic Kubernetes Service (EKS) Cluster has a cluster role with the read-only access to all resources.

   b. Create cluster role binding between the cluster role and a Kubernetes user. For example, read-onlyuser.

   c. Create an AWS IAM role with the policy EKSReadOnly.

   d. Associate the IAM role with the Kubernetes user in one of the following ways:

      - In the cluster, edit the aws-auth ConfigMap.
      - Run the command:

        ```
        eksctl create iamidentitymapping --cluster yourClusterName --arnarn:aws:iam::yourAccountID:role/yourIAMRoleName --username read-only-user
        ```

2. Run Amazon Elastic Kubernetes Service (EKS) cluster discovery in one of two ways: Using the AWS Command Line Interface (CLI) or without using the AWS CLI. First, set the system property *sn_itom_pattern.k8s_aws_cli_to_generate_token* to use the model you choose. This system property is set to **true** by default.

   - Set this system property to **true** to use AWS CLI to generate a token.
   - Set this system property to **false** to use Assume Roles to generate a token.

   a. Run Amazon Elastic Kubernetes Service (EKS) cluster discovery using AWS CLI:

      i. Set up the MID Server with the AWS CLI configured. Configuring AWS CLI credentials grants access to the Amazon Elastic Kubernetes Service (EKS) cluster.

      Note: The user logged in to the system must be the same as the MID Server user.

      ii. To generate the Bearer token, While running the Amazon Elastic Kubernetes Service (EKS) schedule, use the AWS CLI command:`aws eks get-token --cluster-name <cluster_name>`.

      Configuring the AWS CLI user/role in the MID Server instance grants access to the Amazon Elastic Kubernetes Service (EKS) cluster to generate the token.

   b. Run Amazon Elastic Kubernetes Service (EKS) cluster discovery without using AWS CLI:

      Note:

      This feature is supported from Discovery and Service Mapping Patterns version 1.0.96 - December 2022.

   Refer to the following KB for detailed instructions: KB1182188: EKS cluster discovery using STS AssumeRoles (Without AWS CLI)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Navigate to `sys_properties.list` and, using the admin role, configure the following properties:
   - *sn_itom_pattern.k8s_midserver*: Specify a valid MID Server or MID Server cluster name (MID Server cluster support starting with Discovery and Service Mapping Patterns version 1.35.0).

     Note: To target a specific service account or cluster, you can configure the *sn_itom_pattern.k8s_<service_account_id>_midserver* and the *sn_itom_pattern.k8s_<service_account_id>_<clustername>_midserver* properties. The most specific property takes precedence.

   - *sn_itom_pattern.k8s_create_schedule_enabled*: Set the value to **true**.

     Note: Enabling the *sn_itom_pattern.k8s_create_schedule_enabled* property automatically creates a serverless schedule for your cloud clusters, eliminating the need for manual scheduling. If you have an existing manual schedule and want to convert it to an automatic one, enable the property. Your manual schedule will be updated; no additional schedule will be created. An automatically created schedule has no "Max run time" defined.

   - *sn_itom_pattern.k8s_entry_pattern*: To use the per-namespace large-payload (LP) discovery pattern, set the value to **Kubernetes Cluster - Per-Namespace LP** (starting from Discovery and Service Mapping Patterns version 1.35.0).

4. Create and run an AWS Cloud Discovery schedule.

For Microsoft Azure Kubernetes Services (AKS)- Kubernetes cluster discovery, perform the following:

1. Update to the latest Discovery and Service Mapping Patterns version.
2. In the ServiceNow instance, configure the Azure Service Account with valid Azure credentials and permission.
3. Navigate to `sys_properties.list` and, using the admin role, configure the following properties:
   - *sn_itom_pattern.k8s_midserver*: Specify a valid MID Server or MID Server cluster name (MID Server cluster support starting with Discovery and Service Mapping Patterns version 1.35.0).

     Note: To target a specific service account or cluster, you can configure the *sn_itom_pattern.k8s_<service_account_id>_midserver* and the *sn_itom_pattern.k8s_<service_account_id>_<clustername>_midserver* properties. The most specific property takes precedence.

   - *sn_itom_pattern.k8s_create_schedule_enabled*: Set the value to **true**.

     Note: Enabling the *sn_itom_pattern.k8s_create_schedule_enabled* property automatically creates a serverless schedule for your cloud clusters, eliminating the need for manual scheduling. If you have an existing manual schedule and want to convert it to an automatic one, enable the property. Your manual schedule will be updated; no additional schedule will be created. An automatically created schedule has no "Max run time" defined.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

.

- *sn_itom_pattern.k8s_entry_pattern*: To use the per-namespace large-payload (LP) discovery pattern, set the value to **Kubernetes Cluster - Per-Namespace LP** (starting from Discovery and Service Mapping Patterns version 1.35.0).

4. If you don't have local accounts with Kubernetes RBAC and want to improve pattern efficiency, navigate to **MID Server** > **Properties** and set the *sn_itom_pattern.aks_fetch_local_ad_token* property to **false**.
5. Run an Azure cloud discovery schedule.
6. Configure the MID Server in the Discovery schedules according to the cluster account type. If you don't have Local accounts with RBAC, you can ignore this step.

Cluster account type	Discovery schedule MID Server
MS Entra ID auth with Kubernetes RBAC.	Any MID Server.
MS Entra ID authentication with Azure RBAC.	Any MID Server.
Local accounts with Kubernetes RBAC.	Select the MID Server with the Azure Command Line Interface (CLI) configured. Configuring the Azure CLI credentials grants access to the AKS cluster. To fetch the Bearer token while running the AKS Kubernetes schedule, use the Azure CLI command: `az aks get-credentials --name <cluster_name> --overwrite-existing --resource-group <resourceGroup_name> --file -`.

Note:

- The user logged in to the system must be the same as the MID Server user.
- For detailed information about AKS Cluster Discovery configuration, see the AKS Cluster Discovery Configuration Details [KB1220553] article in the Now Support Knowledge Base.

Other Supported System configuration

	Property name	Property description	Type	Default value
	sn_itom_pattern.manifest_digest_image_id		Boolean	false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Other Supported System configuration (continued)

	Property name	Property description	Type	Default value
				Note: Before setting this property to **true** and running discovery: avoid duplicate records from being created by deleting all Docker image records.
	sn_itom_pattern.k8s_create_schedule_enabled	The feature flag that can be enabled/ disabled under the system properties, which is responsible to control the pattern execution. When enabled, it creates discovery schedules despite the new property value. An automatically created schedule has no "Max run time" defined.	Boolean	false
	sn_itom_k8s_run_cloud_discovery	When disabled this property executes cloud Kubernetes patterns, discovering Kubernetes clusters without creating auto schedules.	Boolean	false
MID Server	sn_itom_pattern.k8s_default_mid_server	[Default]	String	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Other Supported System configuration (continued)

	Property name	Property description	Type	Default value
		Example- Valid MID Server or MID Server cluster name (MID Server cluster support starting with Discovery and Service Mapping Patterns version 1.35.0). Applies to newly created discovery schedules only; existing schedules aren't updated.		
	sn_itom_pattern.k8s_service_account_id>_midserver*	[Based on Service Account Level] Example- Valid MID Server or MID Server cluster name (MID Server cluster support starting with Discovery and Service Mapping Patterns version 1.35.0).	String	
	sn_itom_pattern.k8s_service_account_id>_<clustername>_midserver*	[Based on Cluster name] Example- Valid MID Server or MID Server cluster name (MID Server cluster support starting with Discovery and Service Mapping Patterns version 1.35.0).	String	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Other Supported System configuration (continued)

	Property name	Property description	Type	Default value
	sn_itom_pattern.kubernetes_collect_volume	When the property is set to **True**, the data for Kubernetes Volume [cmdb_ci_kubernetes_volume] gets populated.	String	false
	sn_itom_pattern.k8s_add_workload_to_image_relation	Starting from Discovery and Service Mapping Patterns version 1.30.2, the Kubernetes patterns create an indirect-only relationship between Docker Image and workload CIs through Kubernetes Pods. Setting the property to true also creates direct relationships between Docker Image and the following workload CI types: Deployment, DaemonSet, ReplicaSet, StatefulSet, and ReplicationController.	Boolean	false
Credential Alias	sn_itom_pattern.k8s_cred_alias	[Default] Example- credential alias name	String	
	sn_itom_pattern.k8s_<service_account_id>_alias	[Based on Service Account Level]	String	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Other Supported System configuration (continued)

	Property name	Property description	Type	Default value
		Example- Valid credential alias name.		
	sn_itom_pattern.k8s_service_account_id>_<clustername>_alias	[Based on Cluster name] Example- Valid credential alias name.	String	
Prometheus Url	sn_itom_pattern.k8s_prometheusUrl	[Default] Example- Valid Prometheus Url	String	
	sn_itom_pattern.k8s_service_account_id>_prometheusUrl	[Based on Service Account Level] Example- Valid Prometheus Url	String	
	sn_itom_pattern.k8s_service_account_id>_<clustername>_prometheusUrl	[Based on Cluster name] Example- Valid Prometheus URL	String	
	sn_itom_pattern.k8s_run	[Supported Discovery Schedule run- Daily, On Demand, Weekdays, Weekends, Month Last Day, Calendar Quarter End] Example- Daily	String	
	sn_itom_pattern.k8s_batch_count	[Refers how many schedules to run in batch – default set to 5]	Integer	5

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Other Supported System configuration (continued)

	Property name	Property description	Type	Default value
		Example- 5 (Number of schedules to run in on batch)		
	sn_itom_pattern.k8s_sensor_batch_delay	[keeps the track of the time difference between two batches value contains in sec] Example- 300 (in seconds)	Integer	
	sn_itom_pattern.k8s_sensor_time	[keeps track of the current time for a batch] If this property is set, then you can use the same or you can use the dynamic timing, which will be 5 min after the system current timing. Values contains in HH:MM:SS format Example- 10:11:12 (HH:MM:SS)	String	
	sn_itom_pattern.bring_discovery_container	Available starting with Discovery and Service Mapping Patterns version 1.31.0. This property controls whether the Kubernetes, Kubernetes Event, Docker Pattern, and Amazon AWS - ECS patterns discover both Docker container	Boolean	true

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Other Supported System configuration (continued)

	Property name	Property description	Type	Default value
		and Docker image CIs, or only Docker image CIs. Check your entitlements to determine whether you have access to 2026 Packaging SKU.		
	sn_itom_pattern.k8s_entry_path	Controls which entry pattern Kubernetes cloud discovery uses. When set to **Kubernetes Cluster - Per-Namespace LP**, discovery identifies the cluster and then discovers each namespace separately, improving reliability on large clusters. The Kubernetes Cluster - Per-Namespace LP pattern is available starting from Discovery and Service Mapping Patterns version 1.35.0.	String	Kubernetes

* If a MID Server and a MID Server cluster share the same name, the MID Server cluster takes precedence. *sn_itom_pattern.k8s_<service_account_id>_midserver* takes precedence over *sn_itom_pattern.k8s_midserver* for that account. *sn_itom_pattern.k8s_<service_account_id>_<clustername>_midserver* takes precedence over both for that account and cluster.

Note: `<service_account_id>` is the account ID name under Cloud Service Accounts. For more information, see: Create Discovery schedules for cloud resources

Kubernetes Credential-less or mid-in-cluster discovery

Prerequisites for Kubernetes Credentials-less discovery:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Deploy the containerized MID Server to the Kubernetes cluster. Configuring Kubernetes credentials is unnecessary since the MID Server in Kubernetes cluster automatically discovers the API server and authenticate.

Configuring execution pattern attributes for Credentials-less discovery

Field	Description
URL	Enter any one the of following value in URL field: https://cluster Or https://kubernetes.default.svc
namespace	The namespaces that the system passes in the Kubernetes Discovery Configuration. Enter one of the following values: • Individual namespace: enter the namespace and then "kube-system". For example: `dev,kube-system` • The default value. Enter:`default,kube-system` • Multipile namespaces: enter the namespaces, use a comma (,) to separate the values, and then enter "kube-system". For example: `automation,application,test,kube-system` • All namespaces: Use an asterisk (*) to enter all namespaces.
cluster_name	Enter Unique name.

Data collected by Discovery during horizontal discovery

Kubernetes pattern

Kubernetes Cluster [cmdb_ci_kubernetes_cluster]

Field	Description
Name [name]	The name of the kube-controller-manager leader.
Kubernetes UID [k8s_uid]	The kube-system namespace UID [supported versions: 1.0.92 and later]
IP Address [ip_address]	The identifier for the host IP of the Kubernetes apiserver.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes Cluster [cmdb_ci_kubernetes_cluster] (continued)

Field	Description
Port [port]	The identifier for the Kubernetes apiserver port.
Namespace [namespace]	This value shows the namespaces the system passed in the Kubernetes Discovery Configuration.
Event Timestamp [event_timestamp]	The timestamp of the latest event created on this Kubernetes cluster at the time of the discovery.

Kubernetes Node [cmdb_ci_kubernetes_node]

Field	Description
The virtual aspect of the Kubernetes node. Data relating to the physical aspect of the Kubernetes node is stored under Linux server.	
Name [name]	The name of the Kubernetes node. The format can be only the name of the machine or the full name consisting of the name and the hostname: `<name>.<hostname>`.
Kubernetes UID [k8s_uid]	The identifier for the Kubernetes node UUID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Operational status [operational_status]	The operational status of the Kubernetes node.

Kubernetes Service [cmdb_ci_kubernetes_service]

Field	Description
Name [name]	The name of the Kubernetes service.
Selector [selector]	A comma delimited list of the label selectors specified in the Kubernetes configuration that are used to select target pods.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes service belongs.
Kubernetes UID [k8s_uid]	The Kubernetes service UUID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes Pod [cmdb_ci_kubernetes_pod]

Field	Description
Name [name]	The name of the Kubernetes pod.
Kubernetes UID [k8s_uid]	The Kubernetes pod UUID.
Resource version [resource_version]	The resource version of the Kubernetes pod.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes pod belongs.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
State [state]	The Kubernetes pod status: Pending, Running, Succeeded, Failed, and Unknown.

Kubernetes Cronjob [cmdb_ci_kubernetes_cronjob]

Field	Description
Name [name]	The name of the Kubernetes cronjob.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes cronjob belongs.
Kubernetes UID [k8s_uid]	The Kubernetes cronjob UUID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.

Kubernetes Job [cmdb_ci_kubernetes_job]

Field	Description
Name [name]	The name of the Kubernetes job.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes job belongs.
Kubernetes UID [k8s_uid]	The Kubernetes job UUID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.

Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset]

Field	Description
Name [name]	The name of the Kubernetes daemonset.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes daemonset belongs.
Kubernetes UID [k8s_uid]	The Kubernetes daemonset UUID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset] (continued)

Field	Description
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Pods Available [pods_avail]	The number of available pods.
Pods Failed [pods_failed]	The number of pods in the Failed phase.
Pods Running [pods_running]	The number of pods in the Running phase.
Pods Succeeded [pods_succeeded]	The number of pods in the Succeeded phase.
Pods Waiting [pods_waiting]	The number of pods in the Waiting phase.

Kubernetes Ingress [cmdb_ci_kubernetes_ingress]

Field	Description
Name [name]	The name of the Kubernetes ingress.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes ingress belongs.
Kubernetes UID [k8s_uid]	The Kubernetes ingress UID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.

Kubernetes Deployment [cmdb_ci_kubernetes_deployment]

Field	Description
Name [name]	The name of the Kubernetes deployment.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes deployment belongs.
Kubernetes UID [k8s_uid]	The Kubernetes deployment UID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Total Replicas [total_replicas]	The number of replicas in this deployment.
Desired Replicas [desired_replicas]	The number of replicas in the desired phase.
Available Replicas [available_replicas]	The number of available replicas.
Unavailable Replicas [unavailable_replicas]	The number of replicas in the unavailable phase.
Updated Replicas [updated_replicas]	The number of updated replicas.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]

Field	Description
Name [name]	The name of the Kubernetes replicaset.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes replicaset belongs.
Kubernetes UID [k8s_uid]	The Kubernetes replicaset UID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Total Replicas [total_replicas]	The number of replicas in this replicaset.
Desired Replicas [desired_replicas]	The number of replicas in the desired phase.
Available Replicas [available_replicas]	The number of available replicas.
Unavailable Replicas [unavailable_replicas]	The number of replicas in the unavailable phase.
Updated Replicas [updated_replicas]	The number of updated replicas.

Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]

Field	Description
Name [name]	The name of the Kubernetes replication controller.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes replication controller belongs.
Kubernetes UID [k8s_uid]	The Kubernetes replication controller UID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Total Replicas [total_replicas]	The number of replicas in this replication controller.
Desired Replicas [desired_replicas]	The number of replicas in the desired phase.
Available Replicas [available_replicas]	The number of available replicas.
Unavailable Replicas [unavailable_replicas]	The number of replicas in the unavailable phase.
Updated Replicas [updated_replicas]	The number of updated replicas.

Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]

Field	Description
Name [name]	The name of the Kubernetes statefulset.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset] (continued)

Field	Description
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes statefulset belongs.
Kubernetes UID [k8s_uid]	The Kubernetes statefulset UID.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Total Replicas [total_replicas]	The number of replicas in this statefulset.
Desired Replicas [desired_replicas]	The number of replicas in the desired phase.
Available Replicas [available_replicas]	The number of available replicas.
Unavailable Replicas [unavailable_replicas]	The number of replicas in the unavailable phase.
Updated Replicas [updated_replicas]	The number of updated replicas.

Docker Container [cmdb_ci_docker_container]

Field	Description
The component that runs the docker image.	
Container id [container_id]	The unique identifier for the Kubernetes docker container. In cases where duplicate records are created, deduplication tasks appear once discovery runs. For information on how to resolve these tasks, see the Making docker container identifier independent [KB1443042] article in the ServiceNow® Knowledge Base.

Docker Image [cmdb_ci_docker_image]

Field	Description
An executable package of an application and its related software that can be instantiated by a docker container.	
Image id [image_id]	The identifier for the Kubernetes docker image.
Name [name]	The name of the Kubernetes docker image.

Linux Server [cmdb_ci_linux_server]

Field	Description
The server that hosts the Kubernetes node.	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux Server [cmdb_ci_linux_server] (continued)

Field	Description
Name [name]	The name of the Linux server powering the Kubernetes node.
Host name [host_name]	The hostname of the Linux server.
Operating System [os]	The operating system deployed on this Linux server.
Kernel Release [kernel_release]	The version of the Linux kernel operating system deployed on this Linux server.
RAM (MB) [ram]	The size of RAM installed on this Linux server.
IP Address [ip_address]	The IP address of the Linux server.
CPU type [cpu_type]	The CPU architecture of the Linux server hosting the Kubernetes node.
CPU count [cpu_count]	The number of CPUs on the Linux server hosting the Kubernetes node.
Serial number [serial_number]	The serial number of the Linux server hosting the Kubernetes node.

Key Value [cmdb_key_value]

Field	Description
This configuration item contains Kubernetes labels. Labels are key/value pairs that are attached to objects, such as pods.	
Key [key]	The key of the Kubernetes pod or Kubernetes service *`Key Value`* parameter.
Value [value]	The value of the Kubernetes pod or Kubernetes service *`Key Value`* parameter.
Configuration item [configuration_item]	References one of the following CI tables, based on the key-value pair: • Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset] • Kubernetes Deployment [cmdb_ci_kubernetes_deployment] • Kubernetes Ingress [cmdb_ci_kubernetes_ingress] • Kubernetes Namespace [cmdb_ci_kubernetes_namespace] • Kubernetes Node [cmdb_ci_kubernetes_node]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
	• Kubernetes Pod [cmdb_ci_kubernetes_pod] • Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset] • Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller] • Kubernetes Service [cmdb_ci_kubernetes_service] • Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]

Kubernetes Volume [cmdb_ci_kubernetes_volume]

Field	Description
Kubernetes UID [k8s_uid]	The Kubernetes volume UUID.
Mount Path [mount_path]	The path for accessing this Kubernetes volume.
Name [name]	The name of the Kubernetes volume.
Namespace [namespace]	The Kubernetes namespace to which this Kubernetes volume belongs.
Kubernetes Cluster [cluster]	The name of the cluster that contains this resource.
Volume ID [volume_id]	The ID of the Kubernetes volume.

Kubernetes Namespace [cmdb_ci_kubernetes_namespace]

Field	Description
Name [name]	The name of the Kubernetes namespace.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
State [state]	The Kubernetes namespace phase: Active or Terminating.

Collect OpenShift info pattern extension

OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]

Field	Description
Name [name]	The name of the OpenShift deployment configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OpenShift Deployed Config [cmdb_ci_openshift_dep_conf] (continued)

Field	Description
Namespace [namespace]	The name of the namespace containing the deployment configuration.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Url [url]	The URL of the OpenShift deployed configuration, available only for Kubernetes versions earlier than 1.16.

OpenShift Build Config [cmdb_ci_openshift_build_conf]

Field	Description
Name [name]	The name of the OpenShift build configuration.
Namespace [namespace]	The name of the OpenShift namespace containing the build configuration.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Url [url]	The URL of the OpenShift build configuration, available only for Kubernetes versions earlier than 1.16.

OpenShift Source2Image [cmdb_ci_openshift_source_2_image]

Field	Description
Name [name]	The name of the OpenShift source image.
To [to]	Related image.
Parent ID [parent_id]	The ID of the OpenShift source image.

OpenShift Route [cmdb_ci_openshift_route]

Field	Description
Name [name]	The name of the OpenShift route.
Namespace [namespace]	The name of the namespace containing the OpenShift route.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Kubernetes Cluster [cluster]	References the Kubernetes Cluster [cmdb_ci_kubernetes_cluster] table.
Url [url]	The URL of the OpenShift route, available only for Kubernetes versions earlier than 1.16.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OpenShift Route [cmdb_ci_openshift_route] (continued)

Field	Description
Host [host]	The target host of the OpenShift route.
Port [port]	The target port of the OpenShift route.

OpenShift Group [cmdb_ci_openshift_group]

Field	Description
Name [name]	The name of the OpenShift group.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Url [url]	The URL of the OpenShift group, available only for Kubernetes versions earlier than 1.16.

OpenShift User [cmdb_ci_openshift_user]

Field	Description
Name [name]	The name of the OpenShift user.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Url [url]	The URL of the OpenShift user, available only for Kubernetes versions earlier than 1.16.
Full Name [full_name]	The full name of the OpenShift user.

OpenShift Project [cmdb_ci_openshift_project]

Field	Description
Name [name]	The name of the OpenShift project.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Url [url]	The URL of the OpenShift project, available only for Kubernetes versions earlier than 1.16.

OpenShift Image [cmdb_ci_openshift_images]

Field	Description
Name [name]	The name of the OpenShift image.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Url [url]	The URL of the OpenShift image, available only for Kubernetes versions earlier than 1.16.
Docker Image Metadata ID [docker_image_metadata_id]	The ID of the docker image.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OpenShift Image [cmdb_ci_openshift_images] (continued)

Field	Description
Docker Image Metadata Parent ID [docker_image_metadata_parent_id]	The ID of the image parent.
Architecture [arch]	The architecture of the image.
Size [size]	The image size.
Hostname [hostname]	The hostname related to the image.

OpenShift Image Stream [cmdb_ci_openshift_images_stream]

Field	Description
Name [name]	The name of the OpenShift image stream.
Kubernetes UID [k8s_uid]	The Kubernetes UID of this resource.
Url [url]	The URL of the OpenShift image stream, available only for Kubernetes versions earlier than 1.16.
Namespace [namespace]	The name of the namespace containing the OpenShift image stream.

OpenShift Docker Image Repository [cmdb_ci_openshift_docker_images_repository]

Field	Description
Name [name]	The name of the OpenShift docker image repository.
Parent ID [parent_id]	The ID of the parent system.

Key Value [cmdb_key_value]

Field	Description
This configuration item contains OpenShift labels. Labels are key/value pairs that are attached to objects, such as deployed configurations and routes.	
Key [key]	The key of the OpenShift deployed configuration or route *`Key Value`* parameter.
Value [value]	The value of the OpenShift deployed configuration or route *`Key Value`* parameter.
Configuration item [configuration_item]	References one of the following CI tables, based on the key-value pair:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
	• OpenShift Deployed Config [cmdb_ci_openshift_dep_conf] • OpenShift Route [cmdb_ci_openshift_route]

Collect Container Repository pattern extension

Container Repository [cmdb_ci_container_repository]

Field	Description
Name [name]	The name of the container repository.

Container Repository Entry [cmdb_ci_container_repository_entry]

Field	Description
Name [name]	The name of the container repository entry.
Category [category]	The category of the container repository entry.

Kubernetes deployment diagrams

The graphic illustrates CIs that are part of Kubernetes discovery.

Note: This Dependency Views map was simplified for clarity. Your Kubernetes deployments may contain many more CIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Components of the Kubernetes deployment

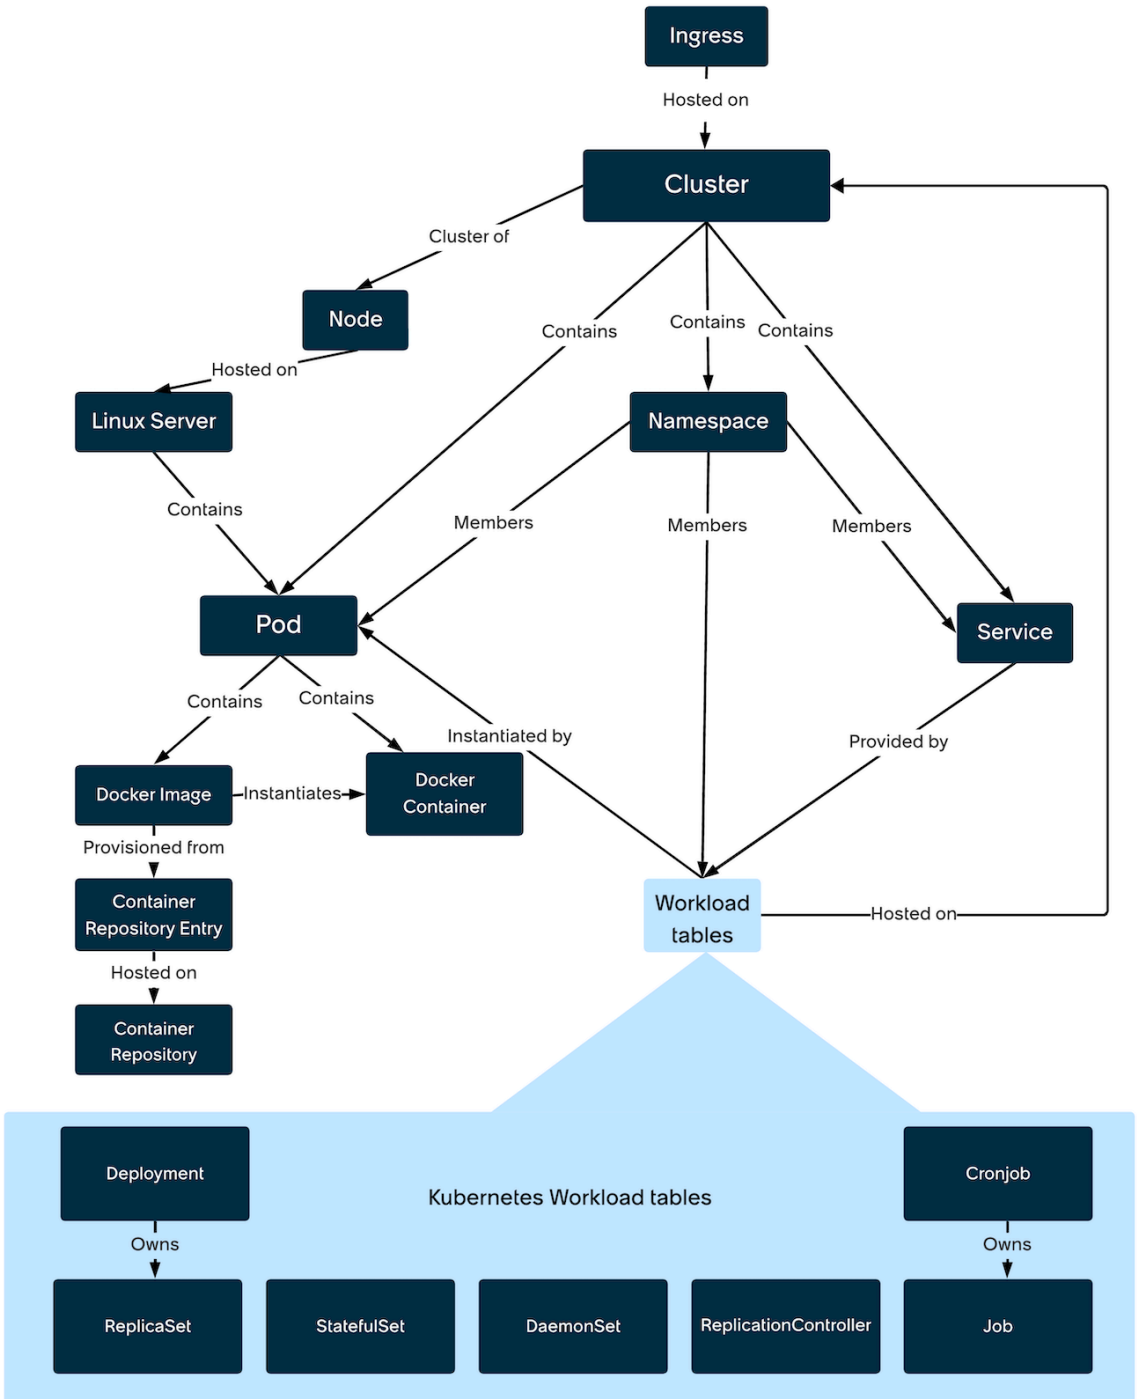

Components of the Kubernetes deployment including OpenShift

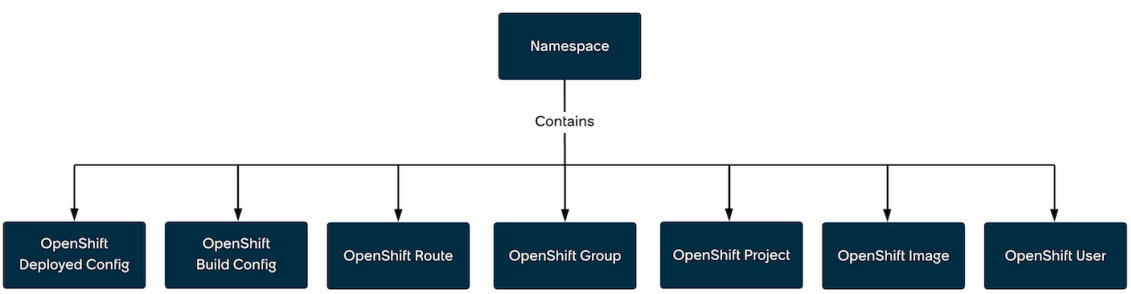

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships and references created by the Kubernetes pattern

These relationships and references are created by the Kubernetes pattern. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Docker Container [cmdb_ci_docker_container]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]*
Docker Image [cmdb_ci_docker_image]	Instantiates::Instantiated by	Docker Container [cmdb_ci_docker_container]*
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	Kubernetes Ingress [cmdb_ci_kubernetes_ingress]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	Kubernetes Namespace [cmdb_ci_kubernetes_namespace]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Cluster of::Cluster	Kubernetes Node [cmdb_ci_kubernetes_node]*
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	Kubernetes Pod [cmdb_ci_kubernetes_pod]*
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	Kubernetes Service [cmdb_ci_kubernetes_service]*
Kubernetes Cronjob [cmdb_ci_kubernetes_cronjob]	Hosted on::Hosts	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Cronjob [cmdb_ci_kubernetes_cronjob]	Owns::Owned by	Kubernetes Job [cmdb_ci_kubernetes_job]
Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset]	Hosted on::Hosts	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset]	Instantiates::Instantiated by	Docker Image [cmdb_ci_docker_image]
Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset]	Instantiates::Instantiated by	Kubernetes Pod [cmdb_ci_kubernetes_pod]
Kubernetes Deployment [cmdb_ci_kubernetes_deployment]	Hosted on::Hosts	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Deployment [cmdb_ci_kubernetes_deployment]	Instantiates::Instantiated by	Docker Image [cmdb_ci_docker_image]
Kubernetes Deployment [cmdb_ci_kubernetes_deployment]	Instantiates::Instantiated by	Kubernetes Pod [cmdb_ci_kubernetes_pod]
Kubernetes Deployment [cmdb_ci_kubernetes_deployment]	Owns::Owned by	Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]
Kubernetes Deployment [cmdb_ci_kubernetes_deployment]	Owns::Owned by	Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]
Kubernetes Ingress [cmdb_ci_kubernetes_ingress]	Instantiates::Instantiated by	Kubernetes Service [cmdb_ci_kubernetes_service]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Kubernetes Job [cmdb_ci_kubernetes_job]	Hosted on::Hosts	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Job [cmdb_ci_kubernetes_job]	Instantiates::Instantiated by	Kubernetes Pod [cmdb_ci_kubernetes_pod]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes Cronjob [cmdb_ci_kubernetes_cronjob]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes Deployment [cmdb_ci_kubernetes_deployment]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes Ingress [cmdb_ci_kubernetes_ingress]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes Job [cmdb_ci_kubernetes_job]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes Service [cmdb_ci_kubernetes_service]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]
Kubernetes Node [cmdb_ci_kubernetes_node]	Hosted on::Hosts	Linux Server [cmdb_ci_linux_server]*
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Contains::Contained by	Docker Container [cmdb_ci_docker_container]*
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Contains::Contained by	Docker Image [cmdb_ci_docker_image]*
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Contains::Contained by	Kubernetes Volume [cmdb_ci_kubernetes_volume]*
Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]	Hosted on::Hosts	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]	Instantiates::Instantiated by	Docker Image [cmdb_ci_docker_image]
Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]	Instantiates::Instantiated by	Kubernetes Pod [cmdb_ci_kubernetes_pod]
Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]	Hosted on::Hosts	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]	Instantiates::Instantiated by	Docker Image [cmdb_ci_docker_image]
Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]	Instantiates::Instantiated by	Kubernetes Pod [cmdb_ci_kubernetes_pod]
Kubernetes Service [cmdb_ci_kubernetes_service]	Provided By::Provides	Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset]
Kubernetes Service [cmdb_ci_kubernetes_service]	Provided By::Provides	Kubernetes Deployment [cmdb_ci_kubernetes_deployment]
Kubernetes Service [cmdb_ci_kubernetes_service]	Provided By::Provides	Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]
Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]	Hosted on::Hosts	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]	Instantiates::Instantiated by	Docker Image [cmdb_ci_docker_image]
Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]	Instantiates::Instantiated by	Kubernetes Pod [cmdb_ci_kubernetes_pod]
Linux Server [cmdb_ci_linux_server]	Contains::Contained by	Kubernetes Pod [cmdb_ci_kubernetes_pod]*
Linux Server [cmdb_ci_linux_server]	Managed by::Manages	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]*

* These relationships are also created by the Kubernetes Event pattern.

CI references

CI	Field	Referenced CI
Container Environment Variables [cmdb_container_environment_variables]	Configuration item [configuration_item]	Docker Container [cmdb_ci_docker_container]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	References one of the following CI tables, based on the key-value pair: • Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset] • Kubernetes Deployment [cmdb_ci_kubernetes_deployment] • Kubernetes Ingress [cmdb_ci_kubernetes_ingress] • Kubernetes Namespace [cmdb_ci_kubernetes_namespace]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
		• Kubernetes Node [cmdb_ci_kubernetes_node] • Kubernetes Pod [cmdb_ci_kubernetes_pod] • Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset] • Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller] • Kubernetes Service [cmdb_ci_kubernetes_service] • Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]
Kubernetes Cronjob [cmdb_ci_kubernetes_cronjob]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes DaemonSet [cmdb_ci_kubernetes_daemonset]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Deployment [cmdb_ci_kubernetes_deployment]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Ingress [cmdb_ci_kubernetes_ingress]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Job [cmdb_ci_kubernetes_job]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Node [cmdb_ci_kubernetes_node]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Service [cmdb_ci_kubernetes_service]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes StatefulSet [cmdb_ci_kubernetes_statefulset]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Volume [cmdb_ci_kubernetes_volume]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	Linux Server [cmdb_ci_linux_server]

CI relationships created by the Kubernetes cloud patterns

These additional relationships are created by the Google Cloud Platform (GCP) – Get Kubernetes Clusters, Amazon AWS Cloud - Get Kubernetes Clusters, and Microsoft Azure - Get Kubernetes Clusters patterns.

CI	Relationship	CI
Azure Datacenter [cmdb_ci_azure_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]

CI relationships and references created by the Kubernetes extension sections

The following extension sections of the Kubernetes pattern identify additional relationships and references. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

Collect OpenShift info

CI relationships

CI	Relationship	CI
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	OpenShift Docker Image Repository [cmdb_ci_openshift_docker_images_repository]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	OpenShift Group [cmdb_ci_openshift_group]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	OpenShift Image [cmdb_ci_openshift_images]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	OpenShift Project [cmdb_ci_openshift_project]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	OpenShift Source2Image [cmdb_ci_openshift_source_2_image]
Kubernetes Cluster [cmdb_ci_kubernetes_cluster]	Contains::Contained by	OpenShift User [cmdb_ci_openshift_user]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	OpenShift Build Config [cmdb_ci_openshift_build_conf]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	OpenShift Image Stream [cmdb_ci_openshift_images_stream]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained by	OpenShift Route [cmdb_ci_openshift_route]
Kubernetes Service [cmdb_ci_kubernetes_service]	Provided By::Provides	OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]
OpenShift Build Config [cmdb_ci_openshift_build_conf]	Contains::Contained by	Tracked Configuration file [cmdb_ci_config_file_tracked]
OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]	Contains::Contained by	Tracked Configuration file [cmdb_ci_config_file_tracked]
OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]	Owns::Owned by	Kubernetes ReplicaSet [cmdb_ci_kubernetes_replicaset]
OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]	Owns::Owned by	Kubernetes ReplicationController [cmdb_ci_kubernetes_replicationcontroller]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	References one of the following CI tables, based on the key-value pair: • OpenShift Deployed Config [cmdb_ci_openshift_dep_conf] • OpenShift Route [cmdb_ci_openshift_route]
OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
OpenShift Route [cmdb_ci_openshift_route]	Kubernetes Cluster [cluster]	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]

Collect Container Repository

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Docker Image [cmdb_ci_docker_image]	Provisioned From::Provisioned	Container Repository Entry [cmdb_ci_container_repository_entry]
Container Repository Entry [cmdb_ci_container_repository_entry]	Hosted on::Hosts	Container Repository [cmdb_ci_container_repository]

Istio Service Mesh

Prerequisites for Istio Service Mesh extension:

- Verify that Istio Service Mesh and Prometheus components are configured on the Kubernetes cluster.
- Verify that Prometheus discovers the service connection information using the `queryistio_requests_total` command.
- Verify that the application services are connected, and verify service-to-service traffic flow in Kiali graph.

For more information on the Bookinfo application, see: https://istio.io/latest/docs/examples/bookinfo/

CI	Relationship	CI
Kubernetes Service [cmdb_ci_kubernetes_service]	Connects to::Connected by	Kubernetes Service [cmdb_ci_kubernetes_service]

Data collected by Service Mapping during tag-based discovery

Service Mapping uses tag-based discovery to create application service maps including the Kubernetes components. Service Mapping comes with the following preconfigured CI relationships used for tag-based discovery. These CI relationships are available from the 1.0.68 release on ServiceNow Store.

CI	Relationship	CI
Kubernetes Service [cmdb_ci_kubernetes_service]	Contained By::Contains	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
OpenShift Project [cmdb_ci_openshift_project]	Contained by::Contains	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Contained by::Contains	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Cluster::Cluster of	Kubernetes Service [cmdb_ci_kubernetes_service]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Build Config [cmdb_ci_openshift_build_conf]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Route [cmdb_ci_openshift_route]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Image Stream [cmdb_ci_openshift_images_stream]

Kubernetes dashboard

After Discovery finishes discovering components of the Kubernetes deployment, you can navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** and use the Kubernetes Explorer dashboard to view the Kubernetes environments and resources of your organization. To use the enhanced Kubernetes dashboard, verify you have Discovery Admin Workspace starting from version 1.3.1 (August 2024 Store). For more information about Kubernetes Explorer, see Kubernetes Explorer.

Troubleshooting

If the mapping process does not proceed as you expected, follow the following suggestions.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again. **Note:** If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

To run the Kubernetes pattern in Debug mode, refer to KB0832567 for operational information.

Container image scanning for software decomposition

The ITOM Visibility apps, Discovery and Service Mapping Patterns and Kubernetes Visibility Agent integrate with Aqua Trivy to collect data on container images and OS packages. You can increase your control over container deployment by having visibility to the container components.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

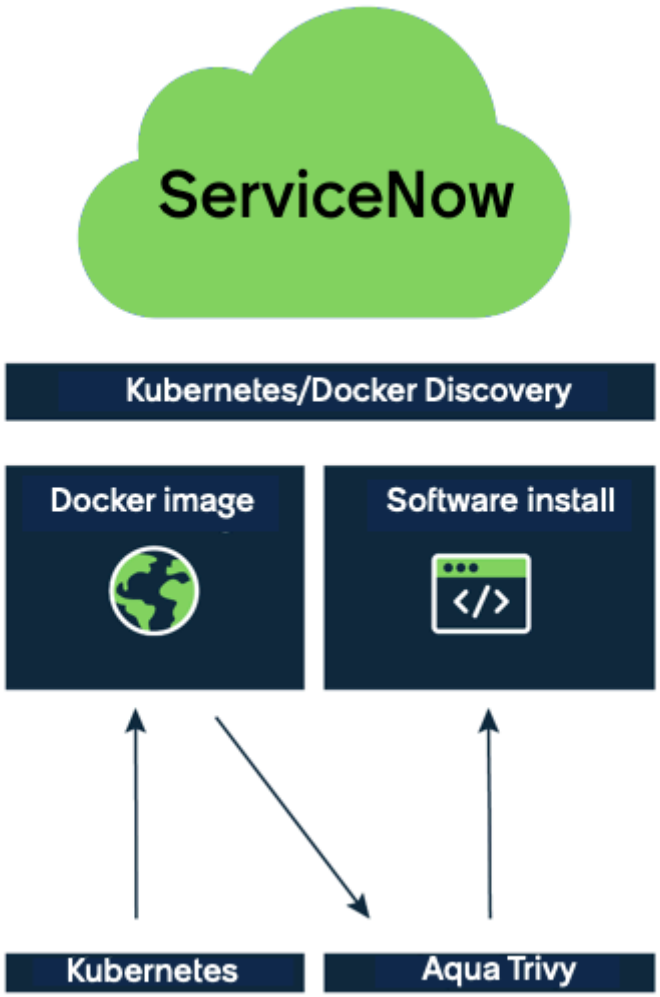

Benefits of image scanning

Scanning your containers gives you visibility into what's inside Kubernetes or Docker containers or OS packages. Scanning images can assist you in multiple ways.

- It helps you identify software installed in containers for regulatory and compliance use cases.
- It helps you adhere to company policies like usage of golden images, outdated software, mandatory labels, or configuration policies.
- It also helps you manage licensed software running in containers.
- You can also get the service context by using tags, and service mesh to understand their impact on your organization.

Image scanning use cases with ITOM Visibility

You can use two ITOM Visibility apps to scan container images, Discovery and Service Mapping Patterns and Kubernetes Visibility Agent. Patterns is a feature set used by Discovery, Cloud Discovery, and Service Mapping. Kubernetes Visibility Agent is a feature of Agent Client Collector. While Kubernetes Visibility Agent (formerly known as CNO-V) is more suitable for Kubernetes and dynamic containerized workloads, pattern-based discovery is more suitable for non-Kubernetes Docker containers.

Use case # 1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Once an application has been packaged up in container images, a security professional can scan the base image and the final image for vulnerabilities. The scan identifies OS packages, software dependencies, and application records. This is specifically for Containerized MSSQL Server.

Visibility methods for use case # 1

Visibility methods	Method characteristics	What's discovered

Discovery and Service Mapping Patterns and Aqua Trivy:
• Best suited for self hosted or cloud Kubernetes deployments with access to bearer tokens and credentials.
• Supports public and self-hosted repository image scanning.
	• Pattern-based discovery without Cloud discovery:
	  ○ Uses a bearer token.
	  ○ Manually created Kubernetes discovery schedule per cluster.
	• Pattern-based discovery with Cloud discovery:
	  ○ No bearer token required.
	  ○ Uses cloud credentials.
	  ○ Automatic creation of Kubernetes discovery schedule.
	• For more information on scanning images using Aqua Trivy, see Scan container images.
		Discovered using Discovery and Service Mapping Patterns:
		• Kubernetes clusters
		• Kubernetes Services
		• Kubernetes topology
		• Docker containers and images
		• Kubernetes deployment including OpenShift
		• Namespace
		• Labels and tags
		• Software in containers
		• Account region details can only be discovered by Cloud Discovery
		• Docker Pattern collects data about specific objects in a Docker engine, running on a Linux host
		For more information, see:
		• Container image discovery
		• Kubernetes discovery using patterns
		• Docker virtualization

Kubernetes Visibility Agent and Aqua Trivy: • Best suited for deployment by cloud native app teams. • Optional ability to monitor Kubernetes with AIOps. • For cloud environments, only AWS ECR (Public and Private) is supported.	Kubernetes Visibility Agent -based discovery doesn't require credential set up, and no need for MID Server. Access is through ServiceAccount/ClusterRole. The installation is via Helm Chart or Kubernetes YAML file. The discovery is run near real-time.	Discovered using Kubernetes Visibility Agent • Kubernetes Namespaces • Nodes and Pods • Deployments • Statefulsets • Daemonsets • Replicasets • Jobs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Visibility methods for use case # 1 (continued)

Visibility methods	Method characteristics	What's discovered
	Use Kubernetes Explorer to download SBOM.	• Cronjobs • Services • Docker container • Docker image • Container repository • Account region details can only be discovered by Cloud Discovery

Use case #2

A compliance officer can generate an#SBOM#o obtain a detailed list of the dependencies of the container image. This ensures that the software complies with industry regulations.

Visibility methods for use case #2

Visibility method	Method characteristics
Kubernetes pattern or Docker pattern	SBOM creation is part of the container scanning.
Kubernetes Visibility Agent	SBOM creation is also a part of the container scanning, but using ACC is best suited for organizations that need flexibility to perform both full and continues discovery.

Use case #3

An engineer found a defect in a custom-built image and needs to find all Kubernetes pods that are running using that image.

Visibility methods for use case #3

Visibility method	Method characteristics	What's discovered
Kubernetes pattern	Aqua Trivy container scanning isn't required. You can identify the pods using Patterns.	• Kubernetes Clusters • Kubernetes Containers • Kubernetes Services • Labels • Pods • Images • Tags

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Visibility methods for use case #3 (continued)

Visibility method	Method characteristics	What's discovered
Kubernetes pattern with Cloud discovery	Aqua Trivy container scanning isn't required. You can identify the pods using Patterns.	All the of the above and account or region details

Use case #4

An engineer finds a defect in a custom-built image and needs to find all Docker containers (non Kubernetes) that are running using that image.

Visibility method	Method characteristics	What's discovered
Horizontal Discovery of VM running Docker (Docker pattern)	Aqua Trivy container scanning isn't required. You can identify the pods using Patterns.	See: Docker virtualization

Image scanning with Discovery and Service Mapping Patterns

Kubernetes and Docker patterns integrate with the Aqua Trivy tool and run scheduled jobs to discover container images and OS packages at fixed intervals of 10 images per minute. During the scan, the pattern indicates the scanning status. The pattern discovers OS packages that are related to an image. Then, it finds the image command attributes like the CI class. Based on the command attributes the pattern creates application records. In addition, the pattern uses enriched scripts to add details to the application records. After that, the pattern maps the relations between the OS packages and the containers.

Part of the data is populated in CMDB tables and part of it in transformation tables (non-CMDB temporary tables). The transformation tables are installed with the pattern. For example, the information you get by scanning includes origin registry, software name, and version.

Related topics

Cloud Native Operations for Visibility

Link software package information to images only

Link software package information to images rather than containers to avoid empty records that associate to short-lived containers.

Before you begin
Role required: admin

About this task
Software package information is linked to both images and containers when the *sn_itom_pattern.image_scan_process_container* system property is set to its default value of **true**. To link this data to images only, set the property value to **false**.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** field, search for `sn_itom_pattern.image_scan_process_container`.
3. Select the *sn_itom_pattern.image_scan_process_container* property.
4. In the **Value** field, enter `false`.
5. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Scan container images

Container image discovery

Map MID Server to container image repository

If network access or datacenter location requires a specific MID Server, map MID Server to a private container image repository.

Before you begin

- Verify that you have at least version 1.30.2 of Discovery and Service Mapping Patterns.
- Verify that the Linux MID Server has the following requirements:
  - 8GB RAM with 4GB for the Java Virtual Machine (JVM).
  - MID Server **Capability** configured to **Scan Container Images**. For more information on configuring capabilities, see MID Server capabilities.
  - Can access the image repository.
  - Has Aqua Trivy, starting 0.44.0, installed on the MID Server. The last version validated is 0.68.2. For more information on installing Aqua Trivy, see the Container Image scan feature overview [KB1218504] article in the Now Support Knowledge Base.
- Verify that the container image repository credentials are configured for the private repositories. For more information, see Create and test your credentials and Container image repository credentials.
- Verify that a container image discovery has run for Docker, Kubernetes, or Amazon Elastic Container Service (Amazon ECS) discovery. For more information, see:
  - Docker virtualization
  - Kubernetes discovery using patterns
  - Amazon ECS resource discovery with Patterns
- Verify that the Container Repository table contains the repository names required for this mapping. You can view the container repository list by entering `cmdb_ci_container_repository_list.do` in the navigation filter.

Role required: discovery_admin

About this task

By default, container image scanning selects a MID Servers with the **Scan Container Images** capability from the available MID Servers. You can configure the Image Registry to Mid Server Mapping table to map specific MID Servers to private repositories.

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. In the **Label** column, enter `Image Registry to Mid Server Mapping`.
3. Select the **Image Registry to Mid Server Mapping** table.
4. In the **Related Links** section, select **Show List**.
5. Select **New**.
6. In the form, fill in the fields:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Image Registry to Mid Server Mapping fields

Field	Description
Registry URL	Registry domain or host from the Container Repository table, up to the first slash. For example, if the full path is `registry.example.com/project/image`, enter `registry.example.com`.
Mid server	Linux MID Server for this repository that has the **Scan Container Images** capability configured.
Active	Option for enabling this mapping.

7. Select **Submit**.

What to do next
Scan the container images. For more information, see Scan container images.

Configure proxy bypass for container registries

Configure a proxy bypass for container image scans to reach specific internal or private registries that require a direct connection.

Before you begin
- Verify that you have at least version 1.30.2 of Discovery and Service Mapping Patterns.
- Verify that you have configured a proxy on the MID Server. For more information, see MID Server parameters.

Role required: discovery_admin

About this task
When a proxy is configured on the MID Server, all container image scanning goes through that proxy by default. If some internal or private registries aren't reachable through the proxy, use the *sn_itom_pattern.container_image_scan_no_proxy* system property to specify registries that the MID Server contacts directly.

Procedure
1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `sn_itom_pattern.container_image_scan_no_proxy`.
3. Select the **sn_itom_pattern.container_image_scan_no_proxy** system property.
4. In the **Value** field, enter the hosts, domains, or IP ranges that bypass the proxy, separated by commas.
   Examples:
   - Hostname: internal.registry.mycompany.com
   - Domains: internal.mycompany.com,.internal.mycompany.com
   - IP ranges: 10.0.0.0/8,172.16.0.0/12
   - Hostname, domain, and IP ranges combined: internal.registry.mycompany.com,internal.mycompany.com,.internal.mycompany.com,10.0.0.0/8,172.16.0.0/12

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: To stop using the bypass, clear the values in this field.

5. Select **Update**.

What to do next

Scan the container images. For more information, see Scan container images.

Scan container images

Enable container image scanning using the Aqua Trivy tool.

Before you begin

Note: The pattern supports only the Linux operating system and runs on a Linux MID Server.

Verify the following setup:

- Discovery and Service Mapping Patterns, starting 1.7.0 (August 2023 store release).
- CMDB CI Class Models, starting 1.43.0.
- Visibility Content, starting 6.7.2.
- For the container images scan to run successfully, verify that all the MID Servers in your environment are configured for specific capabilities and are not set to **All**. For more information on configuring capabilities, see MID Server capabilities.
- MID Server requirements:
  - 8GB RAM with 4GB for the Java Virtual Machine (JVM)
  - MID Server **Capability** configured to **Scan Container Images**
  - Can access the image repository
- Aqua Trivy starting 0.44.0 is installed on the MID Server. The last version validated is 0.68.2. For more information on installing Aqua Trivy, see the Container Image scan feature overview [KB1218504] article in the Now Support Knowledge Base.
- To collect Docker or Kubernetes container image and environment records, you must first run the Docker or Kubernetes discovery before enabling the Container Image Scan feature. For more information, see Kubernetes discovery using patterns and Docker virtualization.
- Starting with Discovery and Service Mapping Patterns version 1.18.0, the Container Image Scan feature also supports the following:
  - Using a proxy MID Server to scan the container images in the following repositories: public, self-hosted private, and Amazon Elastic Container Registry (Amazon ECR) repositories. To use a proxy MID Server, ensure the MID Server proxy parameters are correctly configured. For more information on proxy parameters, see MID Server parameters. For information on bypassing the proxy MID Server for specific internal or private container registries, see Configure proxy bypass for container registries.
  - Scanning container images in self-hosted private repositories. To scan these repositories, you must create container image repository credentials. For more information, see Create and test your credentials and Container image repository credentials.
  - Collecting image records from Amazon ECR in both public and private repositories. To collect these Amazon ECR image records, you must first run the Amazon ECS resource discovery before enabling the Container Image Scan. The AWS cloud service account created for the Amazon ECS resource discovery is also used for the container image scan. For more information about Amazon ECS resource discovery, see Amazon ECS resource discovery with Patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Starting with version 1.27.0, Discovery and Service Mapping Patterns enables you to control whether to link software packages to containers or only to images. For more information, see Link software package information to images only.
- Starting with Discovery and Service Mapping Patterns version 1.30.2, you can map MID Servers to specific repositories. Use this option when network access or datacenter location requires specific MID Servers for certain repositories. For more information, see Map MID Server to container image repository.

Role required: admin

Procedure

1. Enable the Discovery and Service Mapping Patterns Container Image Scan by setting the system property.
   a. In the Navigation filter, enter `sys_properties.list`.
   b. In the **Name** search field, enter `sn_itom_pattern.container_image_scan`.
   c. Set the value to **true**.
2. In the URL bar enter: `<instance>.service-now.com/sn_itom_pattern_container_image_scan_status_list.do` to review the scan status of the images.
3. Select the image record to review the Container Image OS Packages table.

Result

The Scan Container Image pattern finds the image OS packages and creates application records based on the image CI class. The Container Image OS Packages table is located in the image record. The pattern also creates relationships between the configuration items. For more information on the data collected, see Container image discovery.

Configure the SBOM generation

Enable the Software Bill of Materials (SBOM) creation and avoid duplicate SBOM files from being generated by setting the system properties.

Before you begin

Verify the following setup:

- ServiceNow AI Platform, at least the Utah release
- Discovery and Service Mapping Patterns is up to date.
- CMDB CI Class Models, is up to date.
- Visibility Content is up to date.
- MID Server requirements: 8GB RAM with 4GB JVM and the Scan container images capability. For more information, see Scan container images

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Select the SBOM format.
   a. In the navigation filter, enter `sys_properties.list`.
   b. On the **System Properties** list, search for the `sn_itom_pattern.sbom_format` and select the record.
   c. Set the **Value** field to your desired format.
   The available formats are:
      - cyclonedx
      - spdx-json
      - spdx

      Note: Setting the value to **None** disables the creation of the SBOM.
2. Set the *sys_attachment* system property to avoid duplicate SBOM files from being generated.
   a. Navigate to **All** > **System Definition** > **Tables**
   b. On the **Tables** form in the **Label** field, enter `*attach`.
   c. Select the **Attachment** record.
   d. Select the **Application Access** tab.
   e. Select the **Can delete** check box.
   f. Select **Update**.

What to do next

Download a software bill of materials (SBOM)

Download a software bill of materials (SBOM)

Download a software Bill of Materials (SBOM) to gain visibility to the components of the container image such as the operating system packages that are installed.

Before you begin

The SBOM must have been configured. For more information, see Configure the SBOM generation.

Role required: discovery_admin

Procedure

1. In the navigation filter, type `sn_itom_pattern_container_image_scan_status_list.do`.
2. Choose the image for which you need the SBOM in the Container Image Scan Status list.
   If the scan status is **None** rather than **Scanned**, wait for the status to change before selecting the image record.
3. Download the SBOM by selecting **download** in the **Manage attachments** section.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Container image discovery

The Discovery and Service Mapping Patterns application uses the Scan Container Image pattern to discover Docker images and OS packages data. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

For information about performing a container image scan, see Scan container images.

The Scan Container Image pattern supports Aqua Trivy starting with version 0.44.0. The last version validated is 0.68.2.

Starting with Discovery and Service Mapping Patterns version 1.18.0, the Scan Container Image pattern supports scanning container images in the following repositories:

- Public repositories
- Self-hosted private repositories
- Amazon Elastic Container Registry (Amazon ECR), both public and private repositories

Starting with version 1.27.0, Discovery and Service Mapping Patterns enables you to control whether to link software packages to containers or only to images. For more information, see Link software package information to images only.

Container image data model

The following CI classes are part of the data model used by this pattern.

CI class hierarchy

CI class	Extends from
Docker Container [cmdb_ci_docker_container]	Operating-system-level Virtualization Container [cmdb_ci_oslv_container]
Operating-system-level Virtualization Container [cmdb_ci_oslv_container]	Configuration Item [cmdb_ci]
Container Environment Variables [cmdb_container_environment_variables]	Key Value [cmdb_key_value]
Container Enrich Scripts [sn_itom_pattern_container_enrich_scripts]	Application File [sys_metadata]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery for container image scans

Discovery populates the data in the CMDB when running the Scan Container Image pattern.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the container application. The MSSQL application record required name format: ApplicationTableName@containerName Example: MSFT SQL Instance@/sql1

Container Environment Variables [cmdb_container_environment_variables]

Field	Description
Container [container]	Name of the container.
Key [key]	Name of the Container environment variable.
Value [value]	Container environment variable value.

Temporary tables for container image scans

Note: If you're using the 1.0.98 version to collect the data, the enriched scripts are supported only with MSSQL.

Data collected during horizontal discovery

Field	Description
Container image scan Status [sn_itom_pattern_container_image_scan_status]	
Image [image]	Name of the container image.
Message [message]	Errors or issues with the scanning process.
CI Class [ci_class]	The image CI class based on the image command details.
Discovery status [discovery_status]	The discovery status record of the image scan.
Scan Status [scan_status]	The scan status. The available values are: • None - The image isn't scanned yet. • In Progress - The image scan is in progress. **Note:** If one or more images are in progress, the next scan won't trigger. • Scanned - The image was scanned successfully.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected during horizontal discovery (continued)

Field	Description
	• Error - A problem occurred during the image scan. Check the message column for details. • Skipped - The image URL isn't formed properly or isn't reachable.
Container image OS packages [sn_itom_pattern_container_image_os_packages]	
Image [image]	Name of the container image.
Package Name [package_name]	Name of the software package.
Package Version [package_version]	Version of the software package.
Package Maintainer [package_maintainer]	Name of the package maintainer.
Container Enrich Scripts [sn_itom_pattern_container_enrich_scripts]	
Active [active]	Whether the enrich script state is active.
CI Type [ci_type]	The CI type to which the enrich script is applicable.
Order [order]	Execution order of a particular script.
Script [script]	The enrich script name.

CI relationships

The Scan Container Image pattern creates the following relationships and references to support container image discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Operating-system-level Virtualization Container [cmdb_ci_oslv_container]
MSFT SQL Instance [cmdb_ci_db_mssql_instance]	Runs::Runs on	Docker Container [cmdb_ci_docker_container]

CI references

CI	Field	Referenced CI
Container Environment Variables [cmdb_container_environment_variables]	Configuration item [configuration_item]	Docker Container [cmdb_ci_docker_container]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Docker virtualization

Kubernetes discovery using patterns

Disable Docker container CI discovery

Configure Docker discovery to collect image CIs only, instead of both image and container CIs.

Before you begin

- Verify that you have at least version 1.31.0 of Discovery and Service Mapping Patterns.
- Check your entitlements to determine whether you have access to 2026 Packaging SKU.

Role required: discovery_admin (granted automatically with 2026 Packaging SKU)

About this task

In large environments, container instances can generate a high volume of CIs. The *sn_itom_pattern.bring_discovery_container* system property controls whether the Kubernetes, Kubernetes Event, Docker Pattern, and Amazon AWS - ECS patterns discover both Docker container and Docker image CIs, or only Docker image CIs. By default, the property is set to true and both Docker Container [cmdb_ci_docker_container] and Docker Image [cmdb_ci_docker_image] are populated. When set to false, only Docker Image [cmdb_ci_docker_image] is populated.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `sn_itom_pattern.bring_discovery_container`.
3. Select the *sn_itom_pattern.bring_discovery_container* system property.
4. In the **Value** field, enter `false`.
   To discover both Docker container and Docker image CIs, set the value to `true`.
5. Select **Update**.

What to do next

Run discovery again to apply the changes.

Kubernetes Visibility Agent

Kubernetes Visibility Agent detects changes on resources in a Kubernetes cluster. It performs continuous discovery, reports any changes back to your instance, and updates the Configuration Management Database (CMDB) with the latest data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore	Configure	Reference
High-level overview of Kubernetes Visibility Agent.	Kubernetes Visibility Agent installation and configuration tasks.	Reference information, such as collected data, the Kubernetes Visibility Agent support matrix, and advanced configuration options.

Helpful resources

Some ServiceNow resources that can provide you with helpful information are:

ServiceNow Store

ServiceNow Store

Support

Exploring Kubernetes Visibility Agent

Kubernetes Visibility Agent enables you to gain visibility into on-premises Kubernetes clusters and the various Cloud deployments.

Kubernetes Visibility Agent detects changes on resources in a Kubernetes cluster. It performs continuous discovery, reports any changes back to your instance, and updates the Configuration Management Database (CMDB) with the latest data. For the latest information on supported cloud deployments, see the Kubernetes Visibility Agent (formerly CNO for Visibility) Support Matrix [KB1700730] article in the Now Support Knowledge Base.

How it works

When you deploy Kubernetes Visibility Agent, Kubernetes creates a Deployment resource in the cluster with the latest data. This resource uses a secret stored in Kubernetes to connect to your ServiceNow instance.

The Kubernetes Visibility Agent Deployment resource contains a pod called Informer, which connects to the Kubernetes API server and receives events on the resources in the cluster from it. The Informer sends the collected data to the instance through the External Communication Channel (ECC) Queue table, using the ServiceNow Table API to read from and write to the queue. The backend part of Kubernetes Visibility Agent (KVA) then updates the appropriate tables in the CMDB.

The main informer image and the daemonset image used for traffic discovery are FIPS-140 compliant for use in regulated markets. For more information, see Kubernetes Visibility Agent support matrix.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If the Informer is unable to report the changes, for example due to a network problem, the resources that were added to the cluster during the event are added to the CMDB after the next full discovery cycle. The resources that were removed from the cluster during the event are marked as Absent and deleted after two full discovery cycles.

For more information about the Kubernetes resources on which the Informer collects data and the CMDB tables it populates, see Data collected by Kubernetes Visibility Agent.

Initial and periodic discovery

In its initial discovery, the Informer finds all the resources in the Kubernetes cluster and reports them to your instance. Every 30 seconds, the Informer sends up to 1 MB of data to the instance. It typically takes up to two minutes to report data on a cluster containing 1,000 pods and another minute for every additional 1,000 pods. A single Informer pod can handle a cluster with tens of thousands of pods. If the Informer exits for any reason, Kubernetes restarts it automatically.

After the initial discovery, the Informer continuously monitors the addition, updating, and deletion of resources in the cluster. Resources that were deleted from the cluster are marked with install_status=Absent and deleted from the CMDB within hours in a regular cleanup.

During each full discovery cycle, the system can optionally check for inactive clusters and automatically update their status. When the auto-retirement feature is enabled, clusters that meet specific conditions are retired. A cluster is considered inactive when its associated Informer is in the Down state and the cluster CI has not been touched for a configured period (default: 60 days). The system determines this based on the *sys_updated_on* field, which indicates when the cluster was last observed by ServiceNow tools, regardless of whether any field values changed. When these conditions are met, the system updates the *install_status* field on the Kubernetes cluster CI and all associated CIs (pods, namespaces, deployments, containers, and other resources). For information about configuring auto-retirement, see Enable automatic retirement for inactive Kubernetes cluster CIs. For details about the configuration properties, see Kubernetes cluster auto-retirement properties.

Impact of the Informer on the Kubernetes API server

The Informer has minimal impact on the Kubernetes API server. It fetches the complete list of relevant resources only once and saves it to memory. From then on, it synchronizes with the Kubernetes API server and never pulls the complete list again. During the periodic and on-demand full discovery cycles, the Informer resends the saved list of resources to the instance.

Kubernetes Visibility Agent performance and scalability benchmark

For Kubernetes Visibility Agent benchmarks, see the Performance results for Kubernetes Visibility Agent [KB1555851] article in the Now Support Knowledge Base.

Configuring Kubernetes Visibility Agent

Set up Kubernetes Visibility Agent by performing the described procedures.

For information about advanced configuration options, see the Administrating and Troubleshooting Kubernetes Visibility Agent (formerly CNO for Visibility) [KB1532614] article in the Now Support Knowledge Base.

Prepare for Kubernetes Visibility Agent deployment

Prepare your ServiceNow instance and the Kubernetes cluster for Kubernetes Visibility Agent (KVA) deployment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that the following applications are installed and activated:
  - Discovery and Service Mapping Patterns
  - Kubernetes Visibility Agent (KVA) application
- Verify that the Kubernetes command line tool kubectl is installed and configured to communicate with your cluster. For more information, see the kubectl documentation.
- If you plan to use a Helm chart for the installation, verify that the Helm tool is available. For more information, see the Helm documentation.
- If you plan to deploy KVA in the OKE (Oracle Kubernetes Engine) environment, you need to create a Kubernetes Cluster on Oracle Cloud Infrastructure. For more information, see the Oracle documentation.

Role required: admin for steps performed on the ServiceNow instance.

Procedure

1. On the ServiceNow instance, navigate to **All** > **User Administration** > **Users**.
2. Choose or create a user with at least the mid_server role.
   For information about adding a new user to your instance, see Create a user. For information about granting a role to a user, see Assign a role to a user.
3. In the Kubernetes cluster, choose or create a namespace in which you want to deploy KVA.
   To create a namespace:
   a. Open the kubectl command line tool.
   b. Run the following command, after replacing NAMESPACE with the relevant value:

   ```
   kubectl create namespace NAMESPACE
   ```

4. Create a Kubernetes secret that contains the ServiceNow credentials for accessing your ServiceNow instance.

   Note: Ensure that the credentials correspond to the user you created or identified in the earlier step.

   a. Open the kubectl command line tool.
   b. Run the following command, after replacing INSTANCE_NAME, USERNAME, PASSWORD and NAMESPACE with the relevant values:

   ```
   kubectl create secret generic k8s-informer-cred-INSTANCE_NAME
    --from-literal=.user=USERNAME
    --from-literal=.password=PASSWORD -n NAMESPACE
   ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- If the KVA Informer connects to the ServiceNow instance through a proxy server, you can provide the proxy user and password by adding the .proxyUser and .proxyPassword keys to the secret. For example:

```
kubectl create secret generic
 k8s-informer-cred-INSTANCE_NAME
 --from-literal=.user=USERNAME
 --from-literal=.password=PASSWORD
 --from-literal=.proxyUser=PROXY_USER
 --from-literal=.proxyPassword=PROXY_PASSWORD -n
 NAMESPACE.
```

  If you don't provide .proxyUser, KVA assumes that authentication is not needed. If you provide .proxyUser but not .proxyPassword, the system returns an error in the logs and aborts the KVA Informer.

- If your organization uses Amazon Elastic Kubernetes Service (EKS), you can store the secret in AWS Secrets Manager. The KVA Informer then pulls the ServiceNow credentials for accessing your instance from the AWS Secrets Manager. For more information, see the Kubernetes Visibility Agent (formerly CNO for Visibility): Storing the instance credentials in AWS Secrets Manager [KB1581074] article in the Now Support Knowledge Base.
- If your organization uses Azure Kubernetes Service (AKS), you can store the secret in the Microsoft Azure Key Vault. The KVA Informer then pulls the ServiceNow credentials for accessing your instance from the Azure Key Vault. For more information, see the Storing the instance credentials in Microsoft Azure Key Vault [KB1647736] article in the Now Support Knowledge Base.
- If your organization uses Google Kubernetes Engine (GKE), you can store the secret in Google Cloud Secret Manager. The KVA Informer then pulls the ServiceNow credentials for accessing your instance from the Google Cloud Secret Manager. For more information, see the Storing credentials in Google Cloud Secret Manager [KB1709597] article in the Now Support Knowledge Base.
- If your organization uses a custom root CA (Certificate Authority), you can mount a custom CA into the KVA Informer pod. This enables the Informer to communicate with the ServiceNow instance. For more information, see the Using a custom root certificate authority when connecting the Informer to the instance [KB1710906] article in the Now Support Knowledge Base.
- As an alternative to basic authentication, the KVA Informer can use OAuth 2.0 authorization to connect to the ServiceNow instance for enhanced security. When the Informer uses OAuth 2.0, the instance issues an access token with an expiry time rather than requiring login credentials with each resource request. For more information, see the Connecting Kubernetes Visibility Agent (formerly CNO for Visibility) using the OAuth Protocol [KB1648198] article in the Now Support Knowledge Base.

5. **Optional:** If your organization's policy requires it, place the KVA Informer docker image in your organization's image repository.

Note: Starting in KVA version 3.9.0 (Informer version 2.3.0), released in November 2024, the docker image supports both arm64 and amd64 architectures. Upgrading from the previous image to the new one will not cause any disruptions. However, the new image requires more storage space in your image repository than the previous one did.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Pull the image from Docker Hub and push it to your organization's repository.

(Optional) Replace `VERSION` with the number of the most recent release, provided in the Kubernetes Visibility Agent (formerly CNO for Visibility) Helm Chart and Kubernetes YAML file releases [KB1564347] article in the Now Support Knowledge Base.

```
Docker pull docker.io/servicenowdocker/informer:VERSION
 Docker tag docker.io/servicenowdocker/informer:VERSION
 COMPANY_REPO:VERSION
 Docker push COMPANY_REPO:VERSION
```

b. If the image repository requires authentication, create a secret named k8s-informer-repo-cred in the designated namespace.

For example:

```
kubectl create secret docker-registry k8s-informer-repo-cred
 --docker-server https://index.docker.io/v2/
 --docker-username DOCKER_USERNAME --docker-password
 DOCKER_TOKEN --docker-email=user@servicenow.com -n NAMESPACE
```

6. **Optional:** If outgoing traffic from your cluster flows through a proxy, identify the proxy host name and port used in your cluster.

Note: Ask your organization's Kubernetes team for this information. You need to provide it in the installation process.

7. **Optional:** Enable KVA to process more data simultaneously.
By default, the Informer's back-end side can handle data from around 300 clusters at the same time. If you have more clusters or your clusters undergo frequent changes, you can increase the number of concurrent worker threads used for processing the data. For more information, see the Controlling the number of "Kubernetes Visibility Agent" (formerly CNO for Visibility) concurrent worker threads [KB1580241] article in the Now Support Knowledge Base.

Result
Your ServiceNow instance and the Kubernetes cluster are ready for KVA deployment.

What to do next
Install Kubernetes Visibility Agent (KVA) Informer.

Install Kubernetes Visibility Agent (KVA) Informer

You can install Kubernetes Visibility Agent Informer using either a Helm chart or the Kubernetes YAML file.

Before you begin
Role required: none

About this task
This procedure describes the standard Kubernetes Visibility Agent Informer installation process. Typically, you would use a Helm chart to perform the installation, but you can also configure and run the Kubernetes YAML file provided by ServiceNow.

Kubernetes Visibility Agent supports Helm 4. Existing releases created with Helm 3 can be upgraded to Helm 4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

Install Kubernetes Visibility Agent Informer using either a Helm chart or the Kubernetes YAML file provided by ServiceNow.

- Install Kubernetes Visibility Agent Informer using a Helm chart.

  a. Edit the Helm install command.

  i. In the following command, replace NAMESPACE, INSTANCE_NAME, and CLUSTER_NAME with the relevant values and replace the HELM_URL with the URL of the most recent version, provided in the Kubernetes Visibility Agent (formerly CNO for Visibility) Helm Chart and Kubernetes YAML file releases [KB1564347] article in the Now Support Knowledge Base.

```
helm install -n NAMESPACE --set acceptEula=Y
 --set instance.name=INSTANCE_NAME
 --set clusterName="CLUSTER_NAME" {--set
 image.repository=COMPANY_REPO --set image.tag=COMPANY_TAG}
 {--set proxyUrl=http://proxy_host:proxy_port} k8s-informer
 HELM_URL
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- The curly brackets indicate that the code segments they surround are optional. Use these code segments if the image must be pulled from your organization's image repository or if outgoing traffic from your cluster flows through a proxy.
- By passing `--set acceptEula=Y`, you agree to the End-User License Agreement. You can read the End-User License Agreement by extracting the EULA.pdf from informer-helm-<VERSION_NUMBER>.tgz.
- If the Informer is installed in OpenShift, add the *openShift* parameter to the command and set it to `true`:

```
helm install -n NAMESPACE --set acceptEula=Y
 --set instance.name=INSTANCE_NAME
 --set clusterName="CLUSTER_NAME" {--set
 image.repository=COMPANY_REPO --set
 image.tag=COMPANY_TAG} {--set
 proxyUrl=http://proxy_host:proxy_port} {--set
 openShift=true} k8s-informer HELM_URL
```

  Note: If the Informer is not installed in OpenShift, there is no need to set this parameter.

- If you use the Hybrid service maps to map resources both inside and outside the Kubernetes cluster, set the following parameters:

```
--set bringIngressAddresses=true
 --set bringNodeExternalIPs=true --set
 discoverK8sServiceExternalIps=true --set
 openShift=true
```

- Upgrading Existing Kubernetes Visibility Agent Informer from Helm 3 to Helm 4: When Helm 4 is installed on your system and you want to migrate an existing Kubernetes Visibility Agent Informer from Helm 3 to Helm 4, use the Server-Side Apply flag:

```
helm upgrade k8s-informer <CHART_URL>
 --server-side=true [other flags] -n $NAMESPACE
```

  The `--server-side=true` flag explicitly migrates the Informer to Helm 4's Server-Side Apply model for improved conflict resolution.

ii. In the command, configure additional settings as needed:

- Configure informer pod memory limit and memory request
- Change the full discovery frequency in Kubernetes Visibility Agent
- Deactivate continuous discovery in Kubernetes Visibility Agent
- Suppress frequent changes in the Kubernetes cluster to reduce the load on your instance. When this feature is active, the Informer does not report changes before the configured threshold rate back to the instance. For more information, see the Suppressing frequent changes in Kubernetes clusters reported by Kubernetes Visibility Agent (formerly CNO for visibility) [KB1580745] article in the Now Support Knowledge Base.
- Control the maximum number of full discoveries that informers can run on multiple clusters at the same time. For more information, see the Scheduling Full Discoveries in

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes Visibility Agent (formerly CNO for Visibility) [KB1580250] article in the Now Support Knowledge Base.

- Enabling application service maps
- Enable container image scanning with Syft scanner by adding the *runSyftScanner* parameter and setting it to `true`:

```
--set runSyftScanner=true
```

Note: Container image scanning with Syft scanner is only supported when installing the Informer using the Helm chart method. For more information, see Enabling container image scanning with Syft scanner.

b. Run the command.

- Install Kubernetes Visibility Agent Informer using the Kubernetes YAML file.

a. Download the Kubernetes YAML ZIP file provided in the Kubernetes Visibility Agent (formerly CNO for Visibility) Helm Chart and Kubernetes YAML file releases [KB1564347] article in the Now Support Knowledge Base.

b. Extract the k8s_informer.yaml and EULA.pdf files from the ZIP file.

c. Edit the k8s_informer.yaml file.

- Change the value of ACCEPT_EULA to `"Y"`, as follows:

```
- name: ACCEPT_EULA
  value: "Y"
```

Note: By changing the value to `"Y"`, you agree to the End-User License Agreement included in the EULA.pdf file.

- Replace all occurrences of <NAMESPACE> with the namespace in which you want to install the Informer.
- Replace all occurrences of <INSTANCE_NAME> with the name of your instance, without the domain name.
- Replace <CLUSTER_NAME> with the name of your cluster as it appears in the CMDB.
- If the Informer is installed in OpenShift, change the value of the OPENSHIFT environment variable to `true`.
- Enabling application service maps

d. If you use the Hybrid service maps to map resources both inside and outside the Kubernetes cluster, set the following parameters:

```
- name: BRING_INGRESS_ADDRESSES value: "true" -
 name: BRING_NODE_EXTERNAL_IPS value: "true" - name:
 DISCOVER_K8S_SERVICE_EXTERNAL_IP value: "true" - name:
 OPENSHIFT value: "true"
```

e. Run the command:

```
kubectl apply -f k8s_informer.yaml
```

Result

Kubernetes Visibility Agent is deployed on the Kubernetes cluster and begins to stream data into your instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

For information about Kubernetes Visibility Agent administration and troubleshooting, see the Administrating and Troubleshooting Kubernetes Visibility Agent (formerly CNO for Visibility) [KB1532614] article in the Now Support Knowledge Base.

Configure informer pod memory limit and memory request

Set the memory limit and memory request of the Kubernetes Visibility Agent Informer pod.

Before you begin

Role required: none

About this task

Informer pod memory consumption depends mainly on the number of resources contained in the cluster. Memory limit defines the maximum memory the informer container is allowed to use, while memory request defines how much memory to reserve when scheduling the informer pod. Set the Informer pod's memory request and the Informer pod's memory limit to at least (number of pods)/8MB, with a default minimum of 200MB. Memory request must not exceed the memory limit. For example:

Informer pod memory requirements

Estimated pod count	1000	5000	30000
Minimum memory limit	200Mi	625Mi	3.75Gi
Memory request	200Mi	625Mi	3.75Gi

Procedure

Do one of the following:

- When using a Helm chart, in the Helm install command, add the following command line arguments:

```
--set memoryLimit=MEM_LIMIT\
--set memoryRequest=MEM_REQUEST
```

  For example:

```
--set memoryLimit=625Mi\
--set memoryRequest=625Mi
```

- When using the k8s_informer.yaml, under the line containing "limits:" add the line `memory: MEM_LIMIT`.

  In the following example, MEM_LIMIT and MEM_REQUEST are 625Mi.

```
  limits:
    cpu: 100m
    memory: 625Mi
  resources:
    cpu: 100m
    memory: 625Mi
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

Setting `memoryRequest` equal to `memoryLimit` reserves exactly the memory your ACC agent needs. It prevents the ACC agent from being shut down if the system runs low on memory. This guarantees stable, uninterrupted agent operation. If you leave memoryRequest at the default 200Mi while setting a higher memoryLimit (such as 625Mi), the pod is burstable. It schedules against 200Mi but can grow to 625Mi.

Change the full discovery frequency in Kubernetes Visibility Agent

Customize how often you want the Kubernetes Visibility Agent Informer to run a full discovery.

Before you begin

Role required: none

About this task

By default, the Informer runs a full discovery every 24 hours (1440 minutes). You can change this frequency by performing the following procedure. For example, change the value to 2880 to make Kubernetes Visibility Agent run a full discovery every 48 hours (2880 minutes).

Note: When the auto-retirement feature is enabled, the system checks for inactive clusters during each full discovery cycle. Clusters with Informers in the Down state that have not been updated for the configured period are automatically retired. For more information, see Enable automatic retirement for inactive Kubernetes cluster CIs.

Procedure

Do one of the following:

- When using a Helm chart: In the Helm install command, add the following command line argument:

  ```
  --set fullDiscoveryMin=<Frequency in minutes>
  ```

- When using the k8s_informer.yaml: Replace the default value with the required one in the line under FULL_DISCOVERY_MIN.

Deactivate continuous discovery in Kubernetes Visibility Agent

Switch off continuous discovery by Kubernetes Visibility Agent if all you need is periodic snapshots of your cluster resources. If you have multiple clusters with frequent changes, deactivating continuous discovery reduces the load on your instance.

Before you begin

Role required: none

Procedure

Do one of the following:

- When using a Helm chart, in the Helm install command, add the following command line argument: `--set continuousDiscovery=false`
- When using the k8s_informer.yaml, enter the value “false” in the line under CONTINUOUS_DISCOVERY.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Display the Kubernetes cluster version in the CMDB

Make the Kubernetes Visibility Agent Informer populate the relevant field in the cmdb_ci_kubernetes_cluster CI to display the Kubernetes cluster version.

Before you begin

Use ServiceNow Discovery and Service Mapping Patterns application version 1.11.0 or higher.

Role required: none

Procedure

Do one of the following:

- When using a Helm chart, in the Helm install command, set the *getClusterVersion* parameter to true:

  ```
  --set getClusterVersion=true
  ```

- When using the k8s_informer.yaml file, set the environment variable GET_CLUSTER_VERSION to true.

Create a cmdb_ci_linux_server CI for each Kubernetes node

Configure if you want the Kubernetes Visibility Agent Informer to create a cmdb_ci_linux_server CI for each Kubernetes node.

Before you begin

Role required: none

About this task

By default, the Informer creates a cmdb_ci_linux_server CI for every Kubernetes node. If this CI is redundant or interferes with other flows in your organization, you can set the associated configuration parameter to false.

Procedure

Do one of the following:

- When using a Helm chart, in the Helm install command, set the *createServerCi* parameter to false:

  ```
  --set createServerCi=false
  ```

- When using the k8s_informer.yaml file, set the environment variable CREATE_SERVER_CI to false.

Enable automatic retirement for inactive Kubernetes cluster CIs

Enable automatic retirement to update the status of inactive Kubernetes cluster configuration items (CIs) and all associated resources during full discovery cycles, so your CMDB reflects only active infrastructure.

Before you begin

Confirm you have the following setup:

- Kubernetes Visibility Agent (KVA) Plugin application release 3.15.0
- Informer version 2.6.0

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin

About this task

Enable automatic retirement to update the status of inactive Kubernetes cluster CIs and all associated resources (pods, namespaces, deployments, containers, and other resources) during full discovery cycles. This feature is disabled by default.

Procedure

1. Enable automatic retirement for inactive cluster CIs by performing the following substeps.

   a. Navigate to **All > System Properties > All Properties**.

   b. In the **Name** field, enter `sn_acc_visibility.auto_retire_k8s_clusters`.
   - If the property is available, double-click the **Value** field to perform inline editing, enter `true`.
   - If the property doesn't exist on your instance, select **New** and create the `sn_acc_visibility.auto_retire_k8s_clusters` property with Type set to True/False, then double-click the **Value** field to perform inline editing, and enter `true`.

   Auto-retirement is enabled. By default, clusters are retired after 60 days of inactivity with the *`install_status`* set to Retired (value 7).

   c. **Optional:** Customize the inactivity period by setting the *`sn_acc_visibility.auto_retire_period_days`* property to the number of days after which inactive clusters should be retired.
   The default value is 60 days.

   d. **Optional:** Customize the install status value by setting the *`sn_acc_visibility.auto_retire_install_status`* property to the numerical value of the desired install status.

   (Optional) The default value is 7 (Retired). Other possible values include 100 (Absent).

   > ⚠ **Warning:** If you set the value to 100 (Absent), the cluster CI and all associated CIs are deleted by a table cleanup job. This action is irreversible.

   e. **Optional:** Configure the install status for Linux server CIs associated with Kubernetes nodes by setting the *`sn_acc_visibility.node_deletion_set_server_install_status`* property to the numerical value of the desired install status.
   If this property does not exist or the value is empty, no action is taken on Linux server CIs when nodes are retired.

2. To manage the lifecycle of retired CIs, see https://www.servicenow.com/docs/r/servicenow-platform/configuration-management-database-cmdb/cmdb-data-management.html.

Define include and exclude lists of Labels and Annotations

In Kubernetes Visibility Agent, define include and exclude lists of Labels and Annotations in Kubernetes resources that the Informer pulls into the Configuration Management Database (CMDB).

Before you begin

Role required: none

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

By default, that is if *includeLabelsAndAnnotations* and *excludeLabelsAndAnnotations* aren't set, the KVA discovers and includes all the available Kubernetes labels and annotations on the supported Kubernetes resource.

Empty label and annotation values and values longer than 255 characters aren't sent to the CMDB. Duplicate key and value pairs for the same CI are also ignored.

- Use *includeLabelsAndAnnotations* when you want to restrict discovery to a specific list of label or annotation keys.
- Use *excludeLabelsAndAnnotations* when you want to prevent specific label or annotation keys from being sent to the CMDB.

If both options are configured, *excludeLabelsAndAnnotations* takes precedence. If the same key appears in both lists, the key is excluded and is not sent to the CMDB.

Procedure

Create an include or exclude list by using the following procedure.

Task	Procedure
Create an include list	• When using a Helm chart, in the Helm install command, add a comma-separated include list of Labels and Annotations. For example: `--set IncludeLabelsAndAnnotations="label1,label2` • When using the k8s_informer.yaml file, add values under the environment variable IN CLUDE_LABELS_AND_ANNOTATIONS.
Create an exclude list	• When using a Helm chart, in the Helm install command, add a comma-separated exclude list of Labels and Annotations. For example: `--set ExcludeLabelsAndAnnotations="label1,label2` • When using the k8s_informer.yaml file, add values under the environment variable EX CLUDE_LABELS_AND_ANNOTATIONS.

Add custom Labels and Annotations to Kubernetes resources

Add custom Labels and Annotations to all resources deployed by Kubernetes Visibility Agent in the Kubernetes cluster.

Before you begin

Role required: none

Procedure

Do one of the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- When using a Helm chart, in the Helm install command, add the custom attributes to the values.yaml file.

  Use the format shown in this example:

  ```
  --set commonLabels.mylabel1=value1 --
  set commonLabels.mylabel2=value2 --set
  commonAnnotations.myannotation1=value3.
  ```

- When using the k8s_informer.yaml file, add custom Labels and Annotations manually to the relevant resources in the file.

Note:

Label and Annotation names must be alphanumeric and contain no spaces or special characters. For limitations on Label values, see the Kubernetes documentation.

Upgrade Kubernetes Visibility Agent Informers remotely

Upgrade Kubernetes Visibility Agent Informer pods in Kubernetes clusters remotely from the ServiceNow Instance to avoid dependence on your Kubernetes admin. You can upgrade a single Informer or multiple Informers together.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Kubernetes Visibility Agent** > **Properties**.
2. In the **Kubernetes Visibility Agent Image** field, set the image to which you want to upgrade Informers that are currently running, and then select **Save**.
   Use the following format: repository:tag. For example: `servicenowdocker/informer:2.1.1`.
3. Navigate to **All** > **Kubernetes Visibility Agent** > **Home**.
4. In the Kubernetes Informers list, perform the relevant steps to upgrade one Informer or several Informers concurrently.

Note: Only Informers with the value Upgrade Pending in the **Upgrade Status** field can be upgraded. The running status of the Informers must be either Up or Paused.

Upgrade	Steps
One Informer	**a.** Select the Informer in the list to display its form. **b.** In the Related Links section of the form, se lect **Upgrade Informer**.
Multiple Informers	**a.** Select multiple Informers in the list. **b.** From the **Actions on selected rows...** list menu, select the **Upgrade Informers** menu item.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The upgrade process begins, and the **Upgrade Status** field for each selected Informer displays the value Upgrading. The system downloads the new image and starts the new pod for each selected Informer.

When an upgraded pod connects to the Instance, the **Informer Version** field displays the new version number. The **Upgrade Status** field displays Desired image in use.

Override Informer parameters from the Instance

Control Kubernetes Visibility Agent Informer execution parameters from the ServiceNow Instance to avoid dependence on your Kubernetes admin.

Before you begin

Role required: discovery_admin

About this task

The Informer execution parameters are configured during deployment using the Kubernetes Yaml file or the Helm command. You can change these settings later directly from the Instance. This avoids having to ask your Kubernetes team to update or reinstall the deployment.

When you add or modify a parameter, it is synchronized with the informer through the ECC Queue. Depending on the parameter, the Informer pod may need to restart. Parameters requiring a restart must have the SELF_PATCHING_ALLOWED environment variable set to true for changes to be effective upon restart. If this variable is set to false, the Informer is not restarted and continues to use the old values.

Note: Informer-specific parameters only affect a specific Informer. If you add a parameter without assigning it to an Informer, it defaults to a global parameter and affects all Informers in the system. When both global and Informer-specific parameters are present, Informer-specific parameters take precedence.

You can update the following parameters from the Instance:

Parameter	Restart upon change
VERBOSE_LOGGING	No
OPENSHIFT	Yes
CREATE_SERVER_CI	No
INCLUDE_LABELS_AND_ANNOTATIONS	No
EXCLUDE_LABELS_AND_ANNOTATIONS	No
FULL_DISCOVERY_MIN	Yes
SEND_INTERVAL_SEC	Yes
CONTINUOUS_DISCOVERY	Yes
CLUSTER_RESOURCE_ID	No
CLUSTER_NAME	No

The *CREATE_CONTAINERS* parameter determines which Docker containers are deployed to the instance. It is managed automatically and is not exposed in the Configuration Parameters list. A new system property, *sn_itom_pattern.bring_discovery_container*, which is available when the Discovery and Service Mapping Patterns plugin is installed, governs *CREATE_CONTAINERS* automatically. When *sn_itom_pattern.bring_discovery_container* is set to `true` (default),

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CREATE_CONTAINERS is set to `non-init-containers`. When set to `false`, *CREATE_CONTAINERS* is set to `none` and no Docker containers are brought to the instance. Each change triggers the ECC Queue sync to propagate the updated value with the Informer.

For information about the purpose of each parameter, see the Kubernetes Visibility Agent (formerly CNO for Visibility) Advanced Configuration Options [KB1648891] article in the Now Support Knowledge Base.

Procedure

1. Navigate to **All** > **Kubernetes Visibility Agent** > **Configuration Parameters**.
2. Select a parameter in the **Kubernetes Informer Parameters** list.
3. Change the configured value as needed, and then select **Update**.

Enabling container image scanning with Syft scanner

Container image scanning with Syft scanner provides visibility into software packages within container images, helping you identify installed software for compliance, licensing, and security purposes without requiring a MID Server.

KVA integrates with the Syft open-source tool to scan container images and discover software packages. This feature provides an alternative to the existing Aqua Trivy integration that requires a MID Server. For customers already using KVA, Syft scanner offers a simpler deployment without additional infrastructure or manual tool installation.

Note: For more information about this feature, see the Software decomposition of container images using Kubernetes Visibility Agent and Syft [KB3149735] article in the Now Support Knowledge Base.

Benefits of Syft scanner integration

The Syft scanner integration provides the following advantages:

- No MID Server required
- No manual installation of scanning tools
- Automatic integration with KVA Informer
- Scans container images directly from the Kubernetes cluster
- Populates the same data model as the Aqua Trivy integration

Installation method

The Syft scanner is only supported when installing KVA Informer using the Helm chart method. YAML-based installation does not support the Syft scanner feature.

To enable Syft scanner during Helm installation, add the following parameter to the Helm install command:

```
--set runSyftScanner=true
```

For complete installation instructions, see Install Kubernetes Visibility Agent (KVA) Informer.

Enabling container image scanning

After installing the Informer with Syft scanner support, enable container image scanning by setting the following system properties to `true`:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- *sn_itom_pattern.container_image_scan* - Enables the container image scanning mechanism
- *sn_acc_visibility.image_scan_enabled* - Directs the system to use the Syft scanner in KVA instead of Aqua Trivy on a MID Server

Note: You cannot use both scanning methods simultaneously. When *sn_acc_visibility.image_scan_enabled* is set to `true`, the system uses only the Syft scanner method.

Informer status fields

When an Informer is configured with Syft scanner support, the following fields appear on the Informer form:

- **Image Scanner Enabled** - Indicates whether the Informer supports container image scanning (true or false)
- **Image Scan Registries** - Comma-separated list of image registries the Informer can access through credential-less trust relationships (ecr, acr, google)

A new UI action **Get Syft Scanner Logs** is available on Informer records that have image scanning enabled. Selecting this action retrieves the Syft scanner logs and attaches them to the Informer record as `k8s_informer_syft.log`.

Data collected by container image scans

When the Syft scanner successfully scans a container image, the system populates the following fields on the Docker Image CI [cmdb_ci_docker_image]:

Docker Image CI fields populated by Syft scanner

Field	Description
Operating System [os]	Operating system name
OS Version [os_version]	Operating system version
Architecture [architecture]	Processor architecture. Examples: amd64, arm64
Size (bytes) [size_bytes]	Image size in bytes
OS Family [os_family]	Operating system family. Examples: rhel, centos, alpine, azurelinux, windows
Image created [image_created_at]	Timestamp when the image was created
Command [command]	Default command executed when the container starts

Note: Not all fields are populated for every image. Field availability depends on the image metadata and scan results.

Software packages

The Syft scanner identifies software packages contained in the container image and creates records in the Container Image OS Packages [sn_itom_pattern_container_image_os_packages] table. This table appears as a related list on the Docker Image CI form.

Each software package record includes:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Package Name
- Package Version
- Package Maintainer

Note: Package maintainer and version information depends on the data available in the image. The Syft scanner reports what is present in the image metadata.

Monitoring scan status

The Container Image Scan Status [sn_itom_pattern_container_image_scan_status] table maintains the scan status for each container image. Access this table by navigating to the URL: `<instance>.service-now.com/sn_itom_pattern_container_image_scan_status_list.do`

Possible scan status values:

- None - Image has not been scanned
- In Progress - Scan is currently running
- Scanned - Scan completed successfully
- Error - Scan failed (check the Message field for details)
- Skipped - Image URL is malformed or unreachable

Related topics

Container image scanning for software decomposition

Enabling application service maps

Application service maps give you visibility into how workloads communicate in a Kubernetes cluster, helping you detect dependencies and monitor traffic flows in real time.

KVA collects traffic data between workloads and reports it to the ServiceNow instance as part of its installation. Three methods are available: a ServiceNow DaemonSet, Istio or Linkerd service meshes, or Cilium. Each method is suited to different cluster configurations. For deployment prerequisites and installation steps, see Install Kubernetes Visibility Agent (KVA) Informer.

- If you use a ServiceNow DaemonSet, a pod runs on each Kubernetes node and reports new connections it detects to the main Informer pod every 60 seconds. To find out how to enable the application service maps using a ServiceNow DaemonSet, see Enable service maps using DaemonSet.
- If you use service meshes and have a Prometheus server, provide the Prometheus URL. The Informer then queries the Prometheus server periodically for connection data reported by the service mesh. To find out how to enable the application service maps using Istio or Linkerd service meshes, see Enable service maps using service meshes.
- If Cilium is already running in the cluster, the Informer connects to the Cilium agent through its Hubble layer to pull traffic data. A DaemonSet pod is deployed with minimal permissions to collect and report connection data. This method requires fewer permissions than the ServiceNow DaemonSet method. To find out how to enable the application service maps using Cilium, see Enable service maps using Cilium.

Enable service maps using DaemonSet

Enable application service maps based on the traffic between the workloads in Kubernetes by using a ServiceNow DaemonSet as part of Kubernetes Visibility Agent (KVA) installation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: discovery_admin.

About this task
If you use a ServiceNow DaemonSet, a pod is running on each Kubernetes node and reports to the main Informer pod every 60 seconds on new connections detected. For more information, see Install Kubernetes Visibility Agent (KVA) Informer

Procedure

Set parameters either during Helm chart archive installation or by modifying the Kubernetes YAML file provided by ServiceNow.

Method	Actions

During Helm chart archive installation	**1.** In the Helm `install` command, add the fol lowing parameter:
	`--set connectionsDiscovery.method=servicenow`
	2. If your organization's policy requires it, pull the DaemonSet image from Docker hub and place it in your organization's image reposito ry.
	Replace the placeholder X characters with the number of the most recent release pro vided in the Kubernetes Visibility Agent (for merly CNO for Visibility) Helm Chart and Ku article in the Now Support Knowledge Base.
	`Docker pull servicenowdocker/informer_ds:X.X.X`
	`Docker tagservicenowdocker/informer_ds:X.X.X COMPANY_REPO:X.X.X`
	3. In the `Helm install` command, add the following parameters:
	`--set image.dsRepository=COMPANY_REPO –set image.dsTag=X.X.X`

Modifying the `k8s_informer.yaml` file	**1.** In the `k8s_informer.yaml` file, uncom ment the DaemonSet, ServiceAccount (used by the DaemonSet), and Service.Deploy ment resources according to the instruc tions in the file. **2.** Uncomment the lines containing environ ment variables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Actions
	3. If your Prometheus server URL isn't the default one, uncomment the lines containing the environment variable *PROMETHEUS_URL* and change the value. 4. Run the following command: `kubectl apply -f k8s_informer.yaml`

What to do next
Create service maps

Enable service maps using service meshes

Enable application service maps based on the traffic between the workloads in Kubernetes by using Istio or Linkerd or service meshes as part of Kubernetes Visibility Agent (KVA) installation.

Before you begin
You must have a Prometheus server and provide the Prometheus URL.

Role required: discovery_admin.

About this task
If you use a service mesh, a pod running on every node (Informer) communicates with a Prometheus server to get the traffic information. For more information, see Install Kubernetes Visibility Agent (KVA) Informer

Procedure

Set parameters either during Helm chart archive installation or by modifying the Kubernetes YAML file provided by ServiceNow.

Method	Actions

During Helm installation	In the `Helm install` command, add the following parameter:
	`--set connectionsDiscovery.method=<METHOD>`
	where <METHOD> is the URL for either Istio (http://prometheus.istio-system:9090) or Linkerd (http://prometheus.linkerd-viz:9090).
	Note:
	If your Prometheus server URL is not the default one, use the following parameter:
	`--set connectionsDiscovery.prometheusUrl`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Actions

Modify the `k8s_informer.yaml` **script**	**1.** Access the `k8s_informer.yaml` script.
	2. Set the environment variable *`k8s_informer.yaml`* by entering either Istio or Linkerd in quotation marks:
	`-- name: GET_CONNECTIONS_METHOD value: ""`
	3. Run the following command:
	`kubectl apply -f k8s_informer.yaml`

What to do next
Create service maps

Enable service maps using Cilium

Enable application service maps based on the traffic between the workloads in Kubernetes by connecting to a Cilium agent already running in the cluster.

Before you begin
Cilium must already be installed and running in the target Kubernetes cluster. ServiceNow does not install Cilium.

This feature is not supported on OpenShift.

Role required: discovery_admin.

About this task
Cilium is open-source software that you can run in your Kubernetes cluster to control traffic and apply network policies. If Cilium is already running in the cluster, KVA can connect to the Cilium agent through its Hubble layer to pull traffic data. A DaemonSet pod is deployed with minimal permissions to collect this data and report it to the main Informer pod, which then sends it to the ServiceNow instance.

This method requires fewer permissions than the ServiceNow DaemonSet method and reduces the effort required to get security approval in environments with strict permission policies. The outcomes are the same as the other service map methods. For more information, see Install Kubernetes Visibility Agent (KVA) Informer.

Note: Only the Helm chart installation method is supported for this feature. The `k8s_informer.yaml` file method is not supported.

Procedure

In the Helm `install` command, add the following parameter:

```
--set connectionsDiscovery.method=cilium
```

What to do next
Create service maps

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create service maps

Create a service map that maps application services based on traffic between the workloads in Kubernetes using Istio or Linkerd service meshes or a ServiceNow DaemonSet.

Before you begin

You should first enable the Service Maps, by using Istio or Linkerd service meshes or a ServiceNow DaemonSet as part of the Kubernetes Visibility Agent (KVA) installation. For more information, see Enabling application service maps.

Role required: discovery_admin

About this task

Service maps created using this procedure map application services within a single Kubernetes cluster.

Procedure

1. Navigate to **All** > **Configuration** > **Kubernetes** > **Services** in the navigation filter.
2. Select **New** and fill in the fields on the form.

Kubernetes Service Record form

Field	Description
Name	Name of the Kubernetes service. If the service does not exist, the service map is created with the default name in the following format: `Name@Namespace@Cluster Name`. You can modify this name.
Namespace	Kubernetes namespace that the service belongs to.
Selector	Label selector used to identify the pods that the service routes traffic to.
Kubernetes UID	Unique identifier assigned to the Kubernetes service by the cluster.

If the service for this CI root already exists, the message `Service <service-name> already exists and has the same root CI.` is displayed.

3. Select **Submit**.
4. Open the newly created Kubernetes service record.
5. Select **Create Service Map**.
6. **Optional:** View the map.
   a. Select **Open in CMDB Workspace**.
   b. In the CMDB workspace, select **Open Map**.

Result

The Service Map is created and is visible in the CMDB workspace.

Create hybrid service maps

Create hybrid service maps that extend outside the Kubernetes cluster and map other related service resources.

About this task

For more information, see Service Mapping for containerized environments using KVA

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

The Kubernetes cluster must be fully discovered using KVA and configured to discover connections between Kubernetes resources. For more information, see Enabling application service maps.

The following must be installed and up-to-date:

- KVA version 3.14.x or later. See Prepare for Kubernetes Visibility Agent deployment.
- KVA Informer version 2.7.x or later, with the additional settings required for hybrid maps. See Install Kubernetes Visibility Agent (KVA) Informer.
- For the latest Helm chart and Kubernetes YAML file release information, see the Kubernetes Visibility Agent (formerly CNO for Visibility) Helm Chart and Kubernetes YAML file releases [KB1564347] article in the Now Support Knowledge Base.

For Service Mapping prerequisites, including MID Server configuration and required applications, see Prerequisites for performing top-down discovery using Service Mapping.

Role required: service_mapping_admin

Procedure

1. Enable the unified map property.
   For the procedure, see Access the Unified Map feature from the Service Mapping Workspace
2. Navigate to **All** and enter `Service mapping` in the navigation filter, in the menu that appears choose **Service Instances**.
3. Select **New**.
4. Choose a resource type, choose a resource name, and enter the resource's /URL to use as an entry point for this map.
5. Select **Create Service Maps**.
6. **Optional:** View the map.
   a. Select **Open in CMDB Workspace**.
   b. In the CMDB workspace, select **Open Map**.

Result

The hybrid service map is created and is visible in the CMDB workspace.

Pulling additional resources from Kubernetes clusters into the CMDB

You can configure the Kubernetes Visibility Agent Informer to pull extra resources from Kubernetes clusters into the Configuration Management Database (CMDB), in addition to the resources it sends to the database by default. By pulling in additional resources beyond the default set, you can enhance visibility into your Kubernetes clusters.

The Informer collects data on Kubernetes resources, sends the collected data to the instance, and updates the appropriate tables in the CMDB. For a list of the resources the Informer pulls into the CMDB by default, see Data collected by Kubernetes Visibility Agent.

For every additional Kubernetes resource you want the Informer to pull into the CMDB, the system creates a Configuration Item (CI) in the cmdb_ci_kubernetes_component target table. Alternatively, extra resources can be pulled to custom tables that extend the cmdb_ci_kubernetes_component table. For example, you could create a u_cmdb_ci_kubernetes_configmap table for all ConfigMap resources in the cluster. The system creates a relation and a reference from the table to the cluster CI for every additional resource. If the resource is namespaced, a relation to the namespace CI is created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, the Informer pulls the following information from the Kubernetes cluster for each additional resource:

- Name
- Namespace (for namespaced resources)
- Kubernetes ID
- API version
- Kind
- Labels and Annotations

Along with this default data, you can make the Informer pull further information on the additional resources into the CMDB. By default, the system stores this extra data in the attributes field of the cmdb_ci_kubernetes_component table. You can also create a custom field per additional resource for storing the extra information.

Kubernetes Visibility Agent enables you to filter additional resources by defining filtering conditions. If a condition shows a match, the Informer sends the resource to the CMDB.

For more detailed information and procedures for pulling additional resources from Kubernetes clusters into the CMDB, see the Kubernetes Visibility Agent (formerly CNO for Visibility) extensibility and customization [KB1638668] article in the Now Support Knowledge Base.

Upgrade Kubernetes Visibility Agent

You can upgrade Kubernetes Visibility Agent using either a Helm chart or the Kubernetes YAML file.

Before you begin

Role required: none

About this task

This procedure describes the standard Kubernetes Visibility Agent upgrade process. Typically, you would use a Helm chart to perform the upgrade. Alternatively, you can upgrade Kubernetes Visibility Agent by configuring and running the Kubernetes YAML file provided by ServiceNow. In both cases, there is no need to uninstall before upgrading.

You can also upgrade Kubernetes Visibility Agent Informer pods in Kubernetes clusters remotely from the ServiceNow Instance. For more information, see Upgrade Kubernetes Visibility Agent Informers remotely.

Procedure

Upgrade Kubernetes Visibility Agent using either a Helm chart or the Kubernetes YAML file provided by ServiceNow.

- Upgrade Kubernetes Visibility Agent using a Helm chart.

  a. Edit the Helm upgrade command.

  i. In the following command, replace NAMESPACE, INSTANCE_NAME, and CLUSTER_NAME with the relevant values and replace the HELM_URL with the URL of the most recent version, provided in the Kubernetes Visibility Agent (formerly CNO for Visibility) Helm Chart and Kubernetes YAML file releases [KB1564347] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
helm upgrade -n NAMESPACE --set acceptEula=Y
 --set instance.name=INSTANCE_NAME
 --set clusterName="CLUSTER_NAME" {--set
 image.repository=COMPANY_REPO --set image.tag=COMPANY_TAG}
 {--set proxyUrl=http://proxy_host:proxy_port} k8s-informer
 HELM_URL
```

Note:

- The curly brackets indicate that the code segments they surround are optional. Use these code segments if the image must be pulled from your organization's image repository or if outgoing traffic from your cluster flows through a proxy.
- By passing `--set acceptEula=Y`, you agree to the End-User License Agreement. You can read the End-User License Agreement by extracting the EULA.pdf from informer-helm-<VERSION_NUMBER>.tgz.
- If the Informer is installed in OpenShift, add the *openShift* parameter to the command and set it to `true`:

```
helm upgrade -n NAMESPACE --set acceptEula=Y
 --set instance.name=INSTANCE_NAME
 --set clusterName="CLUSTER_NAME" {--set
 image.repository=COMPANY_REPO --set
 image.tag=COMPANY_TAG} {--set
 proxyUrl=http://proxy_host:proxy_port} {--set
 openShift=true} k8s-informer HELM_URL
```

Note: If the Informer is not installed in OpenShift, there is no need to set this parameter.

ii. In the command, configure additional settings as needed:

- Configure informer pod memory limit and memory request
- Change the full discovery frequency in Kubernetes Visibility Agent
- Deactivate continuous discovery in Kubernetes Visibility Agent
- Suppress frequent changes in the Kubernetes cluster to reduce the load on your instance. When this feature is active, the Informer does not report changes above the configured threshold rate back to the instance. For more information, see the Suppressing frequent changes in Kubernetes clusters reported by Kubernetes Visibility Agent (formerly CNO for visibility) [KB1580745] article in the Now Support Knowledge Base.
- Control the maximum number of full discoveries that Informers can run on multiple clusters at the same time. For more information, see the Scheduling Full Discoveries in Kubernetes Visibility Agent (formerly CNO for Visibility) [KB1580250] article in the Now Support Knowledge Base.

b. Run the command.

- Upgrade Kubernetes Visibility Agent using the Kubernetes YAML file.

a. Download the Kubernetes YAML zip file provided in the Kubernetes Visibility Agent (formerly CNO for Visibility) Helm Chart and Kubernetes YAML file releases [KB1564347] article in the Now Support Knowledge Base.

b. Extract the k8s_informer.yaml and EULA.pdf files from the zip file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Edit the k8s_informer.yaml file.

- Change the value of ACCEPT_EULA to `"Y"`, as follows:

```
- name: ACCEPT_EULA
  value: "Y"
```

> **Note:** By changing the value to `"Y"`, you agree to the End-User License Agreement included in the EULA.pdf file.

- Replace all occurrences of <NAMESPACE> with the namespace in which you want to install the Informer.
- Replace all occurrences of <INSTANCE_NAME> with the name of your instance, without the domain name.
- Replace <CLUSTER_NAME> with the name of your cluster as it appears in the CMDB.
- If the Informer is installed in OpenShift, change the value of the OPENSHIFT environment variable to `true`.

d. Run the command:

```
kubectl apply -f k8s_informer.yaml
```

What to do next

For information about Kubernetes Visibility Agent administration and troubleshooting, see the Administrating and Troubleshooting Kubernetes Visibility Agent (formerly CNO for Visibility) [KB1532614] article in the Now Support Knowledge Base.

Kubernetes Visibility Agent Reference

Reference information about Kubernetes Visibility Agent.

Data collected by Kubernetes Visibility Agent

The Kubernetes Visibility Agent Informer collects data on Kubernetes resources and populates various CMDB tables.

For the information about the data Kubernetes Visibility Agent Informer collects on the Kubernetes resources and OpenShift resources, and about the CMDB tables it populates, see the KB article KB2669848

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships created by Kubernetes Visibility Agent (KVA)

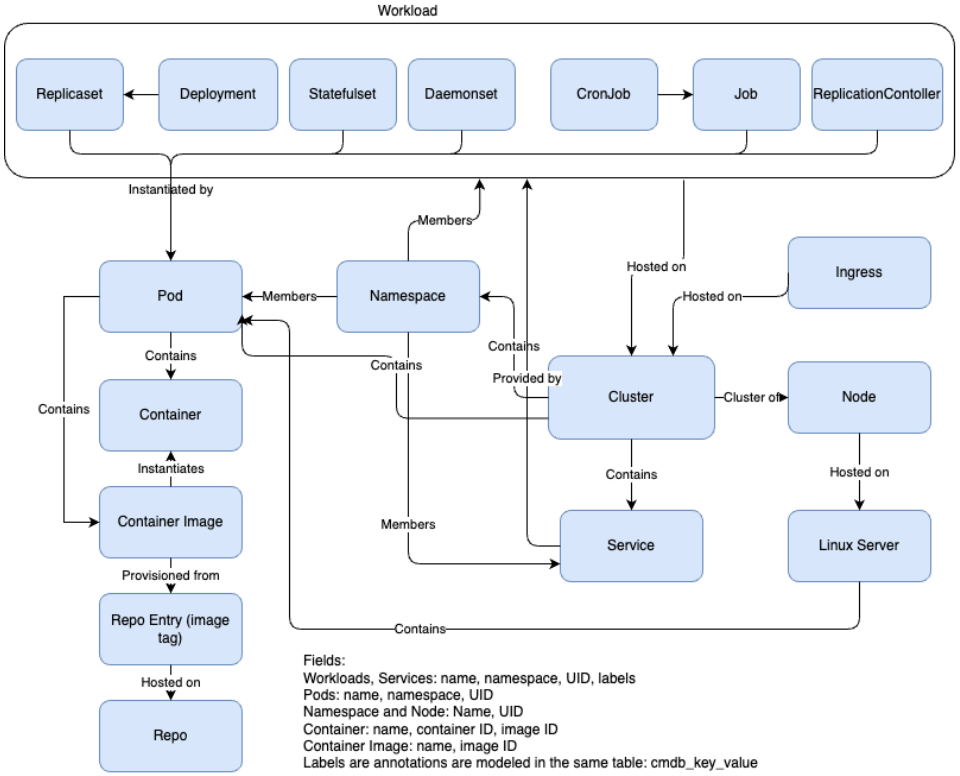

Kubernetes Visibility Agent Informer status fields

The fields in the Kubernetes Visibility Agent Informers table `sn_acc_visibility_kubernetes_informer` describe the status of the Informer pods deployed in your Kubernetes clusters.

Kubernetes Visibility Agent Informers table

Field	Description
Informer name	System-generated name of the Informer, based on the unique ID of the cluster. When a new cluster CI is created, the Informer name is identical to the name you configured for the cluster. Select this field to view the available actions on the selected Informer.
Kubernetes cluster	Reference to the cluster CI. Select this field to view details of the cluster.
Full discovery status	Status of the current full discovery. Possible values are: Empty, In Progress, Queued, Completed, Failed, Skipped, and Timed out.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes Visibility Agent Informers table (continued)

Field	Description
	Note: Full discovery is performed either periodically or on demand. By default, it is run once a day. Full discovery status is considered Skipped if, when the time came to run full discovery by a given informer, the informer was in the Down state and was not able to communicate with the instance. The full discovery status is considered Timed Out if the instance doesn't receive confirmation that the full discovery is completed within the expected time window.
Last full discovery	The time that the last full discovery was completed.
Informer version	The version of the Informer pod that is running in the cluster.
Status	The status of the Informer. Possible values are: Up, Down, Paused. **Note:** The status of an Informer is considered Down when no message has been received in the last 15 minutes.
Continuous Discovery Status	The status of continuous discovery. Possible values are: On or Off.
Upgrade Status	The upgrade status of Informers in the cluster. Possible values are: Upgrade Pending, Upgrading, Desired image in use, Upgrade using kubectl/Helm. **Note:** When the upgrade status of an Informer is Upgrade using kubectl/Helm, contact your Kubernetes admin to upgrade the informer manually.

Kubernetes Visibility Agent support matrix

The Kubernetes Visibility Agent component is supported on several versions of Kubernetes and OpenShift and in various Kubernetes environments.

Note: To deploy KVA in the OKE (Oracle Kubernetes Engine) environment, you need to create a Kubernetes Cluster on Oracle Cloud Infrastructure. For more information, see the Oracle documentation.

For the latest information, see the Kubernetes Visibility Agent (formerly CNO for Visibility) Support Matrix [KB1700730] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

FIPS-140 compliance

In regulated markets that require FIPS-140 compliant code, the Kubernetes Visibility Agent Informer provides the following compliance:

- The main informer image is FIPS-140 compliant.
- The daemonset image used for traffic discovery is FIPS-140 compliant.

Kubernetes Visibility Agent (KVA) Performance Benchmark

Performance tests help you estimate the load, memory, and CPU on the Instance and Kubernetes cluster.

For the performance results from multiple scenarios that are tested, see the Performance results for Kubernetes Visibility Agent [KB1555851] article in the Now Support knowledge base.

Discovering Kubevirt Virtual Machines

KubeVirt provides a unified development platform where developers can build, modify, and deploy applications in both Application Containers and Virtual Machines.

To enable and configure the KubeVirt feature, see the Retrieving Kubevirt Virtual Machines from Kubernetes Clusters using Kubernetes Visibility Agent [KB2613344] article in the Now Support knowledge base.

Running Kubernetes Visibility Agent (KVA) at Scale

When you run Kubernetes Visibility Agent (KVA) (formerly CNO for Visibility) on more than 100 clusters, configure your system properties to help prevent loss of data from the informers.

For KVA at scale settings, see the Running Kubernetes Visibility Agent (CNO for Visibility) at scale [KB1781212] article in the Now Support knowledge base.

Kubernetes Visibility Agent advanced configuration

This section refers to advanced configuration options for the Kubernetes Visibility Agent component running in a Kubernetes cluster.

For the latest information, see the Kubernetes Visibility Agent (formerly CNO for Visibility) Advanced Configuration Options [KB1648891] article in the Now Support Knowledge Base.

Kubernetes cluster auto-retirement properties

Configure system properties to control the automatic retirement of inactive Kubernetes cluster configuration items (CIs) and their associated resources, so your CMDB reflects only active infrastructure.

The auto-retirement feature updates the *install_status* field to retired on inactive Kubernetes cluster CIs and all associated CIs (pods, namespaces, deployments, containers, and other resources). A cluster is inactive when its associated Informer is in the Down state and the cluster CI has not been updated for a configured period.

The system checks these conditions during each full discovery cycle. When the time comes to run full discovery and the Informer is in the Down state, the system evaluates the conditions and updates the status accordingly.

Configuration properties

The following system properties control the auto-retirement behavior. Properties are located in the sys_properties table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Auto-retirement configuration properties

Property name	Type	Default value	Descript
sn_acc_visibility.auto_retire_k8s_clusters	Boolean	false	Determ the sys inactive cluster by Kube Visibilit Set to t auto-rel
sn_acc_visibility.auto_retire_period_days	Integer	60	Numbe of inact which t retires a Kubern The sys *sys_u* field to when th was las
sn_acc_visibility.auto_retire_install_status	String	7	The nur value th writes t *insta* field of Kubern CI and i CIs. Val include • 7 (Rel • 100 (**Wa** If y valu (Ab clu all a CIs by cle Thi irre
sn_acc_visibility.node_deletion_set_server_install_status	String	No value (empty)	The nur value th writes t *insta* field of cmdb_c CIs ass

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Auto-retirement configuration properties (continued)

Property name	Type	Default value	Descript
			Kubern The sys this valu nodes a from the when th retired. doesn't empty, takes n Linux se

Install Kubernetes Visibility Agent (KVA) using the Terraform Helm provider

You can install Kubernetes Visibility Agent using the Terraform Helm provider.

For more information, see the Install Kubernetes Visibility Agent (formerly CNO for Visibility) using the Terraform Helm provider [KB1540556] article in the Now Support Knowledge Base.

Configuring Discovery

Configure the elements that Discovery requires to investigate your network, such as credentials, schedules, and IP addresses.

Configuration overview

Discovery can run on a regular, configurable schedule, or can be launched manually. The IP address ranges you provide tell Discovery which specific devices to investigate. To retrieve useful information, Discovery needs credentials for devices within a particular range, so that Discovery can connect to and run various probes on the devices it finds. Discovery compares the devices it finds with configuration items (CI) in the CMDB and updates any matching devices. If Discovery does not find a matching CI in the CMDB, it creates a CI.

Discovery configuration procedures

Use the following links to configure Discovery for your environment. You don't need to perform all these procedures to run a Discovery. The platform provides many defaults you can use to explore your network that are suitable for most discoveries. To get started quickly with Discovery, use the ITOM Configuration Console. The Configuration Console offers a faster way to install and update apps, with interactive configuration pages, task assignment, and automated implementation by ServiceNow Otto. For more information, see ITOM Configuration Console for Discovery.

Configuration procedures

Discovery IP address configuration	Discovery classifiers	Discovery status
Domain Separation and Discovery	Discovery classification parameters	Discovery Configuration Console
SNMP support for Discovery	Discovery identifiers	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request Discovery

Request and activate the Discovery (com.snc.discovery) plugin to enable Discovery capabilities on your instance.

Before you begin

To purchase a Discovery subscription, contact your ServiceNow account manager or sales representative.

Role required: admin

About this task

Discovery is available as a separate subscription from the rest of the ServiceNow AI Platform and requires the Discovery (com.snc.discovery) plugin. When the Discovery plugin is activated, several plugins and applications are also installed. For more information, see Plugins or applications installed with ITOM Visibility. For additional details about plugins, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Related topics

Discovery setup

Discovery monitoring and issue resolution

Discovery setup

After you activate the Discovery application, you have several ways to get started.

Note: If the Discovery (com.snc.discovery) plugin is activated on the instance and a language plugin (for example, Spanish) is also activated, make sure that the MID Server language is also set to English/None on the user record.

Select your method of setup:

Setup method	Description
ITOM Configuration Console for Discovery	Use this method for a faster way to install and update apps, with interactive configuration pages, task assignment, and automated implementation by ServiceNow Otto. This method requires additional plugins and configuration. For more information, including prerequisites, see Access the ITOM Configuration Console.
Discovery Guided Setup	Use this method if you want to set up everything you need to use Discovery, including a MID Server. ITOM Discovery Guided Setup provides a sequence of tasks that help you configure Discovery on your ServiceNow instance. To access the guided setup, navigate to **All** > **ITOM Guided Setup** > **ITOM Discovery Guided Setup**. For more information about using the guided setup interface, see Using guided setup .
Without Guided Setup	Use this method if you want to set up each component of setup separately. You can configure MID Servers, credentials, and roles separately, then configure one of the discovery types from the Discovery Schedule.
Discovery Quick Start	Use this wizard to configure everything you need to perform subnet and CI Discovery in your environment. Before you can use Quick Start, you must first install and validate your MID Servers and create the necessary credentials.

Related topics

Discovery monitoring and issue resolution

ITOM Configuration Console for Discovery

The ITOM Configuration Console gives administrators a single place to complete all Discovery setup steps, from assigning administrators to scheduling discovery runs. The console contains two modules: Discovery and Event Management.

Discovery guides you through the end-to-end setup of your IT infrastructure scanning. It covers assigning administrators, deploying and validating MID Servers, configuring IP and cloud credentials, scheduling discovery runs, and reviewing schedules. It also includes diagnostics

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

to identify discovery errors and detect anomalies, ensuring your environment is fully validated before you promote your configuration to production. Starting with ITOM Configuration Console v28.3.1, you can view configuration guidance as you set up each task. The guidance shows the actions to take and the actions to avoid, so you can configure Discovery correctly and avoid common setup mistakes. To view available tasks in the Discovery module, see Admin in ITOM Configuration Console.

Event Management helps you take control of your IT alerts from start to finish. It begins with assigning administrators, setting up integration accounts, and installing the necessary integrations to connect your monitoring tools. From there, you can review and enrich incoming alerts, automatically create incidents, and reduce noise by grouping, ignoring, or delaying alerts. The platform also provides structured training for your operations team on key tools such as Service Operations Workspace, Express List, and Service Dashboard. Finally, built-in dashboards help you report on outcomes and measure progress before promoting your configuration to production.

Admin in ITOM Configuration Console

Administrators can configure Discovery settings, assign roles, manage credentials, and review Discovery schedules, errors, and anomalies through the ITOM Configuration Console.

The following table describes the administrative tasks available in the Discovery module.

Discovery administrative tasks

Task	Description	Required Role
Assign Discovery admins	Assign or remove the discovery_admin role from users to manage Discovery administration permissions.	admin
MID Server user	Create a dedicated user with the mid_server role to authenticate the MID Server to your ServiceNow instance.	admin
Deploy MID Server	Use MID Server to securely connect ServiceNow behind your firewall, enabling data flow, automated discovery, and seamless integration.	admin
Validate MID Server	Validate that your MID Servers are trusted to connect to your ServiceNow instance.	admin
IP credentials	Create IP credentials for Discovery to securely connect to devices in your environment.	discovery_admin
Cloud Credentials	Create cloud credentials to manage access to cloud-based applications.	discovery_admin
Quick discovery	Scan one or more IPs to see what's running in your environment without creating on going schedules.	discovery_admin
Discovery schedules	Configure IP-based or cloud-based Discovery schedules in Discovery Admin Workspace. For cloud-based schedules, you can configure discovery for AWS, Azure, GCP, Alibaba Cloud, IBM, OCI, OpenStack, VMware, and oVirt.	discovery_admin
Schedule review	Access the Schedules page in Discovery Admin Workspace to view and manage your Discovery schedules.	discovery_admin
Discovery errors	Access the Diagnostics page in Discovery Admin Workspace to review errors found from your discoveries.	discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery administrative tasks (continued)

Task	Description	Required Role
Anomaly detection	Access Diagnostics page in Discovery Admin Workspace to review anomalies found in your Discovery schedules.	discovery_admin

Install ServiceNow Otto for ITOM

Install the ServiceNow Otto for IT Operations Management (ITOM) plugin to enable generative AI and agentic AI on your instance.

Before you begin

Role required: admin

Procedure

1. Navigate to **Admin** > **Admin Home**.
   The Admin Home page displays.

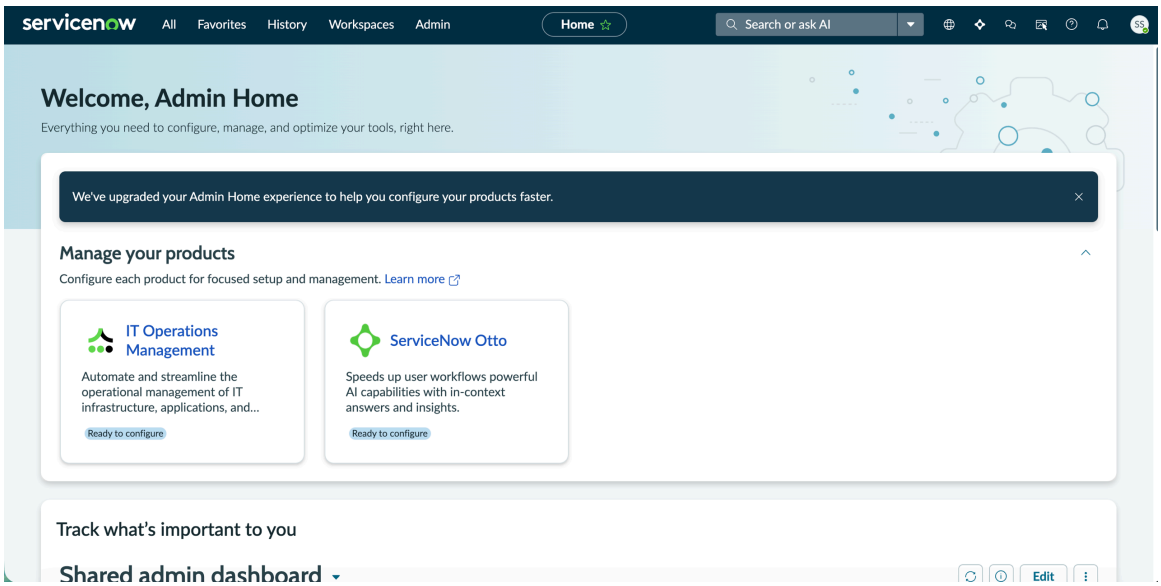

2. Select the IT Operations Management tile.
   The IT Operations Management page in the Product Hub displays.
3. Select the **Not Installed** tab.
4. On the ServiceNow Otto for IT Operations Management (ITOM) card, select the **Application Manager** icon ( ).
   The Application Manager page displays.
5. Select **Install**.
   Once the installation completes, ServiceNow Otto for IT Operations Management (ITOM) displays in the Installed tab of the Product Hub.

Access the ITOM Configuration Console

Access the IT Operations Management Configuration Console to set up Discovery.

Before you begin

Verify the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You're using the Zurich Patch 8 or later version of the ServiceNow AI Platform.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

Role required: admin

Procedure

1. Navigate to **Admin** > **Admin Home** on your instance.
2. Select the IT Operations Management tile.
   The IT Operations Management page in the Product Hub displays.

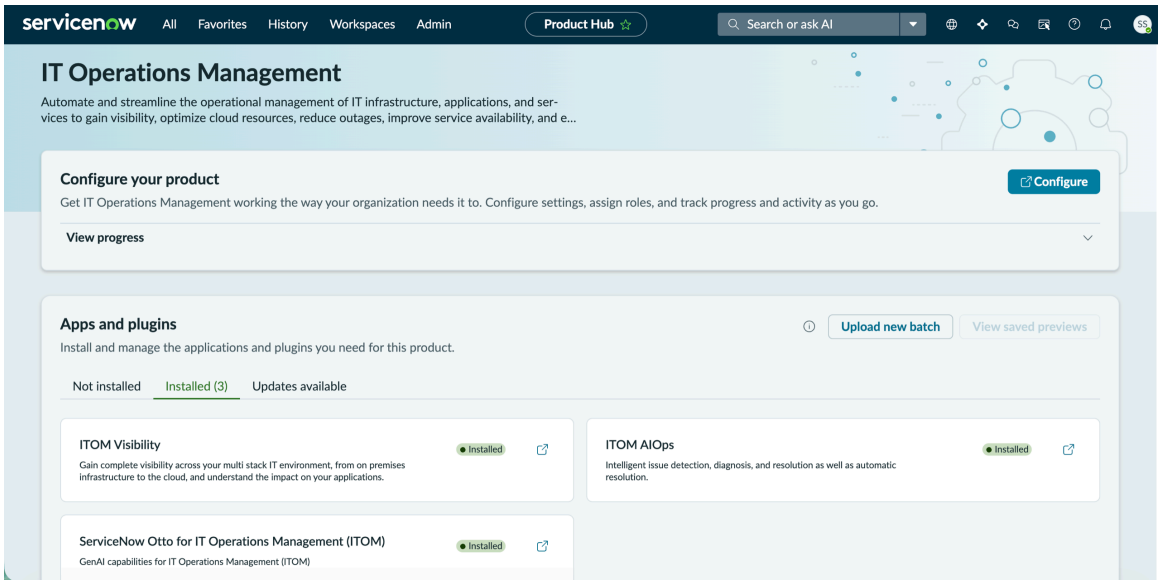

3. Select **Configure**.

   The Configure IT Operations Management page displays.

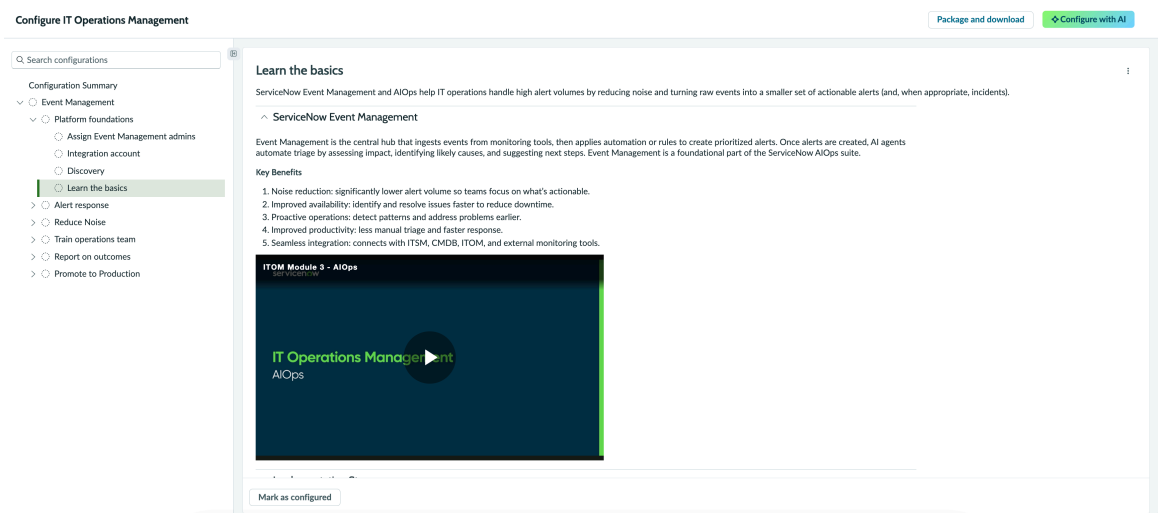

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Assign Discovery admins

Assign the discovery_admin role to users who will be in charge of Discovery configuration and operational control.

Before you begin

Verify the following:

- You're using the Zurich Patch 8 or later version of the ServiceNow AI Platform.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install ServiceNow Otto for ITOM.
- You're on the Configure IT Operations Management page of the Configuration Console. For more information, see Access the ITOM Configuration Console.

Role required: admin

About this task

The discovery_admin role provides users with the necessary permissions to configure and manage Discovery settings, including MID Servers, credentials, IP ranges, and Discovery schedules.

Procedure

1. Navigate to **Configuration Summary** > **Discovery** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **Assign Discovery admins**.
4. In the **User** field, select the user to assign the role to.

   The **Role** field is read-only and displays discovery_admin.
5. Select **Assign role**.
   The role is assigned to the selected user.
6. To complete the setup, select **Mark as configured**.

Create a MID Server user

Create a user account with the mid_server role so MID Servers can authenticate and communicate with the ServiceNow instance.

Before you begin

Verify the following:

- You're using the Zurich Patch 8 or later version of the ServiceNow AI Platform.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install ServiceNow Otto for ITOM.
- You're on the Configure IT Operations Management page of the Configuration Console. For more information, see Access the ITOM Configuration Console.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

MID Servers require a MID Server user account to authenticate with the instance. Assigning the mid_server role allows the MID Server to connect and pass validation checks. ServiceNow monitors MID Server user settings to help prevent validation failures. You can use a dedicated user per MID Server or share a user across multiple MID Servers.

Note: Using the same logged#in user across multiple MID Servers can generate issue records when more than one MID Server is **Up**. Configure a unique logged#in user per MID Server. See (KB1552863) MID Server Unique Logged In User for more information and remediation steps.

Procedure

1. Navigate to **Configuration Summary** > **Discovery** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **MID server user**.
4. Select **Create user**.
   The Create MID server user page displays.
5. Complete the following fields.

Field	Description
User ID	Enter a user name for the MID Server user. This name is specified in the *mid.instance.username* parameter of the configuration file that the MID Server installer creates. For details, see MID Server parameters. **Note:** Avoid using the same MID Server logged in user across multiple MID Servers.
Password	Enter a password for the MID Server user. This password is specified in the *mid.instance.password* parameter of the configuration file that the MID Server installer creates.

6. Select **Create user**.
   You're returned to the MID server user page.

   Note: The MID Server user is limited to MID Server setup. To complete Discovery setup through the ITOM Configuration Console, you must have the discovery_admin role.

7. To complete the setup, select **Mark as configured**.

Related topics

Create the MID Server user and grant the role

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Deploy a MID Server

Deploying a MID Server lets ServiceNow securely communicate with systems inside a private network by installing and validating a trusted runtime that performs discovery, integrations, and automation tasks without requiring inbound network access.

Before you begin

Verify the following:

- The ServiceNow AI Platform is running on the Brazil, Australia, or Zurich release starting with Patch 8.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install ServiceNow Otto for ITOM.
- You're on the Configure IT Operations Management page of the Configuration Console. For more information, see Access the ITOM Configuration Console.

Role required: admin

About this task

Discovery requires a MID Server to securely detect and collect data from your private network, identifying infrastructure and dependencies without exposing internal resources to the internet.

Procedure

1. Navigate to **Configuration Summary** > **Discovery** > **Platform foundations**.
2. Select **Deploy MID server**.
3. Select the Application manager icon ( ).
   The ITOM Infra Services Workspace displays.
4. Download the MID Server for the appropriate operating system.
   For more information, see Install a MID Server on Linux and Install a MID Server on Windows.
5. Select **X** to close the window and return to the Configuration Console.
6. To complete the setup, select **Mark as configured**.

Validate a MID Server

Validate a newly installed MID Server so it can communicate with the ServiceNow instance. During validation, you can optionally define which applications, capabilities, and IP ranges the MID Servers is allowed to use.

Before you begin

Verify the following:

- You're using the Zurich Patch 8 or later version of the ServiceNow AI Platform.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install ServiceNow Otto for ITOM.
- You're on the Configure IT Operations Management page of the Configuration Console. For more information, see Access the ITOM Configuration Console.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Validating a MID Server establishes trust between ServiceNow and the MID Server. A validated MID Server can securely access credentials and perform automation and Discovery tasks, and prevents untrusted systems from executing work on behalf of the instance.

Procedure

1. Navigate to **Configuration Summary** > **Discovery** > **Platform foundations**.
2. Select **Validate MID server**.
3. Select the Application manager icon ( ).
   The ITOM Infra Services Workspace displays.
4. In the Total MID Servers list, select the check box next to the MID Server you want to validate.
5. Select the **Server Actions** drop-down list and select **Validate**.
6. Select **X** to close the window and return to the Configuration Console.
7. To complete the setup, select **Mark as configured**.

Create IP credentials

Create IP credentials for Discovery to securely connect to devices in your environment. Discovery uses these credentials to log in to devices and retrieve configuration data during the discovery process.

Before you begin

Verify the following:

- You're using the Zurich Patch 8 or later version of the ServiceNow AI Platform.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install ServiceNow Otto for ITOM.
- You're on the Configure IT Operations Management page of the Configuration Console. For more information, see Access the ITOM Configuration Console.

Role required: discovery_admin

About this task

IP credentials store the authentication information that Discovery needs to access devices in your network. When Discovery runs, it tries each credential in the order you specify until it finds one that works for a given device. Set the credential order to control which credentials Discovery tries first. This is especially important when security policies lock out users after failed login attempts.

For more information about credential types and requirements, see Create and test your credentials .

Procedure

1. Navigate to **Configuration Summary** > **Discovery** > **Discovery Setup**.
2. Select **IP credentials**.
3. From the **Credential type** drop-down list, select a credential type.
   The following credential types are available for IP credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credential type	Description
Applicative Credentials	Credentials for discovering applications running on host machines.
Basic Auth Credentials	Credentials that use basic authentication with a user name and password.
CIM Credentials	Credentials for devices that use the Common Information Model protocol.
Kubernetes Credentials	Credentials for discovering Kubernetes clusters and resources.
oVirt Credentials	Credentials for oVirt virtualization environments.
Secondary Bot Hash Token Credential	Token credentials for secondary bot integrations that use existing tokens.
Secondary Bot Static Token Credential	Static token credentials for secondary bot integrations.
SNMP Community Credentials (Password Only)	Credentials that use an SNMP community string for network devices.
SNMPv3 Credentials	Credentials that use the SNMPv3 protocol with authentication and privacy options.
SSH Credentials	Credentials for Linux, UNIX, and other SSH-enabled devices.
SSH Private Key Credentials	Credentials that use an SSH private key for authentication instead of a password.
SSL Keystore Credentials	Credentials that use an SSL keystore for secure communication.
VMWare Credentials	Credentials for VMware vCenter environments.
Windows Credentials	Credentials for Windows devices that use WMI or WinRM protocols.

4. Select **Create new**.
5. Complete the form.
   The fields in the form vary based on the credential type that you selected. The following fields appear in most credential types.

Field	Description
Name	Enter a unique and descriptive name for the credential.
User name	Enter the user name for this credential. Avoid leading or trailing spaces in user names.
Active	Select this check box to make the credential available for Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Credential alias	Select a credential alias to associate with this credential. Credential aliases enable you to define a credential once and reference it across multiple configurations.
Order	Enter the sequence in which Discovery tries this credential when it attempts to log in to devices. The smaller the number, the higher the priority. If all credentials have the same order number, Discovery tries them in a random order. **Note:** Establish credential order when you use large numbers of credentials or when security policies lock out users after failed login attempts.

6. Select **Submit**.
   You're returned to the IP credentials page.
7. To complete the setup, select **Mark as configured**.

Related topics

Get started with credentials

Create cloud credentials

Create cloud credentials to manage access to cloud-based applications, including Amazon Web Services, Microsoft Azure, and Google Cloud. Discovery uses these credentials to connect to your cloud provider accounts and discover cloud resources.

Before you begin

Verify the following:

- You're using the Zurich Patch 8 or later version of the ServiceNow AI Platform.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install ServiceNow Otto for ITOM.
- You're on the Configure IT Operations Management page of the Configuration Console. For more information, see Access the ITOM Configuration Console.

Role required: discovery_admin

About this task

Cloud credentials store the authentication information that Discovery needs to access your cloud provider accounts. Before you create cloud credentials, verify that you have the required access and permissions configured in your cloud provider environment. The specific configuration requirements vary by cloud provider.

For more information about cloud provider setup and permissions, see Discovery for cloud environment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Configuration Summary** > **Discovery** > **Discovery Setup**.
2. Select **Cloud credentials**.
3. From the **Credential type** drop-down list, select a credential type.
   The following credential types are available for Cloud credentials.

Credential type	Description
Alibaba Cloud API	API credentials for discovering Alibaba Cloud resources. Requires an AccessKey ID and AccessKey secret from the Alibaba Cloud Console.
AWS	Credentials for discovering Amazon Web Services cloud resources.
AWS SSM Instance User	User credentials for AWS Systems Manager (SSM) Agent discovery on EC2 instances. Requires only a user name for the alternate user credential configured on the EC2 instance.
Azure Enterprise Agreement	Credentials for Azure Enterprise Agreement enrollments.
Azure SAS	Credentials that use a shared access signature token for Azure storage access.
Azure Service Principal	Service principal credentials for Azure cloud discovery. Requires a tenant ID, client ID, and secret key from the Azure portal.
Azure SSH Certificate	SSH certificate credentials for discovering Azure Linux virtual machines. Requires a linked Azure Service Principal credential.
Cloud Management	General credentials for cloud management operations.
CloudFoundry	Credentials for Cloud Foundry environments.
Google API	API credentials for Google Cloud services.
Google Cloud	Credentials for discovering Google Cloud Platform resources.
IBM	Credentials for discovering IBM Cloud Platform resources. Requires a SoftLayer API user and key, or a Bluemix API key.
OpenStack	Credentials for OpenStack cloud environments.
Oracle API	API credentials for Oracle Cloud Infrastructure discovery.
oVirt	Credentials for oVirt virtualization environments.
VMWare	Credentials for VMware vCenter cloud environments.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Select **Create new**.
5. Complete the form.
   The fields in the form vary based on the credential type that you selected. The following fields appear in most credential types.

Field	Description
Name	Enter a unique and descriptive name for the credential.
User name	Enter the user name for this credential. Avoid leading or trailing spaces in user names.
Active	Select this check box to make the credential available for Discovery.
Credential alias	Select a credential alias to associate with this credential. Credential aliases enable you to define a credential once and reference it across multiple configurations.
Order	Enter the sequence in which Discovery tries this credential when it attempts to log in to devices. The smaller the number, the higher the priority. If all credentials have the same order number, Discovery tries them in a random order. **Note:** Establish credential order when you use large numbers of credentials or when security policies lock out users after failed login attempts.

6. Select **Submit**.
   You're returned to the Cloud credentials page.
7. To complete the setup, select **Mark as configured**.

Related topics

Get started with credentials

Run a Quick Discovery

Scan one or more IP addresses to discover devices in your network without creating ongoing schedules.

Before you begin

Verify the following:

- You're using the Zurich Patch 8 or later version of the ServiceNow AI Platform.
- You have installed the ITOM Visibility plugin. For more information, see Install ITOM Visibility using ServiceNow Otto for Setup.
- You have installed the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install ServiceNow Otto for ITOM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You have at least one MID Server that is validated and in an **UP** state. The MID Server must be able to reach the IP ranges that you want to discover. For more information, see Validate a MID Server.
- You have configured the credentials needed to access devices in your target IP ranges. For more information, see Create IP credentials.
- You're on the Configure IT Operations Management page of the Configuration Console. For more information, see Access the ITOM Configuration Console.

Role required: discovery_admin

About this task

Use quick discovery to run an on-demand, IP-based discovery without configuring a full Discovery schedule. Quick discovery is useful for verifying that Discovery can reach and identify devices before you set up recurring schedules.

Procedure

1. Navigate to **Configuration Summary** > **Discovery** > **Discovery Setup**.
2. Select **Quick Discovery**.
3. In the **Enter IP list** field, enter the IPv4 addresses that you want to discover.
   You can enter IP addresses in the following formats.

Format	Example
IP network addresses	`10.0.1.0/24`
IP address range	`10.0.2.1-10.0.2.15`
Comma-separated IP addresses	`10.0.3.176,10.0.3.222,2001::100,2600:1bdc:`

Note: You can combine multiple formats in a single comma-separated entry.

4. From the **MID selection method** drop-down list, select how to assign a MID Server for the discovery.
   The following MID Server selection methods are available.

Option	Description
Auto-select MID Server	Discovery automatically selects an appropriate MID Server to run the discovery.
Specific MID Server	Select a specific MID Server from the list. Only MID Servers that are validated and in an **Up** state are available.
Specific MID cluster	Select a MID Server cluster. Clusters provide failover protection and load balancing between MID Servers.

5. Select **Run discovery**.
   Discovery scans the specified IP addresses and identifies devices in your network.
6. To complete the setup, select **Mark as configured**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Get started with credentials

Discovery Guided Setup

ITOM Discovery Guided Setup provides a sequence of tasks that help you install and get started with Discovery. If you're setting up Discovery for the first time, this is a good way to get your environment up and running quickly.

Note: This guided setup has been replaced by the ITOM Configuration Console (available in Zurich Patch 8 or later versions of the ServiceNow AI Platform), which offers a faster way to install and update apps, with interactive configuration pages, task assignment, and automated implementation by ServiceNow Otto. In ITOM Guided Setup v27.3.1, the Learn phase includes a new task that provides direct access to the ITOM Configuration Console. For more information, see ITOM Configuration Console for Discovery.

To access the ITOM Discovery Guided Setup, navigate to **All** > **ITOM Guided Setup** > **ITOM Discovery Guided Setup**. You're then prompted to enter a configuration name, which enables your progress to be saved and resumed later. If you're returning to the setup, select **Cancel** followed by **Exit** in the confirmation dialog to view the list of in-progress setup configurations.

The ITOM Discovery Guided Setup leads you through a structured series of configuration activities to promote accurate visibility in the Configuration Management Database (CMDB) and establish a strong foundation for IT operations.

Phases

The setup is organized into the four distinct phases: Learn, Crawl, Walk, and Run. Each task within a phase includes built-in resources, such as contextual embedded help and links to detailed documentation on the ServiceNow product documentation site. In some cases, you can even complete the task directly within the guided setup.

Learn

The Learn phase is the first step in the guided setup. It covers core principles, techniques, and prerequisites for Discovery.

Crawl

This phase walks you through the initial technical setup of Discovery, including installing the necessary plug-ins, configuring credentials, and preparing your instance for scheduling discoveries.

Walk

In this phase, you configure and activate Discovery schedules by setting up various schedule types, adding credentials, and installing and validating your MID Servers.

Run

This phase focuses on monitoring and optimizing ongoing Discovery operations, leveraging tools like Discovery Admin Workspace and Content 360 to track performance and promote continued success.

Key features

The ITOM Discovery Guided Setup provides step-by-step configuration support along with intuitive navigation to track progress and move smoothly through each phase.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Guided Setup landing page features

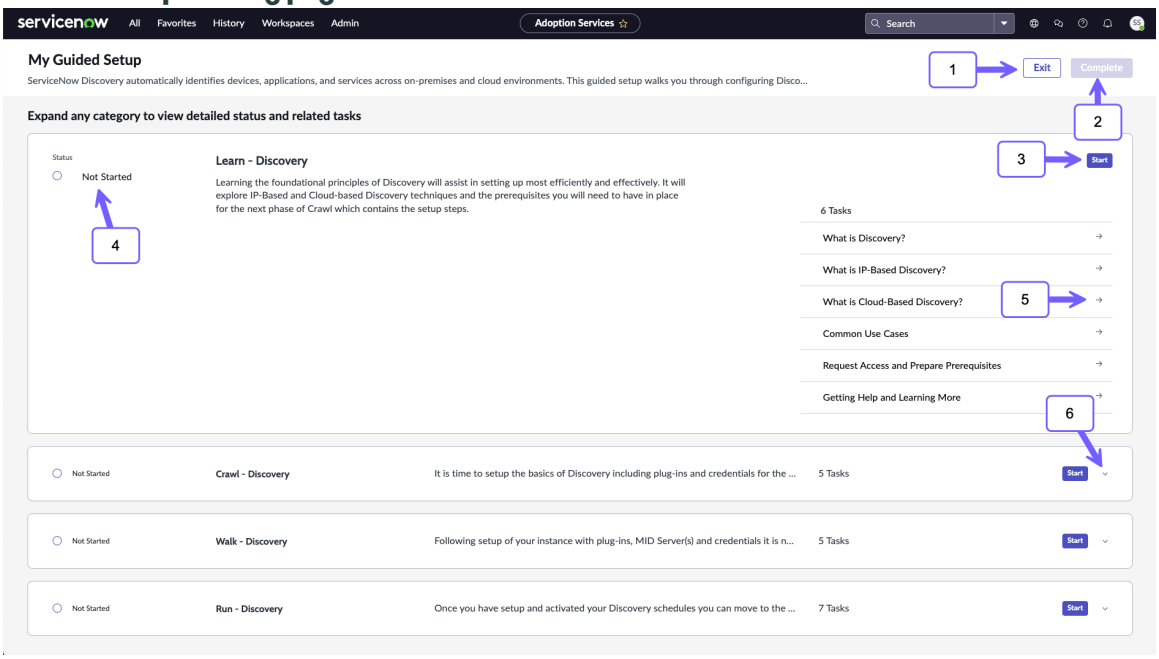

Interactive features within the landing page

Feature	Description
1. Exit	Exit the Guided Setup. **Note:** Selecting **Exit** in the confirmation dialog takes you back to the list of setup configurations currently in progress.
2. Complete	Mark the Guided Setup complete. **Note:** This option is only available once you have completed all the tasks in each phase of the setup.
3. Start	Launch the Guided Setup phase.
4. Status	Track your progress of the Guided Setup phase.
5. Task	Jump to specific activities within a phase.
6. Phase drop-down	View and access tasks within another phase.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Guided Setup in-phase features

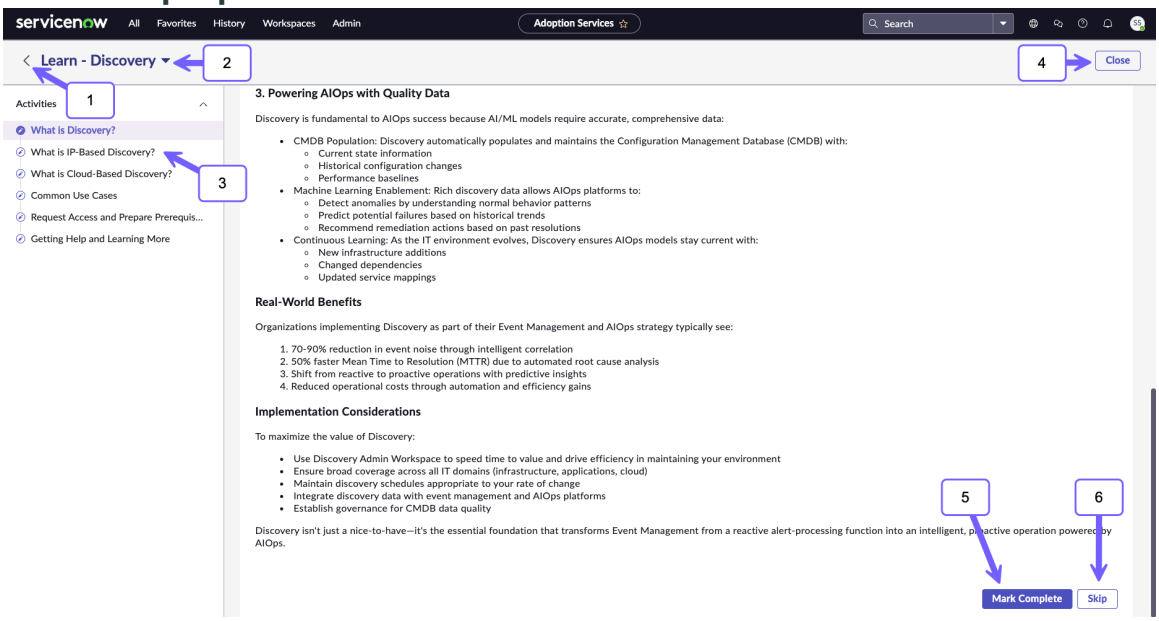

Interactive features within each phase

Feature	Description
1. Back arrow	Return to the Guided Setup landing page.
2. Phase drop-down	Toggle between different phases.
3. Activities	Jump to other activities within the current phase.
4. Close	Navigate back to the Guided Setup landing page.
5. Mark Complete	Track your progress and continue to the next activity in the phase.
6. Skip	Skip directly to the next activity in the phase.

Set up Discovery without Guided Setup

If you aren’t using Guided Setup, you must perform several configuration steps manually to activate the application, set up the MID Server, and then set up Discovery.

Before you begin

- Grant Discovery role: Confirm that users who are expected to configure and execute Discovery in your network have the discovery_admin role. This role grants access to the tables in the Discovery application.
- Designate a computer to host the MID Server.
- Gather the login credentials that the MID Server must use to access the devices it’s expected to discover.
- Determine which IP addresses Discovery must scan. In a highly complex network, group IP ranges into a range sets, which are reusable for different schedules. Discovery doesn’t scan anything outside of these ranges.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

https://player.vimeo.com/video/1113359093?h=41945233fc&badge=0&autopause=0&player_id=0&app_id=58479

About this task

Procedure

1. Install one or more MID Servers on physical or virtual servers that meet the MID Server system requirements and configure them to communicate with the appropriate ServiceNow instance. Install at least two MID Servers at first and assign them to different schedules and IP ranges to help complete discoveries quicker. Starting with the Istanbul release, you must make sure that the MID Server is assigned the Discovery application or the **ALL** application.
2. Verify MID Server connectivity by navigating to **Discovery** > **MID Servers**.
   If the MID Servers are configured properly, they appear in the list of MID Servers.
3. If a MID Server doesn't appear as a choice in the instance, perform the following checks:
   - Verify that the URL in the Configuration Parameters related list provided is correct.
   - If the MID Server is installed on Windows, make sure the ServiceNow MID Server service is configured properly with the correct login credentials and is running.
   - Check the MID Server log for errors.
     - On the instance, navigate to **Discovery** > **MID Servers** > **[MID Server]** > **Logs**.
     - On the MID Server host, navigate to the `agent\logs\agent0.log0` directory.
   - If Basic Authentication is enabled, a user name and password must be provided.
   - The MID Server might not have outbound access on port 443 (SSL) or a proxy server might be preventing TCP communication to the instance.
   - Make sure that no firewalls are blocking communication between the MID Server and the instance.
4. Set the credentials on the instance for all the computers, printers, and network gear in your environment.
   Credentials for Windows devices (using the WMI protocol) are provided by the login configured for the MID Server service on the Windows server host. Credentials for UNIX, vCenter, and SNMP must be configured on the instance. Discovery automatically determines which credentials work for a particular computer or device.
5. Define and run Discovery schedules.

   The Discovery Schedule is the control point for running discoveries. The schedule controls when Discovery runs, defines the MID Server to use, the type of Discovery that should run, and the IP addresses to query. Create as many schedules as necessary, using different types of discoveries, and configure them to run at any time. Let Discovery run on its configured schedule or manually execute Discovery at any time. You can set up a single schedule for the entire network or separate schedules for each location or VLAN. If you don't know the IP address to scan in your network, run a Network Discovery first to discover the IP networks. Once discovered, you can convert these networks into IP address range sets that you use in a Discovery Schedule.

   Note: For advanced discoveries, such as those requiring load balancing or scanning across multiple domains, use Discovery behaviors.

6. Validate the Discovery results.
   Open the ECC queue and analyze the XML payload, then check the Discovery log. For more information, refer to the section for validating results in Schedule a horizontal discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Quick Start

Discovery Quick Start is a wizard that helps you get up and running with Discovery quickly. Use Discovery Quick Start tasks to discover IP ranges, or subnets, in your environment and then to identify the configuration items (CI) in those subnets.

Before using Discovery Quick Start, install and validate MID Servers. See MID Server installation for details.

Configure and schedule subnet discovery

Subnet discovery finds all the subnets in your environment to determine available IP ranges and to verify that there are MID Servers available to access these ranges.

Before you begin

Role required: discovery_admin

About this task

The goal of this phase is to discover all the subnets in your network. To do this, Discovery must have the correct SNMP credentials to identify the routers it encounters. Using the credentials you provide, Discovery begins by reading the router tables of the router connected to the MID Server to identify connected routers. Discovery then moves outward through the network from router to router until it has discovered all the local subnets.

After Discovery identifies the subnets, it automatically assigns the IP ranges of these subnets to qualified MID Servers in preparation for the CI discovery phase.

Procedure

1. Navigate to **All** > **Discovery** > **Home**.
   The first time you access Discovery Home, a welcome screen appears. After you have configured subnet discovery and run the schedule, Quick Start shows the subnet results information.
2. Select **Discovery Quick Start** to begin configuration.

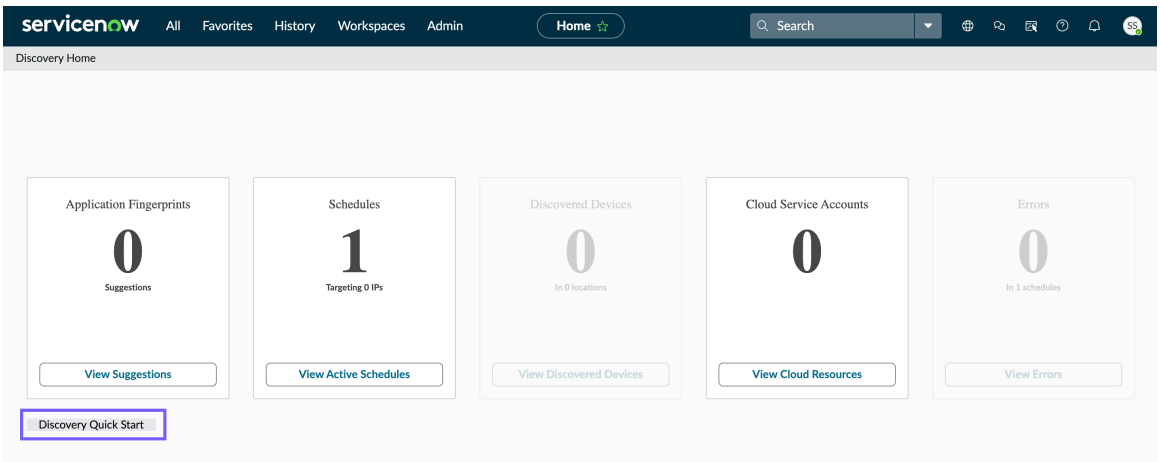

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you did not set up a MID Server yet, select **MID Server Guided Setup** when prompted to create the MID Server user, download and install the MID Server, and validate it. Links in Guided Setup return you to Discovery Quick Start to continue configuring subnet discovery. For more information about using the guided setup interface, see Using guided setup.

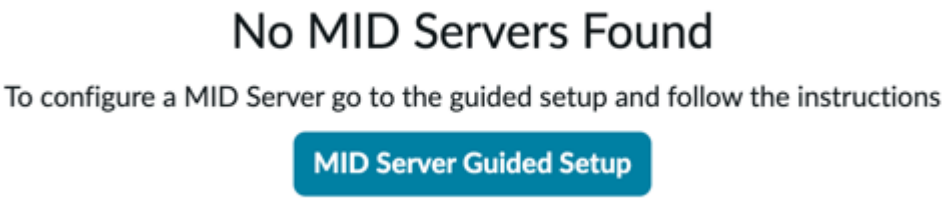

3. In the **Select MID Servers** tab, select qualified MID Servers for subnet discovery.

   a. Move the MID Servers you want to use for subnet discovery from the **Available** column to the **Selected** column.
   Available MID Servers must meet these requirements:

   - Be configured and in an **Up** state.
   - Be validated.
   - Have the Discovery application attached or **ALL** applications configured.

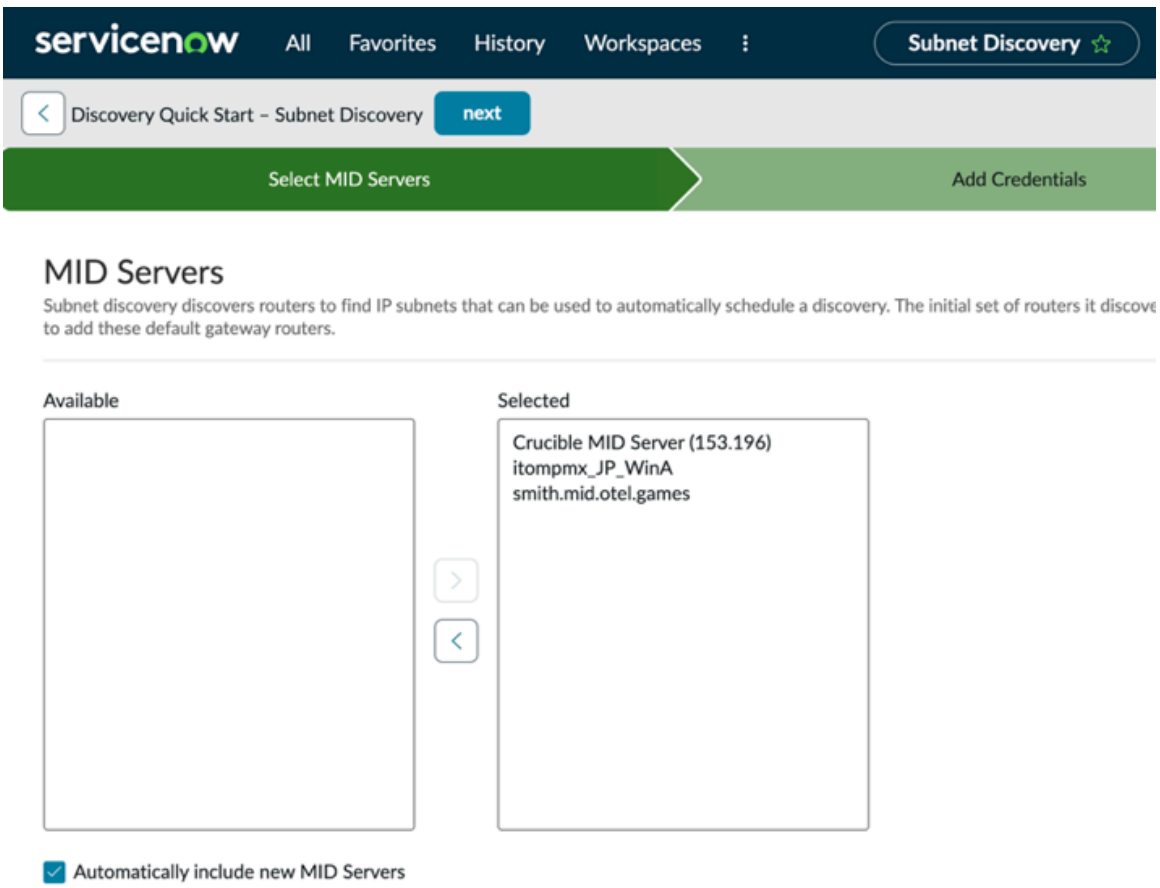

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. If you want Discovery to use all new MID Servers that meet the setup criteria, select the **Automatically include new MID Servers** option.

c. Select **next**.

4. In the **Add Credentials** tab, add credentials for the routers you're going to discover. This tab shows any SNMP credentials that are already configured. Routers that appear in the **Targeted Routers** list are default gateway routers to which the MID Server host is connected.

Note: You can return to this phase after creating a schedule to add credentials.

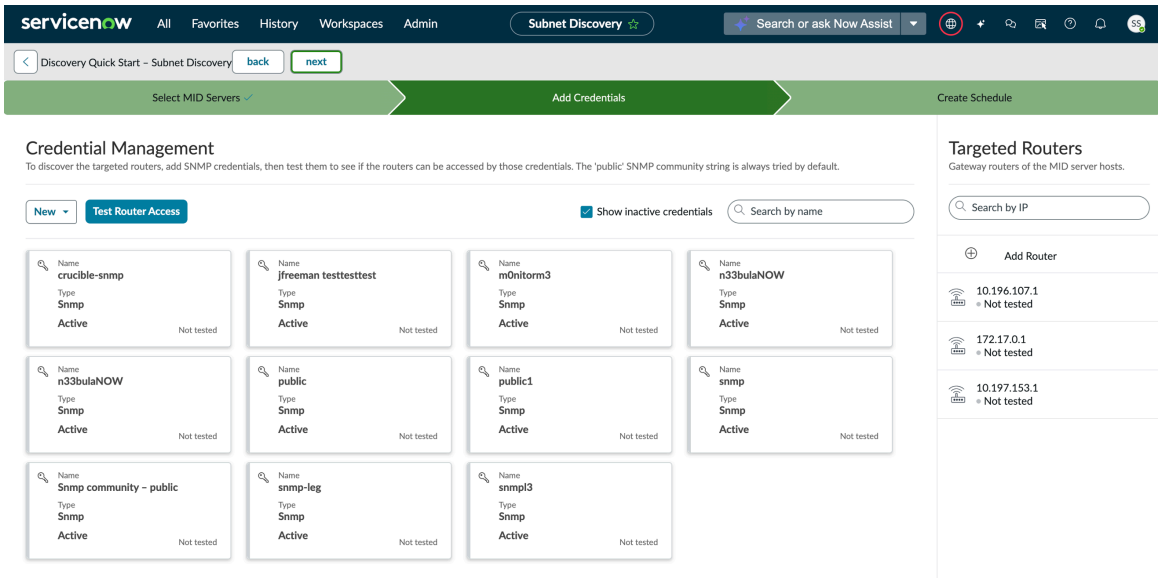

a. Select **New** to add SNMP community or SNMP V3 credentials, if necessary.

b. If the list of targeted routers does not include all the routers you expect to use, select the plus icon and add the router by IP address.

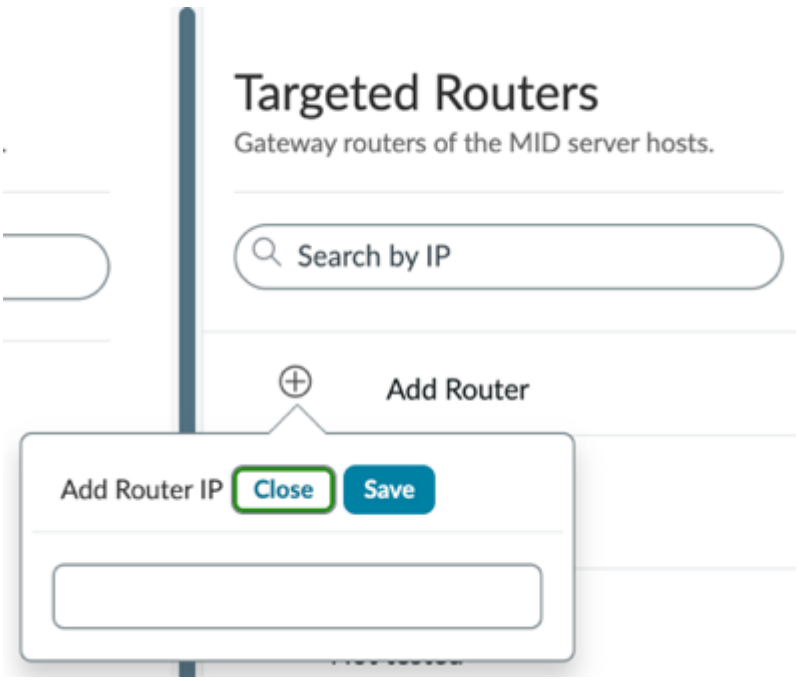

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. After you configure all necessary credentials, select **Test Router Access** to verify that Discovery can use your credentials to access routers.
The gateway routers appear in the **Routers** column, regardless of whether the credentials can access the router. Those routers that the credentials can successfully access appear with a **Success** indicator. The routers that cannot be reached, appear with a **Failed** indicator.

You can filter the list using the **Search by IP** field or by using the **Reached Routers** or **Unreached Routers** options from the **View** choice list.

d. Select **next**.

5. In the **Create Schedule** tab, create a schedule to run subnet discovery regularly.
Execution of this schedule performs subnet discovery and triggers IP subnet auto-assignment, in which the discovered subnets are assigned to qualified MID Servers.

a. Complete the schedule using the fields from the list.

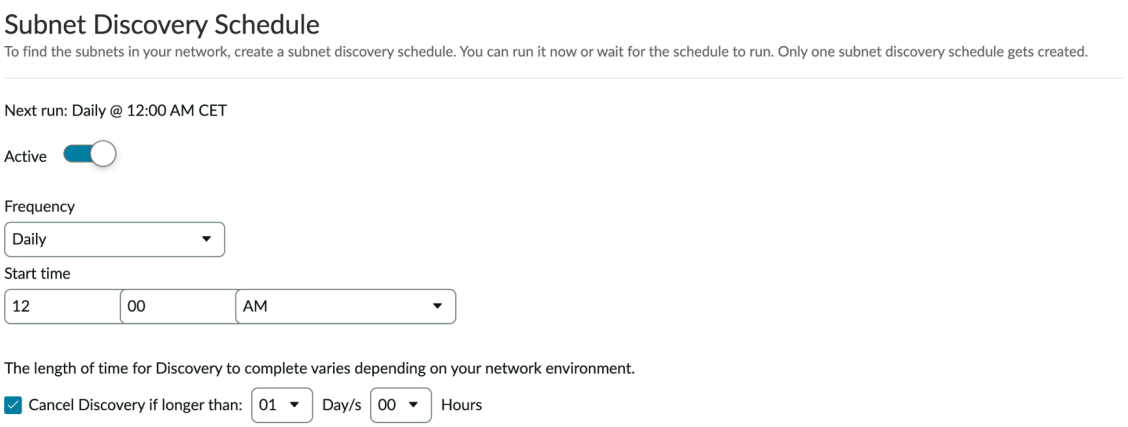

- **Active**: Enables or disables this schedule.
- **Frequency**: Select how often the schedule should run:
  - **Daily**: Run Discovery at a specified time each day.
  - **Weekly**: Run Discovery on one designated day of each week.
  - **Monthly**: Run Discovery on one designated day of the month. If the designated day doesn't occur in the month, the schedule doesn't run that month. For example, if you designate day 30, the schedule doesn't run in February.
  - **Once**: Run Discovery once on the designated date and time.
  - **Periodically**: Run Discovery every designated time. Use the **Repeat Interval** field to define the time in days, hours, minutes and seconds. The first discovery runs at the point in time defined in the **Starting** field. The subsequent discoveries run after each **Repeat Interval** period passes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Cancel Discovery if longer than**: Terminate the Discovery process if it exceeds the specified duration. Discovery might run for an excessive amount of time based on multiple factors, including the size of your network.

  Note: Any MID Server processing taking place when Discovery is canceled is allowed to complete. However, all output records that MID Servers haven't picked up are canceled.

b. Select one of the following options:

- **Finish and Run** saves the schedule and runs discovery immediately.
- **Finish** saves the schedule. Discovery then runs at the interval specified in the schedule.

6. View subnet discovery results.

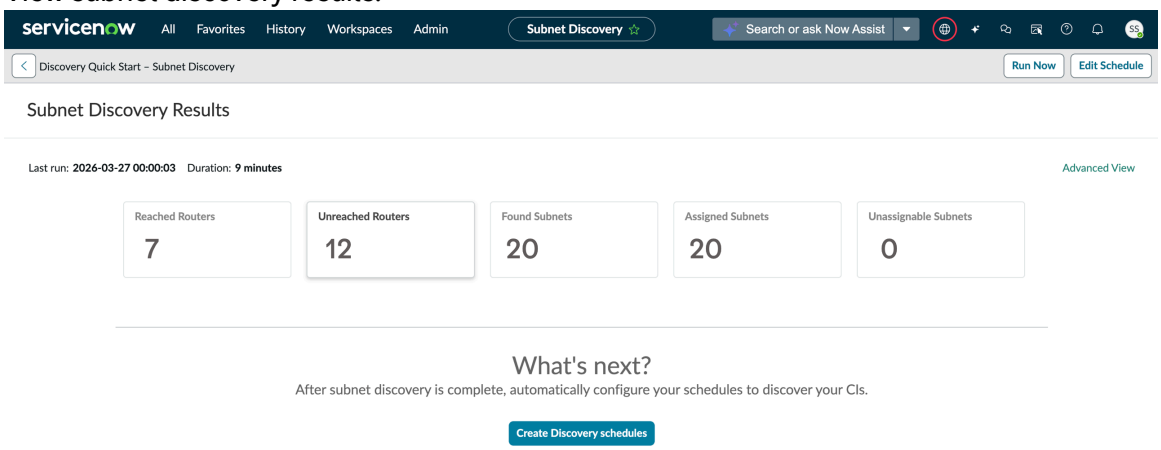

a. Select each tile for details.
If subnet Discovery did not yet run, each of these tiles shows **0**.

- **Reached Routers**: The number of routers that Discovery was able to access and discover.
- **Unreached Routers**: The number of routers that Discovery was not able to access.
- **Found Subnets**: The number of subnets that Discovery found.
- **Assigned Subnets**: IP subnets that were assigned to at least one MID Server.
- **Unassignable Subnets**: IP subnets that were not assigned to any MID Servers.

b. Select **Edit Schedule** to return to the **Select MID Servers** tab.
Resolve any issues with subnet discovery by returning to the Router Access page and making sure each credential is configured correctly. Also test router access for each credential.

c. Select **Advanced View** to see the Automation Status Set form.
This form displays detailed results for both subnet discovery and the automatic assignment of IP ranges.

Configure and schedule CI discovery

Using the credentials you provide, the Quick Start wizard finds all the configuration items (CI) for the IP ranges identified during subnet discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

You should perform this procedure after you complete Configure and schedule subnet discovery in the Discovery Quick Start.

Role required: discovery_admin

About this task

The goal of this phase is to create the schedules necessary to discover the CIs on each subnet. Discovery Quick Start does that automatically, and then runs these schedules using the credentials you provide. You can add additional credentials at any time. Optionally, you can define locations for each of the subnet routers, and Discovery creates separate schedules for each location.

You must be familiar with the types of credentials that you need for the CIs in your network. Each CI type requires different credentials. For example, if Discovery uses patterns instead of probes, the patterns might require applicative credentials.

Review the requirements for the CI types you intend to discover. See Data collected by ITOM Visibility, and then locate the specific type of device, computer, or application.

Procedure

1. From the Subnet Discovery Results page of the Discovery Quick Start, select **Create Discovery schedules**.

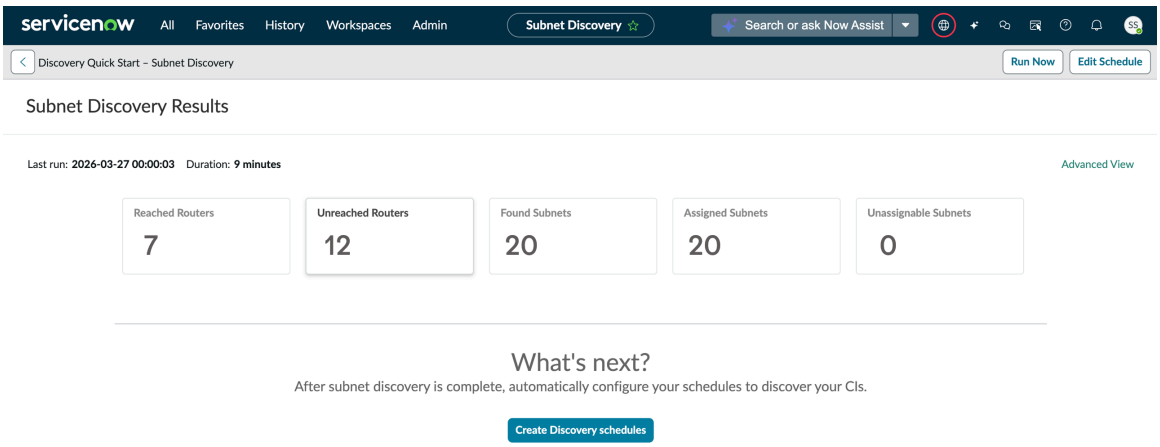

2. Create the credentials necessary to discover CIs in your network.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select the credential type you want to manage from the **View** drop-down list.
All the credentials currently configured for that CI type display in individual tiles.

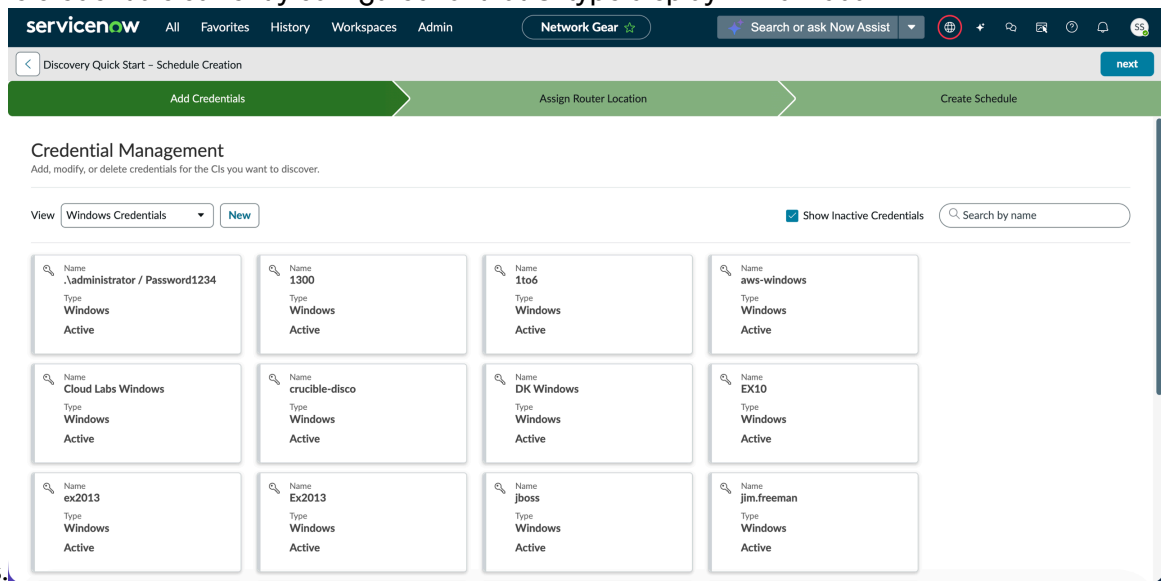

b. Optional: To add a credential, select **New**.
The Add Credential form displays, enabling you to specify your new credential.

c. Select **Next** to go to the **Router Location** phase.

3. Optional: In the **Assign Router Location** tab, assign a location to each router that subnet discovery found.
When you configure router locations, you have the option of creating a schedule for each location. If you do not assign locations, a single Discovery schedule is created for all your subnets. If you assign a time zone to a location, Discovery creates a schedule that runs in that location's time zone. If you do not assign a time zone to a location, the schedule for that location runs in the user's local time zone.

Note: A subnet can belong to multiple routers. If schedules are created by location, Discovery finds subnets managed by multiple routers using multiple schedules.

The list of gateway routers shows those devices not currently assigned to a location.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

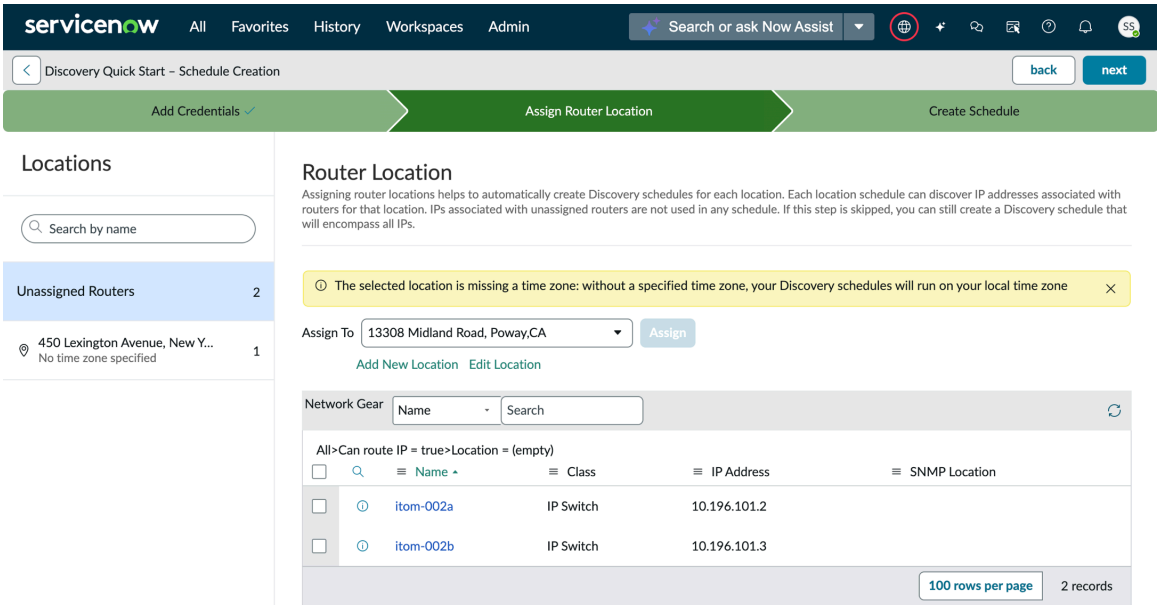

a. Select the check boxes for the routers you want to assign a location.

b. Select the location in the **Assign to** list.

c. If the location is not already in the list, select **Add New Location** and add the location.

d. Verify that the location has a time zone defined.
Timezones enable you to create a schedule for each location based on the time zone.

e. Select **Assign**.

> **Note:** Routers that are unassigned cannot be added to any schedules that are based on location.

f. Select **next**.

4. Create a CI Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Complete the schedule using the fields in the list.

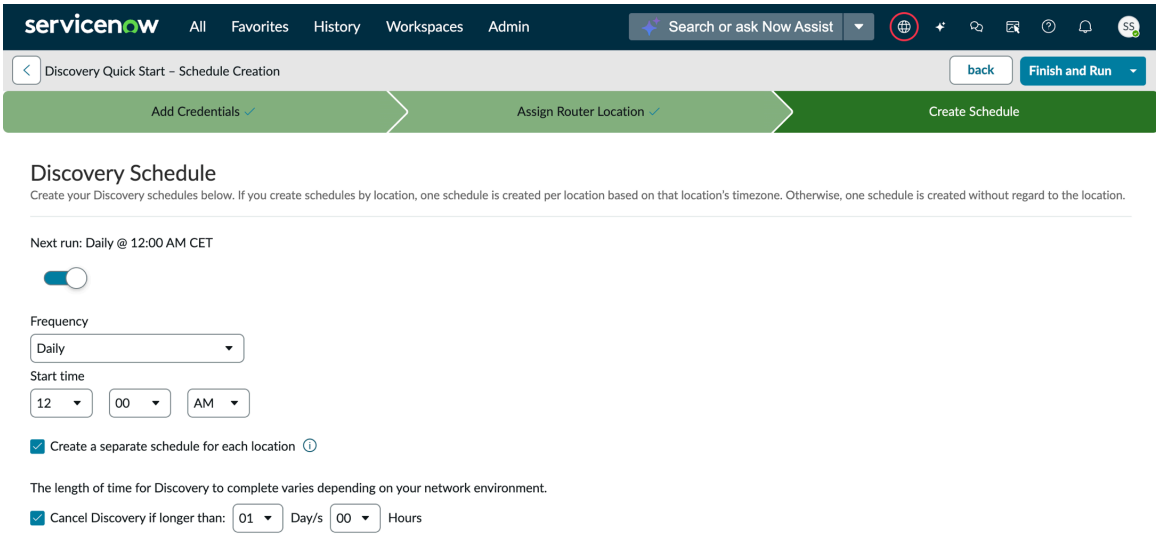

- **Active**: Enables or disables this schedule.
- **Frequency**: Select how often the schedule should run:
  - **Daily**: Run Discovery at a specified time each day.
  - **Weekly**: Run Discovery on one designated day of each week.
  - **Monthly**: Run Discovery on one designated day of the month. If the designated day doesn't occur in the month, the schedule doesn't run that month. For example, if you designate day 30, the schedule doesn't run in February.
  - **Once**: Run Discovery once on the designated date and time.
  - **Periodically**: Run Discovery every designated time. Use the **Repeat Interval** field to define the time in days, hours, minutes and seconds. The first discovery runs at the point in time defined in the **Starting** field. The subsequent discoveries run after each **Repeat Interval** period passes.
- **Create a separate schedule for each location**: Automatically create schedules for any router locations that you've configured. These schedules run on the day and time that you specify, but in their own local time zone. Clear this check box to create a single schedule for all your locations that runs in the user's local time zone. This option is only available if you have configured router locations.
- **Cancel Discovery if longer than**: Terminate the Discovery process if it exceeds the specified duration. Discovery might run for an excessive amount of time based on multiple factors, including the size of your network.

  Note: Any MID Server processing taking place when Discovery is canceled is allowed to complete. However, all output records that MID Servers haven't picked up are canceled.

b. Select one of the following options:

- **Finish and Run** saves the schedule and runs a discovery immediately.
- **Finish** saves the schedule. Discovery runs at the interval you specified in the schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Migrate from CAPI to Patterns

Migrate from Cloud API (CAPI)-based Cloud Discovery to Patterns-based Cloud Discovery. The task requires supported instance and few plugins. Migration works for Amazon Web Services (AWS) and Microsoft Azure. Administrator can perform this task after the initial instance is set up.

Before you begin

Ensure that your ServiceNow AI Platform has the following:

- The release version is Paris or later.
- The Discovery and Service Mapping Patterns (sn_itom_pattern) plugin.

Role required: discovery_admin

About this task

When the migration is triggered, the script disables the appropriate CAPI discovery steps for the cloud environment to prevent them being triggered during cloud discovery. Also, the script enables the appropriate cloud patterns for the environment that is used to run future cloud discovery. After the migration, discovery uses the latest patterns instead of the legacy discovery method, CAPI.

Note:

The CAPI-based discovery method is going to be deprecated in the future.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Migrate CPG to Pattern**.
2. On the CAPI to Pattern Migration page, read the instructions and select **Run Prerequisites Check**.
3. After the Prerequisite scan is passed, you can either choose as follows:
   - Migrate All—Migrate both AWS and Azure together.
   - Migrate AWS or Migrate Azure—Migrate each instance one at a time.
4. Select Migrate and confirm to proceed migration.

Note:

For any issues while migrating, see the troubleshooting steps documented in the CAPI to Pattern Migration: Procedure for switching from CAPI-based Cloud Discovery to pattern-based Cloud Discovery [KB0827153] article in the Now Support Knowledge Base.

Discovery properties

Discovery properties enable you to control several aspects of the horizontal discovery process.

For information about the Discovery properties that Agent Client Collector uses, see Using push-based and horizontal IP-based Discovery together.

Edit the following Discovery properties by navigating to **Discovery Definition** > **Properties**.

IP properties

Field	Description
glide.discovery.exclude_ip_sync_classes	Defines CI classes whose IP addresses shouldn't be substituted if the address

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	returned by Discovery doesn't match one of the devices' NICs. Use a comma-separated list to define multiple classes. By default, the system uses the management IP of a load balancer returned by Discovery in the CI record, rather than substituting it for the IP address of one of the load balancer's NICs. • **Type**: string • **Default value**: cmdb_ci_lb
glide.discovery.enforce_unique_ips	Enforce unique IP addresses: Ignores the IP address after Discovery encounters subsequent devices that use the same IP address. Each time a computer, printer, or network device with a valid IP address is discovered, the IP address field is cleared on any other device sharing that address. If inactive, stores the IP address for each device. • **Type**: true | false • **Default value** false

Network properties

Field	Description
glide.discovery.bgp_router_disable	BGP router exploration disable: Disables running the **SNMP – Routing** probe or pattern when discovering a router running the BGP protocol. This exploration is inactive by default because BGP routing tables are large. These routers typically operate at the edge of large networks, where further network discovery is not relevant. Set this value to false if your organization uses BGP routers as edge routers between relatively small networks. For example, between buildings on a single campus. If you must populate the [discovery_device_neighbors] table during horizontal layer 2 discovery of the bgp-enabled devices, set the **BGP router exploration disable** property to **No**. Notice that enabling this property can cause performance issues including out-of-memory issues on the MID Server. • **Type**: true | false • **Default value**: true

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
glide.discovery.debug.network_discovery	Network discovery debugging: Enables extensive logging of all Network Discovery activities on the instance. • **Type**: true | false • **Default value**: true • **Learn More**: Discovery monitoring and issue resolution
glide.discovery.discoverable.network.max.netmask.bits	Maximum netmask size for discoverable networks (bits): The maximum number of bits in a regular netmask for networks that are discovered by Network Discovery. A regular netmask is a netmask that can be expressed in binary as a string of ones followed by a string of zeroes. For example, 255.255.255.0 is regular and 255.255.255.64 is irregular. Regular networks are commonly expressed like: 10.0.0.0/24, which means a network address of 10.0.0.0 with a netmask of 255.255.255.0. Larger bit numbers mean networks with smaller numbers of addresses in them. For example, the network 10.128.0.128/30 has four addresses in it: one network address (10.128.0.128), one broadcast address (10.128.0.131), and two usable addresses (10.128.0.129 and 10.128.0.130). Small networks are commonly configured in network gear to provide loopback addresses or networks used strictly by point-to-point connections. Since these sorts of networks generally don't need to be discovered by Network Discovery, it would be useful to filter them out. By setting this property to a value of 1 through 32, you can limit the sizes of regular networks that are discovered. Setting it to any other value causes all networks to be discovered. Irregular networks are always discovered. The default value is 28, which means that regular networks with 8 or fewer addresses cannot be discovered. • **Type**: integer • **Default value**: 28
glide.discovery.network_owner_method	Network router selection method: This property controls the method used to decide (during Network Discovery) which router should be selected as the router to be associated with a given IP Network. The possible values are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **First Router**: The first router that discovers the network is associated • **Last Router**: The last router that discovers the network is associated • **Most Networks**: The router with the most attached networks is associated • **Least Networks**: The router with the least attached networks is associated • **Type**: choice list • **Default value**: Most Networks • **Additional options**: First Router, Last Router, Least Networks
glide.discovery.physical_interface_types	Physical interface types: A comma-separated list of interface types that are considered physical for the purposes of network discovery. If a router (or device capable of routing) has an interface of this type, the networks connected to that interface are considered locally connected to that device. The default interface types include Ethernet, 802.11, and Token Ring types. Interface type numbers are defined in the SNMP MIB-2, specifically in OID 1.3.6.1.2.1.2.2.1.3. • **Type**: string • **Default value**: 6,117,9,71,209
glide.discovery.switch_interface_types	Switch interface types: A comma-separated list of interface types that are considered. Interface type numbers are defined in the SNMP MIB-2, specifically in OID 1.3.6.1.2.1.2.2.1.3. Devices with any interface types that do NOT appear in this list are classified as routers, when they have routing. A complete list of the interface type numbers may be found on the IANA web site, in the section "ifType definitions". • **Type**: string • **Default value**: 7,8,9,26,53,62,69,71,78,115,117,209

Other Discovery properties

Field	Description
com.snc.use_external_credentials	Allows the use of external storage of credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: true
com.glide.codesigning.tracking.unsupported_script_tracking.enabled	Controls script tracking for code signing on Cloud Discovery. When set to true and code signing is enabled, ECC queue records inserted through unsupported scripts are not notarized, which can cause Cloud Discovery operations to fail. Set this property to **false** to enable Cloud Discovery on code-signed instances. • **Type**: true | false • **Default value**: true • **Learn more**: Configure Cloud Discovery for code-signed instances
com.glide.queue.discovery.cloud.sensors.processing_job_priority	Manages the background job's worker priority for Cloud-based Discovery jobs when using Event Framework. Increasing the value decreases the priority. • **Type**: integer • **Default value**: 100 **Note:** This behavior only applies when *discovery.use.event.processing* is set to **true**.
com.glide.queue.discovery.sensors.processing_job_priority	Manages the background job's worker priority for IP-based Discovery jobs when using Event Framework. Increasing the value decreases the priority. • **Type**: integer • **Default value**: 100 **Note:** This behavior only applies when *discovery.use.event.processing* is set to **true**.
discovery.event.pin.jobs	Enables event processors to execute on specific nodes. When set to **true**, jobs are assigned to nodes to distribute the work evenly. If set to **false**, jobs are randomly assigned, which may increase discovery throughput.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: false **Note:** This behavior only applies when *discovery.use.event.processing* is set to **true**.
discovery.event.separate.cloud.events	Reduces processing time by separating the *discovery.sensor* and *discovery.cloud.sensor* event queues for cloud and IP-based discoveries. When this property is set to **false**, event processing will only use the *discovery.sensor* queue. • **Type**: true | false • **Default value**: true **Note:** This behavior only applies when *discovery.use.event.processing* is set to **true**.
discovery.use.event.processing	Enables Discovery to use the Event Framework when processing jobs. • **Type**: true | false • **Default value**: false
glide.discovery.active_processes_filter	Active Processes Filter: Optimization for application dependency mapping. Filters the active processes returned by Discovery to only those that have a match in the Process Classification table. • **Type**: true | false • **Default value**: false
glide.discovery.adme.aggregation_interval	ADME - Aggregation interval (sec): How often to aggregate instantaneous data into chunks in seconds. It must be a minimum of 60 seconds and it must be a multiple of ADME Sampling interval. • **Type**: integer • **Default value**: 3600
glide.discovery.adme.base_dir_unix	ADME - Unix Base Dir: An existing directory on the target UNIX machines to be used as a workspace. Must be an absolute path to the directory.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: false
mid.powershell.target_base_dir	Powershell - Target Base Dir: A network share on the target Windows machines to be used as a workspace. • **Type**: string • **Default value**: admin$\temp
glide.discovery.adme.max_total_samples	ADME - Max total samples: The limit of samples being taken by the script. The script will stop running after this number of samples are taken. • **Type**: integer • **Default value**: 100
glide.discovery.adme.rolling_window_size	ADME - Rolling window size: How many recently aggregated chunks to keep. An aggregated chunk consists of ADME Aggregation interval divided by ADME Sampling interval samples. Only the newest chunks are kept. • **Type**: integer • **Default value**: 24
glide.discovery.adme.sampling_interval	ADME - Sampling interval (sec): How often to sample process and connection data in seconds. Must be a minimum of 5 seconds. • **Type**: integer • **Default value**: 120
glide.discovery.adm.update.applications	Controls if application CIs should be updated based on running process status. When Discovery rediscovers the server and notices that a previously running process is no longer active on cmdb_running_process record, Absent is set to **true**. In addition, on related application CI, install_status is updated from installed to absent. On subsequent scans, if the running process is back online on the related application CI, install_status is reverted to installed. **Note:** This behavior will only work for application CIs that are discovered using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: true, on instances where Discovery is first activated beginning with the Tokyo release. **Note:** On instances where Discovery was active before the Tokyo release, this property isn't created automatically during the upgrade. To enable this behavior on those instances, an administrator must add the `glide.discovery.adm.update.applications` property and set it to **true**.
glide.discovery.allow_loopback_adapters	Loopback adapter filter: Allows network loopback adapters on SNMP devices to be discovered instead of being filtered out during Discovery. When this property is set to **true**, network loopback addresses are discovered and added to the CMDB. • **Type**: true | false • **Default value**: false
glide.discovery.application_mapping	Application mapping: Enables the application mapping portion of Discovery **Important:** Disabling this property disables the creation of relationships between applications but doesn't disable the ADM probes and sensors, nor does it prevent process classification. • **Type**: true | false • **Default value**: true • **Learn More**: Application Dependency Mapping (ADM) for Discovery
glide.discovery.certs.cert_admin_user_id	Certificate Inventory and Management: Holds the user ID of the user who will be referenced in the certificate task and incident, created via Scheduled Job. • **Type**: string (40 characters maximum) • **Default value**: Certificate Administrator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
glide.discovery.certs.days_before_expiration_to_create_renewal_task	Certificate Inventory and Management: Number of days before certificate expiration that a renewal task should be created. • **Type**: integer • **Default value**: 60
glide.discovery.certs.enable_incident_creation_for_expired_certificates	Certificate Inventory and Management: Enables the Scheduled Job to create the incidents for expired certificates, for example, certificates whose validity is past current system date or time. • **Type**: true | false • **Default value**: false
glide.discovery.certs.enable_renewal_task_creation_for_discovered_certificates	Certificate Inventory and Management: Enables renewal task creation for all discovered certificates. By default, this property is enabled. • **Type**: true | false • **Default value**: true
glide.discovery.certs.slack_channel_id	Certificate Inventory and Management: Indicates the Slack channel ID where certificate notifications are sent. • **Type**: string • **Default value**: empty
glide.discovery.count.use_optimized_counter	Executes optimized code for updating discovery_status.started and discovery_status.completed counts. If set to **true**, business rules on discovery_status won't execute when "started" or "completed" is changed. There are two exceptions: **1.** When a request to increment starting counts gets triggered and discovery_status.state value is still in 'Starting'. **2.** When discovery_status.completed changes to the same value as discovery_status.started.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: true
glide.discovery.debug.ci_identification	CI identification debugging: If set to **true**, enables debug logging (into the CI Identification Log) for CI Identification. • **Type**: true | false • **Default value**: false
glide.discovery.discover_aws_ec2_host_metadata	When doing IP-based discovery against a given host, also run probes that retrieve AWS EC2 metadata. • **Type**: true | false • **Default value**: true
glide.discovery.discover_azure_host_metadata	When doing IP-based discovery against a given host, also run probes that retrieve Azure metadata. • **Type**: true | false • **Default value**: true
glide.discovery.discover_software	Discover software packages: Enable the discovery of software packages. • **Type**: true | false • **Default value**: true • **Learn More**: General software package discovery
glide.discovery.domain.name.nbt	Set OS domain name by NBT or WMI. If set to **true**, Windows domain name is set by NBT. Otherwise it's set by WMI. • **Type**: true | false • **Default value**: true
glide.discovery.enable_adme	ADME: Enable enhanced ADM probe. If set to **true**, the ADM Enhanced probe is triggered and only fall back to the ADM probe as needed. • **Type**: true | false • **Default value**: false
glide.discovery.enable_file_tracking	Enable configuration file tracking as part of the Pattern based Horizontal Discovery. If set to **true**, the system enables you to specify paths of configuration files for applications that are discovered by patterns. You can

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	track and compare tracked configuration files. • **Type**: true | false • **Default value**: true
glide.discovery.enable_mac_address_verification	Enables MAC address verification to determine if a device has changed its IP address during a discovery. If set to **true**, Discovery passes the MAC address of a device to the MID Server. Probes use this to determine whether the IP address changed during discovery. If the device has changed IP addresses, then Discovery stops processing that IP address and updates the Discovery log with a warning message. Discovery doesn't stop processing other IP addresses. • **Type**: true | false • **Default value**: false Open the System Properties [sys_properties] table to find this property.
glide.discovery.enable.software_filter	Windows software filter: Enable Windows software filtering for Discovery. • **Type**: true | false • **Default value**: true
glide.discovery.enable.software_simplify	Enables cleanup for software's name and version. If set to **true**, it removes any appended version or commonly found phrases from the name and removes leading 0s in the version. If set to **false**, raw data is persisted in the table. Use with the **glide.discovery.enable.software_simplify_sccm** property to avoid software data discrepancy between Discovery and SCCM. • **Type**: true | false • **Default value**: true **Note:** Once the property is set to **false**, to disable the cleanup, duplicate records are generated for a given software. Name and Version are the primary identifiers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
glide.discovery.enable.software_simplify_sccm	Enables cleanup for SCCM software. If set to **true**, Discovery changes the Name of the software to be the same as SCCM. If set to **false**, Discovery doesn't clean up the Name. Use with the **glide.discovery.enable.software_simplify** property to avoid software data discrepancy between Discovery and SCCM. • **Type**: true | false • **Default value**: false
glide.discovery.enable_ssm	Enables AWS SSM-based discovery. • **Type**: true | false • **Default value**: false
glide.discovery.ssm.enable_windows	Enables AWS SSM-based discovery on Windows machines as a high-level system user that can run any command. • **Type**: true | false • **Default value**: false
glide.discovery.file_tracking.max_file_size	File Tracking: The maximum file size for tracked configuration file content. • **Type**: integer • **Default value**: 500000
glide.discovery.file_tracking.max_files_per_ci	File Tracking: The maximum number of tracked configuration files per CI. • **Type**: integer • **Default value**: 50
glide.discovery.file_tracking.change_limit_count	File Tracking: Number of changes allowed on a tracked configuration file in the defined time window. • **Type**: integer • **Default value**: 4
glide.discovery.file_tracking.change_limit_days	File Tracking: Time window (in days) that limits the number of changes on a tracked configuration file. • **Type**: integer • **Default value**: 7
glide.discovery.fqdn.regex	DNS Host Name And Domain Name Regex: The default parsing of FQDN (Fully Qualified

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Domain Name) is to pick the first name separated by dots as the host name and the rest of the names as the domain name. For example, "machine1.testlab.service-now.com" has host name of "machine1" and domain name of "testlab.service-now.com". The property allows regex with two capturing groups with the first group representing the host name and the second group the domain name. • **Type**: string • **Default value**: ^([^.]+)\.((?:[^.]+\.)+[^.]+)$
glide.discovery.hostname.always_update	Always update host name. If set to **true**, discovery always updates the host name with the most recently discovered value contingent upon the source being trusted. Note that this may result in hand-entered values being overwritten. This property is only applicable for Probe-based Discovery. • **Type**: true | false • **Default value**: true
glide.discovery.hostname.case	Host name case. If **Lower case** is selected, always translate the host name into lower case. If **Upper case** is selected, always translate the host name to upper case. If **No change** is selected, leave the host name intact. This primarily affects host names discovered with NETBIOS, though some non-standard DNS systems may also return some or all of the name in upper case. • **Type**: choice list • **Default value**: Lower case • **Additional options**: Upper case, No change
glide.discovery.hostname.dns_nbt_trusted	DNS or NetBIOS is trusted host name source: If set to **true**, trust the device name discovered via DNS or NetBIOS. If checked, CI's host name found via DNS or NBT is used. **Note:** For Linux, AIX, and Solaris patterns, do not use the DNS name as the host name. • **Type**: true | false • **Default value**: true
glide.discovery.hostname.include_domain	Includes domain name in host name: If set to **true**, include the domain name as part of

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	the host name. For example, "bosco.service-now.com" instead of "bosco". **Note:** For AIX and Solaris patterns, do not use the DNS name as the host name. • **Type**: true | false • **Default value**: false
glide.discovery.hostname.snmp_trusted	SNMP is trusted host name source: If set to **true**, trust the device name discovered via SNMP. If checked, any device name found via SNMP is used instead of the name found by a reverse DNS lookup. • **Type**: true | false • **Default value**: false (For new Discovery customers starting in Brazil, default value is true.)
glide.discovery.hostname.ssh_trusted	SSH is trusted host name source: If set to **true**, trust the device name discovered via SSH. If checked, any device name found via SSH is used instead of the name found by a reverse DNS lookup. **Note:** For Linux, AIX, and Solaris patterns, do not use the DNS name as the host name. • **Type**: true | false • **Default value**: false (For new Discovery customers starting in Brazil, default value is true.)
glide.discovery.hostname.wmi_trusted	WMI is trusted host name source: If set to **true**, trust the device name discovered via WMI. If checked, any device name found via WMI is used instead of the name found by a reverse DNS lookup. • **Type**: true | false • **Default value**: false (For new Discovery customers starting in Brazil, default value is true.)
glide.discovery.ip_based.active	This allows scans on all IP Address of the same device in the same Discovery schedule to go through full cycle. This is set to **true** for probe-to-pattern migrated instances, Discovery doesn't attempt extra IP Address detection.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: true (For probe-to-pattern migrated instances.)
glide.discovery.ip_service_affinity	IP service affinity: If set to **true**, IP service affinity is enabled. IP service affinity allows Discovery to remember the last port of the IP address that was discovered. **Note:** You must enable ip_service_affinity records if you're using multiple SSH ports. • **Type**: true | false • **Default value**: false (For new Discovery customers starting in Brazil, default value is true.) • **Learn More**: Check IP service affinity for Discovery and Orchestration
glide.discovery.L3_mapping	Map servers and network devices to routers and layer-3 switches If the "L3 mapping" property is enabled, it maps servers and network gears to its associated routers and layer-3 switches. • **Type**: true | false • **Default value**: true • **Learn More**: Network discovery
glide.discovery.log_debug	Controls debug logging for Discovery. When set to **false** (default), Discovery debug statements aren't logged. Setting the property to **true** enables debug logging for Discovery, allowing detailed debug statements to be written to the logs. • **Type**: true | false • **Default value**: false
glide.discovery.log_message_length	Limits the maximum message length that is displayed in Discovery Log table. A value of 0 or any negative number disables this limit. • **Type**: integer • **Default value**: 200 • **Learn More**: Logs for horizontal discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
glide.discovery.mark_esx_servers_as_absent	If this property is set to true, the ESX Servers are marked as absent instead of retired if the ESX server isn't found in the payload. • **Type**: true | false • **Default value**: false
glide.discovery.max_concurrent_invocations_per_schedule	Maximum concurrent invocations per schedule: Prevents an unbounded number of invocations from inundating the system when a schedule takes longer than the time between invocations. The value is an integer defining the maximum number of automated invocations of the same schedule that may proceed at one time. If the limit has been reached subsequent scheduled invocations are canceled. The default value is 3. A value of 0 or any negative number disables this restriction. • **Type**: integer • **Default value**: 3 **Note:** This property does not apply to schedules that have a 'Run after' configuration set to 'Even if canceled'.
glide.discovery.network_discovery.functionality	Networks discovery functionality: The Functionality used to discover networks. Usually this should be **SNMP only**. • **Type**: string • **Default value**: SNMP only
glide.discovery.perf.metrics.enable_collection	Whether to collect Discovery performance metrics. • **Type**: true | false • **Default value**: true
glide.discovery.perf.metrics.rollup_by_build	Whether to aggregate individual Discovery performance metrics into a roll-up table that groups by discovery build. • **Type**: true | false • **Default value**: true
glide.discovery.perf.metrics.rollup_by_status	Whether to aggregate individual Discovery performance metrics into a roll-up table that groups by discovery status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: true
glide.discovery.perf.metrics.rollup_by_target	Whether to aggregate individual Discovery performance metrics into a roll-up table that groups by target IP address. • **Type**: true | false • **Default value**: true
glide.discovery.populate_software_instance_table	If set to **true**, Discovery populates these tables for software discovery: `cmdb_software_instance` and `cmdb_sam_sw_install` • **Type**: true | false • **Default value**: false Use this property only if you were using Software Asset Management with Discovery in pre-Jakarta releases and want to keep populating the `cmdb_software_instance` table. Open the System Properties [sys_properties] table to find this property.
glide.discovery.retire_stale_accounts	When set to **true**, Discovery automatically removes stale service account references from the Discovery Config - Logical Datacenters [cmp_discovery_ldc_config] table when previously discovered cloud service accounts have been retired or deleted. This prevents Discovery from attempting to use invalid service accounts and helps avoid canceled Discovery runs caused by missing service account references. This property applies to Cloud Discovery schedules that use master/member account relationships. • **Type**: true | false • **Default value**: false
glide.discovery.roundingInterval.cpu	CPU speed rounding: Enter the number to round the CPU speed to. The units are in MHz. • **Type**: integer • **Default value**: 1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
glide.discovery.roundingInterval.ram	Memory rounding: Enter the number to round the computer RAM to. The units are in MB. • **Type**: integer • **Default value**: 1
glide.discovery.save_pattern_log	Controls whether pattern logs are saved during discovery execution. This property helps enhance performance by reducing database interactions. • **Type**: string • **Default value**: true If set to **false**, the system saves the pattern log only when it is important to display information (for example, when a failure occurs). This further improves performance by minimizing unnecessary saves. **Note:** Regardless of the property value, pattern logs are always saved if the pattern status is failure.
glide.discovery.schedule.autogenerate.debug	Enables detailed logging for troubleshooting. If set to **true**, the job's detailed logging output will appear in the System Log [sys_log] table with the source DiscoveryIPAMIntegration. • **Type**: true | false • **Default value**: false • **Learn more**: IPAM Discovery integration
glide.discovery.schedule.autogenerate.group	Groups auto-generated IPAM Discovery schedules by a specific tag or attribute. For example, entering "location" as the value would group the schedules by their associated locations. • **Type**: string • **Default value**: empty • **Learn more**: IPAM Discovery integration
glide.discovery.schedule_ipv4_limit	Defines the maximum number of IPv4 addresses that can be included in a single auto-generated IPv4 Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: integer • **Default value**: 65,536 • **Learn more**: IPAM Discovery integration
glide.discovery.schedule_ipv6_limit	Defines the maximum number of IP6 addresses that can be included in a single auto-generated IPv6 Discovery schedule. • **Type**: integer • **Default value**: 20000 • **Learn more**: IPAM Discovery integration **Note:** On instances below ZP3 and YP9, the maximum number of IPv6 addresses per schedule must not exceed 5,000, which is also the default value in Yokohama and earlier releases.
glide.discovery.schedule_monitoring.dry_run	Determines whether the job will run in "dry run" mode. If set to **true**, the job will only log the decision in the Discovery Log but won't change the status's state, otherwise the job will change the state of the status if needed. • **Type**: true | false • **Default value**: false
glide.discovery.schedule_monitoring.schedule_type_to_monitor	Determines the type of statuses to track. The value can be found on the "Discover" attribute on the schedule. • **Type**: string • **Default value**: CIs
glide.discovery.schedule_monitoring.stale_check_seconds	Determines the number in seconds in which a status is considered stale. A stale status is a status that was not updated for the last X seconds. • **Type**: integer • **Default value**: 1200 (20mins). If nothing is configured, the default is 2x the interval of the job.
glide.discovery.sensors.fire_ssh_probe	In case the VIPs (Load Balancer Virtual Services) didn't return as part of the SNMP Probe payload, fire the SSH Probe. Select this option to use the SSH port on the load balancer, rather than just using SNMP.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: true | false • **Default value**: true
glide.discovery.sensors.netapp_native_7mode	For 7-Mode NetApp storage servers, use native discovery instead of SMI-S.
glide.discovery.sensors.netapp_native_cluster_mode	For Cluster Mode NetApp storage servers, use native discovery instead of SMI-S.
glide.discovery.sensors.save_attachments	Determines whether discovery sensors save ECC queue attachments after successful sensor processing. When this property is **true**, sensors preserve the attachments. When it's **false**, sensors delete the attachments after they finish processing. Preserving attachments lets you inspect the data that sensors return. • **Type**: true | false • **Default value**: true
glide.discovery.shazzam_ranges_json	Use JSON for IP ranges in Shazzam: If set to **true**, Discovery encodes Shazzam's IP ranges as JSON, dramatically reducing the payload size. • **Type**: true | false • **Default value**: true
glide.discovery.shazzam_simplify_ranges	If set to **true**, Shazzam detects IPs included in multiple ranges of the same Discovery schedule and removes duplicates from all but one range. Deduplication is limited to schedules with a single behavior. Schedules with multiple behaviors work, but duplicate IPs aren't removed. • **Type**: true | false • **Default value**: true
glide.discovery.software_filter_keys	Windows software filters: If Windows software filtering is enabled, when the name of a discovered software contains any of the comma-separated values, it's filtered out. • **Type**: string • **Default value**: Language Pack, Hotfix Security Update

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
glide.discovery.software_sccm_managed	Windows software is SCCM managed: If set to **true**, Discovery won't populate software for computer CIs also managed by SCCM. • **Type**: true | false • **Default value**: false • **Learn More**: Discovery and SCCM together
glide.discovery.use_cmdb_identifiers	CMDB Identifiers: If set to **true**, identification and reconciliation will be handled by the CMDB API instead of through the old Discovery implementation. • **Type**: true | false • **Default value**: true
glide.discovery.use_probe_results_cache	Use probe results cache: If set to **true**, the cache is checked to see if the results of the probe must be processed by a sensor. It only needs to be processed if the results have changed from the last discovery run. • **Type**: true | false • **Default value**: false • **Learn More**: Create or modify a probe **Note:** In previous releases, this property was **true** by default. The value in this property isn't changed during an upgrade to Kingston.
glide.discovery.virtual_interface_types	Virtual interface types: List of interface types (comma-separated) that can be considered "virtual" for the purposes of network discovery. In other words, if a router (or device capable of routing) has an interface of this type, the networks connected to that interface are considered virtually connected to that device. The default interface types include the propVirtual type. Interface type numbers are defined in the SNMP MIB-2, specifically in OID 1.3.6.1.2.1.2.2.1.3. • **Type**: integer • **Default value**: 53
glide.discovery.warn_minor_version	Warn on Minor Version Mismatch: If set to **true**, warnings are logged when

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	minor_version mismatches are detected during Discovery sensor processing. • **Type**: true | false • **Default value**: false
glide.ecc_agent.validated.override	An override for the ecc_agent validated field: If set to true, ecc_agent records will always be treated as valid. • **Type**: true | false • **Default value**: false Navigate to **System Properties** > **All Properties** to find this property.
glide.eccprobe.awsrestprobe.max_retries	Set the maximum number of times the AWS request will be retried if a rate limit exceeded error occurs. Maximum value is 10. Note that retries use exponential backoff (2^x). Setting negative values is equivalent to setting the maximum value. • **Type**: integer • **Default value**: 4
glide.eccprobe.awsrestprobe.max_wait	Set the maximum total amount of time (seconds) the AWS request attempts to retry. Maximum value: 2046 seconds. Setting negative values is equivalent to setting the maximum value. • **Type**: integer • **Default value**: 30
glide.eccprobe.max_queued_probes_per_run	Set the maximum number of probes a node can pick up during a run. • **Type**: integer • **Default value**: 20
glide.eccprobe.node_agent_id	ECC queue entries with "mid.server.${VALUE}" will be picked up and processed by the nodes. This must be unique and not the same as any MID server name. • **Type**: string • **Default value**: NODE_AGENT
sn_disco_firewall.default.rule.task.policy.owner.group	Firewall Audits and Reporting: Indicates the name of the assignment group that owns or is assigned to firewall rule tasks

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: string (40 characters maximum) • **Default value**: Task Owner Group Name
sn_disco_firewall.default.rule.task.approval.group	Firewall Audits and Reporting: Indicates the name of the approval group that is used to approve the firewall rule tasks. • **Type**: string (40 characters maximum) • **Default value**: Approval Group Name
sn_disco_certmgmt.enable_csr_generation	Enables automatic CSR generation on the Certificate Inventory and Management **Request New Certificate** and **Request New Ceriticate (Automated)** forms. This property also controls whether the PKCS12 certificate format is available, because a PKCS#12 key store requires a private key that is generated with the CSR. This feature is disabled by default for security. • **Type**: true | false • **Default value**: false
sn_disco_certmgmt.private_key_size	Changes the private key size when using automatic CSR generation on the Certificate Inventory and Management **Request New Certificate** and **Request New Ceriticate (Automated)** forms. This property accepts values: 512 | 1024 | 2048 | 4096. • **Type**: integer • **Default value**: 2048
sn_disco_workspace.discovery_schedule_anomaly_detection_disabled	Discovery Admin Workspace: Disables the anomaly detection feature. When set to **true**, the system doesn't detect anomalies in Discovery schedules. • **Type**: true | false • **Default value**: false ⚠ **Warning:** Turning off this feature significantly impacts the Discovery experience across multiple pages, charts, and tables.
sn_disco_workspace.discovery_schedule_anomaly_detection_method	Discovery Admin Workspace: Determines the approach used for anomaly detection in Discovery schedules. The machine learning (ML)-based approach uses historical scan data and AI algorithms to detect anomalies

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	and outliers. The Stats-based approach uses historical scan data and statistical standard deviation metrics to detect anomalies and outliers. • **Type**: choice list • **Default value**: ML-based
sn_disco_workspace.discovery_schedule_anomaly_discovery_type	Discovery Admin Workspace: Determines the type of Discovery schedules for which anomaly detection is enabled, IP-based (CIs) and/or Cloud Discovery (Cloud Resources). To disable anomaly detection for a specific type of discovery, exclude it from the value. For example, remove "Cloud Resources" from the string if you only want to detect anomalies for IP-based Discovery schedules. • **Type**: string • **Default value**: CIs, Cloud Resources
sn_disco_workspace.discovery_schedule_anomaly_sensitivity	Discovery Admin Workspace: Adjusts the sensitivity for ML-based Discovery schedule anomaly detection. Changing the sensitivity can impact detection accuracy. Lower sensitivity results in more anomalies reported, which could increase false positives. Higher sensitivity results in fewer anomalies reported, which could increase false negatives. • **Type**: choice list • **Default value**: low
sn_disco_workspace.discovery_schedule_cis_discovered_anomaly_threshold	Discovery Admin Workspace: Sets the threshold percentage for detecting anomalies in the number of configuration items (CIs) discovered during a Discovery schedule execution. A Discovery status is marked as anomalous if the CI count decreases from the average by more than this threshold. • **Type**: integer • **Default value**: 20
sn_disco_workspace.discovery_schedule_error_anomaly_threshold	Discovery Admin Workspace: Sets the threshold percentage for detecting anomalies in Discovery schedule error counts. A Discovery status is marked as anomalous if the total error count deviates from the average error count by more than this threshold.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: integer • **Default value**: 20
sn_disco_workspace.discovery_schedule_runtime_anomaly_threshold	Discovery Admin Workspace: Sets the threshold percentage for detecting anomalies in Discovery schedule execution time. A Discovery status is marked anomalous if the run time deviates from the average run time of the schedule by more than this threshold. • **Type**: integer • **Default value**: 20
sn_disco_workspace.discovery_stats_time_scale	Discovery Admin Workspace: Configures the time scale for all relevant charts in Discovery Admin Workspace. Setting the time scale enables you to view the health of your Discovery schedules over your preferred time span. • **Type**: choice list • **Default value**: 7
sn_disco_workspace.enable_auto_create_schedules	Discovery Admin Workspace: Enables automatic creation and management of Discovery schedules based on IPv6 network infrastructure data in IPAM. When set to **true**, Discovery uses your IPAM data to generate schedules that target IPv4 and IPv6 networks without requiring manual schedule configuration. • **Type**: true | false • **Default value**: false
sn_disco_workspace.enable_error_framework	Discovery Admin Workspace: Controls whether the Error Framework is active for Discovery. When set to **true**, Discovery uses the Error Framework to capture and display errors in Discovery Admin Workspace. When set to **false**, Discovery uses the legacy Diagnostics experience. • **Type**: true | false • **Default value**: true
sn_disco_workspace.visibility_diagnostic_utility.enabled	Discovery Admin Workspace: Enables a Virtual Agent to retrieve diagnostic information for troubleshooting issues with Discovery. Setting this property to **true**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	displays a Virtual Agent on the **Diagnostics** page. • **Type**: true | false • **Default value**: false

Discovery IP address configuration

Use one or more of these methods in any combination to define the network or network segment for Discovery to query. You can include or exclude specific IP ranges from your query.

Note: If you don't know the IP addresses in the network, run Network discovery first to determine the IP networks. Then, convert the IP networks into IP address range sets.

If you have integrations which populate sys_metadata and sys_update_xml tables, make sure to clear the update and metadata records after the discovery_range_item or discovery_range_item_ip import occurs.

Important: IPv6 addresses aren't supported for IP address range or IP network.

There are three types of IP collections:

IP collection type	Description
IP address list	Use IP address lists to add individual addresses to query. These addresses aren't included in any existing IP range or IP network. You can enter the IP address of the device or a host name (DNS name). If you enter a host name, it must be mapped to an IP address.
IP address range	You can define arbitrary ranges of IP addresses to query. This process is a good way to include selected segments of a network or subnet. However, Discovery has no way of knowing if the IP range includes addresses for private networks or broadcast addresses, and scans all the addresses in the range. If the network and broadcast addresses are included, then the results are inaccurate. Discoveries configured to detect IP networks are more accurate than discoveries configured for IP address ranges. Only those IP addresses in your range that are reserved for manageable devices on the public network should be included. **Note:** To avoid performance issues, limit Discovery schedules to a maximum range of /16 (approximately 65,536 IPs). For better performance, consider splitting the range into smaller segments.
IP network	You can also scan an entire IP network. An IP network includes the range of available IP addresses in that network. The scan also includes the network address (the lowest address in the range) and the broadcast address (the highest address in the range). After you run network discovery, convert the IP networks that were found into range setsfor use in discovering other devices. IP networks are represented in CIDR notation. Examples of CIDR notation include:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP collection type	Description
	• 192.168.0.0/24 • 10.11.128.192/26 Discovery will not scan the network or broadcast addresses for a network. The equivalent ranges for the two networks are: • 192.168.0.1 – 192.168.254 • 10.11.128.193 – 10.11.128.254 To avoid introducing errors from broadcast addresses, configure Discovery to exclude them by using IP network definitions in CIDR notation. Discovery automatically avoids scanning network and broadcast addresses when using IP networks. Including broadcast addresses in Discovery scans can lead to significant errors in the data, such as redundant or invalid device entries. By excluding broadcast addresses, Discovery helps prevent these errors. This built-in control makes IP networks the best method of defining which IP address ranges to query.

After you define your IP collections, use Discovery generic attributes to automatically set field values on CIs discovered within a schedule, range set, or IP address range. For example, assign different locations to CIs based on the IP range they were discovered in. For more information, see Define CI field attributes.

For more information about viewing and managing your IP collections, see Discovery Admin Workspace IP inventory.

Related topics

- IP address selection properties
- Add IP data
- Add IP data to a range set
- Create a Quick IP range for a Discovery schedule
- Import IP ranges into Discovery schedules with import sets
- Exclude IP ranges from a Discovery range set
- Use Global Excludes List for IP addresses and ranges

IP address selection properties

You can use system properties to control the selection of IP address for specified configuration item (CI) classes.

Use these properties to determine whether Discovery should replace the IP address in a device's CI record. Replacement occurs when the returned address doesn't match a network interface (NIC) on the device.

This matters for devices whose management IP differs from the IP addresses associated with one or more NICs. Because the management IP is used in the Discovery schedule, that is the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

address Discovery returns. Use these properties to determine which IP address to use for CIs of any class.

Field	Description
glide.discovery.exclude_ip_sync_classes	Defines CI classes whose IP addresses shouldn't be substituted if the address returned by Discovery doesn't match one of the devices' NICs. Use a comma-separated list to define multiple classes. By default, the system uses the management IP of a load balancer returned by Discovery in the CI record, rather than substituting it for the IP address of one of the load balancer's NICs. • **Type**: string • **Default value**: cmdb_ci_lb
glide.discovery.enforce_unique_ips	Enforce unique IP addresses: Ignores the IP address after Discovery encounters subsequent devices that use the same IP address. Each time a computer, printer, or network device with a valid IP address is discovered, the IP address field is cleared on any other device sharing that address. If inactive, stores the IP address for each device. • **Type**: true | false • **Default value** false

Note: After you define your IP collections, you can use Discovery generic attributes to automatically set field values on CIs discovered within a schedule, range set, or IP address range. For example, you can assign different locations to CIs based on which IP range they were discovered in. For more information, see Define CI field attributes.

Add IP data

Use the New IP data dialog in the IP inventory page to add IP addresses, ranges, networks, or range sets to your discovery configuration.

Before you begin

Confirm the following:

- Discovery Admin Workspace v1.19.0 is installed.
- The ServiceNow AI Platform is running on the Brazil, Australia, or Zurich release starting with Patch 8.

Role required: discovery_admin

About this task

The New IP data dialog is shared across the **IP Address Lists**, **IP Ranges**, **IP Networks**, and **Discovery Range Sets** tabs. When you open the dialog from a tab, the check box for that tab is selected automatically. The system removes duplicate entries automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For collection types and format details, see Discovery IP address configuration. To view and manage your entries, see Discovery Admin Workspace IP inventory.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules** > **IP-based discovery**.
2. Select **View inventory** from the IP inventory section.
3. Select the tab that contains the IP data to add, such as **IP Address Lists**.
4. Select **New**.
   A New IP data dialog displays.

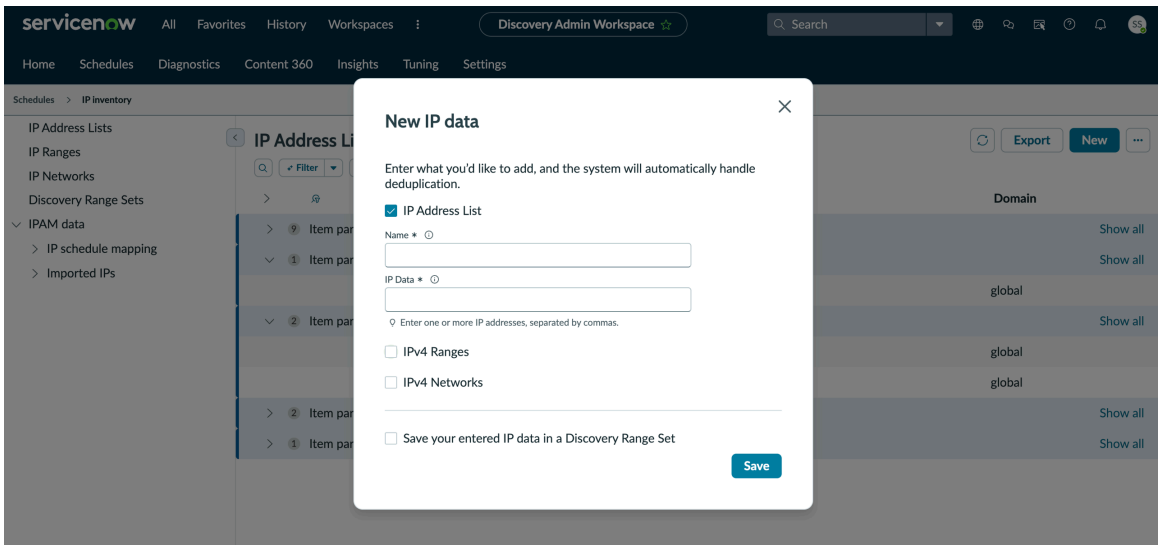

5. Enter the required information for each IP data type you're adding.

IP data type	What to enter
IP address list	Enter a name, then enter one or more IP addresses separated by commas.
IP ranges	Enter each range as startIP-endIP, separated by commas. For example: `10.0.0.1-10.0.0.254`
IP networks	Enter each network in CIDR notation, separated by commas. For example: `192.168.1.0/24`
Discovery range set	Enter a name, then enter any mix of addresses, ranges, or networks, separated by commas.

6. **Optional:** Select the check box for each additional IP data type to add.

   Note: If you open the New IP data dialog from the **Discovery Range Sets** tab, you can add data only as a range set.

7. For an IP range, IP network, or range set, select the **Active** toggle switch to turn the entry on or off.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. **Optional:** To save your entered data in a range set, select the check box, then either search for an existing range set or select **Create a new Range Set** and enter a name.
9. Select **Save**.
   The New IP data dialog closes.
10. Select the **Refresh** icon ( ) to view the new data in the IP Address Lists [discovery_range_item_ip] table.

Related topics

Add IP data to a range set

Add IP data to a range set

Use the IP inventory page in Discovery Admin Workspace to add IP ranges or networks to a new or existing range set.

Before you begin

Confirm the following:

- Discovery Admin Workspace v1.19.0 is installed.
- The ServiceNow AI Platform is running on the Brazil, Australia, or Zurich release starting with Patch 8.

Role required: discovery_admin

About this task

The **Add to Range Set** action is available on the **IP Ranges** and **IP Networks** tabs of the IP Inventory page in Discovery Admin Workspace. Select it to add existing IP data to a new or existing range set. The system removes duplicate entries automatically.

Note: To avoid performance issues, limit Discovery schedules to a maximum range of /16 (approximately 65,536 IPs). For better performance, consider splitting the range into smaller segments.

For collection types and format details, see Discovery IP address configuration. To view and manage your range sets, see Discovery Admin Workspace IP inventory.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules** > **IP-based discovery**.
2. Select **View inventory** from the IP inventory section.
3. Select the tab that contains the IP data to add, such as **IP Ranges**.
4. Select the check box for each row to add.
5. Select **Add to Range Set**.
   The Add to Range Set dialog displays.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

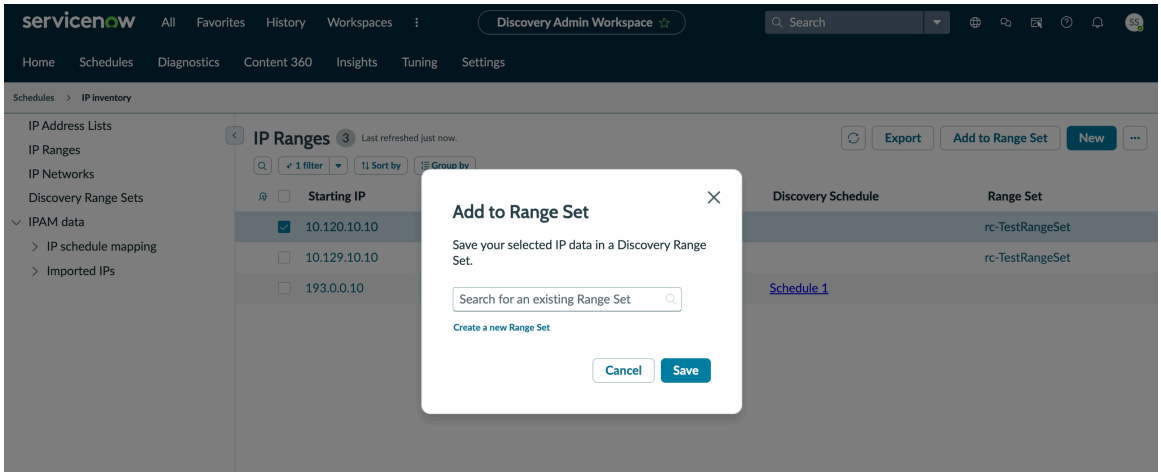

6. Search for and select an existing range set.
   To create a range set instead, select **Create a new Range Set** and enter a name.
7. Select **Save**.
   The Add to Range Set dialog closes.
8. Select the **Refresh** icon ( ) to view the new data in the IP Ranges [discovery_range_item] table.

Related topics

Add IP data

Create a Quick IP range for a Discovery schedule

Quick ranges enable administrators to define IP addresses to scan in a single comma-delimited string without creating separate records.

Before you begin

Only MID Servers that are up and validated are used with quick ranges. The MID Servers must specify the Discovery application (or ALL applications) and have IP ranges configured if you use the auto-select feature on the Discovery schedule.

Role required: discovery_admin

About this task

You can enter IP addresses in one of the following formats:

- An IP range defined by a slash and the number of bits in the subnetwork. For example, the string 10.10.10.0/24 scans 24 bits of IP addresses from 10.10.10.0 to 10.10.10.254.
- An IP range defined by a dash. For example, the string 10.10.11.0-10.10.11.165 scans the IP addresses from 10.10.11.0 to 10.10.11.165.
- A comma-separated list of specific IP addresses. For example, the string 10.0.2.1,10.0.2.15 scans the IP addresses 10.0.2.1 and 110.0.2.15.

Note: IPv6 address ranges and networks aren't supported and will be ignored. Any other entries that can't be interpreted will also be ignored.

Procedure

1. Select the **Quick Ranges** related link on the Discovery Schedule form.
2. Enter the IP ranges and specific IP addresses to scan.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Make Ranges**.

Note: The Quick Range interface is for entering IP addresses only and can't be used to edit IP addresses that have already been submitted.

Entering a quick range from a Discovery schedule

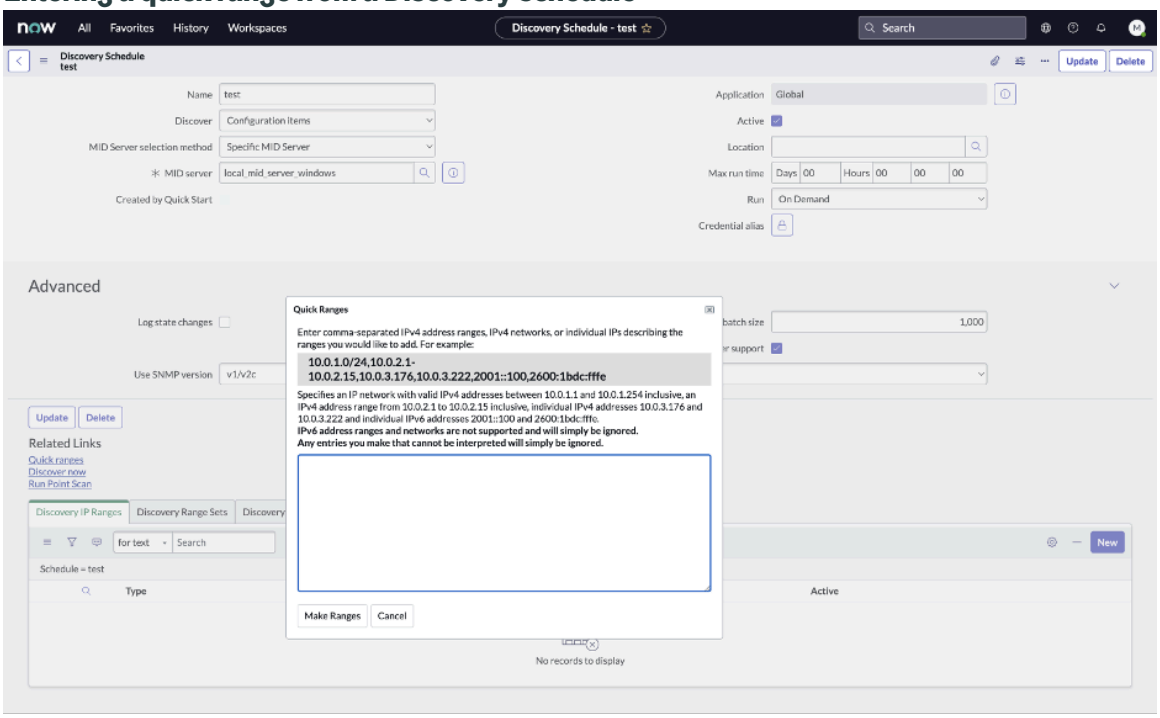

The instance automatically displays the entries in the proper format.

4. For changes to IP address ranges, select the IP address records in the **Discovery Range Sets** related list.

Note: Overriding behavior is not applicable on Discovery range sets.

Import IP ranges into Discovery schedules with import sets

One method of entering large numbers of IP networks into Discovery schedules is by using import sets.

Before you begin

Role required: discovery_admin

About this task

Common groups of IP addresses, known as ranges can be used in Advanced Discovery schedules.

Note: You can also use IPAM integration for entering large numbers of IP networks into Discovery schedules. See IPAM Integration for more information.

Use a data source that can be mapped. Include these fields:

- Start IP: the first IP address in the range (inclusive).
- End IP: the last IP address in the range (inclusive).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Import Sets** > **Load Data**.
2. Identify the file or data source that contains the desired information.
3. Create a table name, such as `ipnetworks`.
4. Select **Upload an Excel file** and browse to the source file.
5. Click **Go** to import the file.

Data Source

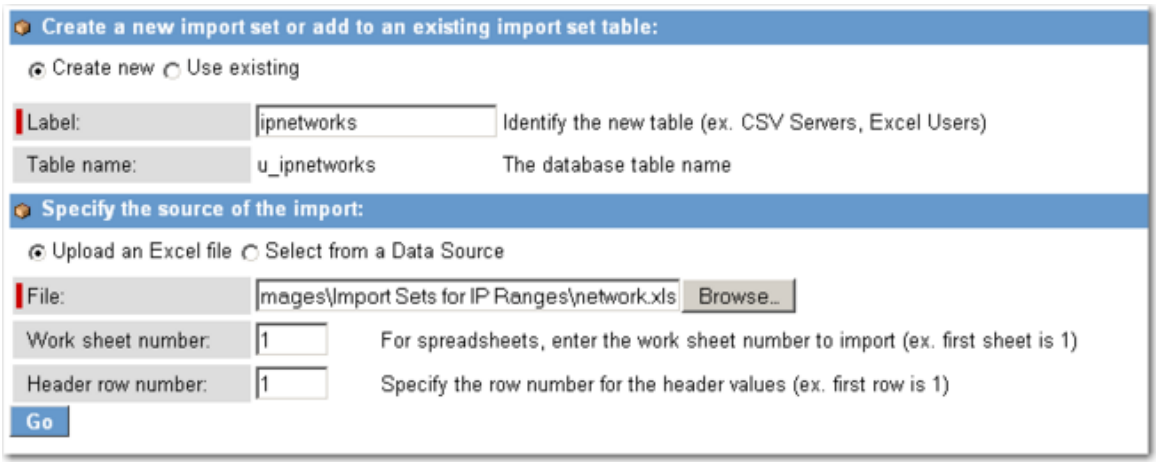

6. Navigate to **System Import Sets** > **Create Transform Map** and map the items in the Excel spreadsheet to the fields of the CMDB in the target table IP Range `[ip_address_range]` table.
7. Give the Transform Map a unique and descriptive name.
8. Submit the form, and then click **New** in the **Field Maps** Related List.
9. Map the fields from the Excel spreadsheet to the fields in the IP Range `[ip_address_range]` table.
   The fields you need values for are the **Start IP** and **End IP** addresses.
10. Click the **Mapping Assist** Related Link and use the lists that appear to resolve the fields between the table and the data source (the Excel spreadsheet in this example).
11. Click **Save**.

    The view returns to the Table Transform Map form.

12. Click **Transform** in the Related Links to move the data into the proper fields in the IP Range `[ip_address_range]` table.

    The imported IP ranges are available now for use in any advanced Discovery schedule.

Exclude IP ranges from a Discovery range set

You can specify a range of IP addresses that you want to exclude from your Discovery query.

Before you begin

Before a Discovery schedule is run, the number of excluded IP addresses is totaled. If there are more than 500,000 excluded IP addresses, the Discovery is cancelled and an error is logged through DiscoveryLogger regarding its status.

Note: Exclude IPs is not supported for IPv6 ranges or subnets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin or agent_admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Range Sets**.
2. Select one of the range sets.
3. Select one of the types from the **Discovery IP ranges** related list.
4. Click **New** from the **Discovery Range Item Excludes** related list.
5. Choose one of the following options from the **Type** drop-down list.

Option	Description
IP Address List	Enables you to exclude non-consecutive IP addresses by listing individual IP addresses in the **Discovery Range Item IPs** related list. **Note:** After you select **IP Address List** as the **Type**, you must right-click the header and select **Save** before you can begin adding IPs to the related list.
IP Address Range	Enables you to exclude a range of IPs by pro viding the starting and ending IP addresses.
IP Network	Enables you to exclude an IP network by pro viding the Network IP and Network mask.

6. Click **Submit**.

Use Global Excludes List for IP addresses and ranges

Global Excludes List allows administrators to define global IP exclusions lists that work across Discovery schedules that discover configuration items, IP addresses, or networks. The list helps to prevent access to sensitive IPs as it blocks discovery when the IP is on the exclusion list. This feature is only applicable for Horizontal Discovery starting in the Rome release.

Before you begin

Global IP exclude ranges are active by default. To deactivate, uncheck the **Active** check box. This makes the records inactive and the entries aren't excluded from Discovery. You can add a single IP or an IP range to the Global Excludes list. The IP exclusion list [ip_exclusion] table references existing IP collection tables and supports three types of collections: IP address list, IP address subnet, and IP address range.

Note: Use of Global Excludes List for IP addresses and ranges is not supported for IPv6 addresses. While the UI still allows users to add IPv6 addresses to the list, these addresses get ignored and are skipped from the exclusion process during Discovery.

Role required: discovery_admin or agent_admin role

Procedure

1. Navigate to **All** > **Global IP Exclusion** > **IP Exclusion**.
2. Click **New** to create a new IP exclusion record.
3. In the IP Excluded box, click the search icon to bring up the IP collections window.
4. Select one of the IP collections listed or click **New** to bring up the IP Collections wizard.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. From the IP Collections wizard, select one of the three types of collections you want to create and fill in the necessary details.
   For each type, you can add a description on why you want to exclude particular IPs or which IPs are excluded.
   - IP Address Range

     Fill in the details and then click **Submit**. This new range is added to the IP Exclusion list.
   - IP Address Subnet

     Fill in the details and then click **Submit**. This new subnet is added to the IP Exclusion list.
   - IP Address List

     Enter the name and then **Save**. The related list then shows at the bottom of the page. From the collections list, pick an existing item or click **New** to create a new list. Click **Update**. Remember the name you created. This new name is added to the collection list. Select this name from the list and click **Submit** again. This address list is added to the IP Exclusion list.
6. Navigate to **Discovery** > **Discovery Schedules** and select a schedule.
7. If the schedule is for discovering configuration items, IP addresses, or networks, you can view the Global IP Exclusion tab at the bottom of the page.

Result

When the scheduled Discovery runs, it skips the discovery of all IPs that are part of the active Global IP Exclusion record. All the rest of the IPs should be discovered.

Note: Overriding behavior is not applicable when the Discovery schedule has active Global IP exclusions.

If you try to run a Quick Discovery that includes an excluded IP, you will see an error message and Discovery will not be triggered.

Discovery generic attributes

Discovery generic attributes enable you to define custom attributes at the schedule, range, and range set level. These attributes attach metadata that can influence CMDB field population and Discovery runtime behavior, providing granular control based on the scope of the discovery.

How it works

Discovery generic attributes use a declarative, rule-based approach for applying metadata during Discovery. Each attribute defines a scope, action, target, and value. You can attach attributes at different levels of the Discovery hierarchy to control how metadata is applied. When attributes exist at multiple scope levels, Discovery applies them using a precedence hierarchy. Attributes defined at a more specific scope take precedence over those at a broader scope. This hierarchy enables you to set default attribute values at a broad level while overriding them for specific scopes where different values are needed.

Attribute scope hierarchy

Discovery reconciles attributes from multiple scope levels by applying a precedence order. Attributes at a more specific scope override attributes with the same key at a broader scope. The available scope levels and their precedence depend on the use case. For more information, see Automatic CI field population.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use case

You can use Discovery generic attributes to automatically populate CI fields based on the Discovery scope. This enables you to apply consistent metadata to discovered CIs without manual intervention. For more information, see Define CI field attributes.

Automatic CI field population

Discovery generic attributes can automatically set configuration item (CI) field values during discovery. Attributes follow a scope hierarchy, where more specific scopes override broader ones, enabling you to define defaults at the schedule level and apply precise values at the range level.

Benefits

Using generic attributes to populate CI fields provides the following benefits:

- Eliminate the need to create separate Discovery schedules for each location or group of CIs that require different field values.
- Set CI field values at the schedule, range set, or IP address range level, providing precise control over location data, asset tags, and other fields.
- Reduce manual CI updates after discovery by automatically populating fields during the discovery process.
- Use the attribute hierarchy to define broad defaults at the schedule level and override them with more specific values at the range set or range level as needed.

How it works

Previously, CI field values such as location could only be set globally on a Discovery schedule. This meant that all CIs discovered within a schedule inherited the same value, even when more specific information was available for individual IP addresses or range sets.

Generic attributes solve this problem by enabling you to define field values at multiple levels of granularity. When discovery runs, it evaluates the attributes defined at each scope level and applies the most granular matching attribute to each discovered CI.

Scope levels

You can define generic attributes to populate CI fields at the following scope levels. Each level represents a different degree of granularity.

Scope	Description
Schedule	Applies the attribute value to all CIs discovered within the entire schedule.
Range set	Applies the attribute value to all CIs discovered within a specific range set associated with the schedule.
Range	Applies the attribute value to all CIs discovered within a specific IP network, IP address range, or IP address list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attribute hierarchy

When attributes are defined at multiple levels, Discovery applies the most granular value to the resulting metadata, which can influence both CMDB field population and Discovery runtime behavior. The scope hierarchy from least to most granular is: schedule, range set, IP address range. A range-level attribute overrides a range set-level attribute, which overrides an IP network-level attribute, and so on. Additionally, attributes can target specific CMDB CI classes. When attributes with the same key exist for different classes within the same class hierarchy, the more specific class takes precedence. If no attribute is defined at a more granular scope or specific class, the CI inherits the value from the next available level.

For example, set the **Location** field to Maryland at the schedule level, Baltimore at the range set level, and a specific street address at the range level. A CI discovered from an IP address in that range inherits the street address. A CI discovered from a different IP address in the same range set inherits Baltimore, because no range-level location is defined.

Field value types

Discovery generic attributes support two types of field values.

Type	Description
Reference	A value that references a record in another table. For example, the Location field references a record in the Location [cmn_location] table.
Static	A string value that you enter directly. For example, you can set the Asset tag field to a custom string value.

Target table filtering

Each attribute includes a Target Table field that specifies which CI class the attribute applies to. You can set the target table to the base Configuration Item [cmdb_ci] table to apply the attribute to all discovered CIs. To restrict the attribute to a specific class and its child classes, select a more specific class, such as Linux Server [cmdb_ci_linux_server].

Example

The following example illustrates how the attribute hierarchy works across three scope levels.

A Discovery schedule contains a range set with two IP address ranges. You define the following attributes:

- **Schedule level: Location** = Maryland (applies to all CIs in the schedule)
- **Range set level: Location** = Baltimore (target table: Linux Server)
- **IP address (10.0.0.52): Location** = 123 Main St, Baltimore
- **IP address (10.0.0.117): Asset tag** = tag-test (static value)

After discovery runs:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The CI discovered at 10.0.0.52 has its **Location** field set to **123 Main St, Baltimore** because the range-level attribute is the most granular.
- The CI discovered at 10.0.0.117 has its **Asset tag** field set to **tag-test**. Because no range-level **Location** attribute is defined for this IP, the CI inherits **Baltimore** from the range set-level attribute.

Define CI field attributes

Define attributes on a Discovery schedule, IP network, IP range set, or IP address range to automatically populate configuration item (CI) field values during discovery.

Before you begin

Confirm the following

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- You have an active Discovery schedule with at least one range set and IP address range configured. For more information, see Create an IP-based Discovery schedule in Discovery Admin Workspace and Discovery IP address configuration.
- You have installed and configured Visibility Content v6.30.0. For more information, see Install Visibility Content.

Role required: discovery_admin

About this task

Generic attributes enable you to attach metadata to discovered CIs based on the scope of the discovery. Currently, you can use attributes to assign field values to CIs. You create attributes in the Discovery Attributes related list, which is available on Discovery schedule, range set, and IP address range records. When discovery runs, it evaluates the attributes and applies the most granular matching value to each discovered CI.

Note: You can define attributes at three scope levels. A range-level attribute overrides a range set-level attribute, which overrides a schedule-level attribute.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Select the **Discovery Attributes** tab from the related lists.
3. Select **New**.
4. Complete the fields on the form.

Field	Description
Name	The name of the generic attribute. Enter a unique name for the generic attribute.
Attribute Scope Table	The record that the attribute is associated with. This field is populated automatically based on the scope that you select and the record from which you create the attribute.
Attribute Scope	The level at which the attribute applies. This field is set automatically based on the record

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	from which you create the attribute and shouldn't be edited.
Action	The type of action to perform. CI field is the only available option for this release and mod ifies a field value on discovered CIs.
Target Table	The CI class that the attribute applies to. Select **Configuration item** to apply the at tribute to all CI types. Or, select a specific class such as Linux Server to restrict the at tribute.
Target Field	The CI field to populate. For example, select **Location** to set the location on discovered CIs.
Value	The value to assign to the selected field. For reference fields, select a record from the refer enced table. For string fields, enter a static val ue.

5. Select **Submit**.

Result

When the Discovery schedule runs, Discovery evaluates the attributes defined at each scope level and applies the appropriate values. If attributes are defined at multiple levels, the most granular value takes precedence.

What to do next

To define more granular attributes, navigate to the range set or IP address range record. From the Discovery Attributes related list on that record, repeat the previous steps.

IPAM Discovery integration

The IP Address Management (IPAM) to Discovery integration feature enables your organization to automatically create and manage Discovery schedules based on your IPv6 network infrastructure data stored in IPAM. This integration keeps your discovery processes synchronized with your IPv6 network changes, providing a complete and up-to-date view of your network environment.

How it works

The IPAM to Discovery integration automates schedule management through the following process:

- Identifying resources marked as operational in your IPAM tables
- Processing only the latest changes from your IPAM system
- Creating Discovery schedules, organized by a tag key you specify
- Disabling schedules that no longer have valid network targets

Benefits

The following examples highlight the primary advantages of integrating IPAM with Discovery:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Automated schedule management**: Eliminates the need to manually create and maintain Discovery schedules for IPv4 and IPv6 networks.
- **Synchronization with network changes**: Automatically adjusts Discovery schedules when networks, subnets, or IP addresses change.
- **Incremental updates**: Processes only resources that changed since the last synchronization, improving performance.
- **IPv4 and IPv6 support**: Handles both IPv4 and IPv6 networks with separate capacity limits and mapping approaches.
- **Audit capability**: Maintains history of synchronization activities and changes through execution records.

Requirements

This feature requires the following:

- Service Graph Connector Central (SGC Central) v2.4.0 with sgc_admin access.
- Service Graph Connector for Infoblox v1.5.0 with sgc_admin access.
- CMDB Workspace v9.0.0.
- Discovery Admin Workspace v1.13.0 with discovery_admin access.
- Australia, ZP8 or later, or YP13 or later version of the ServiceNow AI Platform.

Before using the integration, ensure that your IPAM tables contain the following data:

- Managed networks in the Managed Network [cmdb_ci_managed_network] table with operational status set to operational.
- IP subnets in the IP Network Subnet [cmdb_ci_ip_network_subnet] table with operational status set to operational.
- IP addresses in the Allocated IP Address [cmdb_ci_allocated_ip_address] table with operational status set to operational.
- IPv4 subnets should have valid Classless Inter-Domain Routing (CIDR) notation for address calculation.
- Relationships between managed networks, subnets, and IP addresses must be established.

Configuration

The following system properties control the behavior of the integration:

System properties

Property name	Description	Default
glide.discovery.schedule.autogen.enable.debug	Enable detailed logging for troubleshooting.	false
glide.discovery.schedule.autogen.scope.group	Create auto-generated IPAM Discovery schedules by a specific tag or attribute. The value of this property represents a key that IPAM retrieves and stores in the Key Values [cmdb_key_value] table. For example, entering	empty

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

System properties (continued)

Property name	Description	Default
	"location" as the value would group the schedules by their associated locations.	
glide.discovery.schedule_ipv4_limit	Defines the maximum number of IPv4 addresses that can be included in a single auto#generated IPv4 Discovery schedule.	65,536
glide.discovery.schedule_ipv6_limit	Defines the maximum number of IPv6 addresses that can be included in a single auto#generated IPv6 Discovery schedule.	20,000 **Note:** On instances below ZP3 and YP9, the maximum number of IPv6 addresses per schedule must not exceed 5,000, which is also the default value in Yokohama and earlier releases.

Generated schedules

The integration creates Discovery schedules with the following characteristics:

- **Maintenance source**: The maintenance source is set to Discovery-IPAM Integration.
- **Initial state**: The initial state is set to inactive for review before activation.
- **MID Server selection**: The MID Server selection is set to MID Cluster.
- **Description**: Includes managed network, grouping criteria, and navigation links.

Execution tracking

Each synchronization creates an execution record that captures performance metrics, configuration settings, and any errors. The integration helps prevent multiple simultaneous executions to avoid conflicts. Review these execution records to monitor integration health.

Optimization

Follow these guidelines to get the most out of the IPAM Discovery integration:

- **Assign subnets to groups**: In your IPAM product, assign subnets to groups so the integration can organize schedules based on those groupings.
- **Regular maintenance**: Periodically review generated schedules and adjust configuration as needed.
- **Test with debug logging**: Enable debug logging during initial implementation to understand behavior.
- **Verify data quality**: Verify IPAM data quality before running synchronization.
- **Monitor execution records**: Review execution records regularly to catch issues early.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure auto-created Discovery schedules

Enable the system to automatically create Discovery schedules based on IP data from your IP Address Management (IPAM) connections.

Before you begin

Confirm the following:

- You have installed and configured Service Graph Connector Central (SGC Central) v2.4.0. For more information, see Configuring SGC Central.
- You have installed and configured Service Graph Connector for Infoblox v1.5.0. For more information, see Configure Service Graph Connector for Infoblox using SGC Central.
- You're using Discovery Admin Workspace v1.13.0.
- You're using the Brazil, ZP9 or later, or YP13 or later version of the ServiceNow AI Platform.

Role required: discovery_admin

About this task

The IPAM to Discovery integration monitors your IPv6 network structure by tracking changes in managed networks, subnets, and IP addresses. Automatic schedule creation uses your network topology to generate Discovery schedules and organizes them by location when location data is available. Additional checks maintain the appropriate schedule size by evaluating the number of IP addresses in each schedule and updating them when IPs are added, removed, or deactivated.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Settings** > **IP Address Management (IPAM)**.
2. Select the **Enable auto-created schedules** toggle.
3. Enter a tag key to determine how IPs are assigned to schedules.
   Discovery evaluates the key that you provide, retrieves the corresponding value from your IPAM data by looking it up in the Key Values [cmdb_key_value] table, where it searches only subnet records. IPs that share values are assigned to the same schedule, allowing the schedule structure to follow meaningful attributes such as location or environment.

   Note: You can enter a single space in the **Tag key** field if you don't want to provide a key.

4. Configure when the automatic synchronization job runs.
   a. Select a frequency from the **Update frequency** drop-down list.
   b. Enter the time of day that you want the IP data to refresh.
   c. Choose a start date.
5. Select **Save**.
   The system automatically runs the synchronization job according to your schedule, keeping your discovery targets updated with the latest IPv6 information from IPAM.
6. **Optional:** To synchronize your IP data immediately instead of waiting for the next refresh time, select **Refresh IP data**.

What to do next

After you enable auto-created schedules and configure the IP data refresh settings, IPAM generates a corresponding Discovery schedule. The schedule is inactive when it's created, so make sure to activate it before it runs. For more information, see Activate auto-created Discovery schedules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

IPAM Discovery integration

Activate auto-created Discovery schedules

Activate schedules that were created automatically via the Discovery and IP Address Management (IPAM) integration.

Before you begin

Confirm the following:

- You have at least one validated MID Server with a status of **Up**. The MID Server must reach the IP ranges that you intend to discover. You can install the MID Server by using Use MID Server guided setup or by manually downloading and running the installer. For details, see the MID Server installation instructions for Linux and Windows.
- You have installed and configured Service Graph Connector Central (SGC Central) v2.4.0. For more information, see Configuring SGC Central.
- You have installed and configured Service Graph Connector for Infoblox v1.5.0. For more information, see Configure Service Graph Connector for Infoblox using SGC Central.
- You have enabled auto-created schedules in Discovery Admin Workspace Settings. For more information, see Configure auto-created Discovery schedules.
- You're using Discovery Admin Workspace v1.15.0.
- You're using the Australia, ZP8 or later, or YP13 or later version of the ServiceNow AI Platform.

Role required: discovery_admin

About this task

The IPAM to Discovery integration monitors your IPv6 network structure by tracking changes in managed networks, subnets, and IP addresses. Automatic schedule creation uses your network topology to generate Discovery schedules and organizes them by location when location data is available. Auto-created schedules are inactive by default and must be manually activated before they can run.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules** > **IP-based discovery**.
2. Select **View and activate schedules** from the **Needs attention** section.
3. Select the check box next to the name of the auto-created schedule that you want to activate. Alternatively, you can activate all the schedules by selecting the check box next to the **Name** column.
4. Select **Activate**.
   An Activate schedule? dialog displays.
5. Select the **MID Server** drop-down list and choose a MID Server selection method.

Option	Description
Specific MID Server	Select this option to search for an existing MID Server. Only MID Servers that are validated and up display in the search results.
Specific MID Cluster	Select this option to search for an existing MID Server cluster. Clusters provide failover protection and load balancing between MID

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	Servers. See for more information.

6. Select the **Run settings** drop-down list and configure when the schedule will run.

Option	Description
On Demand	The schedule only runs when triggered manu ally.
After Discovery	The schedule runs after a specified Discovery schedule completes.
Weekly	The schedule runs weekly on a specified day, or days, and time.
Monthly	The schedule runs once a month on a speci fied day and time.
Periodically	The schedule runs at a specified repeat inter val, such as every number of days, hours, or minutes.
Once	The schedule runs one time only on a speci fied date and time.
Daily	The schedules runs once a day at a specified time.
Calendar Quarter End	The schedule runs on the last day of each quarter at a specified time.

7. **Optional:** Set a time limit for running this schedule.
   When the configured time elapses, the remaining tasks for the discovery are canceled, even if the scan isn't complete. Use this field to limit system load to a desirable time window. If no value is entered in this field, this schedule runs until complete. If the schedule exceeds the maximum runtime, it's canceled.
8. Select **Save and activate**.
   The system automatically runs the synchronization job according to your schedule, keeping your discovery targets updated with the latest IPv6 information from IPAM.

Related topics

IPAM Discovery integration

Domain Separation and Discovery

Domain separation is supported in Discovery. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Standard

- Includes all aspects of **Basic** level support.
- Application properties are domain-aware as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Business logic: The service provider (SP) creates or modifies processes per customer. The use cases reflect proper use of the application by multiple SP customers in a single instance.
- The instance owner must configure the minimum viable product (MVP) business logic and data parameters per tenant as expected for the specific application.

Sample use case: An admin must be able to make comments required when a record closes for one tenant, but not for another.

For more information on support levels, see Application support for domain separation.

Domain separation overview

Service Providers (SPs) use domain separation to segregate data for each customer. Users in a given domain have visibility only to the data in their own domains or in child domains. SPs typically control the top-level domain, which gives them visibility to data associated with all domains. Given that Discovery domain separation support is considered **Standard** there is no delegated administration to the child domains. The SPs must retain administrative control.

How domain separation works in Discovery

Multiple domains can be supported by a single MID Server. In releases prior to Kingston, each MID Server could support only a single domain. In newer releases, segregating domains by MID Server is useful when the domain is large, or when the domain's resources are held in a customer's data center rather than the SP's. For Discovery on MID Servers supporting a single domain, the discovered CIs are assigned to the domain of the MID User used to authenticated against the ServiceNow instance. In multi-domain MID Servers, the discovered CIs are assigned to the domain of the user who created the Discovery schedule.

Discovery implements data domain separation through the MID Server by impersonating the MID Server user during sensor processing. Discovery uses the domain, that the MID Server user is in, to determine which domain the discovered data should be put into. Discovery configuration information, including classifiers, identifiers, probes, and sensors, is not domain separated.

Service providers generally use IP-based Discovery. In cases where the SP controls the network addressing, they divide the address space among their customers to ensure that each domain has a distinct IP address space. The SP assigns one or more subnets to a customer or domain and creates Discovery schedules for those subnets.

If the SP is remotely managing their customer's data center, there will often be some overlap between address spaces different customers use. In these cases, the SP can use network address translation (NAT) on the IP range and run a Discovery schedule.

Once the CIs are assigned to the correct domain, the visibility and read/write access control are provided by the platform through the domain hierarchy. Schedules are visible to users in their respective domains. Cross-domain schedule visibility is not possible, except for the SP who controls the parent domain and has visibility to all domains.

Domain separation for MID Server files

You can create versions of these specific MID Server policy records that only a MID Server from the same domain can use. This process separation is supported for records in tables that extend MID Server Synchronized Files [ecc_agent_sync_file]:

- MID Server MIB File [ecc_agent_mib]
- MID Server JAR File [ecc_agent_jar]
- MID Server Script File [ecc_agent_script_files]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, all records in these tables are members of the global domain. A user can override the default global domain and create a version of these policies for use in the user's own domain.

Note: Attachments on MIB or JAR file records might not appear as they did in a non-domain separated environment. The attachments don't appear because the **Attachments** [sys_attachment] table is data separated. When data is separated between domains, a record in a child domain can't access records in a parent domain.

See MID Server domain separation for instructions on setting up domain separation through the MID Server.

Domain-separated tables

Records in all tables that extend the Base Configuration Item [cmdb] table can be domain-separated. In addition, records in these tables can also be domain-separated:

- Serial Number [cmdb_serial_number]
- TCP Connection [cmdb_tcp]
- Fibre Channel Initiator [cmdb_fc_initiator]
- Fibre Channel Targets [cmdb_fc_target]
- IP Address to DNS Name [cmdb_ip_address_dns_name]
- Service [cmdb_ip_service_ci]
- KVM Virtual Device [cmdb_kvm_device]
- Load Balancer Service VLAN [cmdb_lb_service_vlan]
- Load Balancer VLAN Interface [cmdb_lb_vlan_interface]
- Switch Port [cmdb_switch_port]

Related topics

Domain separation for service providers

Domain separation recommended practices for service providers

Discovery behaviors

Configure domain-separated Discovery schedules

Configuring the "Run as" user in a Discovery schedule directs discovered Configuration Items (CIs) to the correct domain and maintains data isolation.

Before you begin

Role required: discovery_admin

About this task

In domain-separated environments, the domain of the user who creates a Discovery schedule typically determines where the discovered data is stored. Assigning a "Run as" user enables the discovery process to run under a different user's domain and permissions. This enables administrators to control which domain stores the data and what resources the discovery process can access, regardless of who created the schedule.

Note: This procedure applies only to existing Discovery schedules. For information about creating a Discovery schedule, see Schedule a horizontal discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. In the **Name** field, enter the name of the schedule you want to configure.
3. Double-click the **Run as** field to perform inline editing.
   By default, the **Run as** column doesn't display in the Discovery Schedules [discovery_schedule] table. For information about customizing columns in tables, see Personalize a list.
4. Enter the name of the user that you want the schedule to run as.
5. Select the check mark ( ) icon.

Microsoft Just Enough Administration (JEA) for Discovery

Using Microsoft JEA with Discovery improves security by forcing the MID Server to run remote Windows commands through a constrained endpoint, which validates commands on the target before execution.

Microsoft JEA enables role-based administration through PowerShell Remoting, which uses Windows Remote Management (WinRM) to manage communication and authentication. This framework provides a secure and reliable method for managing computers that use the HTTP protocol. PowerShell Remoting uses two total ports (5985, 5986) for HTTP and HTTPS, which is easier to secure than the multiple ports used in WMI dynamic port mapping. For more information about Microsoft JEA, see Just Enough Administration.

Requirements for Discovery with JEA

- A ServiceNow instance running on the Rome release or later.
- The MID Server and target server must be part of a Windows domain.
- The JEA credentials with non-administrator rights must be domain level credentials.
- PowerShell 5.0 or Windows Management Framework 5.1 must be installed on the target Windows machines.
- PowerShell Remoting must be enabled on the target Windows machines.

Note: For security enhancement, starting from Rome there's a new profile called JEA v2. Microsoft doesn't recommend specifying any other language mode than **NoLanguage** in the JEA profile. JEA v2 explicitly sets the session type to **RestrictedRemoteServer** and the language mode to **NoLanguage** to help prevent users running arbitrary scripts at the endpoint and bypassing security restrictions. ServiceNow no longer supports the existing sample profile in KB0782125. Follow the instructions in KB0965705 to set up and deploy JEA v2 profiles.

JEA Profiles

Discovery with JEA requires profiles composed of a PowerShell Sessions Configuration and one or more PowerShell Role Capabilities files. You can create multiple PowerShell Role Capability files and multiple user groups to assign the roles to different groups as necessary. A sample profile is provided in KB0965705 as a reference implementation and to serve as a starting point. The configuration file in the KB supports all out-of-the-box horizontal Windows patterns at the time that it was created. ServiceNow isn't responsible for the JEA profile deployment and setup on remote machines.

The following sample profile improves security by blocking attackers if a Windows service account becomes compromised. It restricts the MID Server user to execute only scripts signed

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

by a trusted MID Server. After the profile is deployed, the connection between a MID Server and a Windows target is secured by script and parameter signatures created by the MID Server. This provides one layer of security; but monitoring on the ServiceNow instance and the MID Server itself is also essential to promoting a secure deployment.

Microsoft has detailed documentation at the following links:

- JEA Role Capabilities
- JEA Session Configuration

Basic Discovery with the sample JEA Profile

The sample JEA profile provided in KB0965705 was configured to discover many basic configuration items (CIs) and attributes. The profile can be modified and should only serve as a baseline for Discovery with JEA.

Basic Discovery finds the following key attributes of Windows servers (cmdb_ci_win_server) or Windows desktops (cmdb_ci_computer):

- Hostname
- DNS Name
- Serial Number
- Operating System
- OS Version
- OS Service Pack
- Disk Space
- RAM
- CPU Core Count
- CPU Count
- CPU Manufacturer
- CPU Type

It includes the following CIs:

- Network Adaptors (cmdb_ci_network_adapter)
- File System (cmdb_ci_file_system)
- Storage Devices (cmdb_ci_disk)
- Software Installed (cmdb_software_instance)
- Running Processes (cmdb_running_process)
- Memory Modules (cmdb_ci_memory_module)
- Serial Numbers (cmdb_serial_numbers)
- TCP/IP Connections (cmdb_tcp)
- CI IPs (cmdb_ci_ip_address)
- DNS Names for CIs (cmdb_ci_dns_name)
- Windows Clusters (cmdb_ci_win_cluster, cmdb_ci_win_cluster_node, cmdb_ci_win_cluster_resource)
- Tracked Configuration Files (cmdb_ci_config_file_tracked)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following application CIs can also be discovered:

- MSSQL DB on Windows (cmdb_ci_db_mssql_instance)
- MySQL DB on Windows (cmdb_ci_db_mysql_instance)
- Oracle DB on Windows (cmdb_ci_db_ora_instance)
- WebSphere on Windows (cmdb_ci_app_server_websphere)

The following probes and patterns are used for Discovery with the sample profile:

- Windows - Classify (probe)
- Windows OS - Servers (pattern)
- Windows OS - Desktops (pattern)
- Windows - Installed Software (probe)
- Windows - ADM (multiprobe)
- MySQL server On Windows (pattern)
- MS SQL DB On Windows (pattern)
- Oracle DB on Windows (pattern)
- Windows - WebSphere - Cell (probe)
- Windows - WebSphere - Web Applications (probe)
- Windows - WebSphere - Web Services (probe)

Prepare the instance for Discovery with JEA

To configure the ServiceNow® instance for Discovery with Microsoft Just Enough Administration (JEA), define the Windows credential with the domain name and set the MID Server configuration parameters appropriately.

Before you begin

Role required: discovery_admin or admin

Procedure

1. Navigate to **All** > **Discovery** > **Credentials** and click **New**.
2. Select **Windows Credentials** from the list of available credential types.
3. Create the credentials for the non-administrator, using this format for the **User name**: *domain\user name*.
4. Submit the record.
5. Navigate to **All** > **MID Server** > **Servers**.
6. Select a **MID Server** to configure from the list of MID Servers.
7. Select the **Configuration Parameters** related list.
8. Set the following MID Server config parameters as indicated:

   a. **mid.windows.management_protocol:** This parameter is required for Discovery with JEA. Its default value is WMI, but it must be set to WinRM on MID Servers using Discovery with JEA.

   b. **mid.powershell.jea.endpoint**: This parameter is required for Discovery with JEA.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This parameter specifies a JEA endpoint name which the MID Server connects to on remote hosts. The endpoint name is created when registering a configuration file, and should not be confused with the name of the configuration file itself. The setting affects the entire MID Server, including any WinRM remote sessions created by Discovery on the MID Server which go to that endpoint.

For example, the PowerShell command **Register-PSConfiguration -name JEA_DISCO_V2 -path <session_configuration_file>** sets the endpoint name to JEA_DISCO_V2. In that case, **mid.powershell.jea.endpoint** must be updated to JEA_DISCO_V2.

9. **Optional:** Use the following MID Server property and System property to troubleshoot issues:

   a. **mid.probe.collect_debug_info**: This is an optional MID Server property to collect debug information.
   When this property is set to true, the MID Server collects credential debug info and puts it in the payload of ECC input message. It doesn't affect how JEA works.

   b. **glide.discovery.log_debug_info**: This is an optional system property to collect debug information.
   When this property is set to true, the discovery sensor extracts the debug info from ECC input message and writes it to the discovery log table. The debug information is visible when inspecting discovery status.

SNMP support for Discovery

Discovery supports SNMP versions 1, 2c, and 3. If you have an active SNMP v3 credential, valid options for SNMP version are v3 or All. If you have an active SNMP v1 or v2 credential (community string), valid options for SNMP version are v1, v2c, or All. Default is All.

MID Servers support all SNMP protocol versions by default. You can set a MID Server to support only specific versions of SNMP.

For SNMP credentials information, see SNMP credentials.

SNMP probe parameters can be used while discovering an SNMP device. See SNMP probe parameters for more information.

Add an SNMPv3 user credential in Discovery

Set up your credentials so Discovery can access SNMPv3 targets.

Before you begin
Role required: admin

About this task
A MID Server parameter is also available to control SNMP versions. See MID Server parameters.

Procedure

1. Navigate to **All** > **Discovery** > **Credentials**.
2. Click **New**.
3. Click **SNMPv3 Credentials**.
4. Fill in the form, using the fields in the table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SNMPv3 credentials form

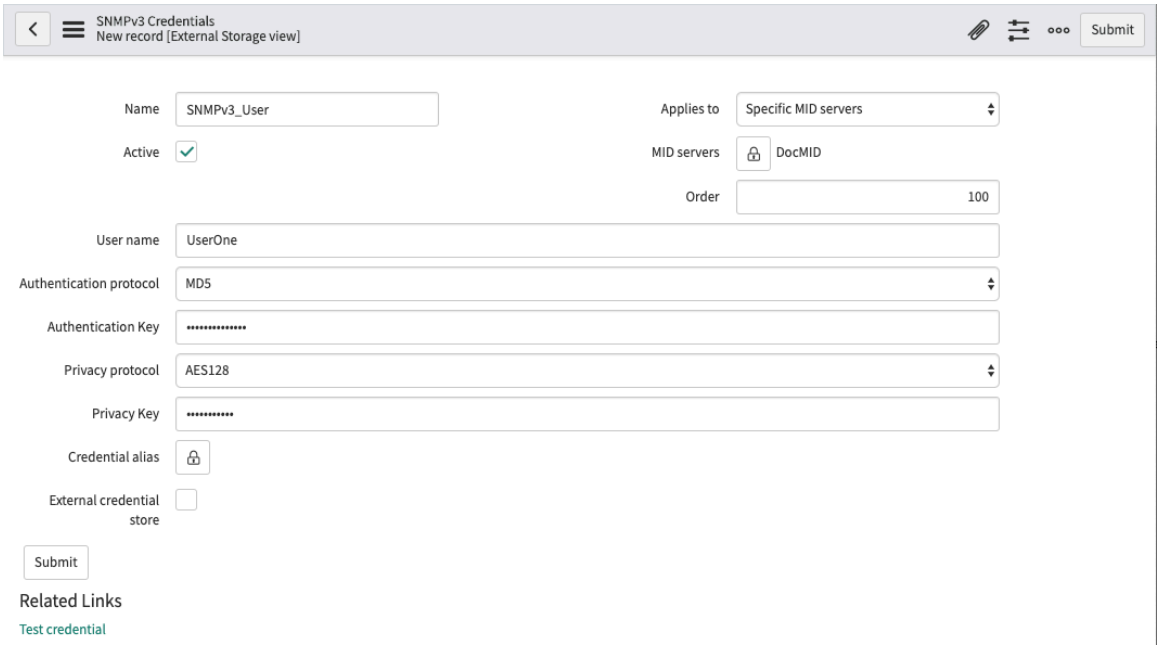

SNMPv3 credential fields

Field	Input value
Name	Unique and descriptive name for this credential. For example, you might call it **SNMP Community Atlanta**.
Active	Enable or disable these credentials for use.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
User name	Enter the SNMP user name. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name.
Authentication protocol	Select the authentication type to use for this credential. The choices are: ○ MD5 ○ SHA-1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input value
	○ SHA-224 ○ SHA-256 ○ SHA-384 ○ SHA-512
Authentication Key	Enter the authentication key to use for this credential.
Privacy protocol	Select the encryption protocol for this credential. The choices are: ○ 3DES ○ AES128 ○ AES192 ○ AES256 ○ DES
Privacy key	Enter the key associated with the selected privacy protocol.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential provider. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Privacy Credential ID	Enter the privacy account key configured for SNMPv3 credentials in CyberArk. If you are using a privacy protocol in CyberArk, this field must have the same value as the **Name** field for the CyberArk SNMPv3 Privacy Account. This field is only visible for SNMPv3 when you select **CyberArk** from the **Credential Store Type** field. If you are not using a privacy key for CyberArk, leave this field empty.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow. You can also assign different credentials to each occurrence of the same activity type.
External credential store	Select this check box to use an external credential storage system. When external storage is enabled, the **Credential ID** field appears. External credential storage is only available when the External Credential Storage plugin is activated and the **External Storage** view is selected.
Credential Store Type	Select the external storage provider. Select **CyberArk** only if you are using a CyberArk SNMPv3 privacy key. The **Privacy Credential ID** field appears to allow entry of the key.
Use Context	Select this check box to add a context value for this credential. This field is visible in the **Discovery** view. Contexts are not currently supported for external credential storage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input value
	Note: A context is a collection of management information accessed by an SNMPv3 credential that references a specific OID. Contexts are sometimes referenced to collect information about the device that cannot be accessed by the normal credential. A context can be provided by the manufacturer or configured separately. If you have multiple SNMPv3 credentials with the same user name and keys, but some of your devices need a context and some do not, then you will need to create separate records for each device.
Context Name	Enter the context name value for this credential. This should only be used if you have devices that require this value for full access. This field is visible when the **Use Context** check box is selected.

SNMPv3 security options

Option	Fields to configure
noAuthNoPriv security	○ **Authentication protocol:** None ○ **Privacy protocol:** None ○ **Authentication Key:** Leave blank ○ **Privacy Key:** Leave blank
authNoPriv security	○ **Authentication protocol:** Set ○ **Privacy protocol:** Leave blank ○ **Authentication Key:** Set ○ **Privacy Key:** Leave blank
authPriv	○ **Authentication protocol:** Set ○ **Privacy protocol:** Set ○ **Authentication Key:** Set ○ **Privacy Key:** Set

Set SNMP version on the Discovery schedule

You must also set the SNMP version after choosing the appropriate credential.

Before you begin

Role required: admin

About this task

A MID Server parameter is also available to control SNMP versions. See MID Server parameters.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery schedules**> *your Discovery schedule.*
2. In the **Use SNMP Version** field, select the appropriate version.

   The default value is **ALL**. You can change this to **v1**,**v2c**, or **v3**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and SCCM together

Use these guidelines to avoid common issues when you use Discovery and System Center Configuration Manager (SCCM) together.

You must decide if you will use Asset Intelligence (AI) or non-AI software data sources. Because the software correlation data is tracked differently for each of these data sources, you should avoid switching back and forth between the two. If you must switch between data sources, remove all the software installation records imported by SCCM before the switch occurs.

When Software Asset Management is not enabled, software installation records are stored in the Software Instance [cmdb_software_instance] table. When Software Asset Management is enabled, software installation records are stored in the Software Installation [cmdb_sam_sw_install] table.

Note: If records aren't removed before the switch, duplicate records may exist. In the event that AI and non-AI data becomes mixed, clear the Software Installation table.

Overview of Service Graph connector for Microsoft SCCM

Use the Service Graph connector for Microsoft SCCM to pull data from SCCM into your ServiceNow instance. See Service Graph connector for Microsoft SCCM for more details.

Note: If you upgrade your instance to a version of the ServiceNow AI Platform that does not provide support for your version of SCCM, you can continue to use that version. However, all new instances require the use of a supported SCCM version.

Upgrading to a later version

The ServiceNow SCCM integrations are self-contained and can exist independently. They each use their own import set tables, data sources and transform maps. However, all SCCM integrations will transform data into the same tables within the ServiceNow CMDB. To avoid the data being overwritten by another source:

- Use one SCCM integration and disable all other SCCM scheduled imports.
- Perform a full import to clear the cmdb_software_instance table, the cmdb_sam_sw_install table, and other tables of old SCCM data.

Collect software data with either SCCM or Discovery

When Discovery and SCCM are both enabled on a system, the software records found through both tools could overwrite each other.

Before you begin
Role required: admin

About this task
To determine how software data is collected, configure the *glide.discovery.software_sccm_managed* property.

Procedure

1. Navigate to **All** > **Discovery Definition** > **Properties**.
2. Set the **Windows software is SCCM managed** property (*glide.discovery.software_sccm_managed*) appropriately.
   If this property is set to **Yes**, software for computer CIs is populated in the CMDB by SCCM. An internal process determines that the computer is managed by SCCM, and Discovery does

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

not populate the software records for Windows software data. The following information level message is written to the Discovery log file.

```
Skipping software population because the CI is managed by SCCM
```

If the check box is cleared, the property is set to **No**, and the system manages software data with Windows.

3. Click **Save**.

> **Important:** If SCCM Integration is active before Discovery is enabled and the property is enabled, Discovery ignores the population of software for any CIs that are also imported through SCCM. If Discovery is enabled before the SCCM Integration, it is possible for software installation data from both sources to be mixed.

What to do next

Use these Discovery properties to avoid software data discrepancy between Discovery and SCCM: *glide.discovery.enable.software_simplify* and *glide.discovery.enable.software_simplify_sccm*. See Discovery properties for more information.

Configure for VMware Discovery

Configure for VMware Discovery by setting up alternate port for vCenter and configuring your ESXi resource pools.

Alternative ports and resource pools aren't managed by ServiceNow ServiceNow directly. They're set up and managed by your security teams and VMware admins. Discovery captures information related to VMware use.

In many environments, security teams block the standard ports, and vCenter/ESXi services run on custom ports. You must configure Discovery for these alternate ports, or it will fail. For more information, see Configure an alternate port for vCenter and Configure ESXi resource pools

> **Note:** Applications that access VMware cloud resources need access to VMware credentials. For more infomration, see VMware credentials .

Configure an alternate port for vCenter

Specify an alternate port for the VMware - vCenter datacenters probe.

Before you begin

Role required: Discovery admin or VMware admin responsible for vCenter and ESXi hosts

About this task

To establish proper connectivity for Discovery to vCenter, alternate ports are required if default ports are unavailable or customized.

By default, vCenter discovery initiates on these ports:

- **vmapp6_https**: 9443
- **vmapp_https**: 5480

When these ports are detected, the VMWare - vCenter datacenters probe is triggered.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Hard code the port from your ServiceNow instance.
   a. Navigate to **All** > **Configuration** > **VMware** > **vCenter**.
   b. Select the specific instance from the list of instances.
   c. Edit the **URL** field and include the port.
   For example, `https://10.0.0.1:444`
   d. Select **Update**.
2. Specify an alternate port number for vCenter in the VMWare - vCenter datacenters probe.
   a. Navigate to **All** > **Discovery Definition** > **Probes**.
   b. Open the VMware - vCenter Datacenters probe record.
   c. In the Probe Parameters related list, select **New**.
   d. In the **Name** field of the Probe parameter record, enter `vcenter_port`.
   e. In the **Value** field, enter your alternate port number.
   f. Select **Submit**.
3. Activate the VMWare - vCenter Classify probe.

   Important: This configuration is only required when vCenter is on port 443.

   a. Navigate to **All** > **Discovery Definition** > **Port Probes**.
   b. Open the http port probe record.
   c. In the Trigger Probes related list, select the **false** link for the VMWare - vCenter Classify classification probe.
   d. Select the **Active** check box.
   e. Select **Update**.
   The VMWare - vCenter Classify probe establishes a connection between the MID Server and the vCenter instance. If the connection is successful, it triggers the VMWare - vCenter datacenters probe.

Configure ESXi resource pools

The ESXi server has a default resource pool called Resources that defines normal resources for a virtual machine.

Before you begin

Role required: admin

About this task

ESXi resource pools must be configured so Discovery can identify virtual machines and their relationships correctly. Resource levels are dynamically generated from shares of the total resources allocated to virtual machines on the ESXi server. For details about how these resources are calculated, review the VMware knowledge base https://www.vmware.com/.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow Discovery finds this default resource pool and adds a record to the ESXi Resource Pools module automatically.

If Discovery isn't running on the ServiceNow instance, create a record for the **Resources** pool. Verify that the **Owner** field is correct and leave the resource fields empty. If a provisioner selects the **Resources** pool when provisioning a virtual server, the ESXi server creates a virtual machine for use under a normal load.

Procedure

1. Navigate to **All** > **VMware Cloud** > **Configuration** > **ESX Resource Pools**.
2. Select **New** in the list.
3. Create a record for each resource pool in the ESXi server, verifying that the **Name** and **Owner** fields are correct.
4. Select the **CPU expandable** and **Memory expandable** check boxes.

   These features enable for expansion of the CPU and memory limits when needed, if those resources are available on the ESXi server. When granted, these extra resources can be revoked to provision other virtual machines. The additional fields on the form are for display purposes only.

5. Select **Submit**.

ESX Resource Pool form

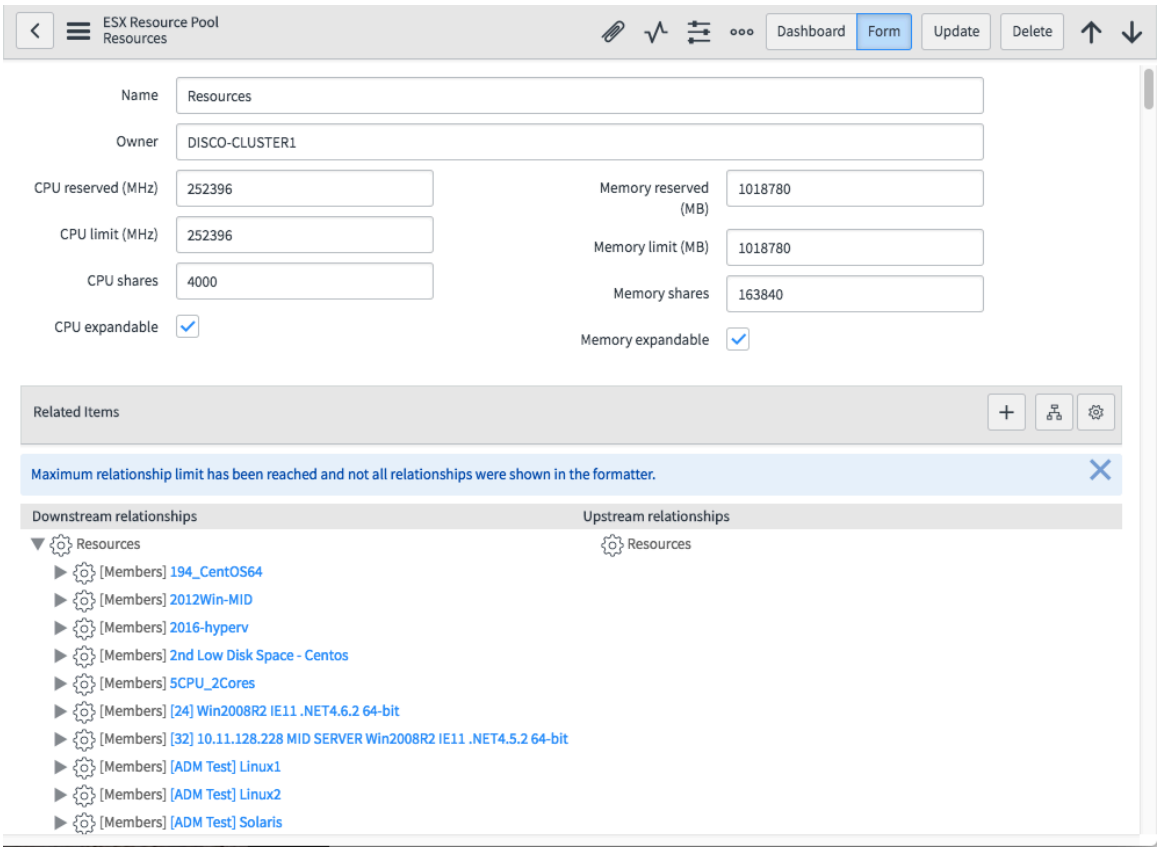

Discovery classifiers

A classifier tells Discovery which probes to trigger for the identification and exploration phases of discovery. Classifiers can also trigger the Horizontal Pattern probe, which launches a pattern, rather than additional probes, for identification and exploration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The classifier essentially starts the identification stage. Discovery uses it after the classification probe returns important parameters to the instance that tell Discovery what to do next.

In most cases, you do not need to create a classifier or modify a classifier. But if you are having trouble with horizontal discovery, you might want to check the conditions that determine when a classifier runs based on the parameters the classification probe returns to the instance. Or if you want to discover a new type of CI that Discovery does not already find, you can create your own classifier.

Device, process, and IP address classification

Discovery classification can be broken down into three types: device classification, process classification, and IP address (or IP scan) classification:

Device classification

The classification of actual device types, such as a computer running Windows, a computer running a flavor or UNIX or LINUX, a router, a switch, or a load balancer, and so on.

When Discovery identifies a computer CI, it triggers an active processes probe to explore the computer CI further. Discovery compares the results of the active processes probe to the process classification conditions to determine if there is a match.

Starting with the Madrid release, the horizontal discovery process can classify devices using HTTP.

Of all protocols that Discovery uses (including WMI, SSH, and SNMP), HTTP is the lowest priority by default. Discovery uses HTTP classification only if:

- Shazzam determines that the ports for HTTP (80) and HTTPS (443) traffic are open.
- The horizontal discovery process fails for the higher priority port probes (WMI, SSH, and SNMP) if the ports for those protocols are not open, or if discovery for those protocols fail. The horizontal discovery process can fail, for example, if SSH and SNMP credentials are not configured or are incorrect.

> **Note:** See Port Probes for more details on how port scanning works and to see the priorities of the different protocols.

HTTP classification launches the **HTTP Classify** probe to classify the device. The **HTTP - Classify** probe runs a GET request for each device for each HTTP classification. The URL of the request is built as follows: `PROTOCOL://IP:PORT/PATH`, where:

For more information on port probes, see Port probes. For instructions on creating an HTTP classifier, see Create an HTTP classification.

Process classification

The classification of applications based on the processes that are running.

Discovery classifies processes during the last phase of discovery: the exploration phase, after identifying devices in the Computer [cmdb_ci_computer] table and its extensions. Just like device classification, process classification has its own classification criteria and also has the ability to launch probes. Unlike device classification, process classification creates child configuration items (CI) with *Runs on::Runs* relationships. By default, Discovery includes classifications for most common processes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If a process matches the classification criteria, Discovery determines whether to run the process handler script. The process handler script modifies the parameter data to help Discovery identify whether the process represents an existing or new application CI. Discovery process handlers prevent the creation of duplicate CIs by filtering out parameters known to have inconsistent values before process classification occurs. Every time Discovery adds or updates an application CI, it also determines the application dependency mapping of the application CI to other CIs in the CMDB.

IP address (IP scan) classification

IP address discovery is credential-less, meaning that it attempts to identify devices and software based on just the open ports and banners it finds without requiring you to create credentials. If the classification criteria are met for a device in the IP Scan mode, Discovery automatically updates the CI in the CMDB. After a device is properly classified, Discovery launches the exploration probes configured for that class of device and begins gathering detailed information about the CI.

In the default Discovery system, the Linux classifier triggers eleven exploration probes that return information such as disk size, memory, and the number of current connections. The data from these probes returns at different times and is stored in the ECC Queue until processing is complete.

This diagram shows the processing flow for classifying and probing devices with an IP scan (no identifiers):

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP scan classification

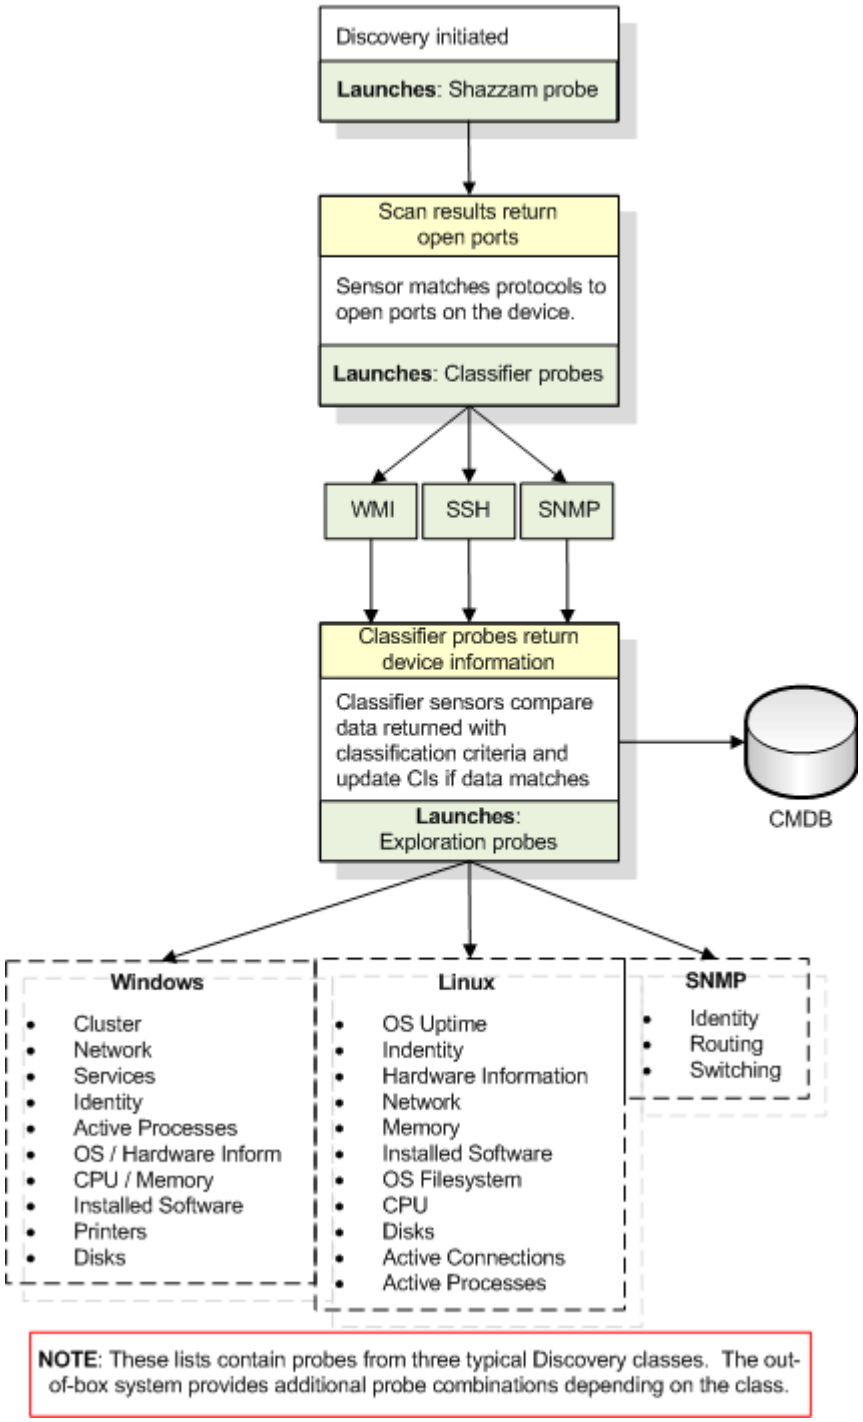

See Classification for IP address discovery for more details about the parameters available to classifiers for this type of discovery.

Classifier criteria

Classifiers also provide criteria that you can use to specify when Discovery should use the classifier under the conditions that you define. The criteria is based on the parameters that a classify probe returns to Discovery. Criteria is constructed with the parameter, an operator, and a value.

Note: Condition filters in process classifiers are case-sensitive.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifiers and patterns

Discovery can use patterns, rather than probes, to identify and explore CIs. Discovery triggers patterns from the Horizontal Discovery probe, which can be specified on a classifier. You can create you own patterns and add them, via the Horizontal Discovery probe, to a classifier. See Add the Horizontal Pattern probe to a classifier for instructions. You might already be using one of the out-of-box patterns that are provided with Discovery. You can verify this by looking at the classifier to see if the Horizontal Pattern Probe is specified.

Logging classification debugging information

To log debugging information about classifications, add the following system property. The resulting log entries list the name of each classifier that runs, along with all the names and values that are available to the criteria in the classifier.

System property	Description
`glide.discovery.debug.classification`	Enables debugging information for process classification. • **Type**: true | false • **Default Value**: false • **Location**: Add to the System Properties [sys_properties] table. For more information, see Add a system property.

What you can do with Discovery classification

- Create or modify a discovery classifier if you want to classify CIs that Discovery does not already classify, or trigger other probes that are not already on a classifier. You can modify classifiers that Discovery uses in standard CI discovery, process classifiers for applications, and classifiers based on IP address scans.

  Before you modify any classifiers, review the parameters that are available for each type of classifier.

- If Windows machines are on your network, you can use the WinRM protocol, rather than WMI, for more efficient lightweight data transfer and remote command execution. By default, Discovery uses WMI. For instructions on the classifier modifications you can make to use WinRM, see Use Windows Remote Management for classification.
- If you have Windows computers that are acting as servers and you want them to be classified by their function rather than by the operating system, you can make changes to the criteria of the Windows classifier. See Reclassify a Windows Workstation machine as a server for instructions.

Create a Discovery CI classification

A CI classification allows Discovery to discover most common operating systems, network devices, and processes.

Before you begin
Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Warning: Adding SNMP OID classifications with a Net-SNMP module OID starting with 1.3.6.1.4.1.8072.3.2 or a Windows OS OID starting with 1.3.6.1.4.1.311.1.1.3 can result in misclassification of devices and incorrect model names and is no longer allowed. These are template system OIDs that can remain on a device if the manufacturer fails to replace them with OIDs specific to their own model. For example, a Check Point Smart-1 3050 Firewall, which uses Linux as its OS, might have the OID for a Linux server, 1.3.6.1.4.1.8072.3.2.10. This results in the firewall being classified incorrectly. If two or more devices have the same OID, configure classification criteria that differentiate the devices from one another. For details, see KB0751287.

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification**.
2. Select the type of classification to create.
3. Click **New** in the list of classifications for the type you selected.
4. Complete the form using the fields in the table.

Classification form fields

Field	Input Value
Name	Name of the configuration item (CI).
OID	SNMP system OID for matching this device. When the OID value matches that of an SNMP device, Discovery uses the information to populate the CMDB with the specified manufacturer, model, and classifier.
Operator	Operator for determining how to match and SNMP OID. The choices are **Is** and **Starts with**. This field is available only for SNMP System OIDs classifiers.
Active	Enables or disables this classifier. When a classifier is disabled, the system stops classification at this level and does not launch classifiers of a lower order. For example, when the classifier for Windows 2008 Server is disabled, the system stops Discovery at this point and does not launch the Windows 2012 Server classifier.
Order	Order (sequence) in which the platform runs this classifier.
Table	Table for this classification. For example, if this record classifies a Linux server, select the Linux Server [cmdb_ci_linux_server] table.
Classifier	Classification of an SNMP device, such as **A10 Load Balancer**. This field is available only for SNMP System OIDs classifiers.
Relationship type	Type of relationship for this classifier, such as **Runs on::Runs**. This field is available only for application and process classifiers.
Match criteria	Criteria that must match to classify this device. The choices are**Any** of the parameters or **All** of the parameters.
Manufacturer	Name of the manufacturer of a network device. This field is available only for SNMP and SNMP System OIDs classifiers.
Model	Model number of a network device. This field is available only for SNMP and SNMP System OIDs classifiers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input Value
On classification script	Script that runs if classification criteria are met. Use this script to perform any special tasks after a device is classified. It is possible to use the `g_probe_parameters` hashmap from within a classification script to set probe parameters for any configured, triggered probes. For example, this code sets a '*node_port*' parameter to 16001 for all triggered probes: `(JS), g_probe_parameters['node_port'] = 16001; //`
Related lists	
Classification Criteria	Criteria formed from specific parameters and the values that they must contain to match devices that Discovery finds in the network with CIs in the CMDB. For a detailed discussion of these parameters, see Discovery classification parameters.
SNMP OID Classifications	Unique "fingerprints" of all the SNMP devices that ship with the base Discovery product. Users can add OIDs for SNMP devices not in this list. This related list is available only for SNMP devices.
Triggers probes	Exploration probes that Discovery launches to gather detailed information about a CI that it has classified in the network. If you want to use patterns for horizontal discovery, add the Horizontal Pattern probe in the **Probe** column, and then specify your pattern in the **Pattern** column. ⚠ **Warning:** Do not specify your pattern in **Probe** column. Choose the Horizontal Pattern probe, which launches the specified pattern.
Versions	Lists versions of this classifier. A new version is created whenever you modify the classifier record. To revert to a previous version, open that record and select **Revert to this version** under **Related Links**.

This example shows a completed CI classification form with the Horizontal Pattern probe defined. For instruction on creating probes, see Discovery probes and sensors. Probes that are

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

defined are launched when the device is properly classified, unless Discovery is configured to stop after classification.

Discovery classification form

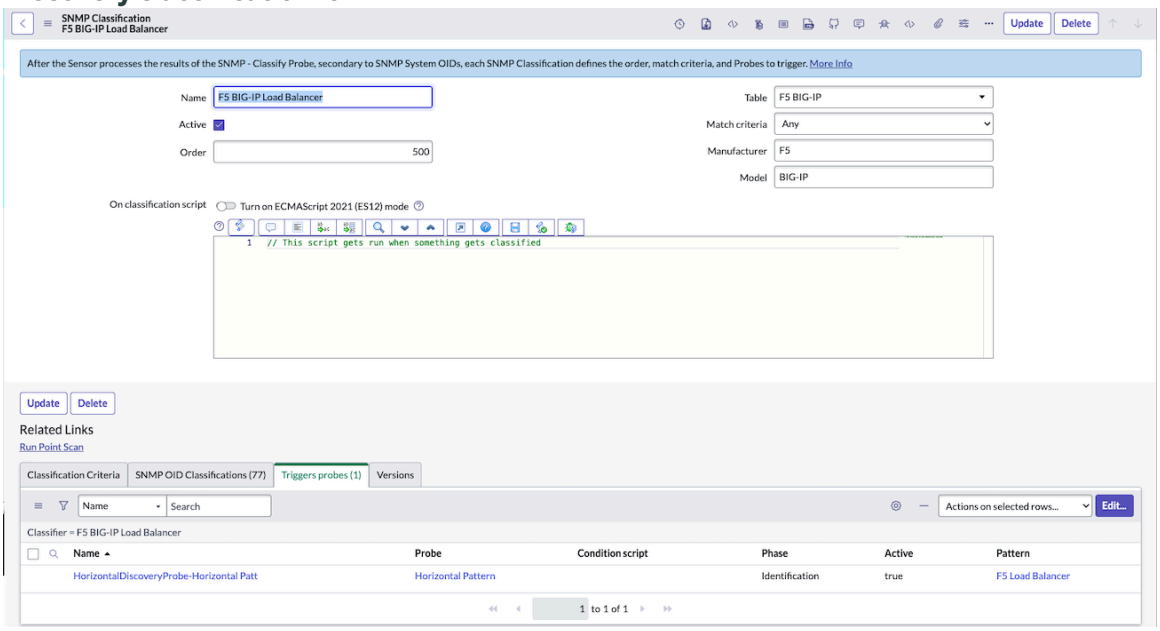

Create a Discovery process classification

A process classification allows Discovery to create a particular CI type from information gathered during the identification and exploration phases.

Before you begin

Role required: none

About this task

When a process matches the classification criteria, Discovery uses the process classification record to create a CI. You can also have Discovery update existing CIs or ignore certain processes by creating a process handler.

Note: There is no support for dot-walk on the parent running the process command.

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **Processes**.
2. Click **New**.
3. Enter the classification fields (see table).

Classification fields

Field	Input value
Table	Select the table where this classification generates CI records. This table must be an extension of the Applications [cmdb_ci_appl] table.
Relation type	Select the CI relationship type for this classification. The relationship field is only available for Process and Scan Application classifications. Discovery process classifications typically use one of these relationship types:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input value
	○ *Runs on::Runs*: Defines the relationship of an application to the host on which it runs. This relationship is expressed from the perspective of the host and the application. For example: My database application runs on server001::server001 runs my database application. ○ *Depends on::Used by*: Defines the relationship of an application that communicates with another application. This relationship is expressed from the perspective of each application. For example: The Tomcat application depends on the MySQL database:: The MySQL database is used by Tomcat. ○ *Virtualized by::Virtualizes*: Defines the relationship of a virtual machine to its host. This relationship is expressed from the perspective of the virtual machine and of the host. For example: server001 is virtualized by Server ESX::Server ESX virtualizes server001.
Active	Select this option to enable the process classification record. Only active process classifications can create application CI records.
Order	Enter the order in which Discovery should run this process classification when there are multiple classifications available for a table. Discovery runs process classifications from the lowest to highest order.
Test with	Lists the host CI where an automatically generated process classification conditioncons were met. Use this field to test changes to the process classification to ensure that the updated classification behaves as expected.
Condition	Use the condition builder to create the match and classification criteria for the process classification. This field replaces both the **Match criteria** field and **Classification Criteria** related list. The upgrade process converts all existing classification criteria into conditions. **Note:** Condition filters in process classifiers are case-sensitive. **Note:** Only these conditions work: ○ name ○ PID ○ PPID ○ command ○ parameters ○ listening_on_ports ○ connecting_to_ports ○ connecting_to ○ key_parameters ○ classify
On classification script	Enter a script to run when the condition and classification criteria are met. Use this script to perform any special tasks after a device is classified. It is possible to use the `g_probe_parameters` hashmap from within a classification script to set probe parameters for any configured, triggered probes. For

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input value
	example, this code sets a `'node_port'` parameter to 16001 for all triggered probes. `g_probe_parameters['node_port'] = 16001;` See On classification script objects for Discovery for more examples and for a list of the objects you can use in this kind of script.
Triggers probes	Select the exploration probes you want Discovery to launch. These probes gather detailed information about a classified CI. Discovery will not launch these probes when it is configured to stop after classification. If you want to use patterns for horizontal discovery, add the Horizontal Pattern probe in the **Probe** column, and then specify your pattern in the **Pattern** column. ⚠ **Warning:** Do not specify your pattern in the **Probe** column. You must choose the Horizontal Pattern probe which launches the specified pattern.
Applications	Use this related list to view the application CIs that match this process classification.
Test results	Use this related list to view the how Discovery classifies processes on the **Test with** host and build better classification conditions.
Parameters	Use this related list to view the parameters associated with this process and build better classification conditions. See Discovery classification parameters for a list of the parameters you can use.
Versions	Use this related list to view previous versions of the process classification record.

4. Right-click the header bar and click **Save**.
5. Enter items from related lists.
6. Click **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example:

An example process classification record

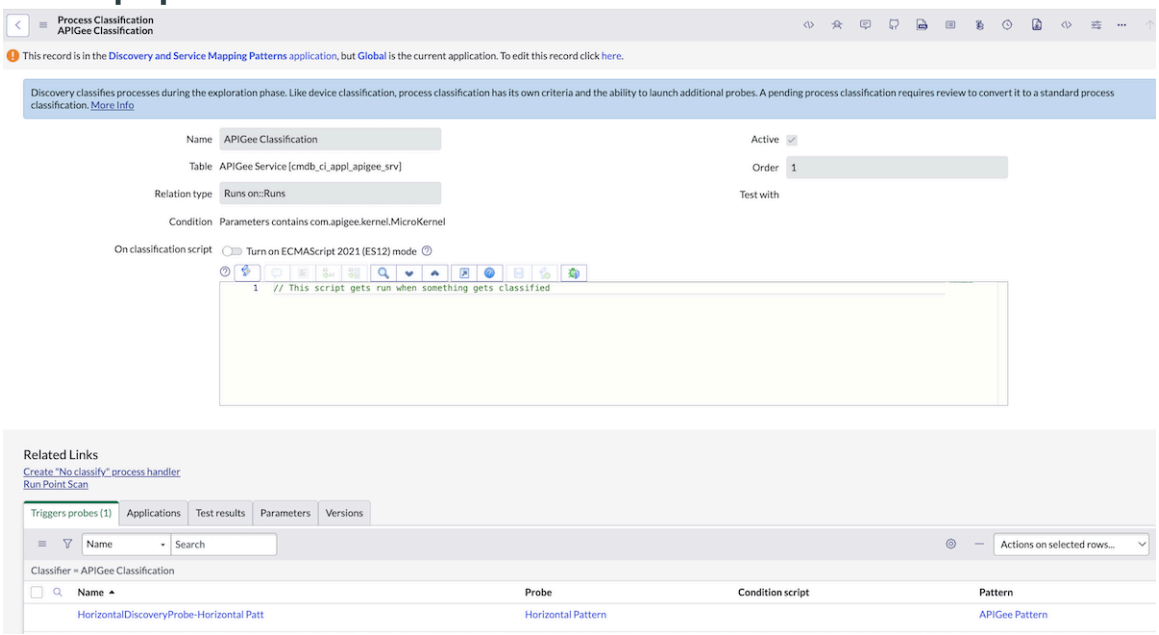

What to do next

Run a discovery through the Discovery Schedule to search for applications. You can also create or modify process handlers that prevent the creation of duplicate CIs based on process classifiers if that is an issue for your environment.

On classification script objects for Discovery

Use an On classification script in a process classifier to customize an application record.

Renaming the default application name

By default, application names are in this format: `<name of the process classifier>@<the name of the computer CI where the process resides>;`

For example, for a MySQL server running on a computer called machineA, the application is named `mysql@machineA`.

You can use the **On classification script** field in the process classifier record to change the default application name to match your business needs. For example, the following script changes the default application name to include a suffix after the process classifier:

```
var computerName = g_sensor.deviceGR.name;
var processClassifierName = g_classification.name;
current.name = processClassifierName + "999" + "@" +
 computerName;
```

In this example, the name of the application record becomes `mysql999@machineA`.

Another common technique is to set the application name based on the name, command, and parameter variables. For example, an Eclipse process might have the following values in these variables:

name	"eclipse"
command	"/glide/eclipse/Eclipse.app/Contents/MacOS/eclipse"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

parameter	"-psn_0_1884620"

If an Eclipse application runs on a computer called machineA, ServiceNow names the application eclipse@machineA. The following script appends the parameter value as part of the application name.

```
var computerName = g_sensor.deviceGR.name;
var processClassifierName = g_classification.name;
current.name = processClassifierName + parameters + "@" +
 computerName;
```

In this example, the name of the application record becomes eclipse-psn_0_1884620@machineA.

Sometimes it is useful to pass values to the triggered probes in the process classification. You can do this by creating a custom script that defines a name/value pair for the `g_probe_parameters` object. For example:

```
g_probe_parameters['processCommand'] = command;
```

In this example, when a classification record triggers a probe, the script passes the probe a parameter called `processCommand` with the value of the `command` variable.

Script objects

Use these objects in the script:

Script object	Description
current	Points to a JavaScript object with its [property:value] pair to update the application record. It is not an actual GlideRecord object of the application.
g_sensor	Points to the DiscoverySensor object defined in the DiscoverySensor script include. This object contains a deviceGR object, which points to the computer CI record on which the process resides.
g_classification	Points to the process classifier record. This object is set in the Application Dependency Mapping sensor and is available in the **On Classification script** field. Use it to access the classified process properties: name, command, parameters, PID, and table name.
name	Points to the process name.
command	Points to the process command.
parameters	Points to the process parameters.
g_probe_parameters	A JavaScript object that enables parameter passing to the triggered probes.

Related topics

Create a Discovery process handler

Create a Discovery process handler

Process handlers prevent the creation of duplicate CIs by filtering out parameters known to have inconsistent values before process classification occurs. You can create new classifiers or edit existing ones.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: none

Procedure
1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **Process Handlers**.
2. Click **New**.
3. Complete the fields on the Process Handler form.

Process Handler form

Field	Description
Name	Enter a unique name for the process handler record.
Active	Select the check box to run the process handler.
Classify	○ Select the check box to enable classification of any Running Process [cmdb_running_process] record that matches this Process Handler's conditions. ○ Clear the checkbox to disable classification of any Running Process [cmdb_running_process] record that matches this Process Handler's conditions.
Condition	Select the conditions that trigger the process handler to run the script. In most cases, this condition contains either specific executable names or the presence of certain parameters.
Script	Enter the JavaScript you want to run on the current record in the Running Process [cmdb_running_process] table when the **Condition** is met. The `current` variable is a reference to a Running Process GlideRecord. The script should examine `current.parameters` for certain values, perform string replacement to manipulate these values, and save the result to `current.key_parameters`. ServiceNow uses the **key_parameters** field, together with the process name, to determine whether the process is unique on a specific machine. See On classification script objects for Discovery for additional options.

4. Click **Submit**.

What to do next
Run a discovery through the Discovery Schedule to search for applications and then verify the results in the CMDB for the applications that are found.

Create an HTTP classification

An HTTP classifier enables the horizontal discovery process to find devices via the HTTP protocol.

Before you begin
Role required: discovery_admin

About this task
If you want to classify devices via HTTP/REST, you can create additional HTTP classifications, rather than by another protocol. Use HTTP classification for device discovery only, not for application discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: The HTTP Classify probe no longer attempts credentials over the HTTP protocol by default. To override this behavior, you can enable *mid.http_classy.allow_credentials_over_http*. However, enabling this setting can expose credentials to man-in-the-middle (MitM) attacks. Therefore, it's strongly recommended to keep this property set to **false** and use HTTPS whenever possible.

By default, the F5 HTTP classifier is provided to classify F5 load balancers using REST. The **F5 LB REST** pattern is also provided to explore the device. You don't need to create an HTTP classifier for F5 load balancers.

Of all protocols that Discovery uses (including WMI, SSH, and SNMP), HTTP is the lowest priority by default. Discovery uses HTTP classification only if:

- Shazzam determines that the ports for HTTP (80) and HTTPS (443) traffic are open.
- The horizontal discovery process fails for the higher priority port probes (WMI, SSH, and SNMP) if the ports for those protocols are not open, or if discovery for those protocols fail. The horizontal discovery process can fail, for example, if SSH and SNMP credentials are not configured or are incorrect.

Note: See Port Probes for more details on how port scanning works and to see the priorities of the different protocols.

HTTP classification launches the **HTTP Classify** probe to classify the device. The **HTTP - Classify** probe runs a GET request for each device for each HTTP classification. The URL of the request is built as follows: `PROTOCOL://IP:PORT/PATH`, where:

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **HTTP**.
2. Select **New**.
3. Fill out the form fields (see table):

The F5 HTTP classification

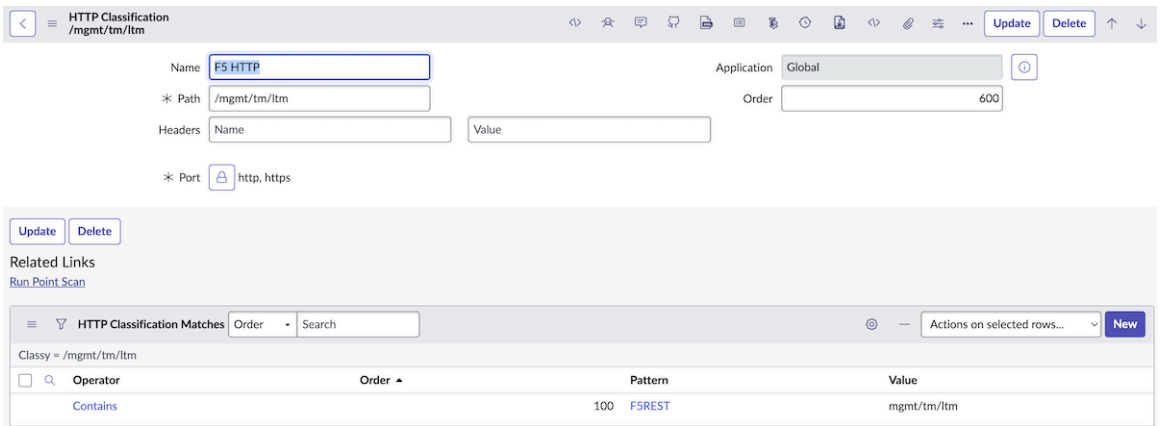

Field	Description
Name	Enter a name for the classifier.
Path	Enter the path for the HTTP GET request.
Headers	Add an option header with a value. You can use a variable in the format `${variable}` instead of a static value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Supported headers depend on the REST web service provider that you want to connect to. See the documentation for your web service provider to identify which headers are valid or required.
Port (IP Service)	If you want this classifier to use a custom port that is not already a default value in the IP Service [cmdb_ip_service] table, do the following: **a.** Select the lookup icon to see the list of IP Services. **b.** Select **New** and create a port-to-protocol mapping. **c.** Select that port-to-protocol mapping for this **Port** field. Many commonly used protocols are already mapped to ports by default, such as `http` to port 80 and the `https` to port 443. If you want to use one of these default port-to-protocol mappings, don't enter any values in the **Port** field. The HTTP port probe uses the correct one by default.
Order	Enter a number that represents the priority of this classifier. If Discovery finds more than one classifier that applies, it uses the classifier with lower **Order** number.
Protocol [optional] This field isn't visible by default. You must configure the form to add it.	If you configure a unique port-to-protocol mapping for the **Port** field, enter an optional protocol override that Discovery uses with this classifier. For example, if you want to use port 5000 and the HTTP protocol, add the 5000/HTTP port-to-protocol mapping to the **Port** field, and then choose **http** in the **Protocol** field.

4. Right-click the header and select **Save**.

   A record is created in the HTTP Classifications [discovery_classy_http] table.

   The HTTP Classification Match related list appears. This related list shows criteria that you can configure to match against the body of the response to the GET Request that Discovery sends.

5. In the HTTP Classification Match related list, define the criteria and specify the pattern to launch for this classifier.
   The criteria must be met for Discovery to use this HTTP classifier.

   a. Select **New**.

   b. On the HTTP Classification Match form, fill in the form fields (see table).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

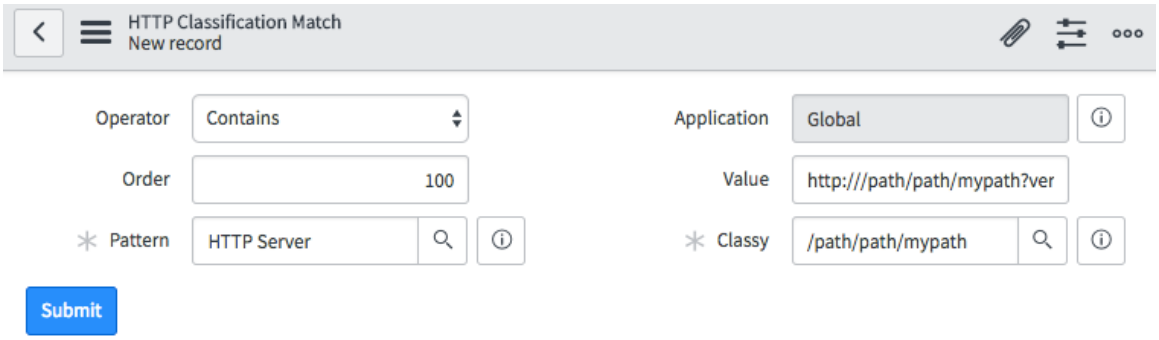

Field	Description
Operator	Choose the operator that the system uses to evaluate the condition: ▪ **Contains** ▪ **Equals** ▪ **Regex**
Order	Enter a number that represents the priority of this condition. If Discovery finds more than one condition that applies, it uses the condition with lower Order number.
Pattern	Specify the pattern to launch upon a correct match. This pattern must use HTTP to discover the device.
Value	Enter the value that the body of the returned REST method must have.

c. Select **Submit**.

d. If there are different patterns for different versions of the device you're trying to discover, define as many HTTP Classification Matches as necessary.

Classification for IP address discovery

Discovery provides a way to classify devices it finds through IP address discovery, even when no credentials are available.

When you run a discovery for IP addresses, as opposed to a configuration item (CI) discovery, the Discovery application makes certain assumptions about devices and the applications running on those devices from the ports that it finds open. Classification parameters for this type of Discovery are generated differently from scans in which credentials are available.

The syntax for creating parameters is derived from the fields returned by the Shazzam probe when conducting a Discovery for IP addresses. Parameters for CIs and applications are formed in the same way. The Shazzam probe creates an XML file containing the following fields:

- **name**
- **port**
- **portprobe**
- **protocol**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **result**
- **service**

Note: Optional fields that can be used to form parameters appear as child tags beneath the default fields. Example of these are the **sysDescr** and **banner_text** fields.

Parameters are expressed in the form of <portprobe.service.field>. The value for *field* can come from any of the fields or child tags in the XML file. For example, the following parameters classify a device as a UNIX server and detect an installation of MySQL:

- `ssh.ssh.result`
- `mysql.mysql.result`

These parameters were derived from the values in the following XML file generated by a Shazzam probe conducting an IP Scan. The *result* field returned a value of *open* for ports 22 and 3306 on the target device. The *service* field indicates the services that normally communicate over those ports.

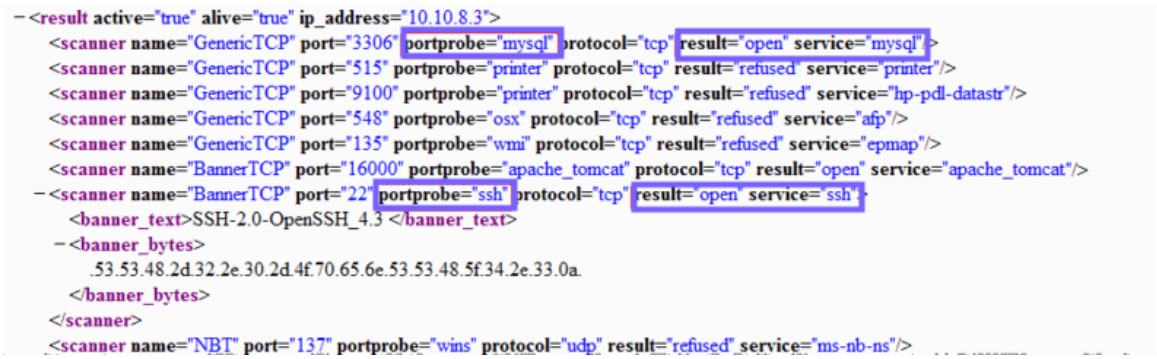

The *sysDescr* field can provide additional information about devices, depending upon the manufacturer. This XML file from the Shazzam probe reveals the following about port 161 on the device at IP 10.10.11.149:

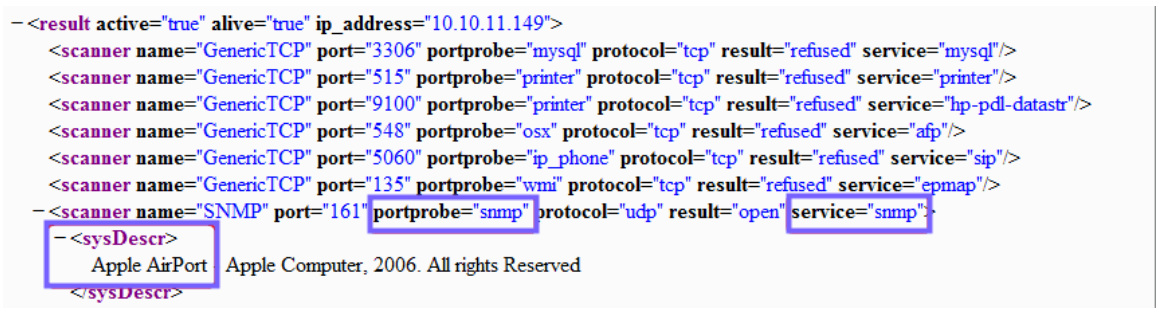

In the classification criteria, we can construct the following parameter with **sysDescr** that returns an Apple AirPort wireless router:

```
snmp.snmp.sysDescr  contains  Apple AirPort
```

Modify classifiers for IP address discovery

When you run an IP address type of discovery, port probes scan devices without the use of credentials, and then Discovery can determine which classifiers to use. You can add port probes and additional classifiers for IP address discovery.

Before you begin

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Although no credentials are required to scan for Windows or UNIX devices with this type of scan, credentials are still required for SNMP devices. By determining which ports are open on the devices that it scans, IP address classification can discover such things as the type of device (computer, UPS, etc.), operating system, running applications, and version numbers.

Note: IP address classification attempts to classify devices when no credentials are available; however, Discovery will use credentials when they are available, even when IP address classification is configured.

To use IP address classification, follow these steps:

Procedure

1. Determine what ports to use for classification and run a scan program such as *Nmap* on specific IP addresses to decide which ports reveal the desired information about a device or application.

   The scan can reveal several pieces of data that are useful for configuring classification parameters. An Nmap scan displays port numbers, their state (open or closed), their service names, and any version information it can find. From the port information returned in the example below, we can construct criteria to classify UNIX servers (port 22), MySQL (port 3306), and Apache Tomcat (port 16000).

```
Host is up (0.035s latency).
Not shown: 65521 closed ports
PORT      STATE    SERVICE     VERSION
22/tcp    open     ssh         OpenSSH 4.3 (protocol 2.0)
| ssh-hostkey: 1024 d9:77:8b:68:1a:3a:b3:80:c3:a4:60:ba:61:7c:6f:65 (DSA)
|_2048 12:ed:36:6c:aa:dc:51:a0:ce:59:45:c9:c8:42:9b:85 (RSA)
80/tcp    open     http        Apache httpd 2.2.3 ((Red Hat))
|_html-title: Test Page for the Apache HTTP Server on Red Hat Enterprise Linux
111/tcp   open     rpcbind     2 (rpc #100000)
| rpcinfo:
| 100000  2    111/udp  rpcbind
| 100024  1    655/udp  status
| 100000  2    111/tcp  rpcbind
|_100024  1    658/tcp  status
443/tcp   open     ssl/http    Apache httpd 2.2.3 ((Red Hat))
|_html-title: Test Page for the Apache HTTP Server on Red Hat Enterprise Linux
658/tcp   open     status      1 (rpc #100024)
1720/tcp  filtered H.323/Q.931
2000/tcp  filtered cisco-sccp
3306/tcp  open     mysql       MySQL (unauthorized)
5060/tcp  filtered sip
10000/tcp open     http        Webmin httpd
|_html-title: Site doesn't have a title (text/html).
|_http-favicon: Unknown favicon MD5: 5B9D90CA66BB985FA0D070E0E02D93F4
16000/tcp open     http        Apache Tomcat/Coyote JSP engine 1.1
16001/tcp open     http        Apache Tomcat/Coyote JSP engine 1.1
16003/tcp open     http        Apache Tomcat/Coyote JSP engine 1.1
| robots.txt: has 1 disallowed entry
|_/
|_html-title: DBSymphony
```

2. Add an IP Service and port probe.

   The out-of-box system supplies probes for some of the most common ports, but additional port probes will be needed for effective IP address scanning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **Discovery Definition** > **IP Services** and click *New*.

b. Create a new IP Service record using the port number and service from the Nmap scan. In this example, we associate the *mysql* service with port 3306 and add the CI (*sanops02*) on which the service runs to the *Available on* Related List.

Discovery IP Service

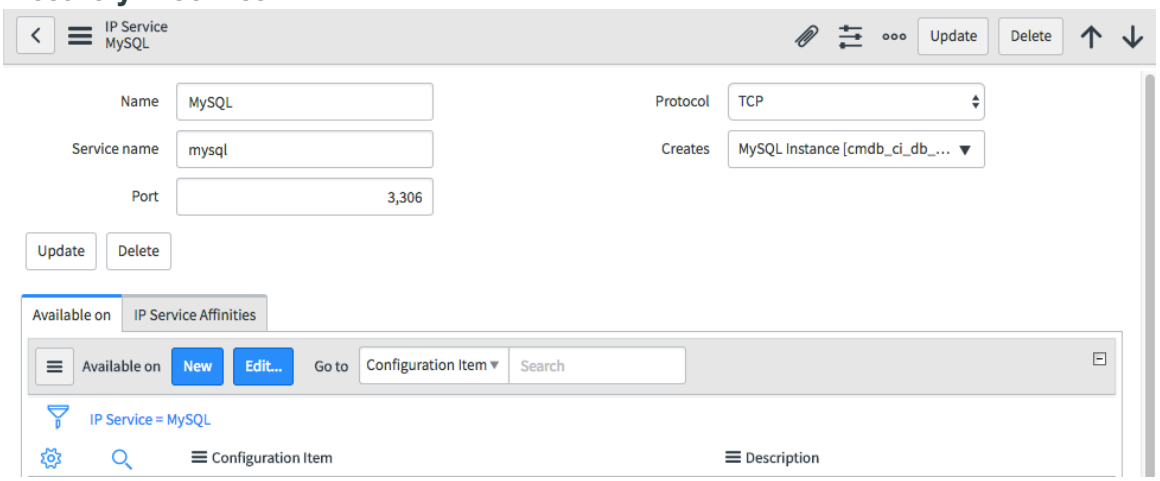

c. To use Basic Discovery, navigate to **Discovery Definition** > **Functionality Definition** and select the record for **All**.

d. Add the new port probes to the list. This tells Discovery which port probes to run for IP address scans.

Discovery functionality definition

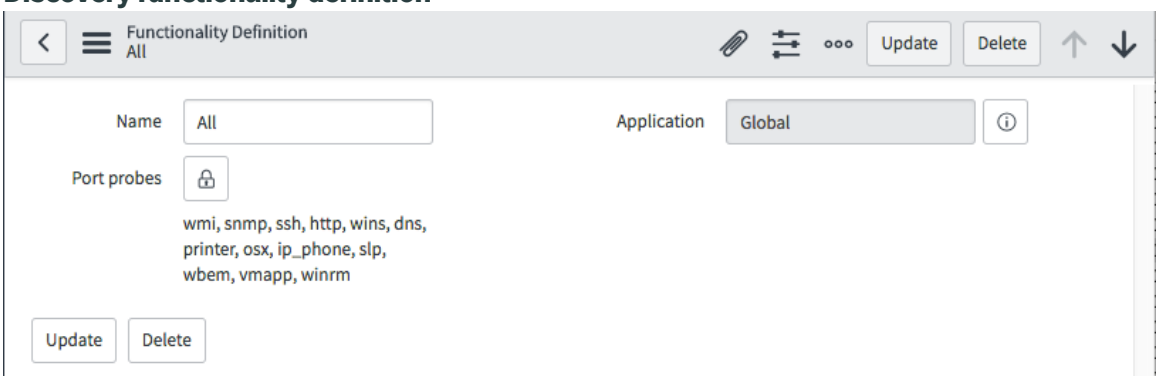

e. Save the record and navigate to **Discovery Definition** > **Port Probes** and click **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

f. Create a port probe using the new IP Service you just defined.

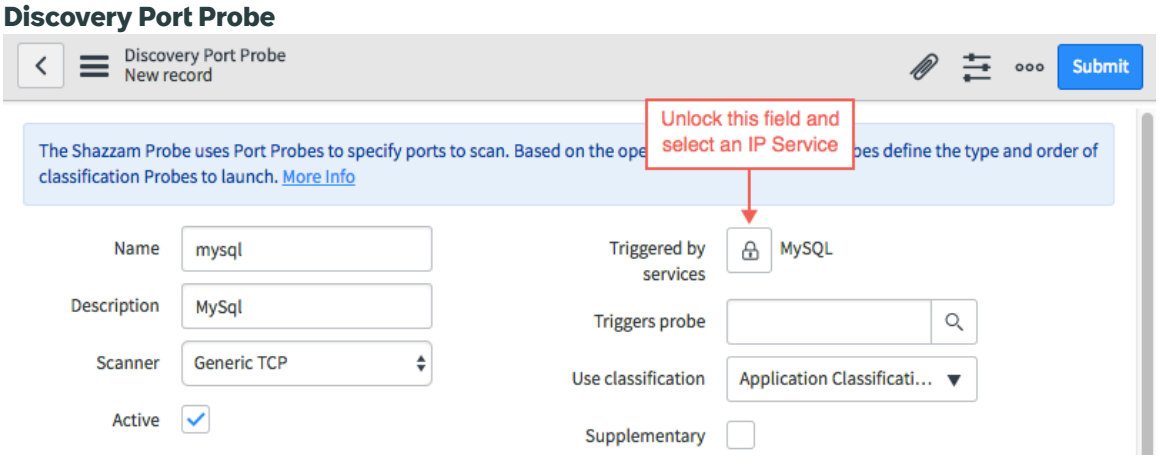

3. Create a new classification and add the parameter for IP address scanning.

In this example, we have created an application classifier that will discover Apache Tomcat, based on the port information we received from the Nmap scan. See the following section for details about forming parameters for IP address scans.

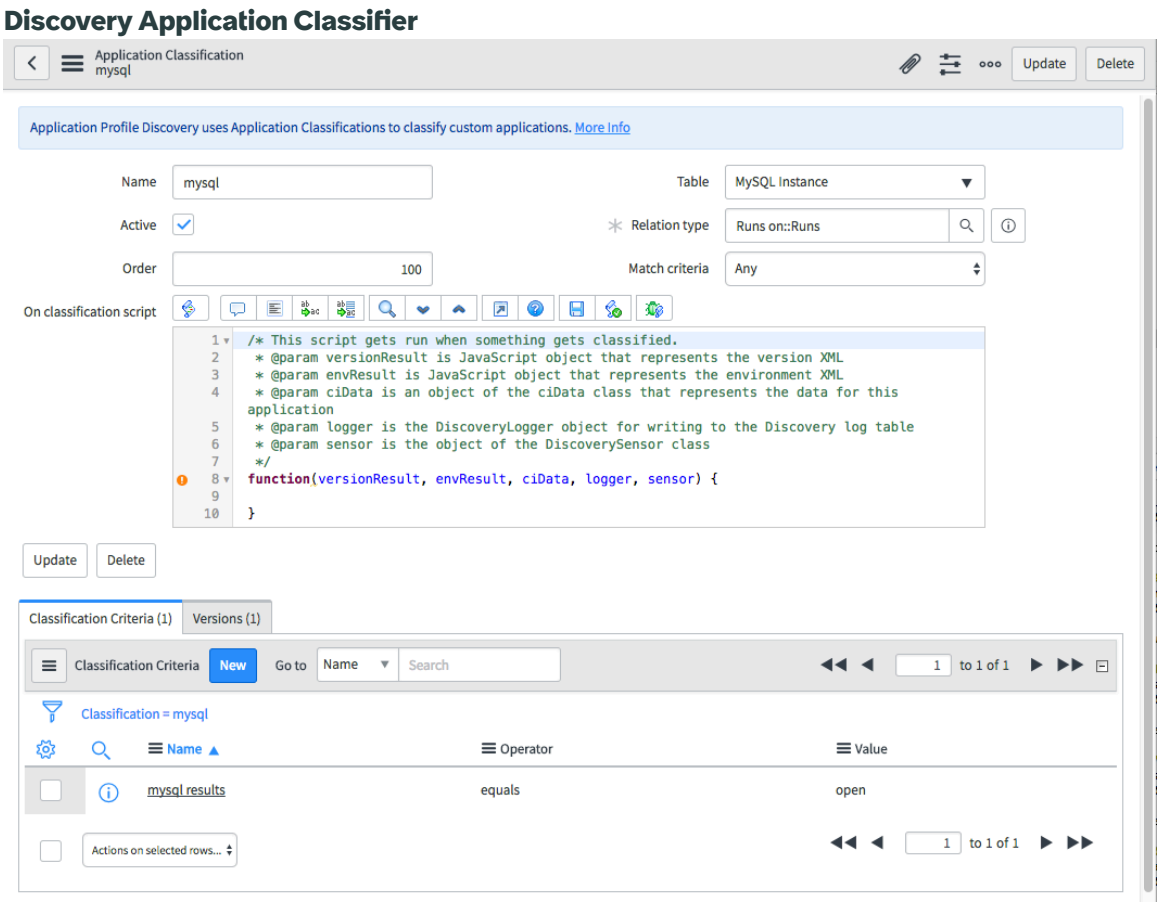

4. Optional: In the Classification Criteria related list, create a criteria filter that determines when this classifier applies to the discovered devices.
See the IP address classification parameters for a list of the parameters you can use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next
Run a IP address discovery through the Discovery Schedule to search for devices.

Use Windows Remote Management for classification

You can configure the discovery of Windows hosts using the Windows Remote Management (WinRM) protocol.

Before you begin

- Install and configure a Windows MID Server on the local network.
- WinRM service must be enabled on all discoverable Windows hosts.

Roles required: discovery_admin, agent_admin, admin

About this task
By default, the system uses the WMI protocol for device classification of Windows hosts. Administrators can instead use the WinRM protocol for more efficient lightweight data transfer and remote command execution.

Procedure

1. Enable the WinRM service on all Windows hosts you want to discover.
2. Navigate to **Discovery** > **MID Servers**.
3. Select the MID Server you will use for discovery of Windows hosts.
   The system displays the MID Server record.
4. From the Configuration Parameters related list click **New**.
   The system displays a MID Server Configuration Parameter record.
5. In the **Parameter name** field, select the **mid.windows.management_protocol** parameter from the choice list.
6. Enter a value of **WinRM**.

MID Server configuration parameter

7. Click **Submit**.
   The system displays the MID Server record.
8. **Optional:** Add other Windows Remote Management protocol parameters as needed.

Additional WinRM parameters

Parameter	Description	Default	Requires MID Server restart
mid.powershell_api.winrm.port	Specifies the communications port the MID Server uses to communicate with the WinRM protocol.	5985	No

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description	Default	Requires MID Server restart
mid.powershell_api.session_pool.target.max_size	Specifies the maximum number of sessions allowed in the pool per target host.	2	Yes
mid.powershell_api.session_pool.max_size	Specifies the maximum number of sessions allowed in the session pool.	25	Yes
mid.powershell_api.idle_session_timeout	Specifies the timeout value of idle Powershell sessions in seconds.	60	Yes
mid.powershell_api.winrm.additional_session_options	Specifies additional WinRM PSSession options the MID Server applies when creating a session.	None	
mid.powershell_api.winrm.always_taskkill	Specifies whether the MID Server always sends a TaskKill command to close the PowerShell process. Enable only when unexpected PowerShell processes remain when used with WinRM.	false	
mid.powershell_api.winrm.https_port	Specifies the HTTPS port that WinRM uses to connect to remote hosts.	5986	
mid.powershell_api.winrm.skip_ssl_check	Specifies whether the MID Server skips the SSL certificate check with WinRM.	false	
mid.powershell_api.winrm.skip_ssl_check_options	Specifies the options used to skip the SSL certificate check with WinRM.	`-SkipCACheck -SkipCNCheck -SkipRevocationCheck`	
mid.powershell_api.winrm.use_reverse_dns_lookup	Specifies whether the MID Server uses the FQDN when the IP address is not in the list of TrustedHosts for WinRM.	true	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description	Default	Requires MID Server restart
mid.powershell_api.winrm.use_ssl	Specifies whether the MID Server uses SSL with WinRM.	false	
mid.powershell_api.winrm.transport_mode	Specifies the WinRM working mode. Possible values: http, https.	http	

What to do next
Run a discovery from the Discovery schedule to find Windows machines on your network.

Reclassify a Windows Workstation machine as a server

By default, Discovery automatically classifies computers using certain Windows operating systems as workstations. However, you might want specific computers in your network that are acting as servers to be classified by their function and not their operating system.

Before you begin
Role required: none

About this task
Use the following variables, preceded by **cidata**, to construct a reclassification condition. For example, to reclassify based on a machine's IP address, use *`cidata.ip_address`*.

- name
- dsn_domain
- os_domain
- ip_address
- serial_number

The following procedure reclassifies any Windows workstation operating system (Windows Vista, XP, or Windows 7) that is acting as a server.

Procedure
1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **Windows**.
2. Create a new classification record, such as Windows XP Server.
3. Select Windows Server `[cmdb_ci_win_server]` as the table.
4. Right-click in the header bar and select **Save** from the context menu.
   The **Classification Criteria** and **Triggers Probes** Related Lists appear.
5. Configure the following **Classification Criteria**:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

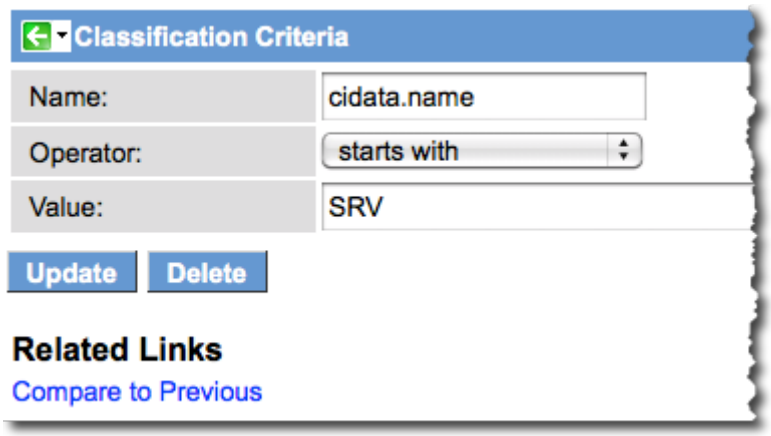

Name	Select a variable to use as the classification criteria from the list above. For example, to reclassify a machine by name, enter `cidata.name`. This works for servers that have a uniform naming convention, such as SRV001, SRV002, etc., regardless of operating system.
Operator	Select an operator for the classification condition. In networks containing servers named with a specific convention, you might select **starts with** or **contains**.
Value	Enter the value for the condition. In our example of a network with a server naming convention, this value would be the root of that convention, such as **SRV** This condition will classify all computers as servers if their machine name is SRVXXX.

6. Select the **Triggers Probe** related list and add the appropriate probes.

   a. Copy the list of probes from another Windows server classification, including the *`Condition`* scripts.

   b. Ensure that the **Windows - Identity** probe has a phase of **Identification** (the default is **Exploration**).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The completed form looks like this:

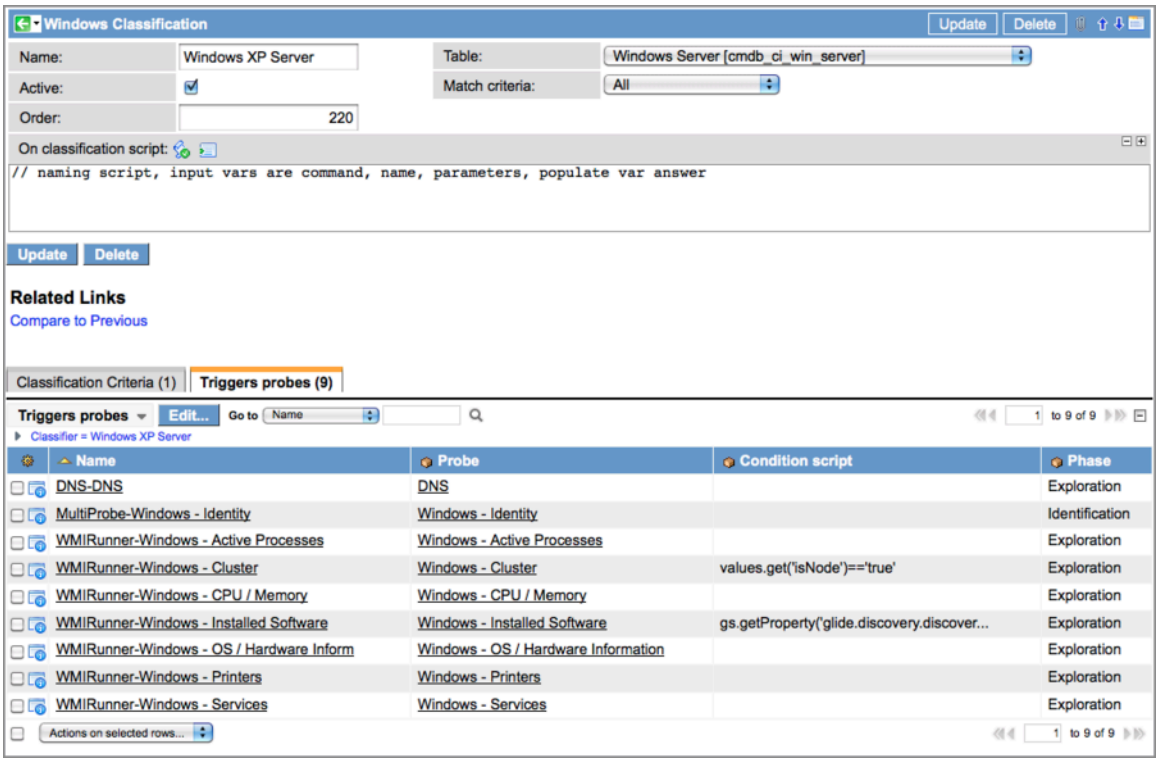

What to do next

Run a discovery from the Discovery schedule to find Windows machines on your network, and then check the cmdb_ci_win_server table and related tables to see how data is populated in the CMDB.

Discovery classification parameters

These parameters are available for Discovery classifiers.

Note: Condition filters in process classifiers are case-sensitive.

Unix parameters

The UNIX parameters define the characteristics of several types of computers, such as Linux, Solaris, and HP-UX, communicating with SSH protocol, version 2.

Parameter	Description
output	The raw output of the classifier probe (uname -a).
type	Returns the classification UNIX.
IP address	Returns the IP address of the device being discovered.
name	Name of the operating system for this UNIX CI. For example, Linux or HP-UX.

Windows parameters

Windows parameters identify Windows computers communicating with the WMI protocol.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
isNode	Indicates if this instance is a node.
type	Returns the classification *windows*.
isVIP	Indicates if this CI is a virtual machine, with a virtual IP address.
ip_address	Returns the IP address of the device being discovered.
name	Name of the Windows version.
cidata.name	Hostname of the Windows OS

SNMP parameters

The SNMP parameters can define the characteristics of several types of devices, such as routers, switches, and printers.

Parameter	Description
powering	A value of *true* indicates that this device is an uninterruptible power supply (UPS).
hosting	A value of *true* indicates that this device can host programs. Hosts are general purpose computers such as servers.
netware	A value of *true* indicates that this device is running the Netware operating system.
routing	A value of *true* indicates that this device has network routing capabilities.
ip_address	Returns the IP address through which the device is being discovered. A device can have multiple IP addresses.
sysdescr	Required descriptive field on any SNMP device that can contain useful classification data, such as the operating system and its version.
vlans	A value of *true* indicates that this device can host a virtual local area network.
hint_router	A value of *true* indicates that Discovery has determined that this device is a router. This field only applies to devices that can be used as both a router and a switch.
block_router_exploration	If this parameter is *true*,Discovery will not launch exploration probes for routers it detects. This parameter is used for network Discovery only.
switching	A value of *true* indicates that this device has network switching capabilities.
mfr_apc	A value of *true* indicates that this device is an uninterruptible power supply (UPS) manufactured by American Power Conversion (APC).
printing	A value of *true* indicates that this device has printing capabilities.
block_switch_exploration	If this parameter is *true*, Discovery will not launch exploration probes for switches it detects. This parameter is used for network Discovery only.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Process parameters

Process parameters identify processes such as those used by LDAP, Apache Server, and JBoss Server.

Parameter	Description
parameter	The parameters used to run the process.
command	The command used to run the process.
output	The complete output of the current line of the process probe.
type	Indicates the process type (e.g. *unix* or *windows*).
PID	The *process ID* generated by the operating system of a device to identify a running process. Generally, this parameter is not a practical classification criteria, because the value does not remain static, except in the case of processes running on an appliance that is never restarted.
name	Name of the process being discovered. In some cases, this parameter is not reliable, since several process might be given the same name. In Windows, for example several processes return *scvhost.exe* for this parameter.

Check IP service affinity for Discovery and Orchestration

You can check the IP Services table for a list of IP addresses that are associated with a protocol.

Before you begin

Role required: admin

About this task

The IP Services table maps a port to a protocol. Several mappings are provided by default for commonly used port-protocol combinations, such as port 80 for HTTP, port 22 for SSH, and port 161 for SNMP.
A system property called `glide.discovery.ip_service_affinity` allows Discovery to remember the last port of the IP address that was discovered.

Important: You should not modify IP services unless your organization uses custom ports.

Procedure

1. Navigate to **All** > **Discovery Definition** > **IP Services**.
2. Filter the list to find the appropriate IP service.
3. Select the name of the service to go to that IP service page.
4. Select the **IP Service Affinities** tab for the list of IP addresses associated with that service.

Discovery identifiers

After Discovery classifies a configuration item (CI), it uses identifiers to determine if the device already exists in the Configuration Management Database (CMDB).

Discovery launches special identity probes that accumulate identification data for each device and feed that data into the identifiers. The identifiers determine the action that Discovery must take for each device. Identifiers accurately determine the identity of the device to avoid the creation of duplicate CIs. This identification step only takes place for the Configuration item type of discovery, not for the other types of discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The identity probe in the base Discovery system can be configured to ask the device for information such as its serial numbers, name, and network identification. The results of this scan are processed by an identity sensor, which then passes the results to the identifier. The identifier then attempts to find a matching device in the CMDB. If the identifier finds a matching CI, the identifier either updates that CI or does nothing. If the identifier cannot find a matching CI, it either creates a new CI or does nothing. If Discovery is configured to continue, the identifier launches the exploration probes configured in the classification record to gather additional information about the device. Exploration probes can be multiprobes or simple probes.

Important: Serial numbers are necessary for accurate asset tracking. If you modified baseline probes, sensors, or patterns, verify that they still discover serial numbers. In addition, do not configure sensors or patterns to modify the serial number syntax, such as adding a custom prefix. Non-standard serial numbers can lead to inaccurate asset tracking.

CMDB identifier tables

Table	Description
Identifier [cmdb_identifier]	Stores all identifier rules.
Identifier Entry [cmdb_identifier_entry]	Stores all the identifier attributes.

Identifier rules

The default Discovery system contains these identifier rules. Each rule is associated with a specific CI type (the **sys_class_name** field on the CI record) or the table in the **Applies to** field, and contains the appropriate attributes for discovering CIs from that table. Where necessary, tables from related lists (**Search on** tables) are included in the rule. For more information, see Create or edit a CI identification rule.

CMDB identifier rules

Rule	Applies to table/attributes	Search on table/attributes
ESX Server Rule	ESX Server [cmdb_ci_esx_server]: correlation_id	none
Hardware Rule	Hardware [cmdb_ci_hardware] • serial_number • serial_number_type • name • ip_address • mac_address	• Serial Number [cmdb_serial_number] ◦ serial_number ◦ serial_number_type • Network Adapter [cmdb_ci_network_adapter] ◦ ip_address ◦ mac_address
Storage Server Rule	Storage Server [cmdb_ci_storage_server] • cim_object_path • name • serial_number • serial_number_type	• Serial Number [cmdb_serial_number] ◦ serial_number ◦ serial_number_type

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CMDB identifier rules (continued)

Rule	Applies to table/attributes	Search on table/attributes

	• mac_address
	• ip_address
		• Network Adapter [cmdb_ci_network_adapter]
		  ○ ip_address
		  ○ mac_address

WBEM Service Rule	WBEM Service [cmdb_ci_wbem_service]: cim_object_path	none

Matching strategy for the hardware rule

The **sys_class_name** cannot be an attribute for independent rules, such as cmdb_ci_hardware. If your Discovery identification strategy depends on matching a CI with a specific class, create a rule for each class you want to use for matching. Specify that class in the **Applies to** field of the Identifier form.

For example, you can create an identifier for a Linux server with different attributes than the Hardware Rule. You might want to use the machine name, IP address, and MAC address for identification. This is a solution for networks that use NIC bonding or teaming to increase available bandwidth. Bonded interfaces appear to be the same physical device and share the same IP and MAC addresses. The use of the **name** attribute allows Discovery to differentiate between the individual interfaces in the bonded channel.

Important: If you create an identifier with the **name** attribute, avoid changing adapter names. Discovery will be unable to resolve existing CIs for renamed adapters. Discovery labels the Install Status of that CI as "Absent" and creates another CI.

Your new rule would look like this:

Linux identifier rule

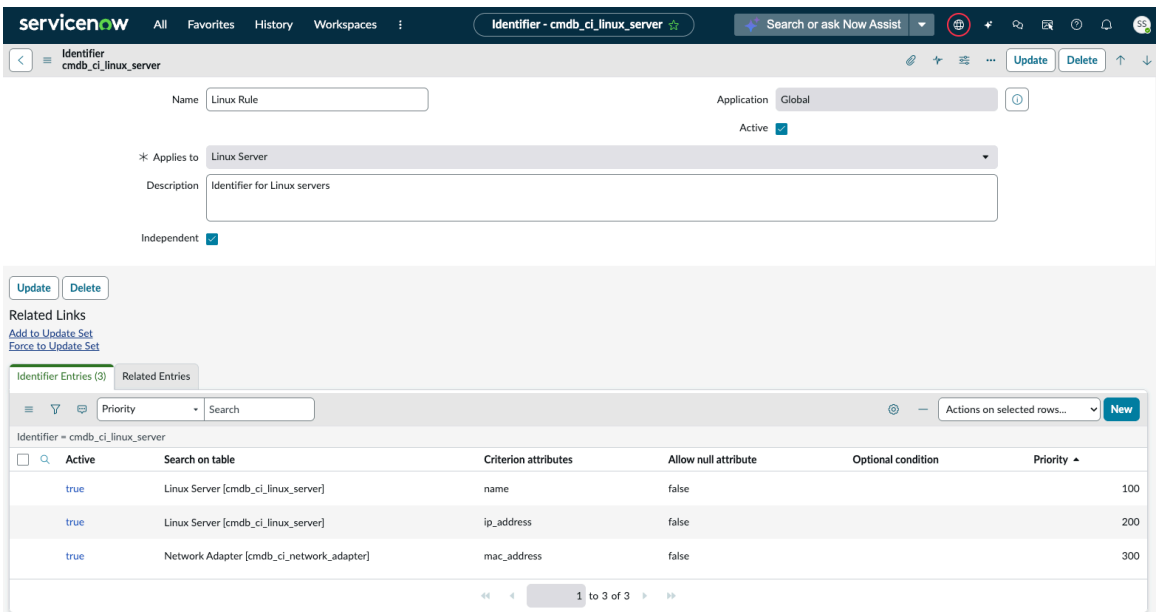

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Evaluation order for Discovery identifiers

Custom identifiers must have different *Order* values than those of the default identifiers. Discovery parses identifiers and attributes in sequence from low order numbers to high. You can create identifiers to run before or after the default identifiers, or mixed in with the identifiers from the base system. To avoid any identifier or rule from running, disable it by clearing the **Active** check box. The evaluation order for CMDB identifiers is established within each rule and only controls the parsing order of the attributes in that rule.

Evaluation order in CMDB identifier rules

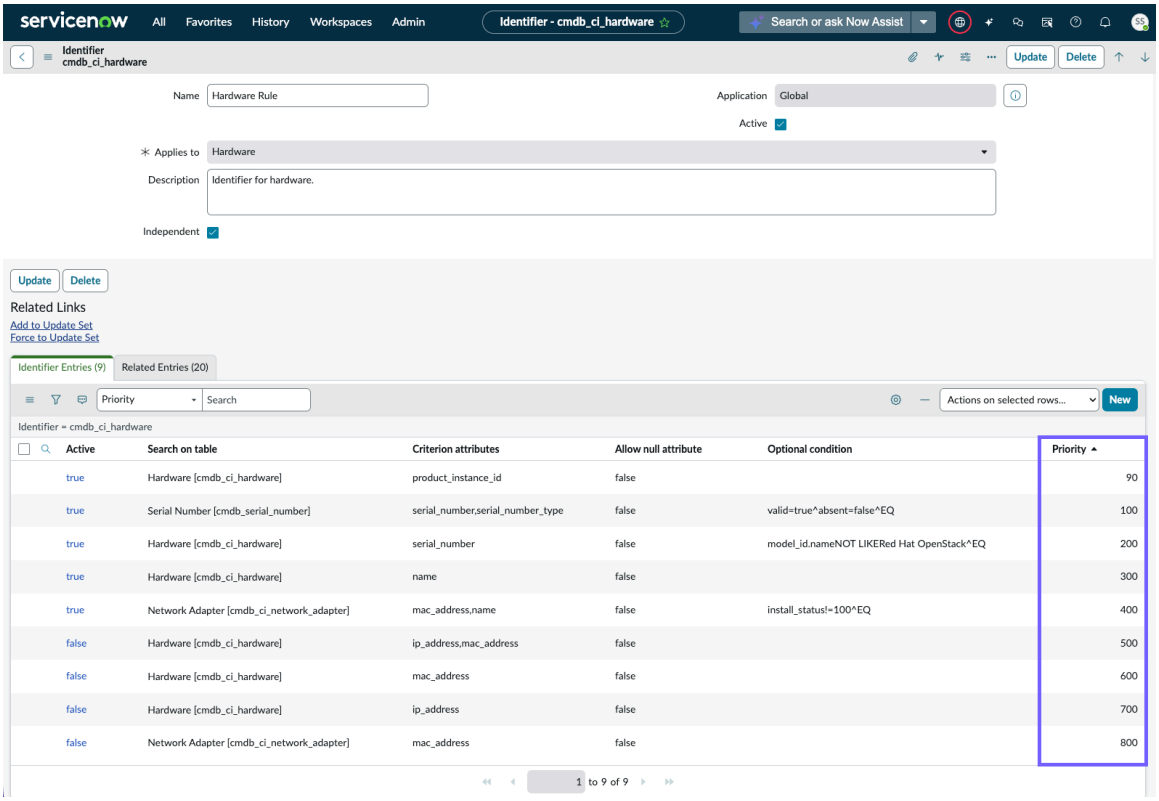

Properties for processing duplicate CIs

You can control how Discovery handles duplicate CIs with properties installed with Identification and Reconciliation. Use the `glide.identification_engine.skip_duplicates` and `glide.identification_engine.skip_duplicates.threshold` properties. For more information, see Properties for Identification and Reconciliation.

Properties that control identifier versions

All instances use identifiers from the CMDB Identification and Reconciliation framework. Upgrades from pre-Geneva versions still preserve the legacy identifiers, but you can switch to the new identifiers using a property: `glide.discovery.use_cmdb_identifiers`. If you upgraded from a pre-Geneva version, you must manually add this property and set it to `true` to use the new identifiers. If you upgraded from Geneva or later releases, this property is available in the System Properties [sys_properties] table. To preserve functionality in custom legacy identifiers, convert them to the new CMDB identifier rules format before enabling this property. The system does not reconfigure your custom identifiers to the new framework automatically.

Note: When Service Mapping is active, the new identifiers from the CMDB Identification and Reconciliation framework are always used regardless of the property value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How Discovery identifiers work

When Discovery has determined the device's class, it launches an identity probe that is configured to run one or more commands with a single authentication.

The identity probe in the base Discovery system can be configured to ask the device for information such as its serial numbers, name, and network identification. The results of this scan are processed by an identity sensor, which then passes the results to the identifier. The identifier then attempts to find a matching device in the Configuration Management Database (CMDB). If the identifier finds a matching CI, the identifier either updates that CI or does nothing. If the identifier cannot find a matching CI, it either creates a new CI or does nothing. If Discovery is configured to continue, the identifier launches the exploration probes configured in the classification record to gather additional information about the device. Exploration probes can be multiprobes or simple probes.

Note: When you use patterns for Discovery, the identity probe isn't used. Discovery uses the appropriate identifier rules based on the CI type that you're trying to discover as specified in the pattern operations. The operations in the pattern along with these identifier rules perform the identification and exploration phases of discovery.

Configure Discovery identity probes

Identity probes are multi-probes, which contain one or more simple probes configured to extract specific information from manageable devices. You can create your own identity multi-probe to identify CIs that Discovery does not already identify.

Before you begin

Role required: discovery_admin or admin

About this task

Identity probes return such information as device serial numbers (there can be several per device), the computer name, and network identification (MAC address). For example, the *Linux - Identity* multiprobe provided in the system, contains two simple SSH probes configured to return hardware and network information about the device. After the probe has discovered this information, it passes the data to a multisensor configured to process identity information.

Note: A multiprobe cannot contain another multiprobe.

Procedure

1. Navigate to **All** > **Discovery Definition** > **Probes**.
2. Select a probe from the list whose **Class** is *Multiprobe*.
3. Make sure the scope is set to **Global**.
4. In the **Includes Probes** Related List, click **Edit** to select simple probes to include in this multiprobe.
5. Click **New** to create a new simple probe.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example:

Discovery identity probes

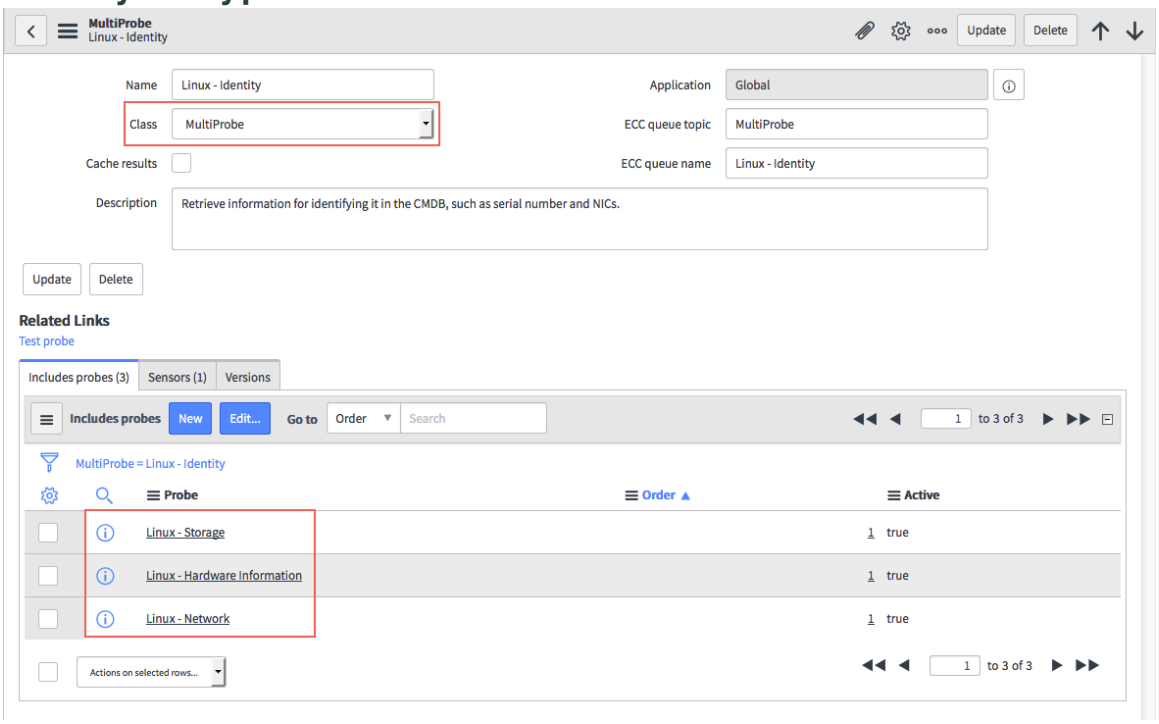

What to do next

Configure Discovery identity sensors that are part of the identify probes.

Configure Discovery identity sensors

If you customize an identify multi-probe, you can also configure a multi-sensor for it.

Before you begin

Role required: discovery_admin or admin

About this task

These multi-sensors pass the data returned by the identity probes to the Discovery identifiers. The identifiers then search the CMDB for matching CIs. When the identity of a device is resolved, the identifiers pass the result state for the device to the multi-sensors. The multi-sensors react accordingly, either by launching exploration probes or stopping Discovery for that device.

Procedure

1. Navigate to **All** > **Discovery Definition** > **Sensors**.
2. Select an existing identity multi-sensors from the list.

   Note: To avoid confusion, the multi-sensors provided in the instance must have the same names as their matching multi-probes.

   The **Responds to Probes** related list shows the simple probes which pass their data to this multi-sensors. These are the same simple probes that appear in the Includes Probes list in the matching multi-probes record.

3. Click the **Script** link for each probe in the list to see the script that the multi-sensors runs to process the data from the probe.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example:

Discovery identity sensors

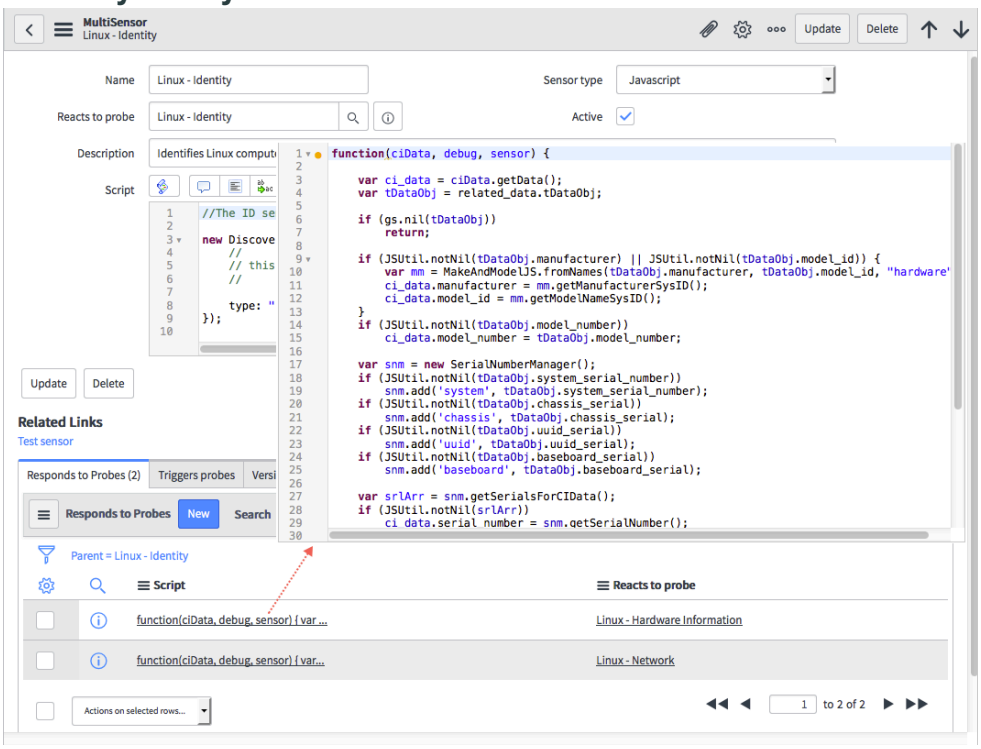

What to do next

Run a discovery through the Discovery Schedule to search for CIs and verify that they are identified correctly in the CMDB.

Serial number types for identification

As Discovery finds CIs, their serial numbers are listed in the Serial Number [cmdb_serial_number] table so they are easy to identify. Serial number types vary depending on the CI, as described in the following examples.

Important: Serial numbers are necessary for accurate asset tracking. If you modified baseline probes, sensors, or patterns, verify that they still discover serial numbers. In addition, do not configure sensors or patterns to modify the serial number syntax, such as adding a custom prefix. Non-standard serial numbers can lead to inaccurate asset tracking.

Linux

In the result of dmidecode | cat, the value on the left side is what Discovery looks for. The value on the right side is how it is stored in the Serial Number [cmdb_serial_number] table.

```
Serial Number : system_serial_number

UUID : uuid_serial

Serial Number : baseboard_serial

Serial Number : chassis_serial
```

Windows

For Win32 WMI classes, the value on the left is the name by which it is stored in the Serial Number [cmdb_serial_number] table. The value on the right is the WMI value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
('system' : Win32_ComputerSystemProduct.IdentifyingNumber);

('uuid' : Win32_ComputerSystemProduct.UUID);

('chassis' : Win32_SystemEnclosure.SerialNumber);

('bios' : Win32_BIOS.SerialNumber);

('baseboard' : Win32_BaseBoard.SerialNumber);
```

SNMP

For SNMP, the mapping below is based on the code. Physical types of serial numbers are from all instances of iso.org.dod.internet.mgmt.mib-2.entityMIB.

```
'cisco_stack' :
 'iso.org.dod.internet.private.enterprises.cisco.workgroup.cisco
StackMIB.chassisGrp.chassisSerialNumberString'

'cisco_chassis' :
 'iso.org.dod.internet.private.enterprises.cisco.temporary.chass
is.chassisId'

'foundry' :
 'iso.org.dod.internet.private.enterprises.foundry.products.swit
ch.snChassis.snChasGen.snChasSerNum'

'apc_pdu' :
 'iso.org.dod.internet.private.enterprises.apc.products.hardware
.masterswitch.sPDUIdent.sPDUIdentSerialNumber'

'printer' :
 'iso.org.dod.internet.mgmt.mib-2.printmib.prtGeneral.prtGeneral
Entry'

'standard' :
 'iso.org.dod.internet.private.enterprises.apc.products.hardware
.ups.upsIdent.upsAdvIdent.upsAdvIdentSerialNumber'
```

Discovery status

The Discovery status provides a summary of a Discovery launched from a schedule. You can also cancel a Discovery that is in progress from the status form.

Note: Discovery schedules that have no status updates for a over a defined number of minutes are analyzed automatically by the Discovery Status Monitor job. By default, this job applies to Discovery schedules that discover configuration items. For more information, see the Discovery Status Monitoring Job article in the Now Support Knowledge Base.

To access the Discovery Status form, navigate to **Discovery** > **Status** and open the status record for a Discovery.

Each record in the Status list represents the execution of a Discovery by a schedule and displays such high-level information as the date of the Discovery, the mode, the number of probe messages sent to devices, and the number of sensor records that were processed. A status record contains data that can help you troubleshoot a failed discovery. Use this data to troubleshoot the behavior of individual probes and sensors or even run those elements separately. Use the status controls to enter probe/sensor threads at any point for a specific

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery, and then follow the process in either direction. You can also clearly see which classification probe ran and discovered a particular device by viewing the classification probe column in the device list.

- **Refresh**: Refreshes the related tabs, including the **Discovery Log** and **ECC Queue** for discoveries still in progress.
- **Discovery timeline**: The Discovery timeline is a graphical display of a discovery, including information about each probe and sensor that was used.
- **Discovery log entries**: The Discovery Log shows information such as classification failures, Configuration Management Database (CMDB) updates, and authentication failures. A Discovery Log record is created for each action associated with a discovery status.
- **ECC queue entries**: Entries in the ECC queue provide you with a connected flow of probe and sensor activity, as well as the actual XML payload that is sent to or from an instance.
- **Device history**: The device history provides a summary of all the devices scanned during discovery, and what action sensors took on the CMDB.

Note: By default, only 30 days of Discovery records are displayed in the status list at a time.

The following Discovery status record shows the discovery of a load balancer using a pattern.

Example Discovery status record

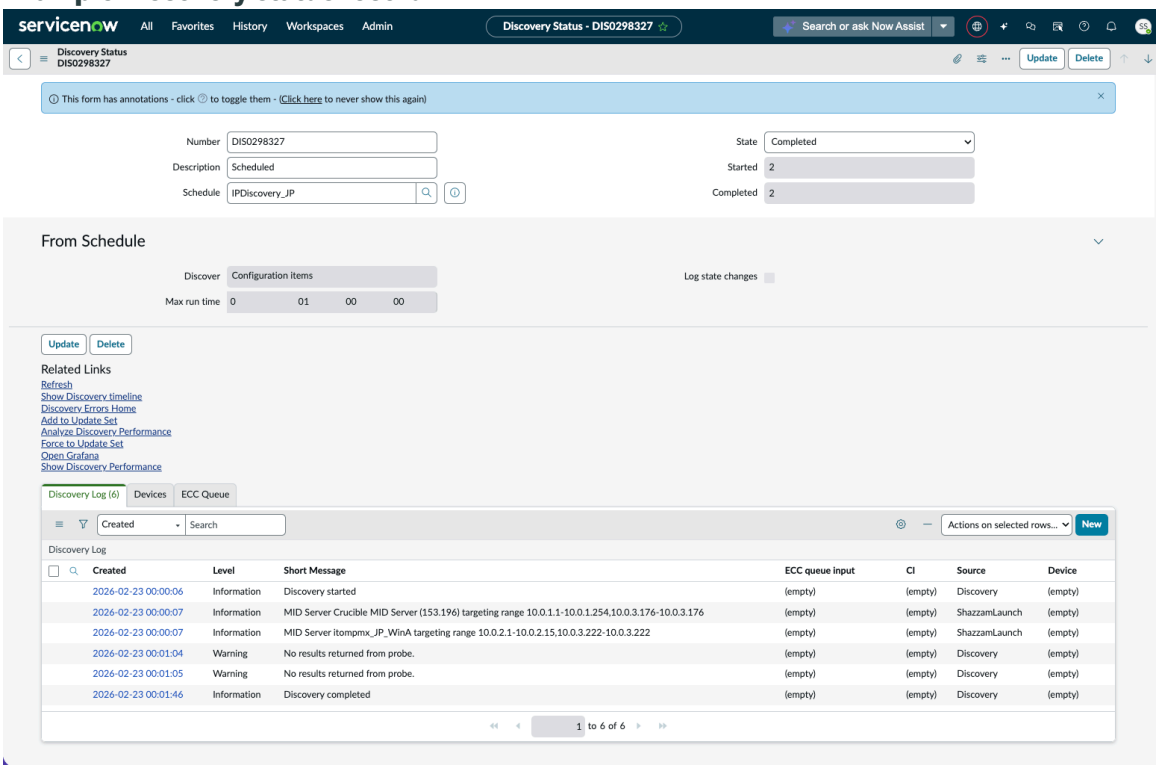

Field	Description
Number	An auto-generated number for the record.
Description	How this discovery was run. Typically, the description is *Scheduled*, but if you ran discovery manually from a Schedule, the record would show *Discover Now* in the *Description* field.
Schedule	The name of the Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
State	The state of the discovery: • **Starting** • **Active** • **Completed** • **Canceled**
Started and Completed	The number of probes started and completed during discovery.
Source	Optional read-only field that shows the source of the discovery, such as **Schedule Discovery**, or **Quick Discovery**. If the source of the Discovery was from an API call, the source is **DiscoveryAPI**.
From Schedule (Read-only schedule data)	
Discover	Shows the Discovery type. The possible types are: **Configuration items**, **IP addresses**, **Networks**, or **Processes/Connections** (ADM).
Max Run Time	Displays the maximum amount of time Discovery was permitted to run on this schedule.
Log state changes	Indicates that state changes were logged during this Discovery. These states can be seen in the *Last* and *Current* fields in the**Discovery Devices** list in this form.

Related topics

Schedule a horizontal discovery

Cancel a Discovery from a Discovery Status

You can manually cancel a discovery while it is in progress from the Discovery Status record.

Before you begin

Role required: discovery_admin or discovery_operator

About this task

Procedure

1. Navigate to **All** > **Discovery** > **Status**.
2. Click a **Number** to open the status record.
3. Click **Cancel Discovery** in the **Related Links**.
   Cancel Discovery is only available for an active Discovery.

Result

When a Discovery status cancels, any associated sensor transactions are immediately terminated and any scheduled sensor jobs are deleted from the system. After cancellation, the cleanup Status shows the **Completed count** and the cancellation is logged in the Discovery log. In the Queue [ecc_queue] table, any records belonging to sensors exceeding the Transaction Quota Rule are set to the **Error** state.

Discovery timelines

A Discovery timeline generates a graphical display of a Discovery Status record, including information about each probe and sensor that was used in the discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use Discovery timelines to display the following:

- The flow of probes and sensors, including those that are used with patterns through a discovery.
- The duration of each probe and sensor that ran during a discovery, and the proportion of time required for queuing and processing.
- Tooltips containing additional data about a probe and sensor.
- Records from the ECC Queue.

Size of discoveries in the timeline

View timelines for an entire discovery or for individual devices in a discovery. By default, the maximum size discovery that you can display in a timeline is 300 entries in the ECC Queue (150 probes and 150 sensors). To display larger discoveries, change the default setting in the `glide.discovery.timeline.max_entries` property.

Example Discovery timeline

This example shows the Discovery timeline for the first two phases of Discovery (Port scan and Classification):

Discovery (early phases) on the timeline

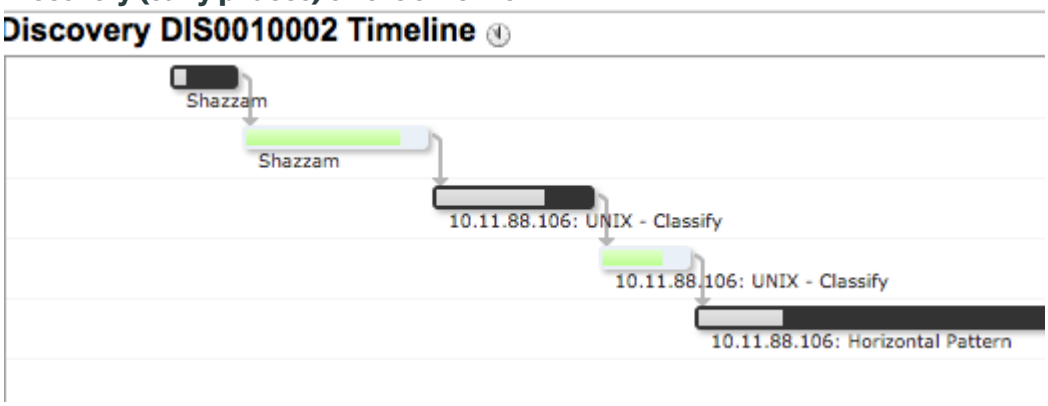

In this example, the Shazzam probe and sensor run, followed by the Unix classifier probe and sensor. The classifier calls the Horizontal Pattern probe, which runs a specific pattern.

View the Discovery timeline

A Discovery Timeline generates a graphical display of a Discovery Status record, including information about each probe, sensor, and pattern running.

Before you begin

Role required: discovery_admin or admin

About this task

Procedure

1. Navigate to **All** > **Discovery** > **Status**.
2. Select a Discovery from the record list.
3. Click the **Show Discovery Timeline** Related Link.
4. Clear the warning, and then select the **Devices** Related List.
5. Click the IP address of a device.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. In the Device record, click the **Show Discovery Timeline** Related Link.

   The Discovery timeline for that device appears.

   Discovery Timeline

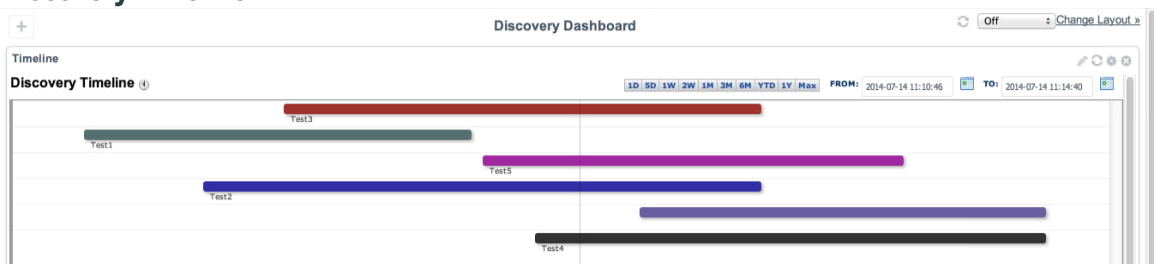

7. Use the pink slider at the bottom of the timeline to change the perspective.

   a. Move the slider from right to left to view all the tasks on a long timeline.

   b. Adjust the end points of the slider to change the magnification.
   A narrow slider zooms in on the spans and provides a more detailed view of complex timelines. A wide slider pulls the view out and makes more of the timeline visible on the screen.

8. Use the selector range at the top of the screen to adjust the visible time frame.
9. To limit the timeline to the length of the Discovery, click **Max**.

   The time scale adjusts automatically to the length of the Discovery. The available time scale range is from one day to 1 year.

   Tooltips

   Discovery timelines display probe and sensor performance data and CI information in tooltips. Hover your cursor over a span to view this data.

   Probes are displayed by black spans. The queue time for a probe is shown as a silver bar within the span, and the processing time is represented by the remaining space.

   Sensor spans are red, and the queue time is shown as a green bar. Selected spans of any type display in yellow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Timeline Shazzam Sensor

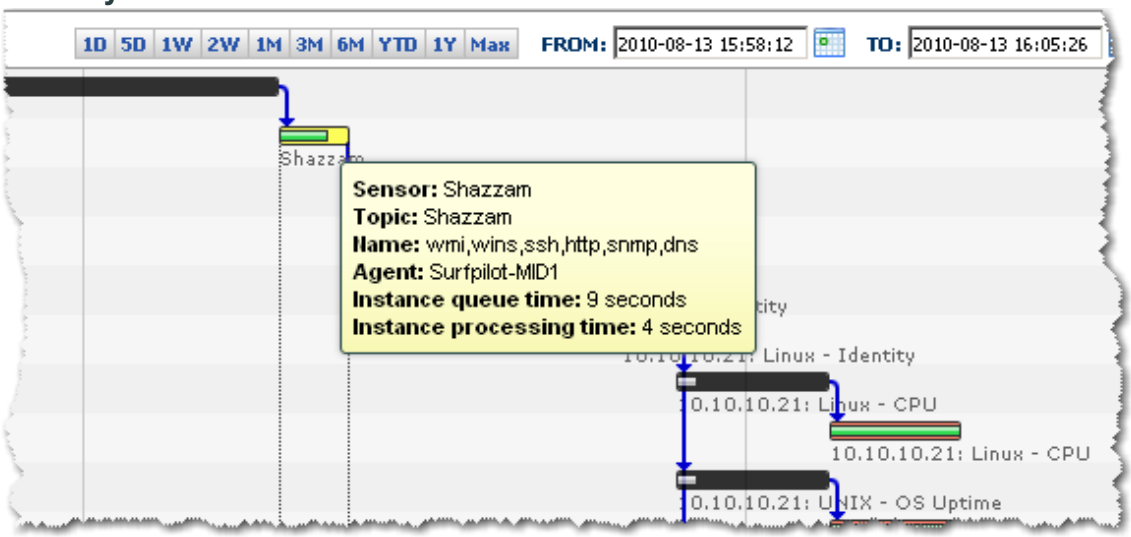

ECC Queue

Double-click a span to open the ECC Queue record for that probe or sensor.

Discovery Timeline ECC Queue

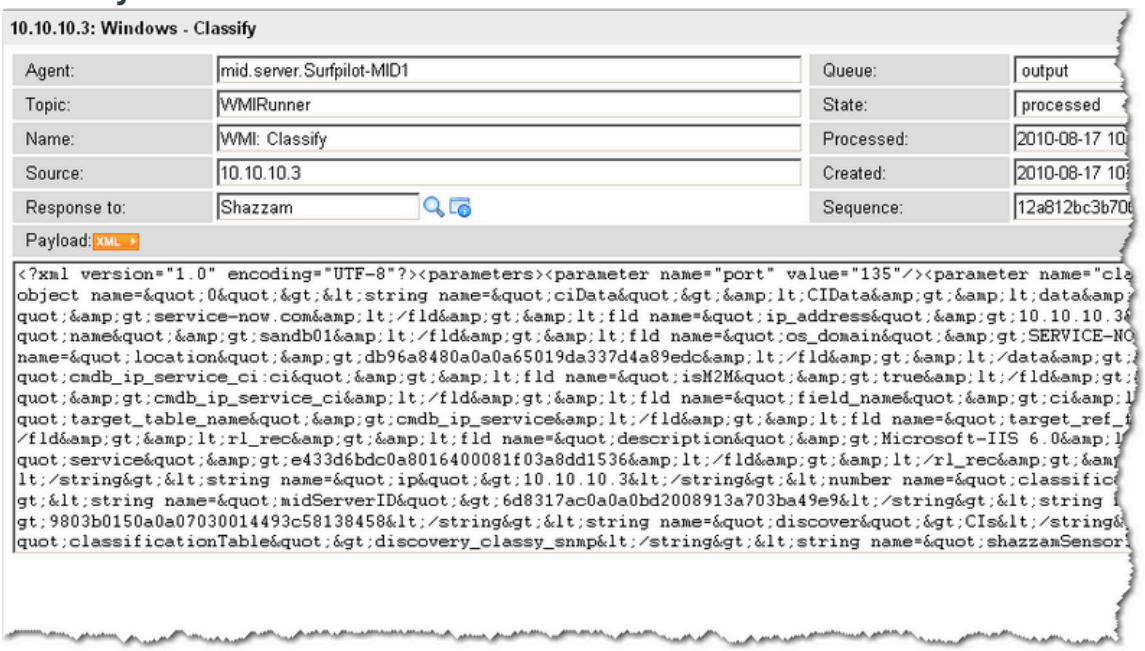

The ECC queue for Discovery

The External Communication Channel (ECC) Queue is a connection point between your ServiceNow instance and other systems that integrate with it, most commonly a MID Server.

Contents of the ECC Queue

The ECC queue contains records that are created during all stages of discovery. Each record is a message from the instance to system your are targeting for discovery, classified as **output**, or a message from that system to the instance, classified as **input**.

ECC queue records provide you with a connected flow of probe and sensor activity along with the XML payload sent to or from an instance and the assets you're targeting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Automatic deletion of ECC queue records

As your IT Asset Discovery runs, the size of the ECC queue can grow to several gigabytes. Most of the “completed tasks” records are unnecessary, but some might be important for issue resolution.

For example, if Discovery is not properly capturing the disk drives on a particular Windows server, you can look in the ECC Queue at the data returned by the Windows - System Information probe. For this reason, you should retain ECC Queue data from Discovery for at least a month.

By default, records in the ECC Queue older than 7 days are deleted automatically. You can set the deletion schedule by updating the ECC queue (as we will explain in the next section) and the table rotation schedule. The table rotation names are:

- `ecc_queue_event`
- `ecc_queue`

Accessing the ECC queue

You can access the ECC queue from either of these locations:

- **Discovery** > **Output and Artifacts** > **ECC Queue**
- **ECC** > **Queue**

The following image is an example of a record in the ECC queue. This record shows that a ConnectorProbe was instructed to run and has been processed. Note the XML payload it returned as the discovery task collected data. See The MID Server ECC queue for a description of each field.

An example ECC queue record

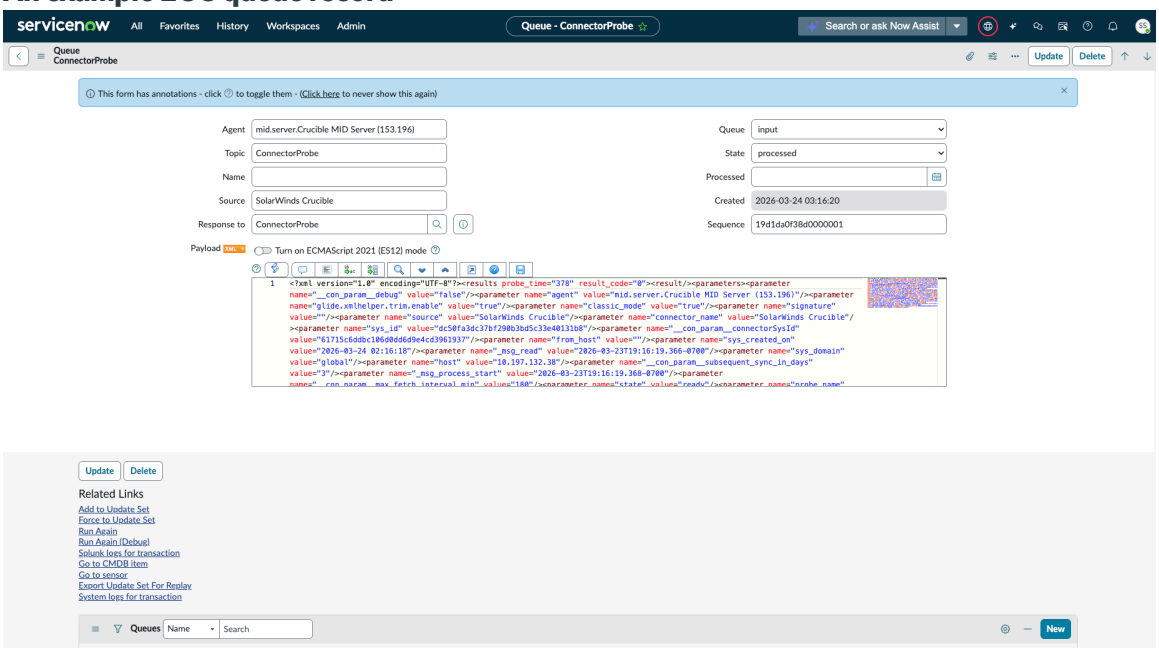

ECC queue states

Input and output messages keep processing until the discovery is finished or terminated. You can monitor the ECC queue to see records populate during discovery and see the states of each record change.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ECC queue input records are processed by business rules (BR), but not all business rules direct these ECC queue inputs to process. As a result, some ECC queue inputs remain in a **ready** state.

The process is as follows:

1. When new probe instructions are available for the MID Server, the system creates an **output** record with the state of **ready**.
2. The MID Server then starts processing the probe instructions and the state changes to **processing** before it finally changes to **processed**.
3. When a result returns to the MID Server and is ready to be sent to the instance, the system creates an **input** record with the state of **ready**.
4. While the system processes the payload, the state changes to **processing** before it finally changes to **processed**.

   If any errors are encountered, the state changes to **error**. You can review errors as a starting point for troubleshooting a failed discovery.

Debug processing of ECC queue records

To use a script debugger, process the payload on the same session as the debugger. Review the business rule which processes the particular ECC queue record and the actual code which processes the payload.

Properties and parameters that affect the ECC queue

These properties are available to control aspects of the ECC queue:

Property	Description
Save ECC queue attachments	The normal behavior for Discovery sensors is to delete attachments to ECC queue entries on successful sensor processing. Enabling this property overrides this behavior, and forces attachments to be preserved. This is typically done for debugging purposes. To enable this property, navigate to **Discovery Definition** > **Properties**.

Related topics

Manage ECC Queue content for a MID Server

Discovery device history

The **Devices** Related list in the Discovery Status form provides a summary list of all the devices scanned.

During a Discovery, the list tracks current and completed activity and displays an incremental scan counter. When Discovery is finished for a device, the final disposition is displayed in the *Completed activity* column. Successful discoveries that result in updated or created configuration items (CIs) are highlighted in green. When device history or duplicate IP record is created, the classification_probe is also populated. To view the log entries for errors (such as connection failure) on a specific device, select the **Details** link in the *Issues* column.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery device list

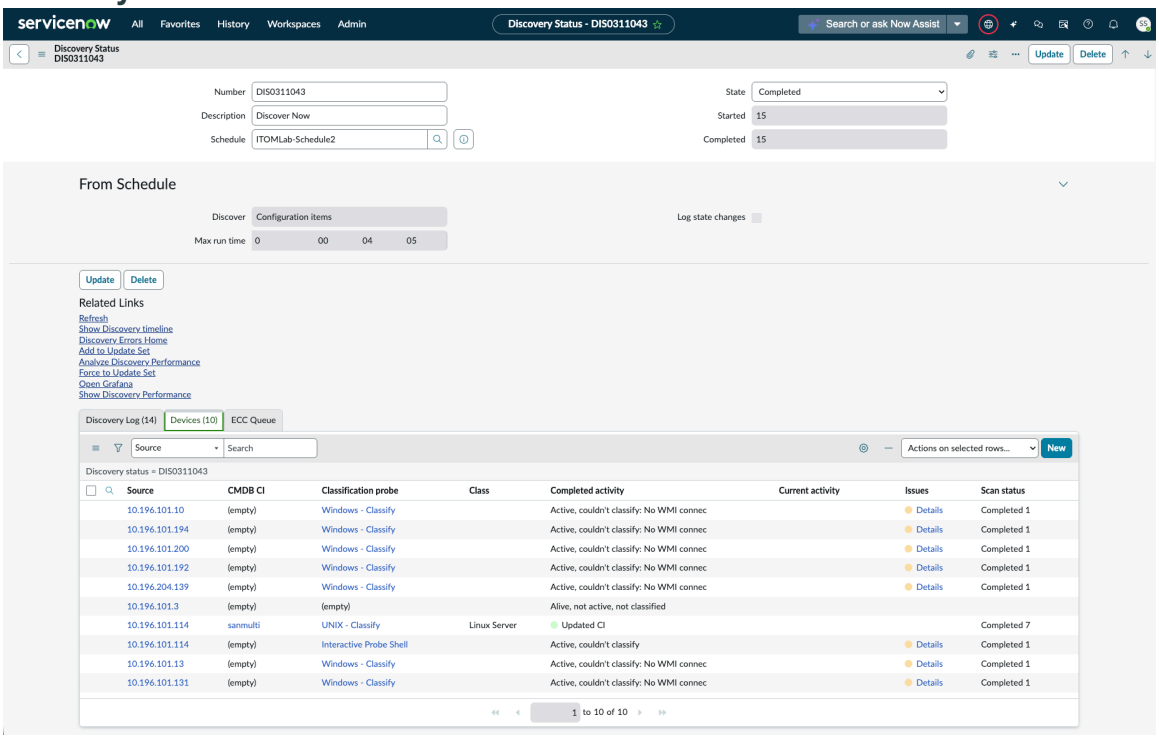

Select the IP address of a device in this list for details about that device. The log results for that device are displayed in the list at the bottom of the form.

Discovery device

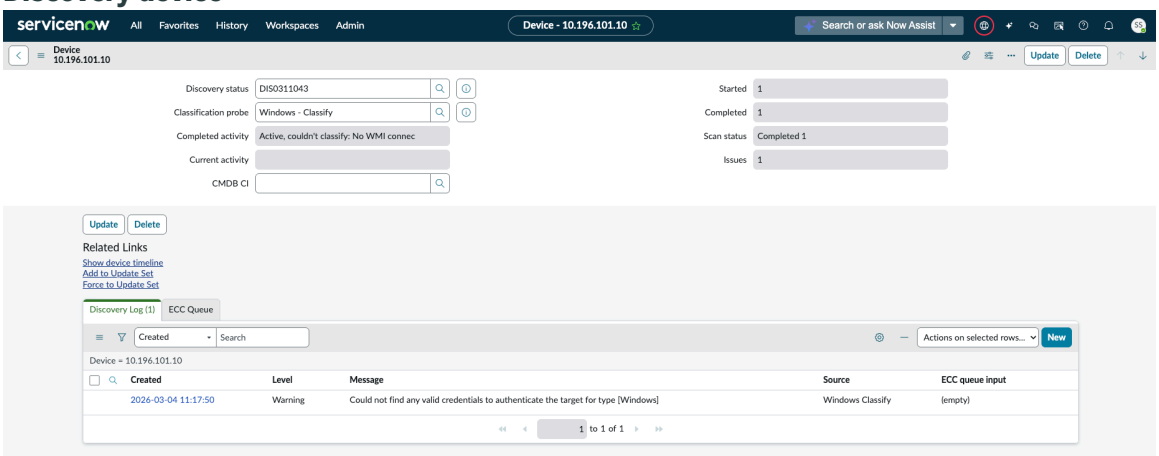

The following fields are available in this form:

Field	Input value
Source	The IP address of the device discovered.
Completed activity	Indicates the outcome of Discovery for this device or the last completed activity for a Discovery in progress, such as *Identified CI*. Successful outcomes are indicated in green.
Current activity	The current scanning activity for this device for a Discovery in progress, such as *Updating CI*.
CMDB CI	The name of this device as it appears in the Configuration Management Database (CMDB).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input value
Started	The number of device-specific probes run. This number does not include the universal probes, Shazzam and Ping, that run initially.
Completed	The number of sensor records created from the device-specific probes that were run.
Scan status	Shows the final scan count of a completed Discovery or an incremental scan counter for a Discovery in progress. For example, Scan 17 of 19.
Issues	Displays the number of issues encountered during Discovery of this device. Select the **Discovery Log** Related list to view these issues.
Fields that can be added by configuring the form:	
DNS Names	Displays DNS names for each discovered device.

During different phases of Discovery, you may see different values in the **Completed Activity** field on a discovery_device_history record.

Potential Values for Completed Activity in discovery_device_history

Value	When you might see this value	Meaning
Common examples		
Alive, not active, not classified	After Shazzam	Device is responding, but it does not have any open (active) ports that we can run classification against.
Active	After Shazzam	Device has open (active) ports that we potentially can run classification against.
Classified	After Classification Probe	We were able to classify the device type and will be proceeding to the Identification stage.
Active, couldn’t classify	After Classification Probe	We were unable to classify the device type. Sometimes, this value will also include a reason for why we were unable to classify. For example, Active, couldn't classify: No WMI connection.
Created CI	After Identification Probe/ Pattern or after Discovery - Device Complete Business Rule runs (for new CI)	New CI has been created based on identification results.
Identified CI	After Identification Probe/ Pattern	Existing CI has been found based on identification results.
Identified, ignored extra IP	After Identification Probe/ Pattern	Device has already been discovered via a different IP in same Discovery status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Potential Values for Completed Activity in discovery_device_history (continued)

Value	When you might see this value	Meaning
Updated CI	After certain Exploration Probes/Patterns or after Discovery - Device Complete Business Rule runs (for existing CI)	CI has been updated with additional information.
Less common examples		
Unclassified	After Shazzam (only for credential-less Discovery and the Discover value on the Discovery Schedule is set to IP Addresses)	For the credential-less Discovery of the target IP, we were unable to match to an existing CI on the instance.
Classified	After Shazzam (only for credential-less Discovery and the Discover value on the Discovery Schedule is set to IP Addresses)	For the credential-less Discovery of the target IP, we were able to match to an existing CI on the instance.
Active, Windows Cluster VIP	After Windows - Classify	Device is a Windows Cluster VIP and will be updated as such.
Classified, but not disabled	After Classification Probe	We made a match to a Classifier type, but the classifier itself is disabled.
Classified, couldn't identify	After Identification Probe/ Pattern	Couldn’t create or match to existing CI record.
Identified, not updating CI	After Identification Probe/ Pattern	We were able to properly identify a CI record, but we are choosing not to update the CI.

Any device that includes the following classes or one of its child classes will be displayed in the Discovery device history table.

- Computer [cmdb_ci_computer]
- Network Gear [cmdb_ci_netgear]
- PDU [cmdb_ci_pdu]
- Printer [cmdb_ci_printer]
- UPS [cmdb_ci_ups]

If there were issues, or if Discovery failed to complete, select the **Details** link to view the log records for the issue. The failure of any probe is considered an issue, even if the device was eventually classified properly and updated in the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery details about a device

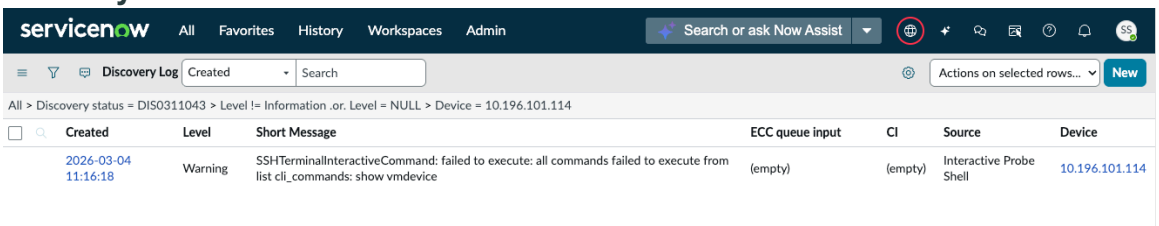

Address scan data

When Discovery scans for IP addresses only (without credentials or identifiers), no updates are made to the CMDB. All IP addresses discovered appear on this list, including multiple IPs on the same device. The results of IP address scans include slightly different information than the results of a CI scan. Since there is no CMDB activity associated with the IP address scan, the *Completed activity* column displays only the classification status.

Possible statuses are:

- *Classified*
- *Unclassified*
- *Alive, not classified*

For Classified devices, Discovery might identify the type of device in the *Current activity* column. For example, Network Gear might be classified as Cisco Network Gear, and a Computer might be classified as a Windows Computer.

Logs for horizontal discovery

The system collects logs to reflect the activity that takes place during a horizontal discovery based on both patterns and probes. Use the logs to fine-tune or troubleshoot the discovery process.

For horizontal discovery the system generates logs of the following types:

Probe-based

The logs for probe-based horizontal discovery cover classification failures, CMDB updates, and authentication failures. A Discovery Log record is created for each action associated with a discovery status.

Note: The Discovery log also displays errors that occur during identification by taking log entries from the Identification Engine log. See Examine run logs for more information about the identification engine and possible errors.

Pattern-based

The logs for pattern-based horizontal discovery include the information about discoveries that were performed with patterns. A horizontal discovery log record is created for an entire horizontal discovery run, which includes the results of all the operations specified in the pattern.

Review the logs for horizontal discovery

Access the logs to understand the discovery process and fine-tune the discovery process.

Before you begin

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Discovery** > **Status**.
2. Open a Discovery Status record.
   The **Discovery Log** related list displays the logs for probe-based discovery. See Discovery log details.

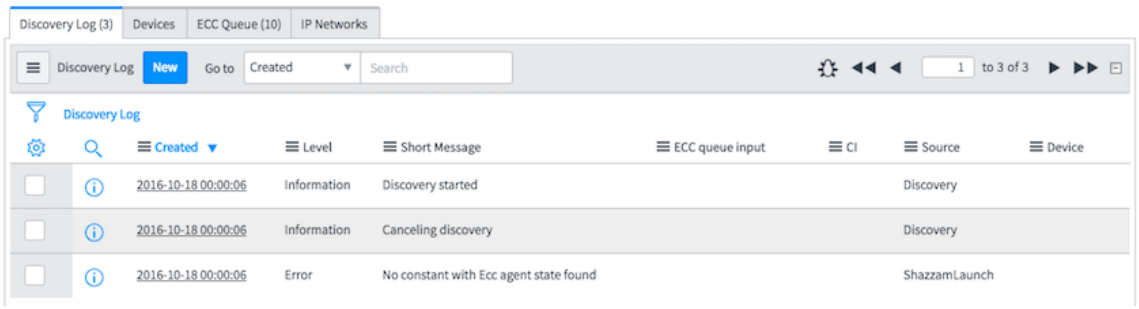

3. **Optional:** Find out the information related to the relevant log entry:
   - Select the link in the **CI** column to open the record for that CI.
   - Select the link in the **Created** column to view the log record.
   - Select the IP address link in the **Device** column to view the log records for a particular device.
4. To view logs for pattern-based discovery, select the **Here** link for a log entry that has **HorizontalDiscoveryProbe** in the **ECC queue input**.

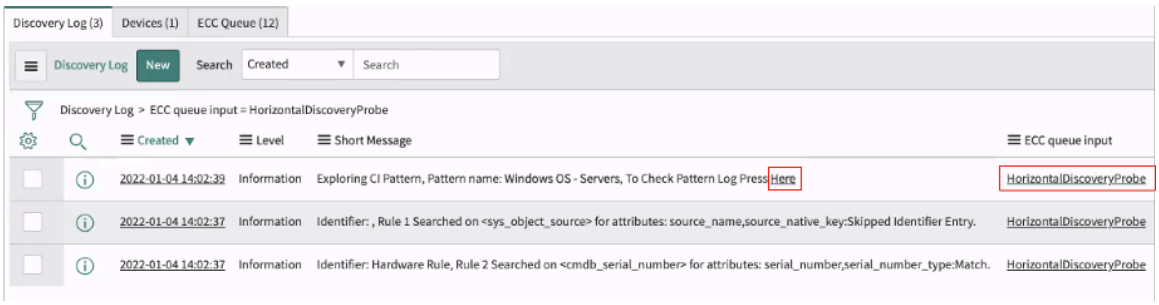

The Horizontal Discovery Log window opens showing the list of pattern-based log entries. For information on logged discovery phases, see Discovery log details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Horizontal Discovery log

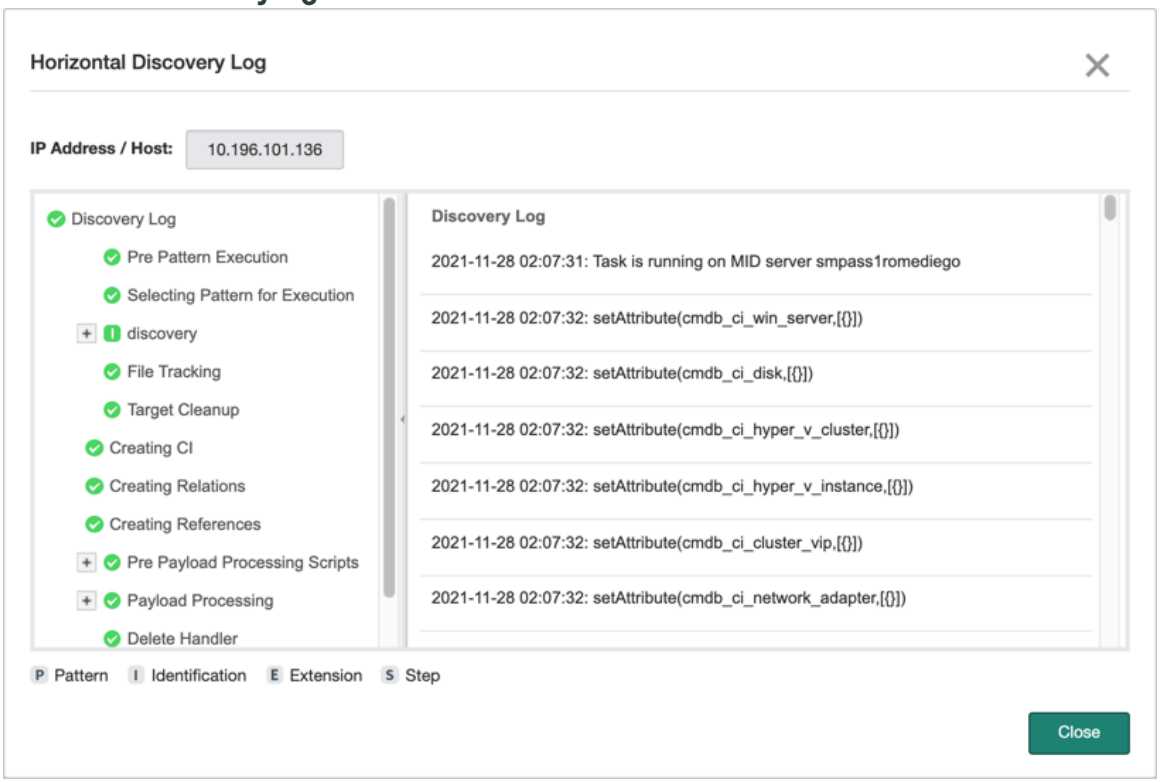

5. Select an item in the left-hand column to see more information about it in the right-hand column.
6. **Optional:** To debug the pattern, select the name of the Identification Section, and then select the **Debug** button.
7. To return to the log record, select **Close**.

What to do next

- To fine-tune or troubleshoot the pattern-based discovery, retrieve the payload as described in KB0998923: Retrieve Payload Feature.
- Validate commands used in pattern-based discovery, if there is one of the following errors:
  - Authentication issue
  - Authorization/Permission issue
  - No Valid credentials present

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- IP not reachable from mid
- Command with wrong syntax

- If there is an error in payload processing, select the **IRE Error Documentation** button to learn about the Identification and Reconciliation Engine (IRE) error that occurred and how to resolve it.

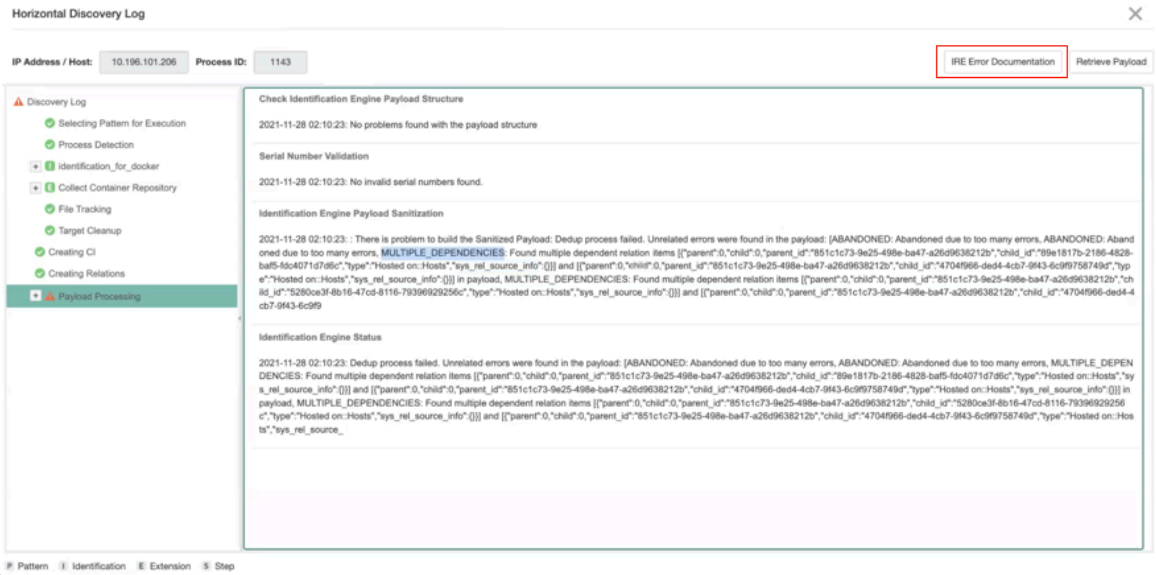

Related topics

Identification and Reconciliation engine (IRE)

KB0854603: Pattern Debugger Session Timed Out

Change retention time for probe-based discovery logs

You can configure the amount of time that Discovery retains log entries for the probe-based horizontal discovery.

Before you begin

Role required: admin

About this task

By default, the probe-based discovery log retains information for 30 days, or 2,592,000 seconds.

Procedure

1. Open the Auto Flushes table by typing `sys_auto_flush.list` in the navigation filter.
2. Open the **discovery_status** record.
3. Change the value in the **Age in seconds** field.
4. For the change to take effect immediately, run the Table Cleaner Job in the Schedule [sys_trigger] table.
5. Click **Update**.

Discovery commands for probes and patterns

Discovery commands are used for both probe and pattern-based discovery to access configuration items (CIs) in your environment.

For information about Service Mapping commands, see Prerequisites for performing top-down discovery using Service Mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information about MID Server privileged commands, see MID Server privileged commands.

Probe-based discovery

Discovery runs several different types of commands during probe-based horizontal discovery. Some of these commands require elevated privileges to run.

Resources for information about commands and permissions for probe-based discovery

Command/Permission	Resource
SSH commands requiring a privileged user	Privileged SSH commands for probe-based discovery
SSH commands not requiring a privileged user	Non-privileged SSH commands during probe-based discovery
Windows commands and permissions	Windows probes and permissions
Additional permissions	Discovery probe permissions

Pattern-based discovery

Discovery runs several different types of commands during pattern-based horizontal discovery. Some of these commands require elevated privileges to run. For pattern-based discovery, you can view all the commands used by a specific infrastructure or application pattern on the ServiceNow AI Platform and validate them before running discovery. For cloud resource patterns, you can view all the required permissions in the Cloud Discovery patterns spreadsheet.

Resources for information about viewing and validating pattern-based discovery commands and permissions

Action	Resource
View commands used by infrastructure and application patterns	• View Patterns commands through the Discovery Patterns module* • View Patterns commands through the Command List module • Command List form for pattern-based discovery
Validate commands used by infrastructure and application patterns	• Validate commands used in pattern-based discovery* • View command validation task results in pattern-based discovery • Command Validation Tasks form for pattern-based discovery
View permissions required by cloud resource patterns	Cloud Discovery patterns spreadsheet

* These tasks are exclusive to the discovery_admin role and can't be performed by users with the pde_viewer role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

KB1123625: Command Validation Tool

View Patterns commands through the Discovery Patterns module

View all the commands required for an infrastructure or application pattern to verify you have sufficient permissions to run discovery.

Before you begin

- Confirm that the following applications have been installed:
  - Pattern Designer Enhancements (sn_itom_pde)
  - Discovery and Service Mapping Patterns (sn_itom_pattern)
- If you're viewing commands for the first time, have upgraded Discovery and Service Mapping Patterns, or have customized any pattern, refresh the command list by navigating to **All** > **Pattern Designer** > **Command Validation Tool** > **Commands List** and selecting **Refresh Command List**.

Role required: discovery_admin

About this task

The discovery_admin role is required to access the **Discovery Patterns** module. Starting with Pattern Designer Enhancements version 3.9.0, the new pde_viewer role can view pattern commands through the **Command List** module but can't edit or modify them. For more information, see View Patterns commands through the Command List module.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. In the **Name** field, search for the relevant pattern.
3. Select and hold (or right-click) the name of the pattern and select **List Pattern Commands**.

Result

You're redirected to the Command List form filtered for the selected pattern.

For more information on the Command List form fields, see Command List form for pattern-based discovery.

Related topics

Discovery commands for probes and patterns

Validate commands used in pattern-based discovery

View command validation task results in pattern-based discovery

View Patterns commands through the Command List module

View and filter infrastructure or application pattern commands by various criteria, including pattern, CI type, or command, to confirm you have all the required permissions.

Before you begin

- Verify that Pattern Designer Enhancements (sn_itom_pde) is installed.
- If you have the discovery_admin role and match one of these conditions, refresh the command list:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You are viewing commands for the first time
- You have customized any pattern
- You have upgraded Discovery and Service Mapping Patterns

You can refresh the command list by navigating to **All** > **Pattern Designer** > **Command Validation Tool** > **Commands List** and selecting **Refresh Command List**.

Note: The pde_viewer role can't refresh commands.

Role required: One of the following.

- discovery_admin
- pde_viewer (starting with Pattern Designer Enhancements version 3.9.0)

About this task

Starting with Pattern Designer Enhancements version 3.9.0, the new pde_viewer role can view commands through the **Command List** module but can't edit or modify them.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Command Validation Tool** > **Commands List**.
2. Filter the command list according to your requirements.
   For more information about the filtering options, see Command List form for pattern-based discovery.

Related topics

View Patterns commands through the Discovery Patterns module

Validate commands used in pattern-based discovery

View command validation task results in pattern-based discovery

Discovery commands for probes and patterns

Validate commands used in pattern-based discovery

Validate pattern commands to verify that the MID Server can successfully run them. Typically, commands might fail if you haven't configured the credentials necessary to run these commands on your ServiceNow instance. Another common reason of command failure is that the IP addresses used for discovery aren’t reachable.

Before you begin

- Verify that you have the following applications:
  - Pattern Designer Enhancements (sn_itom_pde)
  - Discovery and Service Mapping Patterns (sn_itom_pattern)
- Check the log for horizontal discovery based on patterns, as described in Review the logs for horizontal discovery.
- If you are viewing commands for the first time, have upgraded Discovery and Service Mapping Patterns, or have customized any pattern, refresh the command list by navigating to **All** > **Pattern Designer** > **Command Validation Tool** > **Commands List** and selecting **Refresh Command List**.
- To be able to test credentials on an IP range or IP network, install the ParseIPRanges script include attached to the KB article: KB1123625: Command Validation Tool.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Whether you perform discovery using patterns or probes, the MID Server executes commands to access configuration items (CIs). Most commands require credentials to run them. In addition, you must assume a role of a user with elevated rights to run some of these commands. If the necessary credentials are missing, discovery fails.

For pattern-based discovery, you can validate all commands used by a certain pattern simultaneously.

You can also validate discovery commands to check that the IP addresses you're using for discovery are reachable. Sometimes IP addresses are incorrect or the servers using these IP addresses are down.

If you validate commands that contain applicative credentials, the Command Validation Tool checks the Connections and Credentials module and uses relevant applicative credentials configured there.

Starting with Pattern Designer Enhancements version 3.9.0, the new pde_viewer role can view the command validation task results through the **Command Validation Tasks** and **Command Validation Task Results** modules but can't edit or modify them. For more information, see View command validation task results in pattern-based discovery.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Command Validation Tool** > **Create New**.
2. **Optional:** To validate all commands used by a pattern, select **Validate All Pattern Commands**.
3. Fill in the fields as described in Command Validation Tasks form for pattern-based discovery.
4. Select and hold (or right-click) the form header and select **Save**.
5. **Optional:** If the commands you selected contain temporary variables, add the values for these variables before validating such commands.

   a. Select **Refresh Temporary Variable Mappings** under **Related Links**.
   The tool creates a list of all temporary variables for the chosen commands and IPs.

   b. Select the **Temporary Variable Mappings** tab and enter the values for the temporary variables.

Temporary Variable Mappings tab

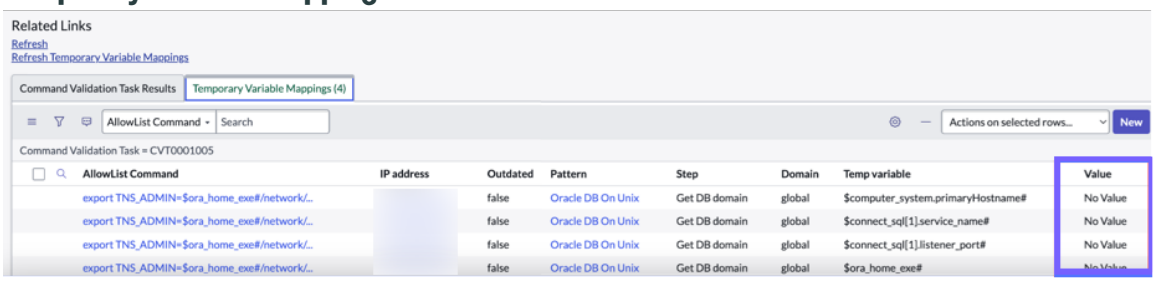

6. Select **Run Command Validation**.
7. Check if **FAILURE** appears in the **Result** column in the **Command Validation Task Results** section.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

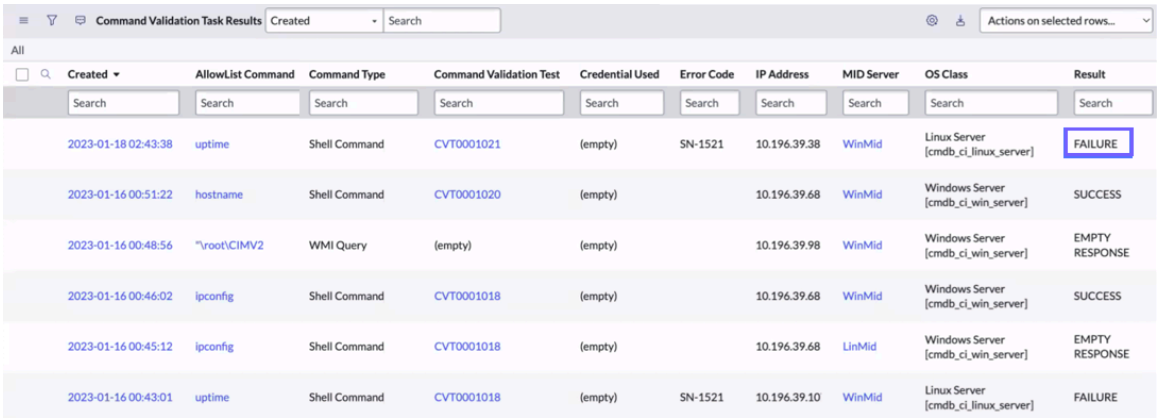

8. Select the link in the **Command Validation Test** column.
9. Review the description in the **Results Details** field.

What to do next

1. Select the **Open suggestions to troubleshoot failure** link.

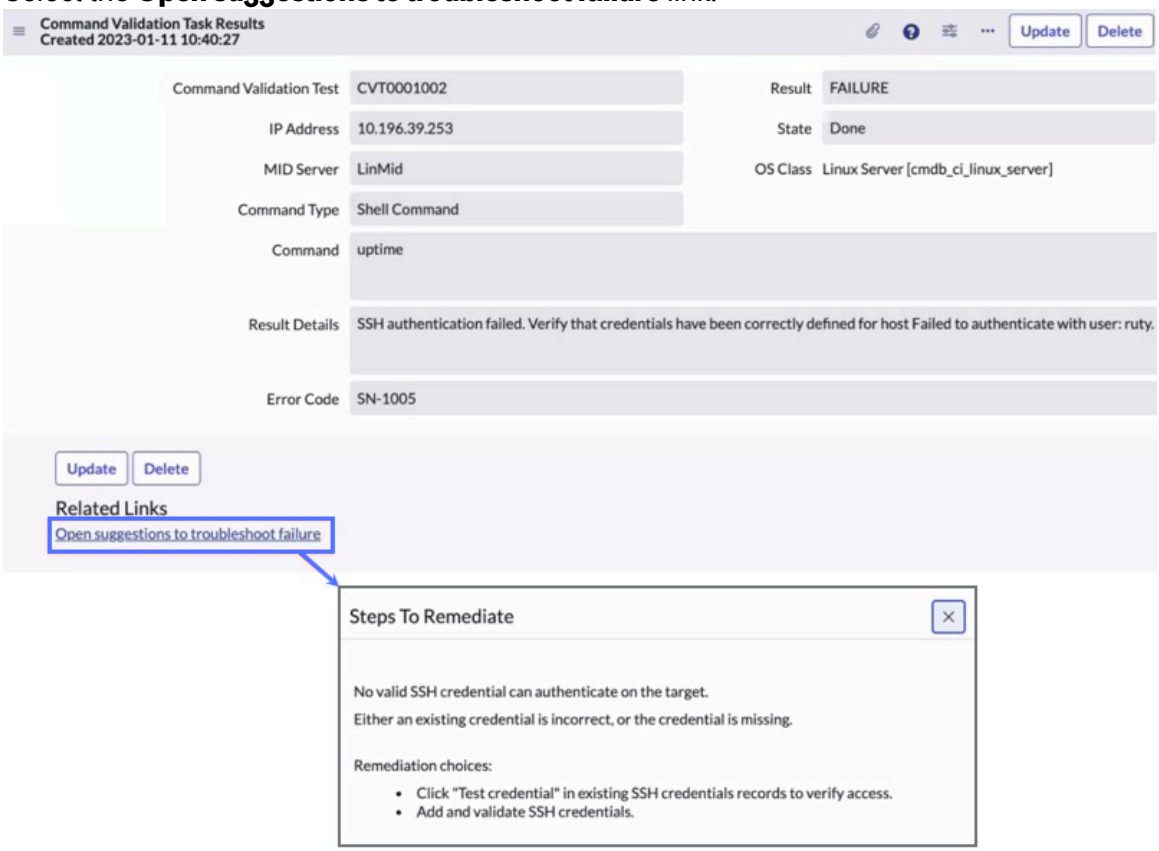

2. Follow the instructions in the **Steps to Remediate** window.

Related topics

- Discovery commands for probes and patterns
- Getting started with credentials
- Command List form for pattern-based discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View command validation task results in pattern-based discovery

View the results of command validation tasks to verify you have the correct credentials and permissions to run pattern-based discovery successfully.

Before you begin

- Verify that Pattern Designer Enhancements (sn_itom_pde) is installed.
- Verify that a command validation task has been created. For more information, see Validate commands used in pattern-based discovery.

Role required: One of the following.

- discovery_admin
- pde_viewer (starting with Pattern Designer Enhancements version 3.9.0)

About this task

Starting with Pattern Designer Enhancements version 3.9.0, the new pde_viewer role can view the command validation task results through the **Command Validation Tasks** and **Command Validation Task Results** modules but can't edit or modify them.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Command Validation Tool** > **Command Validation Tasks**.
2. In the **Number** column, search for the relevant command validation task.
3. Select the relevant command validation task.
4. Verify that the **Status** field is set to **Complete**.
5. In the **Command Validation Task Results** tab, check if **FAILURE** appears in the **Result** column.
6. If any of the commands failed, select the link in the **Command Type** column.
7. Review the failure description in the **Results Details** column.

What to do next

If you're using the discovery_admin role and a command fails, follow the remediation suggestions in the **Open suggestions to troubleshoot failure** related link.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Select the **Open suggestions to troubleshoot failure** link.

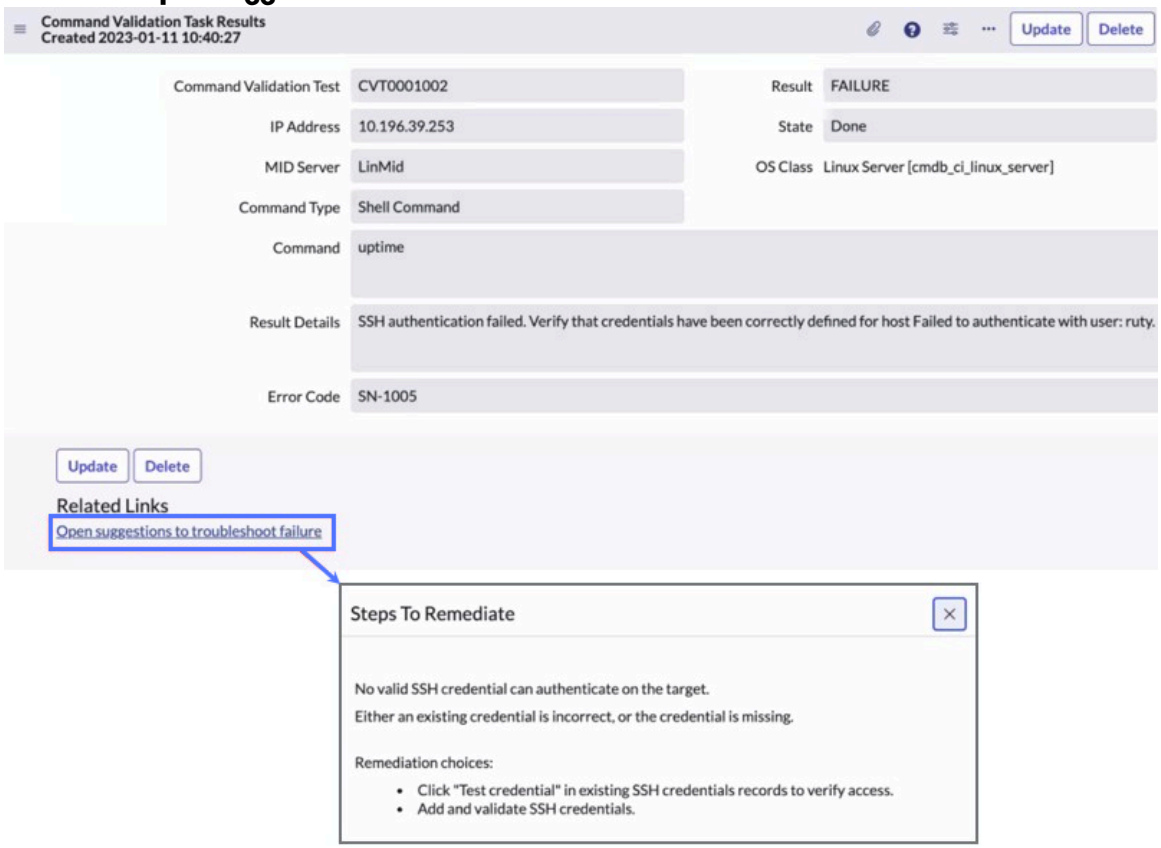

2. Follow the instructions in the **Steps to Remediate** window.

Note: This functionality isn't available to users with the pde_viewer role.

Related topics

Command Validation Tasks form for pattern-based discovery

Discovery commands for probes and patterns

Command List form for pattern-based discovery

Discovery on Code Signing instances

Discovery is supported on instances with Code Signing enabled. Code signing validates the integrity of Discovery components before execution to help prevent unauthorized scripts from running on MID Servers and target machines.

How it works

When Code Signing is enabled, Discovery validates signatures on components that contain scripts or commands before execution. Validation occurs at both the instance level and the MID Server level.

The instance validates signatures on probes and probe parameters. If a probe is tampered with, the instance rejects it and doesn't send it to the MID Server for processing. If probe parameters aren't trusted, they aren't added to the probe. The probe still launches, but without the untrusted parameters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The MID Server validates signatures on patterns, MID script includes, and MID script files. If a pattern's signature is invalid, the MID Server won't download it. If the pattern was previously downloaded, it's deleted. Discovery can't run that pattern on any machine until it's re-signed.

Signatures are included automatically with Discovery applications and plugins. When your platform and store apps meet the requirements, code signing for Discovery works without manual signing. For more information about how code signing works, see Code Signing.

Benefits

The following examples highlight the primary advantages of using code signing:

- **Security protection**: Blocks unsigned or tampered payloads from executing on MID Servers and target machines, helping prevent malicious actors from running dangerous commands even if an instance is compromised.
- **Authenticity verification**: Confirms that probes, parameters, and sensors originate from a trusted source and haven't been modified since signing.
- **Compliance support**: Strengthens audit capabilities by helping prevent unauthorized script execution without impacting discovery performance.
- **Integrity assurance**: Validates that sensitive fields in Discovery components haven't been altered after signing.

Requirements

Code signing for Discovery requires the following dependencies:

- You must be using the Brazil version of the ServiceNow AI Platform.
- You must have the following ServiceNow Store applications and versions installed:
  - Discovery Admin Workspace v1.13.0
  - Visibility Content v6.13.10
  - Discovery and Service Mapping Patterns v1.30.2
- You must have Code Signing enabled with the Circle of Trust established. For more information, see Configuring Code Signing.

Supported discovery types

Code signing for Discovery supports the following discovery methods:

- IP-based Discovery
- Cloud Discovery

> **Note:** Cloud Discovery requires additional configuration. For more information, see Configure Cloud Discovery for code-signed instances.

Signature validation scope

Code signing validates signatures on Discovery tables that contain scripts or commands executed on target machines or MID Servers. This includes probes, probe parameters, sensors, patterns, and more.

Limitations

Pattern debugger isn't available on code-signed instances.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Code Signing

Advanced Discovery configuration

You can configure several additional components of Discovery such as Application Dependency Mapping, the ECC queue, and extensions for the MID server.

gMSA configuration for Discovery

Group managed service accounts (gMSAs) are managed domain accounts that you use to help secure services. gMSAs can be used for credential-less Discovery.

Benefits

After you configure Discovery to use gMSA, password management for that account is handled by the Windows operating system. So, you can run Windows Discovery without sharing credentials with the ServiceNow instance. Benefits include the following:

- You don’t have to handle gMSA passwords on your own.
- You can choose the cycle of gMSA password rotation for better security.
- You don't need to store the password on the ServiceNow instance.
- The gMSA user doesn't need to be member of a domain admin group.
- The gMSA user used as the MID Server service account doesn't need to be on the local admin group of the MID Server.

Configure gMSA for Discovery

Use Group Managed Service Accounts (gMSA) to securely run MID Servers and perform Windows discovery without storing passwords locally. This configuration improves security and simplifies credential management by leveraging Active Directory for automatic password rotation and centralized control.

Before you begin

- Verify that the gMSA account is created and configured in Active Directory.
- Add the gMSA accounts to the MID Server host's local Administrator's group.
- For target servers, add the gMSA account to their local Administrator's group.

Role required: agent_admin, discovery_admin, or admin

Procedure

1. On the PowerShell command line, create a KDS Root key on a Domain Controller using the following commands:

   ```
   Add-KdsRootKey -EffectiveImmediately
   ```

   or

   ```
   Add-KdsRootKey –EffectiveTime ((Get-Date).AddHours(-10))
   ```

2. Set up your gMSA and security groups using the group managed services account information https://docs.microsoft.com.

   Note: It's not usually necessary to add discovered Windows servers to the GMSAGroup. Member hosts obtain the current and previous password values directly from a domain controller, and Discovery doesn't require this step.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Start the MID Server with the gMSA account following the directions here on using gMSA: Install a MID Server on Windows
4. Create a Windows credential on the instance and select the check box **Use MID Server Service Account**.
5. Launch a Discovery on the server hosting the MID Server and another computer.

Credential-less discovery with Nmap

When authentication failure prevents configuration item (CI) identification, Discovery or Service Mapping can run selected Nmap commands through a MID Server. This collects basic CI information without credentials.

A MID Server administrator can install Nmap on individual MID Server instances running on a Windows host. Those MID Server instances can then discover some basic information about CIs in your network when normal authentication fails.

Important: Self-hosted users whose network security doesn't permit downloads from `install.service-now.com` must install and configure Nmap manually on their system. Refer to Install Nmap on a self-hosted system for instructions.

Credential-less discovery can create or modify host and application CIs when credentials are missing or misconfigured. If a credential-based discovery is performed successfully after Nmap creates a CI, the system reconciles the information gathered from each type of discovery.

What Nmap can discover

The Nmap commands executed during credential-less discovery can:

- Perform reverse DNS name resolution to identify the host from the IPv4 address.
- Return the MAC address of the host if that host is on the same subnet as the host executing the Nmap command.
- Detect applications installed on a target host.
- Detect the operating system of a target host and the OS version.

Note: Credential-less discovery classifies routers and switches as hardware. It does not create or update CIs specifically for them. Use it only on known subnets where credentials aren't viable, and don't use it long term.

Nmap credential-less discovery scans in cloud computing platforms

Running Nmap scans to or from resources within cloud computing services such as Amazon Web Services (AWS), Microsoft Azure, IBM Cloud, or Google Cloud Platform (GCP) is often against the terms of service. For example, the AWS environment is tightly regulated and requires the permission of AWS through the AWS Vulnerability/Penetration Testing Request form. Unauthorized tests against AWS services or AWS-owned resources are prohibited. For this reason, credential-less discovery within a cloud computing service environment is not appropriate, and if a violation of their policy occurs, could result in expulsion from the service. Contact your platform service provider for information on limitations or permission requirements for running Nmap.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Components installed with Nmap

The Discovery - IP Based plugin (com.snc.discovery.ip_based) that provides the Nmap functionality is activated automatically when either Discovery or Service Mapping is active. These Nmap components are provided by the Discovery - IP Based plugin:

Component	Description
System properties	• *mid.discovery.credentialless.enable*: Enables or disables Nmap for all MID Server instances on which Nmap is installed that are connected to the instance. This property is installed with the Discovery plugin and is enabled by default. It is configurable by a system administrator. • *mid.discovery.credentialless.alt_port_opti* Starting from Discovery and Service Mapping Patterns version 1.31.0, controls the port range that Nmap scans during credential-less discovery. Set to F to use Nmap fast mode (top 100 ports), or T to scan the top 1,000 ports (Nmap default). By default, ports are collected from the IP Services [cmdb_ip_service] table based on records where the Credentialless Discovery [cl_discovery] field is set to true. • *mid.discovery.credentialless.include_os_sc* Controls whether Nmap can guess the operating system when it can't confirm an exact match. Set to **true** to expand the possibility of an operating system guess, or **false** to report an operating system only on an exact match. This property isn't installed by default and a system administrator must add it to the MID Server Property [ecc_agent_property] table. Its default value is false. Enable it if hosts aren't returning an operating system during credential-less discovery. • *mid.discovery.credentialless.include_udp_s* Controls whether the scan checks UDP ports and TCP ports. Set to **true** to scan both UDP and TCP ports, or **false** to scan TCP ports only. This property isn't installed by default and a system administrator must add it to the MID Server Property [ecc_agent_property] table. It's turned off by default. Its default value is false. Enable

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Component	Description
	it if a host requires a UDP response to be identified. **Note:** Enabling UDP scanning increases scan time, especially if you widen the port range by setting *mid.discovery.credentialless.alt_port_c* to T. A wide scan can exceed the command timeout and fail. The timeout comes from the *sa.local_command_timeout_ms* MID Server parameter, which defaults to 60 seconds when it isn't set. If UDP scans are timing out, increase this value, because a wide scan can take up to about 30 minutes.
MID Server properties	These properties, from the MID Server Property [ecc_agent_property] table, aren't intended to be configured: • *mid.nmap.version*: Version of Nmap that is installed on MID Server instances in your environment. This field is visible on the MID Server [ecc_agent] form after Nmap is installed. • *nmap.safe.scripts*: Defines the list of Nmap scripts that are classified as safe for use during execution of Nmap's Application Version Detection phase (-sV command option). • *nmap.npcap.version*: The version of Npcap that is installed with Nmap. The Nmap installer can only perform upgrades of existing Npcap installations it encounters.
Fields	• **Credentialless Discovery Port** [cl_port]: Optional field on the Application [cmdb_ci_appl] table that displays the number of a port scanned by credential-less discovery. This port number is used to determine whether an application returned by Nmap has a matching CI in the CMDB or if a new CI must be created. • **Discovery source** [discovery_source]: Optional field in the Configuration Item [cmdb_ci] table to which the **CredentiallessDiscovery** choice is added. This option shows that credential-less discovery was used to create a CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Component	Description
Nmap MID Server capability	The Nmap MID Server capabilities is added to the MID Server when Nmap is installed and removed automatically when Nmap is uninstalled. Only MID Server instances with this capability can perform credential-less discovery. A system administrator can't add or remove this capability manually. Self-hosted users who have the maint role can modify or delete the Nmap capability, but shouldn't do so. Service Mapping doesn't check for the presence of the **Nmap** capability and selects the MID Server based on the IP address only. To prevent Service Mapping from selecting a MID Server without the **Nmap** capability, install Nmap on all MID Servers assigned to the IP address ranges for which you want credential-less discovery to be available. If Service Mapping selects a MID Server for credential-less discovery that doesn't have Nmap capabilities, this error message appears in the map, at the site of the CI being discovered: `Nmap is not installed on MID Server. Verify all MIDs configured to handle selected IP Address have Nmap Capability. Nmap root directory path does not exist: <path>` **Note:** The ALL MID Server capability does not include the Nmap capability.
Npcap	Npcap is Nmap's packet capture library for Windows. Npcap enables Nmap to perform port scans quickly and to identify the family of the operating system running on the target. Only one copy of Npcap is installed per MID Server host. Because Npcap can be used by other applications, uninstalling Nmap does not automatically uninstall Npcap. You must uninstall Npcap manually, after determining that no other dependencies exist.
Patterns	• Credentialless Discovery Network Device: Scans a host IP address using an Nmap command to identify the host. This pattern launches the **Credentialless Discovery Network Device - PreLaunch** script to retrieve the list of ports to explore from the IP Service [cmdb_ip_service] table. Don't modify this script.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Component	Description
	• Credentialless Discovery Application: Scans a port at an IP address using an Nmap command to identify the application service actively listening on that port. Service Mapping launches this pattern when all credential-based port classification steps fail. Discovery creates a CI in the Application [cmdb_ci_appl] table if the port is open and it can identify the service by name and product. If the service does not respond to any scan attempts, Nmap consults its nmap-services registry. Nmap then guesses which service is most likely running on that port. If Nmap guesses the application on a scanned port, the Credentialless Discovery Application pattern does not create or update an application CI.
MID Server script includes	• `SetCredentialLessDeviceClassName`: Determines which host CI to create or update after the successful execution of the Nmap command. don't modify this script. • `CredentialLessApplicationClassNameMapper`: Maps the service product, service name, and extra service information supplied by Nmap for the scanned port to a supported application table in the instance. System administrators can modify this script. • `SetCredentialLessApplicationClassName`: Verifies that the *`CredentialLessApplicationClassNameMapper`* script is invoked only once. don't modify this script.
System script include	The *`CredentiallessDiscoveryAjax`* script include runs on the instance and handles the installation and uninstallation of Nmap on Windows MID Server instances, executed from UI actions on the form. don't modify this script.

Credential-less host Discovery

Credential-less host discovery occurs when a scanned host is found to be alive, but not active, or when all configured credential-based classification probes have failed.

How the host Discovery pattern is launched

If the Shazzam probe scans a host that is alive but not active, and if credential-based classification probes fail, horizontal Discovery launches the Credentialless Discovery Network Device pattern to gather host information. If the host being scanned does not have a CI defined, Service Mapping launches the HorizontalDiscoveryProbe probe, which in turn launches the Credentialless Discovery Network Device pattern. This pattern attempts to create a new CI if

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

one does not already exist for the scanned host or to update an existing CI in the Hardware [cmdb_ci_hardware] table.

Important: To allow the Credentialless Discovery Network Device pattern to launch, ensure that the *mid.discovery.credentialless.enable* system property is set to **true**. To disable, ensure this property is set to **false**.

ECC Queue entries

The system creates these entries in the ECC queue during execution of the HorizontalDiscoveryProbe.

Credential-less Discovery

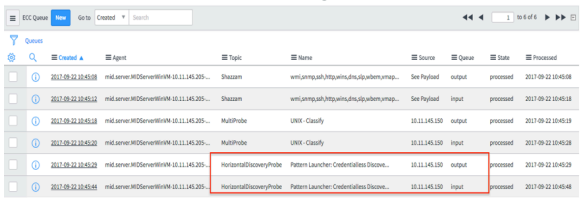

Log messages

These log messages are published during execution of the HorizontalDiscoveryProbe.

Credential-less host identifier

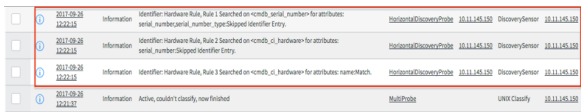

The Nmap command

The Nmap command executed on the MID Server determines if the host is up. Using the IP address and a list of ports passed to Nmap by the pattern, Nmap performs reverse DNS name resolution to identify the host name associated with the IP address. If multiple host names are configured, the first name returned by Nmap is used. If no host name is configured, then the IP address of the remote host is used to create the host CI. The Npcap packet capture library, included with the Nmap installation, identifies the host operating system family. If the scanned host is located on the same subnet as the Windows MID Server host that executes Nmap, the remote host's MAC address is returned.

Note: The list of ports that Nmap is configured to scan during credential-less host Discovery is stored in the IP Service [cmdb_ip_service] table, which is editable. By default, all ports are available for scanning. To block the use of any port for an Nmap scan, set the value in the **Credentialless Discovery [cl_discovery]** column to **false**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credential-less port list

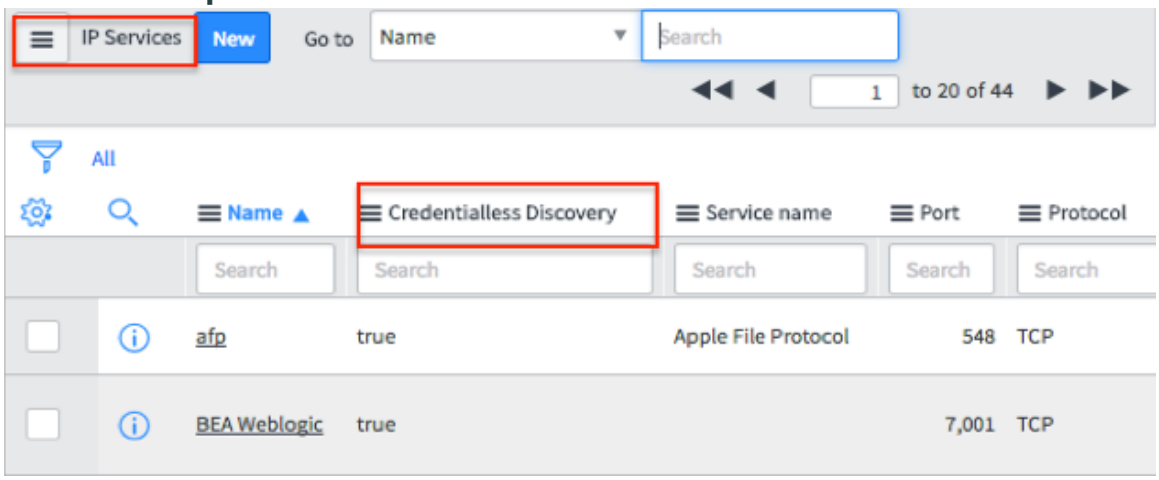

Creating or updating host CIs

After successful execution of the Credentialless Discovery Network Device pattern, the *SetCredentialLessDeviceClassName* MID Server script runs to identify the operating system family of the discovered host. The system class of the host CI is based on the operating system family that is returned by Nmap. If the OS family matches one of the six supported server operating systems, then the system uses a server class derived from the Hardware [cmdb_ci_hardware] base class, such as Linux Server [cmdb_ci_linux_server]. If no match is found, the system uses the base class.

Note: When the issue with the credentials is resolved and Discovery runs again, the instance uses the serial number, host name, and system class provided by credential-based discovery to update the host CI that was created by credential-less Discovery.

Supported OS families

OS family	CI table
AIX	cmdb_ci_aix_server
HP-UX	cmdb_ci_hpux_server
Linux	cmdb_ci_linux_server
Solaris	cmdb_ci_solaris_server
OS X or iOS	cmdb_ci_osx_server
Windows	cmdb_ci_win_server
Undefined	cmdb_ci_hardware

Hardware identification

The Discovery - IP Based [com.snc.discovery.ip_based] plugin adds an identifier to the Hardware Rule for the Hardware [cmdb_ci_hardware] base table that matches on the host CI **name** for Nmap scans. The Hardware Rule is used by both credential-based and credential-less Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credential-less hardware rule

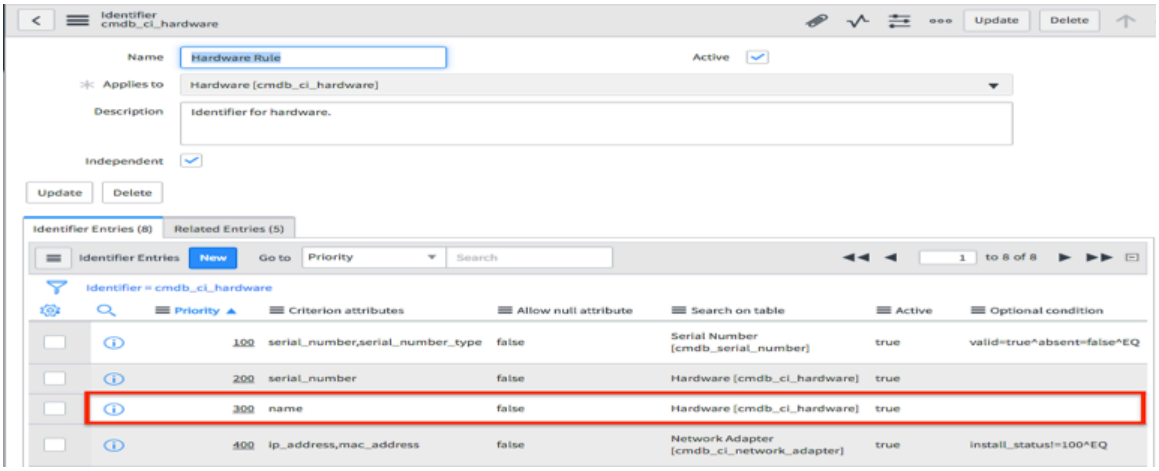

Credential-less Application Discovery

Credential-less Application Discovery attempts to identify an application service actively listening on a specific port at a given IP address.

The application discovery pattern

Service Mapping launches the Credentialless Discovery Application pattern when all credential-based port classification steps fail. The pattern executes an Nmap command on a Windows MID Server with Nmap installed, that is configured to perform application/version detection against a specific remote host IP address and port. If the port being scanned by Nmap is open, the pattern executes the `CredentialLessApplicationClassNameMapper` MID Server script include, which maps the service product, service name, and any extra information supplied by Nmap to a supported ServiceNow application table. If the script can map the returned product to an appropriate table derived from the base Application [cmdb_ci_appl] table, the script passes this information to the pattern. The pattern passes the match to the Discovery identifier for eventual CI creation or reconciliation. If the information returned by Nmap does not match any derived table, then the instance uses the base Application [cmdb_ci_appl] table to create the CI.

Important: To allow the Credentialless Discovery Application pattern to launch, ensure that the `mid.discovery.credentialless.enable` system property is set to **true**. To disable, ensure this property is set to **false**.

Example scan

This information was returned by an Nmap Application/Version Detection port scan on a Linux test system and illustrates the type of application data Nmap scans can return.

Information	Port
Port state	open
Service name	ssh
Service product	OpenSSH
Extra service information	Protocol 2.0

Default application mappings

The `CredentialLessApplicationClassNameMapper` MID Server script include is configured with a subset of the most common application tables available for Discovery and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping. A user with the agent_admin role can edit this script include to add CI tables that credential-less application Discovery can use for mapping to a derived application CI class.

Shown in this table are examples of close matches returned by Nmap on a test system that CMDB Identification and Reconciliation was able to resolve into defined products. In many cases, the service name returned by Nmap was not needed to determine a match. Service names that appear in the table were required to determine a match.

Products returned by Nmap that cannot be resolved into defined Discovery products use the following naming format: **<serviceProduct>:<serviceExtrainfo>:<serviceName>**. If any value in this string is null, it is dropped from the name.

- **serviceProduct**: Service product information returned by Nmap.
- **serviceExtrainfo**: Any additional information that Nmap returns about the application that might help identify it, such as protocol information.
- **serviceName**: The installed name of the service or daemon of the product.

Examples of close matches that were resolved

Nmap response			Script response - Identifier input	
Service product	**Service name**	**Service extra information**	**Discovered product**	**CI application table**
Apache Tomcat/ Coyote JSP engine 1.1	N/A	NULL	Tomcat	cmdb_ci_app_server_tomcat
Apache httpd 2.2.10 ((Linux/ SUSE))	N/A	NULL	Apache Web Server	cmdb_ci_apache_web_server
IBM HTTP Server	N/A	Derived from Apache	Apache Web Server	cmdb_ci_apache_web_server
IBM DB2 Database Server (QDB2/LINUX)	N/A	NULL	DB2 Instance	cmdb_ci_db_db2_instance
Microsoft Exchange smtpd	smtp	NULL	Exchange Client Access Server	cmdb_ci_exchange_cas
Microsoft Exchange 2010 log copier	msexchange-logcopier	NULL	Exchange Mailbox	cmdb_ci_exchange_mailbox_serve
JBoss service httpd	N/A	NULL	JBoss	cmdb_ci_app_server_jboss
Microsoft IIS httpd 6.0	N/A	NULL	Microsoft iis Web Server	cmdb_ci_microsoft_iis_web_serve
Microsoft SQL Server 2005 9.00.4035; SP3	N/A	NULL	Microsoft SQL Server	cmdb_ci_db_mssql_instance
MongoDB 2.5.1	N/A	NULL	MongoDB Instance	cmdb_ci_db_mongodb_instance
MySQL 5.5.51	N/A	NULL	MySQL Instance	cmdb_ci_db_mysql_instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Examples of close matches that were resolved (continued)

Nmap response			Script response - Identifier input	
Service product	**Service name**	**Service extra information**	**Discovered product**	**CI application table**
nginx 1.4.6 (Ubuntu)	N/A	NULL	Nginx Web Server	cmdb_ci_nginx_web_server
PostgreSQL DB	N/A	NULL	PostgreSQL Instance	cmdb_ci_db_postgresql_instance
Oracle WebLogic Server	N/A	NULL	Weblogic	cmdb_ci_app_server_weblogic
IBM WebSphere MQ 6.0	N/A	NULL	IBM WebSphere MQ	cmdb_ci_appl_ibm_wmq
IBM WebSphere Application Server 6.1	N/A	NULL	IBM Websphere	cmdb_ci_app_server_websphere
OpenSSH : ssh	N/A	NULL	OpenSSH	cmdb_ci_appl
Oracle Instance	N/A	NULL	Oracle Database	cmdb_ci_db_ora_instance
Oracle Instance	N/A	NULL	Oracle TNS Listener	cmdb_ci_db_ora_instance
product-A	service-B	NULL	product-A:service-B	cmdb_ci_appl
product-A	service-B	extrainfo-C	product-A:extrainfo-C:service-B	cmdb_ci_appl

Examples of applications not uniquely matched

In this example, the information returned by Nmap does not match any derived table, and the instance must use the base Application [cmdb_ci_appl] table to create the CI.

Scanned application	Nmap response			Script response - identifier input	
	Service product	**Service name**	**Service extra information**	**Discovered product**	**CI application table**
ExchangeHub	Microsoft Windows RPC	msrpc	null	Microsoft Windows RPC:msrpc	cmdb_ci_appl
HAProxy Load Balancer	IBM HTTP Server (Derived from Apache)	http	null	IBM HTTP Server (Derived from Apache):http	cmdb_ci_appl
SharePoint	Oracle Database	http	null	Oracle Database:http	cmdb_ci_appl

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scanned application	Nmap response			Script response - identifier input	
	Service product	Service name	Service extra information	Discovered product	CI application table
SharePoint	Oracle Instance	N/A	null	Oracle Database	cmdb_ci_appl

Application identification

The Discovery - IP Based [com.snc.discovery.ip_based] plugin adds an identifier to the Application Rule for the Application [cmdb_ci_appl] table that matches on sys_class_name and cl_port for Nmap scans.

Nmap identifier for the Application Rule

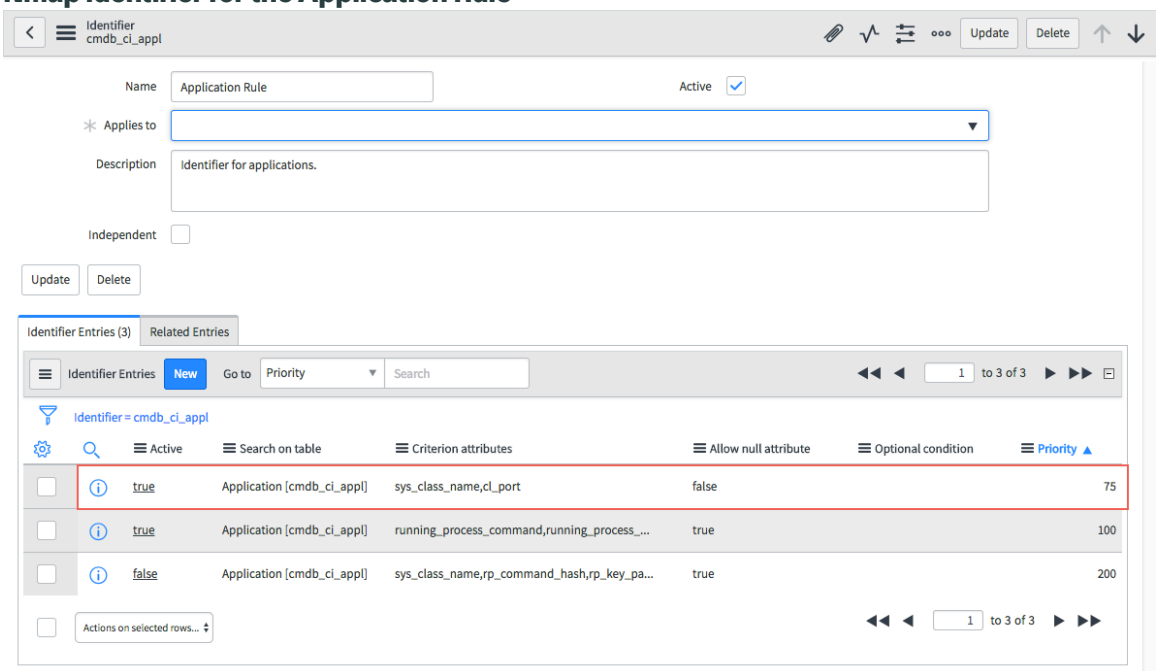

Nmap commands and data collected with credential-less Discovery

Nmap executes in phases when collecting data and runs a controlled set of safe commands with two patterns for exploring applications and devices.

Data returned by Nmap from host Discovery

Nmap returns the following data for a host it determines is up:

- **Operating system family**: For one of the six supported server operating system families, a host CI is created or updated using the appropriate server class. For unsupported operating system families, a host CI is created or updated using the base Hardware [cmdb_ci_hardware] class.
- **MAC address**: Nmap returns the remote host's MAC address if the scanned host is on the same subnet as the Windows MID Server host running the Nmap commands.

Commands run by Nmap during host Discovery

The MID Server passes the IP address of a remote host to the Nmap command and a set of ports, taken from the IP Service table. The host must be reported by Nmap to be up. The host is up if the port state returned by Nmap is **open** or **open/filtered**. If the host is reported to be up, then the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

operating system family returned is examined to determine the class used to create or update the host CI.

Command	Description
-PS	TCP SYN option determines if the host is online. This option that sends an empty, raw TCP packet with the SYN flag set to all the specified ports on the scan list. Nmap reports the host is online if a SYN/ACK TCP packet is returned and if an RST is returned, indicating that the port is closed. The Nmap host machine tears down the nascent connection with an RST, instead of an ACK, since the SYN/ACK response is unexpected from its perspective. If multiple port probes are specified, they are sent in parallel. Requirements: • Raw packet privileges. • Must run as root on UNIX. • Administrator privileges on Windows.
-PA	Detects packet-filtering devices. This option is a TCP ACK host Discovery, using a specified list of ports. If no list is defined, the default port 80 is used. Nmap sends an empty, raw TCP packet with the ACK flag set. This action acknowledges data over an established TCP connection, causing the remote host to respond with an RST packet. If the remote host does respond with an RST packet, Nmap knows the host is online. The ACK ping probe may get past packet filtering of TCP SYN packets by stateless packet filtering, network entities, such as firewalls, routers, and host routing tables/firewall software. Stateful firewalls drop unexpected packets, causing the ACK packet to be dropped, which increases the likelihood of the SYN TCP probe getting through. Requirements: • Raw packet privileges. • Must run as root on UNIX. • Administrator privileges on Windows . **Note:** Both Ping SYN (-PS) and Ping ACK (-PA) packets types can be configured together in the same scan attempt.
-O	Enables OS detection against a remote hostm providing the vendor name, underlying OS, OS generation, and device type, such as a router or a switch. Nmap uses TCP/IP stack fingerprinting to determine the OS family running on the host. Nmap sends a series of TCP and UPD packets to the host and examines the responses. After testing the responses thoroughly, Nmap compares results to its **nmap-os-db** database and prints the OS details, if there is a match. If Nmap cannot guess the OS, it checks for at least one open port and one closed port. When both are found, it generates a fingerprint and a URL to submit to `nmap.org`. OS detection does not use scripts from the Nmap Scripting Engine (NSE). OS detection runs after Application/Service Version detection,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Description
	which helps the accuracy of OS Detection. However, the results of the two scan types might be contradictory. OS Detection runs at the TCP transport layer, while Application/Service Version Detection runs at a higher level in the TCP stack. Requirements: • Raw packet privileges. • Must run as root on UNIX. • Administrator privileges on Windows .
-p	List of ports to scan.
-T4	Nmap aggressive timing template.
-v	Enable verbose detail. This command is normally used to provide progress indication to viewers of standard output. However, it can cause the Nmap Scripting Engine (NSE) scripts to provide additional detail.
-Pn	Disable the default ping. This command skips the initial Nmap host discovery scan phase and treats all hosts as if they are online. Use this command if you want a port scan even if the set of ports run by the default Discovery probes will not find an open port.
-r	Scan ports consecutively, without randomizing.
--reason	Explains why Nmap classified a port with a specific state.
--system-dns	Use the DNS resolver of the operating system for reverse DNS lookup, rather than the DNS name servers. This option is slower than using DNS servers, since it can only resolve one IP address at a time, while multiple rDNS requests are performed in parallel. However, the value returned matches the value returned by the MID Server that is used to create host CIs.
-X -	Redirects `stdout` to a XML file.
2> <file-path>	Redirects errors messages written by Nmap to `stderr` to the specified file.

Data returned by Nmap during application Discovery

Nmap returns the following application data from a host it determines is up:

- **Port state**:
  - **open**: An application is actively accepting TCP (or UDP) packets on the port.
  - **open|filtered**: Nmap cannot determine if the port is open or filtered. An open port may not return a response when using Nmap's UDP, IP protocol, FIN, NULL and Xmas scans.
- **Service name**: Name of the service actively accepting TCP (or UDP) packets on the service port.
- **Service product**: Name of the application running on the open port, as identified by Nmap.
- **Extra service information**: Any additional information returned by Nmap that Discovery can use to identify the application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Commands run by Nmap during application Discovery

The MID Server passes one IP address for a remote host to the Nmap command and one port to scan. Nmap attempts to identify the application actively listening on the port and then determines whether it should create a new CI or update an existing one. The host must be reported by Nmap to be up.

Nmap uses this information to determine which course to take:

- Service product
- Service name
- Extra service information
- State of the scanned port

Note: The service product and service name values cannot be NULL, and the service name cannot end with a question mark. Nmap appends a question mark to a service name when it guesses at the application listening on the scanned port. The instance does not create or update an application CI if Nmap cannot retrieve the application name from its static `nmap-services` registry file.

Command	Description
-sS	TCP SYN (stealth) port scan. This is the default scan option and never completes TCP connections. Scans can differentiate between ports that are **open**, **closed**, or **filtered**. TCP SYN scan implements a half-open reset that aborts the request before the connection can be established. The port is **open** (the application is listening) if a SYN/ACK is returned. The port is **closed** if an RST is returned. If no response is received after multiple retransmissions, the port is marked **filtered**. The port is also marked filtered if an ICMP unreachable error (type 3, code 1,2,3,9,10,13) is received. Requirements: • Raw packet privileges • Must run as root • Windows administrator privileges
-sV	Probes open ports to determine the service protocol, application name, version number, hostname, device type, OS Family and miscellaneous details such as the SSH protocol version. When Nmap receives a response it cannot match to the database it prints out a fingerprint and URL for submission to **nmap.org**. If looks for the port number in its static `nmap-services` registry file to determine the name of an application service that might be actively listening on the port. The service name returned by Nmap is appended with a question mark, indicating that Nmap had to guess at its identity. The instance does not create or update an application CI when Nmap guesses at the application running on an open port. **Note:** This command does not require raw-packet privileges.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Description
--datadir	Path to the directory containing the safe set of Nmap scripts to be run during Application/Version detection.
-p	Single port to scan.
-T4	Nmap aggressive timing template.
-v	Enable verbose detail. This command is normally used to provide progress indication to viewers of standard output. However, it can cause the Nmap Scripting Engine (NSE) scripts to provide additional detail.
-Pn	Disable the default ping. This command skips the initial Nmap host discovery scan phase and treats all hosts as if they are online. Use this command if you want a port scan even if the set of ports run by the default Discovery probes will not find an open port.
-r	Scan ports consecutively, without randomizing.
--reason	Explains why Nmap classified a port with a specific state.
--system-dns	Use the DNS resolver of the operating system for reverse DNS lookup, rather than the DNS name servers. This option is slower than using DNS servers, since it can only resolve one IP address at a time, while multiple rDNS requests are performed in parallel. However, the value returned matches the value returned by the MID Server that is used to create host CIs.
-X -	Redirects `stdout` to a XML file.
2> <file-path>	Redirects errors messages written by Nmap to `stderr` to the specified file.

PowerShell for Discovery and Service Mapping

MID Servers use PowerShell and PowerShell Remoting for accessing configuration items (CIs) during horizontal and top-down discovery. Review MID Server parameters and script includes, probe parameters, and credentials for using PowerShell.

PowerShell is used to control and automate the administration of Windows servers and applications.

MID Servers can use PowerShell to directly communicate with Windows servers using both WMI and WinRM protocols. For Windows services using the WinRM protocol, the PowerShell process establishes a secure PSSession (PowerShell Remoting session) that stays open until the MID Server finishes querying a Windows server. For Windows servers using the WMI protocol, the PowerShell process sends every PowerShell command with credentials.

PowerShell is also the preferred method for performing discovery over multiple Windows domains. PowerShell allows a single MID Server to authenticate on servers on different domains using credentials stored on the instance.

If you do not configure MID Servers to use PowerShell and PowerShell Remoting, MID Servers use WMI.

How PowerShell Discovery works

The following descriptions explain how MID Servers use PowerShell to deploy probes.

Probe and sensor

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When a Windows machine is classified with PowerShell, and an MSSQL instance is detected, a probe called *Windows - MSSQL* is launched. The probe returns the SQL database catalogs and version to a matching sensor.

Probe parameter

The *WMI_ActiveConnections.ps1* probe parameter contains a script that runs `netstat.exe` on a target server when PowerShell is enabled. This script extracts the information on Windows server connections, such as process IDs, ports, and IP addresses.

Credentials

Discovery uses Windows PowerShell credentials from the Credentials [discovery_credentials] table or the domain administrator credentials of the MID Server service. If Discovery cannot find PowerShell credentials in the Credentials table of the type (Windows) it uses the login credentials of the MID Server service.

MID Server Script Includes

The following script includes were added for PowerShell discoveries. These scripts run on the MID Server to generate the scripts that Discovery uses for WMIRunner and PowerShell.

- GenerateWMIScriptJS: Generates a Javascript script for the WMIRunner probe.
- GenerateWMIScriptPS1: Generates a PowerShell script for PowerShell discovery.

MID Server parameters for PowerShell

Optional parameters for the MID Server can be found at MID Server parameters for PowerShell. After changing the setting for any parameter, be sure to restart the MID Server service.

PowerShell version requirements

MID Servers using PowerShell must be installed on a supported Windows operating system. ServiceNow supports these PowerShell versions:

Version 3.0

- Regular Discovery
- Application Dependency Mapping (ADM)
- File-based Discovery
- PowerShell version 3.0 does not support Windows Server 2003.

Version 4.0

- Regular Discovery
- Application Dependency Mapping (ADM)
- File-based Discovery

Version 5.0

- Regular Discovery
- Application Dependency Mapping (ADM)
- File-based Discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: PowerShell version 6.0 is **not** supported. Many of the cmdlets that discovery relies on have been removed from this version. For example, only cmdlets using WinRM are available for remote operations.

Windows PowerShell execution policies

Windows PowerShell has four different execution policies. Customers can set the script execution policies with their group policy settings. For more information, see the Microsoft website for PowerShell documentation.

- Restricted: No scripts can be run. Windows PowerShell can be used only in interactive mode.
- AllSigned: Only scripts signed by a trusted publisher can be run.
- RemoteSigned: Downloaded scripts must be signed by a trusted publisher before they can be run.
- Unrestricted: No restrictions; all scripts can be run.

Note: If you have any policy other than Unrestricted, the script needs to be signed.

Set up MID Servers to use PowerShell

Configure MID Servers in your organization to use PowerShell for horizontal and top-down discovery of Windows servers.

Before you begin

Role required: agent_admin, discovery_admin, or admin

Note: This is enabled by default.

PowerShell version 2.0 is not supported, therefore it is recommended to upgrade to a higher version. If the MID Server detects that it has version 2.0, it may work partially. The MID Server creates a record in the mid_server issue table with the following message: PowerShell version 2 is deprecated. Upgrading PowerShell to version 5 is recommended.

Procedure

1. Download PowerShell and install it on each MID Server configured to discover Windows computers.
   If a compatible PowerShell is available on the Host, you don't need to do anything else.
2. Determine which credentials MID Servers use for PowerShell.
   - To discover Windows computers using credentials from the credentials table, set the *mid.powershell.use_credentials* parameter to **true**. Using credentials form the credentials table is the default behavior when PowerShell is enabled.
   - To force Discovery and Service Mapping to use the credentials of the MID Server service user, set the *mid.powershell.use_credentials* parameter to **false** on the MID Server. The MID Server service must have domain admin credentials to have access to the Windows machines in the domain.
3. **Optional:** If necessary, configure advanced parameters for PowerShell and PowerShell Remoting on each MID Server separately.
   Although you can disable PowerShell by changing the *mid.use_powershell* and *mid.sa.prefer_powershell*, this method is deprecated and should not be used.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PowerShell remoting for Discovery

Probe developers can use the PowerShell remote execution framework to automatically handle remote execution of scripts on target devices. The unified framework removes inconsistencies in remote execution, increases efficiency, and improves stability.

Requirements

To use PowerShell remote execution framework, the following requirements must be met:

- The MID Server must be able to write to and read from the network share of the target.
- The remote target must have PowerShell 3 or higher (up to 7).
- For MID Servers using WinRM or WMI that choose to copy the script to the remote target, the MachinePolicy and UserPolicy scopes must be set to Undefined. If the script isn't copied, the execution policy can be any other setting up to Restrictive.

See Set up MID Servers to use PowerShell for more information.

Application discovery

The PowerShell remote execution framework has options for copying files to the remote target when executing a scan. Copying files to the target is important for probes like Windows — File discovery because its script invokes itself on the remote target to spawn off a new process. MID Servers using WMI to execute scripts remotely can experience an error with launchProcess if the script is too long. Copying the script to the remote target resolves this error. Copying a script to a remote target may result in anti-virus software on the target flagging the script. To avoid issues with anti-virus software, add the scripts to the permitted list in the anti-virus app.

Probe configuration

The following is the configuration page for the Windows — Active Connections probe, which is included in the Windows — ADM multiprobe.

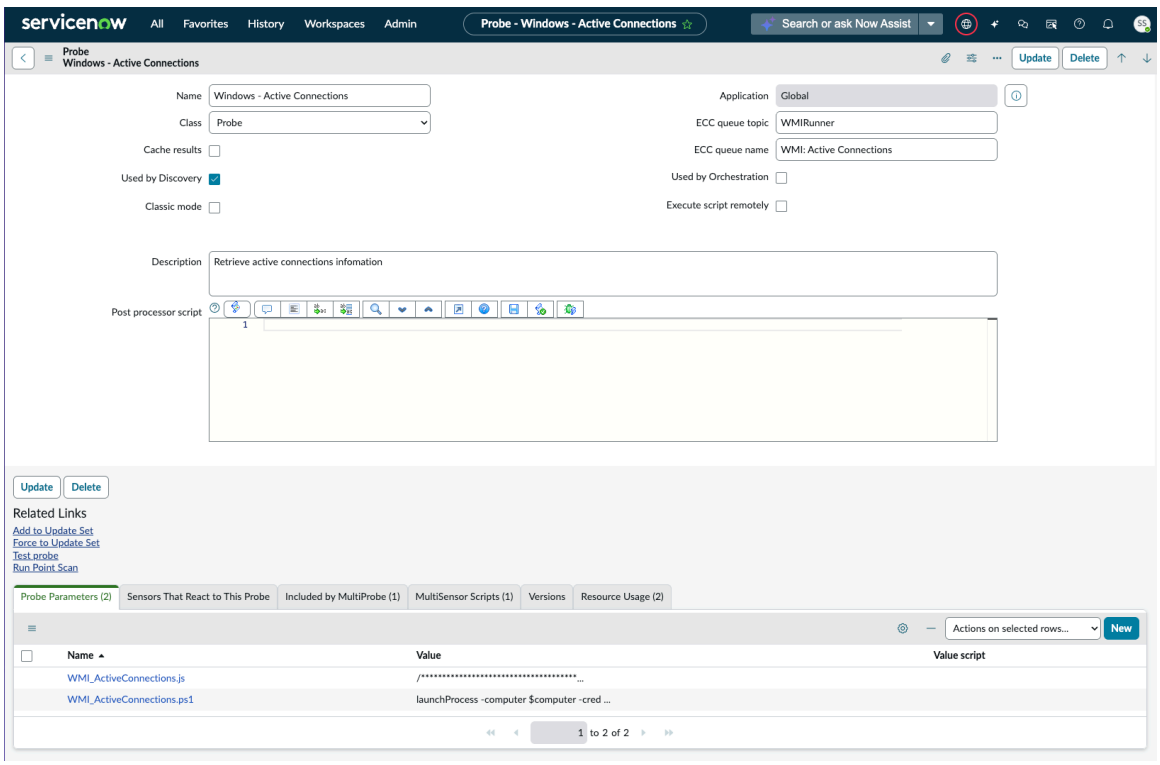

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Execute script remotely** check box is visible when the ECC queue topic is WMIRunner or PowerShell. When checked, the script executes on the remote target. Otherwise the script executes on the MID Server.

The **Copy script to target** check box is visible when **Execute script remotely** is checked. If **Copy script to target** is checked, then the script is copied to, and run on, the target. If **Copy script to target** is unchecked, then the script is executed on the target without copying it.

PowerShell probe development

PowerShell remote execution framework is a unified method of executing PowerShell scripts, contained in a probe parameter, on a remote target server. The framework removes the need for probe developers to write their own remote execution code, which can cause inconsistencies between developers. The probe developer writes the script as if the probe was collecting information locally because the remote execution framework automatically handles remote script execution.

The framework handles remote execution whether the MID Server is configured to use WMI or WinRM. If the MID Server is configured for WMI, the probe uses launchProcess to execute commands on the remote target. Using launchProcess complicates remote execution and can cause failures. However, a MID Server configured for WinRM does not use launchProcess, and so is more efficient and stable.

PowerShell commands run by Discovery

These are the PowerShell cmdlets and their parameters that Discovery runs to control and automate the administration of Windows servers and applications. Included are the probe parameters and MID Server scripts that contain the commands.

Cmdlet/Parameters	Description
Add-Content • Path • Encoding • Value	Appends content to a specified item or file. You can specify the content by typing the content in the command or by specifying an object that contains the content. **Probe** Windows - File discovery (filebaseddiscovery.ps1) **MID Server script** CollectConnectionsAndProcessesInfo.ps1
Add-Member • NotePropertyName • NotePropertyValue • MemberType • Name • Value	Lets you add properties and methods to an instance of a Windows PowerShell object. For instance, you can add a NoteProperty member that contains a description of the object or a ScriptMethod member that runs a script to change the object. **Probe** Windows - File discovery (filebaseddiscovery.ps1) **MID Server script** CollectConnectionsAndProcessesInfo.ps1
Add-Type • TypeDefinition • AssemblyName • MemberDefinition	Adds a Microsoft .NET Framework class in your Windows PowerShell session. You can then use the New-Object cmdlet to instantiate objects and use those objects, just as you would use any .NET Framework object. If you add an Add-Type command to your Windows PowerShell profile, the class is available in all Windows PowerShell sessions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
• Name • Namespace • PassThru	**Probes** • Windows - MSSQL (sqlinfo.ps1) • Windows - Azure (is.azure.ps1) **MID Server script** • Credentials.psm1 • RegistryAPI.psm1 • LaunchProc.psm1
Compare-Object • Property • PassThru	Compares two sets of objects. One set of objects is the reference set, and the other set is the difference set. The result of the comparison indicates whether a property value appeared only in the object from the reference set or, if the IncludeEqual parameter is specified, in both objects. **MID Server script** DiagnosticsUtil.psm1
ConvertTo-SecureString • String • AsPlainText • Force	Converts encrypted standard strings into secure strings. It can also convert plain text to secure strings. It is used with ConvertFrom-SecureString and Read-Host. The secure string created by the cmdlet can be used with cmdlets or functions that require a parameter of type SecureString. **Probes** Windows - Storage 2012 - PS (storage.2012.ps1) **MID Server script** Credentials.psm1
ConvertTo-Json	Converts any object to a string in JavaScript Object Notation (JSON) format. The properties are converted to field names, the field values are converted to property values, and the methods are removed. **MID Server script** CollectConnectionsAndProcessesInfo.ps1
Copy-Item • Path • Destination • Force • Recurse • PassThru	Copies an item from one location to another location in the same namespace. For instance, it can copy a file to a folder, but it cannot copy a file to a certificate drive. **MID Server script** • ADMELauncherWMI.psm1 • ExecuteRemote.psm1 • LaunchProc.psm1 • NmapInstallation.ps1
ForEach-Object - process	Performs an operation against each item in a collection of input objects. The input objects can be piped to the cmdlet or specified by using the InputObject parameter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
	Probes • Windows - JBoss Find web.xml List (indwebxml.ps1) • Windows - MSSQL (sqlinfo.ps1) **MID Server script** • CollectConnectionsAndProcessesInfo.ps1 • DiagnosticsUtil.psm1 • PSScript.ps1
Format-List	Formats the output of a command as a list of properties in which each property is displayed on a separate line. You can use Format-List to format and display all or selected properties of an object as a list. **MID Server script** RegistryAPI.psm1
Get-ChildItem • Path • File • Recurse • Filter • Directory • Include	Gets the items in one or more specified locations. If the item is a container, it gets the items inside the container, known as child items. You can use the Recurse parameter to get items in all child containers. **MID Server script** • CollectConnectionsAndProcessesInfo.ps1 • DiagnosticsUtil.psm1 • ExecuteRemote.psm1 • LaunchProc.psm1 • NmapInstallation.ps1 • PSScript.ps1 • RegistryAPI.psm1
Get-Content • Raw • Path	Gets the content of the item at the location specified by the path, such as the text in a file. It reads the content one line at a time and returns a collection of objects, each of which represents a line of content. **Probes** Windows - ADM Enhanced **MID Server script** • ADMELauncherWMI.psm1 • ADMELauncherWinRM.psm1 • ScriptAPI.psm1
Get-Date - Uformat	Gets a DateTime object that represents the current date or a date that you specify. It can format the date and time in several Windows and UNIX formats. You can use Get-Date to generate a date or time character string, and then send the string to other cmdlets or programs. **Probes**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
	Windows - File Discovery (filebaseddiscovery.ps1) **MID Server script** • DiagnosticsUtil.psm1 • LaunchProc.psm1
Get-Item	Gets the item at the specified location. It does not get the contents of the item at the location unless you use a wildcard character (*) to request all the contents of the item. **Probes** Windows - ADM Enhanced **MID Server script** • ExecuteRemote.psm1 • ADMELauncherWinRM.psm1 • ADMELauncherWMI.psm1
Get-ItemProperty • Path • Name • ErrorAction	Gets the properties of the specified items. For example, you can use this cmdlet to get the value of the LastAccessTime property of a file object. You can also use this cmdlet to view registry entries and their values. **MID Server script** • RegistryAPI.psm1 • ScriptAPI.psm1 • NmapInstallation.psm1
Get-Member - MemberType	Gets the members, the properties and methods, of objects. To specify the object, use the InputObject parameter or pipe an object to Get-Member. To get information about static members, the members of the class, not of the instance, use the Static parameter. To get only certain types of members, such as NoteProperties, use the MemberType parameter. **Probes** Windows - ADM Enhanced **MID Server script** CollectConnectionsAndProcessesInfo.ps1
Get-Process • Id • ErrorVariable (alias EV) • PSProvider	Gets the processes on a local or remote computer. Without parameters, this cmdlet gets all of the processes on the local computer. You can also specify a particular process by process name or process ID (PID) or pass a process object through the pipeline to this cmdlet. **Probes** Windows - File Discovery (filebaseddiscovery.ps1) **MID Server script**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
	• CollectConnectionsAndProcessesInfo.ps1 • Get-PEB.psm1
Get-PSDrive • Name • ErrorAction • PSProvider	Gets the drives in the current session. You can get a particular drive or all drives in the session. **Probes** Windows - ADM Enhanced **MID Server script** • ADMELauncherWMI.psm1 • Credentials.psm1
Get-PSSession	Gets the user-managed Windows PowerShell sessions (PSSessions) on local and remote computers. **MID Server script** Credentials.psm1
Get-Variable - Name	Gets the Windows PowerShell variables in the current console. You can retrieve just the values of the variables by specifying the ValueOnly parameter, and you can filter the variables returned by name. **MID Server script** • ExecuteRemote.psm1 • PSScript.ps1 • DiagnosticsUtil.psm1
Get-WmiObject • ComputerName • Filter • Credential • Impersonation • Authentication • ErrorAction (alias EA) • List • Namespace • Class	Gets instances of Windows Management Instrumentation (WMI) classes or information about the available WMI classes. To specify a remote computer, use the ComputerName parameter. If the List parameter is specified, the cmdlet gets information about the WMI classes that are available in a specified namespace. If the Query parameter is specified, the cmdlet runs a WMI query language (WQL) statement. **Probes** • Windows - MSSQL (sqlinfo.ps1) • Windows - Storage 2008 - PS (storage.ps1) **MID Server script** • ExecuteRemote.psm1 • ADMELauncherWinRM.psm1 • ADMELauncherWMI.psm1
Get-WSManInstance - ResourceURI	Retrieves an instance of a managemenet resource that is specified by a resource Uniform Resource Identifier (URI). The information that is retrieved can be a complex XML information set, which is an object, or a simple value. **MID Server script**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
	ScriptAPI.psm1
Group-Object - NoElement	Displays objects in groups based on the value of a specified property. Group-Object returns a table with one row for each property value and a column that displays the number of items with that value. **MID Server script** RegistryAPI.psm1
Import-Module • Name • DisableNameChecking • Global	Adds one or more modules to the current session. The modules that you import must be installed on the local computer or a remote computer. **Probes** Windows - Storage 2012 - PS (storage.2012.ps1) **MID Server script** • Credentials.psm1 • PSScript.ps1
Invoke-Command • Session • ScriptBlock • ErrorAction • ErrorVariable • ArgumentList (alias is args) • FilePath	Runs commands on a local or remote computer and returns all output from the commands, including errors. By using a single Invoke-Command command, you can run commands on multiple computers. **Probes** Windows - Storage 2012 - PS (storage.2012.ps1) **MID Server script** • ADMELauncherWinRM.psm1 • Credentials.psm1 • ExecuteRemote.psm1 • LaunchProc.psm1 • Get-PEB.psm1 • WMIFetch.psm1
Invoke-Expression	Evaluates or runs a specified string as a command and returns the results of the expression or command. Without Invoke-Expression, a string submitted at the command line would be returned (echoed) unchanged. **MID Server script** • NmapInstallation.ps1 • ScriptAPI.psm1
Invoke-WmiMethod • Name • ComputerName	Calls the methods of Windows Management Instrumentation (WMI) objects. **MID Server script**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
• Argumentlist • Credential • ErrorAction	• LaunchProc.psm1 • ADMELauncherWinRM.psm1 • ADMELauncherWMI.psm1
Measure-Object - Sum	Calculates the property values of certain types of object. Measure-Object performs three types of measurements, depending on the parameters in the command. **MID Server script** CollectConnectionsAndProcessesInfo.ps1
New-Item • ItemType • Path • Force • ErrorAction	Creates a new item and sets its value. The types of items that can be created depend on the location of the item. For example, in the file system, New-Item creates files and folders. In the registry, New-Item creates registry keys and entries. **MID Server script** • CollectConnectionsAndProcessesInfo.ps1 • ExecuteRemote.psm1 • LaunchProc.psm1 • NmapInstallation.ps1 • ADMELauncherWinRM.psm1 • ADMELauncherWMI.psm1
New-Object • TypeName • Property • ArgumentList	Creates an instance of a Microsoft .NET Framework or COM object. **Probes** • Windows - AWS Relationship (AWS_InstanceID.ps1) • Windows - Azure (is.azure.ps1) • Windows - MSSQL (sqlinfo.ps1) • Windows - File Discovery (filebaseddiscovery.ps1) • Windows - Storage 2012 - PS (storage.2012.ps1) **MID Server script** • ADMELauncherWMI.psm1 • CollectConnectionsAndProcessesInfo.ps1 • Credentials.psm1 • Get-PEB.psm1 • LaunchProc.psm1 • WMIFetch.psm1 XMLUtil.psm1 • ExecuteRemote.psm1
New-PSDrive • PSProvider • Root	Creates temporary and persistent drives that are mapped to or associated with a location in a data store, such as a network drive, a directory on the local computer, or a registry key, and persistent

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
• ErrorAction • Credential • Name • Scope	Windows mapped network drives that are associated with a file system location on a remote computer. **MID Server script** • ADMELauncherWMI.psm1 • ScriptAPI.psm1
New-Variable • Namer • Value • Force	Creates a new variable in Windows PowerShell. You can assign a value to the variable while creating it or assign or change the value after it is created. **MID Server script** PSScript.ps1
New-PSSession • Credential • SessionOption	Creates a Windows PowerShell session (PSSession) on a local or remote computer. When you create a PSSession, Windows PowerShell establishes a persistent connection to the remote computer. **MID Server script** PSRemoteSession.psm1
New-PSSessionOption	Creates an object that contains advanced options for a user-managed session. You can use the object as the value of the SessionOption parameter of cmdlets that create a PSSession, such as a New-PSSession, Enter-PSSession, and Invoke-Command. **MID Server script** PSRemoteSession.psm1
New-Timespan • Start • End	Creates a timespan object which represents an interval. You can use to add or subtract time from a DateTime object. **Probes** Windows - File Discovery (filebaseddiscovery.ps1)
Out-File - FilePath	Sends output to a file. You can use this instead of > operator when you need to use its parameters. **Probes** Windows - File Discovery (filebaseddiscovery.ps1)
Out-Null	Sends output to NULL, in effect, deleting it. **MID Server script** Get-PEB.psm1
Out-String	Converts the objects that Windows PowerShell manages into an array of strings. By default, Out-String accumulates the strings and returns them as a single string, but you can use the stream parameter to direct Out-String to return one string at a time. This cmdlet lets you search and manipulate string output as you would in traditional shells when object manipulation is less convenient. **MID Server script** Get-PEB.psm1
Receive-Job	Gets the results of Windows PowerShell background jobs, such as those started by using the Start-Job cmdlet or the AsJob parameter

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
	of any cmdlet. You can get the results of all jobs or identify jobs by their name, ID, instance ID, computer name, location, or session, or by submitting a job object. **MID Server script** LaunchProc.psm1
Remove-Item • Recurse • Path • Force	Deletes one or more items. Because it is supported by many providers, it can delete many different types of items, including files, folders, registry keys, variables, aliases, and functions. **MID Server script** • ExecuteRemote.psm1 • NmapInstallation.ps1 • LaunchProc.psm1 • CollectConnectionsAndProcessesInfo.ps1
Remove-Job	Deletes Windows PowerShell background jobs. You can start jobs by using the Start-Job or the AsJob parameter of any cmdlet. You can use this cmdlet to delete all jobs or delete jobs based on their name, ID, instance ID, command, or state, or by passing a job object to Remove-Job. Without parameters or parameter values, Remove-Job has no effect. **MID Server script** LaunchProc.psm1
Remove-PSDrive - Name	Deletes temporary Windows PowerShell drives that were created by using the New-PSDrive cmdlet. **MID Server script** • ADMELauncherWMI.psm1 • ScriptAPI.psm1
Remove-PSSession - Session	Closes Windows PowerShell sessions (PSSessions) in the current session. It stops any commands that are running in the PSSessions, ends the PSSession, and releases the resources that the PSSession was using. If the PSSession is connected to a remote computer, this cmdlet also closes the connection between the local and remote computers. **Probes** Windows - Storage 2012 - PS (storage.2012.ps1) **MID Server script** Credentials.psm1
Remove-Variable • Name • Value	Deletes a variable and its value from the scope in which it is defined, such as the current session. You cannot use this cmdlet to delete variables that are set as constants or those that are owned by the system. **MID Server script** PSScript.ps1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
Rename-Item • Path • NewName	Changes the name of a specified item. This cmdlet does not affect the content of the item being renamed. You cannot use Rename-Item to move an item, such as by specifying a path together with the new name. To move and rename an item, use the Move-Item cmdlet. **Probes** Windows - File Discovery (filebaseddiscovery.ps1)
Select-Object • ExpandProperty (alias Expand) • Property • ExcludeProperty	Selects specified properties of an object or set of objects. It can also select unique objects, a specified number of objects, or objects in a specified position in an array. Alias is "Select". **Probes** Windows - MSSQL (sqlinfo.ps1) **MID Server script** • CollectConnectionsAndProcessesInfo.ps1 • RegistryAPI.psm1
Select-String • Pattern • NotMatch	Searches for text and text patterns in input strings and files. You can use it like Grep in UNIX and Findstr in Windows. You can type Select-String or its alias, sls. **MID Server script** CollectConnectionsAndProcessesInfo.ps1
Select-XML • Xpath • Content • Namespace	Allows you to use XPath queries to search for text in XML strings and documents. Enter an XPath query, and use the Content, Path, or Xml parameter to specify the XML to be searched. **MID Server script** ScriptAPI.psm1
Set-Content • Path • Value • Force • Encoding	Writes or replaces the content in the specified item, such as a file. You can type the content in the command or send content through the pipeline to Set-Content. **MID Server script** • ADMELauncherWinRM.psm1 • CollectConnectionsAndProcessesInfo.ps1 • ExecuteRemote.psm1
Set-Location - Path	Sets the working location to a specified location. That location could be a directory, a sub-directory, a registry location, or any provider path. **MID Server script** Credentials.psm1
Set-StrictMode - Version	Configures strict mode for the current scope and all child scopes, and turns it on and off. When strict mode is on, Windows PowerShell generates a terminating error when the content of an expression, script, or script block violates basic best-practice coding rules. **MID Server script** Get-PEB.psm1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
Set-Variable • Name • Value • Option • Scope	Assigns a value to a specified variable or changes the current value. If the variable does not exist, the cmdlet creates it. **MID Server script** • PSScript.ps1 • CollectConnectionsAndProcessesInfo.ps1 • ExecuteRemote.psm1
Sort-Object - Property	Sorts objects in ascending or descending order based on the values of properties of the object. You can specify a single property or multiple properties (for a multi-key sort), and you can select a case-sensitive or case-insensitive sort. You can also direct Sort-Object to display only the objects with a unique value for a particular property. **Probes** Windows - MSSQL (sqlinfo.ps1) **MID Server script** DiagnosticsUtil.psm1
Split-Path • Leaf • Path • Parent	Returns only the specified part of a path, such as the parent folder, a subfolder, or a file name. It can also get items that are referenced by the split path and tell whether the path is relative or absolute. **Probes** • Windows - Weblogic - Find (findconfig.ps1) • Windows - Storage 2012 - PS (storage.2012.ps1) **MID Server script** • Credentials.psm1 • LaunchProc.psm1 • PSScript.ps1 • ExecuteRemote.psm1
Start-Sleep • S (alias for seconds) • M (alias for millisecond)	Suspends the activity in a script or session for the specified period of time. You can use it while waiting for an operation to complete or pausing before repeating an operation. **Probes** Windows - Active Connections Multiple (WMI_ActiveConnections.ps1) **MID Server script** • CollectConnectionsAndProcessesInfo.ps1 • LaunchProc.psm1 • NmapInstallation.ps1
Stop-Job	Stops Windows PowerShell background jobs that are in progress. You can use this cmdlet to stop all jobs or stop selected jobs based on their name, ID, instance ID, or state, or by passing a job object to Stop-Job.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
	MID Server script LaunchProc.psm1
Test-Path • Path • PathType	Determines whether all elements of the path exist. It returns $True if all elements exist and $False if any are missing. It can also tell whether the path syntax is valid and whether the path leads to a container or a terminal or leaf element. **MID Server script** • ExecuteRemote.psm1 • LaunchProc.psm1 • NmapInstallation.ps1 • PSScript.ps1 • RegistryAPI.psm1 • ADMELauncherWinRM.psm1 • ADMELauncherWMI.psm1 • CollectConnectionsAndProcessesInfo.ps1
Test-WSMan • Computer • Port	Submits an identification request that determines whether the WinRM service is running on a local or remote computer. If the tested computer is running the service, the cmdlet displays the WS-Management identity schema, the protocol version, the product vendor, and the product version of the tested service. **MID Server script** ScriptAPI.psm1
Where-Object	Selects objects that have particular property values from the collection of objects that are passed to it. For example, you can use the Where-Object cmdlet to select files that were created after a certain date, events with a particular ID, or computers that use a particular version of Windows. **MID Server script** • CollectConnectionsAndProcessesInfo.ps1 • DiagnosticsUtil.psm1 • WMIFetch.psm1
Write-Error • Message • Category	Declares a non-terminating error. By default, errors are sent in the error stream to the host program to be displayed, along with output. **Probes** • Windows - MySQL Configuration (mysqlconfig.ps1ps1) • Windows - MySQL Version (mysqlinfo.ps1) • Oracle - Instance Pfile (oraclepfile.ps1) • Oracle - Instance Version (oracleversion.ps1) • Windows - MSSQL (sqlinfo.ps1) **MID Server script**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
	• ADMELauncherWinRM.psm1 • ADMELauncherWMI.psm1 • ExecuteRemote.psm1 • LaunchProc.psm1 • NmapInstallation.ps1
Write-Host - NoNewLine	Writes customized output to a host. You can specify the color of text by using the ForegroundColor parameter, and you can specify the background color by using the BackgroundColor parameter. The Separator parameter lets you specify a string to use to separate displayed objects. The particular result depends on the program that is hosting Windows PowerShell. **Probes** • Windows - Network ARP Table (arpTable.ps1) • Windows - Weblogic - Find (findconfig.ps1) • Windows - JBoss Get (findjbossservicexml.ps1) • Windows - WebSphere - Web Services () • Windows - JBoss Find (findwebxml.ps1) • Windows - Weblogic - Get () • Windows - WebSphere - Cell (getcellxml.ps1) • Windows - WebSphere - Web Applications (getserverindexxml.ps1) • Oracle - Instance PFile (oraclepfile.ps1) • Oracle - Instance Version (oracleversion.ps1) • Windows - MSSQL (sqlinfo.ps1) • Windows - Storage 2012 - PS (storage.2012.ps1) • Windows - Storage 2008 - PS (storage.ps1) • Windows - Active Connections Multiple (WMI_ActiveConnections.ps1) **MID Server script** • ADMELauncherWinRM.psm1 • ADMELauncherWMI.psm1 • PSScript.ps1 • RegistryAPI.psm1 • ScriptAPI.psm1 • WMIFetch.psm1 • DiagnosticsUtil.psm1 • NmapInstallation.ps1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cmdlet/Parameters	Description
Write-Output	Sends the specified objects to the next command in the pipeline. If the command is the last command in the pipeline, the objects are displayed in the console. **Probes** • Windows - Azure (is.azure.ps1) • ExecuteRemote (testExecuteRemote.ps1) • Windows - File Discovery (filebaseddiscovery.ps1) **MID Server script** Get-PEB.psm1
Write-Verbose - Message	Writes text to the verbose message stream in Windows PowerShell. Typically, the verbose message stream is used to deliver information about command processing that is used for debugging a command. **MID Server script** • Get-PEB.psm1 • LaunchProc.psm1 • ExecuteRemote.psm1

MID Server parameters for PowerShell

Parameters control the behavior of a particular MID Server and have lower precedence than MID Server properties.

The following parameters are optional.

Note: After changing the setting for any parameter, be sure to restart the MID Server service.

Name	Description
mid.powershell_api.session_pool.max_size	Specifies the maximum number of sessions allowed in the session pool. **Note:** Setting or changing this parameter requires restarting the MID Server. • Type: Integer • Default value: 25
mid.powershell_api.session_pool.target.max_size	Specifies the maximum number of sessions allowed in the pool per target host. **Note:** Setting or changing this parameter requires restarting the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
	• Type: Integer • Default value: 3
mid.powershell_api.winrm.use_ssl	Requires the use of SSL certificates for HTTPS connections using WinRM. • Type: True | False • Default value: false
mid.powershell_api.winrm.additional_pssesion_options	Controls advanced options for a PSSession. For more information about advanced PSSession options, see https://docs.microsoft.com/en-us/powershell/module/microsoft.powershell.core/new-pssessionoption?view=powershell-6.
mid.powershell_api.winrm.remote_https_port	Configures the port for connecting to Windows servers using PowerShell over HTTPS. • Type: Integer • Default value: 5986
mid.powershell_api.winrm.remote_port	Configures the port for connecting to Windows servers using PowerShell over HTTP. • Type: Integer • Default value: 5985
mid.powershell_api.winrm.skip_ssl_cert_check	Skips the SSL certificate check when using WinRM for HTTPS connections. • Type: True | False • Default value: false
mid.powershell_api.winrm.skip_ssl_cert_check_options	Skips specific SSL certificate checks when using WinRM for HTTPS connections. Configure the MID Server to skip checks for certificates from a Certification Authority (CA), from the Common Name (CN) that identifies the host associated with the certificate, and for revoked certificates.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
	• Type: String • Default value: -SkipCACheck -SkipCNCheck -SkipRevocationCheck
mid.powershell_api.wmi.use_reverse_dns_lookup	Enable this property to utilize Kerberos authentication when employing the WMI protocol to communicate with Windows target devices. When activated, the system performs a reverse DNS look-up of IP addresses, enabling the use of Fully Qualified Domain Names (FQDN) instead of IP addresses. This can facilitate Kerberos authentication. **Note:** Note: If the reverse DNS look-up fails, the system may revert to using the IP address. Additionally, performing reverse DNS look-ups on IP addresses may impact performance. • Type: True | False • Default value: False
mid.powershell.prefer_7_plus	Sets the MID Server to prefer PowerShell 7 or later, if available, when running File-based Discovery and Windows ADME scripts locally on the MID Server host. • Type: true | false • Default value: false For more information, see Configure MID Server parameters for PowerShell 7.
mid.powershell.prefer_7_plus_remote	Sets the MID Server to prefer PowerShell 7 or later, if available, on remote target hosts when running File-based Discovery and Windows ADME. • Type: true | false • Default value: false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
	For more information, see Configure MID Server parameters for PowerShell 7.
mid.powershell.use_credentials	Determines the credentials to use for Discovery with PowerShell. A setting of *true* directs the MID Server to run probes with the Windows credentials from the credentials table. To run probes with the credentials of the user for the MID Server service, set this parameter to *false*. • Type: true | false • Default value: true
mid.use_powershell	Enables or disables PowerShell for Discovery. Restart the MID Server after changing the value. If PowerShell is not installed or the version installed is less than version 2.0, Discovery reverts to using WMIRunner. • Type: true | false • Default value: true
mid.powershell.path	Enables an administrator to point to a specific PowerShell on a MID Server in cases where more than one PowerShell is installed. Supply the path to the folder containing the PowerShell executable, for example, `C:\mypowershell` or `C:\mypowershell\`. ServiceNow automatically appends the string powershell.exe to the path. Configure this parameter when both a 32-bit and 64-bit PowerShells are active on the same MID Server, and it becomes necessary to launch the correct PowerShell for the context. 64-bit Windows employs file system redirection and the MID Server runs as a 32-bit application. If trying to specify a path in %WinDir%\System32, Windows automatically redirects to %WinDir%\SysWOW64. To avoid redirection, specify the path as %WinDir%\Sysnative. For example, instead of C:\WINDOWS\system32\WindowsPowerShell\v1.0\, specify C:\WINDOWS

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
	\sysnative\WindowsPowerShell \v1.0\. **Note:** On a 64-bit version of Windows XP, a Microsoft hotfix may be required to enable this. To discover applications running on a 64-bit Windows machine, the MID Server must be running on a 64-bit Windows host machine. • Type: string (path) • Default value: none
mid.powershell.enforce_utf8	Enable this parameter to force commands on a target Windows system to return UTF-8 encoded output. Disabling it allows the target system to use its default encoding. This parameter is only valid when PowerShell is enabled. Setting this value to false may result in incorrect values in the CMDB when non-ASCII characters are returned by a probe. • Type: true | false • Default value: true
mid.powershell.local_mid_service_credential_fallback	Specifies the login credentials the MID Server uses if all other credentials fail. • Type: True | False • Default value: true
mid.powershell_api.idle_session_timeout	Specifies the timeout value of idle PowerShell sessions in seconds. **Note:** Setting or changing this parameter requires restarting the MID Server. • Type: Integer • Default value: 60
mid.powershell.command.parameter_passing	Enable this parameter to allow passing PowerShell parameters from the command line.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
	• Type: True | False • Default value: false
mid.powershell.command.script.parameter_passing	Enable this parameter to allow passing PowerShell scripts from the command line. • Type: True | False • Default value: false
mid.windows.management_protocol	Enables administrators to select the Windows management protocol used for device and process classification. Options include: • WMI • WinRM • Type: String • Default value: WMI
mid.windows.probe_timeout	Sets the timeout interval for all Windows probes on a specific MID Server. This value is overridden by the values configured for individual probes with the wmi_timeout probe parameter. • Type: Integer • Default value: 600

Application Dependency Mapping (ADM) for Discovery

Application Dependency Mapping (ADM) discovers CIs by detecting TCP connections between devices.

When ADM discovery is enabled, Discovery always runs the ADM probes during the exploration stage to find the TCP traffic on your network. Using these TCP connections, Discovery can find additional CIs and create relationships between them.

IPv6 Is supported for ADM.

Note: Currently, ADM does not support multi-languages. If values returned are not in English, the returned data cannot be parsed properly and the discovery will fail.

Standard and enhanced ADM discovery

Standard ADM discovery

Triggers an ADM probe to take a single snapshot of TCP connections on target machines using `netstat` and `lsof` commands. However, it does not detect new TCP connections that occur between these 24-hour intervals.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Standard ADM discovery is enabled by default.

Enhanced ADM (ADME) discovery

Triggers an enhanced ADM probe that installs a script on host machines on which Discovery takes regular samples of active TCP connections. You can use a set of properties and probe parameters that control several aspects of enhanced ADM discovery.

ADME discovery is not enabled by default.

Service Mapping uses the ADM probes for traffic-based discovery. See Discovery based on Predictive Intelligence for more information.

Requirements

Both standard and enhanced ADM discovery support the following operating systems:

- Windows
- Unix
- Solaris (On Solaris machines, enhanced ADM is available in global zones. Only standard ADM is available for local zones.)
- AIX
- HP-UX

To use the WMI protocol, version 3.0 of PowerShell is required.

Note:

If you set your Windows server Execution Policy to RemoteSigned or AllSigned, ADME can run PowerShell scripts. You must import the ServiceNow certificate to the trusted publisher in order for the scripts to execute properly.

TCP tracking file for enhanced ADM

The first time that Discovery launches the enhanced probe, no TCP information is gathered. Rather, the enhanced probe only deploys the TCP tracking file, in which information about detected TCP connections is saved. After Discovery runs again, the enhanced probe relays the contents of the TCP tracking file back to the instance via the MID Server. The information is saved in the CMDB.

The TCP tracking file is saved in these locations by default:

- Windows: `admin$\temp\{your_instance_name}`. While TCP connections are being monitored, the information is stored in memory and then later saved to the file.
- Linux: `/tmp/{your_instance_name}`.

You can change the location and other aspects of the tracking file with ADM properties.

If you disable enhanced ADM, Discovery cleans up the folders from the host file system the next time Discovery runs, but the TCP tracking file remains.

Requirements for Enhanced ADM on Linux machines

Modify the sudoers file to include these commands:

Windows

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Utility tool support: netstat
- Required sudoer commands: None
- Additional requirements: None

AIX

- Utility tool support:
  - netstat
  - lsof (if netstat does not work)
- Required sudoer commands:
  - `Cmnd_Alias ADME_CMDS=/usr/bin/netstat -Aan, /usr/sbin/lsof -iTCP -n -P`
  - `discoUser ALL=(root)NOPASSWD:ADME_CMDS`
  - `Defaults:discoUser ! requiretty`
- Additional requirements: None

Solaris

- Utility tool support:
  - netstat (for Solaris version 11.2)
  - lsof (for Solaris versions prior to 11.2)
- Required sudoer commands:
  - `Cmnd_Alias ADME_CMDS = /usr/bin/netstat,/opt/csw/bin/lsof -iTCP -n -P -zglobal`
  - `discoUser ALL=(root) NOPASSWD:ADME_CMDS`
  - `Defaults!ADME_CMDS !requiretty`
  - `Cmnd_Alias ADME_PS = /usr/ucb/ps`
  - `discoUser ALL=(root) NOPASSWD:ADME_PS`
  - `Disco ALL=(root) /usr/bin/pfiles`
  - `Disco ALL=(root) /usr/bin/pgrep`
  - `Defaults!ADME_PS !requiretty`

  Note: You must add **-zglobal** for lsof to run.

- Additional requirements: Add the lsof path to the secure path. For example, if the path to lsof is `/opt/csw/bin`, the secure path should be `secure_path = /usr/bin:/usr/sbin:/bin:/sbin:/opt/csw/bin`

Note: The commands with PS apply to standard ADM as well as enhanced ADM.

HP-UX

- Utility tool support: lsof
- Required sudoer commands:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `Cmnd_Alias ADME_CMDS = /usr/local/bin/lsof -iTCP -n -P`
- `discoUser ALL=(root) NOPASSWD:ADME_CMDS`
- `Defaults!ADME_CMDS !requiretty`

Note: Netstat is not supported.

- Additional requirements: None

Centos and Ubuntu

- Utility tool support: None
- Required sudoer commands:
  - `Cmnd_Alias ADME_CMDS = /usr/sbin/lsof -iTCP -n -P, /usr/bin/netstat -antp, /sbin/ss -tlnp, /sbin/ss -tenp`
  - `discoUser ALL=(root) NOPASSWD:ADME_CMDS`
  - `Defaults!ADME_CMDS !requiretty`
- Additional requirements: None

Note the following for the sudo commands:

- `ADME_CMDS` is a variable referring to a list of commands.
- `/usr/bin/` is the directory where your lsof and netstat resides.
- `discoUser` is a user that must have root access.

Setup and use enhanced ADM

To set up and use enhanced ADM, perform these steps:

1. Enable enhanced ADM with the glide.discovery.enable_adme Discovery property.
2. Determine the base directory where the TCP tracking file is placed on the host computers. If you want to change the location, specify it in these Discovery properties:
   - `glide.discovery.adme.base_dir_unix`
   - `mid.powershell.target_base_dir`
3. Do either of the following based on the operating systems of the host computers:
   - Linux: modify the sudoers file to include the necessary commands.
   - Windows: Set the `mid.windows.management_protocol` parameter on the MID Server to use WinRM or WMI. This parameter sets the protocol that your MID Server uses to communicate with remote Windows hosts.
4. Create a schedule and run Discovery.
5. View the results in the TCP tables.
6. Optionally control both the size of the TCP tracking file and the payload using parameters on the enhanced ADM probes. See TCP connection discovery for the names of the probes and parameters.

Configure ADM discovery parameters and properties

Using properties, you can enable or disable ADM discovery, and you can set the location of the TCP tracking file for enhanced ADM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: discovery_admin or admin

About this task
You can also control both the size of the TCP tracking file and the payload using parameters on the enhanced ADM probes. See TCP connection discovery for the names of the probes and parameters.

Procedure
1. Navigate to **All** > **System Properties** > **Discovery** to set the system properties.
2. Navigate to **Discovery Definition** > **Probes** to set the probe parameters.
3. Configure the following properties and probe parameters.
   See the knowledge article KB0822767 for more information.

ADM properties and probe parameters

Probe parameter	System property	Description
	glide.discovery.enable_adme	ADME: Enable enhanced ADM probe. If "yes", the ADM Enhanced probe is triggered and only fall back to the ADM probe as needed. ○ **Type**: true | false ○ **Default value**: false
	glide.discovery.application_mapping	Enable the application mapping feature. When this property is set to true, all discoveries trigger the relevant ADM probe during the exploration phase if Discovery finds a server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe parameter	System property	Description
		Important: Disabling this property disables the creation of relationships between applications but doesn't disable the ADM probes and sensors, nor does it prevent process classification. ○ **Type**: true | false ○ **Default value**: true
sampling_interval	glide.discovery.adme.sampling_interval	ADME - Sampling interval (sec): How often to sample process and connection data in seconds. Must be a minimum of 5 seconds. ○ **Type**: integer ○ **Default value**: 120
base_dir	glide.discovery.adme.base_dir_unix	ADME - Unix Base Dir: An existing directory on the target UNIX machines to be used as a workspace. Must be an absolute path to the directory. The UNIX and Linux credentials that you use to connect to the device must also be able to access this folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe parameter	System property	Description
		○ **Type**: string ○ **Default value**: / tmp
powershell_param_base_dir	mid.powershell.target_base_dir	Powershell - Target Base Dir: A network share on the target Windows machines to be used as a workspace. If you change the directory, it must also be a network shared folder. ○ **Type**: string ○ **Default value**: admin$\temp
aggregation_interval (Unix) powershell_param_aggregation_interval (Windows)	glide.discovery.adme.aggregation_interval	ADME - Aggregation interval (sec): How often to aggregate instantaneous data into chunks in seconds. It must be a minimum of 60 seconds and it must be a multiple of ADME Sampling interval. ○ **Type**: integer ○ **Default value**: 3600
rolling_window_size (Unix) powershell_param_rolling_window_size (Windows)	glide.discovery.adme.rolling_window_size	ADME - Rolling window size: How many recently aggregated chunks to keep. An aggregated chunk consists of ADME Aggregation interval divided by ADME Sampling interval samples. Only the newest chunks are kept.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe parameter	System property	Description
		○ **Type**: integer ○ **Default value**: 24
max_total_samples (Unix) powershell_param_max_total_samples (Windows)	glide.discovery.adme.max_total_samples	ADME - Max total samples: The limit of samples being taken by the script. The script will stop running after this number of samples are taken. ○ **Type**: integer ○ **Default value**: 100

4. Click **Save**.

Discovery behaviors

Discovery behaviors determine the probes that Shazzam launches, and from which MID Servers these probes are launched.

Unlike a scan performed by a single MID Server on a designated IP address range, a behavior can assign different tasks to multiple MID Servers on the same IP address segment or on different network segments. Behaviors are available in Discovery schedules for discoveries in which configuration items (CI) are updated in the CMDB.

Behaviors can be used in the following scenarios:

- Load balancing: A behavior enables load balancing in systems that use multiple MID Servers deployed across one or more domains.
- Multiple protocols in multiple domains: Configure one MID Server to scan for all protocols on one domain and another MID Server to perform a WMI scan on a second domain.
- Access Control Lists (ACL): Discovery can scan SNMP devices protected by an ACL if the MID Server host machine is granted access by that ACL. Use a behavior to configure a MID Server to scan devices protected by an ACL.
- Devices running two protocols: Some devices might have two protocols running at the same time. Examples of this are the SSH and SNMP protocols running concurrently on one device (most common). A behavior can control which of the two protocols is explored for certain devices. The behavior then prevents the other protocol from being explored.

Behaviors also enable the efficient Discovery of SSH and SNMP devices and WMI devices running on multiple Windows domains, using multiple MID Servers.

For example, an organization has two Windows domains in its network and a variety of UNIX computers and SNMP devices. The challenge is to discover all the devices efficiently, without duplicating effort. Each domain contains a Windows MID Server which is used to scan the IP addresses from the two domains specified in the Discovery Schedule, as well as the SSH and SNMP devices. We need a Behavior that divides the work appropriately to avoid scanning anything twice. In this example, we assume that both domains are in the same geographical location, and that a single schedule is sufficient.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The preferred method for running Discovery over multiple Windows domains is to use PowerShell, which allows a single MID Server to authenticate on machines on different domains using credentials stored on the instance.

Available functionality definitions

When creating a behavior, you select a functionality definition that determines which protocols Shazzam scans. The following table lists the available options.

Discovery functionalities

Functionality	Description
Windows only (WMI)	Scans for Windows devices using WMI protocol. Requires a Windows MID Server and functionality criteria to specify the domain.
Windows, DNS, and WINS	Scans for Windows devices using WMI protocol and resolves the domain using DNS and WINS. Requires a Windows MID Server and functionality criteria.
SNMP only	Scans for SNMP devices only. Use this when scanning devices protected by an ACL that restricts SNMP access to specific MID Server hosts. Does not require functionality criteria.
All except SNMP	Scans for all protocols except SNMP, including WMI and SSH. Use this in combination with SNMP only functionality when you need separate MID Servers for SNMP and other protocols.
All except Windows (no WMI)	Scans for all protocols except WMI, including SSH and SNMP. Use this in combination with Windows only functionality when scanning multiple domains or when non-Windows devices require a different MID Server.

Create a Discovery behavior

Create a Discovery behavior to determine which probes Shazzam launches and which MID Server is used.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery Definition** > **Behavior**.
2. Select **New**.
3. Enter a name.
4. Right-click the form header and select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

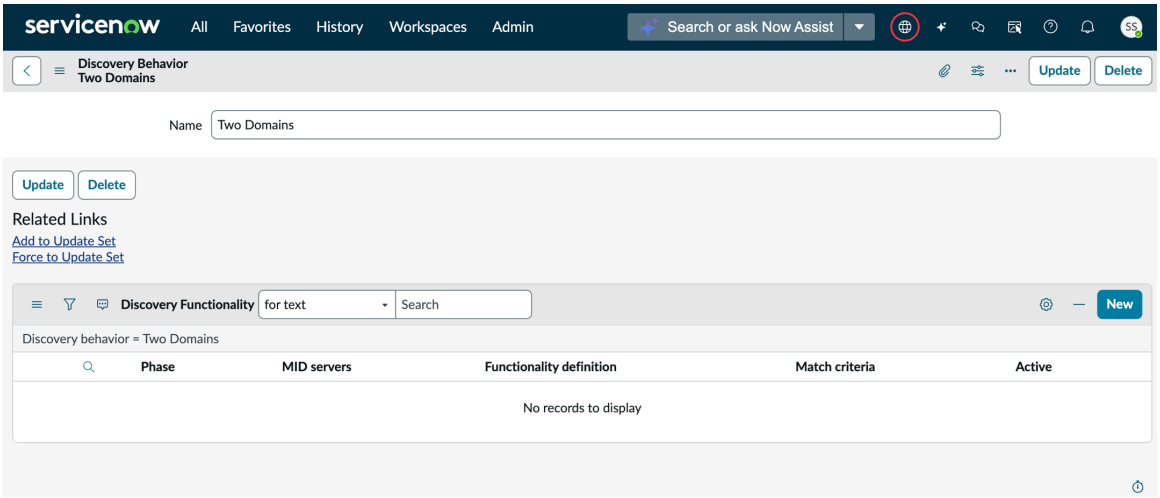

5. In the Discovery Functionality related list, select **New**.
   Discovery Functionality defines what each MID Server in this behavior must do, specifically which protocols to detect.
6. Fill out the form fields:

Field	Description
Phase	Enter an integer that represents an arbitrary phase. The phase is used to group one or more functionalities together. All the functionalities within a specified phase are executed together, and all phases are executed in numerical order. All functionalities in a behavior can have the same phase. The Shazzam probe runs once for each phase in a behavior, which makes fewer phases desirable. Run multiple phases for behaviors only when devices in the network are running multiple protocols, such as SSH and SNMP. In that example, set one phase for the SSH scan and another phase for the SNMP scan.
Active	Keep this option selected to apply the discovery functionality.
Functionality definition	Select the lookup icon, and then select a pre-configured functionality that defines the protocol or list of protocols that each MID Server scans.
Match criteria	Define criteria here for Windows MID Servers.
MID Servers	Select one or more MID Servers to perform this functionality for the following Discovery types:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ IP Scan ○ CI Scan Discovery automatically balances the load when multiple MID Servers are selected.

7. Right-click the form header and select **Save**.
8. To add criteria that the functionality must meet to be triggered, select **New** in the Functionality Criteria related list.

   Note: Functionality criteria are required for Windows MID Servers only, and only when the behavior controls Discovery across multiple domains. When the instance launches the Shazzam probe for a Discovery in which a behavior defines multiple MID Servers to scan multiple domains, the functionality criteria determine which MID Server process the results of the probe.

9. Fill out the form fields:

Field	Input Value
Name	The name in the criteria is the variable that passes the following information: ○ `mid_server`: MID Server that processes the results from the Shazzam probe. Enter the actual name of the MID Server in the Value field. ○ `win_domain`: Windows domain of the target device. Enter the actual name of the Windows domain in the Value field. ○ `mid_win_domain` Windows domain of the MID Server that is processing the Shazaam results. Enter the actual domain name in the Value field.
Operator	Select a logical operator.
Value	Enter the actual name of the MID Server (`mid_server`) or domain (`win_domain`) to pass to Discovery for this criteria. This field can also have a value of `mid_domain`, which defines the Windows domain of the MID Server that is processing the Shazzam results.

The following graphic shows an example of functionality criteria.

10. Select **Submit**.

What to do next

Schedule a horizontal discovery of type **Configuration Item**, and select **Use Behavior** for the MID Server selection method.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up a load balancing behavior

When multiple MID Servers are configured to scan the same protocol, you can set up load balancing behavior to automatically balance the work between MID Servers.

Before you begin

Role required: agent_admin, discovery_admin, or admin

About this task

Behaviors enable load balancing when multiple MID Servers are configured to scan the same protocol.

For example, an organization has 50 remote locations, varying in size from 10 devices to several hundred devices. All the satellite offices are part of the same Windows domain, which is administered from a central location. There are three MID Servers installed at this central location: Two to scan all the Windows devices and one to scan the remaining devices. Since remote offices are in different time zones, you must create different schedules to run Discovery at off-peak hours in each time zone. You can use the same behavior for each schedule, and the behavior load balances the two Windows MID Servers automatically.

Procedure

1. Create the behavior record.

   a. Create a Discovery behavior record and name it something descriptive, such as *LoadBalanced*.

   b. Select and hold (or right-click) the header bar and select **Save** from the pop-up menu.

      You save the record and create the Discovery Functionality related list in the form.

2. Define the Windows functionality.

   a. In the Discovery Functionality related list, select **New**.

   b. Create a record using the fields described in the following table.

   c. Select and hold (or right-click) the header bar and select **Save** from the pop-up menu.

Field	Input Value
Phase	Enter a phase number of `1` in this field.
Functionality	Select *Windows only (WMI)* from the list. This functionality defines the protocol that is scanned. Because you selected to scan WMI, you must select Windows MID Servers for this functionality.
MID Servers	Select the two MID Servers you want to share the load for the WMI probes. By entering multiple MID Servers in this field, you tell Discovery to balance the load between these servers automatically. If you create separate functionality for each server, load balancing wouldn't occur.
Active	Select this check box to enable this behavior.
Match criteria	Leave the default criteria of *Any*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Functionality Form

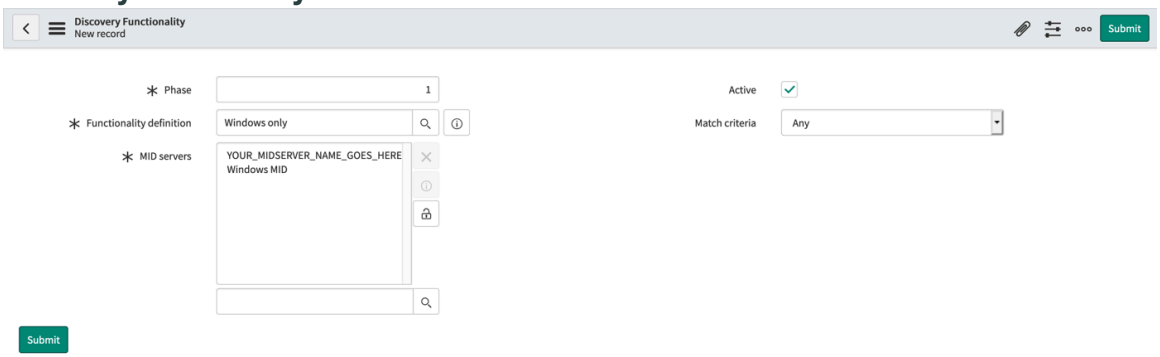

3. Create Functionality Criteria.

   Note: All Windows functionality requires criteria to identify the Windows domain and define any special behavior for the MID Servers named.

   a. In the related list, select **New**.

   b. Enter the following values:

Field	Input Value
Name	Enter `win_domain` to name the Windows domain that Discovery scans with the MID Servers you have defined.
Operator	Select *equals* as the operator.
Value	The name of the Windows domain that these MID Servers scans for devices.
Active	Enable the criteria by selecting this check box (*`true`*).

4. Define the functionality for the remaining scans.

   Create a second record to scan for all the remaining protocols (SSH, SNMP, and so on) using the following settings:

Field	Input Value
Phase	Enter a phase number of `1` in this field. This phase runs concurrently with WMI scans using the same Shazzam probe. The same phase number is used because the target IP ranges don't include devices running multiple protocols (such as SSH or SNMP). If any devices were running multiple protocols, additional phases (second or third) should be specified to identify the correct protocol for each device before proceeding.
Functionality	Select *All except Windows (no WMI)* from the list. This functionality causes the MID Server to scan all remaining protocols after Discovery has run the WMI scans.
MID Servers	Enter the name of the MID Server that scans all remaining devices. To enable automatic load balancing, add an additional MID Server in this field.

5. Create a Discovery schedule for each time zone and name the behavior you created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **Discovery** > **Discovery Schedules** and select **New** in the record list.

b. Select a **Discover** type of *Configuration items*.

c. In the **MID Server selection method** field, select *Use Behavior*.

> **Warning:** If this option isn't selected, the required value for the next step may not be available.

d. Select the search icon and choose the behavior to use.
In this example, select *LoadBalanced*.

e. Select the appropriate time to run Discovery for this location.

f. Select **Quick Ranges** and define the IP address ranges, networks, or lists to scan for this location.

g. Save the record.

h. Create additional schedules for each time zone or region in the network and select the same behavior.

Discovery Schedule

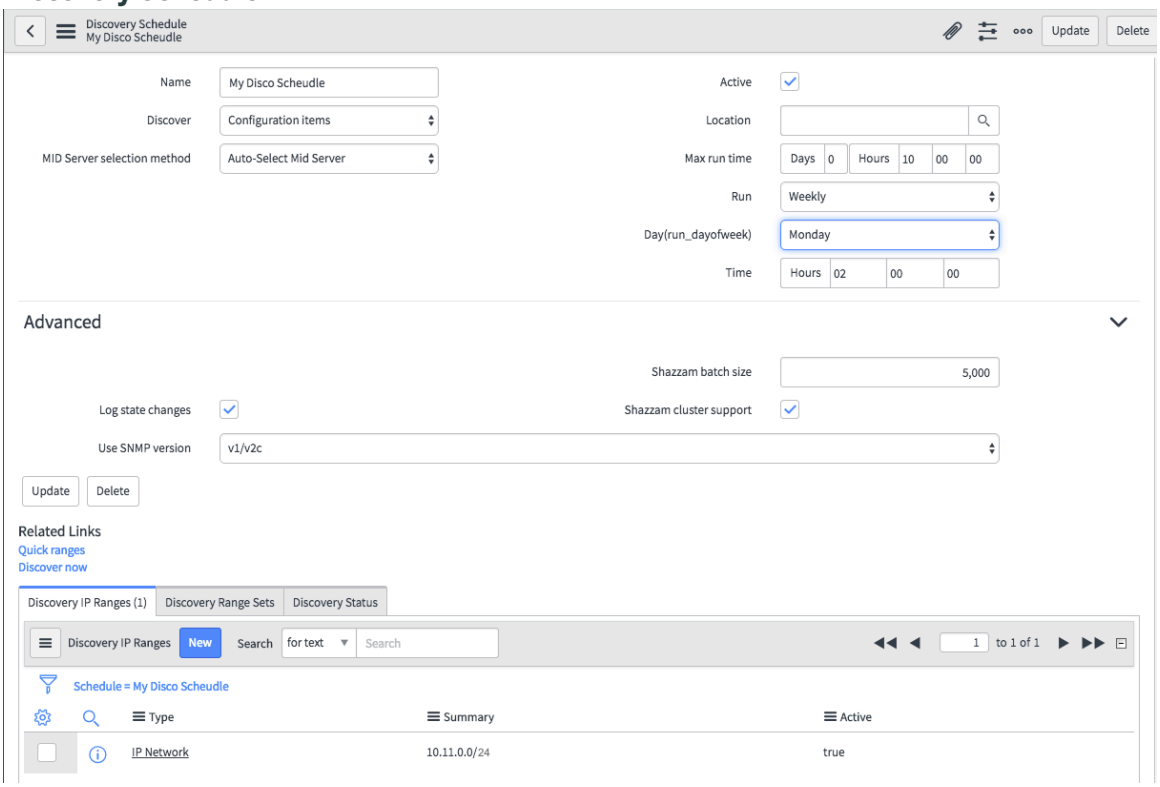

Examples of Discovery behavior functionalities

This example of a Discovery behavior requires three functionalities for the behavior.

We will create three functionalities for this Behavior: one MID Server to scan Domain A for Windows devices only; a second functionality for the same MID Server to scan for all SSH and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SNMP devices; and a third functionality that names a second MID Server to scan Domain B for Windows devices. The rationale for this is as follows:

- A Windows MID Server can only discover Windows machines on the Windows domain to which it is joined. This is entirely due to the way Windows authentication works. For this reason, we need a WMI functionality for each domain.
- A Windows MID Server, provided with the correct credentials, can discover SSH and SNMP devices anywhere; however, we cannot combine WMI, SSH, and SNMP functionalities across Windows domains. This is because the functionality criteria for the WMI scans locks in the Discovery to one specific domain. For this reason, SSH and SNMP discoveries require a separate functionality.
- We want to scan each machine only once.

Functionality 1: WMI Scanning on Domain A

We configure a MID Server to scan for the WMI protocol on Domain A. WMI scans authenticate on Windows machines using the domain credentials of the Windows MID Server machine. Windows MID Servers cannot scan for the WMI protocol outside their own domains.

Create the first functionality using the following values:

Field	Input Value
Phase	Type a phase number of **1** in this field. All functionalities in this example use the same phase number, which launches a single Shazzam probe for all the functionalities in that phase. A single phase, when feasible, is the most efficient use of the Shazzam probe.
Functionality	Select **Windows, DNS, and WINS** from the list. This functionality defines the WMI protocol that will be scanned and resolves the domain. Because we selected to scan for WMI, we must select a Windows MID Server for this functionality.
MID Servers	We select a Windows MID Server from Domain A - in this case **sandb01-358**.
Active	Make sure this check box is selected to enable this behavior.
Match criteria	Change the criteria to **All**.

Create Functionality Criteria

All Windows functionality requires criteria to identify the domain and the MID Server. In our example, we will create two criteria for this functionality. To create Functionality Criteria, click **New** in the Related List and enter the following values for the MID Server doing the WMI scanning on Domain A:

Field	Input Value
Name	Create the following criteria: • Enter `mid_server` to name the MID Server that will execute the WMI scans on Domain A. • Enter `win_domain` to name the Windows domain that Discovery will scan with the MID Server defined.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input Value
Operator	Select **equals** as the operator in this criteria.
Value	• For the *mid_server* value, enter the name of the MID Server that will scan Domain A for Windows devices. • For the *win_domain* value, enter the name of Domain A that this MID Server will scan for Windows devices.
Active	Be sure to enable the criteria by selecting this check box (true).

The completed criteria appear in the Discovery Functionality form for this behavior.

Functionality 2: SSH and SNMP

In our network, we want to scan for UNIX computers and netgear, but we don't want to classify these devices twice. One of our MID Servers will be configured to classify SSH and SNMP using a different functionality than it does for WMI scans. We do not need to create criteria for non-WMI functionality.

Create the second functionality using the following values:

Field	Input Value
Phase	Type a phase number of `1` in this field. All functionalities in this example use the same phase number, which launches a single Shazzam probe for all the functionalities in that phase. A single phase, when feasible, is the most efficient use of the Shazzam probe.
Functionality	Select **All except Windows (no WMI)** from the list. This functionality will scan SSH and SNMP protocols only.
MID Servers	We select the MID Server from Domain A - in this case **sandb01-358**.
Active	Make sure this check box is selected to enable this behavior.
Match criteria	Leave the default criteria of *Any*. Criteria are not used for non-WMI functionalities.

Functionality 3: WMI Scanning on Domain B

All that remains is to create a functionality for the WMI scans on Domain B. Because of the Windows authentication mechanism, we must configure a Windows MID Server to scan Domain B that is a member of that domain.

Create the third functionality using the following values:

Field	Input Value
Phase	Type a phase number of `1` in this field. All functionalities in this example use the same phase number, which launches a single Shazzam probe for all the functionalities in that phase. A single phase, when feasible, is the most efficient use of the Shazzam probe.
Functionality	Select **Windows, DNS, and WINS** from the list. This functionality defines the WMI protocol that will be scanned and resolves the domain. Because

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input Value
	we selected to scan for WMI, we must select a Windows MID Server for this functionality.
MID Servers	We select a Windows MID Server from Domain B - in this case **disco-win2003**.
Active	Make sure this check box is selected to enable this behavior.
Match criteria	Change the criteria to *All*.

Create Functionality Criteria

All Windows functionality requires criteria to identify the Windows domain and the MID Server. In our example, we will create two criteria for this functionality. To create Functionality Criteria, click **New** in the Related List and enter the following values for the MID Server doing the WMI scanning on Domain B:

Field	Input Value
Name	Create the following criteria: • Enter `mid_server` to name the MID Server that will execute the WMI scans on Domain B. • Enter `win_domain` to name the Windows domain that Discovery will scan with the MID Server defined.
Operator	Select **equals** as the operator in this criteria.
Value	• For the *mid_server* value, enter the name of the MID Server that will scan Domain B for Windows devices. • For the *win_domain* value, enter the name of Domain B that this MID Server will scan for Windows devices.
Active	Be sure to enable the criteria by selecting this check box (true).

The completed criteria appear in the Discovery Functionality form for this behavior.

Access an ACL-protected SNMP device

Access an SNMP device protected by an ACL using a Discovery behavior.

Before you begin

Role required: discovery_admin or admin

About this task

A specific type of Access Control List (ACL) on an SNMP device can prevent Discovery from identifying that device. This list defines host machines by IP addresses that are permitted to run agents on SNMP devices. In this example, we want to classify, identify, and update all the devices in an IP range, including the SNMP devices protected by an ACL. To do this, install a MID Server with access to the protected SNMP devices. Then create a Behavior that scans for all protocols without missing devices or doing extra work.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Install a MID Server on a host permitted by the ACL.

   To scan the SNMP devices in a network protected by an ACL, install a MID Server on a host machine specified in the ACL. The ACL must list the host's IP address as having access to the SNMP devices. Because this MID Server scans SNMP devices only, it can be installed on any supported operating system. The other MID Server, configured to discover WMI and SSH, can be installed on any Windows host that has visibility to the specified IP ranges.

2. Install additional MID Servers if you intend to configure either functionality to load balance.
3. Create the behavior record.

   a. Create a new Discovery Behavior record, and name it something descriptive, such as *Southwest-SNMP*.

   b. Right-click in the header bar and select **Save** from the pop-up menu.
   This action saves the record and creates the Discovery Functionality related list in the form.

   c. To complete this behavior (as shown in Figure 1) we must create functionalities using the MID Servers we installed for that purpose.

   For the functionality that scans for the WMI protocol, we must define criteria that specify the Windows domain and MID Server being used. Because this Discovery is being performed on one Windows domain, we can configure a functionality for WMI and SSH (**All except SNMP**) using the same MID Server.

Completed Discovery Behavior

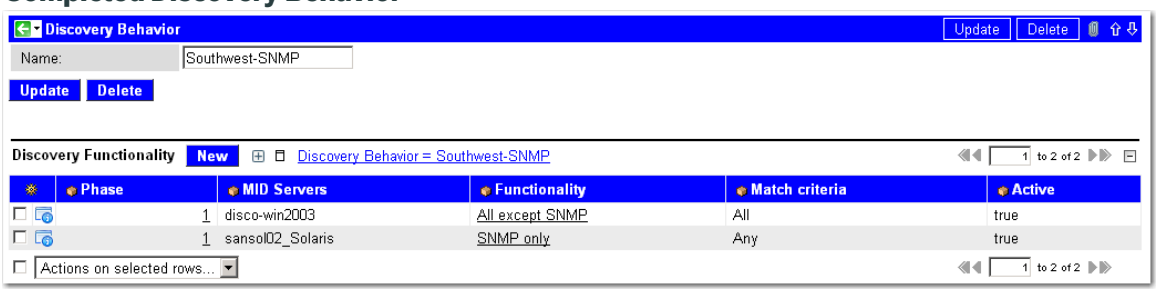

4. Define the functionalities.
   For this behavior, create two functionalities. Configure one for SNMP on the MID Server installed on the host specified in the ACL. Configure another for WMI and SSH on a second MID Server.

   a. Click **New** in the Discovery Functionality related list to add functionalities.

   b. **Note:** We configure the MID Server installed on the privileged host machine to scan for the SNMP protocol only.

   Create the first functionality using the following values:

Field	Input Value
Phase	Type a phase number of 1 in this field. Both functionalities in this example use the same phase number, which launches a single Shazzam probe for all

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input Value
	the functionalities in that phase. A single phase, when feasible, is the most efficient use of the Shazzam probe.
Functionality	Select *SNMP only* from the list. This functionality defines the protocol that will be scanned.
MID Servers	We select a MID Server (*sansol02_Solaris*) installed on a Solaris host that is granted access to the SNMP devices by the ACL. You can select multiple MID Servers if you want Discovery to load balance this functionality automatically.
Active	Make sure this check box is selected to enable this behavior.
Match criteria	Leave the default criteria of *Any*. Criteria are not used for SNMP functionalities.

Scan for SNMP Only

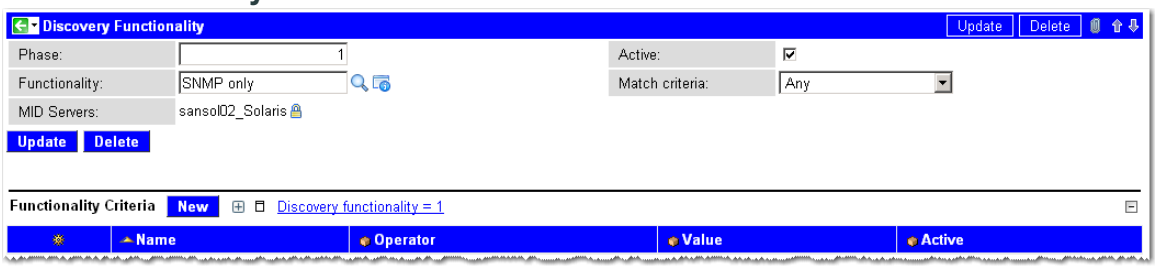

c. **Note:** In the second functionality, Discovery will scan our domain for WMI and SSH protocols. Because WMI is one of the protocols, the MID Server for this functionality must be installed on a Windows machine and must have criteria.

Create the second functionality using the following values:

Field	Input Value
Phase	Type a phase number of 1 in this field. All functionalities in this example use the same phase number, which launches a single Shazzam probe for all protocol scans. A single phase, when feasible, is the most efficient use of the Shazzam probe.
Functionality	Select *All except SNMP* from the list. This functionality defines the protocols for which scanning will occur. Because WMI is one of the protocols selected, we must use a Windows MID Server for this functionality.
MID Servers	We select a Windows MID Server from our domain. In this case we select *disco-win2003*. You can select multiple MID Servers if you want Discovery to load balance this functionality automatically.
Active	Make sure this check box is selected to enable this behavior.
Match criteria	Change the criteria to *All*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scan for All Except SNMP

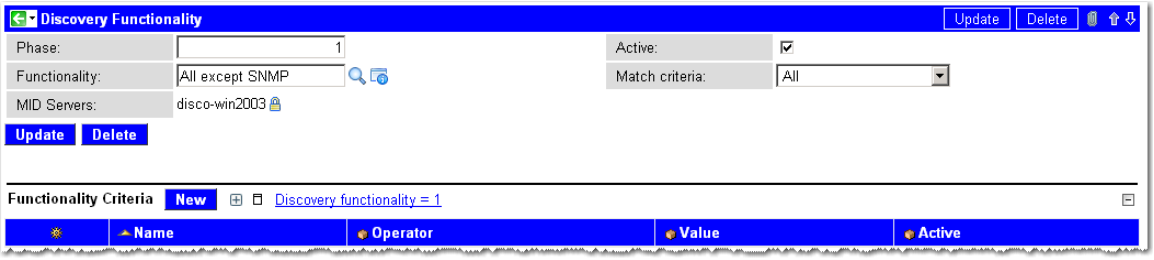

d. **Note:** Functionalities that scan for WMI require criteria to identify the domain and the MID Server. In our example, we will create two criteria for this functionality.

To create functionality criteria for a WMI scan, click **New** in the related list.

e. Enter the following values for the MID Server doing the WMI scanning on our domain:

Field	Input Value
Name	Create the following criteria: ▪ Enter `mid_server` to name the MID Server that will execute the WMI scans. ▪ Enter `win_domain` to name the Windows domain that Discovery will scan with the MID Server defined.
Operator	Select *equals* as the operator in this criteria.
Value	▪ For the *mid_server* value, enter the name of the MID Server that will scan our domain for Windows devices. ▪ For the *win_domain* value, enter the name of our Windows domain.
Active	Be sure to enable the criteria by selecting this check box (*`true`*).

The completed criteria appear as follows in the Discovery Functionality form for this behavior.

Functionality Criteria for a WMI Scan

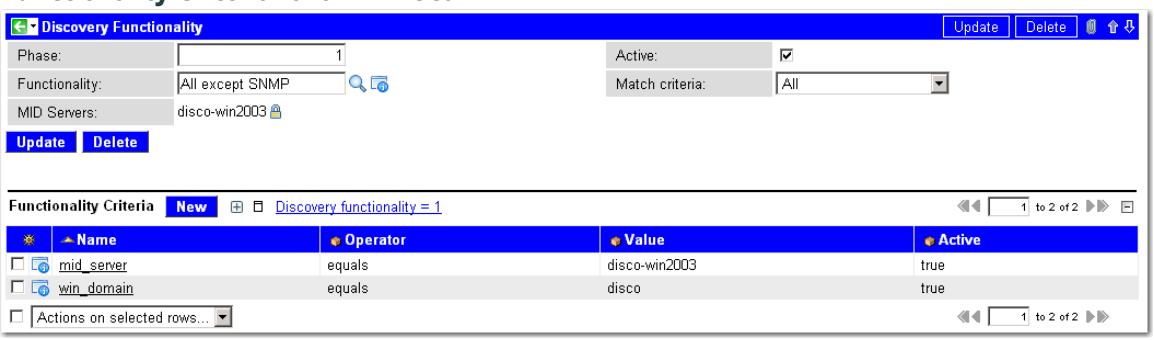

5. Create the schedule.

a. Create a Discovery Schedule record and name the behavior we just created.

b. Navigate to **Discovery** > **Discovery Schedules** and click **New** in the record list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Select a **Discover** mode of *Configuration items*.
This action displays the **Behavior** field.

d. Click the magnifier icon and select the behavior to use.
In our example, we select *Southwest-SNMP*.

e. Select the appropriate time to run Discovery for this location.

f. Click **Quick Ranges** and define the IP address ranges, networks, or lists to scan for this domain.

g. Save the record.

Discovery Schedule Record

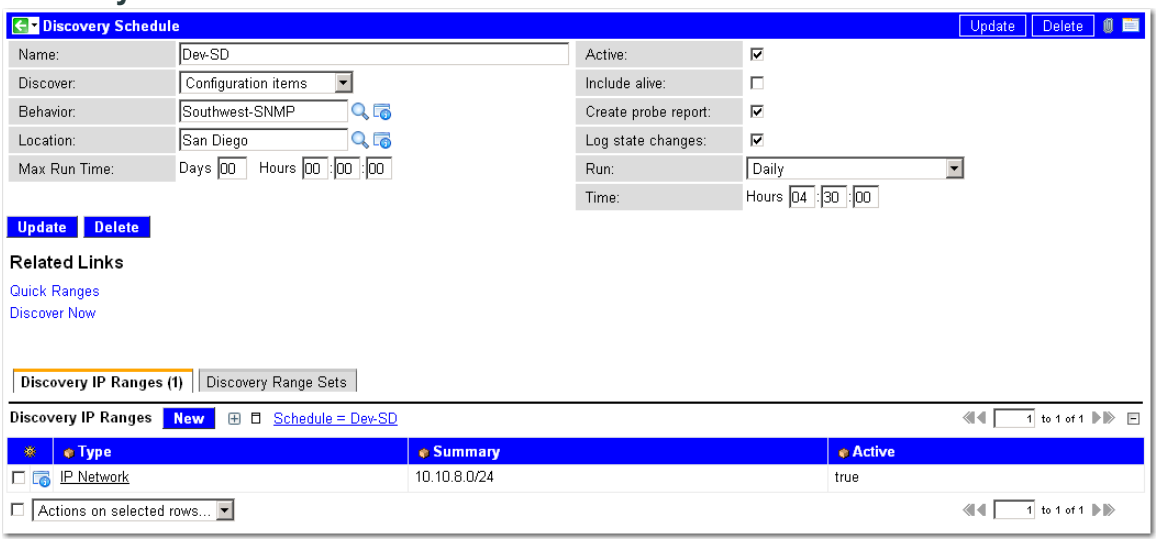

CI deletion strategies for pattern discovery

When you perform discovery with a pattern, you can choose what to do with configuration items (CIs) that are in the Configuration Management Database (CMDB) but Discovery can no longer find.

Note: Deletion strategies are available only with patterns, not probes/sensors.

When you use a pattern to perform discovery, you identify a CI Type, which is the main CI you're trying to find, such as an IIS web server or Tomcat server. Discovery can also find related CIs and populate the CMDB accordingly. For example, after Discovery finds an IIS server with the IIS pattern, it can also find these related CIs:

- The Windows machine that hosts the web server
- The configuration files for the IIS server
- The website that the IIS server hosts

The dependency map on the main CI also shows the related CIs, as illustrated by the following graphic:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

An example of an IIS server main CI and its related CIs

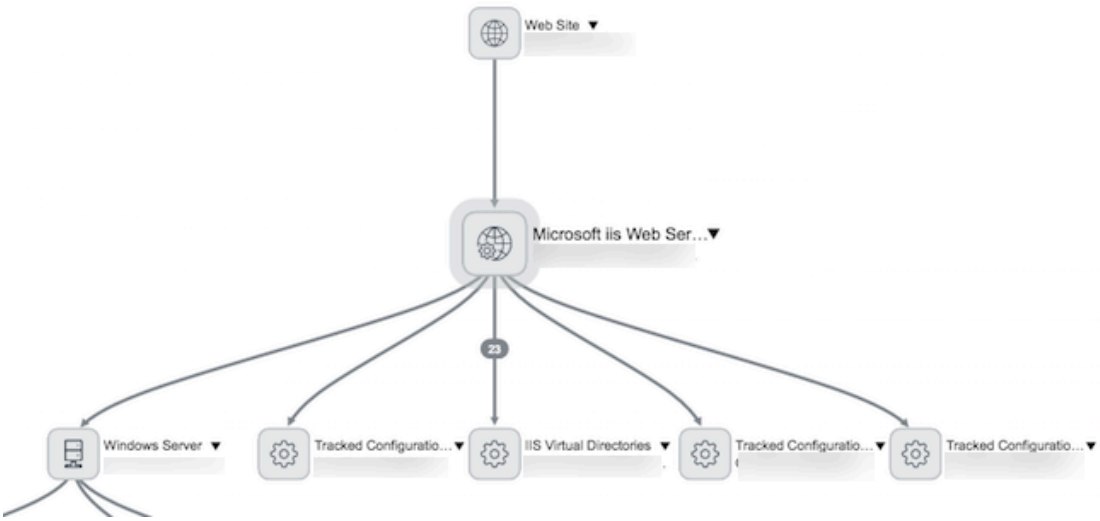

After you discover the main CI and its related CIs for the first time using a pattern, you must manually configure an entry in the Related CI Types [sa_ci_to_pattern] table for each relation between the main CI and a related CI type. These records are where you determine the deletion strategy for each related CI based on its type. Note that you cannot set a deletion strategy for the main CI for Infrastructure pattern type.

You can view a list of the relations as the following image illustrates:

Example CIs related to an IIS server

Related CI Types | CI type | Search

All > Pattern Name contains iis

CI type	Deletion Strategy	Pattern
Search	Search	*iis
cmdb_ci_microsoft_iis_web_server	Keep	IIS
cmdb_ci_config_file_tracked	Keep	IIS
cmdb_ci_iisdirectory	Keep	IIS
cmdb_ci_web_site	Keep	IIS

The deletion strategy on a related CI takes effect when both of these conditions occur:

- After you run Discovery with the pattern again.
- Discovery successfully finds the main CI but cannot find the related CI.

These deletion strategy options are available:

Value	Description
Keep	Preserves the configuration file CI record and makes no other changes to the record. This option is the default setting.
Delete	Deletes the configuration file CI record from the CMDB, and the relationship to the main CI.
Mark as absent	Marks the **Status** [install_status] field of the configuration file CI record as **Absent**, meaning that Discovery cannot find the tracked configuration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Value	Description
	file. This status does not instruct the system to delete the actual CI or the relationship. **Note:** There are two tables that do have their CIs deleted if marked as absent: cmdb_ci_network_adapter and cmdb_ci_ip_address.
Delete relations	Deletes only the CI relationships between the related CI and the main CI.
Mark as retired	Marks the **Status** [install_status] field of the configuration file CI record as **Retired**, meaning that Discovery no longer uses this configuration file. This status does not instruct the system to delete the actual CI or the relationship.

Note: Deletion strategies are not available for related entries, which are rules that are based on related CIs. For more information, see Related entries.

Set a deletion strategy

Set a deletion strategy when you want to take action on a related configuration item (CI) that Discovery can no longer find through pattern discovery. You can delete the main CI only when it is an Application or Cloud Resource pattern type.

Before you begin

- CIs: a main CI and related CIs that the discovery process found through pattern discovery
- Role required: cloud_admin, service_mapping_admin, discovery_admin, or admin. Starting from Visibility Content version 6.30.0, pd_admin and pd_mid are also supported.

About this task

You can set a deletion strategy for all discovered CIs, including configuration file CIs. However, the way you access the deletion strategy for tracked configuration files differs from the procedure in this topic. See Set the deletion strategy for tracked configuration files for instructions.

Procedure

1. Open the list of related CI types by entering `sa_ci_to_pattern.list` in the application filter.
2. Filter the list to find the CI type that is related to a main CI through pattern discovery.
   If you want to set a deletion strategy for a cloud database, search the list for *`cmdb_ci_cloud_database`* CI type.

   Notice that the *`cmdb_ci_cloud_database`* CI types have the **Is Main CI** value set to **true**. This indicates that you can set a deletion strategy for both AWS and Azure databases, even if they are the main CI that cloud discovery finds.

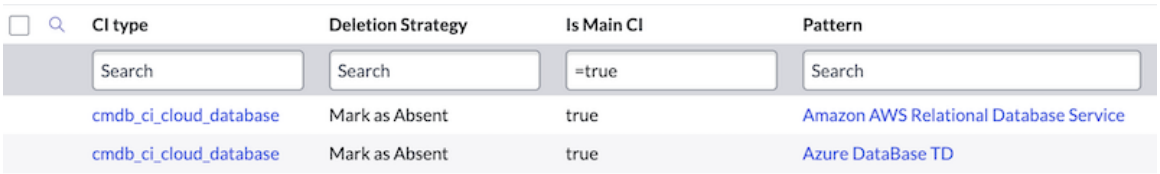

CI type	Deletion Strategy	Is Main CI	Pattern
Search	Search	=true	Search
cmdb_ci_cloud_database	Mark as Absent	true	Amazon AWS Relational Database Service
cmdb_ci_cloud_database	Mark as Absent	true	Azure DataBase TD

3. From the list view, double click the value in the field in the **Deletion Strategy** column.
4. Select a new value:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Value	Description
Keep	Preserves the configuration file CI record and makes no other changes to the record. This option is the default setting.
Delete	Deletes the configuration file CI record from the CMDB, and the relationship to the main CI.
Mark as absent	Marks the **Status** [install_status] field of the configuration file CI record as **Absent**, meaning that Discovery cannot find the tracked configuration file. This status does not instruct the system to delete the actual CI or the relationship. **Note:** There are two tables that do have their CIs deleted if marked as absent: cmdb_ci_network_adapter and cmdb_ci_ip_address.
Delete relations	Deletes only the CI relationships between the related CI and the main CI.
Mark as retired	Marks the **Status** [install_status] field of the configuration file CI record as **Retired**, meaning that Discovery no longer uses this configuration file. This status does not instruct the system to delete the actual CI or the relationship.

Configuration file tracking

The horizontal discovery process can find configuration files that belong to certain applications and add those configuration files to the CMDB. You can track the changes to these files by comparing them to previous versions.

Warning: Configuration files contain sensitive system information. To prevent unauthorized access, ensure that access control lists (ACL) are placed on the Tracked Configuration file table [cmdb_ci_config_file_tracked]. Only allow authorized users to view this table or uncheck the **Save Content** setting.

Components for configuration file tracking

CI type

All applications and hosts in your organization must have a corresponding configuration item (CI) type, which is necessary for discovering and processing applications and hosts correctly. In a base system, many CI types have configuration file paths defined for them. You can add new or modify existing definitions for tracking configuration files. See Modify tracking changes in configuration files for instructions.

Patterns

Configuration file tracking is available for patterns that discover applications. On the pattern, you can create tracked file definitions that specify the CI type to which the application CI belongs and the path of the configuration file. Specify as many tracked file definitions as needed. You can also specify whether you want to save the contents of configuration files so you can view and compare the contents of different versions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

Configuration file tracking is not available for discoveries performed by traditional probes and sensors.

The classifier that triggers the pattern must specify the Horizontal Pattern probe, which in turn, must specify the pattern. If you upgrade your instance to the current version, not all classifiers are configured to use patterns for discovery by default.

CMDB

All configuration files are saved as a CI in the Tracked Configuration file [cmdb_ci_config_file_tracked] table. If you enable the content to be saved, these CI records provide the contents of the configuration files, including previous versions. From the configuration file CI record, you can compare different versions.

Dependency maps and application service maps

Both dependency maps and service instance maps display tracked configuration files. The relationship between a configuration file and its host is a contains relationship. The application contains the configuration file.

For example, this IIS web server contains three tracked configuration files:

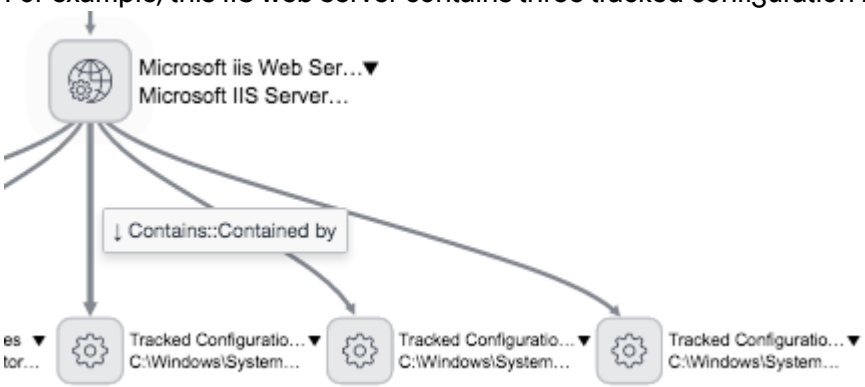

Sometimes you organize CI types as a main CI type and its related CI types. On a service instance map, Service Mapping shows changes to configuration files of related CIs for the main CIs in inclusions. In inclusions, the system treats applications hosted on a server as independent objects. For example, the Tomcat WAR CI appears separate from its host, the Tomcat CIs. In this case, Service Mapping shows changes to configuration files of Tomcat WAR when you select Tomcat. In addition, Service Mapping displays changes to configuration files of the hardware server hosting inclusions. In this example, it is a Linux server:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Map showing an inclusion with a host

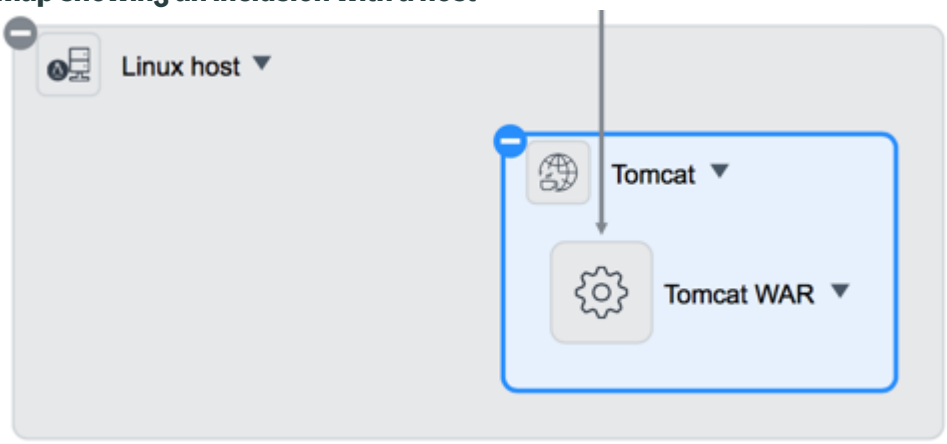

Deletion strategy

You can specify what you want to do with tracked configuration file CI records when discovery can no longer find them. You can keep the configuration file CI record, automatically delete it, delete only the CI relationships to it, or mark it absent.

Discovery patterns that support configuration file tracking by default

These patterns provide tracked file definitions by default:

Classifier	Pattern	CI Type	File path of tracked file
Apache Server	Apache On Unix Pattern Apache On Windows Pattern	Apache Web Server [cmdb_ci_apache_web_server]	`$config_file`
MySQL Server	MySQL server On Windows and Linux Pattern	MySQL Instance [cmdb_ci_db_mysql_instance]	`$config_file`
Microsoft IIS Server	IIS	Microsoft iis Web Server [cmdb_ci_microsoft_iis_web	`EVAL(javascript: var rtrn = '';var winDir = CTX.getCommandManager().shellCommand(" %WinDir%", false, null, null, CTX);rtrn = winDir.trim() + '\ \System32\\Inetsrv\\Config \\*.config';)`
		IIS Virtual Directory [cmdb_ci_iisdirectory]	`$install_directory + "\*.config"`
TIBCO BusinessWorks and EMS	ActiveMatrix BusinessWorks	ActiveMatrix Business Works [cmdb_ci_appl_tibco_matrix]	`$config_file`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier	Pattern	CI Type	File path of tracked file
	Enterprise Message Service	Tibco Enterprise Message Service [cmdb_ci_appl_tibco_message]	`$config_file`
Oracle	Oracle DB on Windows Pattern	Oracle Instance [cmdb_ci_db_ora_instance]	`$install_directory + "\network\admin\*.ora"`
		Oracle Instance [cmdb_ci_db_ora_instance]	`$install_directory + "\dbs\*.ora"`
	Oracle DB on Unix Pattern	Oracle Instance [cmdb_ci_db_ora_instance]	`$install_directory + "/dbs/*.ora"`
		Oracle Instance [cmdb_ci_db_ora_instance]	`$install_directory + "/network/admin/*.ora"`
Tomcat	Tomcat	Tomcat [cmdb_ci_app_server_tomca	`"/conf/server.xml"`
		Tomcat WAR [cmdb_ci_app_server_tomca	`$install_directory + "/WEB-INF/web.xml"`
WMB	WMB On Unix Pattern	IBM WebSphere Message Broker [cmdb_ci_appl_ibm_wmb]	`$install_directory + "/*/etc/config/*/*.prop"`
	WMB On Windows Pattern	IBM WebSphere Message Broker [cmdb_ci_appl_ibm_wmb]	`$install_directory + "\*\etc\config\*\*.prop"`
WMQ	WMQ On Windows Pattern	IBM MQ Manager (formerly IBM WebSphere MQ) [cmdb_ci_appl_ibm_wmq]	`$install_directory + "\*\config\*"`
	WMQ On Windows Pattern	IBM MQ Manager [cmdb_ci_appl_ibm_wmq]	`$install_directory + "/bin/*.sh"`

What to do

1. Enable configuration file tracking by setting the `glide.discovery.enable_file_tracking` Discovery property to true.

   You can also configure other properties to control the size and number of tracked configuration files, the time window during which changes to configuration files are tracked for a given version, and the number of changes allowed on a configuration file during that time window. See Discovery properties for details.

2. To prevent unauthorized access to the sensitive information from configuration files, perform the following steps:

   a. Ensure that access control rules (ACL) are placed on the Tracked Configuration file table [cmdb_ci_config_file_tracked]. Only allow authorized users to view this table or uncheck the **Save Content** setting.

   b. Ensuring that the tracked_file_reader role that controls access to configuration file data is assigned to correct users and user groups. By default, the itil role contains the tracked_file_reader role giving all users with the itil role access to the configuration file information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Verify that the Horizontal Discovery probe is active on the classifier for the software that you want to discovery. If not, you can enable it, specify the pattern, and then disable the other probes. See Add the Horizontal Pattern probe to a classifier for instructions.
4. Add or modify tracked file definitions to change the CI type or file path.
5. Set the tracked files deletion strategy to specify what you want to do with tracked configuration file CI records when pattern discovery can no longer find them.
6. Run horizontal discovery on the hosts that are running the applications you want to discover with patterns, open the application CI record, and check the Tracked Configuration Files related list.
7. Compare two versions of tracked CI configuration files to see the actual changes made to them.

Modify tracking changes in configuration files

Configure the system to collect information about changes in configuration files belonging to a configuration item (CI). Service Mapping uses this information to notify users that CI configuration files changed and to view actual changes to configuration files directly in the service instance maps.

Before you begin

To enable tracking configuration files for a CI for which the system does track files, configure classification for the CI type to which this CI belongs:

- For CI types representing applications, create a Discovery process classification for the relevant pattern. Add the Horizontal Pattern probe in the **Probe** column, and then specify your pattern in the **Pattern** column.
- For CI types representing SNMP devices, perform configuration as described in Create a Discovery CI classification.

Note: There is no need to create CI classifications for hosts because these classifications are included in the base system.

Role required: cloud_admin, service_mapping_admin, or admin

About this task

The system tracks configuration files as part of horizontal discovery process of a CI to which these files belong. Configuration files contain CI settings and parameters. Service Mapping uses this information to notify users that CI configuration files changed and to view actual changes to configuration files directly in the service instance maps. To learn about how the system tracks configuration files, stores and updates information about these files, see Configuration file tracking.

The system tracks configuration files for CI types with process classifiers that trigger patterns. The patterns contain tracked file definitions, which specify the CI type that the application belongs to and the path of the configuration file. Several patterns, such as the IIS pattern, provide tracked file definitions by default. You can use these default definitions or create your own for any pattern a process classifier triggers.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Select the pattern for the relevant CI.

Important: Enabling configuration file tracking does not modify the contents of a pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select the **Tracked Files** tab.
   One or more tracked file definitions appear by default for Discovery patterns that support configuration file tracking.

   If you don't see the Tracked Files tab, verify that the pattern is specified process classier. See Add the Horizontal Pattern probe to a classifier for instructions.

4. Select **Edit** to modify an existing definition or select **New** to create one.
5. Fill out or modify the form fields (see table):

Field	Description
CI Type	Select either the primary or related CI type to which the configuration file belongs.
File Path	Define the path of the configuration file to track. You can use pattern variables and wildcards. For example, for the following path: <the folder with the "install_directory" parameter>/conf/<any xml file>, enter $install_directory+"/conf/*.xml"
Save Content	Select this check box to make this file available for viewing and comparing directly the ServiceNow instance.
Active	Select this check box to enable tracking of this configuration file.

6. Select **Done**.
7. If necessary, set the tracked files deletion strategy as described in Set the deletion strategy for tracked configuration files

What to do next

Run discovery on the hosts that are running the applications you want to discover with patterns, open the application CI record, and check the Tracked Configuration Files related list.
If you're using Service Mapping, verify that it started collecting information about changes to configuration files:

1. Navigate to the service instance map containing the CIs.
2. On the Changes tab, check that the list includes a record of the change you made. For example, that the configuration file was added, modified or deleted.

Compare versions of CI configuration files

You can compare two versions of tracked CI configuration files to see the actual changes made to them.

Before you begin

If Discovery and Service Mapping don’t track a configuration file you want to track, perform advanced configuration as described in Modify tracking changes in configuration files.

Role required: one of the following roles.

- admin
- service_mapping_admin (starting from Visibility Content version 6.29.0)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Discovery and Service Mapping check configuration files for changes during horizontal discovery using patterns. If a configuration file is modified, the system saves the version of the file. You can use the timestamp to differentiate between versions. For Discovery, you can compare any two versions from the list of saved versions. For Service Mapping, you can compare the latest and previous versions.

Changes to configuration files are associated with CIs to which these files belong. Maps show configuration file changes as changes to related CIs.

Procedure

1. Compare two versions of the CI configuration file from Service Mapping or Discovery.

Product	Steps
Service Mapping	**a.** Navigate to **All** > **Service Mapping** > **Services** > **Business Services**. **b.** Select **View Map** next to the required business service. **c.** On the **Changes** tab, locate the change record for the modified CI configuration file. **Important:** The name of the change must be **File Modified**. 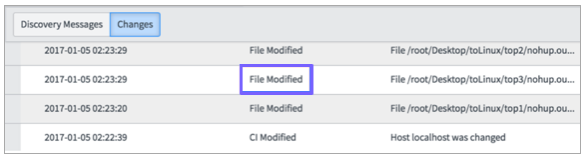 **d.** Select and hold (or right-click) the change record for the modified CI configuration file, and select **View Change in Tracked File Content**. The Tracked Configuration File Comparison tab opens, displaying the most recent and the previous versions of the configuration file.
Discovery	**a.** Navigate to the configuration files list in **Tracked Configuration Files**. ▪ Navigate to **All** > **Discovery** > **Tracked Configuration Files** and sort or filter the list by criteria such as **Related CI** or **Host**. ▪ Open a CI record from the Configuration module for the CI type to which the configuration file belongs. For example, to find a Tomcat server, navigate to **All** > **Configuration** > **Applications Servers** > **Tomcat**, then select the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Product	Steps
	specific Tomcat server. The list of configu ration files appear in the Tracked Configu ration files related list. **b.** Select the configuration file from the list. **c.** In the **Content Versions** related list, select one of the versions in the **Update time** relat ed list. **d.** Select **Compare**. The Tracked Configuration File Comparison tab opens, displaying the configuration file versions you selected.

2. Review actual changes.

   Highlight colors indicate the type of change:

   - Purple — Updated line
   - Pink — New line
   - Gray — Deleted line

3. Navigate between the changes using the Next difference icon and the Previous difference icon .

Related topics

Tracked Configuration file form

Set the deletion strategy for tracked configuration files

You can specify what you want to do with tracked configuration file CI records when pattern discovery can no longer find them.

Before you begin

A pattern that specifies tracked configuration files.

Role required: discovery_admin, service_mapping_admin, or admin. Starting from Visibility Content version 6.30.0, pd_admin and pd_mid are also supported.

About this task

You can set a deletion strategy for all discovered CIs, not just configuration file CIs. However the way you access the deletion strategy for standard CIs differs from configuration file CIs. See Set a deletion strategy for instructions.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns** and open the desired pattern.
2. Click the **Tracked Files** tab.
3. Under **Related Links**, click **Edit deletion strategy for the configuration files of this pattern**. The Related CI Types list appears showing you filtered records in the Related CI Type [sa_ci_to_pattern] table. The filter shows you only records with the CI type of Tracked Configuration files [cmdb_ci_config_file_tracked], which is the table where the system saves configuration file CIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. From the list view, double-click the value in the field in the **Deletion Strategy** column for the tracked configuration file.
5. Select a new value:

Value	Description
Keep	Preserves the configuration file CI record and makes no other changes to the record. This option is the default setting.
Delete	Deletes the configuration file CI record from the CMDB, and the relationship to the main CI.
Mark as absent	Marks the **Status** [install_status] field of the configuration file CI record as **Absent**, meaning that Discovery cannot find the tracked configuration file. This status does not instruct the system to delete the actual CI or the relationship. **Note:** There are two tables that do have their CIs deleted if marked as absent: cmdb_ci_network_adapter and cmdb_ci_ip_address.
Delete relations	Deletes only the CI relationships between the related CI and the main CI.
Mark as retired	Marks the **Status** [install_status] field of the configuration file CI record as **Retired**, meaning that Discovery no longer uses this configuration file. This status does not instruct the system to delete the actual CI or the relationship.

Related topics

CI deletion strategies for pattern discovery

Set a deletion strategy

Change the source name of Discovery results

You have the option of changing the source name of discovery results. This might be desirable if Discovery is running on your network together with another discovery product, and you want to use customized identifiers.

Before you begin

Role required: discovery_admin or admin

About this task

The Source [sys_object_source] table stores information identifying the source of a discovery, the ID of that source, and the date/time of the last scan.

To view this information, configure a CI form and add the **Sources** related list. This table is populated automatically when the Discovery plugin is enabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

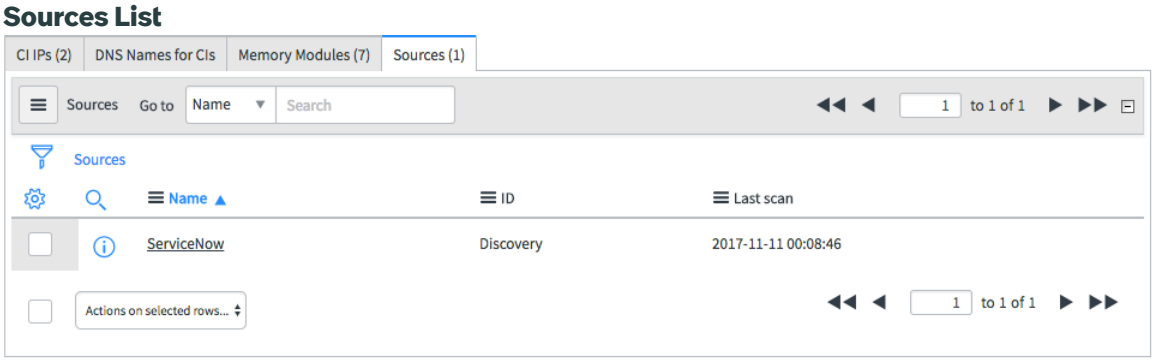

To migrate your Discovery Source to ServiceNow:

Procedure

1. Update the *glide.discovery.source_name* system property to have a value of ServiceNow.
2. Edit the dictionary entry for the Configuration Item [cmdb_ci] table.
3. Update the choice list for the *discovery_source* column to only have the value ServiceNow for a discovery product.
4. Open a list for the Configuration Item [cmdb_ci] table and filter on a *discovery source* field value you want to change.
5. Use the *Update All* option to change the value to ServiceNow.
6. Open a list for the Source [sys_object_source] table and filter on a name field value you want to change.
7. Use the *Update All* option to change the value to ServiceNow.

Discovery API plugin

The Discovery API plugin provides APIs for scoped applications and is loaded when the Discovery plugin is activated.

Details about these Discovery API methods are available on the ServiceNow® Developer Site. They are listed here by class.

Note: Java API methods are not customizable.

DiscoveryAPI - Scoped

The methods in this class launch a quick Discovery of a single IPv4 address and return summaries of previously launched Discovery statuses for a single CI or for all scanned CIs. A MID Server is selected automatically, based on the IP address provided or the application specified.

- *discoverIpAddress()*: Discovers a single IPv4 address.
- *reportCiIpAddressStatus()*: Returns a summary of a configuration item's Discovery status given the specific status sys_id and IPv4 address.
- *reportCiStatus()*: Returns a summary of a CI Discovery status given a specific Discovery Status sys_id.

ReportCiStatusOutputJS

The methods are getters that return specific object properties for the DiscoveryAPI *reportCiIpAddressStatus* method and then convert the information into a JSON string.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- *getCiOperationStatus()*: Used to return the state of the scanned CI.
- *getCmdbCI()*: Used to return the value in the cmdb_ci field from the discovery_device_history table for the CI being scanned.
- *getDiscoveryState()*: Used to return the value from the **State** field in the Discovery Status [discovery_status] table.
- *getIpAddress()*: Used to return the value from the source field in the discovery_device_history table for the CI being scanned.
- *getIssues()*: Used to return the value from the issues field in the discovery_device_history table for the CI being scanned.
- *getIssuesLink()*: Used to return the value from the issues_link field in the discovery_device_history table for the CI being scanned.
- *toJson()*: Used to serialized the *ReportCiStatusOutputJS* object.

Running process filters

Filter any processes, using combinations of name and key parameters, from coming back to the instance. Out-of-the box, you can filter OS processes. Filtering processes helps reduce the number of records created, updated, or deleted in the database which can improve overall database performance.

Before you begin

Role required: discovery_admin or admin

About this task

To add processes to the filter, you must setup configurations using the Running Process filter record [running_process_filter]. These records make sure the filtered processes do not populate the Running process [cmdb_running_process] table.

Procedure

1. Navigate to running_process_filter.list.
2. Select **New** to create a record.
   The form and descriptions are explained below.

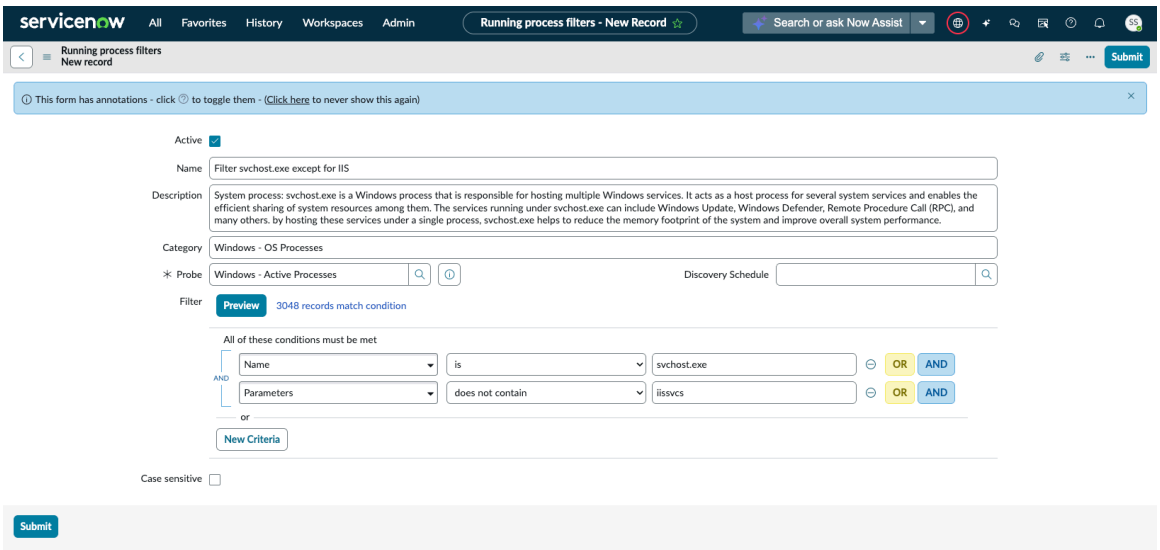

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Active	Check the box to include or exclude.
Name	Name for the record.
Description	Short description of the record.
Category	A name for the configuration type. Running process filters can be grouped however you want so that it's easier to identify and choose in the table.
Probe	Mandatory field. Which probe will use this filter. Often Linux or Windows active processes or ADM enhanced. When the probe executes, the instance sends the running_process_filter records as a probe parameter. The probe then uses this to filter the results before sending the information back to the instance.
Discovery Schedule	You can apply the filter on a specific Discovery schedule. If you leave this field, it applies to all schedules, even Quick Discovery.
Filter	Any process matching the filter will not be added to the Running Processes [cmdb_running_process] table. For example, if you want to filter all svchost.exe processes except the svchost.exe process for IIS classification, the filter definition would be: ○ name is svchost.exe ○ parameters does not contain iissvcs For information on setting filter conditions see: Condition builder
Case sensitive	Applies to the Filter field. If case sensitive is checked, the filter may pull more or less information than you want. Default is off.
OOB filters	There are Linux and Windows filters to choose from. See Out-of-the-box process filters for more details.

3. Enter the details and select **Submit**.

Out-of-the-box process filters

There are a variety of Linux and Windows out-of-the-box running process filters that are turned on and used for filtering processes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters

Name	Description
dhclient	dhclient is a system process in Linux that is responsible for dynamically configuring network interfaces using the Dynamic Host Configuration Protocol (DHCP). It communicates with a DHCP server to obtain network configuration information, such as IP address, subnet mask, default gateway, and DNS servers. dhclient is commonly used in environments where network addresses are assigned dynamically, such as in home and small business networks.
crond	crond is a system service in Linux that provides support for the cron scheduling utility. It is responsible for running scheduled jobs at specified times and intervals. crond is commonly used for automating system maintenance tasks, such as backups and log rotation, as well as for running user-defined scripts and applications at specific times.
sleep	The sleep process is a command that is used to pause the execution of a shell script or a command for a specified amount of time. It is a simple way to delay the execution of a script or command for a specific duration.
agetty	agetty is a system service in Linux that manages virtual console login sessions. It provides a terminal interface for users to log in to the system and access the command line interface. agetty is responsible for initializing the terminal, displaying the login prompt, and authenticating the user's credentials.
dbus-daemon	dbus-daemon is a system service in Linux that implements the D-Bus message bus system. It is responsible for enabling communication between applications and system services on the same host or on a network. dbus-daemon manages the message bus and dispatches messages between applications and services, allowing them to exchange data and request services from each other.
rsyslogd	rsyslogd is a system service in Linux that manages system logs and provides centralized logging capabilities. Administrators can configure and manage log messages generated by various applications and services on a system or across a network. rsyslogd provides advanced features such as log filtering, log forwarding, and log rotation. It supports a variety of output formats and destinations, including files, databases, and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters (continued)

Name	Description
	remote syslog servers. rsyslogd is commonly used in server environments and can be configured using a configuration file or a graphical user interface.
auditd	auditd is a system service in Linux that provides an auditing framework for tracking security-relevant events on the system. It collects and logs security events generated by the kernel and user space applications, which can be used for forensic analysis, compliance monitoring, and intrusion detection. auditd can be configured to generate alerts or take other actions based on predefined rules or policies.
irqbalance	irqbalance is a system service in Linux that balances interrupt requests (IRQs) among processors in a multicore system to improve performance and reduce latency. It detects and monitors IRQ activity and redistributes IRQs to the least busy processor, which can help to reduce processing bottlenecks and improve system responsiveness. irqbalance is commonly used in high-performance computing environments and other systems with high levels of interrupt activity.
systemd-journald	systemd-journald is a system service in Linux that provides a centralized logging system for system and application logs. It is designed to collect and store log data in a structured and efficient manner and allows for advanced querying and filtering of log data. systemd-journald provides a unified interface for accessing and managing logs across different Linux distributions and can be configured using a configuration file or a command-line interface. It is commonly used in server and desktop environments to monitor and troubleshoot system and application issues.
abrtd	abrtd is a system process that runs on Linux systems and is responsible for detecting and reporting application crashes to the system administrator. The process collects information about the crashed application, such as core dumps, and stores them in a report.
atd	atd is a system service in Linux that allows for the scheduling of one-time jobs to be executed at a specific time and date. It is useful for running jobs that need to be executed at a specific time or for automating tasks that need to run once in the future.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters (continued)

Name	Description
lvmetad	lvmetad is a system service in Linux that provides metadata caching for Logical Volume Management (LVM). It caches metadata information about LVM volumes and devices, which can improve the performance of LVM operations and reduce the time needed to scan and initialize LVM devices. lvmetad is commonly used in systems that rely heavily on LVM, such as virtualization environments and storage servers.
polkitd	polkitd is a system service in Linux that provides an authorization framework for controlling access to privileged actions and resources. Non-privileged users can perform specific tasks or access certain resources that require elevated privileges, such as installing software or configuring network settings. Users authenticate using their password or other means. polkitd provides a policy engine that can be configured by administrators to define fine-grained access controls for specific users, groups, or applications.
acpid	acpid is a system service in Linux that listens for ACPI (Advanced Configuration and Power Interface) events generated by the hardware and notifies the operating system. It is responsible for managing power events, such as power button presses, and executing user-defined actions in response to them.
anacron	anacron is a system service in Linux that schedules periodic jobs, similar to cron. It ensures that missed jobs are executed when the system next becomes available. It is useful for running periodic tasks that may be delayed due to system downtime or user inactivity.
(sd-pam)	sd-pam is a system service in Linux that provides support for the Pluggable Authentication Modules (PAM) framework. It is responsible for handling user authentication and authorization requests from various system services and applications.
NetworkManager	NetworkManager is a system service in Linux that manages network connectivity for desktops, laptops, and other devices. It provides a unified interface for configuring and managing wired, wireless, mobile broadband, and other network connections, handling IP addressing, DNS configuration, and other network settings. NetworkManager is commonly used in desktop environments

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters (continued)

Name	Description
	and graphical user interfaces and can be configured using a graphical user interface or command-line interface.
hcid	hcid is a system service in Linux that provides support for Bluetooth connectivity. It is responsible for managing Bluetooth devices, connections, and services on the system. hcid provides an interface for applications and system services to interact with Bluetooth devices. It configures and manages Bluetooth-related settings, such as device pairing, authentication, and authorization.
hidd	hidd is a system process in Linux that provides support for Bluetooth human interface devices (HID), such as keyboards and mice. It is responsible for managing Bluetooth connections and providing a common interface for HID devices to communicate with the system. hidd is commonly used in desktop environments to enable wireless input devices to be used with the system.
iscsid	iscsid is a system process in Linux that provides support for the iSCSI protocol, which is used to access storage devices over a network. It is responsible for discovering iSCSI targets, establishing and managing iSCSI sessions, and providing access to iSCSI devices as if they were locally attached storage. iscsid is commonly used in virtualization environments and other systems that require remote access to storage resources.
klogd	klogd is a system process in Linux that is responsible for logging kernel messages to the system log. It monitors the kernel log buffer for new messages and writes them to the system log file or sends them to a remote logging server. klogd is commonly used for troubleshooting kernel issues and monitoring system activity, as well as for collecting diagnostic information for system administrators and developers.
mingetty tty1	mingetty tty is a system process in Linux that provides a terminal login prompt on a virtual console. It is responsible for initializing the terminal device, displaying the login prompt, and accepting user login credentials. mingetty tty is used for each virtual console in a Linux system. It can be used for system

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters (continued)

Name	Description
	administration and debugging tasks and for running console-based applications.
ModemManager	ModemManager is a system service in Linux that manages mobile broadband (GSM/CDMA/UMTS/LTE) and other modem connections. It provides a unified interface for configuring and managing modem devices, handling connections to the internet or other networks, and monitoring modem activity. ModemManager is commonly used on laptops, tablets, and other mobile devices that require wireless network connectivity.
pcscd	pcscd is a system service in Linux that provides support for smart card readers and smart cards. It enables applications to access smart cards and their content using the PC/SC (Personal Computer/Smart Card) interface standard. pcscd provides a unified interface for accessing various types of smart card readers. Applications such as web browsers, email clients, and authentication systems use it for smart card support.
power-profiles-daemon	power-profiles-daemon is a system service in Linux that manages power profiles and power management settings for laptops and other mobile devices. It provides a unified interface for configuring and managing power-saving features, such as screen brightness, CPU frequency scaling, and suspend modes. These settings are based on user preferences and system usage patterns. power-profiles-daemon is commonly used in desktop environments and graphical user interfaces and can be configured using a graphical user interface or command-line interface.
pulseaudio	pulseaudio is a system service in Linux that manages audio devices, streams, and applications. It provides a unified audio interface for both input and output audio streams and supports advanced features such as volume control, mixing, and audio effects. pulseaudio is commonly used in desktop environments and graphical user interfaces and can be configured using a graphical user interface or command-line interface. It supports a variety of audio hardware and software configurations, including multiple audio devices and network audio streaming.
chronyd	chronyd is a system service in Linux that is responsible for synchronizing the system

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters (continued)

Name	Description
	clock with a reference time source. It implements the Network Time Protocol (NTP) and can be used to synchronize the system clock with a variety of time sources, including GPS clocks and atomic clocks. chronyd is commonly used in environments where accurate time synchronization is important, such as in financial trading or scientific research.
rtkit-daemon	rtkit-daemon is a system service in Linux that provides real-time scheduling and priority management for audio and other time-sensitive applications. It ensures that these applications get access to the CPU and other system resources they need to function properly. This avoids audio glitches, latency, or other performance issues. rtkit-daemon uses the Linux real-time capabilities and provides a simple interface for applications to request real-time scheduling and priority. It is commonly used in desktop environments and audio applications and can be configured using a configuration file or a graphical user interface.
sssd	sssd is a system service in Linux that provides centralized authentication, identity, and access management for a variety of services and applications. It is designed to simplify the management of user accounts and credentials in large, multi-domain environments and provides a unified interface for authentication, authorization, and identity resolution. sssd supports a variety of identity sources, including local files, LDAP directories, and Active Directory domains, and provides advanced features such as offline authentication, caching, and failover. It is commonly used in enterprise environments and can be configured using a configuration file or a graphical user interface.
sssd_kcm	sssd_kcm is a system service in Linux that provides a Kerberos Credential Cache Manager for the System Security Services Daemon (SSSD). It improves the security and performance of Kerberos-based authentication and authorization by caching Kerberos credentials and tickets locally. It provides a unified interface for accessing them. sssd_kcm is commonly used in enterprise environments that rely on Kerberos for centralized authentication and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters (continued)

Name	Description
	access control. It can be configured using a configuration file or a graphical user interface.
systemd-hostnamed	systemd-hostnamed is a system service in Linux that provides hostname and DNS resolution configuration. It is designed to simplify the configuration of network-related settings such as hostname, domain name, and DNS servers. systemd-hostnamed provides a unified interface for configuring these settings across different Linux distributions and can be configured using a configuration file or a command-line interface. It is commonly used in server and desktop environments to configure network settings during system startup or runtime.
systemd-timesyncd	systemd-timesyncd is a system service in Linux that provides time synchronization for the system clock. It synchronizes the system clock with a remote time server over the network and can be configured to use multiple time sources for redundancy. systemd-timesyncd is commonly used in server and desktop environments to ensure accurate timekeeping for various system and application functions.
hald	hald is a system service in Linux that detects and handles hardware events. These events include the insertion and removal of storage devices and other hardware components. It provides a common interface for applications and system services to interact with the hardware. It is used to configure and manage system resources, such as input devices and storage volumes. hald is commonly used in desktop environments to provide plug-and-play functionality for hardware devices.
xfs	xfs is a high-performance, journaling file system for Linux that is optimized for scalability and large files. It was originally developed by Silicon Graphics, Inc. for use in their IRIX operating system and was later ported to Linux. xfs supports features such as snapshots, online resizing, and advanced file permissions. It is commonly used in high-performance computing, media and entertainment, and other industries where handling large files and data sets is critical.
audispd	audispd is a system service in Linux that is responsible for processing audit events generated by the kernel's auditing subsystem.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux out-of-the-box filters (continued)

Name	Description
	It is part of the Audit framework and is used to provide secure logging of system events for auditing and forensic purposes. audispd can be configured to send audit events to various targets, such as syslog or a centralized log server.
automount	automount is a system service in Linux that provides automatic mounting of file systems when they are accessed by users or applications. It allows for on-demand mounting of file systems, reducing the amount of time and system resources spent on unnecessary file system mounts. Automount is particularly useful in environments with large numbers of networked file systems or removable storage devices.
avahi-daemon: chroot helper	avahi-daemon is a system service in Linux that implements the zeroconf protocol for service discovery on local networks. The chroot helper is a component of avahi-daemon that sets up a chroot environment for the service. This increases security by limiting the access that avahi-daemon has to the system.

Windows out-of-the-box filters

Name	Description
System Idle Process	In Windows operating system, "System Idle Process" is not actually a process, but rather a system idle task that runs when the CPU has nothing else to do. It represents the percentage of time the CPU is idle and available to handle other processes. So, it is an important system component that shows the current CPU usage when no other processes are running.
svchost.exe	svchost.exe is a Windows process that is responsible for hosting multiple Windows services. It acts as a host process for several system services and enables the efficient sharing of system resources among them. The services running under svchost.exe can include Windows Update, Windows Defender, Remote Procedure Call (RPC), and many others. By hosting these services under a single process, svchost.exe helps to reduce the memory footprint of the system and improve overall system performance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows out-of-the-box filters (continued)

Name	Description
wmiprvse.exe	wmiPrvSE.exe is a Windows system process that provides management information and control through the Windows Management Instrumentation (WMI) framework. It runs as a service in the background and allows applications and scripts to access system management information, perform administrative tasks, and monitor system performance. WmiPrvSE.exe can be run in both local and remote modes and can be used to manage a wide range of system resources, including hardware, software, and network components. It is an essential component of the Windows operating system and is typically used by system administrators and IT professionals.
conhost.exe	conhost.exe is a Windows process that serves as a host for console applications. It was introduced in Windows 7 and is responsible for managing the command prompt windows and other console-based applications. conhost.exe provides improved compatibility and security features compared to its predecessor, the csrss.exe process, which was used for the same purpose in older versions of Windows.
WUDFHost.exe	WUDFHost.exe is a Windows User-Mode Driver Framework (UMDF) Host Process that manages user-mode drivers running in the context of the Windows Driver Foundation (WDF). It provides a framework for developing and executing drivers in user mode rather than kernel mode, improving stability and reliability of the system. WUDFHost.exe is used for Plug and Play devices and other driver-based hardware components. It runs in the background and typically consumes low system resources.
LogonUI.exe	LogonUI.exe is a Windows system process that manages the Windows logon screen and user authentication. When a user attempts to log in to a Windows computer, LogonUI.exe displays the logon screen and prompts the user to enter their login credentials. Once the user's credentials are entered, LogonUI.exe verifies them and grants the user access to the system. LogonUI.exe is an essential component of the Windows operating system and should not be terminated or disabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows out-of-the-box filters (continued)

Name	Description
csrss.exe	The csrss.exe process is a critical component of the Windows operating system. It is responsible for handling certain functions related to the graphical user interface (GUI), such as managing windows and consoles. The process runs in a separate session from the user's session to ensure stability and security of the operating system. It is an essential system process and should not be terminated unless there is a specific and valid reason to do so.
dllhost.exe	dllhost.exe is a Windows system process responsible for running and managing COM (Component Object Model) components. COM components are reusable software modules that can be accessed by other programs, and they are often used to facilitate inter-process communication between different applications. The dllhost.exe process hosts these COM components and ensures their proper execution and management. It is commonly used by a variety of Windows applications and services.
Fontdrvhost.exe	Fontdrvhost.exe is a Windows system process that manages the installation, removal, and enumeration of font files. It is responsible for loading and rendering fonts in applications that use the Windows Font Cache Service. Fontdrvhost.exe runs in the background and does not have a visible interface. It is important for maintaining the system's font library and ensuring that fonts are available to applications when needed.
wlanext.exe	wlanext.exe is a Windows system process that provides support for wireless LAN (WLAN) configurations and operations. It manages the wireless network adapters and implements the WLAN AutoConfig service functionality. This service enables automatic configuration of wireless networks and supports the Wi-Fi Protected Setup (WPS) protocol. The process runs in the background and can be managed through the Services console or the Command Prompt.
winlogon.exe	winlogon.exe is a system process in Windows that manages the login and logout of user sessions. It authenticates users and loads their user profile upon login. It also closes down any open applications and saves user data when a user logs off or shuts down the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows out-of-the-box filters (continued)

Name	Description
	system. Winlogon.exe runs as a service in the background and is essential for the proper functioning of the Windows operating system.
unsecapp.exe	unsecapp.exe is a Windows system process that is related to the Distributed Component Object Model (DCOM). It runs in the background and acts as a mediator between a client and a server application that communicate using DCOM. When a client sends a request to a server using DCOM, unsecapp.exe handles the security aspects of the communication. It ensures that the request is authorized. It is a legitimate system process and is typically located in the C:\Windows \System32 folder.
w3wp.exe	w3wp.exe is a Windows process that stands for Windows Process Activation Service Worker Process. It is used by the Internet Information Services (IIS) web server to host and manage ASP.NET web applications. w3wp.exe allows multiple web applications to run simultaneously on the same server, and it provides a secure, isolated environment for each application. It is a critical component of the IIS server and is often monitored closely to ensure the stability and performance of web applications.
dwm.exe	dwm.exe is a Windows system process that stands for Desktop Windows Manager. It is responsible for managing graphical effects, such as transparent windows and window animations, in the Windows operating system. The process is also responsible for rendering the desktop, managing windows, and providing visual effects such as Aero Glass. It was introduced with Windows Vista and has since been included in all subsequent versions of Windows.
lsass.exe	lsass.exe is a critical system process in Windows that handles authentication and security-related functions such as local and remote logon, password changes, and security policy enforcement. It is a core component of the Windows operating system and is required for the proper functioning of the system.
System	System is not a Windows process in the traditional sense. Instead, it refers to a collection of processes and services that are critical to the functioning of the operating system. These processes and services are

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows out-of-the-box filters (continued)

Name	Description
	responsible for managing hardware resources, handling system events, and providing essential functions like memory management, process scheduling, and security. In other words, System is the core of the Windows operating system.
smss.exe	smss.exe, or Session Manager Subsystem, is a critical system process in Windows that is responsible for launching user-mode applications during system startup. It also creates system threads, initializes system-wide data structures, and handles system shutdown and reboot requests. It is one of the first processes to start during the Windows boot process.
services.exe	services.exe is a system process in Windows responsible for managing system services, including starting and stopping them, and controlling their configuration settings. It runs in the background and provides essential functionality for the operating system.
wininit.exe	wininit.exe is a system process in Windows that runs during the startup process of the operating system. Its main function is to initialize critical system services and processes that are required for the system to function properly. It is responsible for running the Windows startup sequence, including launching the Session Manager (smss.exe) process.
spoolsv.exe	spoolsv.exe is a system service process in Windows that manages the printing process. It is responsible for queuing and spooling print jobs, communicating with the printer, and ensuring that each job is printed correctly. The process runs in the background and is critical for the proper functioning of the printing subsystem in Windows. If the spoolsv.exe process encounters errors, it can result in printing problems and errors.
taskhostw.exe	taskhostw.exe is a Windows system process that serves as a host for various background tasks that run on the system. It manages and executes tasks such as Windows updates, maintenance tasks, and user-defined tasks. It is a critical system process and terminating it can cause system instability.
SearchProtocolHost.exe	SearchProtocolHost.exe is a Windows process that is responsible for indexing and performing

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows out-of-the-box filters (continued)

Name	Description
	search operations on the system. It is a part of the Windows Search service, which allows users to quickly search for files and folders on their computer. The SearchProtocolHost.exe process runs in the background and indexes files and folders based on user-defined search criteria. When a search query is performed, the process is activated and retrieves the relevant search results.
SearchIndexer.exe	SearchIndexer.exe is a Windows system process that indexes files and folders on local and networked drives to enable faster searching. It runs in the background and continuously indexes new or changed files to maintain an up-to-date index database. The process is essential for Windows Search functionality and can be found on most Windows operating systems.
Explorer.exe	Explorer.exe is a core process in Windows operating systems responsible for providing the user interface, desktop, and file management functions. It is commonly referred to as the Windows Explorer or File Explorer, and allows users to navigate and manage files, folders, and applications on their computer.
SearchFilterHost.exe	SearchFilterHost.exe is a Windows system process that is responsible for indexing and searching files on the computer. It is a component of the Microsoft Windows Search service, which allows users to quickly search for files and documents on their computer. The process runs in the background and can use a significant amount of system resources while indexing files.
SCNotification.exe	SCNotification.exe is a Windows system process that is responsible for showing system notifications, such as alerts and messages, to the user. It is part of the Windows Shell Experience Host (ShellExperienceHost.exe) process, which manages the user interface for the Start menu, taskbar, and Action Center in Windows 10.
sihost.exe	sihost.exe is a system process in Windows that is responsible for launching the Start menu and handling its interactions. It is an essential part of the user interface and provides a graphical interface to interact with the system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows out-of-the-box filters (continued)

Name	Description
MonitoringHost.exe	MonitoringHost.exe is a legitimate system process that is a part of the Windows operating system. It is used by the Windows Error Reporting (WER) infrastructure to collect and send telemetry and diagnostic data to Microsoft. This data is used to improve the stability and performance of Windows.

Configure Discovery to use Event Framework

Configure Discovery to process jobs using the Event Framework method, which enables you to manage discovery throttling effectively and prioritize tasks based on their importance in the queue.

Before you begin

Confirm the following:

- You're using Discovery Admin Workspace v1.9.0 or later.
- You're using the Xanadu Patch 9, Yokohama Patch 4, or later version of the ServiceNow AI Platform®.

Role required: discovery_admin

About this task

When not using the Event Framework method, Discovery creates a one-time Schedule Item [sys_trigger] record for each External Communication Channel (ECC) queue sensor or pattern input record. After the task is processed, the system scheduler deletes the sys_trigger. This approach introduces additional overhead and can lead to significant database churn, while offering limited control over how sensors are processed.

The Event Framework reduces database churn by cycling through multiple event tables over time. Instead of deleting old records one by one, it clears out entire tables when they're no longer needed and reuses them. This approach is much more efficient and helps keep the system running smoothly.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. In the System Properties [sys_properties] table, select the *discovery.use.event.processing* property.
3. Set the **Value** field to **true**.
4. Select **Update**.

Result

When an input record is added to the ECC Queue, a business rule is triggered, which fires an event. The system processes these events at regular intervals with a default priority of 100.

Note: To update a job priority, configure the *com.glide.queue.discovery.sensors.processing_jobs.priority* and *com.glide.queue.discovery.cloud.sensors.processing_jobs.priority* system properties. For more information, see Configure background worker job priority for Event Framework.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Optimizing Discovery load for Event Framework

You can optimize Discovery properties to adjust throughput and enhance safety when pinning jobs.

Event Framework Discovery tuning

Issue	Action	Related task
Your instance is slow	Decrease the scale factor	Adjust the scale factor for Event Framework jobs
Discovery is slow	Increase the scale factor	Adjust the scale factor for Event Framework jobs
Discovery is slow when both Cloud and IP Discovery are running	Use a single queue for both Cloud and IP-based events	Configure a single queue for Cloud-based and IP-based events
Discovery is preempting other work	Reduce the worker priority of the background job	Configure background worker job priority for Event Framework
Nodes are executing too many Discovery jobs simultaneously	Configure event processors to execute on specific nodes	Configure event processors to execute on specific nodes

Adjust the scale factor for Event Framework jobs

Adjust the scale factor to modify the number of worker threads per app node, which increases or decreases throughput.

Before you begin

Confirm the following:

- You're using Discovery Admin Workspace v1.9.0 or later.
- You're using the Xanadu Patch 9, Yokohama Patch 4, or later version of the ServiceNow AI Platform®.
- The *discovery.use.event.processing* system property is turned on. For more information, see Configure Discovery to use Event Framework.
- The **Job configuration type** field of the Queue Registration form is set to **Scale with nodes**. For more information, see Queue Registration form reference.

Role required: discovery_admin

About this task

If your instance becomes slow while Discovery is running, you can adjust the scale factor to modify the number of worker threads per app node. The default scale factor is set to 2, which corresponds to two worker threads per app node.

For example, if you increase the scale factor to 3, Discovery uses three worker threads per app node, increasing throughput. Conversely, reducing the scale factor to 1 limits Discovery to a single worker thread per app node, which decreases throughput.

Procedure

1. Navigate to **All** > **System Policy** > **Queue Registry**.
2. Choose one or more job types for which you want to adjust the scale factor.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Job type	Action
IP-based	Select *discovery.sensors* from the Queue Registry [sysevent_queue] table.
Cloud-based	Select *discovery.cloud.sensors* from the Queue Registry [sysevent_queue] table.

3. In the **Scale factor** field, set a value.

   Note:

   - On the Queue Registry list view, this field is labeled **Job configuration value**. On the record form, it is labeled **Scale factor**.
   - The default maximum value is 3. To increase the scale factor beyond 3, open the Queue Provider Param [sysevent_queue_provider_param] table, filter by **Provider: Event Provider**, and increase the value in the **max_thread_utilization** field. This change affects all jobs in the Queue Registry [sysevent_queue] table where the provider is Event Provider. By default, only `discovery.sensor` and `discovery.cloud.sensors` have scale factor values high enough to be affected.

4. Select **Update**.

Configure a single queue for Cloud-based and IP-based events

Configure Cloud-based and IP-based events to share a single queue, which helps reduce the impact of Discovery on app nodes.

Before you begin

Confirm the following:

- You're using Discovery Admin Workspace v1.9.0 or later.
- You're using the Xanadu Patch 9, Yokohama Patch 4, or later version of the ServiceNow AI Platform®.
- The *discovery.use.event.processing* system property is turned on. For more information, see Configure Discovery to use Event Framework.

Role required: discovery_admin

About this task

When using a single queue, Cloud Discovery jobs frequently have higher priority than IP-based Discovery jobs. Separating Cloud-based and IP-based job queues helps avoid interference between Discovery types. However, balancing Discovery throughput and instance impact involves trade-offs.

Instance impact based on queue configuration

Separate queues	Single queue
More worker threads, which means more database activity	Fewer worker threads, which means less database activity
More impact on jobs unrelated to Discovery	Less impact on jobs unrelated to Discovery
Cloud Discovery doesn't impact IP Discovery	Cloud Discovery can preempt IP Discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. In the System Properties [sys_properties] table, select the *discovery.event.separate.cloud.events* property.
3. Set the **Value** field to **false**.
4. Select **Update**.

Configure background worker job priority for Event Framework

Configure the job priority of background workers to help prevent noncritical tasks from preempting important work, which promotes system efficiency and stability.

Before you begin

Confirm the following:

- You're using Discovery Admin Workspace v1.9.0 or later.
- You're using the Xanadu Patch 9, Yokohama Patch 4, or later version of the ServiceNow AI Platform®.
- The *discovery.use.event.processing* system property is turned on. For more information, see Configure Discovery to use Event Framework.
- The *com.glide.queue.discovery.sensors.processing_jobs.priority* and *com.glide.queue.discovery.cloud.sensors.processing_jobs.priority* system properties are configured. For more information, see Add a system property.

Role required: discovery_admin

About this task

By default, Event Framework jobs have a priority of 100. If Discovery is preempting more important work, you can use the *com.glide.queue.<queue_name>.processing_jobs.priority* property to adjust the priority.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Choose one or more job types for which you want to adjust the priority.

Job type	Action
IP-based	Select *com.glide.queue.discovery.sensors.processi* from the System Properties [sys_properties] table.
Cloud-based	Select *com.glide.queue.discovery.cloud.sensors.pr* from the System Properties [sys_properties] table.

3. In the **Value** field, enter `110`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Increasing the number in the **Value** field lowers the job's priority, while decreasing the number raises it.

4. Select **Update**.

Configure event processors to execute on specific nodes

Configure the *discovery.event.pin.jobs* system property to enable event processors to execute on specific nodes and distribute the work evenly. The scale factor calculates the number of nodes but doesn't pin jobs to specific nodes.

Before you begin

Confirm the following:

- You're using Discovery Admin Workspace v1.9.0 or later.
- You're using the Xanadu Patch 9, Yokohama Patch 4, or later version of the ServiceNow AI Platform®.
- The *discovery.use.event.processing* system property is turned on. For more information, see Configure Discovery to use Event Framework.

Role required: discovery_admin

About this task

The *discovery.event.pin.jobs* system property determines how jobs are assigned to nodes. When that property is set to **true**, jobs are evenly distributed across specific nodes, but scheduling flexibility is limited. When the property is set to **false**, jobs are randomly assigned, which can increase Discovery throughput.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. In the System Properties [sys_properties] table, select the *discovery.event.pin.jobs* property.
3. Set the **Value** field to **true**.
4. Select **Update**.

Post-clone Discovery configuration

When a clone occurs, Discovery schedules are copied from the source instance to the target instance. Additional configuration is necessary for these schedules to function correctly on the target instance, helping you properly configure Cloud-based and IP-based Discovery schedules and maintain optimal performance.

Note:

For general information about instance cloning, see the Clone FAQS-Frequently Asked Questions [KB0715621] article in the Now Support Knowledge Base.

For instructions on how to deactivate or cancel a Discovery schedule after creating a clone, see the A Post-Clone script to Deactivate and Cancel Discovery Schedules [KB0789119] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Post-clone target configuration for IP-based Discovery schedules

After completing the Request a clone process, post-clone configuration is determined by the MID Server selection method you chose in the Discovery schedule on the source instance.

Target instance configuration for IP-based Discovery schedules

MID Server selection method	Next steps
Specific MID	Update the MID Server record for each Discovery schedule.
Cluster	Create the cluster, attach the MID Server to the cluster, and then update the cluster name in the Discovery schedule. Alternatively, remove the MID cluster [ecc_agent_cluster] record from the Clone Exclude Tables [clone_data_exclude]. For more information, see Exclude a table from cloning.
Use Behavior	Create the behaviors, create a functionality or attach an existing functionality, then go back to the Discovery schedule and attach the behavior to the Discovery schedule. **Note:** Discovery functionalities aren't cloned with an instance.
Auto-select	Verify all MID Servers on the target instance are set up with target IP ranges, supported applications, and capabilities.

Post-clone target configuration for Cloud Discovery schedules

After completing the Request a clone process, you must access the Cloud Service Account [cmdb_ci_cloud_service_account] table and update the service accounts with the proper reference to the Discovery credentials. You must also update the MID Servers for the Discovery schedule. Additional post-clone configuration is determined by the MID Server selection method that you chose in the original Discovery schedule.

Target instance configuration for Cloud Discovery schedules

MID Server selection method	Next steps
Specific MID	Update the MID Server record for each Discovery schedule.
Cluster	Create the cluster, attach the MID Server to the cluster, and then update the cluster name in the Discovery schedule. Alternatively, remove the MID cluster [ecc_agent_cluster] record from the Clone Exclude Tables [clone_data_exclude]. For

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Target instance configuration for Cloud Discovery schedules (continued)

MID Server selection method	Next steps
	more information, see Exclude a table from cloning.
Use Behavior	Create the behaviors, create a functionality or attach an existing functionality, then go back to the Discovery schedule and attach the behavior on the target instance. **Note:** Discovery functionalities aren't cloned with an instance.
Auto-select	Verify that all MID Servers are set up with target IP ranges, supported applications, and capabilities.

For more information, see the Cloud discovery schedules fail after the instance clone [KB1759363] article in the Now Support Knowledge Base.

Clone Exclusions

Credential aliases are cloned but their credentials aren't cloned. If Discovery schedules have credential aliases, you must go to the Connection & Credential Aliases [sys_alias] table and confirm the alias points to the correct credentials. After cloning, update the Discovery schedules and add credentials within the target instance.

Additionally, no MID Server related tables are cloned. For more details, see the MID Servers and Clones [KBKB0786475] article in the Now Support Knowledge Base.

Related topics

Managing Instance Clone

Sensitive data filters

The Discovery Sensitive Data Filters [discovery_sensitive_data_filter] table provides a way to help prevent sensitive information from being exposed in the Configuration Management Database (CMDB) by applying redaction rules during data collection.

Discovery collects configuration and operational data from servers and applications to populate the CMDB. Some of this data may include sensitive information such as passwords, tokens, or credentials. Storing these values in CMDB can create security and compliance risks. Sensitive data filters enable administrators to define regex filters that identify sensitive information in probe results. When Discovery runs, the probe collects data and processes it on the MID Server. Before the ECC Queue input payload is sent to the instance, the Post Processor script applies transformations. The script checks for regex filters defined in the Discovery Sensitive Data Filters [discovery_sensitive_data_filter] table. If a match is found, the script redacts the sensitive value in the payload so that only redacted data is transmitted to the instance and stored in CMDB.

For example, if the regex is `(?i)(?:pwd|password|passwd|secret)=(\S+)` and the original content is:

```
user=admin password=MySecret123 host=localhost
```

After redaction, it becomes:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
user=admin password=REDACTED host=localhost
```

Requirements

Discovery must be using version XP11, YP10, ZP4 or later.

Visibility Content must be using version 6.29.0.

Benefits

Benefits of implementing sensitive data filters include:

- Protect confidential information: Prevent credentials or other sensitive values from appearing in logs or output.
- Compliance: Support organizational security and privacy standards by meeting requirements.
- Flexibility: Customize filters to your environment and data sources.

Examples

Examples of data you can protect with sensitive data filters include:

- Tracked configuration files: A MySQL configuration file may contain a password. A regex filter can detect the password and redact it.
- Process parameters: Linux server process arguments may include sensitive tokens. Filters can identify and redact these values.

Using Discovery

Run Discovery schedules and execute Discovery probes, sensors and patterns to find your IT resources and add them to the CMDB.

- Running discoveries in your network
- Discovery probes and sensors
- Patterns and horizontal discovery
- Discovery resource utilization
- Discovery monitoring and issue resolution

Pattern Diagnostic agentic workflow

The Pattern Diagnostic agentic workflow helps Discovery administrators investigate missing CI attributes. It identifies the gap, parses discovery logs, identifies the root cause, and suggests remediation — without manually navigating log files.

When Discovery runs, it populates CI attributes in the CMDB using both probe-based and pattern-based discovery. The Pattern Diagnostic agentic workflow supports investigation into CIs discovered through pattern-based discovery. When an attribute is missing, identifying the cause requires navigating multiple tables and interpreting nested JSON in discovery logs. The Pattern Diagnostic agentic workflow automates this investigation and suggests a remediation action, all from the ServiceNow Otto panel.

Requirements

ServiceNow Otto for IT Operations Management (ITOM) must be installed on your instance. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Discovery and Service Mapping plugin (com.snc.discovery) must be active for this workflow to be available.

Role required: discovery_admin.

Note:

AI Search must be enabled for the workflow to match natural language queries to CI tables. If AI Search is not enabled, you can still use the workflow by providing the direct table name in your query. To enable AI Search, navigate to **All** > **AI Search** > **Enable AI Search**. After enabling, verify that the search status is enabled to confirm indexing is complete.

Required property configuration

The *glide.discovery.save_pattern_log* property controls whether pattern logs are saved after successful pattern execution. By default, this property is set to true, which saves all pattern logs.

If this property is set to false, successful pattern logs are not saved. This can prevent the Pattern Diagnostic agentic workflow from performing missing attribute analysis in cases where the attribute is not failing the pattern. Verify that this property is set to true to enable complete analysis. For information about configuring this property, see Discovery properties.

Pattern Diagnostic agentic workflow overview

A Discovery administrator triggers the workflow from the ServiceNow Otto panel in Discovery Admin Workspace by asking a natural-language question or by selecting the corresponding pill. The workflow supports both exploratory investigation, where it surfaces available attributes for selection, and direct investigation, where the administrator specifies the CI class and attribute in the initial question. The workflow then identifies affected CI records with the selected attribute missing, parses the relevant discovery logs to identify the root cause and, when the root cause is classified, displays the corresponding error code, and suggests a remediation.

The workflow uses two agents:

- Pattern Diagnostic agent: Receives the administrator's query and runs script tools autonomously to investigate the missing attribute.
- Error Framework Remediation agent: Receives the identified root cause and suggests a remediation.

The second agent runs only when the root cause is classified and logged to the Error Framework. If classification or logging doesn't succeed, the workflow ends after the diagnostic summary, which lists the affected CIs, the error messages from the discovery logs, and general troubleshooting guidance.

When the Error Framework plugin is active and remediations are mapped for the identified error, the workflow provides enhanced remediation suggestions.

Scope

The workflow covers attribute coverage gaps for CI classes discovered through pattern-based Discovery. The following conditions apply:

- Only CIs discovered through patterns are supported for analysis.
- Analysis is based on the pattern metadata table (sn_disco_ai_pattern_metadata). This table contains metadata about the patterns populating a CI and its attributes. Only default pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

information is included; customizations made to patterns aren't considered. Metadata population is limited to an agreed set of CI classes.

- Log analysis focuses on the first CI in the sample. If issues are found, the workflow reports them and stops. If no issues are found, it checks the next CI in the list, up to the sample limit of five.
- Log analysis covers the last five days. If no results are found in that period, the search expands to 30 days.

Examples

The following scenarios illustrate how the workflow handles different CI and pattern types.

Scenario 1: Direct infrastructure (Linux Server)

The workflow investigates a missing `cpu_manufacturer` attribute on Linux Server CIs. The workflow resolves the CIs IP address and queries the discovery logs directly. It returns a full report with sample CIs investigated, link to CIs, root cause and remediation guidance.

Scenario 2: Applicative pattern (MySQL database)

The workflow investigates a missing attribute on a MySQL database CI. MySQL CIs are discovered through their host servers. The workflow navigates the Runs on relationship to resolve the host server's IP address and queries the discovery logs on the host, not the MySQL instance. It returns a full report with sample CIs investigated, link to CIs, root cause and remediation guidance.

Scenario 3: Cloud pattern (AWS Auto Scaling Group)

The workflow investigates a missing `cluster_connection` attribute on an AWS Auto Scaling Group CI. Cloud CIs use an SA-LDC identifier instead of a traditional IP address. The workflow navigates the Hosted on relationship to resolve the SA-LDC identifier and searches logs using the identifier and process ID. It returns a full report with sample CIs investigated, link to CIs, root cause and remediation guidance.

Running discoveries in your network

You can run discoveries from schedules or scripts to create configuration items, define subnets, or to find resources in AWS and Azure clouds.

MID Server configuration prerequisites

Ensure that your MID Servers are properly configured before creating a Discovery schedule.

- **Supported applications**: Select the applications that are allowed to use the MID Server. You can use the **ALL** application option to allow any application to use the MID Server.
- **IP ranges**: Define the ranges of IP address the MID Server can scan. To find a MID Server match, the IP range you configure on the Discovery schedule must fall into the ranges that one or more MID Servers can support.

  Note: To improve security, limit the range to IP address you control and exclude unnecessary ranges.

- **Capabilities**: Create the capabilities that the MID Server supports. You can use the **ALL** capability option to allow any application to use the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery configuration prerequisites

Ensure that your MID Servers can authenticate on the devices they find and classify configuration items (CI) properly.

- **Credentials**: Configure the MID Servers with the login credentials they need to query the devices in the network. The MID Server tries all available credentials on each discovered device, then creates an affinity for any successful credentials. For more information, see Credential affinity for Discovery and Orchestration.
- **Classifications**: The device and process classifications provided in the base platform are normally sufficient. Create classifications as needed for the devices, processes, and applications in the network not covered by the default classifiers.

Get started running a discovery

1. Use the Discovery Configuration Console to get started with Discovery. The console provides configuration options which let you choose the types of devices, applications, software files, and software CIs you want Discovery to find. If you select a CI to exclude from scanning, the instance disables the related probe or classifier that Discovery uses to identify the CI. See Discovery Configuration Console to get started.
2. Determine what type of discovery to run:
   - Run a Configuration item (CI) discovery to find the devices, computers, and applications on your network. This is the most common type of discovery. Run CI discovery from the Discovery schedule, where you to set up a recurring schedule or run a discovery on demand. The Discovery schedule also provides configuration options for MID Servers and the Shazzam port probe.
   - Run a Network Discovery to find the internal IP networks within your organization. If you already know the IP address ranges in your network, it is not necessary to run Network Discovery. It is intended for organizations that don't have complete knowledge of the IP addresses available for Discovery in their networks.
   - Run a Serverless Discovery to find applications on host machines without the need to discover the host first. Serverless Discovery relies on infrastructure patterns to explore CIs on a host. This kind of discovery skips the scanning and classification phases of discovery. You need an advanced knowledge of patterns to use this type of discovery. Refer to Patterns and horizontal discovery to get started with patterns.
3. After you run a discovery, monitor the results of the discovery and resolve errors if they occurred:
   - Use the Discovery status to see a summary of a discovery and to access the ECC queue. The ECC queue shows probe and sensor activity and the actual XML payload that is sent to or from an instance.
   - Use the Discovery Admin Workspace to monitor ongoing Discovery operations.
   - Use the Discovery Home page to access details for all schedules, cloud resources (virtual machines), discovered devices, and related errors that might have occurred. Error details include possible remediation steps.

Related topics

Discovery resource utilization

Discovery Configuration Console

Use the Discovery Configuration Console to manage what kind of configuration items (CIs) and CI information you want to discover.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, Discovery finds all the information on your network that is specified in probes and patterns. Use the controls in this console to select the data Discovery adds to the Configuration Management Database (CMDB) that your organization needs. You can control these aspects of Discovery:

- The discovery of entire CI types, such as Windows servers.
- The discovery of specific CI details, such as OS information on Windows servers.
- The discovery of ARP data for Networks, UNIX, and Windows devices.
- The discovery of software packages containing keyword terms that you add, such as Hotfix or Security Update.
- File-based Discovery of software on UNIX and Windows by particular extension or folder, such as .cert or .ext.

Note: Discovery returns all data from the target to the instance before applying console settings. The instance ignores any excluded data when it updates the CMDB.

Requirements and accessibility

If you use Internet Explorer, you must use version 8 or later. The configuration console supports keyboard navigation and screen readers.

Console overview

To open the configuration console, navigate to **Discovery Definition** > **Configuration Console**.

The console is divided into these sections. Use the toggle buttons to enable various options.

- **Devices**: network devices such as printers, storage devices such as storage switches, and Unix and Windows computers
- **Applications**: automation applications such as Puppet, databases such as MSSQL, and web servers such as Tomcat
- **Software Filter**: UNIX and Windows applications that include or exclude keywords that you enter
- **File-based Discovery**: software files by path, folder, and extensions to scan or ignore

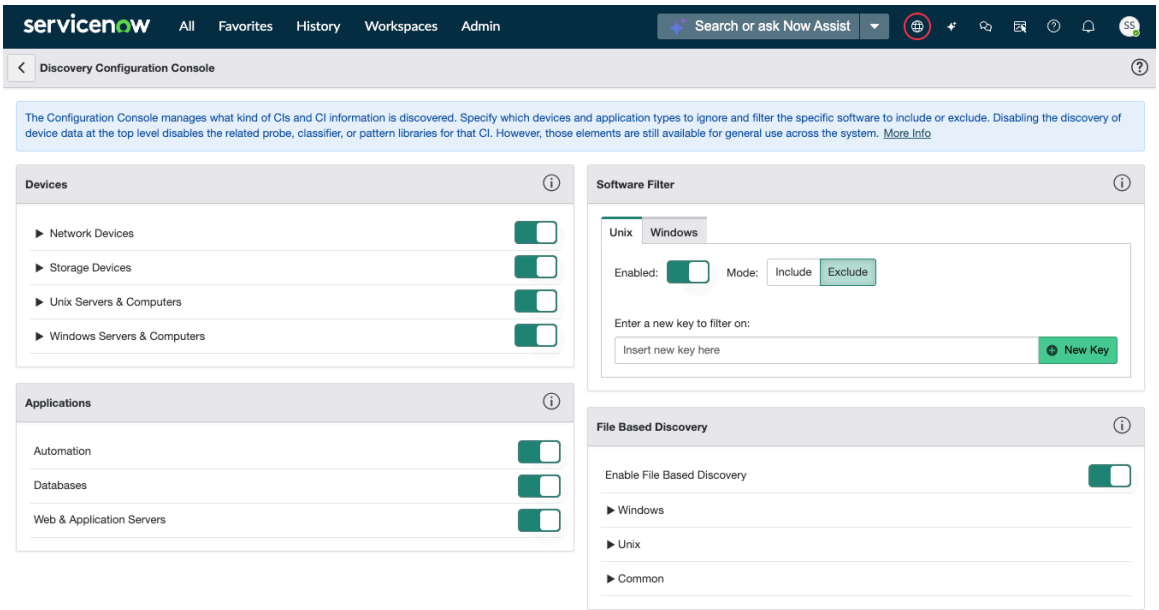

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Device Discovery

Disabling the discovery of device data at the top level disables the related port probes and pattern libraries. For example, blocking Network Devices disables the SNMP port probe. Blocking devices at the second level of the hierarchy has the following effects:

- Device class categories disable classifiers.
- Device info categories disable probes and pattern libraries.

You can enable or disable discover ARP data. This option is turned off by default. After upgrading, ARP data is not be discovered unless you enable it in the Discovery Configuration Console. ARP data can be controlled in three device classes: Network, UNIX, and Windows.

Potential impacts include the following:

- Without ARP data, the network path might not be explored if you're using Service Mapping with L2 or L3 Discovery.
- After upgrading and rerunning Discovery, ARP data does not populate for existing customers. Old data is auto-flushed per the deletion strategy of the out-of-the box (OOB) behavior. This is a one-time activity unless you change the OOB configuration. Therefore, you must run your Discovery schedule based on the device volume.
- If you have turned off the Network Device Exploration for network devices, you should also keep ARP turned off to retain the old behavior before upgrading.

Application Discovery

Disabling the discovery of application data affects all host devices on which the application runs. For example, if you configure Discovery to ignore databases, no information is gathered for either Linux or Windows databases. Conversely, if you configure the system to ignore a device type, such as a Windows server, no databases running on that server are explored, even if they are configured to be discovered. The instance cannot identify applications running on a server until it discovers that server first.

When you exclude an item from the **Applications** section, the system disables the relevant process classifier. The probes that belong to this classifier, including the Horizontal Pattern probe that launches patterns, never launch.

Software Discovery

You can filter the discovery of software packages based on the keywords they contain for Windows and UNIX computers. A keyword can be any term that is present in a software package name. Keywords for filtering software are stored in the Software Filter Keys [discovery_spkg_keys] table. The system does not provide default filtering keywords for UNIX software.

Important: Keyword software filtering occurs after the probe returns its payload to the instance and not on the target host.

The following default keywords are provided for Windows and are filtered out by default:

- Hotfix
- Language Pack
- Security Update

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File-based Discovery

You can configure Discovery to either scan or ignore paths, folders, and extensions for File-based Discovery. Discovery then scans a Windows or UNIX file system with these settings to identify file fingerprints. Installed software products detected with these attributes are matched to products in a comprehensive software normalization library.

Effects on Discovery tables

This table shows you what is inactive in Discovery tables when you use the Discovery Configuration Console.

When you disable	Field inactive
Top-level device categories, such as: • Network Devices • Windows Servers & Computers • Unix Servers & Computers	The **CIs** field is inactive on the Port probes [discovery_port_probe] table. The Port probes [discovery_port_probe] table is referred to from the Protocol Category [discovery_category_protocol] table.
Device classes, such as: • Windows 2008 Servers • Linux Servers	The **Active** field is inactive on the Classifiers [discovery_classy] table for the related classifier. The Classifiers [discovery_classy] table is referred to from the Device Classification Category [discovery_category_device_class] table.
Device info, such as: • Network Device Exploration • Windows Storage • ARP Data	The **Active** field is inactive on the Classifier Probes [discovery_classifier_probe] table for the related probe. It is also inactive on the Discovery Patterns [sa_pattern] table for the related pattern that performs CI identification and exploration. These tables are referred to from the Device Info Category [discovery_category_device_info] table.
Applications	The **Active** field is inactive on the Process Classifiers [discovery_classy_proc] table for the related process classifier. The Process Classifiers [discovery_classy_proc] table is referred to from the Application Category [discovery_category_appl] table.

Note: If you use patterns for discovery, and you disable the discovery of device classes or device information, the system disables the pattern libraries that collect the device information. For example, if you have a Windows pattern in which you create a shared library that collects OS information, and then you turn off the collection of OS data, the shared library that collects OS data is inactive. For more information about the shared library operation, see KB1698521: Create or customize patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Exclude CIs from discovery

Use the Discovery Configuration Console to determine which devices and applications to exclude from discovery.

Before you begin

Role required: discovery_admin or admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Definition** > **Configuration Console**.
2. If you want to exclude specific devices and device information, expand the category for the device types.
3. To see what applications can be disabled, hover your cursor over the application type.
4. Click the toggle button to turn off the CIs you do not want to discover.
   The instance creates an update set record for any change you make to the console.

Customize Applications and Devices to Discover

You can toggle on and off Applications and Devices using the Discovery Configuration Console. Some are listed out-of-the-box, but you may also customize which specific Applications and Devices you want to Discover according to your specific requirements.

Before you begin

Role required: admin

Procedure

1. Customize Applications.
   a. Navigate to **Discovery Definition** > **Application Categories**.
   b. Click **Automation**, **Databases**, or **Web & Application Servers**.
   c. Unlock the Process classifier list.
   d. To delete from the list, select and application then click the **X** to remove.
   e. To add to the list, type in the box and click **Enter**.
   f. Click the **lock** icon.
   g. Click **Update**.
      The next time Discovery runs, it will process those types of specific applications that you configured.
2. Customize Devices.
   a. Navigate to **Discovery Definition** > **Device Info Categories**.
   b. To create a new record, click **New**.
   c. Enter the details and click **Submit**.
   d. To update an existing device, click on a device listed.
   e. Make any changes and click **Update** or **Delete**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Filter software to discover

Use the Discovery configuration console to filter the software you add to the CMDB.

Before you begin
Role required: discovery_admin or admin

About this task
The instance creates an update set record for any change you make to the console.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Definition** > **Configuration Console**.
2. In the Software Filter panel, select the tab for the operating system you want, either **Unix** or **Windows**.
3. Add or delete keywords for filtering software packages.
   a. To add a keyword, enter it in the field provided and click **New Key**.
   b. To delete a keyword, click the red **x** icon to the right of the keyword.
   Discovery looks for software packages containing these terms in the selected operating system and either includes or excludes them from the CMDB. The system does not provide default keywords for filtering UNIX software packages, but provides these default keywords for Windows:
   - Hotfix
   - Language Pack
   - Security Update

   These default keywords exclude patch-related software by default. To enable Security OS Patch Discovery, remove *`Hotfix`* and *`Security Update`* from the exclusion list. This allows Discovery to include patch and hotfix data in the CMDB.
4. Select whether to **Include** or **Exclude** software packages with the keywords in their names.
   - **Include**: Only add software packages that match the keywords.
   - **Exclude**: Add all software packages except those matching the keywords.
5. Disable software filtering for either **UNIX** or **Windows**.

   When the filter is disabled for an operating system, all discovered software packages for that operating system are added to the CMDB, with no filtering applied.

   The instance creates an update set record for any change you make to the console.

Configure File-based Discovery

Use the Discovery Configuration Console to configure several settings to find and manage all of your installed software.

Before you begin
Role required: discovery_admin or admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Definition** > **Configuration Console**.
2. Select the toggle to enable File-based Discovery.
   This feature is turned off by default.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select customizable options specifically for Windows and UNIX machines as well as options common to both types.
   Some options have default values and others, such as folders to ignore or extensions, are empty until you add in values.
4. For Windows and UNIX, select the arrow to expand the target path options.
5. Configure basic selections.

Folders to Scan

Indicates which folders to scan for Windows and UNIX. Default values are listed. You can also add additional paths to scan.

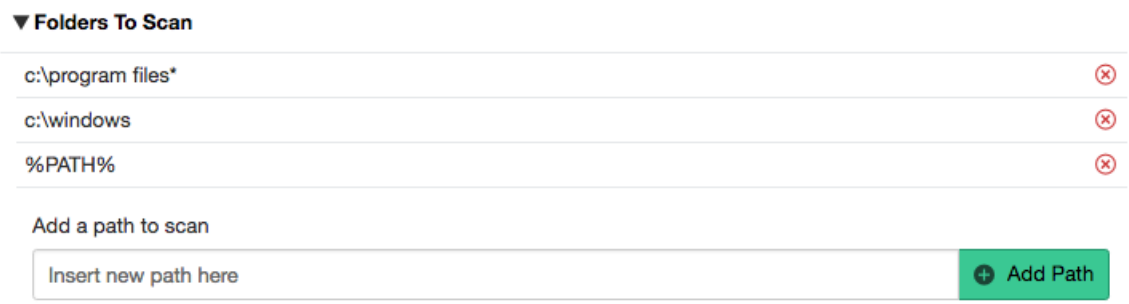

Note: If you want File-based Discovery to scan the entire Windows C: drive (not just the default folders), update the windows Folders to Scan target path to C:/ in the Discovery Configuration Console.

Folders to Ignore

Indicates which folders and child folders/files to ignore from within the list of Windows or UNIX folders that you want to scan. For example, if you have specified "C:\Program Files" as a folder to scan, you could additionally specify "C:\Program Files\Temp" as an ignore folder. When File-based Discovery scans a Windows server, it scans all files and sub folders under C:\Program Files, excluding C:\Program Files\Temp.

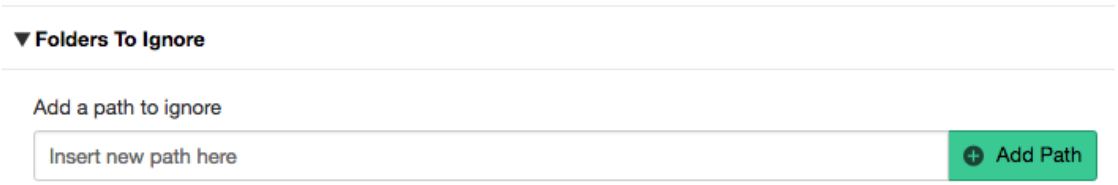

Extensions Exclusion List

Indicate extensions (.txt, .war, .exe, and so on) that you want to ignore for Windows.

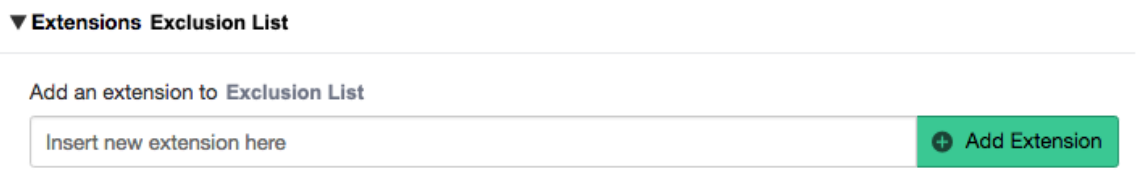

Extensions Wildcards

Indicate extensions (.cert, .crt, and so on) that you always want to scan for Windows. This enables you to scan for extensions such as *.cert where the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

file name itself might not be known ahead of time. You should be able to find certificates (.cert or .crt) and java keystore files (.jks) in this way.

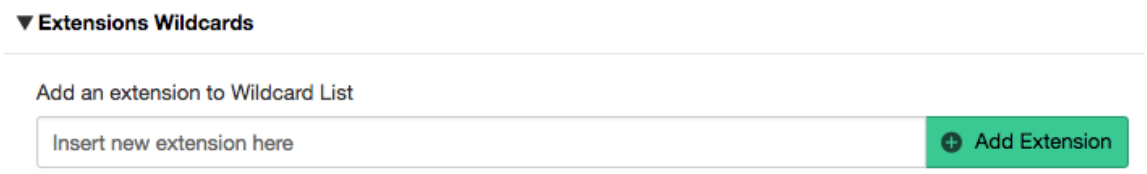

6. Configure **Advanced Windows** settings.

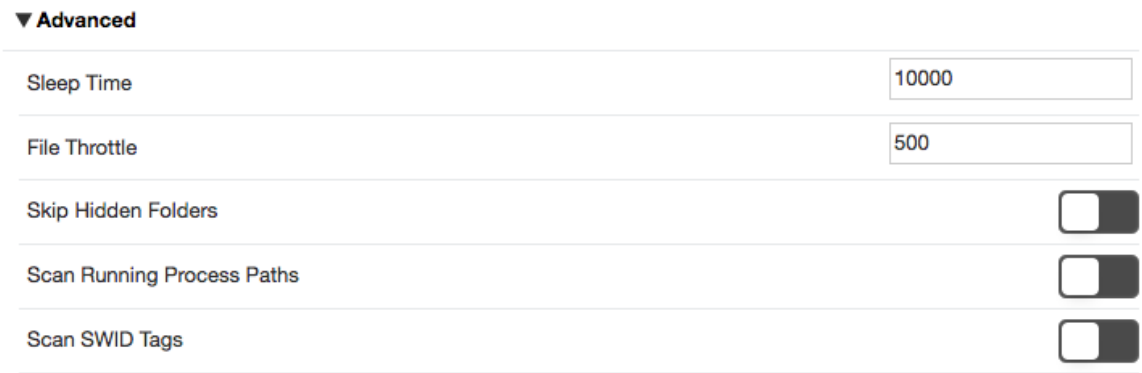

Sleep Time

Sets the length of time in milliseconds to sleep during each throttling interval while scanning for Windows files. The default value is 10000. Don't change this value unless you're an advanced user. If you're an advanced user and must make a change, make sure to set to a non-zero value.

File Throttle

Sets the number of files to scan on Windows before sleeping. The default value is 500. Don't change this value unless you're an advanced user. If you're an advanced user and must make a change, make sure to set to a non-zero value.

Skip Hidden Folders

Skips hidden folders on target computers when scanning for file signatures. For Windows devices, the visibility of a folder is based on a hidden attribute that can be set for each folder.

Scan Running Process Paths

Enables the use of running process paths in the list of targeted paths to scan. This automatically adds paths that contain running processes and therefore most likely contain files of interest for use in determining software installations.

Scan SWID Tags

Scans for and analyzes the software identification (SWID) tags. This option is only visible if SAMP is enabled. When enabled, the software product, publisher, and version are returned in the payload. This data is directly inserted into the software installs table, bypassing the software match API. This feature is turned off by default.

7. Configure **Advanced UNIX** settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

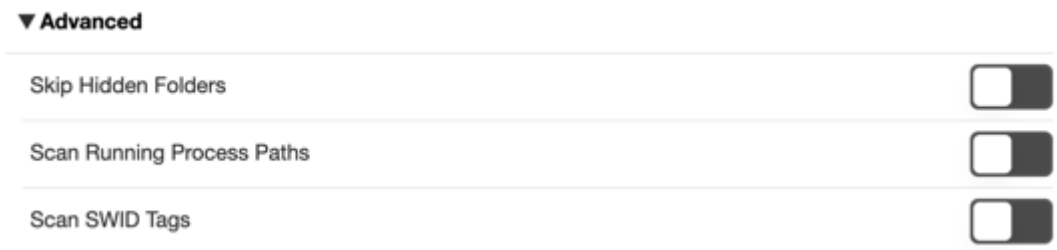

Skip Hidden Folders

Skips hidden folders when scanning for file signatures. For UNIX-based devices, these folders are prefixed by a "." in their name.

Scan Running Process Paths

Enables the use of running process paths in the list of targeted paths to scan. This automatically adds paths that contain running processes and therefore most likely contain files of interest for use in determining software installations.

Scan SWID Tags

Scans for and analyzes the software identification (SWID) tags. This option is only visible if SAMP is enabled. When Scan SWID Tags is enabled, this returns the software product, publisher, and version in the payload to be directly inserted into software installs table. This bypasses the software match API. This feature is turned off by default.

8. Configure **Common** settings for either Windows or UNIX.

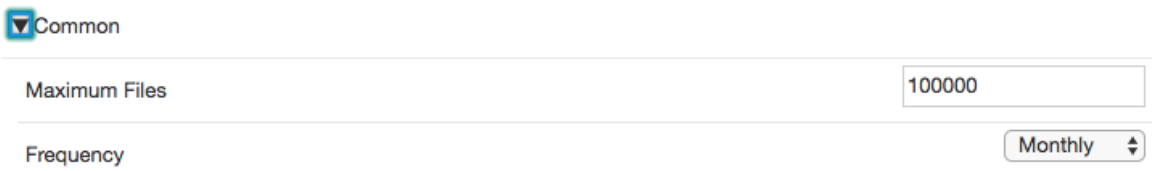

Maximum Files

Sets the maximum number of files that you want to discover. This is the number of files scanned per CI/IP.

Frequency

Sets the frequency that File-based Discovery runs on the CI.

Result

File-based Discovery is now enabled, configured, and ready for you to run. See Run File-based Discovery for more information.

Create Basic Auth credentials

Create credentials to access a ServiceNow® instance.

Before you begin

Role required: system_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Select **New**.
3. Select **Basic Auth Credential**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. On the form, fill in the fields.

Basic Auth Credentials form

Field	Description
Name	Enter a unique and descriptive name for this credential.
User name	Name to identify the user.
Password	Password to use this credential.
Active	Option to enable the use of this credential.
Order	The order (sequence) in which the platform tries this credential while it attempts to log in to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order. Default value: `100`

5. Select **Submit**.

Result
The Basic Auth credentials for use with the server instance are created.

Schedule a horizontal discovery

A discovery schedule determines what horizontal discovery searches for, when it runs, and which MID Servers are used. Create a discovery schedule for your local environment or a schedule for discovering the resources in your cloud service account.

Before you begin
Ensure that your Discovery schedule conforms to security best practices, such as limiting the range of discovery targets and using the most secure credentials.

Make sure to test your credentials before you run a schedule. Bad credentials are a leading cause of failed discoveries.

Roles required: discovery_admin or admin

About this task
You can use a Discovery schedule to launch horizontal discovery, which uses probes, sensors, and pattern operations to scan your network for CIs. Use this procedure to create a schedule manually from the **Discovery Schedules** form.

Service Mapping also provides a Discovery schedule for top-down discovery. See Schedule a top-down discovery by Service Mapping for more information.

Use the **Discovery Schedules** module in the Discovery application to:

- Configure a schedule to discover resources in your cloud service account.
- Configure a schedule to discover certificates from URL scans.
- Configure device identification by IP address or other identifiers.
- Determine if credentials are used in device probes.
- Name the MID Server to use for a particular type of discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Create or disable a schedule that controls when the discovery runs in your network.
- Configure the use of multiple Shazzam probes for load balancing.
- Configure the use of multiple MID Servers for load balancing.
- Run a discovery schedule manually.
- Run a discovery on a single IP address.

Note: To view the run-results of your schedules for both IP-based and Cloud Discovery, use the summaries on the Discovery Home page. The Home page publishes the details of any errors that might have occurred and displays possible actions to take to remediate problems.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules** to create a new record.
2. Select the type of schedule to open:
   - **New**: Creates a new horizontal schedule for discovering components in your network.
   - **Quick Discovery**: Runs a horizontal discovery on a single IP address without requiring a schedule.
   - **Create a Cloud Discovery schedule**: Creates a schedule, using the Discovery Manager wizard, for discovering resources in a cloud service account.
3. Complete the Discovery schedule form, using the fields in the table.
   For more information, see Discovery Schedule form reference.
4. Right-click in the header of the record and select **Save** from the context menu.
5. To create a range of IP addresses to discover, click **Quick Ranges** under **Related Links**.

   Note: To improve security, limit the range of discovery targets to exclude unnecessary networks and devices.

6. Define the frequency of schedule running as described in Run options for discovery schedules.

What to do next

After you create and save a Discovery schedule, you can optionally configure the following settings:

- Define IP addresses or ranges to exclude from all Discovery schedules. See Use Global Excludes List for IP addresses and ranges.
- Define Discovery schedule attributes to automatically set CI field values such as location or asset tag on discovered CIs. You can define attributes at the schedule, range set, or IP address range level. See Define CI field attributes.
- Configure Discovery behaviors to use multiple MID Servers for load balancing or to scan with multiple credentials. See Discovery behaviors.

Run a Quick Discovery

Quick Discovery, or DiscoverNow, allows an administrator to run a CI Configuration discovery on a single IP address without requiring a schedule.

Before you begin

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The platform automatically selects the correct MID Server to use for the discovery if one is associated with the IP address selected. If no MID Server is configured for the network in which that address appears, you can select a MID Server. Use this feature to discover new devices in the network as soon as they are connected to the network, rather than waiting for a regularly scheduled discovery.

To configure the system to automatically determine which MID Server to use, set up the IP range capabilities for each MID Server in your system.

You can run DiscoverNow from a Discovery schedule form or from a script.

Note: Quick Discovery using a IPv6 target address is supported.

Procedure

1. Open Quick Discovery from one of these locations:
   - Navigate to **Discovery** > **Discovery Schedules** and click **Quick Discovery** in the header bar.
   - Navigate to **Discovery** > **Home** and click **Discovery Quick Start** under the Schedules tile.

   A dialog box appears asking for an IP address and the name of the MID Server to use. Only **Up** and **Validated** MID Servers are available.
2. Enter the target IP address for a discovery in the **Target IP** field.

   Note: DiscoverNow does not currently support IP network discovery. Make sure that you enter a single IP address only and not an entire network, such as 10.105.37.0/24.

   When a MID Server is assigned to the subnet containing the target IP address and currently in an operational status of **Up**, the name appears automatically in the **MID Server** field. If multiple MID servers are found, the system selects one for you. The value in the **MID Server** field can be overwritten if you want to select a different MID Server.

   Important: If the selected MID Server is part of a load balanced cluster and becomes unavailable for any reason, the instance does not assign another MID Server from that cluster to the quick Discovery. You must select another MID Server from the list of appropriate MID Servers.

3. If no MID Server is defined for that network, select one from the list of available MID Servers.

Quick Discovery Dialog

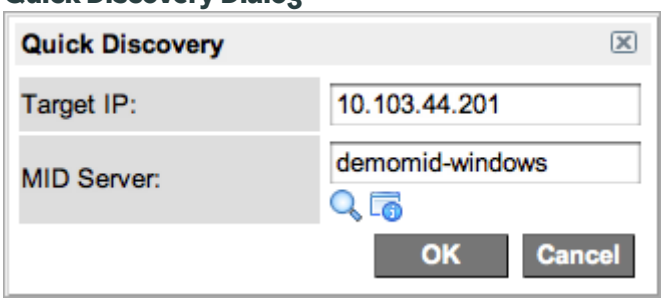

4. Click **OK** to run discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The status record for that discovery appears. The **Schedule** column is empty because no schedule is associated with this discovery.

Quick Discovery Status List

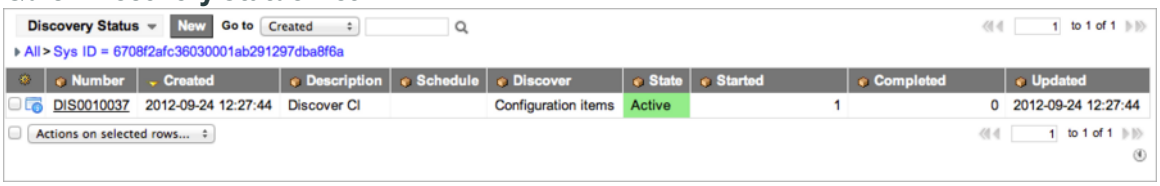

Run DiscoverNow from a script

You can run DiscoverNow from a script, such as a background job, a business rule, or web services.

Before you begin

Role required: admin

Procedure

1. Create the following script:

```
var d = new Discovery();
var statusID = d.discoveryFromIP(TARGET_IP, TARGET_MIDSERVER);
```

The *discoveryFromIP* method takes two arguments: *IP* and *MID Server*. The *IP* argument is mandatory, but the *MID Server* argument is optional.

2. To choose the MID Server, supply either the *sys_id* or name of the MID Server as the argument.
If you do not name a MID Server, the system attempts to find a valid one automatically. A valid MID Server has a status of **Up** and can discover the given IP address. If the system finds a valid MID Server and runs a Discovery, the *discoveryFromIP* method returns the *sys_id* of the Discovery status record. If no MID Server can discover this IP address, the method returns the value **undefined**.

   If you manually specify the TARGET_MIDSERVER, the system validates the given value and ensures that the MID Server table contains the specified MID Server record. If the validation passes, the discoveryFromIP method returns the sys_id of the discovery status record. If the validation fails, the method return the value **undefined**.

Validate discovery results

Validate the results of your discovery by accessing the ECC queue, analyzing the XML payload, and checking the Discovery log.

Before you begin

Role required: discovery_admin or discovery_operator

About this task

Initial discoveries often reveal unexpected results, such as previously unknown devices and processes or failed authentication. Results should also accurately identify known devices and update the CMDB appropriately. Become familiar with the network that is being discovered and the types of data returned for the different types of discoveries. Use the Discovery Log and the ECC Queue to monitor the Discovery process as data is returned from probes or pattern operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. To view the actual payload of a probe, click the **XML** icon in a record in the ECC Queue.

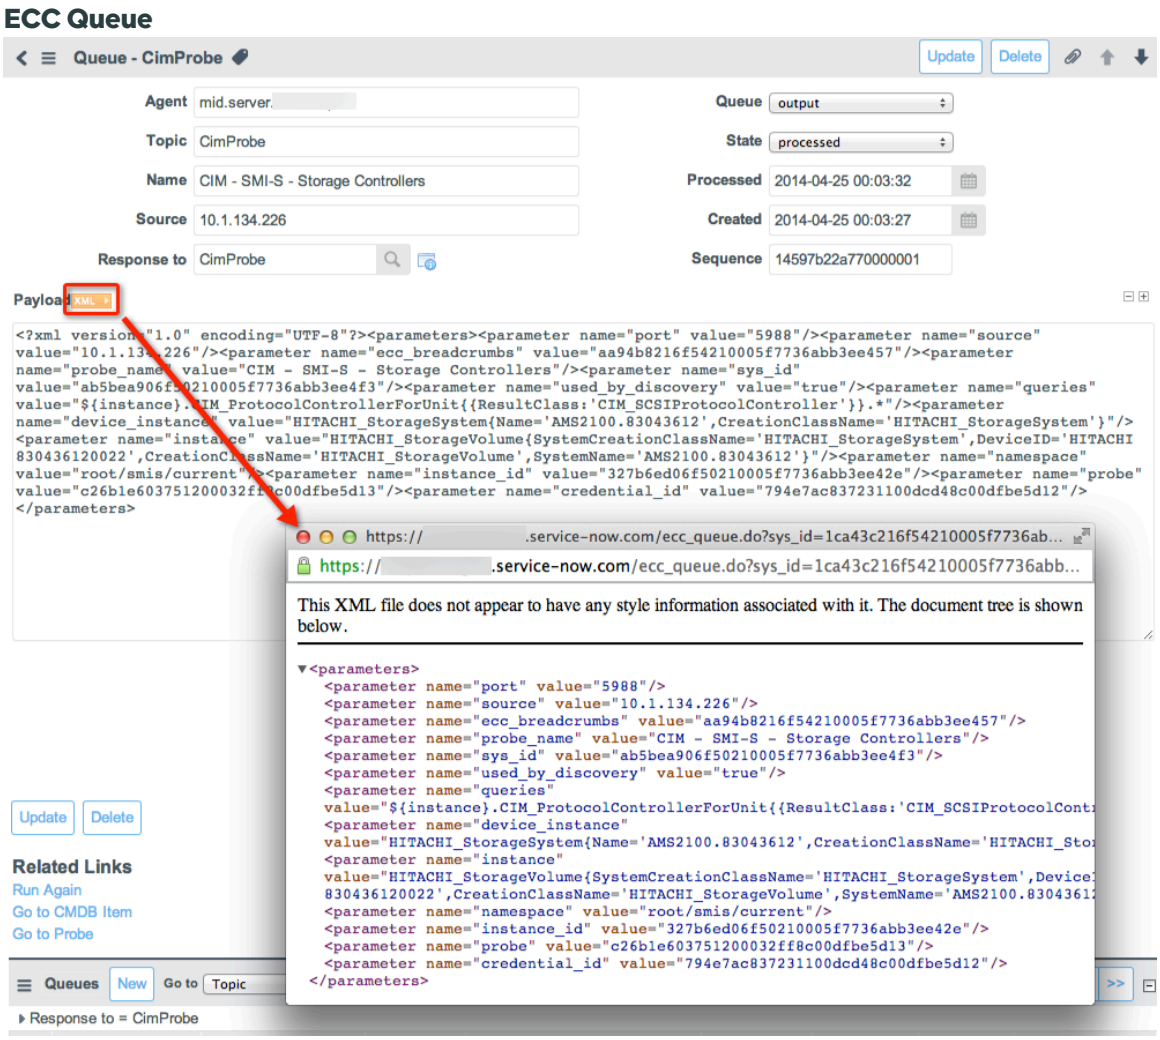

2. To view the actual payload of a probe, click the **XML** icon in a record in the ECC Queue.
3. Use the Discovery Log form for a quick look at how the probes are doing.

   To display the Discovery Log, navigate to **Discovery** > **Discovery Log**.

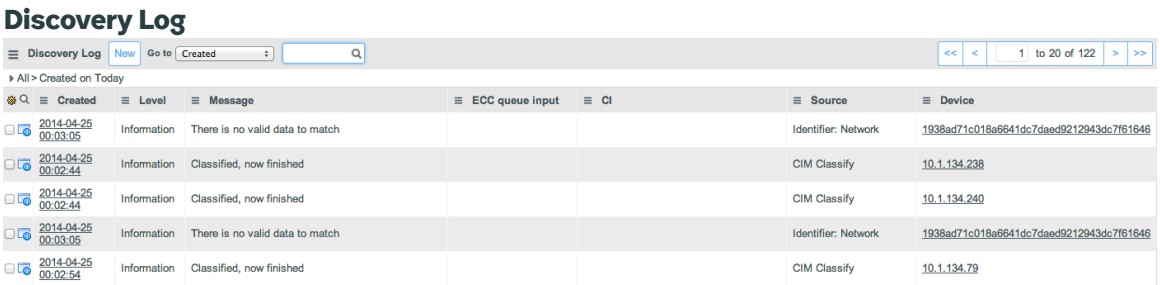

   The Discovery Log provides this information:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Information
Created	Displays the timestamp for the probe launched. Click this link to view the record for the probe launched in this list.
Level	Displays the type of data returned by this probe. The possible levels are: ○ Debug ○ Error ○ Information ○ Warning
Message	Message describing the action taken on the information returned by the probe.
ECC queue input	Displays the ECC queue name associated with the log message.
CI	The CI discovered. Click this link to display the record from the CMDB for this CI.
Source	Displays the probe name that generated the log message.
Device	Displays the IP address explored by the probe. Click this link to examine all the log entries for the action taken on this IP address by this Discovery.

Note: If you cancel an active discovery, note the following information:

- Existing sensor jobs that have started processing are immediately terminated.
- The existing sensor jobs that are in a **Ready** state, but have not started processing, are deleted from the system.

4. View the Discovery Home page for details about all schedules, cloud resources (virtual machines), discovered devices, and related errors that might have occurred. Error details include possible remediation steps.

MID Server selection sequence for Discovery schedules

The Discovery application follows this sequence to find a MID Server.

MID Server auto-selection

Discovery follows this sequence when you select **Auto-Select MID Server** for the **MID Server selection method** on the Discovery Schedule form.

Note: MID Server auto-select isn't supported with IPv6.

1. Discovery looks for a MID Server that also has an appropriate IP range configured.
2. If no MID Servers meet these criteria, it looks for a MID Server that has the **ALL** application that also has an appropriate IP range configured.
3. If more than one MID Server meet the criteria, Discovery chooses the first MID Server with the status of **Up**. If more than one MID Server is up, it randomly picks one.
4. If none are up, it uses the default MID Server specified for the Discovery application, assuming it's up.
5. If no default MID Server is specified, it uses the default MID Server specified for the **ALL** application, assuming it's up.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. If no default MID Server is specified, Discovery cycles through the previous steps and looks for MID Servers with the status of **Paused** or **Upgrading**.

> **Note:** When a MID Server is paused or upgrading, it doesn't actually process commands until it returns to the status of **Up**.

MID Server clusters

These steps are followed when you select **Specific MID Cluster** for the **MID Server selection method** on the **Discovery** form, and the cluster is a load balancing cluster:

1. **Discovery** uses the first MID Server in the cluster that it finds with the status of **Up**.
2. If more than one MID Server is up, it randomly picks one. If it can't find any MID Servers, it looks for MID Servers in the cluster with the status of **Paused** or **Upgrading**.

These steps are followed when the cluster is a failover cluster:

1. **Discovery** uses the MID Server with the lowest **Order** value that also has the status of **Up**.
2. If no MID Servers are found, it looks for MID Servers in the cluster with the status of **Paused** or **Upgrading**, choosing the one with the lowest **Order** value.

> **Note:** **Discovery** ignores the default MID Server for it and **ALL** applications when selecting a MID Server from the cluster.

Port scan (Shazzam) phase

During the port scan phase, Discovery collects all the target IP addresses. It splits them equally between MID Servers matching the criteria (MID Servers are qualified to do the port scan). The Shazzam batch size, which you configured on the Discovery schedule, determines the number of IP addresses that each Shazzam probe can scan. This phase helps determine how much work each MID Server does during the port scan phase.

For example, you have 16,000 IP addresses to scan among three qualified MID Servers, and you use the default Shazzam batch size of 5,000. Two of the MID Servers handle 5,000 IP address scans (one Shazzam probe each). The other MID Server handles 6,000 IP address scans by launching two Shazzam probes.

> **Note:**
>
> Shazzam can process IP lists containing up to 20,000 addresses that include both IPv4 and IPv6 types. If a schedule contains any IPv6 addresses (> 0), the combined total of IPv4 + IPv6 addresses must not exceed 20,000 or else the Discovery fails. If you're using only IPv6 addresses, you must use a list and not an IP address range. IPv6 address ranges and networks aren't supported and are ignored.

Related topics

MID Server pause

MID Server cluster configuration

Create an IP-based Discovery schedule in Discovery Admin Workspace

Use the Discovery Admin Workspace dashboard to create IP-based Discovery schedules.

Before you begin
Store version: 1.8.0.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **IP-based discovery** and select **Continue**.

   Note: This schedule is only for IP-based Discovery. For Cloud Discovery, see Create a discovery schedule in Cloud Discovery Workspace.

4. Enter a name for the Discovery schedule.
5. Choose from available Discovery locations.
6. Choose IP ranges by either selecting from existing sets or by manually entering the ranges.
   - To select from existing IP range sets, select **Use existing IP range set (recommended)**. See Discovery IP address configuration for more information on IP range sets.
   - To manually enter comma-separated IPv4 address ranges, IPv4 networks, or individual IPs that describe the ranges you want to add, select **Enter new IP ranges** and enter the IP ranges.
7. Select **next**.
8. Choose your MID Server.

Option	Description
Cluster of MID servers	Select an existing MID Server or create one. Clusters provide failover protection and load balancing between MID Servers. See Config for more information.
Auto-select MID server	An available MID Server is automatically selected when the Discovery schedule runs. See Automatic MID Server selection for more information.
MID server	Choose a MID Server to run the Discovery schedule. All MID Server are shown, but only ones that are up and validated can successfully run a Discovery schedule.

9. Select **next**.
10. Configure the Discovery schedule.

Option	Description
Run on demand	The schedule only runs when triggered manually. Trigger the schedule by selecting **Finish and run** at the end of this setup, or by navigating to the schedule in the Schedules table and selecting **Discover Now**.
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the schedule runs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Run after series	The schedule runs after another existing Dis covery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all the other associated schedules.
Cancel discovery if longer than maximum runtime	If the schedule exceeds the maximum run time, it's canceled. After toggled, configure the runtime threshold.
Finish and run	After being selected, all the information pro vided is validated. Then, a Discovery sched ule is created in the background, a Discov ery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After being selected, all the information pro vided is validated. A Discovery schedule is created and you're redirected to its entry in the Discovery Schedules [discovery_sched ule] table. You can edit the schedule informa tion or run it by selecting **Discover Now**.

What to do next

After you create and save a Discovery schedule, you can optionally configure the following settings:

- Define IP addresses or ranges to exclude from all Discovery schedules. See Use Global Excludes List for IP addresses and ranges.
- Define Discovery schedule attributes to automatically set CI field values such as location or asset tag on discovered CIs. You can define attributes at the schedule, range set, or IP address range level. See Define CI field attributes.

File-based Discovery

File-based Discovery helps you identify what software is running on your Windows and UNIX servers and devices, even if there's no registration information available. You can then manage and maintain records of your software licenses, check for unlicensed files, detect forbidden or damaged files, and help evaluate any threats from unwanted files.

Required plugins

The File-based Discovery [com.snc.discovery.file_based_discovery] plugin is required for file signature filtering. Your Discovery subscription includes this plugin, but you must request activation. Once the File-based Discovery plugin is active, the Software Asset Management - File Signature Normalization [com.snc.file_signature_normalization] plugin is also activated. For more information on the File Signature Normalization plugin, see File Signature Normalization.

How File-based Discovery works

File-based Discovery enhances the pre-existing discovery of installed software. It scans target servers for a known list of file signatures and processes those files with an established set of rules. The resulting data enhances the identification of installed software and identifies unregistered software products. For information about using Agent Client Collector for Visibility

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Content to perform file-based discovery, see Discover java installation data using Agent Client Collector for Visibility Content file-based discovery.

File-based Discovery is triggered in the exploration phase of normal Discovery. File-based Discovery probes execute a scan searching for specific file extensions or file names in paths that you configure. The resulting file information is returned in the probe payload. The sensor attempts to match the discovered files with installed software, using the file name, size, and version returned by the probe. File-based Discovery uses file signatures to detect software that might not have been registered. This information is then stored in the File Information [cmdb_file_information] table with a reference to the CI of the server. You can view the files found from each CI in a related list on this table. For more information, see Related list of CI components. When Software Asset Management (SAM) is active, if any file matches a software product, Discovery populates the Product and Publisher information for that file. Use this information to understand what software is running on your server and to help evaluate any threats from unwanted files. Discovery uses lists of known file signatures for Windows and UNIX to constrain the scope of the search. The filtering process for Windows and UNIX hosts is executed differently because their signature lists differ greatly in size. The smaller UNIX signature list is included with the Unix - File Discovery probe and processed directly on the target. The Windows signature list is larger and can't be processed on the target. The Windows - File Discovery probe scans the target for specific file extensions and paths and returns these results to the MID Server. The MID Server performs file signature filtering using the entire Windows list. The MID Server then sends all file information back to the instance for normalization and matching.

If SAMP is active on the instance, File-based Discovery creates or updates identified software products in the Software Installation [cmdb_sam_sw_install] table and updates matched software package licenses. Without SAMP, no software records are created. Only the file information goes into the File Information [cmdb_file_information] table.

You can enable SWID tags in the Discovery Configuration Console. With SWID tag enabled, when running File-based Discovery, the SWID tag information then populates the [cmdb_swid_tag] table. Information about the software installed on a particular machine includes name, file information, publisher, version, installed on, and content. The software_installation column in the [cmdb_swid_tag] is a reference to the [cmdb_sam_sw_install] table.

Note: Base64 package is a prerequisite for any UNIX or Linux servers to scan SWID tag files using File-based Discovery.

File-based Discovery inserts any file not matched by the normalization process into the Unidentified File Set [cmdb_unidentified_file_set] table. You can update the records in this table and provide additional details for previously unidentified files. If you provide values for the **Product** and **Publisher** fields for a file, SAMP settings can enable File-based Discovery to use that file for installed software matching in future discoveries.

You can disable File-based Discovery at any time by changing the setting in the Discovery Configuration Console. If you disable File-based Discovery before scan results are returned, the file data is ignored.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

File-based Discovery supports Windows, UNIX, and macOS devices. The UNIX probe is POSIX-compliant and should run on any Linux/Solaris server. Discovery supports Windows versions 2008, 2008R2, 2012R2, 2016, 2019, and later with PowerShell 3.0–7. Discovery also supports AIX versions 5.3, 6.1, and 7.1 and HP/UX 8.11.

To use File-based Discovery on target hosts that run PowerShell 7, configure the MID Server to prefer PowerShell 7. For more information, see Configure MID Server parameters for PowerShell 7.

If you're running File-based Discovery on Ubuntu version 20, modify the default Bourne shell (sh) to point to Bourne Again shell (bash).

Version information is populated only for the files with version information returned from probes. Not all files have versions. Files with extensions such as .exe, .jar, and so on, have versions.

Target system permissions

File-based Discovery needs specific permissions on each UNIX target system to scan files and read software details. The following sections list the required permissions and the sudoers configuration for the discovery user.

The MID Server runs the discovery script on each UNIX system with sudo. When the discovery user has sudo access, File-based Discovery scans every directory in the scan path.

When the Discovery user doesn't have sudo access, File-based Discovery scans only the directories that all users can read. It skips protected directories and doesn't report an error for them.

On Windows systems, File-based Discovery uses the administrative discovery credential and runs a signed script. Windows systems don't use sudoers.

Require permissions

The Discovery user needs the following permissions on each UNIX target system:

Permission	Description
Run the discovery script with sudo	Enables file-based Discovery to scan directories that aren't readable by all users.
Read and run on the directories in the scan path	Enables the scan to move through the directory trees you select.
Read on the files that match the scan list	Enables File-based Discovery to read each file's size and details.
Read on the temporary directory, /tmp by default	Enables the scan to write and read its working files on the target system.
Read on SWID files, when SWID tag scanning is enabled	Enables File-based Discovery to read software identification details from SWID files.

Required operating system tools

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File-based Discovery uses standard operating system tools on each target system. The following tools must be available in the system path:

- find, wc, awk, and grep
- ps and nohup
- sh or bash

When SWID tag scanning is enabled, base64 or uuencode must also be available. File-based Discovery uses uuencode first and falls back to base64.

Sudoers configuration

Add entries to the sudoers file on each UNIX target system. These entries give the discovery user the sudo access that File-based Discovery needs. Replace discoUser with your discovery credential user and adjust the tool paths to match your installation.

Linux, CentOS, Ubuntu, and RHel use the following entries for Linux systems:

```
Cmnd_Alias FBD_SCRIPT
 = /bin/sh /tmp/SNC_File_Discovery_Script_*.sh
Cmnd_Alias FBD_CMDS = /usr/bin/find *, /usr/bin/wc
 *, /bin/cat *
discoUser ALL=(root) NOPASSWD:FBD_SCRIPT, FBD_CMDS
Defaults!FBD_SCRIPT !requiretty
Defaults!FBD_CMDS !requiretty
```

To enable SWID tag scanning on Linux, add the following entry:

```
Cmnd_Alias FBD_SWID = /usr/bin/base64
 *, /usr/bin/uuencode *
discoUser ALL=(root) NOPASSWD:FBD_SWID
Defaults!FBD_SWID !requiretty
```

Use the same script and command entries on AIX, Solaris, and HP-UX systems. For SWID tag scanning on these systems, use uuencode:

```
Cmnd_Alias FBD_SWID = /usr/bin/uuencode *
discoUser ALL=(root) NOPASSWD:FBD_SWID
Defaults!FBD_SWID !requiretty
```

On AIX, use this requiretty syntax instead of the Linux form:

```
Defaults:discoUser ! requiretty
```

Example

A Discovery user scans the /opt and /etc directories on a Linux system. These directories aren't readable by all users.

With the sudoers entries in place, the discovery user runs the scan with sudo. File-based Discovery reads the files in both directories and reports the software it finds.

Without the sudoers entries, File-based Discovery skips /opt and /etc. It reports only the software in directories that all users can read.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Run File-based Discovery

File-based Discovery references

File-based Discovery issue resolution

Run File-based Discovery

Run File-based Discovery to find all of your installed software whether it is registered or not. You can enable and configure File-based Discovery at any time using the Discovery Configuration Console.

Before you begin

Role required: admin

Procedure

1. Set up the PowerShell script.
   By default, the **filebaseddiscovery.ps1** script has a ServiceNow signature. Its certificate chain resolves up to the VeriSign Universal Root Certification Authority which is trusted by Windows by default. Since PowerShell scripts are signed by ServiceNow publisher, add ServiceNow publisher to your trusted publisher repository.

   If your Execution Policy requires you to use your own certificate or if you must make any changes to the script, re-sign the script.

   a. Navigate to the probe "Windows - File discovery."
   b. Open the **filebaseddiscovey.ps1** probe parameter.
   c. Copy the contents of the **value** field into a file.
   d. Make the necessary changes to the file and then remove the (old) signature block at the end of the file.
   e. Re-sign the script: https://docs.microsoft.com/en-us/powershell/module/microsoft.powershell.security/set-authenticodesignature?view=powershell-6.

      After re-signing, the file should have a new signature block at the end of it.

   f. Copy the entire file contents and replace the old value in the **filebaseddiscovery.ps1** probe parameter from step b.

      Windows File-based Discovery should now run using the updated signed script.

   If you're using anti-virus software, make sure to include all of the ServiceNow Windows PowerShell scripts so the anti-virus software does not see it as a threat.

2. Activate the File-based Discovery plugin.
3. Enable and configure File-based Discovery.
   Using the Discovery Configuration Console, you can enable File-based Discovery and configure several components to find and manage all of your installed software.

   To avoid impacting performance for existing customers, File-based Discovery is turned off by default.

4. Configure File-based Discovery scans.
   After Discovery runs and returns file information for a configuration item (CI), it will not execute File-based Discovery again on that target until the interval has expired. Since there is a performance cost when File-based Discovery performs scans, it is important to determine

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

how frequently to scan. However, choosing a more frequent interval than Monthly is not recommended due to performance considerations.

a. Using the Discovery Configuration Console, expand **Common** and then enter the maximum number of files that you want to discover.

> **Note:** Increasing this value can impact performance.

b. Select the **Frequency** that File-based Discovery runs on the CI.
The default is set to Monthly.

> **Note:** File-based Discovery does not trigger until initial Content Data Service (CDS) synchronization occurs. CDS synchronization could take approximately 24 hours for the initial set of data to be synchronized from CDS. For more information see, File Signature Normalization.

Result
Your File-based Discovery should run based on the configurations set.

Related topics

File-based Discovery references

File-based Discovery references

Review this reference material for further details on File-based Discovery.

How File-based Discovery minimizes performance impact

Scanning large networks, particularly for Windows computers, can extend scan time considerably. Moreover, file scanning can create heavy CPU demands on target machines and produce large payloads that greatly increase processing time. To avoid performance issues, File-based Discovery allows you to optimize processing with many of the settings configured using the Discovery Configuration Console.

- **Multi-phase scan**: Probes are triggered in multiple phases to reduce long scan times. The first probe creates the scanning process on the target, which runs in the background. Subsequent probes check the progress of this scan each time Discovery runs and return the results when the scan is complete. For example, if a file scan is initiated via Quick Discovery or from a Discovery schedule, no information is returned if a background scan is still running. If the scan is finished the next time the server is discovered, the probe returns the data.

  > **Important:** The multi-phase scan that File-based Discovery uses does not return results immediately. Do not expect any results from a file-based scan until Discovery runs again on those targets and fetches the list of files discovered by the previous scan.

- **Scan folders**: You can limit the scope of the scan to particular folders as well as scan hidden folders.
- **Scan extensions**: You can retrieve all files for a given list of extensions as well as ignore specific file extensions when scanning.
- **Scan UNIX file names**: The number of unique file names used to identify installed software on UNIX is relatively small. The complete list of file names is included in the UNIX probe, which filters discovered files against the signature list on the target. No additional filtering is required until the results are returned to the instance.
- **Scan Windows file names**: For Windows File-based Discovery, the probe does not initiate a scan on target for full signature list. Instead, Discovery generates a list of file extensions that

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

the probe attempts to match on the target. This reduced list is then filtered on the MID Server against the entire Windows signature list.

- **Reduced payloads**: By filtering the scan results with the file signature list before the results get to the instance, Discovery reduces the size of the payload from target servers. Only those files from the targets that match the files specified in this list are returned to the instance.
- **Scan throttling**: To reduce the impact on Windows CPU and I/O performance on remote targets, Windows File-based Discovery is configured to sleep occasionally during scans. You can fine-tune this behavior with Discovery properties. Throttling Windows scans can make those scans take longer.

File-based Discovery tables

Descriptions of the tables provided by the File-based Discovery [com.snc.discovery.file_based_discovery] plugin for storing data:

File Information [cmdb_file_information]

This table stores all the files, matched or unmatched, that were discovered on a particular CI. The files are available in a related list on CI computer records. The following columns are provided in the default view:

- **Path**: Path in which the file was found on the target.
- **Product**: Software product name. This field is populated when SAMP is activated and a match exists.
- **Publisher**: Name of the publisher of the file, such as Oracle or Microsoft. This is a reference to the Company [core_company] table. This field is populated when SAMP is activated and a match exists.
- **Size**: Size of the file.
- **Software installation**: Name of the installed software associated with this file, from the Software Installation [cmdb_sam_sw_install] table. This field is populated when SAMP is activated and a match exists.
- **Version**: Version of the file. Populated for Windows files only.

Unidentified File Set [cmdb_unidentified_file_set]

This table stores all the files returned by Discovery that cannot be matched to a software product, with or without SAMP. Discovery populates the name, size, and version of each unmatched file in the record. These optional fields can be used to further identify the file and make it available for future file matching:

- Publisher
- Product
- Version
- Edition
- Platform
- Language

Note: A feature in Software Asset Management Professional (SAMP) allows signature data in this table to be shared with Content Data Services (CDS). When you enable data sharing, the system automatically uploads records from the Unidentified File Set [cmdb_unidentified_file_set] table to CDS when you provide the missing information. Uploading the identified files makes them available to improve future file matching. For details, see Add a custom software product.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a full reference of File-based Discovery system properties and script includes, see the File Based Discovery Now Support Knowledge Base article.

File-based Discovery issue resolution

If you're experiencing any issues when setting up or running File-based Discovery, see File-based Discovery issue resolution.

Related topics

File-based Discovery

File-based Discovery issue resolution

If you have any issues while setting up or running File-based Discovery, follow the actions listed here to help resolve them.

Issue	Cause/Action
Error message: File Discovery data not yet synced...	• Allow Content Data Service (CDS) time to synchronize the samp_file_* tables. Initial synchronization can take approximately 24 hours. • Review the Data Services: Download Schedules (cds_client_schedule) table to check if the timestamp has been updated for “samp_file_*” tables.
File-based Discovery is not triggered during discovery.	• Feature has not been enabled. Check the Discovery Definition Configuration Console and verify that File-based Discovery is enabled. • Interval between scans is too long. Check the interval setting on the Discovery Definition Configuration Console.
File-based Discovery is not triggered during discovery.	Check the Internal properties managed by the Discovery [discovery_private_properties] table for the following items. To view this table, enter **discovery_private_properties.list** in the navigation filter and press **Enter**. • File_discovery.file_extensions_windows should contain a list of Windows file extensions to be scanned. • File_discovery.unix_filename_timestamp should be something other than 1970 timestamp. • File_discovery_windows_filename.timestamp should be something other than 1970 timestamp. If either timestamp shows 1970, CDS content hasn't downloaded yet. To trigger

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Issue	Cause/Action
	an immediate download, navigate to the Data Services: Download Schedules (cds_client_schedule) table and search for records where the **Table** field is `samp_file_name`, `samp_file_map`, or `samp_file_set`. Open each record and select **Execute Now**. To verify the download, navigate to **System Logs > Outbound HTTP Requests** and confirm the **Response status** is 200 for requests where the **URL hostname** is `sncdataservices.service-now.com`.
Errors in MID Server log from FileDiscoveryFilter or MID Server issue created.	• Check if the File-based discovery inclusion list exists. The inclusion list is in the MID Server installation folder in `/agent/work`. If it does not exist, the synchronization did not succeed. • Review the MID Server record on the instance to see if a MID Server issue was created. • MID Server uses the GetMIDInfo scripted soap service to synchronize the inclusion list. See if errors are generated from this script. • Select the **File discovery filter refresh** Related Link in the MID Server record to synchronize the file names to this MID Server.
PowerShell errors	• Include the script in your anti-virus software. • Add ServiceNow publisher to your trusted publisher repository.

Related topics

File-based Discovery references

Configure MID Server parameters for PowerShell 7

Configure two MID Server parameters to run File-based Discovery and Windows ADME on target hosts that use PowerShell 7.

Before you begin

Confirm the following:

- The instance is on the Brazil version of the ServiceNow AI Platform.
- PowerShell 7 is installed on the MID Server.
- PowerShell 7 is installed on target hosts.
- On each target host that runs PowerShell 7, *`DisableImplicitWinCompat`* is set to `true` in the PowerShell configuration file. By default, PowerShell 7 loads certain modules through a background Windows PowerShell 5.1 session, which blocks Discovery scripts from running on

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PowerShell 7. For more information, see the "About Windows PowerShell Compatibility" article on the Microsoft PowerShell documentation site.

- File-based Discovery or Windows ADME is active on the instance.

Role required: agent_admin or admin

About this task

By default, File-based Discovery and Windows ADME use PowerShell 5 on target hosts. To run these features on hosts that use PowerShell 7, create two parameters on the MID Server. Setting *mid.powershell.prefer_7_plus* controls the MID Server itself. Setting *mid.powershell.prefer_7_plus_remote* controls the target host. Each parameter defaults to `false` if not set, so any host where the parameter is not configured runs on PowerShell 5.

Procedure

1. Navigate to **All** > **MID Server** > **Servers**.
2. Select the MID Server.
3. In the **Configuration Parameters** tab, select **New**.
4. Create the required parameters in the following table.

Parameter name	Value	Description
mid.powershell.prefer_7_plus_remote	`true`	Prefer PowerShell 7 or later, if available, on remote target hosts. Default: `false`.
mid.powershell.prefer_7_plus	`true`	Prefer PowerShell 7 or later, if available. Default: `false`.

5. Select **Submit** after each.

Result

File-based Discovery and Windows ADME run using PowerShell 7 on the configured target hosts.

Related topics

File-based Discovery

Discover applications based on fingerprints

Discover applications based on suggestions based on ServiceNow Predictive Intelligence. Predictive Intelligence automatically classifies and categorizes the discovered running processes, as application fingerprints, and provides suggestions. ITOM Visibility uses Predictive Intelligence to perform initial analysis of discovered processes and suggest applications that you might want to discover. When using this method, ITOM Visibility automatically creates a Configuration Management Database (CMDB) configuration item (CI) class, a classifier, or a pattern for the new application CI class.

Before you begin

Perform horizontal discovery on all relevant IP ranges as described in Convert IP networks into Discovery range sets. Provide all the relevant credentials for ITOM Visibility access to the applications and devices that you want to discover.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Discovering applications using fingerprints is especially useful when you can't use patterns available by default or from the ServiceNow Store. For example, when you want to discover non-standard, home-grown applications or brand new applications.

The Discovery feature of ITOM Visibility uses application dependency-mapping (ADM) to detect processes running on TCP ports and discover applications that run these processes. When you discover applications based on fingerprints, you take it a step further. The fingerprint-based discovery method analyzes processes that ADM hasn't used for discovery.

Fingerprint-based discovery uses algorithms to identify running processes and organizes them into groups. These process groups become suggested applications or candidates. You review suggested applications and choose which ones to discover. The list of suggested applications is refreshed every hour.

Patterns created during the fingerprint discovery are used for both horizontal and top-down discovery. For Service Mapping to run such a pattern correctly, define when to run this pattern. You do that by defining before or after which other pattern Service Mapping should run this new pattern.

When you discover an application based on an application suggestion, the fingerprint-based discovery performs the following steps:

1. The fingerprint-based discovery creates a CMDB CI class for the new application. The new CI class extends the Application [cmdb_ci_appl] CI class and by default uses the process group name for the suggested CI class display name.
2. The fingerprint-based discovery creates a process classifier for the application group you chose to discover. ADM uses the new process classifier to trigger the new pattern created for the CI class. The priority of the classifier is lower to ensure that the system first uses the default patterns available with ITOM Visibility or custom patterns.
3. The fingerprint-based discovery creates a discovery pattern of the application type for the new CI class. **Enforce Process Classification** is enabled to make sure that the MID Server runs this pattern only if the process identified on a CI matches the classification criteria for this pattern.

   Note: The Service Mapping user interface refers to CI classes as CI types.

4. ADM uses the classifier to trigger the horizontal discovery of the application CI.
5. ADM triggers discovery using the new pattern.
6. After the pattern completes discovery, it writes the discovered application CI into the relevant CMDB table.

By default, the fingerprint-based discovery is enabled.

Procedure

1. Navigate to **All** > **Discovery** > **Home**.
2. Click **Application Fingerprints**.
   The **Application Fingerprints** dashboard opens.
3. Review suggested applications in one of the following ways:
   - On the **Top Suggestions** pie chart, view the top suggestions arranged by the number of processes making up these suggestions. The process count is reflected by the size of the chart segment. Click the segment to see its corresponding suggestions in a filtered list. To remove a suggestion from the chart, click the suggestion name in the chart legend.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To view only suggestions for application servers, click the **Suggested Application Servers** tile. The filtered list of suggestions for application servers is displayed.
- Alternatively, click **All suggestions**. The unfiltered list of all suggestions is displayed.

4. Click the relevant suggestion in the **Suggested Group Name**, for example **python**.
   The suggestion details open in the **Application Suggestion** form.
5. Review the suggestion details, paying attention to the following parameters:

Field	Description
Suggested Group Name	The name ITOM Visibility generates for the suggestion. This field is pre-populated, but you can modify it if necessary.
Extended Group Name	The top 10 key words that provide more information about the process. Predictive Intelligence extracts these key words from the process.
Suggested CMDB CI class	The name of the CMDB CI class that the ADM creates, if you choose to discover the suggested application.
Process Count	The number of discovered processes that this suggestion is based on.
View Processes	Click to view sample processes that this suggestion is based on. The overall number of processes may be larger and starts from one thousand. The number of sample processes you can view by default is one hundred.
Accuracy	A validation algorithm checks the quality of the prediction. The algorithm analyzes the regular expression (ReGex) parameter of the processes in this suggestion and other suggestions (application groups). Pick suggestions with higher rule accuracy. If the accuracy is low, you may adjust the process ReGex definition manually, as described in Adjust regular expression used for fingerprint-based discovery. **Note:** Do not modify the accuracy value manually.
Process Classifier	The suggested regular expression used in the classifier. If defined, ADM uses this regular expression during horizontal discovery.
Ignore Rule	The name of the process that was excluded from discovery using the **Ignore** action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. **Optional:** If this suggestion is not relevant and you do not want to use it for discovery, perform one of the following actions:
   - To exclude processes from horizontal discovery, click **Ignore**.

     This process is added to the Process Handlers [discovery_proc_handler] table with the *`Classify parameter`* set to false. Discovery stops checking the classifier for the excluded process and does not discover the CI on which these processes run.

   - To remove this suggestion from the suggestion list, click **Delete**.

     The suggestion disappears from the list. If Predictive Intelligence adds new processes to a deleted suggestion, it reappears in this list.

7. Click **Discover Application**.
   The fingerprint-based discovery discovers an application CI for the selected suggestion.
8. On the **Application Fingerprints** dashboard, check that the discovered application appears under **Created Applications**.

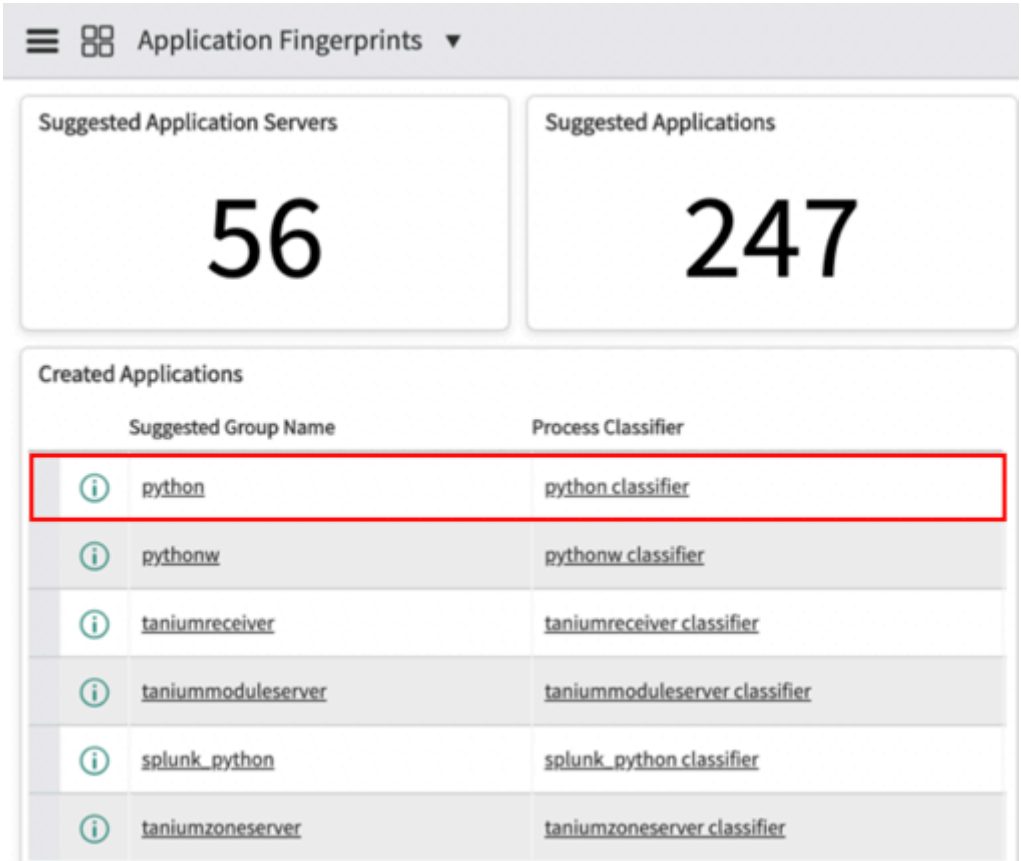

9. Configure the running order for the automatically created pattern:

   a. Click the process classifier next to the new application.
   The Process Classification form opens.

   b. On the Triggers probes tab, click the pattern link.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

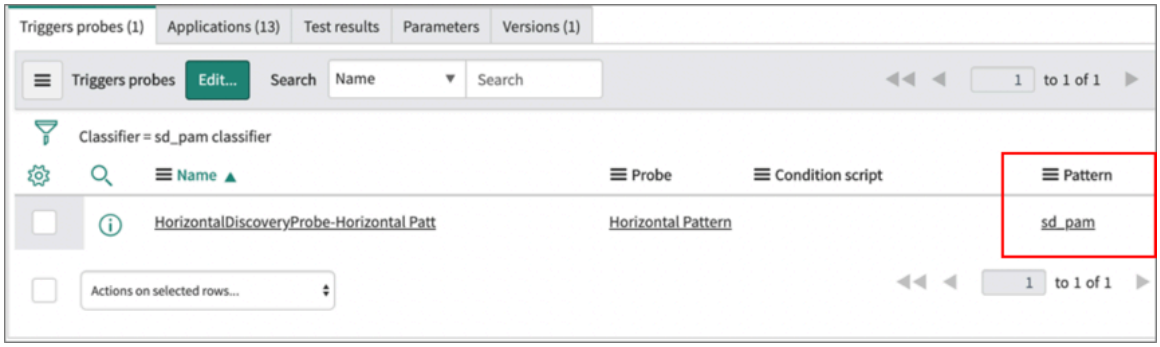

The pattern form opens.

c. Select the order in which this pattern always runs:

- **Before**
- **After**

d. Then select the other applicable pattern.

> **Note:** If you are not sure what the right order is or you are not planning on using this automatically generated pattern for top-down discovery, configure the run order to **After All**.

e. Click **Save**.

10. **Optional:** Click the process classifier next to the new application to see the classifier and the pattern that the fingerprint-based discovery creates for this application CI.
11. **Optional:** If the discovery result is not satisfactory, calibrate the fingerprint-based discovery feature.

Related topics

Predictive Intelligence

Application Fingerprints dashboard

Application Fingerprints dashboard

ITOM Visibility uses Predictive Intelligence to perform initial analysis of discovered processes and suggest applications that you might want to discover. When using this method, ITOM Visibility automatically creates a Configuration Management Database (CMDB) configuration item (CI) class, a classifier, or a pattern for the new application CI class. Use the Application Fingerprints dashboard to review suggested applications.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

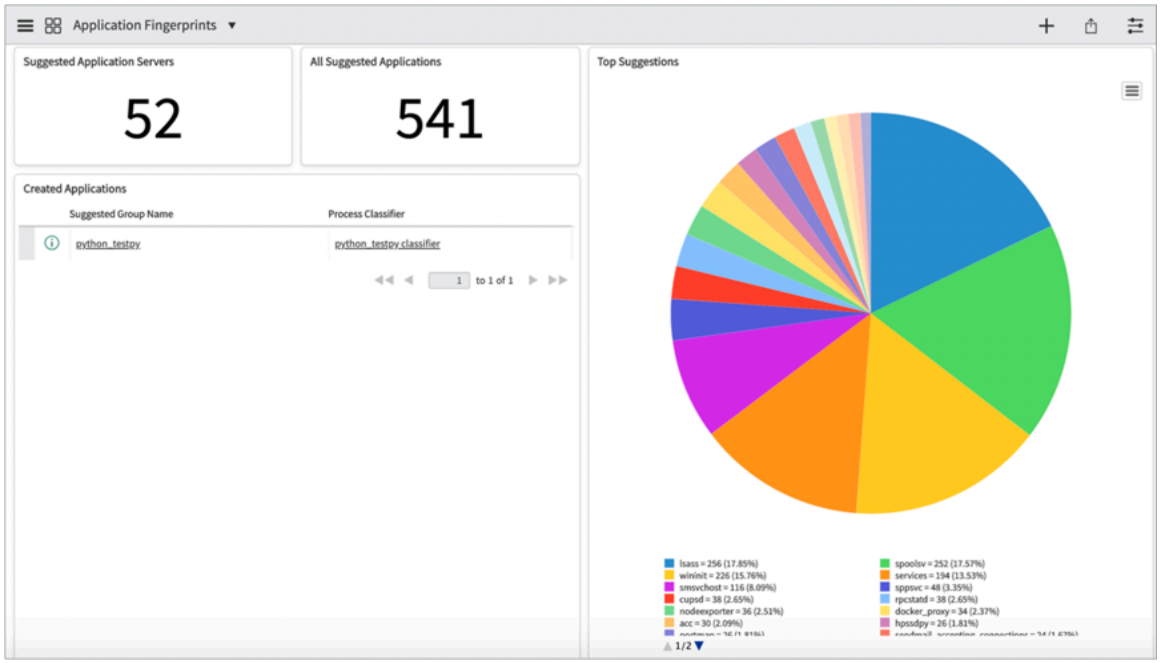

Required ServiceNow AI Platform roles

The discovery_admin, who runs horizontal discovery using Discovery.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
Administrator performing horizontal discovery of the organization IT network.	Review application suggestions and decide which applications to discover based on fingerprints. Performing this action enhances the results of the basic horizontal discovery.

Indicators

Suggested application servers

The filtered list of suggestions for application servers.

Suggested applications

The complete list of suggestions for applications.

Created applications

The list of applications that the system discovered based on fingerprints.

Data visualizations

Title	Type	Source table	Description
Top Suggestions	Pie Chart	Application Suggestion [cmdb_process_groups]	A pie chart report that shows the top suggestions arranged

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
			by the number of processes making up these suggestions. The process count is reflected by the size of the chart segment. Click the segment to see its corresponding suggestions in a filtered list. To remove a suggestion from the chart, click the suggestion name in the chart legend.

Discover applications using application fingerprints or ITOM Content Service

Discover applications using suggestions based on ServiceNow Predictive Intelligence or ServiceNow ITOM Content Service to create configuration items.

Before you begin

Role required: discovery_admin

About this task

Review application suggestions based on application fingerprints or ITOM Content Service and add them to discovery.

- Fingerprint-based discovery uses algorithms to identify running processes and organizes them into groups. These process groups become suggested applications or candidates. You review suggested applications and choose which ones to discover. The list of suggested applications is refreshed every hour. To learn more, see Discover applications based on fingerprints.
- ITOM Content Service identifies the products that you're currently using and delivers new product suggestions every week. To learn more, see ITOM Content Service.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Content 360**.
2. Select an application suggestion to add to discovery based on either application fingerprints or ITOM Content Service.

Suggestion source	Actions
Application fingerprints	**a.** Select the **AI suggested application finger prints** widget. **b.** Select the suggestion under **Suggested Group Name**. **c.** On the Application fingerprints suggestion form that is displayed, verify the informa tion. For information about the form fields, see Application fingerprint suggestion form. **d.** Determine whether to add the suggestion to discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Suggestion source	Actions
	▪ If the suggestion isn't relevant and you don't want to add it, select **Ignore**. ▪ If you want to add it, select **Add to discovery**.
ITOM Content Service	**a.** Confirm that you've opted in to ITOM Content Service. **Note:** If the **Opt-in** option is displayed, you have not yet opted in. For more information, see Share data on ITOM Content Service. **b.** Select the **Content service suggestions** widget. **c.** Manually sync with suggestions from your system by selecting **Sync candidates**. **Note:** If a message appears about the application scope, select **OK** to exit the message. Then, change the scope by selecting the application scope icon ( ) and selecting **Global**. **d.** In the Sync candidates confirmation window, select **Yes**. **e.** Review the list and select the name of the relevant suggestion to see more detailed information. **Note:** You can select multiple suggestions to add to discovery at the same time. Select the check box next to each suggestion and then select **Activate**. **f.** Determine whether to add the suggestion to discovery. ▪ If the suggestion isn't relevant and you don't want to add it, return to the list of suggested applications. ▪ If you have made changes to the suggestion but do not want to add it to discovery, select **Save**. ▪ To add the suggestion to discovery select **Activate**.

Calibrate fingerprint-based discovery

Fine-tune discovering applications based on processes, if the discovery results are not satisfactory.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Modify default settings for fingerprint-based discovery if this feature does not generate any suggestions or does not generate. In addition, you may also fine-tune the generated suggestions to make them relevant and useful.

Verify fingerprint-based discovery generates suggestions

Fingerprint-based discovery relies on Predictive Intelligence for generating suggestions for discovery. If the **Application Fingerprints** dashboard does not display any suggestions, verify that the Predictive Intelligence is configured correctly.

Before you begin

Confirm that fingerprint-based discovery remains enabled, since it is enabled by default. To verify, navigate to **All** > **System Properties** > **All Properties**. Check that the *sys_property process.clustering.appfingerprint.enabled* property does not appear on the list.

Role required: discovery_admin

About this task

If the **Application Fingerprints** dashboard does not display any suggestions, perform the following steps. After every action, check if the issue persists.

Procedure

1. Verify that Predictive Intelligence is installed and active.

   a. Navigate to **System Definition** > **Plugins**.

   b. Search for **Predictive Intelligence**.

   c. Verify that the Predictive Intelligence plugin is installed.

   If Predictive Intelligence is not installed, select the **Install** button.

   d. Select the **Predictive Intelligence** tile.

   e. On the Predictive Intelligence form, verify that the **Status** is **Active**.
   If the status is **Inactive**, select the **Activate/Repair** link under **Related Links**.

2. Verify that the scheduled job controlling the fingerprint-based discovery, is enabled and configured correctly.

   a. Navigate to **System Definition** > **Scheduled Jobs**.

   b. Select the *Applications suggestion - ITOM Autodisco* scheduled job from the list.

   c. Confirm that the **Active** check box is selected.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Check that the discovery frequency is set to one hour as shown in the following example:

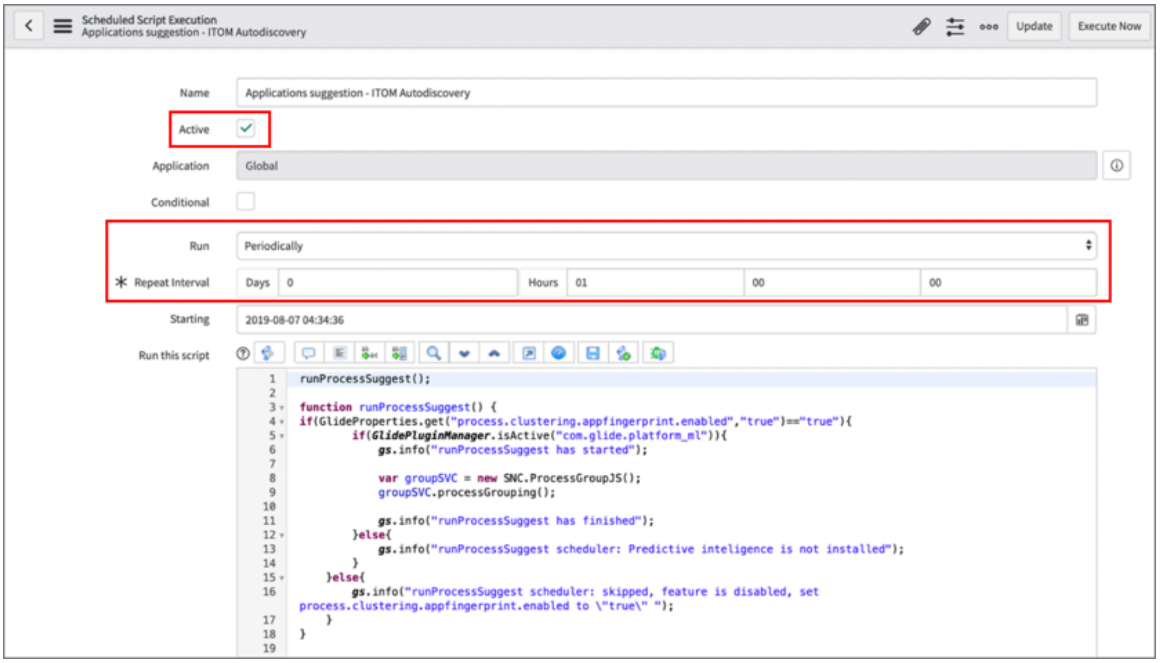

3. Check that the clustering solution responsible for grouping processes into suggestions is active.

a. Enter `ml_solution.list` in the **Filter navigator** field of the **Navigation** panel.
The **ML Solutions** tab displays all solutions, including clustering solutions.

b. Identify clustering solutions by checking the **Table** column.
The system saves clustering solutions in the Virtual running processes [v_cmdb_running_process] table.

c. If there are several clustering solutions, identify the latest solution by the timestamp in the **Update** column.
Alternatively, you can use the number that the system appends to the solution name to identify the latest solution.

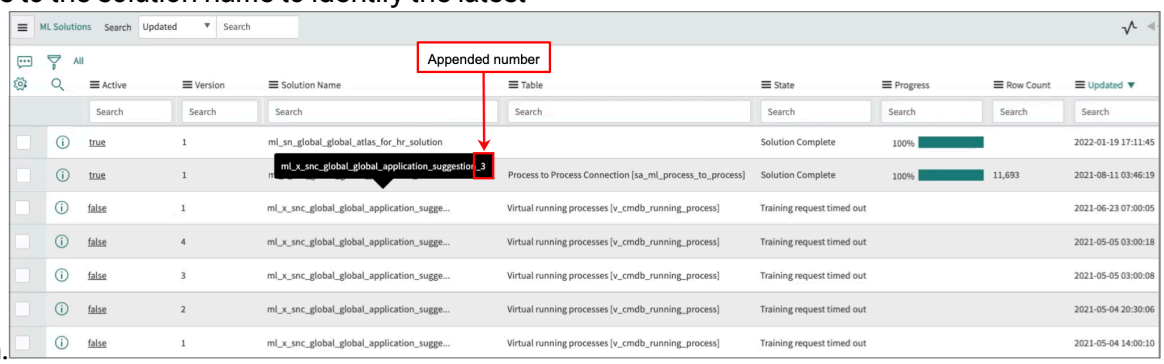

d. Verify that the **State** value is **Solution Complete** and **Progress** is at **100%**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

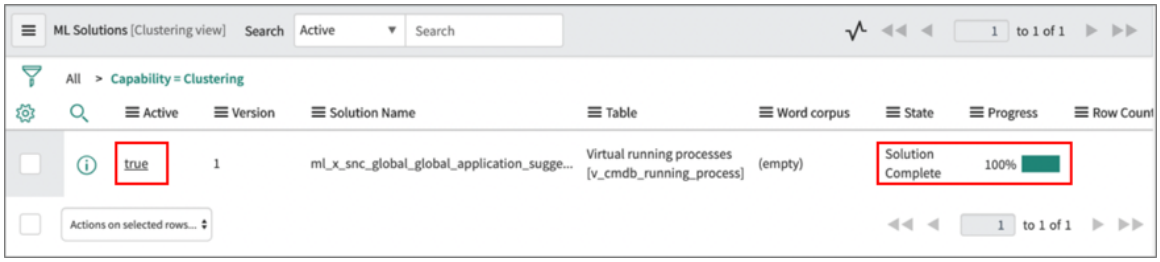

e. If the State value is not Solution Complete, configure Predictive Intelligence as described in Configuration tips for Predictive Intelligence.

4. Check that the clustering solution created enough process groups, which become suggestions.
   a. Navigate to **System Definition** > **Tables** and locate the Virtual running processes [v_cmdb_running_process] table.
   b. Check that there are more than 1000 records in this table.
   c. If there are less than 1000 processes, discover more processes by running horizontal discovery on additional IP ranges or subnets.
   d. If there are still less than 1000 processes after discovering additional IP ranges or subnets, disable the fingerprint-based discovery.
5. Check that the **Application suggestion - ITOM Autodisco** transaction is not slow-loading, as described in View and kill active transaction.
   If this transaction is running for a long time, kill it.

Fine-tune suggestions for fingerprint-based discovery

Adjust definitions of process-based suggestions before starting discovery. Typically, you need to fine-tune automatically generated suggestions if the accuracy is less than 100%. Predictive Intelligence creates suggestions and provides the accuracy estimate for each suggestion. The lack of accuracy may result in the failure to discover required applications or the discovery of wrong applications in addition to the correctly discovered applications.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Home**.
2. Select **Application Fingerprints**.
   The **Application Fingerprints** dashboard opens.
3. Navigate to the list of suggestions by one of the following methods:
   - On the **Top Suggestions** pie chart, view the top suggestions arranged by the number of processes making up these suggestions.
   - To view only suggestions for application servers, select the **Suggested Application Servers** tile. The filtered list of suggestions for application servers is displayed. This list is also filtered to display only suggestions whose accuracy is 100%.
   - Alternatively, select **All suggestions**. The unfiltered list of all suggestions is displayed.
4. In the Application Suggestions list, select the relevant suggestion in the **Suggested Group Name** column.
   The Application Suggestion form opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

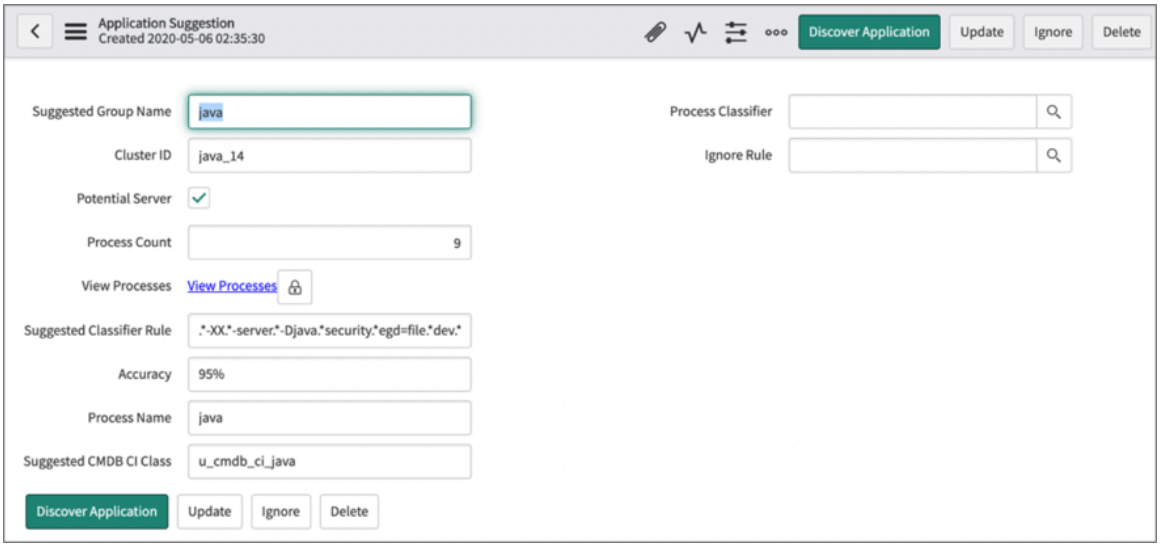

5. To change the name used for the fingerprint-based application, enter the preferred name in the **Name** field.
   For example, many applications run java processes and the suggested group name for these groups is **java**. You may want to assign a name that better describes the application that you want to discover using this suggestion (for example, `java_myweblogic`).
6. To change the name used for the CI class, enter the preferred CI class name in the **Suggested CMDB CI Class** field.
7. To fine-tune the classifier rule, change the regular expressions in the **Suggested Classifier Rule** field.
   For example, add version information to the regular expression for discovering a Tomcat application: `.*tomcat.* -Dopss.*version=12.1.3.*`
8. Click **Update**.
9. If you are not sure how to change the classifier rule, fine-tune it after performing the initial discovery as described in Adjust regular expression used for fingerprint-based discovery.

Adjust the regular expression for fingerprint discovery

Fine-tune the regular expression in the classifier rule to make discovery results more precise. Predictive Intelligence extracts regular expressions from processes and assigns them to classifier rules as part of creating suggestions. If you're not satisfied with the discovery results or the suggestion accuracy is lower than 100%, modify the regular expression.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Application** > **Module**.
2. Select **Discover Application**.
   The fingerprint-based discovery finds the application and creates a CI class, a classifier, and a pattern.
3. On the **Application Suggestions** list, select the newly created process classifier.
   The Process Classification form opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

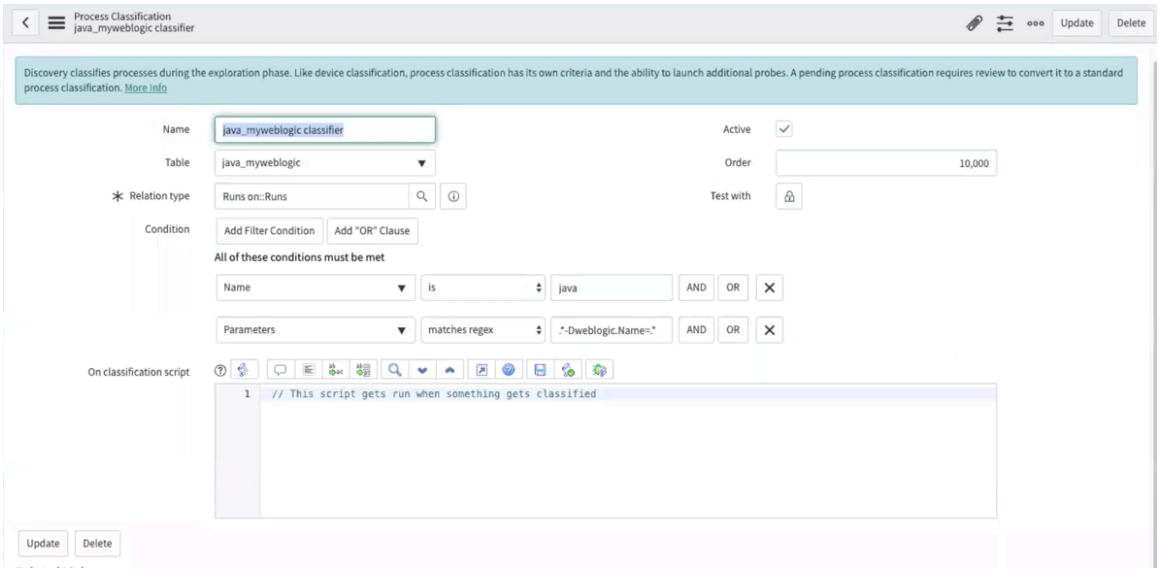

4. In the **Parameters** value field under **Condition**, notice the regular expression automatically selected by Predictive Intelligence.
   In this example, it is **.*"-DweblogicName=.*"**
5. Select **Padlock** icon next to **Test with**.
6. Select the **Look up** icon to select the server on which you want to test additional regular expressions.
   The list of all discovered servers opens.
7. Select the sever that runs the application that you want to discover using the customized classifier.
   The selected server appears in the **Test with** field.
8. Right-click the header of the Process Classification form and select **Save**.
   The system runs a test and displays the results on the **Test results** tab at the bottom of the page.

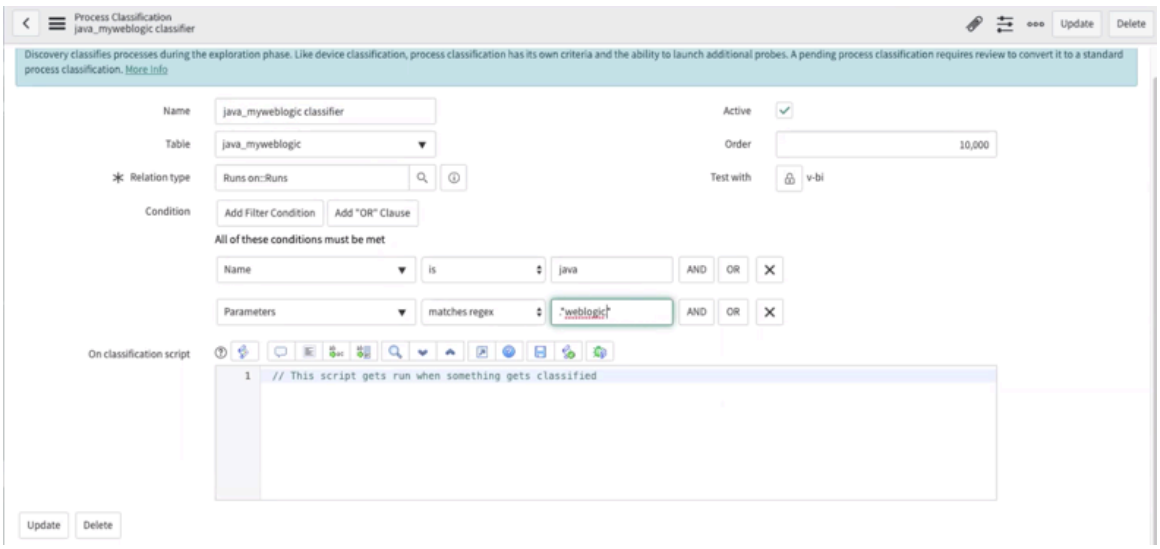

9. On the Test results tab, check the processes that were discovered using the regular expressions configured for this classifier.
10. Adjust the regular expression in the Parameters value field under **Conditions**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, make it more generic: `"weblogic."`

11. Right-click the header of the Process Classification form and select **Save**.
    The system runs the test again and displays the results on the **Test results** tab at the bottom of the page.
12. On the Test results tab, verify that the fingerprint discovery found the applications as you expected.

Disable fingerprint-based discovery

By default, the fingerprint-based discovery is enabled. Disable application discovery based on processes, if Predictive Intelligence does not generate enough processes.

Before you begin

Role required: discovery_admin

Procedure

1. Add the system property *`sys_property process.clustering.appfingerprint.enabled`* as described in Add a system property.
2. To disable, enter `false` in the **Value** field.
3. To enable the fingerprint-based discovery again, change the **Value** field to `true`.

Related topics

KB0826261: Advanced configuration and troubleshooting of fingerprint-based discovery

Troubleshooting Guide for ML Connection Suggestions

Network discovery

Network Discovery discovers the internal IP networks and subnets within your organization.

Note: If you already know the IP address ranges in your network, it's not necessary to run Network Discovery. It's intended for organizations that don't have complete knowledge of the IP addresses available for Discovery in their networks.

How network discovery works

Discovery uses the information it gathers to update routers and Layer 3 switches in the Configuration Management Database (CMDB). A single MID Server begins its scan on a configurable list of starting (or seed) routers performs network Discovery. The starting routers are the default routers used by all the MID Server host machines in the network, but can be any designated routers. The MID Server uses the router tables on the starting routers to discover other routers in the network. The MID Server then spreads out through the network, using router tables it finds to discover other routers. It continues until the MID Server explored all the routers and switches.

After running Network Discovery, convert the IP networks it finds into IP address Range Sets that you use in Discovery schedules to discover configuration items (CI). Configure SNMP credentials. Port 161 must be open for SNMP access.

Network Discovery properties

The default values are correct for most discoveries. Navigate to **Discovery Definitions** > **Properties** to edit the properties.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
glide.discovery.bgp_router_disable	BGP router exploration disable: Disables running the **SNMP – Routing** probe or pattern when discovering a router running the BGP protocol. This exploration is inactive by default because BGP routing tables are large. These routers typically operate at the edge of large networks, where further network discovery is not relevant. Set this value to false if your organization uses BGP routers as edge routers between relatively small networks. For example, between buildings on a single campus. If you must populate the [discovery_device_neighbors] table during horizontal layer 2 discovery of the bgp-enabled devices, set the **BGP router exploration disable** property to **No**. Notice that enabling this property can cause performance issues including out-of-memory issues on the MID Server. • **Type**: true | false • **Default value**: true
glide.discovery.debug.network_discovery	Network discovery debugging: Enables extensive logging of all Network Discovery activities on the instance. • **Type**: true | false • **Default value**: true • **Learn More**: Discovery monitoring and issue resolution
glide.discovery.discoverable.network.max.netmask.bits	Maximum netmask size for discoverable networks (bits): The maximum number of bits in a regular netmask for networks that are discovered by Network Discovery. A regular netmask is a netmask that can be expressed in binary as a string of ones followed by a string of zeroes. For example, 255.255.255.0 is regular and 255.255.255.64 is irregular. Regular networks are commonly expressed like: 10.0.0.0/24, which means a network address of 10.0.0.0 with a netmask of 255.255.255.0. Larger bit numbers mean networks with smaller numbers of addresses in them. For example, the network 10.128.0.128/30 has four addresses in it: one network address (10.128.0.128), one broadcast address (10.128.0.131), and two usable addresses

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	(10.128.0.129 and 10.128.0.130). Small networks are commonly configured in network gear to provide loopback addresses or networks used strictly by point-to-point connections. Since these sorts of networks generally don't need to be discovered by Network Discovery, it would be useful to filter them out. By setting this property to a value of 1 through 32, you can limit the sizes of regular networks that are discovered. Setting it to any other value causes all networks to be discovered. Irregular networks are always discovered. The default value is 28, which means that regular networks with 8 or fewer addresses cannot be discovered. • **Type**: integer • **Default value**: 28
glide.discovery.network_owner_method	Network router selection method: This property controls the method used to decide (during Network Discovery) which router should be selected as the router to be associated with a given IP Network. The possible values are: • **First Router**: The first router that discovers the network is associated • **Last Router**: The last router that discovers the network is associated • **Most Networks**: The router with the most attached networks is associated • **Least Networks**: The router with the least attached networks is associated • **Type**: choice list • **Default value**: Most Networks • **Additional options**: First Router, Last Router, Least Networks
glide.discovery.physical_interface_types	Physical interface types: A comma-separated list of interface types that are considered physical for the purposes of network discovery. If a router (or device capable of routing) has an interface of this type, the networks connected to that interface are considered locally connected to that device. The default interface types include Ethernet, 802.11, and Token Ring types. Interface type numbers are defined in the SNMP MIB-2, specifically in OID 1.3.6.1.2.1.2.2.1.3.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• **Type**: string • **Default value**: 6,117,9,71,209
glide.discovery.switch_interface_types	Switch interface types: A comma-separated list of interface types that are considered. Interface type numbers are defined in the SNMP MIB-2, specifically in OID 1.3.6.1.2.1.2.2.1.3. Devices with any interface types that do NOT appear in this list are classified as routers, when they have routing. A complete list of the interface type numbers may be found on the IANA web site, in the section "ifType definitions". • **Type**: string • **Default value**: 7,8,9,26,53,62,69,71,78,115,117,209

Learn

What is network discovery?

Run network Discovery

Configure network Discovery from a Discovery schedule.

Before you begin

Configure SNMP credentials, and keep port 161 open for SNMP access.

Role required: none

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Click **New** and select **Networks** from the list in the *Discover* field.
3. Select a MID Server.

   This field is mandatory.
4. Complete the form, including the scheduling fields.
5. Right-click in the header bar and select **Save** from the pop-up menu.

   The **Related Links** and related lists appear.
6. Click **Network Discovery Private IPs** in Related Links to view the list of default private IP networks in the *Discovery IP Ranges* Related List.

   The default IP networks in this list are available to every Network Discovery you conduct and are sufficient for most discoveries.
7. If your organization has additional private IP addresses, click **New** to add them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Private IP addresses

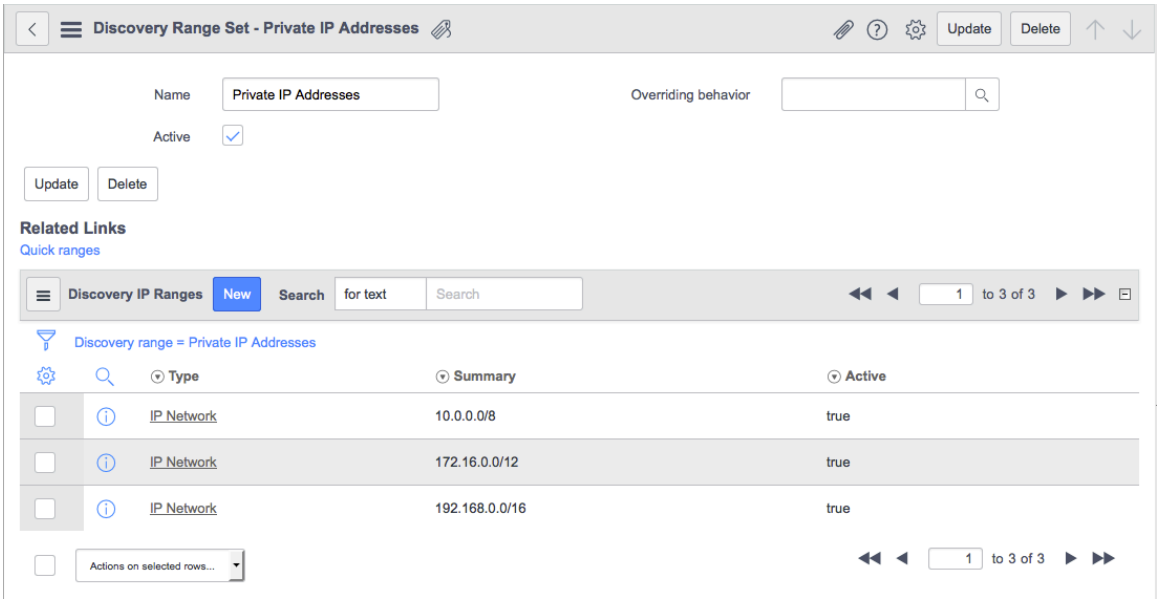

8. Add starting routers to the schedule in the *Discovery Range Sets* list.
   a. Click the **Network Discovery Auto Starting Routers** link to populate the list with the starting router for each MID Server in your network.
   b. Click **Edit** to add or delete routers from the list.

Network Discovery schedule

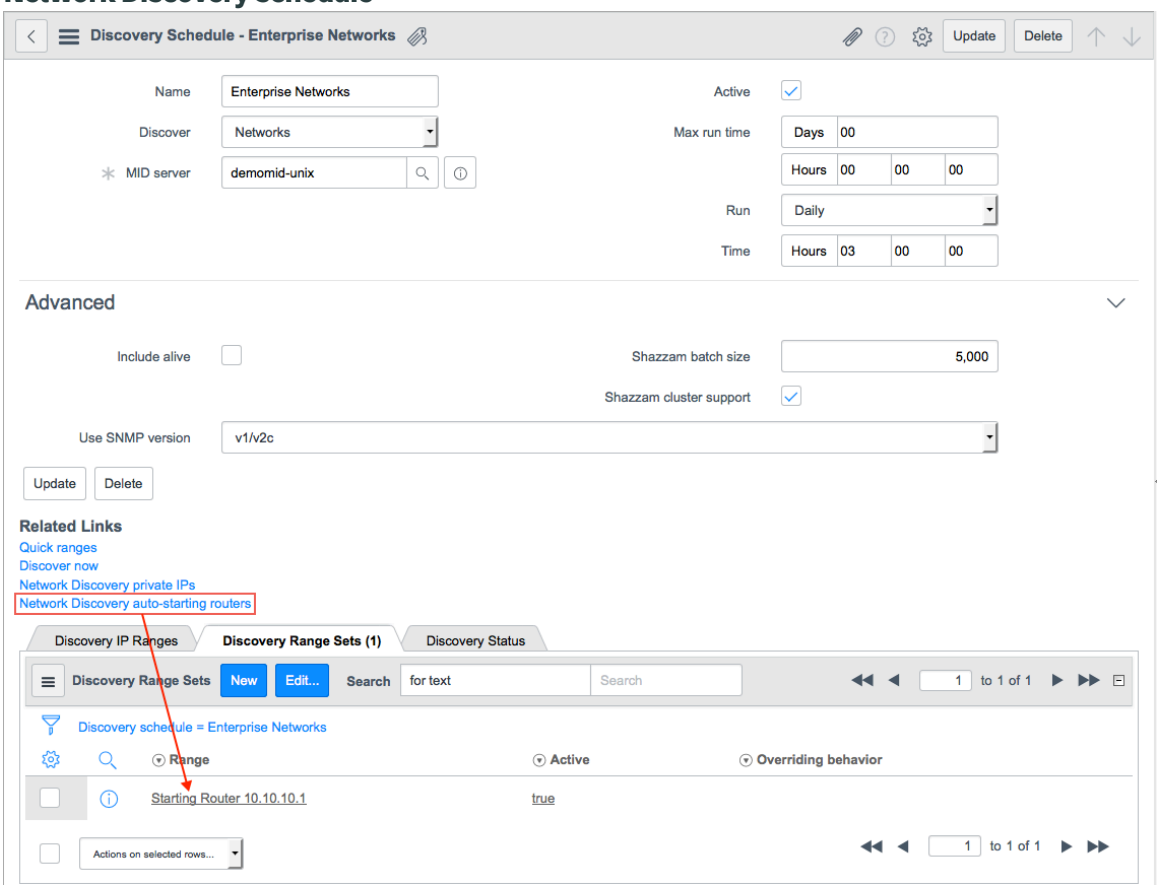

9. Run Discovery manually, or through the scheduler.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Convert IP networks into Discovery range sets

After you conduct a network Discovery, you must convert the IP networks that were found into range sets for use in discovering other devices.

Before you begin

Role required: agent_admin or admin

Procedure

1. Navigate to **All** > **Discovery** > **Output and Artifacts** > **IP Networks**.
2. Click **New** to add an IP network.
3. Ensure that the following field values are provided:
   - **State**: Ready
   - **Discover**: selected (true)
   - **Router**: Select a router from the list. This field must not be empty.

   If your IP Networks were created through network Discovery, then these fields are populated automatically in the *IP Networks* related list. However, if you entered the IP Networks manually, and you want to convert your IP Networks into range sets manually, you must edit these fields accordingly.

IP network form

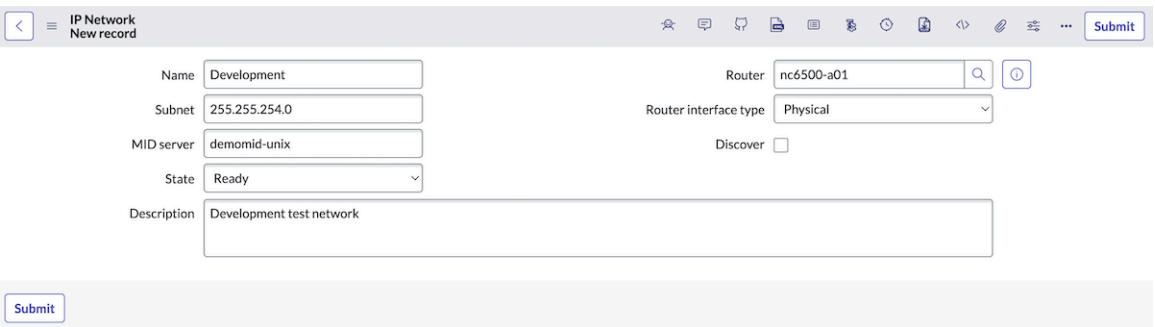

4. Select **Submit** to return to the list of IP Networks.
5. In **Related Links**, select **Create Range Sets**.

   Use this feature to convert all the IP networks in the list into range sets.

IP network list

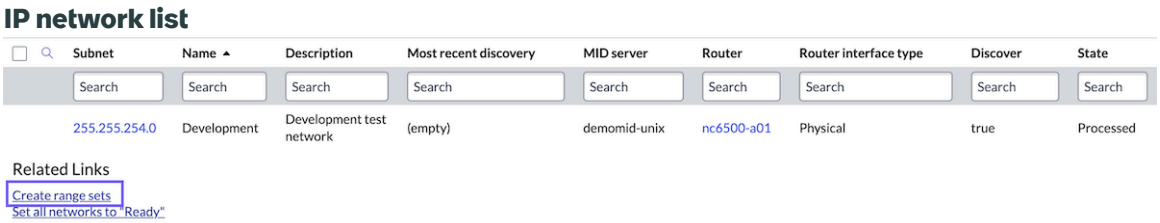

   The Discovery Status page appears, displaying the progress of the conversion. The system increments the **Started** and **Completed** count of IP networks, until all the networks are converted.

What to do next

Add discovery range sets to Discovery schedules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add a Discovery range set to a Discovery schedule

After creating discovery range sets from a network discovery, add the sets to a discovery schedule. This process narrows down the range of IP addresses that Shazzam probes can scan. You can also add a behavior.

Before you begin

Role required: agent_admin or admin

Procedure

1. Open a Discovery schedule.
2. In the **Discovery Range Sets** related list, click **Edit** to add an existing range set and use the slushbucket to add the range set.
3. On the **Discovery schedule** form, click the name of the range set under the **Range** column in the **Discovery Range Sets** related list.
4. On the **Discovery Range Set** form, add a Discovery behavior if necessary.
   You can also activate or deactivate the range set for this schedule.
5. Right-click the header, and then click **Save**.

Serverless Discovery

Discovery can find applications on host machines without the need to discover the host first. This type of Discovery is referred to as serverless Discovery.

Serverless Discovery relies on infrastructure patterns to explore CIs on a host. This kind of Discovery skips the scanning and classification phases of Discovery. There are two types of serverless Discovery: standard and host-based. Both types require an infrastructure pattern.

Standard serverless Discovery

Standard serverless Discovery finds remote application CIs that reside on an unknown host. This type of serverless Discovery uses protocols, such as REST or HTTP, to find the CIs. The IP address of the host machine is not required. You define input parameters in an infrastructure pattern to connect to a remote resource. Use this type of Discovery when you want to discover web-based applications.

Host-based serverless Discovery

Host-based serverless Discovery finds CIs that reside on a host machine that you already discovered. The host acts as a proxy through which you find other CIs. You do not actually run Discovery on the host. Use this type of Discovery when:

- You do not know the IP address of the CIs you want to discover, or the CIs do not actually have IP addresses.
- The CIs reside on an already discovered host CI.

For this type of Discovery, you also need an infrastructure pattern, but you do not need to define input parameters like you do for standard serverless Discovery. The pattern can specify operations, such as a WMI Query, which executes queries on a host. After Discovery connects to the host, it uses the credentials of the host to perform the Discovery of CIs you are exploring in the pattern. You do not need to specify additional credentials. When you run debug mode on the pattern, you can see the details under **Temporary Variables** on the pattern. You can then create pattern operations that make API calls from the host to the CIs you are trying to discover.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Serverless Discovery schedules

The infrastructure patterns that serverless Discovery uses are triggered from a Discovery schedule, rather than from a classifier. Serverless Discovery ignores classifiers. Use a serverless schedule type for both types of serverless Discovery. The serverless schedule allows you to specify:

- Values for input parameters in the pattern if you are using standard serverless Discovery. See Create a schedule for standard serverless Discovery for instructions.
- A host CI if you are using proxy-host based Discovery. See Create a schedule for host-based serverless Discovery for more information.

Create pattern and schedule for standard serverless discovery

To find application CIs that reside on an unknown host, create an infrastructure pattern for standard serverless discovery. After creating the pattern, create a Discovery schedule that triggers the pattern.

Before you begin

Role required: cloud_admin, service_mapping_admin, or admin

About this task

In the pattern, you must specify input parameters to explore the CI.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Click **New**.
3. Specify an **Infrastructure** type pattern with the correct CI type that applies to the CIs you want to discover.
   See Create or customize patterns for an explanation of all form fields.
4. Create an identification section and save the pattern.
5. Click the **Input Parameters** tab, and then click **New**.
6. Enter the parameters that are required to reach the serverless CIs.
   For example, you could enter `url` to reach a remote resource via an HTTP request.

   Note: You specify the actual URL in a serverless discovery schedule.

7. Save the pattern.
8. Click the **Basic** tab, and then click the identification section name to open the pattern designer.
9. Create the steps you need to explore serverless CIs.
   For example, create an **Http Get Call** operation to reach a remote resource through a URL. In this case, you would specify the `$url` parameter in the **URL** field.
10. To test the pattern, click **Debug Mode** and fill in values for the input parameters.
    You do not need to enter a value for **Proxy Host**. In the this example, we are connecting to a website that serves up an XML file. The XML file parsing method is already defined in the Http Get Call operation. To test the pattern operation, click **Debug Mode**, enter the URL, and click **Connect**.
11. After you connect and run the operation, the **Output** field displays the content.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a schedule for standard serverless Discovery

After you create a pattern for standard serverless discovery, create a Discovery schedule and specify the pattern. Discovery triggers the pattern from the schedule rather than from a classifier.

Before you begin

Ensure you have basic authentication credentials for HTTP calls in the pattern.

Role required: cloud_admin, service_mapping_admin, or admin

About this task

You can reuse this schedule to run the same pattern using different parameters.

Procedure

1. Navigate to **All** > **Discovery** > **Schedules**.
2. Click **New**.
3. Select **Serverless** in the **Discovery** field.
   See Schedule a horizontal discovery for explanation of other fields.
4. Right-click the form header and select **Save**.
5. In the **Serverless Execution Pattern** related list, click **New**.
6. Fill out the form fields (see table):

Field	Description
Name	Provide a descriptive name.
Pattern	Select the pattern you created for standard serverless discovery or for proxy host-based discovery.
Proxy Host	Leave this field empty. This field applies to the host proxy-based scenario.
Active	Keep this check box selected to use the pattern execution definition.

7. Right-click the form header and select **Save**.
8. In the **Discovery Pattern Launcher Parameters** related list, add the parameters required by the pattern and provide a value for each one in the **Value** column.
9. Click **Submit**.
10. Define as many serverless execution pattern records as needed.
    You can launch the same pattern, but they should have different parameters.
11. Execute discovery by clicking **Discover Now** or wait until the schedule triggers the discovery.

Result

Check the Discovery status to see the Discovery log, the devices list, which holds the discovered CIs, and the ECC queue records. Note that the Shazzam probe did not launch and Discovery did not trigger any classification probes.

Related topics

Discovery status

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create pattern and schedule for host-based serverless discovery

To find application CIs that reside on a known host in your CMDB, create an infrastructure pattern for host-based serverless discovery. After creating the pattern, create a Discovery schedule that triggers the pattern.

Before you begin

Role required: cloud_admin, service_mapping_admin, or admin

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Select **New**.
3. Specify an **Infrastructure** type pattern with the correct CI type that applies to the CIs you want to discover.
   See Create or customize patterns for an explanation of all form fields.
4. Create an identification section and save the pattern.
5. In the **Input Parameters** tab, add at least one input parameter and value.
   For example:
   - Parameter: `ip_address`
   - Value: IP address of the host
6. Add operations to the pattern to discover the CI that you can access via the host.
7. To test the pattern, select **Debug Mode** and enter the IP address of the host.

Create a schedule for host-based serverless Discovery

Discovery triggers the pattern for host-based, serverless Discovery from the schedule rather than from a classifier.

Before you begin

Ensure you have credentials for the host device that acts as the proxy.

Role required: cloud_admin, service_mapping_admin, or admin

Procedure

1. Navigate to **All** > **Discovery** > **Schedules**.
2. Click **New**.
3. Select the **Serverless** in the Discovery field.
   See Schedule a horizontal discovery for explanation of other fields.
4. Right-click the form header and select **Save**.
5. In the **Serverless Execution Pattern** related list, click **New**.
6. Fill out the form fields (see table):

Field	Description
Name	Provide a descriptive name.
Pattern	Select the pattern you created for host-based discovery.
Proxy Host	If you are using host-based discovery, specify an existing CMDB CI. Discovery uses the credentials of this CI when exploring the CIs on the host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Active	Keep this check box selected to use the pattern execution definition.

7. Right-click the form header and select **Save**.

   Define as many serverless execution pattern records as needed. You can launch the same pattern, but they should have different parameters.

   Note: You do not need to create entries in the **Discovery Pattern Launcher Parameters** related list.

8. Execute discovery by clicking **Discover Now** or wait until the schedule triggers the discovery.

Result
Check the Discovery status to see the Discovery log, the devices list, which holds the discovered CIs, and the ECC queue records. Note that the Shazzam probe did not launch and Discovery did not trigger any classification probes.

Discovery probes and sensors

Discovery probes and sensors perform data collection and update the Configuration Management Database (CMDB).

With each release, patterns are replacing many probes and sensors for Discovery. Consider creating new patterns or editing existing ones if you want to customize what Discovery can find. The information on probes and sensors is intended for customers who are not yet using patterns or who have customized probes retained upon upgrade. See Patterns and horizontal discovery for more information on patterns.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Discovery phases

Discovery always uses probes and sensors during the first two phases of discovery: scanning and classification. For the last two phases, identification and exploration, Discovery can use probes and sensors or patterns. This topic refers to probes and sensors only. See Exploring Discovery for an explanation of these phases. See Patterns and horizontal discovery for more information on patterns.

Probes, sensors, and the ECC queue

The probe collects the information and the sensor processes it. Both get their instructions from the ECC queue. There is a worker job on the MID Server that monitors the queue for work. The monitor checks for any entries where the *Queue* is *output* and the *State* is *ready*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

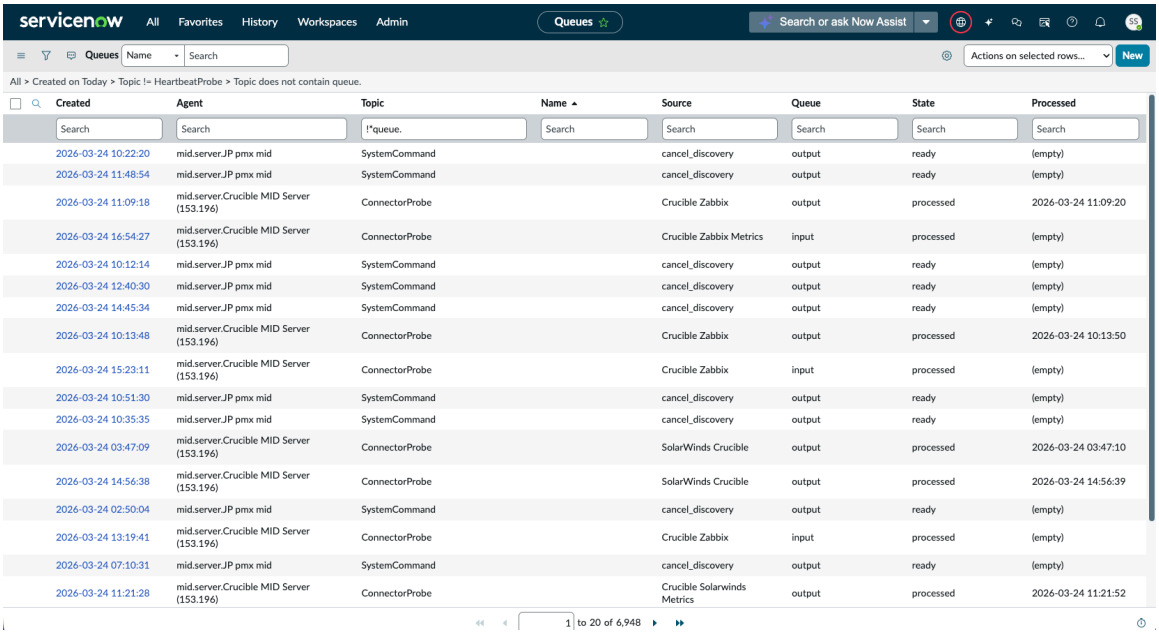

The MID Server then processes all the output ECC messages, runs the necessary probes, and returns the probes results to the ECC queue. These results are put in the ECC queue as *input* entries.

ECC queue input

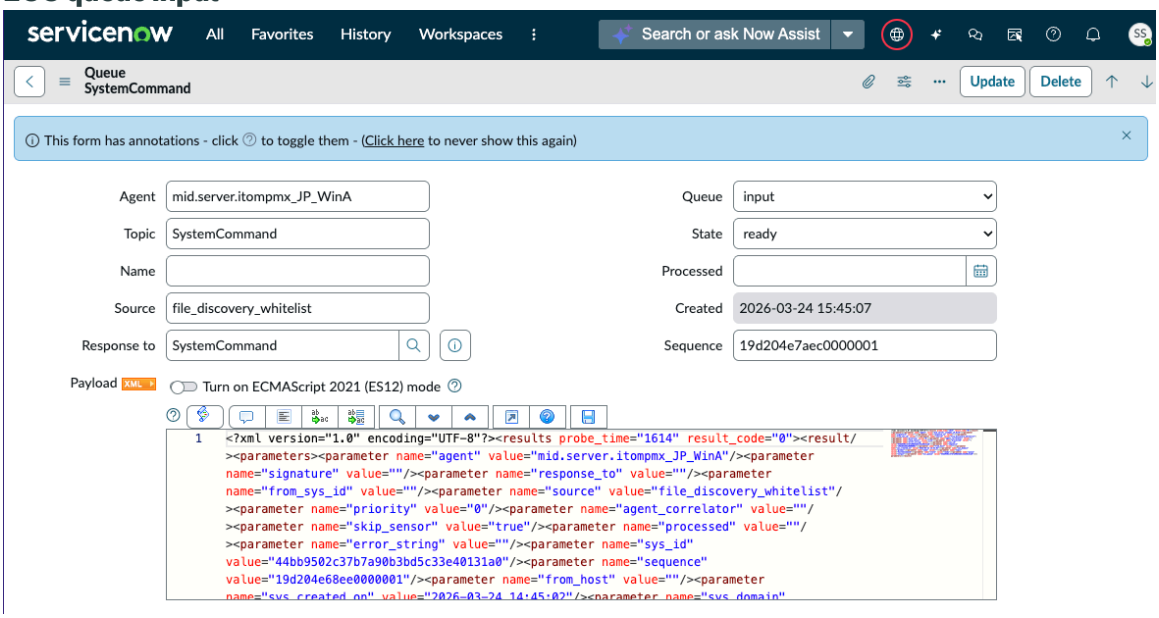

After an entry is inserted in the ECC Queue table, a business rule fires (on insert) that takes that information and runs it through a *sensor processor*. The sensor processor's job is to take the input data, find any sensors interested in that data, and pass it along to be processed. Those sensors ultimately update the CMDB.

Important: Discovery now enforces code signing for probes, probe parameters, and sensors when code signing is enabled on the instance. Unsigned or tampered payloads are blocked, and signatures are validated before execution. For more information about code signing, see Discovery on Code Signing instances.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How probes and sensors work together

The MID Server launches probes to collect information about a device. The probe sends back information to the sensor to be processed. If the probe has a post-processing script defined, the post-processing script does some data processing on the MID Server before data is sent back to the sensor on the ServiceNow instance. Otherwise the probes sends back all the data collected and the sensor performs this data processing. In both cases, the sensor updates the CMDB.

A multi-probe is a probe that contains probes. A multi-sensor processes the data from a multi-probe. To process the data from the multi-probe, the multi-sensor contains individual scripts to process the data returned by each probe contained in the Multiprobe, as well as a main multi-sensor script. The individual scripts pass their processed data to the main multi-sensor script.

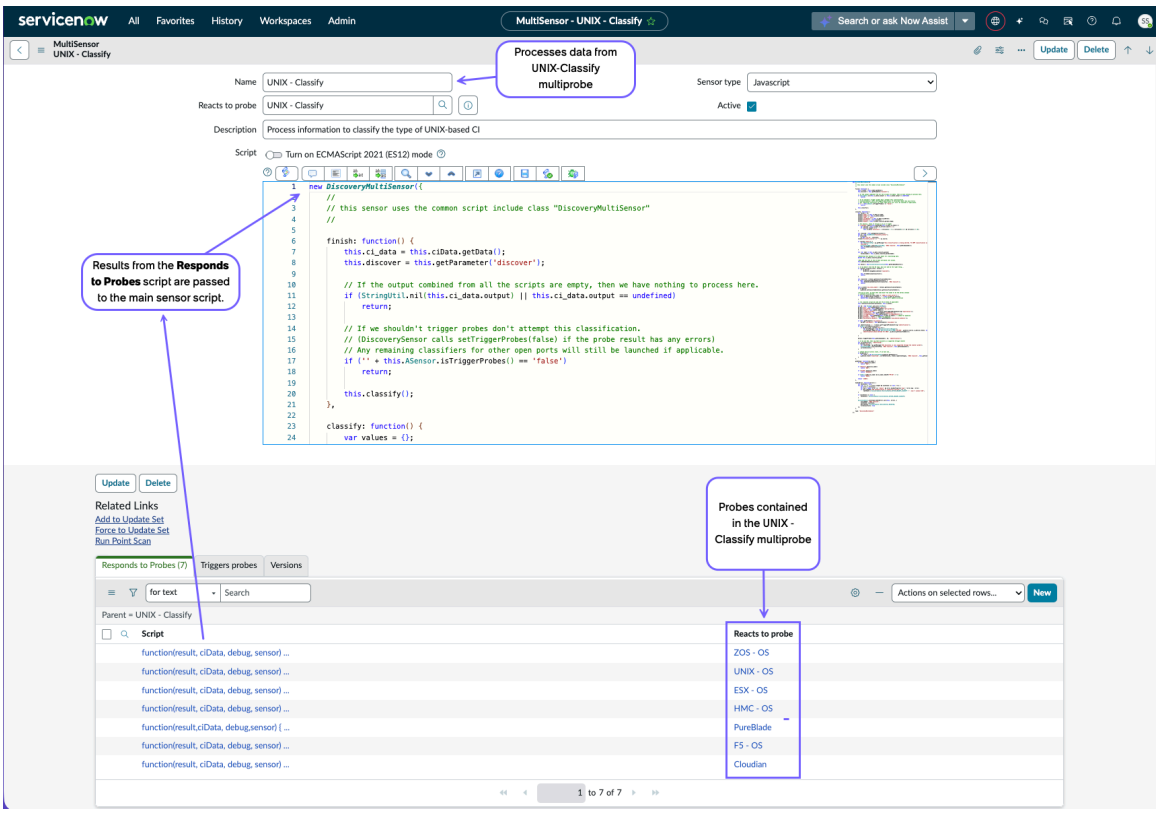

Probe types

Device	Probe Type
Windows computers and servers	Remote WMI queries, shell commands
UNIX and Linux servers	Shell command (via SSH protocol, version 2). Discovery supports any Bourne-compatible shell.
Storage	CIM/WBEM queries
Printers	SNMP queries
Network gear (switches, routers, etc.)	SNMP queries
Web servers	HTTP header examination

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Device	Probe Type
Uninterruptible Power Supplies (UPS)	SNMP queries

List of Discovery probes

A wide variety of probes exist for the Discovery application to detect elements on your network.

These probes are included with the Discovery application.

The PPP script designation in the table indicates whether the probe includes a probe post-processing (PPP) script that runs on the MID Server. The PPP script transforms probe results into a JSON string and returns the string to the ServiceNow instance for sensor processing.

To view probes and their descriptions, navigate to **Discovery Definition** > **Probes**.

CIM probe

The CIM probe uses WBEM protocols to query a particular CIM server, the CIM Object Manager, for a set of data objects and properties.

For instructions on configuring probe parameters, see Set probe parameters.

The following parameters may be passed to the CIM probe:

CIM Probe

Parameter	Description	Default Value
source	[Required] The initial host to connect to.	None
port	The port to connect to. If empty, the value is determined by the "schema" parameter: http = 5988, https = 5989.	5988 or 5989
schema	[Required] The schema to use: '"http"' or '"https"'.	http
namespace	[Required] The CIM namespace. May be overridden by a query.	None
queries	[Required] A semicolon-delimited list of CIM probe queries to process and return results for.	None
retries	The number of times to retry a query if it fails due to network connectivity issues.	2
connection_timeout	The number of milliseconds the probe has to connect to a server.	5000
socket_timeout	The number of milliseconds the probe has to read data.	5000

The CIM Intermediate Query Language (CimIQL) uses keys, filters, and dot-walking to traverse the CIM schema.

Parameter Expansion

The CIM query language supports standard SNC preprocessed probe parameter expansion. Place variables in queries by encapsulating their names like this:

```
${foobar}.CIM_RunningOS[0].Name
CIM_ComputerSystem.${barfoo}
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The text ${foobar} is replaced with the contents of the foobarprobe parameter passed to the CIM probe; likewise for barfoo.

CIMIQL

The CIM Intermediate Query Language (CimIQL) is an intermediate language designed to simplify the process of querying CIM providers.

CimIQL currently supports the standard Web-Based Enterprise Management (WBEM) protocol stack, but others, such as Web Services-Management (WS-MAN), may be added in the future. The query language syntax borrows from elements of Microsoft's WMI query language and UNIX's wbemcli command. The CimIQL library is a pure Java implementation.

Note: CimIQL is pronounced "simicle".

CimIQL syntax

CimIQL syntax consists of several elements, including a query and different tokens.

CimIQL Syntax Element Descriptions

Element	Description
Statement	The most basic element of CimIQL is a valid . A statement contains multiple queries delimited by a period . character.
Query	A represents a single high-level protocol-independent request. Each query is comprised of nested language components and sub-components known as tokens.
Token	A describes a specific lexical aspect of the CimIQL syntax.
Operation Token	The first token of each query must be an , which represents the overall logical operation to be performed.
Component Token	A is a sub-component of an operation token.
Result	Each query is paired with a result, which is then provided as input to the next query in the statement. A is comprised of a set of objects and their properties.

CimIQL operation tokens

The CimIQL probe requires operation tokens.

Each of the following core operations has a counterpart in the CIM Operations over HTTP standard.

CimIQL Operation Token Summary

Return Value	Details	Equivalent CIM Operation over HTTP
class object	Get Object Retrieves a single object of a specific class by specifying all of its unique keys (as key tokens) and any optional parameter tokens, separated by commas.	GetInstance
class object	Enumerate Objects	EnumerateInstances

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CimIQL Operation Token Summary (continued)

Return Value	Details	Equivalent CIM Operation over HTTP
	Retrieves objects that match a set of condition tokens and parameter tokens.	
class object	Enumerate Associated Objects Retrieves objects associated with each result from the preceding query.	Associators
statement results	Substitution A no-op token that feeds the results of a previous named statement as input into the next query of its own statement.	Reference the results of a named statement

Get Object Token

<classname>{<key token>,<parameter token>,...}

- Retrieves a single object of a specific class by specifying all of its unique keys (as key tokens key tokens) and any optional parameter tokens, separated by commas. This token is also known as the .
- The <classname> is the case-sensitive CIM class name of the desired object. By default, objects of the specified class and of any extended classes are retrieved.
- The key and parameter tokens are enclosed by a single pair of curly brackets { ... }.
- This token must only be used as the first query in a statement.
- Returns: class object
- Example:

```
CIM_ComputerSystem{CreationClassName='Linux_ComputerSystem',Name='runtime'}.*
```

Enumerate Objects Token

<classname>{{<condition token>,<parameter token>,...}}<array index token> OR
<classname><array index token>

- Retrieves objects that match a set of condition tokens and parameter tokens. This token is also known as the .
- The condition tokens and parameter tokens are enclosed by two pairs of curly brackets {{ ... }}. The curly brackets are optional if there are no conditions or parameters necessary.
- The <classname> is the case-sensitive CIM class name of the desired objects. By default, objects of the specified class and of any extended classes are retrieved.
- The index token is optional.
- This token must only be used as the first query in a statement.
- Returns: class object
- Example:

```
CIM_ComputerSystem{{Name!='runtime'}}.*
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Enumerate Associated Objects Token

<association classname>{{<property filter token>,<parameter token>,...}}<array index token> OR <association classname><array index token>

- Retrieves objects associated with each result from the preceding query.
- The condition tokens and parameter tokens are enclosed by two pairs of curly brackets {{ ... }}. The curly brackets are optional if there are no properties filters or parameters necessary.
- The <association classname> is the name of the many-to-many or one-to-many class that associates two objects together. By default, objects of the specified class and of any extended classes are retrieved.
- The <parameter token>, ResultClass, may be specified to filter results based on the resulting object's classname.
- The index token is optional.
- This token must not be used as the first query in a statement.
- Returns: class object
- Example:

```
CIM_ComputerSystem{{Name='runtime'}}[2].*
```

Substitution Token

${<statement name>}

- A no-op token that feeds the results of a previous named statement as input into the next query of its own statement.
- Returns: void
- Example:

```
$(lastComputer).ElementName
```

CimIQL component tokens

The CimIQL probe requires component tokens, which are sub-components of operational tokens.

The following tokens are sub-components of operation tokens.

CimIQL Component Token Summary

Token	Details
Properties token	* OR <property name>,<property name>,... Specifies which properties are to be returned for each object of the final result set.
Query Delimiter Token	. (Period) Separates queries.
Index Token	[index] Reduces a preceding query's results to a single object at the specified integer index.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CimIQL Component Token Summary (continued)

Token	Details
Key Token	<key name>='<value>' Matches an object property designated as a key by exact value.
Condition Token	<property name><conditional operator><enclosed value> Matches a single property of an object based on the condition specified.
Parameter Token	<parameter name>:'<value>' Passes a parameter by <parameter name> to the operation being called. The parameter may be consumed during CimIQL pre-processing or by the CIMOM via request, depending on the parameter.

Properties Token

* OR <property name>,<property name>,...

- Specifies which properties are to be returned for each object of the final result set.
- The wildcard * returns all properties available. Otherwise, each property name desired is provided within a comma-separated list.
- This token is required at the end of each statement.
- Example:

```
CIM_ComputerSystem[0].*
```

Query Delimiter Token

. (Period)

- Separates queries.
- Example:

```
CIM_ComputerSystem.PrimaryOwnerContact
```

Index Token

[index]

- Reduces a preceding query's results to a single object at the specified integer index.
- This token is always optional.
- Example:

```
CIM_ComputerSystem[0].*
```

Key Token

<key name>='<value>'

- Matches an object property designated as a key by exact value.
- The <key name> is the name of the property used as a key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Example:

```
CIM_ComputerSystem{CreationClassName='Linux_ComputerSystem',Name='runtime'}.*
```

Condition Token

<property name><conditional operator><enclosed value>

- Matches a single property of an object based on the condition specified.
- The <property name> is the name of the property to match against.
- The <conditional operator> determines how the property's actual value is compared to its expected value. The operators available are *equality* (=) and *inequality* (!=).
- The <enclosed value> should be one of the following:
  - Literal value enclosed in single-quotes ' ... '. For example, foo='bar'
  - Regular expression, enclosed by a pair of slashes / ... /. For example, foo=/bar.*/
- Example:

```
CIM_ComputerSystem{{Name!='runtime'}}.*
```

Parameter Token

<parameter name>:'<value>'

- Passes a parameter by <parameter name> to the operation being called. The parameter may be consumed during CimIQL pre-processing or by the Common Information Model Object Manager (CIMOM) via request, depending on the parameter.
- Example:

```
CIM_ComputerSystem.CIM_RunningOS{{ResultClass:'Win32_ComputerSystem'}}.*
```

CimIQL tutorial

This is a tutorial by example where each example builds on the previous example.

CimIQL Tutorial

Order	CimIQL Statement	Result
1	CIM_ComputerSystem[0].*	Retrieves the firs of all instances CIM_ComputerS and its descend Retrieves all pro
2	CIM_ComputerSystem.PrimaryOwnerContact	Retrieves all instances of CIM_ComputerS and their desce Retrieves only one property, PrimaryOwnerC
3	CIM_ComputerSystem{CreationClassName='Linux_ComputerSystem',Name='runtime'}.*	Retrieves a sing unique instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CimIQL Tutorial (continued)

Order	CimIQL Statement	Result
		CIM_ComputerS and its descend All key tokens m specified within identity token.
4	CIM_ComputerSystem{{Name!='runtime'}}.*	Retrieves all inst and descendan CIM_ComputerS that do not have property of 'runt The filter token { filters out instan that do not cont of the properties specified.
5	CIM_ComputerSystem{{Name=/^run.*$/}}.*	Retrieves all inst and descendan CIM_ComputerS that have a valu matching the re expression cont within the / / ch **Note:** The expression not require quotations. The filter token { filters out instan that do not cont of the properties specified.
6	CIM_ComputerSystem{{Name='runtime'}}[2].*	Retrieves the se result of all insta CIM_ComputerS and its descend where the instar have a property 'runtime'. The order of ope follows the quer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CimIQL Tutorial (continued)

Order	CimIQL Statement	Result
		1. Query server CIM_Compute and descenda 2. Filter results b Name propert 3. Retrieve the s instance that the filter.
7	CIM_ComputerSystem.CIM_RunningOS[0].Name	Retrieves the Na property for the CIM_OperatingS instance of each CIM_ComputerS instance. The middle-toke CIM_RunningOS name of the Ass class, not the en
8	CIM_ComputerSystem.CIM_RunningOS{{Name=/CentOS/}}[0].Name	Retrieves the Na property for the CIM_OperatingS instance of each CIM_ComputerS instance, where CIM_OperatingS instance has a N property contain 'CentOS'.

CimIQL results

CIM Probe results are passed to the probe sensor as an XML document embedded within the <output> element.

The following is a commented example of a CimQuery batch result.

```
<!-- document root -->
 <cimqueryset>
 <!-- A single query and query result. Multiple <cimquery> tags
 may be provided. -->
 <cimquery>
   <!-- The original query, enclosed by CDATA. -->

 <query><!CDATA[[>CIM_ComputerSystem[0].PrimaryOwnerContact<! ]]
></query>
   <!-- The resulting data is enclosed within a single <result>
 tag. -->
   <result>
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    <!-- A single class instance result. Multiple <instance>
 tags may be provided.
         Special tags are prefixed with an underscore
 character. -->
    <instance>
      <!-- The instance's CIM classname -->
      <_classname>Linux_ComputerSystem</_classname>
      <!-- A set of this instances identifying keys. Always
 provided, regardless of property filters.
           Within here, each key is provided as
 <KeyName>VALUE</KeyName> with the VALUE enclosed as CDATA. -->
      <_key>

 <CreationClassName><![CDATA[Linux_ComputerSystem]]></CreationCl
assName>
        <Name><![CDATA[runtime]]></Name>
      </_key>
      <!-- Each property that matches the query's property
 filter will be provided here, in the same format as keys;
           As <PropertyName>VALUE<PropertyName> where VALUE is
 enclosed as CDATA -->

 <PrimaryOwnerContact><![CDATA[root@runtime]]></PrimaryOwnerCont
act>
    </instance>
  </result>
 </cimquery>
 </cimqueryset>
```

DNS probe

DNS probes determine the DNS names for configuration items (CI).

The default DNS Name Resolver probe is a lightweight and fast probe. It achieves its speed by executing in parallel on the MID Server. The probe utilizes DNS lookup by the MID Server host OS to make use of the maintained DNS cache. The probe also leverages the hosts file on the OS.

These performance improvements can be disabled by reverting to the legacy DNS probe with either MID Server properties or parameters. The names of the MID Server properties/parameters are [mid.probe.use_legacy_forward_dns] and [mid.probe.use_legacy_reverse_dns]. For more information, see MID Server parameters and MID Server properties.

DNS probe definitions

Probe	Description
DNS	Determines DNS names for devices. It has the parameter ip_addresses.
DNSLookupProbe	Resolves the IP address of a host.
DNSNameResolver	Resolves the DNS name of an IP address.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Port probes

Probe	Description
dns	Triggered by 'dns'. Uses 'DNS Name Resolver' Scanner and uses classification 'Process Classification [discovery_classy_proc]'

Horizontal Pattern probe

Discovery uses the Horizontal Pattern probe to launch patterns for horizontal discovery.

The Horizontal Pattern probe works with the Horizontal Discovery sensor to enable Discovery to use patterns for discovery. When you see messages in the ECC Queue from this probe, they appear with the ECC queue name *`Pattern Launcher`*, followed by the name of the pattern. The probe contains a sensor named *`Horizontal Discovery Sensor`*, which performs the actual updates of the CMDB based on identification rules.

If you create your own device or process classifier and you want to use patterns for discovery, you must specify this probe in the classifier record. You don't need to modify this probe or the Horizontal Discovery sensor.

Splitting payload

When Discovery uses patterns to discover certain devices such as load balancers or network switches, the MID Server can collect large amounts of data. To better handle this, the Horizontal Pattern probe splits the payload into chunks and create multiple ECC Queue records.

Payload splitting happens in two phases. First, the MID Server checks the total payload against size and count limits. The system checks the size limit first. Second, the MID Server evaluates each configuration item (CI) for the number of related items. When a CI has more related items than the chunk size, the system splits those items into separate ECC queue pages. Each page contains a copy of the parent CI with a subset of its related items.

When both phases interact, the number of ECC queue pages can grow significantly. For example, a network switch with 50,000 related items and a chunk size of 100 creates 500 ECC queue pages for that single CI.

The following MID Server properties control payload splitting.

Property	Default value	Description
mid.discovery.max_total_items_size_per_page	3145728 bytes (3 MB)	Sets the maximum size limit of all items (CIs, related items, relations, and references) per ECC queue page. The system checks this threshold first.
mid.discovery.max_ci_count_per_page	500	Sets the maximum number of CIs per ECC queue page.
mid.discovery.max_relation_count_per_page	10,000	Sets the maximum number of relations per ECC queue page.
mid.discovery.max_related_count_per_page	100	Sets the maximum number of related items per chunk. Each chunk becomes a separate ECC queue page containing a copy of the parent CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Default value	Description
mid.discovery.should_minify_pattern_payload	false	Enables JSON minification for pattern payloads. When set to true, trims JSON whitespace from the pattern payload to reduce its size.

Tuning payload splitting

You can adjust MID Server properties to reduce the number of ECC queue pages that Discovery creates. The following table identifies common issues and the actions to address them.

Issue	Action	Related property
A CI with many related items creates an excessive number of ECC queue pages	Increase the related items chunk size	mid.discovery.max_related_count_per_page
Payload size is the splitting bottleneck	Increase the size limit and the MID Server JVM heap size	mid.discovery.max_total_items_size_per_page
Payload size exceeds the limit but increasing the JVM heap size is not feasible	Enable payload minification to reduce the JSON payload size	mid.discovery.should_minify_pattern_payload
A high CI count per page is causing splitting	Increase the CI count limit	mid.discovery.max_ci_count_per_page
A high relation count per page is causing splitting	Increase the relation count limit	mid.discovery.max_relation_count_per_page

Note:

- These properties are per MID Server. Apply changes to all MID Servers involved in the affected Discovery schedules.
- If necessary, increase *mid.discovery.max_total_items_size_per_page* and the MID Server JVM heap size accordingly. Setting these values too high can cause out-of-memory errors on the MID Server and instance nodes. Set values appropriate for your environment.
- Setting *mid.discovery.should_minify_pattern_payload* to true can reduce payload size without requiring heap size increase.
- For environments where a CI has many related items, but fewer total CIs, the primary bottleneck is likely *mid.discovery.max_related_count_per_page*. Start by increasing that property.

PowerShell probe

The PowerShell Probe executes PowerShell V2 scripts on the MID Server host.

PowerShell scripts are defined as probe parameters with the filename as the parameter name. It is available as a Probe probe type by specifying PowerShell as the probe's ECC queue topic.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PowerShell probe parameters

Parameter name	Description
source	[Required] The initial host to connect to. **Default:** None
<script name>.ps1	[Required] The filename of the PowerShell script to run. Replace *<script name>* with a valid filename prefix. **Default:** None
powershell_command_parameter_passing	Specifies whether to pass script parameters on the command line. Regardless of this parameter's value, ServiceNow makes all script parameters on the command line automatically available to PowerShell scripts as environment variables. **Default:** false
powershell_param_<script parameter name>	Passes additional parameters to the PowerShell script to be executed. Each parameter will appear to the script as an environment variable in the format *$env:SNC_<script parameter name>*. Parameters with this prefix are not considered encrypted and are passed through to the script untouched. Make sure you select the appropriate parameter between *powershell_param_<script parameter name>* and *powershell_<script parameter name>*. Using the wrong prefix results in errors in the PowerShell execution, which is passed back to the instance in the ECC queue input. **Default:** None
powershell_<script parameter name>	Passes additional parameters to the PowerShell script to be executed. Each parameter will appear to the script as an environment variable in the format *$env:SNC_<script parameter name>*. The MID Server assumes that any parameter with this prefix is encrypted and attempts to decrypt it. Make sure you select the appropriate parameter between *powershell_param_<script parameter name>* and *powershell_<script parameter name>*. Using the wrong prefix results in errors in the PowerShell execution, which is passed back to the instance in the ECC queue input **Default:** None

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter name	Description
debug	Enables debug log output during the probe. **Default:** false
credentials_debug	Displays a <credentials_debug> section in the ECC queue, which can help you troubleshoot credentials. If you set this property to true, credential troubleshooting information is output to the ECC queue, even if the credentials succeed. **Default:** false

Scripting requirements

Any custom PowerShell scripts must use environment variables to pass any non-Boolean command line parameter. Replace non-Boolean parameters in the *Param()* portion of the script with script variables of the same name. Define the script variable as part of the environment with an SNC_ prefix. So a string parameter such as this:

```
Param([string]$paramName)
```

Becomes a script variable such as the following:

```
if(test-path env:\SNC_paramName) {
    $paramName = $env:SNC_paramName
}
```

For example, this parameter definition from the PSScript.ps1 script contains several string parameters that need to be redefined as script variables:

```
Param([string]$computer, [string]$script, [string]$user,
 [string]$password, [boolean]$useCred, [boolean]$isDiscovery,
 [boolean]$debug)
```

Defining the non-Boolean parameters as script variables would result in this type of script:

```
Param([boolean]$useCred, [boolean]$isDiscovery, [boolean]$debug)

# Copy the environment variables to the params
if(test-path env:\SNC_computer) {
  $computer=$env:SNC_computer
}

if(test-path env:\SNC_script) {
  $script=$env:SNC_script
}

if(test-path env:\SNC_user) {
  $user=$env:SNC_user
  $password=$env:SNC_password
}
```

Create a custom PowerShell probe

You can create your own PowerShell probe and configure probe parameters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **Discovery Definition** > **Probes**.
2. Click **New**.
3. Enter the following values.
   - **Name**: Any unique name
   - **Probe** Type: Probe
   - **ecc_queue_topic**: Powershell
4. Right-click the form header and click **Save**.
5. From the **Probe parameters** related list, click **New**.
6. Create a probe parameter record for each parameter.
7. Click **Submit**.

SCPRelay probe
The *SCP Relay Probe* copies a single file or the contents of a directory from one host to another, using the MID Server as a relay.

The SCP Relay probe uses the same parameters as SSHCommand. The commands may be sent in or out of the context of a terminal (tty). Commands can be sent with or without sudo for those commands, such as lsof, that require being executed in the context of root to return the required information. When commands are sent in the context of a terminal, the path is automatically widened to include a set of default paths. You can further widen this with the path_override parameter. If the target machine is the local machine, SSH is not used; instead, a local shell is run to execute the command.

For instructions on configuring probe parameters, see Set probe parameters.

The following parameters may be passed to the *SCP Relay probe*:

Parameters

Parameter	Description	Default Value
debug	Enables debug logging.	false
debug_ssh	Enables J2SSH debug logging **agent.log** **Note:** This log file can grow quickly and should be disabled when not needed.	false
timeout	Sets the socket connection timeout, in ms.	60000
path_override	Overrides the default paths set before executing a command.	none
keyboard_interactive	If true, forces the use of keyboard_interactive SSH login mode.	false
must_sudo	If true, forces the command to be executed through sudo.	false
run_in_terminal	If false, disables SSH commands from running in an SSH terminal (this will break many commands).	true

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameters (continued)

Parameter	Description	Default Value
long_runner	If present, indicates a long-running SSH command.	false
set_path	If false, disables setting the path environment variable before running the command.	true
rm_override	If present, overrides the default rm command ("/sbin/rm -f") with the provided value.	none
source	Source host or IP to copy from.	required
source_port	SSH port on the source (defaults to 22).	required
from_file	Name of the file to copy from the source.	required
target	Target host or IP to copy to.	required
target_port	SSH port on the target (defaults to 22).	required
to_file	Name of the file to copy to the target.	required

SNMP probes

The SNMP probes use the SNMP protocol to query a particular device for a list of OIDs, which are then traversed and the results passed back to the sensors. MID Servers support all SNMP protocol versions by default. You can set a MID Server to only support specific versions of SNMP.

Discovery supports SNMP versions 1, 2c, and 3. The default value is ALL. If you select one of the versions, but don't have that version's credentials, the Use SNMP Version reverts back to ALL.

Note: SNMP version credentials must be active=true. If a particular version is inactive, it will also revert back to ALL. A message will appear similar to: `--- No SNMP credentials of type 'snmpv3' are active in this instance. Setting Use SNMP version back to All.`

SNMP patterns

The Network Router pattern replaces the several SNMP-related probes for the discovery of routers and switches. See Available discovery patterns for a list of patterns.

SNMP probe parameters

This list of parameters may be passed to the SNMP probes.

For instructions on configuring probe parameters, see Set probe parameters.

Parameter	Description	Default Value
oid_spec_list	A list of OID specifications, one per line. Each specification must be in one of the following two forms: • *walk {OID}*: Walks the OID and all its children • *table {OID} {OID Children}*: Walks all entries in the table, returning only the given children (for example, "iso.org.dod.internet")	required

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description	Default Value
	{OID Children} refers to a comma-delimited list of child nodes within the entries for the given table. For example, "ifEntry.ifIndex,ifEntry.ifDescr,ifEntry.ifType" are OID children of the table "iso.org.dod.internet.mgmt.mib-2.interfaces.ifTable". As a convenience, the table entry prefix may be left off. (The preceding children could be specified as "ifIndex,ifDescr,ifType".) Any child may include a filter qualifier in parentheses. For example, the child "entPhysicalContainedIn(=0)" specifies returning table entries only if the value of "entPhysicalContainedIn" equals 0. The operators supported in the expression are: • `=` : equals • `!=` : does not equal • `#` : contains If more than one child has a filter expression, a match on any one of the children causes that entry to be read. Any content including and after a "//" is ignored (comments). Any OID that does not start with "1.3.6.1" or "iso.org.dod.internet" automatically prefaces with "1.3.6.1." as a convenience.	
source	The IP address or host name of the device to query SNMP on.	required
index	The index to apply after the community string, for Cisco-style community string indexing (for VLAN interrogation).	0
credential_id	The sys_id of a specific credential that is preferred for use above the rest. This parameter is for internal use only and is not supported.	none
credential_tag	The credential tag that must be used. This parameter is for internal use only and is not supported.	none
timeout	The timeout value (in milliseconds) to wait for a response, instead of the default. You can use this parameter to override the *mid.snmp.request.timeout* SNMP MID Server configuration parameter. **Note:** When *use_getbulk* is set to *true*, the timeout value is for an individual GETBULK request.	1500
established_session_timeout	The timeout (in milliseconds) to wait for a response after at least one response has been received. Longer values can be useful for collecting complete and accurate data. You can use this parameter to override the *mid.snmp.session.timeout* SNMP MID Server configuration parameter.	500
debug	Enables debug logging. Set to true for debug mode.	false
request_interval	The interval (in milliseconds) between successive requests for an OID when a response has not been received, until	400

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description	Default Value
	the timeout (or established_session_timeout) value is reached. If this value is set to at least as long as the timeout (or established_session_timeout) value, then only a single request is sent for any particular OID. If you change the timeout (or established_session_timeout) value, adjust the request_interval at the same time to avoid sending too many or too few requests for the same OID, as appropriate for a given environment.	
request_delay	The interval (in milliseconds) between the receipt of a response and the transmission of the next request. The default is 0 (no delay). This value may be set to slow the overall rate of an SNMP query.	0
result_format	Returns JSON formatted payloads for these probes: • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System • SNMP - Routing This parameter returns data in a more compact format to prevent sensor failure or memory problems on a node when the payload becomes large. Do not change this value or delete this parameter. ⚠ **Warning:** Use of this parameter with any other probes causes the sensor to fail.	JSON
use_getbulk	Enables the use of SNMP GETBULK requests to retrieve tabular data from SNMP devices instead of using multiple SNMP GETNEXT requests. For tabular data, GETBULK is more efficient. Regardless of the request type, certain devices may not return any results when they are busy with other tasks. This parameter is used to configure at the probe level. GETBULK can also be set for an individual MID Server or globally for all MID servers. Settings are listed in the order of precedence: • SNMP probe parameter • MID Server configuration parameter • MID Server properties The *established_session_timeout*, *request_interval*, and *request_delay* parameters are ignored when *use_getbulk* is set to **true**. Instead, the *retries* parameter is available. The timeout configuration is the same one used by *use_getscalar*. By default, the following probes use GETBULK requests (the parameter value is *true*). • SNMP - Switch - Vlan • SNMP - Switch - BridgePortTable • SNMP - Switch - ForwardingTable	false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description	Default Value
	• SNMP - Switch - SpanningTreeTable • SNMP - Network - ArpTable • SNMP - Layer 2 Protocol Caches • SNMP - F5 BIG IP - System (only for Service Mapping customers) **Note:** These probes have a *timeout* value of 5000.	
use_getscalar	Enables the use of simplified retrieval and handling of scalar values from SNMP devices. The *established_session_timeout*, *request_interval*, and *request_delay* parameters are ignored when *use_getscalar* is set to **true**. Instead, the *retries* and *timeout* parameters are available. The timeout configuration is the same one used by *use_getbulk*.	false
retries	The number of additional attempts Discovery makes to complete an individual GETBULK request (see *use_getbulk*) or a GETNEXT request when the *use_getscalar* parameter is set to **true**.	2

Load a MIB module

You can load an additional Management Information Base (MIB) module by creating a new ecc_agent_mib record and attaching the actual MIB file to the record.

Before you begin

Role required: admin

About this task

MIBs are only loaded during MID Server startup. When the MID Server starts up, the following MIB modules are loaded:

- Modules that are bundled with the MID Server software. These MIBs are defined industry standards.
- Modules that are included in each instance. These MIBs are provided by manufacturers for Discovery to extract specific information from a device.

You can view any errors associated with loading a MIB module in the agent log.

Note: Adding an MIB file allows you to request OID values using String/ASN notation in addition to numerical notation, but it doesn't guarantee that any additional information will be retrieved. It simply provides another way to define new OIDs you may want to discover.

Procedure

1. Navigate to **All** > **MID Server** > **SNMP MIBs**.
2. Check whether dependencies are met.
   If your new MIB has dependencies on another MIB, the MIB that fills the dependency must exist before you create your new record. Search the existing MIBs to check that the required MIBs are already loaded. If they are not, use this procedure to also add them to the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click **New** to create a new record.
   The MID Server MIB File form opens to create a new ecc_agent_mib record.
4. Use the following information to fill out the form:
   - **Name**: The name of the MIB.
   - **Version**: The version of the MIB.
   - **Source**: Use this field to note where the MIB was acquired, such as a URL.
   - **Description**: The description that appears in the ecc_agent_mib table.
   - **Active**: This check box denotes whether the MIB module is enabled or disabled in the instance.
5. Click the Add Attachment icon ( ) in the upper right to attach the actual MIB file to the new record.
   The MIB name must begin with an alphabetical character. Remaining characters must be one of the following: alphanumeric, hyphen (-), or underscore (_). The file name must not have an extension. You can reference the existing MIBs for examples. Use the actual name of the MIB for both the MIB record name and the attachment name, but it is not required.
6. Select **Submit**.

What to do next

Update the appropriate pattern or probe to query for the specific OID values defined in the new MIB.

SNMP probe MIB modules

A management information base module (MIB) is a database that is used to manage elements in a network.

Your instance includes object definitions from the most common management information base (MIB) files. Before adding a new object definition, consult the list of MIB modules to ensure that the object definitions are not already available. The instance includes object definitions from the most common management information base (MIB) files. The MID Server MIB File [ecc_agent_mib] table is domain separated. You can create versions of these policies that only a MID Server from the same domain can use. For instructions, see Set up domain separation for MID servers.

The following table contains the MIBs that load automatically. The MIBs that are bundled with the MID Server are not visible in the instance. The MIBs that are in the ecc_agent_mib table are included in the instance by default and can be viewed in the **MID Server** > **SNMP MIBs** module.

Important: MIB files in a ServiceNow instance do not use a file extension. If the **glide.attachment.extensions** system property in your instance contains a list of file extensions permitted for import, MIB files do not load. To ensure that MIB files can load, add a temporary extension to the files prior to loading, such as **.mib**, and then list that extension in the **glide.attachment.extensions** property. When you load the file, a business rule called **MIB filename compliance** strips this extension from the file before it inserts the record into the ecc_agent_mib table.

Management Information Base (MIB) module

Bundled with the MID Server	Included in the instance
ATM-MIB	A10 AX MIB
BRIDGE-MIB	A10 COMMONS MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Management Information Base (MIB) module (continued)

Bundled with the MID Server	Included in the instance
DISMAN-EVENT-MIB	CISCO-SMI
ENTITY-MIB	CISCO-STACK-MIB
EtherLike-MIB	CISCO-TC
FRAME-RELAY-DTE-MIB	CISCO-VTP-MIB
HOST-RESOURCES-MIB	F5-BIGIP-COMMON-MIB
HOST-RESOURCES-TYPES	F5-BIGIP-LOCAL-MIB
IANA-ADDRESS-FAMILY-NUMBERS-MIB	F5-BIGIP-SYSTEM-MIB
IANAifType-MIB	FOUNDRY-SN-AGENT-MIB
IP-MIB	FOUNDRY-SN-ROOT-MIB
IP-FORWARD-MIB	LanMgr-Mib-II-MIB
IPMROUTE-STD-MIB	MSFT-MIB
IPV6-MIB	NetWare-Host-Ext-MIB
ISDN-MIB	NetWare-Server-MIB
RFC-1212	PowerNet-MIB
RFC-1215	QMS-MIB
RFC1155-SMI	
RFC1213-MIB	
RFC1398-MIB	
RFC1406-MIB	
RMON2-MIB	
RSVP-MIB	
SNMPv2-SMI	
SNMPv2-TC	
UPS-MIB	

SSHCommand probe

A probe using the ECC queue topic name *SSHCommand* executes a shell command on the target host, and returns the resulting output to the sensor.

Discovery supports Bourne Shell (sh) and Bourne-again Shell (bash) commands. Enter shell script commands in the probe's *ECC queue name* field. The shell script can use variables and file operations supported by the target UNIX shell.

- The SSH engine is active by default on new instances.
- Customers on upgraded instances can manually enable ServiceNow SSH for a particular probe by setting the use_snc_ssh parameter to true. Alternatively, enable it for all probes on the MID Server by setting the MID Server parameter mid.ssh.use_snc to true.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: To discover network devices, such as routers and switches, use SNMP credentials credentials, not SSH credentials.

SSHCommand parameters

Several parameters are available for the SSHCommand probe.

For instructions on configuring probe parameters, see Set probe parameters.

Parameters Table

Parameter	Description
allow_unsupported_shells	Allows a probe designer to attempt to run a command on a device that does not have a supported shell. With no shell, the following is true: • No path information can be provided. • No scripts can function, because there is no ability to handle script parameters. • The text that is specified in the **ECC queue name** field of the probe form is the raw command that is run on the device. For example, you may design a probe to display the version of a Cisco switch running NX-OS. The command that the switch understands is `show version`. To accomplish this, put the `show version` command in the **ECC queue name** field, and add the parameter name `allow_unsupported_shells` with a value of **true**. This parameter is only effective for `sncssh`. It is not supported with j2ssh. Currently supported shells are sh, bash, ksh, csh, and tcsh. • Type: string, true or false • Default value: false
source	[Required] Specifies the initial host to connect to. • Type: string (URL) • Default value: None
port	Specifies the target port to connect to. • Type: integer (port) • Default value: 22
debug	Specifies whether to write SSH debug information to the log file. The parameter usage depends on whether the ServiceNow SSH client is enabled. When the ServiceNow SSH client is enabled, this parameter functions as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameters Table (continued)

Parameter	Description
	• Type: string • Default value: false The following string values are valid for the ServiceNow SSH client: ○ true: Enables SSH debug information in the log file. ○ false: Disables SSH debug information in the log file. ○ <IP Addresses>: Specifies which IP ranges to enable SSH debug information in the log file. You can enter IP addresses in the following formats: ▪ An IP range defined by a slash and the number of bits in the subnetwork. For example, the string 10.10.10.0/24 scans 24 bits of IP addresses from 10.10.10.0 to 10.10.10.254. ▪ An IP range defined by a dash. For example, the string 10.10.11.0-10.10.11.165 scans the IP addresses from 10.10.11.0 to 10.10.11.165. ▪ A comma-separated list of specific IP addresses. For example the string 10.10.11.200,10.10.11.235 scans the IP addresses 10.10.11.200 and 10.10.11.235. ○ deferred: Specifies to log SSH debug information in memory unless an error or warning occurs. If an error or warning occurs, the platform publishes the debug information to the log file. This ensures that only the part of the log file pertaining to the error or warning is recorded. If no error or warning is detected, the platform deletes the unused log data from memory when the session closes. Each session stores up to 1000 log messages. If the session exceeds 1000 log messages, the deferred log discards the oldest log message to make room for the newest log message. When the ServiceNow SSH client is disabled, this parameter enables or disables SSH debug information in the log file: • Type: true | false • Default value: false
debug_ssh	Specifies whether the legacy SSH client writes debug information into the `agent/logs/agent.log` log file. **Note:** This log file can grow quickly and should be disabled when not needed. The ServiceNow SSH client does not use this parameter. • Type: true | false • Default value: false
timeout	Sets the socket connection timeout for the legacy SSH client.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameters Table (continued)

Parameter	Description
	The ServiceNow SSH client does not use this parameter. • Type: integer (milliseconds) • Default value: 60,000
path_override	Specifies how to change the default paths set before executing a command. Type one or more override paths delimited by a colon (:). The default path is /usr/sbin: /usr/bin: /bin: /sbin. The ServiceNow SSH client accepts the following prefixes in front of the path_overide value: • append: Appends the override path to the end of the host's path. This is the default behavior. • replace: Replaces the host path with the path_overide value. • prepend: Appends the override path to the front of the host path. • Type: string (a colon-separated list of directories) • Default value: None
keyboard_interactive	Determines whether to enforce keyboard_interactive SSH login mode. • Type: true | false • Default value: false
must_sudo	Determines whether SSH commands run through sudo. **Note:** must_sudo probe parameter affects all patterns. On Unix or Linux hosts, pattern execution can run many different commands, some should not have sudo appended. • Type: true | false • Default value: false

run_in_terminal	Determines whether SSH commands run in an SSH terminal.
	• Type: true | false
	• Default value:
	  ○ ServiceNow SSH client: false
	  ○ Legacy SSH client: true

set_path	Determines whether the probe is allowed to alter the session's PATH variable or not. By default, during session setup, the PATH variable is set to `/usr/sbin:/usr/bin:/bin:/sbin`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameters Table (continued)

Parameter	Description
	• Type: true | false • Default value: true
rm_override	Overrides the default remove command (`/bin/rm -f`) with the provided value. • Type: string • Default value: none
use_snc_ssh	Enables the ServiceNow SSH client. The ServiceNow SSH client is active by default on new instances. Enabling the ServiceNow SSH client disables the legacy SSH client. • Type: true | false • Default value: true
command_timeout_ms	Number of milliseconds an SSH command is allowed to run before timing out (default is configurable per MID server). The legacy SSH client does not use this parameter. • Type: integer • Default value: value of the mid.ssh.command_timeout_ms MID Server parameter.
channel_timeout_ms	Specifies the amount of time, in milliseconds, that a MID Server waits for activity during processing of an SSH command. If the MID Server does not detect activity in the specified timeout window, this parameter kills the command. • Type: integer (milliseconds) • Default value: 300000 (5 minutes)
preserve_sudo_environment	Specifies whether to use sudo to preserve the environment for SSH. This parameter is only effective if the sudo environment on the host being probed supports the `-E` switch. • Type: true | false • Default value: false
credentials_debug	Displays a <credentials_debug> section in the ECC queue, which can help you troubleshoot credentials. If you set this property to true, credential troubleshooting information is output to the ECC queue, even if the credentials succeed. See Credentials Troubleshooting for more information.

SSHCommand path

The SSHCommand probe computes the default path from the following sources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- MID Server parameter: mid.ssh.path_override
- SSH Command probe parameter: path_override
- User's default path: Shell configuration file

If you set the MID Server path override parameter, Discovery appends this path to the default path. If you set the probe path parameter, Discovery uses this path instead of the default path. Discovery always appends a user's default execution path to the MID Server and probe parameter paths.

Default Path

By default, the MID Server searches for SSH commands in the following paths:

- /usr/sbin
- /usr/bin
- /bin
- /sbin

Shell script options

The SSHCommand probe supports the following scripting options in the *ECC queue name* field.

Shell Scripting Options Table

Summary	Syntax	Description
Variable	${variable}	Replaces the token with the value of the variable. For example, ${catalina_home} specifies the installation location of a Tomcat server.
Include File	${File:file_name.sh}	Treats the contents of the specified file as a shell script. For example, ${File:findcat.sh} runs the findcat shell script.

Non-privileged SSH commands during probe-based discovery

These tables display the SSH commands run by Discovery probes on target devices during horizontal discovery. These SSH commands don't require elevated privileges to run.

Operating systems

For information on commands used by Service Mapping during the top-down discovery, see Service Mapping commands requiring a privileged user and Service Mapping commands not requiring a privileged user.

AIX commands

Command (ECC queue name)	Parameters
sh	*variable*
lsattr	*-El sys0 -a modelname -F value*
lslpp	*-Lc*
/etc/ifconfig	*-au; netstat -in*
instfix	*-i | grep AIX_ML*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AIX commands (continued)

Command (ECC queue name)	Parameters
oslevel	*-r*
echo `lsattr	*-El sys0 -a systemid | cut ...*
rmsock	
lsdev	• *-C -F 'name,status' -c memory| grep '^mem'* • *-C -F 'name,status' -c processor -S a*
ps	*-aef -o "pid,ppid,args"*

Linux commands

Command (ECC queue name)	Parameters
arp	*-n*
bash	*variable*
cat	• */sys/hypervisor/compilation/compiled_by* • */sys/devices/virtual/dmi/id/product_uuid* • */etc/*release* • */proc/cpuinfo* • */var/lib/dpkg/status* • */proc/mounts* • */proc/scsi/scsi* • */etc/iscsi/initiatorname.iscsi* • */sys/hypervisor/compilation/compiled_by*
grep MemTotal	*/proc/meminfo*
ifconfig	*-a && route -n*
sh	*variable*
rpm	*-q -a --qf*
ls	• *-d /sys/block/* | grep '^/sys/block/\(dm-|sd|sr\)'* • */sys/class/scsi_host/host*/protocol 2> /dev/null*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux commands (continued)

Command (ECC queue name)	Parameters
	• */sys/class/scsi_host/host*/ info 2> /dev/null* • *-d /sys/class/fc_host/host**
shopt	*-s extglob*
findmnt	*-l -n -o SOURCE,TARGET 2>/dev/ null*
find	*/sys/class/iscsi_session -maxdepth 1 -type l*
fcoeadm	*-i*
ip	• *neighbor 2> /dev/null* • *address show && ip route*
hostname	*-f 2> /dev/null*
dnsdomainname	*2> /dev/null*
grep	• *^[[:space:]]*domain /etc/ resolv.conf* • *^[[:space:]]*search /etc/ resolv.conf*
wget	• *--spider -t 1 -T 1* • *-qO-*
curl	*--connect-timeout 1*
python	*-c*
sudo cat	*/sys/devices/virtual/dmi/ id/product_uuid*

HP-UX commands

Command (ECC queue name)	Parameters
ps	*-elx*
echo itick_per_usec/D	*| adb -k /stand/vm...*
sh	*variable*
model	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HP-UX commands (continued)

Command (ECC queue name)	Parameters
uname	• *-i* • *-r*
swlist	*-l product | grep -v 'PH.._'*
echo 'memory_installed_in_machine/D'	*| a...*
netstat	*-in | grep -v Name*
lanscan	*-ia*
typeset	*-i2*
ioscan	*-fnkC processor | grep PROCESSOR | wc -l*

Solaris commands

Command (ECC queue name)	Parameters
sh	*variable*
kstat	*cpu_info*
pkginfo	*-l | egrep 'PKGINST|NAME| CATEGO...*
prtconf	• *| grep Memory* • *| grep 'System Configuration:'* • *-pv 2>/dev/null*
ifconfig	*-a; netstat -pnr -f inet*
arp	*-an*
showrev	*-p*
sneep	*-T*
bash	*variable*
pkg	*list*
iostat	*-En | awk '/^c/'*
zpool	• *list* • *status | awk '$1 == "NAME", $1 == ""'*
df	*-k*
fcinfo	*hba-port*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solaris commands (continued)

Command (ECC queue name)	Parameters
zonename	
zoneadm	
zonecfg	*-z*
hostid	
uname	*-r*
nawk	
ps	• *-o pid,ppid* • *-awwx*
pfiles	

Mac commands

Command (ECC queue name)	Parameters
system_profiler	• *SPHardwareDataType SPMemoryDataType -xml* • *SPSerialATADataType SPSASDataType SPFireWireDataType -xml* • *SPNetworkDataType -xml;ifconfig | grep '|Winet6|W'* • *SPSoftwareDataType -xml* • *SPHardwareDataType -xml* • *SPApplicationsDataType -xml*
last	*| grep -m 1 console*
hostname	*2> /dev/null*
grep	• *'^[\s]*domain' /etc/resolv.conf* • *'^[\s]*search' /etc/resolv.conf*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

UNIX commands

Command (ECC queue name)	Parameters
ps	• *awwxo pid,ppid,command | sed -n '/<defunct>/!p'* • *-ef*
sh	*variable*
cat	*variable*
variable	*-version*
uname	*-a*
df	*-kP*
uptime;env	*LC_ALL=C date '+%a %b %e %H:%M %Z %Y'*
echo	*|echo `uname -I` `hostname` `uname -Iv`.`uname -Ir` 2>&1``*
awk	
grep	
find	
uuencode	*-m*
base64	*-w 0*
wc	*-c*
nohup	
vmware	*-v 2>&1*
systemctl	*status cloudian-s3 2>&1*

HMC OS commands

Command (ECC queue name)	Parameters
echo	*|lshmc -V 2>&1``*

Web servers

Apache commands

Command (ECC queue name)	Parameters
sh	*variable*
httpd_command	*-V*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tomcat commands

Command (ECC queue name)	Parameters
sh	*variable*
cat	*variable*
find -L ${dirname}	*-name web.xml -print 2>/dev/null*

NGINX commands

Command (ECC queue name)	Parameters
sh	• *get-nginxcfg* • *get-nginx-ver*

Databases

Oracle commands

Command (ECC queue name)	Parameters
sh	*variable*
lsnrctl	*status*

SQL commands

Command (ECC queue name)	Parameters
cat	*variable*
variable	*-V*

HBase commands

Command (ECC queue name)	Parameters
cat	*variable*
sh	*variable*
variable	*-lv | grep ^${processID}*
echo	*"version" | hbase shell*

Applications

ESX - OS commands

Command (ECC queue name)	Parameters
sh	*variable*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

F5 - OS commands

Command (ECC queue name)	Parameters
sh	*variable*
tmsh	*show sys version 2>&1*

HAProxy commands

Command (ECC queue name)	Parameters
sh	*variable*
variable	*-v*

JBoss commands

Command (ECC queue name)	Parameters
find	*-name web.xml -print*
cat	*variable*

KVM commands

Command (ECC queue name)	Parameters
sh	*variable*
virsh	*Various capabilities*

Puppet commands

Command (ECC queue name)	Parameters
puppet	• *--version* • *config print*
hostname	*--fqdn*

Storage

Command (ECC queue name)	Parameters
purehw	*list --all --nvp --type ch*
vxprint	*-AQmb*
vxdisk	*path*
sudo vxdisk	*list | awk 'NR > 1'*
pureblade	*list 2>&1*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Discovery commands for probes and patterns

Privileged SSH commands for probe-based discovery

These tables display the SSH commands run by Discovery probes during horizontal discovery. These SSH commands require elevated privileges to run.

Operating system commands requiring elevated rights

These examples assume that the user name is **Disco**. Substitute the actual user name and verify that the paths for the commands match the paths on the system.

Note: Sudo commands don't work with private key credentials, because there's no password to supply to the sudo command. A solution is to add the *NOPASSWD* option to the sudo configuration. For example, you might enter: `disco ALL=(root) NOPASSWD: /usr/sbin/dmidecode,/usr/sbin/lsof,/sbin/ifconfig`.

For information on commands that don't require elevated rights, see Non-privileged SSH commands during probe-based discovery.

For information on commands used by Service Mapping during the top-down discovery, see Service Mapping commands requiring a privileged user and Service Mapping commands not requiring a privileged user.

SSH key not validated

When the MID Server connects to a system, the MID Server doesn't perform host key validation against that system and so treats it as untrusted. If an attacker performs a man-in-the-middle attack and redirects the traffic to a malicious SSH service, the attacker can intercept or modify any data sent over the connection.

Therefore, limit any sensitive information exchanged between the MID Server and the target SSH server. Only use keys or certificates for SSH authentication, and avoid sending system credentials. Configure *NOPASSWD* in the sudoers file for the required privileged commands.

HP-UX

Command	Purpose
adb	Gathers CPU speed and memory. **/etc/sudoers line example**: `Disco ALL=(root) /usr/bin/adb`

All Linux

Command	Purpose
dmidecode	Gathers several pieces of information about the hardware, including the serial number embedded within the motherboard. **/etc/sudoers line example**: `Disco ALL=(root) /sbin/dmidecode`
fdisk	Gathers the disks and size information on the system. **/etc/sudoers line example**: `Disco ALL=(root) /usr/bin/fdisk -l`
multipath	Gathers device mappings for MultiPath Input Output (MPIO).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

All Linux (continued)

Command	Purpose
	/etc/sudoers line example: `Disco ALL=(root) /usr/bin/multipath -ll`

Linux and Solaris

Command	Purpose
dmsetup	Examines a low-level volume. **/etc/sudoers line example** • `Disco ALL=(root) /usr/bin/dmsetup table *` • `Disco ALL=(root) /usr/bin/dmsetup ls`

All UNIX versions

Command	Purpose
lsof	Determines the relationship between processes and the connections being made to the system. **/etc/sudoers line example**: `Disco ALL=(root) /sbin/lsof`
oratab	Grants read access to the oratab file for locating the Oracle Home and pfile.
netstat	Determines the relationship between processes and the connections being made to the system. **/etc/sudoers line example**: `Disco ALL=(root) /bin/netstat`
ss	Determines the relationship between processes and the connections being made to the system. **/etc/sudoers line example**: `Disco ALL=(root) /sbin/ss`

Solaris

Command	Purpose
iscsiadm	Gets iSCSI qualified names (IQNs). **/etc/sudoers line example**: `${sudo:iscsiadm list target -S}`
fcinfo	Gets World Wide Port Names (WWPNs) for ports. **/etc/sudoers line example**: `${sudo:fcinfo remote-port -sl -p $port}`
prtvtoc	Reports information about disk partitions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solaris (continued)

Command	Purpose
	/etc/sudoers line example: `Disco ALL=(root) /usr/bin/prtvtoc`
/usr/bin/ps	Lists running process. As an alternative to running with root access, add a proc_owner role.sola. **/etc/sudoers line example**: `Disco ALL=(root) /usr/bin/ps`
/usr/ucb/ps	Lists running process. As an alternative to running with root access, add a proc_owner role. The use of the `/usr/ucb/ps` command is deprecated as of Solaris 11. Because Discovery requires the use of this command for all Solaris versions, you must install the ucb utility manually on Solaris 11 systems. For instructions, see KB0564262. **/etc/sudoers line example**: `Disco ALL=(root) /usr/ucb/ps`
pgrep	Gets list of process IDs (PIDs) with socket information. **/etc/sudoers line example**: `Disco ALL=(root) /usr/bin/pgrep`
pfiles	For each PID, gets and processes the output for S_IFSOCK. **/etc/sudoers line example**: `Disco ALL=(root) /usr/bin/pfiles`

Related topics

Discovery commands for probes and patterns

vCenter probes and probe parameters

vCenter probes scan virtual machines using VMware's vSphere product suite. Each probe scans for different kinds of data, such as networks, NICs, and tags. The VMware - vCenter probe that discovered all vCenter objects in previous releases is deprecated in the Istanbul release and replaced by multiple probes.

Upgrade changes

In upgraded systems, the vCenter process classifier and the vmapp port probe are configured to trigger the VMware - vCenter Datacenters probe. This probe then triggers the probes that discover individual vCenter objects such as hosts, storage, and datastores. The legacy probe, VMware - vCenter, is still in the system but is not triggered to run in the updated instance.

Available vCenter probes and probe parameters

vCenter probe parameters enable you to disable the probes for the objects you're not interested in discovering. You can also reduce the size of the payload of a probe by specifying a page size.

Important: Before disabling a probe, be aware of any dependencies the probe might have. If the probe you disable triggers another probe, the dependent probe is also disabled, and cannot collect data.

These parameters are available for vCenter probes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: By default, these parameters are not set in the platform, except where noted.

vCenter probe parameters

Probe	Parameters
VMWare - vCenter Datacenters	• datacenter_only (Cloud Provisioning and Governance) • debug • disable_vm_probe • disable_network_probe • disable_datastore_probe • disable_cluster_probe • datacenters_only • fixed_credential_id (Cloud Provisioning and Governance) • disable_vm_nic_probe • disable_host_probe • disable_host_storage_probe • vcenter_timeout: default 10000 milliseconds (10 seconds).
VMWare - vCenter VMs	• disable_vm_nic_probe • disable_vm_nic_vnics (set to **false**) • disable_vm_nic_vdisks (set to **true** for Discovery only and **false** when Cloud Provisioning and Governance is enabled) • disable_vm_tags_probe • refresh_state (Cloud Provisioning and Governance) • page_size (100) **Note:** The settings for the vnics and vdisks parameters preserves the default behavior of vCenter Discovery.
VMWare - vCenter VM Tags*	none
VMWare - vCenter VM NICs	none
VMWare - vCenter Networks	page_size (500)
VMWare - vCenter Datastores	page_size (500)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vCenter probe parameters (continued)

Probe	Parameters
VMWare - vCenter Clusters	• disable_host_probe • disable_host_storage_probe • page_size (1000)
VMWare - vCenter ESX Hosts	• disable_host_storage_probe • page_size (10)
VMWare - vCenter ESX Hosts Storage	page_size (175)
VMWare - vCenter ESX Hosts License	none
VMWare - vCenter VM Performance	This probe relies on the statistics intervals configured for the vCenter. By default, the following intervals are available: • 5 minutes (300 seconds) – this data is saved for 24 hours. • 30 minutes (1800 seconds) – this data is saved for a week • 2 hours (7200 seconds) – this data is saved for a month • 24 hours (86400 seconds) – this data is saved for a year The following properties support this probe: • data_sampling_weeks_frame: The period for which the probe retrieves the data, in weeks. • data_sampling_minutes: The time lapse between taking the data samples. • metrics: The collected metrics (CPU / memory / network / disk). This probe is used only by Cloud Migration Assessment. **Important:** To modify the default values for this probe, coordinate with the administrator responsible for the vCenter management. For more information, see VMware documentation.

*For more information, see VMware tags.

Consider the following when setting these parameters:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Disable**: You disable a probe by setting the parameter in the triggering probe. For example, if you're not interested in discovering storage, set the *disable_host_storage_probe* parameter to **true** in the VMWare - vCenter ESX Hosts probe.
- **Page size**: Page size parameters control the number of CIs to discover with a single probe. Use this parameter to limit payload size by reducing the number of vCenter elements discovered at a time by any probe. The page size expressed in parentheses is the default in the base system.
- **Debug**: Set the debug parameter in the VMWare - vCenter Datacenters probe to enable debugging for all the vCenter probes. Debugging returns the raw vCenter data in each probe payload.

Trigger custom probes with the vCenter Discovery extension

Use the vCenter Discovery extension to collect additional attributes of discovered configuration item (CI) types or attributes of new CI types. The extension triggers custom probes from an existing sensor.

Before you begin

Role required: admin

About this task

vCenter sensor records support triggering custom probes that use conditional scripts for gathering data not typically collected by Discovery or for discovering CI types that you create.

> **Important:** The vCenter extension is intended to be used to create a single custom probe that extends existing Discovery functionality. Unless you're a user who is familiar with advanced scripting, don't attempt to configure your new probe/sensor to launch an additional probe.

Procedure

1. Navigate to **All** > **Discovery Definition** > **Sensors**.
2. Select a vCenter sensor record.
   It is important to trigger a new probe at the correct time. For example, if you want to discover additional information about virtual machines in your network, trigger your probe from the VMWare - vCenter VMs sensor. If you trigger your custom probe from the VMWare - vCenter Datacenters sensor, your new probe and the existing VM probe execute in parallel. This might help prevent VM records from being created or updated.
3. Select the **Triggers probes** related list.
4. Select **New** to create probe to be triggered by this sensor or **Edit** to add an existing probe to the list.
   You can also use an existing probe record as a template for your new probe by opening the record and selecting the **Insert and Stay** option from the context menu.
5. Select the information icon to the right of the probe name to open the **Conditional Probes Triggered by Sensor** record.
6. View the script.
   The script passes the **parms** and **data** arguments to the function, using values passed to them from the triggering sensor. The script configures the **parms** object, from which the probe parameters are derived. The script returns **true** to trigger the probe or **false** to cancel it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In this example, the VMWare – vCenter Datacenters sensor **parms** are all probe parameters that are supported by the Datacenters probe. The **data** argument is an object containing up to 4 arrays, named **vm**, **datastore**, **network**, and **cluster**.

Conditional Probes Triggered by Sensor record

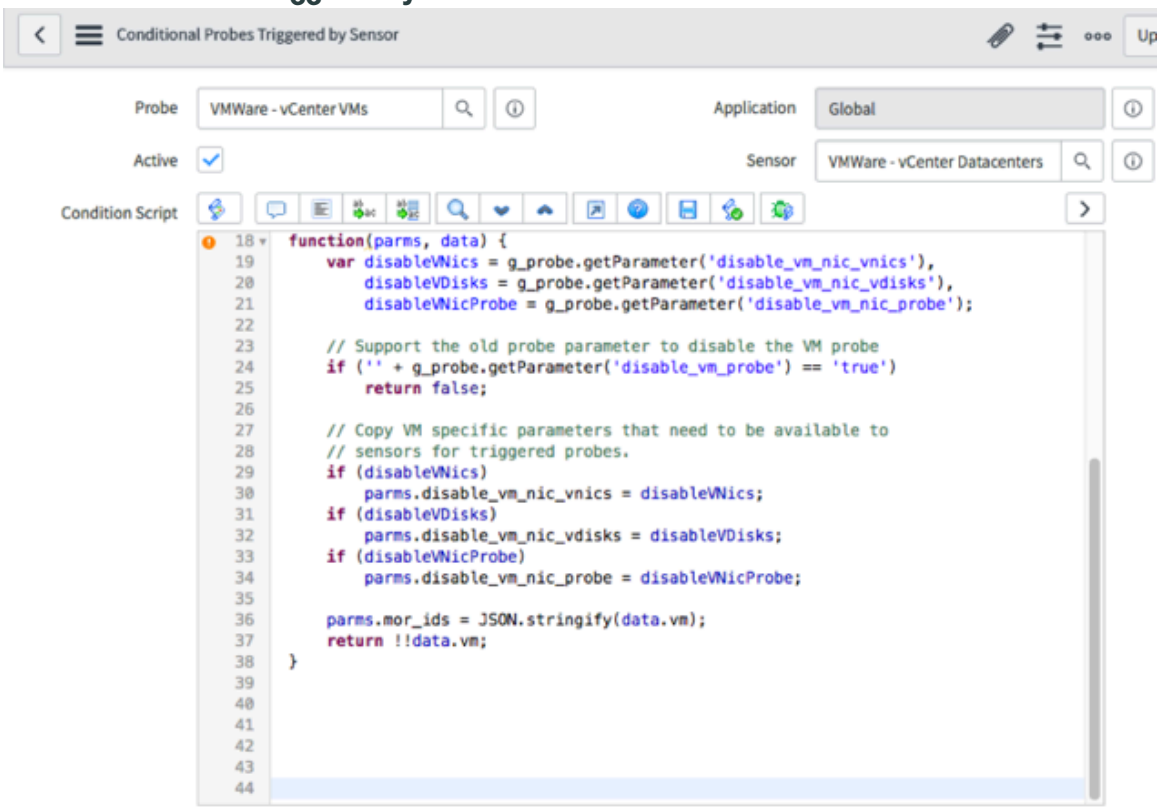

Windows probes and permissions

Discovery accesses devices and software by executing commands as a specific user on Windows computers.

Important: Use patterns to run Discovery on Windows devices. The legacy Windows probes on this page are outdated and should no longer be used for Discovery. To use patterns, see Add the Horizontal Pattern probe to a classifier.

Most probes require access to Windows classes, properties, and registry entries. Certain probes also require additional access to Windows directories and resources. Security policies vary by organization, so there's no one specific role or right to grant. Verify that the Windows user has `local admin` permission for these Windows components.

Administrative shares

Windows administrative shares are hidden server resources that Discovery uses to store temporarily the results of processes run by specific probes. The MID Server script file *`LaunchProc.psm1`* launches the process, writes its output to the administrative share on the machine, and then retrieves the results for the MID Server. Access to administrative shares is restricted to users with administrative rights to the target. For example: Users, local or on the domain, who belong to the local administrators group.

With the New York release, all Windows probes using the WMI protocol call the *`LaunchProc.psm1`* script file and write to the admin$ share folder as default. If another

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

network share is mounted on each Windows target, the folder could be changed by updating the MID Server property *mid.powershell.target_base_dir*.

Windows Classes

Most Windows probes access Windows classes and properties contained in those classes. Some Windows classes don't specify a file namespace path. If an explicit path isn't specified, ServiceNow uses `root\cimv2\<Class>` by default.

Note: Discovery of Cluster VIP isn't supported.

Windows Classes

Probe	Windows Class	Property
Hyper-V - Cluster	`root\MSCluster\MSCluster_ClusterToNode`	**Antecedent**
Hyper-V - Cluster	`root\MSCluster\MSCluster_ClusterToNode`	**Dependent**
Hyper-V - Cluster	`root\MSCluster\MSCluster_Resource`	**Type**
Hyper-V - Resource Pools	`root\virtualization\Msvm_ResourcePool`	**AllocationUnits**
Hyper-V - Resource Pools	`root\virtualization\Msvm_ResourcePool`	**Capacity**
Hyper-V - Resource Pools	`root\virtualization\Msvm_ResourcePool`	**InstanceID**
Hyper-V - Resource Pools	`root\virtualization\Msvm_ResourcePool`	**Name**
Hyper-V - Resource Pools	`root\virtualization\Msvm_ResourcePool`	**ResourceSubType**
Hyper-V - Resource Pools	`root\virtualization\v2\Msvm_ResourcePool`	**AllocationUnits**
Hyper-V - Resource Pools	`root\virtualization\v2\Msvm_ResourcePool`	**Capacity**
Hyper-V - Resource Pools	`root\virtualization\v2\Msvm_ResourcePool`	**InstanceID**
Hyper-V - Resource Pools	`root\virtualization\v2\Msvm_ResourcePool`	**Name**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Hyper-V - Resource Pools	root\virtualization\v2\Msvm_ResourcePool	**ResourceSubType**
Hyper-V - Virtual Machines	root\virtualization\Msvm_ComputerSystem	**Description**
Hyper-V - Virtual Machines	root\virtualization\Msvm_ComputerSystem	**ElementName**
Hyper-V - Virtual Machines	root\virtualization\Msvm_ComputerSystem	**EnabledState**
Hyper-V - Virtual Machines	root\virtualization\Msvm_ComputerSystem	**Name**
Hyper-V - Virtual Machines	root\virtualization\Msvm_ComputerSystem	**TimeOfLastStateChange**
Hyper-V - Virtual Machines	root\virtualization\Msvm_LogicalDisk	**BlockSize**
Hyper-V - Virtual Machines	root\virtualization\Msvm_LogicalDisk	**ElementName**
Hyper-V - Virtual Machines	root\virtualization\Msvm_LogicalDisk	**NumberOfBlocks**
Hyper-V - Virtual Machines	root\virtualization\Msvm_LogicalDisk	**SystemName**
Hyper-V - Virtual Machines	root\virtualization \Msvm_MemorySettingData	**AllocationUnits**
Hyper-V - Virtual Machines	root\virtualization \Msvm_MemorySettingData	**InstanceID**
Hyper-V - Virtual Machines	root\virtualization \Msvm_MemorySettingData	**VirtualQuantity**
Hyper-V - Virtual Machines	root\virtualization \Msvm_ProcessorSettingData	**InstanceID**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Hyper-V - Virtual Machines	root\virtualization \Msvm_ProcessorSettingData	**VirtualQuantity**
Hyper-V - Virtual Machines	root\virtualization \Msvm_SyntheticEthernetPortSettingData	**ElementName**
Hyper-V - Virtual Machines	root\virtualization \Msvm_SyntheticEthernetPortSettingData	**InstanceID**
Hyper-V - Virtual Machines	root\virtualization \Msvm_VirtualSystemSettingData	**BaseBoardSerialNumber**
Hyper-V - Virtual Machines	root\virtualization \Msvm_VirtualSystemSettingData	**BIOSGUID**
Hyper-V - Virtual Machines	root\virtualization \Msvm_VirtualSystemSettingData	**BIOSSerialNumber**
Hyper-V - Virtual Machines	root\virtualization \Msvm_VirtualSystemSettingData	**ChassisSerialNumber**
Hyper-V - Virtual Machines	root\virtualization \Msvm_VirtualSystemSettingData	**InstanceID**
Hyper-V - Virtual Machines	root\virtualization \Msvm_VirtualSystemSettingData	**SystemName**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_ComputerSystem	**Description**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_ComputerSystem	**ElementName**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_ComputerSystem	**EnabledState**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_ComputerSystem	**Name**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_ComputerSystem	**TimeOfLastStateChange**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Hyper-V - Virtual Machines	root\virtualization\v2\Msvm_LogicalDisk	**BlockSize**
Hyper-V - Virtual Machines	root\virtualization\v2\Msvm_LogicalDisk	**ElementName**
Hyper-V - Virtual Machines	root\virtualization\v2\Msvm_LogicalDisk	**NumberOfBlocks**
Hyper-V - Virtual Machines	root\virtualization\v2\Msvm_LogicalDisk	**SystemName**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_MemorySettingData	**AllocationUnits**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_MemorySettingData	**InstanceID**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_MemorySettingData	**VirtualQuantity**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_ProcessorSettingData	**InstanceID**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_ProcessorSettingData	**VirtualQuantity**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_SyntheticEthernetPortSettingData	**ElementName**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_SyntheticEthernetPortSettingData	**InstanceID**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_VirtualSystemSettingData	**BaseBoardSerialNumber**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_VirtualSystemSettingData	**BIOSGUID**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_VirtualSystemSettingData	**BIOSSerialNumber**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_VirtualSystemSettingData	**ChassisSerialNumber**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_VirtualSystemSettingData	**InstanceID**
Hyper-V - Virtual Machines	root\virtualization \v2\Msvm_VirtualSystemSettingData	**SystemName**
Hyper-V - Virtual Networks	root\virtualization\Msvm_ActiveConnection	**Antecedent**
Hyper-V - Virtual Networks	root\virtualization\Msvm_ActiveConnection	**Dependent**
Hyper-V - Virtual Networks	root\virtualization \Msvm_VirtualEthernetSwitch	**ElementName**
Hyper-V - Virtual Networks	root\virtualization \Msvm_VirtualEthernetSwitch	**Name**
Hyper-V - Virtual Networks	root\virtualization\Msvm_VirtualSwitch	**ElementName**
Hyper-V - Virtual Networks	root\virtualization\Msvm_VirtualSwitch	**Name**
Hyper-V - Virtual Networks	root\virtualization \v2\Msvm_ActiveConnection	**Antecedent**
Hyper-V - Virtual Networks	root\virtualization \v2\Msvm_ActiveConnection	**Dependent**
Hyper-V - Virtual Networks	root\virtualization \v2\Msvm_VirtualEthernetSwitch	**ElementName**
Hyper-V - Virtual Networks	root\virtualization \v2\Msvm_VirtualEthernetSwitch	**Name**
Hyper-V - Virtual Networks	root\virtualization\v2\Msvm_VirtualSwitch	**ElementName**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Hyper-V - Virtual Networks	root\virtualization\v2\Msvm_VirtualSwitch	**Name**
Windows - Active Connections Session	Win32_OperatingSystem	**Name**
Windows - Active Processes	Win32_Process	**Caption**
Windows - Active Processes	Win32_Process	**CommandLine**
Windows - Active Processes	Win32_Process	**CreationDate**
Windows - Active Processes	Win32_Process	**Description**
Windows - Active Processes	Win32_Process	**ExecutablePath**
Windows - Active Processes	Win32_Process	**Name**
Windows - Active Processes	Win32_Process	**ParentProcessId**
Windows - Active Processes	Win32_Process	**Priority**
Windows - Active Processes	Win32_Process	**ProcessId**
Windows - Amazon EC2	Win32_ComputerSystemProduct	**UUID**
Windows - Azure Relationship	Win32_ComputerSystemProduct	**UUID**
Windows - Classify	root\MSCluster\MSCluster_Cluster	**Name**
Windows - Classify	root\MSCluster\MSCluster_ClusterToNode	**Antecedent**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Classify	root\MSCluster\MSCluster_ClusterToNode	**Dependent**
Windows - Classify	root\MSCluster \MSCluster_ClusterToResource	**GroupComponent**
Windows - Classify	root\MSCluster \MSCluster_ClusterToResource	**PartComponent**
Windows - Classify	root\MSCluster\MSCluster_Node	**Name**
Windows - Classify	root\MSCluster\MSCluster_Resource	**Name**
Windows - Classify	root\MSCluster\MSCluster_Resource	**PrivateProperties**
Windows - Classify	root\MSCluster\MSCluster_Resource	**Type**
Windows - Classify	root\virtualization\Msvm_ComputerSystem	**Name**
Windows - Classify	root\virtualization \v2\Msvm_ComputerSystem	**Name**
Windows - Classify	Win32_ComputerSystem	**Domain**
Windows - Classify	Win32_ComputerSystem	**Name**
Windows - Classify	Win32_OperatingSystem	**Caption**
Windows - Classify	Win32_OperatingSystem	**Version**
Windows - Cluster	root\MSCluster\MSCluster_Cluster	**Caption**
Windows - Cluster	root\MSCluster\MSCluster_Cluster	**Description**
Windows - Cluster	root\MSCluster\MSCluster_Cluster	**Name**
Windows - Cluster	root\MSCluster\MSCluster_Cluster	**Status**
Windows - Cluster	root\MSCluster\MSCluster_ClusterToNetwork	**GroupComponent**
Windows - Cluster	root\MSCluster\MSCluster_ClusterToNetwork	**PartComponent**
Windows - Cluster	root\MSCluster \MSCluster_ClusterToNetworkInterface	**GroupComponent**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Cluster	root\MSCluster \MSCluster_ClusterToNetworkInterface	**PartComponent**
Windows - Cluster	root\MSCluster\MSCluster_ClusterToNode	**Antecedent**
Windows - Cluster	root\MSCluster\MSCluster_ClusterToNode	**Dependent**
Windows - Cluster	root\MSCluster \MSCluster_ClusterToQuorumResource	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_ClusterToQuorumResource	**PartComponent**
Windows - Cluster	root\MSCluster \MSCluster_ClusterToResource	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_ClusterToResource	**PartComponent**
Windows - Cluster	root\MSCluster \MSCluster_ClusterToResourceGroup	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_ClusterToResourceGroup	**PartComponent**
Windows - Cluster	root\MSCluster\MSCluster_Event	**Name**
Windows - Cluster	root\MSCluster\MSCluster_Event	**Type**
Windows - Cluster	root\MSCluster\MSCluster_Network	**Address**
Windows - Cluster	root\MSCluster\MSCluster_Network	**AddressMask**
Windows - Cluster	root\MSCluster\MSCluster_Network	**Description**
Windows - Cluster	root\MSCluster\MSCluster_Network	**Name**
Windows - Cluster	root\MSCluster\MSCluster_NetworkInterface	**Address**
Windows - Cluster	root\MSCluster\MSCluster_NetworkInterface	**Description**
Windows - Cluster	root\MSCluster\MSCluster_NetworkInterface	**DeviceID**
Windows - Cluster	root\MSCluster\MSCluster_NetworkInterface	**Name**
Windows - Cluster	root\MSCluster\MSCluster_NetworkInterface	**Network**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Cluster	root\MSCluster\MSCluster_Node	**Caption**
Windows - Cluster	root\MSCluster\MSCluster_Node	**Characteristics**
Windows - Cluster	root\MSCluster\MSCluster_Node	**Description**
Windows - Cluster	root\MSCluster\MSCluster_Node	**Flags**
Windows - Cluster	root\MSCluster\MSCluster_Node	**Name**
Windows - Cluster	root\MSCluster\MSCluster_Node	**NodeInstanceID**
Windows - Cluster	root\MSCluster\MSCluster_Node	**PrivateProperties**
Windows - Cluster	root\MSCluster\MSCluster_Node	**State**
Windows - Cluster	root\MSCluster\MSCluster_Node	**Status**
Windows - Cluster	root\MSCluster \MSCluster_NodeToActiveGroup	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_NodeToActiveGroup	**PartComponent**
Windows - Cluster	root\MSCluster \MSCluster_NodeToActiveResource	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_NodeToActiveResource	**PartComponent**
Windows - Cluster	root\MSCluster \MSCluster_NodeToNetworkInterface	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_NodeToNetworkInterface	**PartComponent**
Windows - Cluster	root\MSCluster\MSCluster_Resource	**Caption**
Windows - Cluster	root\MSCluster\MSCluster_Resource	**Description**
Windows - Cluster	root\MSCluster\MSCluster_Resource	**Name**
Windows - Cluster	root\MSCluster\MSCluster_Resource	**PrivateProperties**
Windows - Cluster	root\MSCluster\MSCluster_Resource	**Status**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Cluster	root\MSCluster\MSCluster_Resource	**Type**
Windows - Cluster	root\MSCluster\MSCluster_ResourceGroup	**Description**
Windows - Cluster	root\MSCluster\MSCluster_ResourceGroup	**Name**
Windows - Cluster	root\MSCluster\MSCluster_ResourceGroup	**PrivateProperties**
Windows - Cluster	root\MSCluster\MSCluster_ResourceGroup	**State**
Windows - Cluster	root\MSCluster \MSCluster_ResourceGroupToResource	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_ResourceGroupToResource	**PartComponent**
Windows - Cluster	root\MSCluster \MSCluster_ResourceToDependentResource	**Antecedent**
Windows - Cluster	root\MSCluster \MSCluster_ResourceToDependentResource	**Dependent**
Windows - Cluster	root\MSCluster \MSCluster_ResourceToPossibleOwner	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_ResourceToPossibleOwner	**PartComponent**
Windows - Cluster	root\MSCluster\MSCluster_ResourceType	**Description**
Windows - Cluster	root\MSCluster\MSCluster_ResourceType	**DisplayName**
Windows - Cluster	root\MSCluster\MSCluster_ResourceType	**Name**
Windows - Cluster	root\MSCluster\MSCluster_ResourceType	**PrivateProperties**
Windows - Cluster	root\MSCluster \MSCluster_ResourceTypeToResource	**GroupComponent**
Windows - Cluster	root\MSCluster \MSCluster_ResourceTypeToResource	**PartComponent**
Windows - Cluster	Win32_NetworkAdapterConfiguration	**IPAddress**
Windows - Cluster	Win32_NetworkAdapterConfiguration	**IPEnabled**
Windows - CPU / Memory	Win32_PhysicalMemory	**BankLabel**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - CPU / Memory	Win32_PhysicalMemory	**Capacity**
Windows - CPU / Memory	Win32_PhysicalMemory	**DataWidth**
Windows - CPU / Memory	Win32_PhysicalMemory	**DeviceLocator**
Windows - CPU / Memory	Win32_PhysicalMemory	**FormFactor**
Windows - CPU / Memory	Win32_PhysicalMemory	**Manufacturer**
Windows - CPU / Memory	Win32_PhysicalMemory	**MemoryType**
Windows - CPU / Memory	Win32_PhysicalMemory	**PartNumber**
Windows - CPU / Memory	Win32_PhysicalMemory	**SerialNumber**
Windows - CPU / Memory	Win32_PhysicalMemory	**Speed**
Windows - CPU / Memory	Win32_PhysicalMemory	**Status**
Windows - CPU / Memory	Win32_PhysicalMemory	**Tag**
Windows - CPU / Memory	Win32_PhysicalMemory	**TotalWidth**
Windows - CPU / Memory	Win32_PhysicalMemory	**TypeDetail**
Windows - CPU / Memory	Win32_Processor	**Manufacturer**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - CPU / Memory	Win32_Processor	**MaxClockSpeed**
Windows - CPU / Memory	Win32_Processor	**Name**
Windows - CPU / Memory	Win32_Processor	**NumberOfCores**
Windows - CPU / Memory	Win32_Processor	**NumberOfLogicalProcessors**
Windows - Get IIS Information	root\MicrosoftIISv2\IIsWebServerSetting	**LogFileDirectory**
Windows - Get IIS Information	root\MicrosoftIISv2\IIsWebServerSetting	**Name**
Windows - Get IIS Information	root\MicrosoftIISv2\IIsWebServerSetting	**SecureBindings**
Windows - Get IIS Information	root\MicrosoftIISv2\IIsWebServerSetting	**ServerBindings**
Windows - Get IIS Information	root\MicrosoftIISv2\IIsWebServerSetting	**ServerComment**
Windows - Get IIS Information	root\WebAdministration\Site	**Bindings**
Windows - Get IIS Information	root\WebAdministration\Site	**Id**
Windows - Get IIS Information	root\WebAdministration\Site	**LogFile**
Windows - Get IIS Information	root\WebAdministration\Site	**Name**
Windows - Get Physical Address for VIP	root\MSCluster\MSCluster_NetworkInterface	**Address**
Windows - Get Physical	Win32_NetworkAdapterConfiguration	**IPAddress**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Address for VIP		
Windows - Get Physical Address for VIP	`Win32_NetworkAdapterConfiguration`	**IPEnabled**
Windows - Hardware Information	`Win32_BaseBoard`	**SerialNumber**
Windows - Hardware Information	`Win32_BIOS`	**SerialNumber**
Windows - Hardware Information	`Win32_ComputerSystemProduct`	**IdentifyingNumber**
Windows - Hardware Information	`Win32_ComputerSystemProduct`	**UUID**
Windows - Hardware Information	`Win32_SystemEnclosure`	**ChassisTypes**
Windows - Hardware Information	`Win32_SystemEnclosure`	**SerialNumber**
Windows - Hardware & Network Info	`Win32_BaseBoard`	**SerialNumber**
Windows - Hardware & Network Info	`Win32_BIOS`	**SerialNumber**
Windows - Hardware & Network Info	`Win32_ComputerSystemProduct`	**IdentifyingNumber**
Windows - Hardware & Network Info	`Win32_ComputerSystemProduct`	**UUID**
Windows - Hardware & Network Info	`Win32_IP4RouteTable`	**Destination**
Windows - Hardware & Network Info	`Win32_IP4RouteTable`	**Mask**
Windows - Hardware & Network Info	`Win32_IP4RouteTable`	**NextHop**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Hardware & Network Info	Win32_NetworkAdapter	**Index**
Windows - Hardware & Network Info	Win32_NetworkAdapter	**Manufacturer**
Windows - Hardware & Network Info	Win32_NetworkAdapter	**NetConnectionID**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**Caption**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**DefaultIPGateway**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**DHCPEnabled**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**Index**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**IPAddress**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**IPEnabled**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**IPSubnet**
Windows - Hardware & Network Info	Win32_NetworkAdapterConfiguration	**MACAddress**
Windows - Hardware & Network Info	Win32_SystemEnclosure	**ChassisTypes**
Windows - Hardware & Network Info	Win32_SystemEnclosure	**SerialNumber**
Windows - Network	Win32_IP4RouteTable	**Destination**
Windows - Network	Win32_IP4RouteTable	**Mask**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Network	Win32_IP4RouteTable	**NextHop**
Windows - Network	Win32_NetworkAdapterConfiguration	**Caption**
Windows - Network	Win32_NetworkAdapter	**Index**
Windows - Network	Win32_NetworkAdapter	**Manufacturer**
Windows - Network	Win32_NetworkAdapter	**NetConnectionID**
Windows - Network	Win32_NetworkAdapterConfiguration	**DefaultIPGateway**
Windows - Network	Win32_NetworkAdapterConfiguration	**DHCPEnabled**
Windows - Network	Win32_NetworkAdapterConfiguration	**Index**
Windows - Network	Win32_NetworkAdapterConfiguration	**IPAddress**
Windows - Network	Win32_NetworkAdapterConfiguration	**IPEnabled**
Windows - Network	Win32_NetworkAdapterConfiguration	**IPSubnet**
Windows - Network	Win32_NetworkAdapterConfiguration	**MACAddress**
Windows - OS Information	Win32_ComputerSystem	**Manufacturer**
Windows - OS Information	Win32_ComputerSystem	**Model**
Windows - OS Information	Win32_ComputerSystem	**UserName**
Windows - OS Information	Win32_OperatingSystem	**CSDVersion**
Windows - OS Information	Win32_OperatingSystem	**Description**
Windows - OS Information	Win32_OperatingSystem	**Version**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - OS Information	Win32_Processor	**AddressWidth**
Windows - Printers	Win32_TCPIPPrinterPort	**HostAddress**
Windows - Printers	Win32_Printer	**DeviceID**
Windows - Printers	Win32_TCPIPPrinterPort	**Name**
Windows - Printers	Win32_Printer	**Name**
Windows - Printers	Win32_Printer	**PortName**
Windows - Services	Win32_Service	**AcceptPause**
Windows - Services	Win32_Service	**AcceptStop**
Windows - Services	Win32_Service	**DesktopInteract**
Windows - Services	Win32_Service	**DisplayName**
Windows - Services	Win32_Service	**Name**
Windows - Services	Win32_Service	**PathName**
Windows - Services	Win32_Service	**ProcessId**
Windows - Services	Win32_Service	**ServiceType**
Windows - Services	Win32_Service	**StartMode**
Windows - Services	Win32_Service	**StartName**
Windows - Services	Win32_Service	**State**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**Caption**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**DeviceID**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Storage 2008 - WMI	Win32_DiskDrive	**Index**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**InterfaceType**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**Manufacturer**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**Model**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**SCSIBus**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**SCSILogicalUnit**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**SCSIPort**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**SCSITargetId**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**SerialNumber**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**Size**
Windows - Storage 2008 - WMI	Win32_DiskDrive	**Status**
Windows - Storage 2008 - WMI	Win32_DiskPartition	**DiskIndex**
Windows - Storage 2008 - WMI	Win32_DiskPartition	**Index**
Windows - Storage 2008 - WMI	Win32_DiskPartition	**Size**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Storage 2008 - WMI	Win32_DiskPartition	**StartingOffset**
Windows - Storage 2008 - WMI	Win32_LogicalDisk	**DeviceID**
Windows - Storage 2008 - WMI	Win32_LogicalDisk	**VolumeSerialNumber**
Windows - Storage 2008 - WMI	Win32_LogicalDiskToPartition	**Antecedent**
Windows - Storage 2008 - WMI	Win32_LogicalDiskToPartition	**Dependent**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**Description**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**DeviceID**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**DriveType**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**FileSystem**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**FreeSpace**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**ProviderName**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**Size**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**VolumeName**
Windows - Storage 2008 - WMI	Win32_MappedLogicalDisk	**VolumeSerialNumber**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Storage 2008 - WMI	Win32_MountPoint	**Directory**
Windows - Storage 2008 - WMI	Win32_MountPoint	**Volume**
Windows - Storage 2008 - WMI	Win32_Volume	**Capacity**
Windows - Storage 2008 - WMI	Win32_Volume	**DeviceID**
Windows - Storage 2008 - WMI	Win32_Volume	**DriveLetter**
Windows - Storage 2008 - WMI	Win32_Volume	**DriveType**
Windows - Storage 2008 - WMI	Win32_Volume	**FileSystem**
Windows - Storage 2008 - WMI	Win32_Volume	**FreeSpace**
Windows - Storage 2008 - WMI	Win32_Volume	**Label**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**AllocatedSize**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**BusType**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**FriendlyName**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**Manufacturer**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**Model**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**Number**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**ObjectId**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**OperationalStatus**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**SerialNumber**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage\MSFT_Disk	**Size**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_DiskToPartition	**Disk**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_DiskToPartition	**Partition**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_iSCSISession	**InitiatorNodeAddress**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_iSCSISession	**SessionIdentifier**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_iSCSISession	**TargetNodeAddress**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_iSCSISessionToDisk	**Disk**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_iSCSISessionToDisk	**iSCSISession**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Partition	**AccessPaths**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Partition	**DiskId**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Partition	**ObjectId**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Partition	**Offset**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Partition	**OperationalStatus**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Partition	**PartitionNumber**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Partition	**Size**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_PartitionToVolume	**Partition**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_PartitionToVolume	**Volume**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Volume	**DriveLetter**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Volume	**DriveType**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Volume	**FileSystem**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Volume	**FileSystemLabel**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Volume	**ObjectId**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Volume	**Size**
Windows - Storage 2012 - WMI	root\microsoft\windows\storage \MSFT_Volume	**SizeRemaining**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Storage 2012 - WMI	Win32_DiskDrive	**Name**
Windows - Storage 2012 - WMI	Win32_DiskDrive	**SCSIBus**
Windows - Storage 2012 - WMI	Win32_DiskDrive	**SCSILogicalUnit**
Windows - Storage 2012 - WMI	Win32_DiskDrive	**SCSIPort**
Windows - Storage 2012 - WMI	Win32_DiskDrive	**SCSITargetId**
Windows - Storage 2012 - WMI	Win32_LogicalDisk	**DeviceID**
Windows - Storage 2012 - WMI	Win32_LogicalDisk	**VolumeSerialNumber**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**Description**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**DeviceID**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**DriveType**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**FileSystem**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**FreeSpace**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**ProviderName**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**Size**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Classes (continued)

Probe	Windows Class	Property
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**VolumeName**
Windows - Storage 2012 - WMI	Win32_MappedLogicalDisk	**VolumeSerialNumber**
Windows - Storage 2012 - WMI	Win32_MountPoint	**Directory**
Windows - Storage 2012 - WMI	Win32_MountPoint	**Volume**

Windows Registry entries

Several Windows Registry entries are available for Discovery Windows probes.

Note: The Windows - Installed Software sensor appends a timestamp of 00:00:00 to the install_date retrieved from the registry. The installation time of all Windows software is independent of the time zone and is set to midnight of the day it was installed. For example, an install date of **2.19.2017** in the Windows registry appears as **2.19.2017 00:00:00** in the CMDB.

Probe	Windows Registry Entries
Windows - Active Connections Session	HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services/Tcpip/Parameters/Hostname
Windows - Classify	HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services/Tcpip/Parameters/Hostname
Windows - Classify	HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services/Tcpip/Parameters/Domain
Windows - Get IIS Information	HKEY_LOCAL_MACHINE/SOFTWARE/Microsoft/InetStp/MajorVersion
Windows - Get IIS Information	HKEY_LOCAL_MACHINE/SOFTWARE/Microsoft/InetStp/MinorVersion
Windows - Find APD File Location	HKEY_LOCAL_MACHINE/SOFTWARE/APD/APD/CONFIGPATH
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/InstallDate

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	Windows Registry Entries
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Uninstall/*/DisplayVersion
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows NT/CurrentVersion/ProductId
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/ProductId
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/InstallDate
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Uninstall/*/DisplayName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Uninstall/*/ParentDisplayName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/DisplayVersion
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Uninstall/*/UninstallString
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/Publisher
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Internet Explorer/Version
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/DisplayName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Uninstall/*/Publisher
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/ProductID
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Office/*/Registration/*/ProductID
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/ProductID

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	Windows Registry Entries
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Uninstall/*/DisplayVersion
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/DisplayVersion
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Office/*/Registration/*/DigitalProductID
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/Publisher
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Uninstall/*/DisplayName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Installer/UserData/*/Products/*/InstallProperties/DisplayName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Uninstall/*/UninstallString
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Office/*/Registration/*/DigitalProductID
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Internet Explorer/Registration/ProductId
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Uninstall/*/Publisher
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Internet Explorer/svcVersion
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Office/*/Registration/*/ProductID
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Uninstall/*/InstallDate
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/Windows/CurrentVersion/Uninstall/*/ParentDisplayName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/CurrentVersion/Uninstall/*/InstallDate

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	Windows Registry Entries
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows NT/CurrentVersion/DigitalProductID
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Component Based Servicing/Packages/ Package_*_KB*/InstallName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Component Based Servicing/Packages/ Package_*_KB*/InstallTimeHigh
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Component Based Servicing/Packages/ Package_*_KB*/InstallTimeLow
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Component Based Servicing/Packages/ Package_for_RollupFix*/InstallLocation
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Component Based Servicing/Packages/ Package_for_RollupFix*/InstallName
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Component Based Servicing/Packages/ Package_for_RollupFix*/InstallTimeHigh
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Component Based Servicing/Packages/ Package_for_RollupFix*/InstallTimeLow
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Installer/UserData/*/Products/ */InstallProperties/InstallLocation
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Microsoft/Windows/ CurrentVersion/Uninstall/*/InstallLocation
Windows - Installed Software	HKEY_LOCAL_MACHINE/SOFTWARE/WOW6432Node/Microsoft/ InetStp/SetupString
Windows - Installed Software	HKEY_LOCAL_MACHINE/SOFTWARE/WOW6432Node/Microsoft/ InetStp/VersionString
Windows - Installed Software	HKEY_LOCAL_MACHINE/Software/Wow6432Node/Microsoft/ Windows/CurrentVersion/Uninstall/*/InstallLocation
Windows - Installed Software	HKEY_USERS/*/SOFTWARE/Adobe/Adobe Acrobat/ DC/AVEntitlement/iEntitlementLevel
Windows - Installed Software	HKEY_USERS/*/SOFTWARE/Microsoft/Windows/ CurrentVersion/Uninstall/*/DisplayName

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	Windows Registry Entries
Windows - Installed Software	HKEY_USERS/*/SOFTWARE/Microsoft/Windows/CurrentVersion/Uninstall/*/DisplayVersion
Windows - Installed Software	HKEY_USERS/*/SOFTWARE/Microsoft/Windows/CurrentVersion/Uninstall/*/ParentDisplayName
Windows - Installed Software	HKEY_USERS/*/SOFTWARE/Microsoft/Windows/CurrentVersion/Uninstall/*/Publisher
Windows - Installed Software	HKEY_USERS/*/SOFTWARE/Microsoft/Windows/CurrentVersion/Uninstall/*/UninstallString

Related topics

Discovery commands for probes and patterns

WMIRunner probe

WMI Runner is a probe type that fetches data from Windows operating systems via the Windows Management Instrumentation (WMI) interface. IPv6 supports WMI Runner.

The probe handles multiple user-specified WMI Paths to be queried, using a basic form of native WMI query. Each field to be probed must be uniquely named (within the domain of the probe). The probe results returned to the sensor will provide the data found for each field queried, indexed by its name.

When creating a WMI probe, the probe type must be set to WMI Probe and the ECC Queue Topic must be set to WMIRunner.

For instructions on configuring probe parameters, see Set probe parameters.

The following parameters may be passed to the WMI Probe:

WMIRunner probe parameters

Parameter	Description	Default value
source	Host to connect to.	
port	Port to connect to.	
debug	Enables debug logging.	false
wmi_timeout	Timeout for the actual WMI probe, in seconds. Use this parameter to change the timeout interval for individual Windows probes. This value overrides the value in the windows_probe_timeout MID Server parameter, which sets a timeout for all probes launched by a specific MID Server. Windows - Installed Software probe is configured	300 (seconds)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

WMIRunner probe parameters (continued)

Parameter	Description	Default value
	with a timeout value of 15 minutes.	
process_timeout	Timeout for the process running the script, in seconds. This parameter is for internal use only and is not supported.	wmi_timeout + 10 (seconds)
credentials_debug	Displays a <credentials_debug> section in the ECC queue, which can help you troubleshoot credentials. If you set this property to true, credential troubleshooting information is output to the ECC queue, even if the credentials succeed. See Credentials troubleshooting for more information.	false

Note: The default timeout for WMI/Powershell is 5 minutes, except for the Windows Installed Software probe, which has a default timeout value of 15 minutes. Adding *wmi_timeout* to a probe parameter can change the default timeout of a Windows probe.

Port probes

Port probes are used in Discovery by the Shazzam probe to detect protocol activity on open ports on devices it encounters.

When a port probe encounters a protocol in use, the Shazzam sensor checks the port probe record to determine which classification probe to launch. The common protocols WMI, SSH, SNMP, and HTTP in the base system have priority numbers that control the order in which they are launched.

The priority is as follows:

- 1 - WMI
- 2 - SSH
- 3 - SNMP
- 4 - HTTP

In the base system, the WMI probe is always launched first, and if it is successful on a device, no other port probes are launched for that device. If the WMI probe is not successful, then the SSH probe is launched to gather information on the device. If it is not successful, the SNMP probe is launched. This method allows Discovery to classify a device correctly if the device is running more than one protocol (for example, SSH, SNMP, and HTTP).

Discovery Port Probe form

To access the Port Probe form, navigate to **Discovery Definition** > **Port Probes**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To add multiple classification probes to a port probe, create a link between the port probe and the actual classification probe itself. See the bottom of the form to add additional Trigger probes. That way, if one classification fails, it does not affect the others, thus Discovery performance may be enhanced.

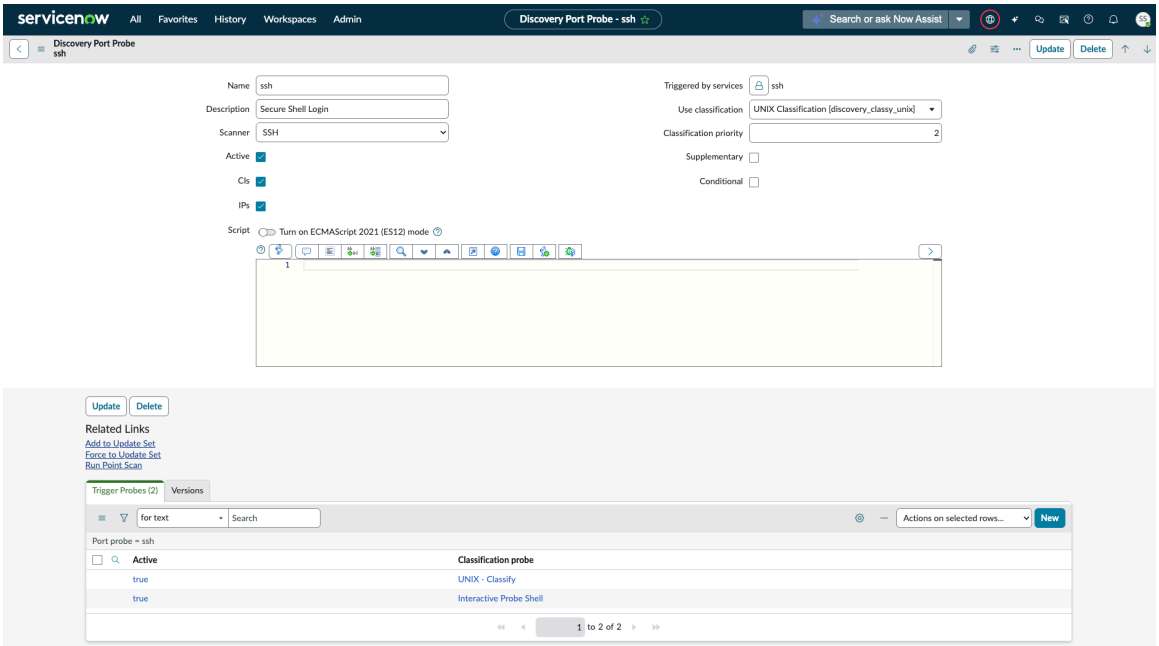

The Port Probe form provides the following fields:

Port probes

Field	Input Value
Name	Simple name for the port probe that reflects its function (for example, SNMP).
Description	Definition of the acronym for the protocol. (For example, SSH is Secure Shell login).
Scanner	Shazzam techniques for exploring a port. Some of these are protocol-specific, and others are generic. For example, a WMI port probe uses a Scanner value of Generic TCP, and the SNMP port probe uses a value of SNMP.
Active	Indicates whether this port probe is active or inactive.
CIs	Indicates whether this port probe is active or inactive for discovering Configuration Items.
IPs	Indicates whether this port probe is active or inactive for discovering IP addresses.
Triggered by services	Indicates which services define the port usage. Use this setting to define non-standard port usage and pair the port number with the protocol.
Use classification	Names the appropriate classification table, based on the protocol being explored.
Classification priority	Establishes the priority in which this port probe runs. If the first port probe fails, then the next probe runs on the device, and so forth, until the correct data is returned. This allows for the proper classification of a device that has two running protocols, such as SSH and SNMP. The default priorities for the Discovery protocols are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Port probes (continued)

Field	Input Value
	• 1 - WMI • 2 - SSH • 3 - SNMP • 4 - HTTP
Supplementary	Launches supplementary classifications after a higher-priority identification succeeds, in order of priority.
Conditional	Runs this port probe if any one of the non-conditional probes returns an open port. The conditional port probes in the base system attempt to resolve the names of Windows devices and DNS names. These ports probes take additional resources and are not used unless activity is detected on open ports.
Script	Script to run.

Shazzam probe, port probes, and protocols

Port scanning is the first step in the discovery process. The Shazzam probe performs port scanning, regardless of whether you use patterns for horizontal discovery. The following table lists the known ports and protocols used by Discovery.

After upgrading to Discovery Admin Workspace version 1.3.1 (August 2024 Store), you can navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** and use the enhanced dashboard.

Several port probes are available in the base system. Each port probe uses an IP Service, which is a record that tells Discovery which port to use for a specific protocol. Review this table before you block any ports with a firewall.

Important: Make sure that you do not block any ports that Discovery needs.

Default port probes and default IP services

Default port probe name	Default classification	Default IP Service and port
dns	Process Classification [discovery_classy_proc]	dns (port 53)
http	HTTP Classification [discovery_classy_http]	http (port 80) and https (port 443)
ip_phone	SNMP Classification [discovery_classy_snmp]	sip (port 5060)
osx	Scan Results Application Classifier [discovery_classy_scan_app]	afp (port 548)
printer	Scan Results Application Classifier [discovery_classy_scan_app]	hp-pdl-datastr (port 9100) and printer (port 515)
slp	Process Classification [discovery_classy_proc]	slp (port 427)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Default port probes and default IP services (continued)

Default port probe name	Default classification	Default IP Service and port
snmp	SNMP Classification [discovery_classy_snmp]	snmp (port 161)
ssh	UNIX Classification [discovery_classy_unix]	ssh (port 22)
tls_ssl_certs		tls_ssl_certs • SSL ports: 443, 8443, 9443, 636 (ldaps), 993 (imaps), 995 (popssl), 989, 990 • StartTLS ports: 25 (smtp), 110, 143, 389, 21, 587 (smtp)
vmapp	Application Classification [discovery_classy_appl]	vmapp_https (port 5480) and vmapp6_https (port 9443)
wbem	CIM Classification [discovery_classy_cim]	wbem_https (port 5989)
winrm	Windows Classification [discovery_classy_windows]	winrm (port 5985) and winrm_ssl (port 5986)
wins	Process Classification [discovery_classy_proc]	ms-nb-ns (port 137)
wmi	Windows Classification [discovery_classy_windows]	epmap (port 135)

This table shows you other common ports and protocols that Discovery uses. All ports listed in the Default port probes and default IP services table are also included in the following table.

Discovery ports and protocols

Name	Service name	Port	Details	Creates	Protocol
afp	Apple File Protocol	548			TCP
BEA Weblogic		7001		cmdb_ci_app_server	TCP
dns	Domain Name Service	53	To resolve the name of each IP Address		TCP/ UDP
epmap	Microsoft RPC (WMI, DCOM)	135	Windows Systems		TCP
ftp		21			TCP
ftps	FTP over SSL (control channel)	990	tls_ssl_certs (SSL)		TCP
ftps-data	FTP over SSL (data channel)	989	tls_ssl_certs (SSL)		TCP
hp-pdl-datastr	Printer PDL Data Stream	9100	HP Printers		TCP

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery ports and protocols (continued)

Name	Service name	Port	Details	Creates	Protocol
http	HyperText Transfer Protocol	80	Web Servers	cmdb_ci_web_server	TCP
https	HyperText Transfer Protocol over Secure Socket	443	Secure Web Servers	cmdb_ci_web_server	TCP
https-alt	HTTPS alternate port	8443	tls_ssl_certs (SSL)		TCP
IBM DB2		50000			TCP
IBM MQSeries		1414			TCP
IBM Websphere		9080			TCP
IBM WebSphere SSL		9443			TCP
imap	Internet Message Access Protocol	143	StartTLS (tls_ssl_certs)		TCP
IMAPS		993			TCP
pip (Internet Print Protocol)	IP Phone/ Session Initiation Protocol	5060			TCP
LDAP		389			TCP
LDAPs		636			TCP
Microsoft netbios		139			TCP
Microsoft-ds		445			TCP
ms-nb-ns		137			UDP
Microsoft SQL server		1433			TCP
MySQL		3306			TCP
Nagios NRPE		5666			TCP
nfs		2049			TCP/ UDP
Oracle TNS		1521			TCP
POP3		110			TCP
popssl	Post Office Protocol 3 over SSL	995	tls_ssl_certs (SSL)		TCP
postgresql		5432		cmdb_ci_database	TCP
printer	Printer	515	Printers		TCP

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery ports and protocols (continued)

Name	Service name	Port	Details	Creates	Protocol
sip	SIP (Session Initiation Protocol)	5060			TCP
slp	Service Location Protocol (SLP)	427			TCP/ UDP
smtp	Simple Mail Transfer Protocol	25			TCP
smtp-submission	SMTP Submission (StartTLS)	587	StartTLS (tls_ssl_certs)		TCP
smux (SNMP multiplexing)		199			UDP
snmp	Simple Network Management Protocol	161	Network Devices		UDP
snmptrap		162			UDP
ssh	Secure Shell Service	22	Unix Systems		TCP
sunrpc		111			TCP
telnet		23			TCP
TIBCO Rendezvous		7500			TCP
Tomcat HTTP		8080			TCP
vmapp6_https		9443			TCP
vmapp_https	vCenter Server Appliance Web Interface using https	5480			TCP
wbem_https	CIM-XML via HTTPS(WBEM)	5989	CIM Classification		TCP
winrm	Windows Remote Management	5985	Windows Systems		TCP
winrm_ssl	Windows Remote Management over SSL	5986	Windows Systems		TCP
wins	Windows Internet Name Service	137	NetBIOS Name Resolver		UDP

Windows and dynamic ports

Supported Windows machines can have dynamic ports ranges: 49152-65535 for both TCP and UDP.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure Shazzam probe

When you run Discovery, the Shazzam probe finds your active network devices by scanning specified ports on specified IP address ranges. If the list of IP ranges being scanned is large, you can configure the Shazzam payload for JSON encoding to reduce its size.

Before you begin

Role required: admin

About this task

You control the behavior of individual Shazzam probes using basic and advanced parameters. For instructions on configuring probe parameters, see Set probe parameters.

Note: The Shazzam payload is controlled by the following MID Server properties. These properties do not need to be added to the MID Server. The Shazzam probe will fall back on their default values.

- The **shazzam.chunk_size** property controls the maximum number of IP addresses Shazzam can scan in parallel.
- The **mid.shazzam.regulator.interval_ms** property sets the interval in which Shazzam can launch packets.
- The **mid.shazzam.regulator.packets_per_interval** property sets the number of packets that Shazzam can launch in that time interval.

Procedure

1. Navigate to **All** > **Discovery Definition** > **Probes**.
2. Select **Shazzam**.
3. Add or edit parameters in the **Probe Parameters** related list.
4. Configure the Shazzam parameters.
   These parameters are available for fine-tuning the Shazzam probe. These values are defined in the probe record only.

Shazzam advanced parameters

Parameter	Description
BannerTCP_waitForConnectMS	Sets the number of milliseconds the BannerTCP scanner waits for a connection and banner. **Default**: 1500
debug	Enables debug logging if set to **true**. **Default**: false
delay_wbem	Delays classification of systems with a WBEM port open until the final Shazzam sensor job. For a large schedule discovering many WBEM ports, delaying classification until the last sensor job could cause the node to run out of memory. Setting this parameter to **false** allows classification of these systems to occur across all Shazzam sensor jobs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
	Default: true
DNS_alternativePort	Deprecated
DNS_waitForResponseMS	Sets the number of milliseconds the DNS scanner waits for a response. **Default**: 1000
GenericTCP_waitForConnectMS	Sets the number of milliseconds the GenericTCP scanner waits for a connection. **Default**: 1000
HTTP_waitForConnectMS	Sets the number of milliseconds the HTTP scanner waits for a connection. **Default**: 500
HTTP_waitForResponseMS	Sets the number of milliseconds the HTTP scanner waits for a response. **Default**: 500
NBT_alternativePort	Deprecated
NBT_waitForResponseMS	Sets the number of milliseconds the NBT scanner waits for a response. **Default**: 500
shazzam_report_inactive	When **true**, reports devices that are alive but inactive. For example, a device has no ports open but refuses at least one port connection request. **Default**: true
scanner_log	Enables scanner logging if set to **true**. This logging information appears in the Shazzam probe response. **Default**: false
shazzam_report_dead	When **true**, reports devices with dead IP addresses. For example, a device that has all ports closed. **Default**: false
SNMP_alternativePort	Deprecated
SNMP_tapIntervalMS	Sets the number of milliseconds the SNMP scanner waits between taps. **Default**: 1000

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
SNMP_taps	Sets the number of taps (requests) the SNMP scanner attempts. **Default**: 2
SNMP_waitForResponseMS	Sets the number of milliseconds the SNMP scanner waits for a response after the last tap. **Default**: 1000
TLS_keepOriginalCertificate	When **true** and Discovery is running, the certificate_file field in the cmdb_ci_certificate is populated with the original certificate. This can increase your payload size. **Default**: false

Control Shazzam payload size

A system property converts Shazzam payloads into JSON strings, which dramatically reduces their size. This setting prevents nodes from running out of memory when a single schedule discovers large numbers of IP ranges.

Before you begin

Role required: admin

About this task

The `glide.discovery.shazzam_ranges_json` property is set to `true` for new instances. This setting encodes the payload as a JSON string. The property is configurable by administrators and is available in the **Discovery Definition** > **Properties** module. The property label is **Use JSON for IP ranges in Shazzam** in the module.

This property is set to `false` in upgraded instances and is not visible by default. Adding the property manually to your upgraded instance enables the feature, but does not add it to the **Discovery Definition** > **Properties** module. To enable JSON encoding and add it to the module, import the update set attached to KB0687626.

Procedure

1. In the navigation filter, type `sys_properties.list` and press `Enter`.
2. In the list of system properties, click **New**.
3. Complete the form, using these field values:
   - **Name**: glide.discovery.shazzam_ranges_json
   - **Description**: Encodes the Shazzam payload in JSON.
   - **Value**: true
4. Click **Submit**.

Discovery sensors

Every probe in Discovery must have a corresponding sensor to process the data returned.

For example, if incoming data is the result of a WMI probe, then the WMI sensor is triggered to process the payload.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you create a multiprobe, you must create a multisensor to process the data returned from this probe. For details, see Multiprobes and Multisensors.

Navigate to **Discovery** > **Discovery Definition** > **Sensors** and edit or create a sensor.

Sensor Fields Table

Field	Description
Name	Enter a unique probe name.
Reacts to probe	Select the probe whose payload this sensor must process.
Sensor Class [sys_class_name]	Select the type of sensor to create: • Import Export Map: This option is not currently used. • MultiSensor: Multisensors process the data returned from multiprobes. Select this type to create a multisensor that responds to the simple probes used in a multiprobe. • Sensor: Select this type if this sensor is a simple sensor that responds to a simple probe.
Description	Enter an optional description of the sensors function.
Sensor type [sensor_type]	Specify how the results from the probe are processed. • Classifier: This field is not currently used. • Java: This field is not currently used. • Javascript: Returned data from the probe is processed in the sensor itself, outside the application, and is visible to the user. This is the most common sensor type. • Multiline Text: This field is not currently used. • XML: The XML data from the probe is broken into pieces. Some pieces can be used to launch other probes that the original sensor needs to complete all the necessary information about a device.
Active	Select this option to enable or disable the sensor.
Script	Enter a script to run when processing the probe. You can use the g_probe_parameters hashmap in a sensor script to set probe parameters for any configured, triggered probes. For example, this code sets a '*node_port*' parameter to 16001 for all triggered probes. `g_probe_parameters['node_port']=16001;`
Responds to probes	Use this related list to specify the simple probe within a multiprobe whose payload this multisensor must process. This list is available for multisensors only.
Triggers probes	Use this related list to specify which probes this sensor can start for additional exploration.
These fields can be added by configuring the form	
Field	Description
Agent	Enter the ECC queue type to match this sensor with.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sensor Fields Table (continued)

Field	Description
Apply Defaults	Select this option to apply default sensor values.
Classification type	Select the classification type associated with this sensor.
Coalesce	Enter the field to match against when determining whether to coalesce.
Condition script	Create a predetermined requirement for this sensor to run.
Data option	Select how to treat imported sensor data. Options include: • Insert and Update • Insert Only • Update Only
External names	Enter a name to tag the sensor data.

Cancel sensor transaction by duration

By default, this rule cancels the sensor transaction if it takes more than 20 minutes to complete.

Before you begin

Role required: admin

About this task

If the Discovery sensor transaction takes more than 20 minutes to complete, the following actions occur:

- The associated ECC Queue input record is set to the error state.
- The *Discovery Status Completed Count* is updated.
- An error message is logged, for example:

```
Sensor error when processing Linux - Storage: Transaction
 cancelled: maximum execution time exceeded
```

- The remainder of the discovery continues to run.

  This behavior is provided by the Discovery Sensors Transaction Quota Rule. This rule is active by default.

  You can disable the rule and you can modify the amount of time that has elapsed before the rule takes effect. Accessing this record requires the admin role.

Procedure

1. To disable the rule, navigate to **System Definition** > **Quota Rules**.
   The **Transaction Quotas** plugin must be active.
2. Search on the name `*Discovery` to filter the list of rules.
3. In the name column, click **Discovery Sensors**.
4. To modify the amount of time, edit the **Maximum Duration** field of the rule.
   The value is in seconds.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery probe management

Several discovery probes and their associated sensors are included with Discovery. You rarely need to modify probes or sensors. But you might need to set parameters to control the behavior of a particular probe or align versions of customized probes.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

What you can do

These are the things you can do with probes and sensors:

Review the base system probes

Review the List of Discovery probes to see the probes that exist in the base system. Probes need to be active on classifiers for Discovery to trigger the probes. However, not all probes are active on classifiers as more patterns replace probes and sensors. To see a list of patterns, see Available discovery patterns.

Create or modify a probe

You can create a new probe to discover additional CIs that Discovery does not find with the base system probes or patterns, or modify an existing probe to collect additional information on the type of CI. After you create or modify a probe, test it. You can also create multiprobes and multisensors. See Create a Discovery multiprobe for more information.

Upgrade to the latest version of a probe or sensor

If you upgrade your instance, the Discovery application is also updated, along with components like probes and sensors. However, if you customized any probes or sensors, they do not upgrade. You need to copy your customizations to a text file, upgrade the probes and sensors, and reapply your customizations. See Align versions of customized probes and sensors for more information.

Set probe parameters

Probe parameters control several aspects about how probes function. With each probe provided in the base system, certain parameters are allowed. These are specified in the list of Discovery probes. See Set probe parameters for instructions on how to set a parameter.

Review probe permissions

Certain probes need permissions to run on the target machines or CIs that you are trying to discover. See Discovery probe permissions for more information.

Create or modify a probe

Create a new probe to discover additional CIs or modify an existing probe to collect additional information. After you create or modify a probe, test it.

Before you begin

Role required: admin

Important: You need an advanced knowledge of scripting to modify probes or their associated sensors. Many existing probes provide parameters that you can set, rather than modifying the probe itself. See Set probe parameters for more information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Definition** > **Probes**.
2. Click the name of the probe that you want to modify.
3. Modify the form fields (see table).

Field	Description
Class	Select the probe for the operating system of the device it will query. ○ AWS Probe: Select this class to query your resources inside an AWS cloud. ○ **CIM Probe**: Select this probe type to query a CIM server using WBEM protocols. ○ **Multiprobe**: Select this probe type to run one or more simple probes of any type or mix probes of different types. This type of probe can make several queries simultaneously that return all the results at the same time. You cannot add multiprobes to other multiprobes. ○ **Probe**: Select this generic type to define a probe class. Specify the name of the probe class in the *ECC queue topic* field. ○ **SNMP Probe**: Select this probe type for network devices, such as routers. ○ **WMI Probe**: Select this probe type for Windows devices.
ECC queue topic	Enter a descriptive term for the function of the probe. The probe uses this label as the **Topic** field for incoming ECC queue messages. The term does not have to be unique. For example, all the UNIX probes might have an ECC queue topic value of **SSHCommand**. By default, probes use the following ECC queue topics: ○ CimProbe ○ Multiprobe ○ Powershell probe ○ SCPRelay ○ SSHCommand ○ SNMP ○ WMIRunner
ECC queue name	Enter either a descriptive name for human use, or the actual command the probe is to run. For example, if the value in the ECC queue topic field is **SSHCommand**, then enter the actual shell command to run in this field.
Cache results	Select this option to cache this probe's results to improve overall discovery performance. The probe results cache should be enabled only for those probes and sensors whose output is unlikely to change. For example, the Linux – CPU sensor is cached because CPU information seldom changes. The cache is turned on by default for base system probes and sensors whose output is unlikely to change. ⚠ **Warning:** Do not turn on the cache for classification and identification probes. Furthermore, do not turn on the cache for probes that trigger additional probes because this may prevent them from being triggered.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Classic Mode	Select this option to cache this probe's results in the Starting in Fuji, select this option to debug the post-processing script on the ServiceNow instance instead of the MID Server. This mode is only valid if a post-processor script exists. Use the Classic mode to debug post-processing scripts in the test environment. ServiceNow does not recommend use of the Classic mode in the production environment.
Post processor script	Define an optional post-processing script that runs on the MID Server. The script accepts probe results as an input and outputs a JSON string that is sent back to the instance for a sensor to use as input. Use this type of script to accomplish tasks like parsing data. All probe post-processing scripts and their corresponding sensors have access to the following global variables: ◦ `current`: Sets the attribute to current. Same functionality that exists in sensors. They are attributes of the device that is discovered. ◦ `related_data`: Any additional data to be passed from the probes post-processing scripts, such as `related_list`. All probe post-processing scripts and their corresponding sensors have access to the following methods: ◦ `log(string)`: Logs the string to the MID Server and the ServiceNow instance log files. ◦ `getParameter(param)`: Gets the parameter from the ECC queue payload.

4. Right-click the form header and select **Save**.
5. Modify these related lists if necessary.

Related list	Description
Probe Parameters	Parameters that control the functionality of the probe. See Set probe parameters for instructions.
Sensors that react to this probe or Sensors	The sensors that this probe utilizes.
Included by MultiProbe	If this probe is included with other probes as a MultiProbe.
MultiSensor Scripts	Scripts that run for multiple sensors.
Versions	The version of the probe. If the **State** is **Current**, the probe is the most up to date with your version of the instance. If not, you do not have the most current version and you might need to realign it. See Align versions of customized probes and sensors for instructions.

6. Click **Test probe**.
7. On the Test Probe window, enter the IP address of the target and select the MID Server. Only **Up** and **Validated** MID Servers are able to be selected.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Click **OK**.
9. Check the ECC Queue for the MID Server to see the results.

What to do next

Add the probe to the **Triggers Probe** related list on the appropriate classifier. See Create a Discovery CI classification for a description of the fields and related lists on the classifier form.

Discovery probe permissions

Several probes require additional permissions to run.

Discovering Active Connections

Discovery uses the Windows - Active Connections probe to access active connection information. The application dependency mapping feature requires this probe to function.

Discovering Active Connections

Probes	Additional Permissions
Windows - Active Connections	Ability to invoke the *\root\CIMV2:Win32_Process* object Read and Write access to the admin$ share

Discovering Application Profiles

Discovery uses these probes to access application profile information.

Discovering Application Profiles

Probes	Additional Permissions
Windows - Get APD Environment Files Windows - Get APD Env File Content Windows - Get APD Version File Content	Ability to invoke the *\root\CIMV2:Win32_Process* object Read and Write access to the admin$ share Write access to the %SystemRoot%\temp directory

Discovering VMware Workstation

Discovery uses the **Windows - Get VMware Workstation** probe to access information about VMware virtual machines installed on Windows.

Discovering VMware Workstation

Probes	Additional Permissions
Windows - Get VMware Workstation	Ability to invoke the *\root\CIMV2:Win32_Process* object Read and Write access to the *admin$ share* Write access to the *%SystemRoot%\temp* directory

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering VMware Workstation (continued)

Probes	Additional Permissions
	Configure the **PATH** variable to include the path to the `vmrun.exe` command. This command is typically found in the VMware Workstation `install` directory.

Discovering MSSQL

Discovery uses the Windows - MSSQL probe to access information about Microsoft SQL Server installed on Windows.

Discovering MSSQL

Probes	Additional Permissions
Windows - MSSQL	• Access to the *Win32_Process* object • Access to the *Win32_Service* object • Access to the *Win32_NetworkAdapterConfiguration* object • Access to the *Microsoft.SqlServer.Management.Smo* object • Access to run the *nbtstat* command • Access to HKLM\SOFTWARE\Microsoft\Microsoft SQL Server\InstalledInstances • Access to HKLM\SOFTWARE\Microsoft\MSSQLServer\$instanceName \SuperSocketNetLib\Tcp\TcpPort where *$instanceName* is an array of possible values.

Related topics

Discovery commands for probes and patterns

Set probe parameters

Use probe parameters to control the behavior of a particular probe every time it is triggered.

Before you begin
Role required: admin

About this task
The following parameters are available for all probe types:

Parameter	Description
`source`	[Required] The initial host to connect to. • Type: string (URL) • Default value: None

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create or select the probe you want to set parameters for.
2. From the *Probe Parameters* related list, click **New**.
3. Fill in the fields, as appropriate (see table).

The Probe parameter form has the following fields.

Probe Parameter Form Fields

Field	Description
Name	Enter the parameter name. Each probe type has its own list of available parameters.
Value	Enter the parameter value or script as required by the parameter.
Value script	[Optional] Enter the script you want the parameter to run if you have not already specified a script in the *Value* field.
Probe	Displays the probe this parameter relates to.

Align versions of customized probes and sensors

If you customized a probe or sensor and upgraded to a new instance version, realign the customized probe and sensor to the most current version.

Before you begin

Role required: admin

About this task

Probes and sensors have a major and a minor version.

Major version

A major version change reflects a change in the expected output of a probe. Examples include the addition or subtraction of a targeted attribute, a format change (such as XML vs. JSON), or a probe parameter script change.

Minor version

A minor version change reflects small fixes that don't impact the result or the processing of the data.

A sensor and its corresponding probe must have the same major version. It is a general guideline that they also have the same minor version. This version matching confirms that the data sent back from the probe is understood and properly processed by the sensor. All members of a multi-probes bundle must have the same major and minor version.

By default, Discovery tracks major version mismatches and displays version mismatch errors in the Active Discovery Errors section in the Discovery Log. You can control whether or not Discovery tracks minor versions mismatches by setting the **Warn on Minor Version Mismatch** (`glide.discovery.warn_minor_version`) Discovery property. Minor version mismatches are tracked in the Discovery log, but aren't displayed on the Discovery Dashboard.

Versions for multi-probes and multi-sensors are checked as follows:

- The versions of the individual probes contained in the multi-probes are compared with the Responds to Probes scripts that process their data.
- The versions of the Responds to Probes scripts and the main multi-sensors script are compared.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If a probe and its corresponding sensor don't have the same major version, a sensor does not process information during a discovery and sends error messages to the log file. Errors also show up on the Discovery Dashboard when you run a discovery job. If the major version is the same, but the minor version is not, a sensor processes information during a discovery.

To use the most up-to-date version of a probe that is available:

Procedure

1. Save the customizations that you made with your probe or sensor into a text file so you can re-apply them later.
2. Navigate to **Discovery** > **Dashboard**.
3. Look for version mismatch errors in the **Active Discovery Errors** section, and note the probe or sensor that caused the error.
4. Navigate to **Discovery Definition** > **Probes** or **Discovery Definition** > **Sensors** and select the probe or sensor that you want to upgrade.
5. Add the **Major version** and **Minor version** fields to the form if you want to display the versions on the form.
6. In a multi-sensor, you can verify the major version with the versions of the probes that use the sensor in the **Responds to Probes** related list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Matching multi-probe and multi-sensor version

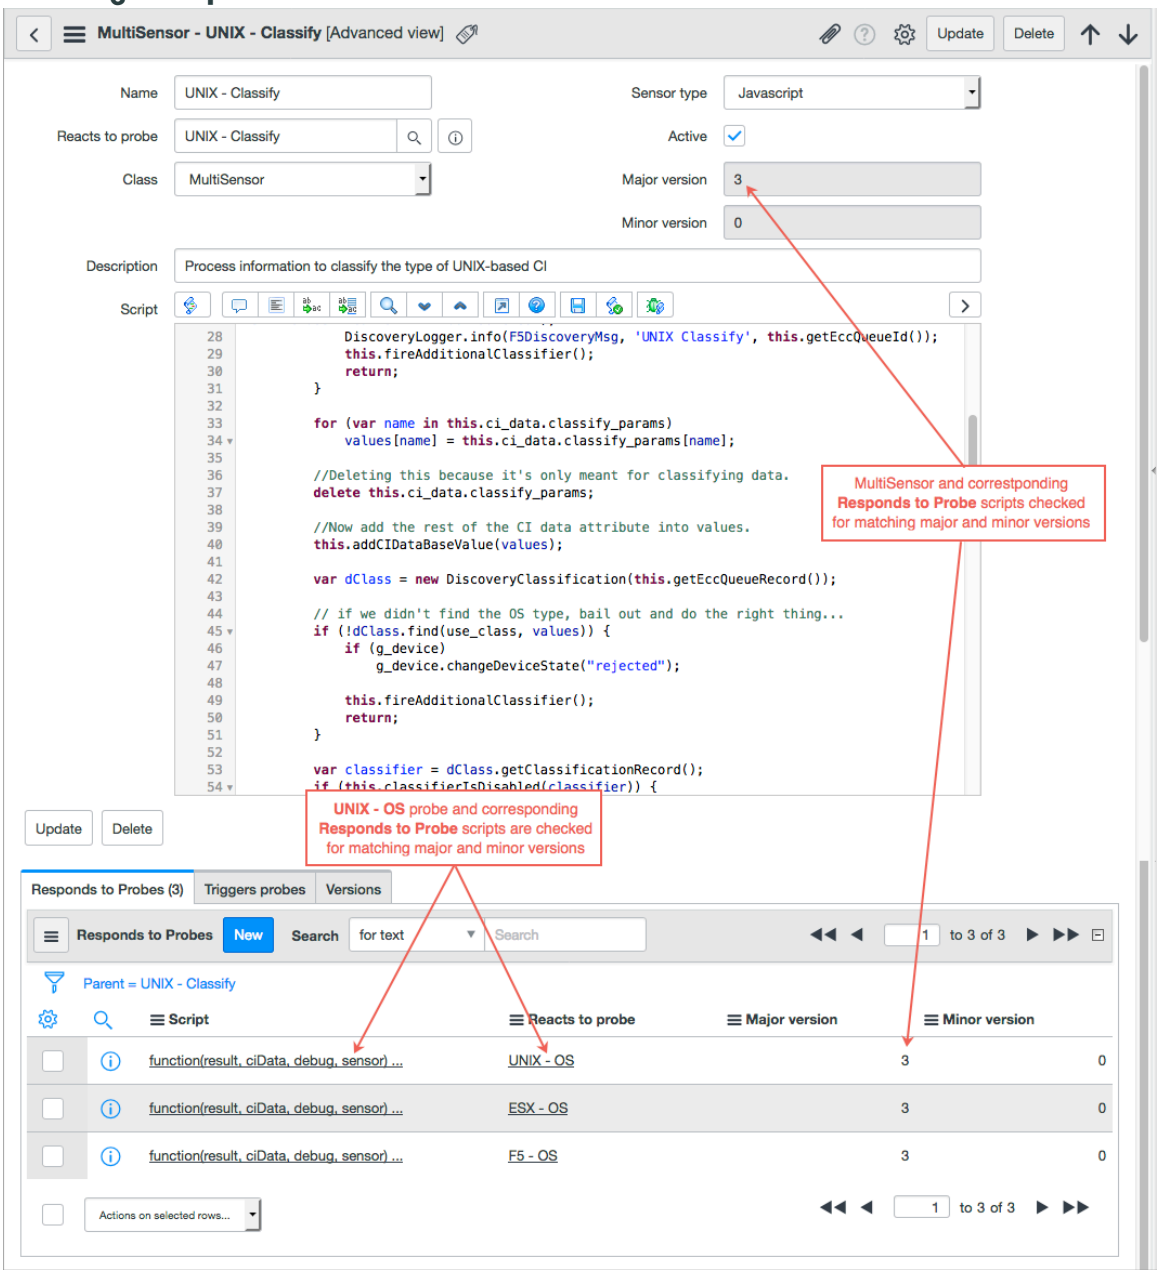

7. Select the **Versions** related list.
8. Look for version where the **State** is **Previous** and the **Source** is **System upgrades**.
9. Select the date link in the **Created** column for that version to open the version record.
10. Select **Revert to this version**.
11. Navigate back to the probe or sensor form and verify that the major version is correct.
12. Reapply your customizations to the probe or the sensor.
13. Verify that your customizations are valid with the new version of the probe or sensor by running a Discovery in your test environment.

Discovery multiprobes and multisensors

Multiprobes contain one or more simple probes configured to extract specific information from manageable devices by executing multiple queries with a single authentication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can schedule multiprobes to run any time in the discovery process to make exploration more efficient. One common use for multiprobes is as identity probes. Identity probes ask a device for information such as its name and serial number. The probes use the results of those queries to update existing CIs in the CMDB.

To process the data returned from a multiprobe, you must create multisensors, which run scripts that process the data returned by the multiprobes. Multisensors typically use the same name as their corresponding multiprobes.

Note: The corresponding multisensors have the same names as the multiprobes.

The following multiprobes are included with the base system.

Multiprobe	Includes These Probes
AIX - ADM	• AIX - Active Processes: Gets active running processes. • Unix - Active Connections: Gets active connections information.
AIX - Identity	• AIX - Network: Determines network interfaces, IPs, and MACs. • AIX - Serial Number: Gets AIX serial numbers.
CIM - Identity	CIM - Computer System: Gets CIM Computer Systems per SMI-S.
CIM - Identity - WBEM Service	CIM - CIMOM: Probes for the WBEM Service, CIMOM.
HP-UX - ADM	• HP-UX - Active Processes: Gets active running processes. • Unix - Active Connections: Gets active connections information.
HP-UX - Identity	• HP-UX - Hardware Serial Number: Gets HP-UX serial numbers. • HP-UX - Network: Determines network interfaces, IPs, and MACs.
Linux - Identity	• Linux - Hardware Information: Gets DMI (BIOS) information. • Linux - Network: Determines network interfaces, IPs, and MACs.
Mac OS X - Identity	• Mac OS X - Network: Determines network interfaces, IPs, and MACs. • Mac OS X - CPU/Memory: Gets CPU and memory information.
SNMP - A10 - Identity	• SNMP - A10 - Identity - Serial: Part of the SNMP - Load Balancer - Identity multiprobe that is used to identify the A10 load balancer. • SNMP - Identity Info: Identifies an SNMP device.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Multiprobe	Includes These Probes
SNMP - ACE - Identity	• SNMP - Identity Info: Identifies an SNMP device. • SNMP - ACE - Identity - Serial: Part of the SNMP - Load Balancer - Identity multiprobe, that is used to identify the ACE load balancer
SNMP - Alteon - Identity	• SNMP - Alteon - Identity - Serial: Part of the SNMP - Load Balancer - Identity multiprobe, that is used to identify the Alteon load balancer. • SNMP - Identity Info: Identifies an SNMP device.
SNMP - AppDirector - Identity	• SNMP - AppDirector - Identity - Serial: Part of the SNMP - Load Balancer - Identity multiprobe, that is used to identify the AppDirector load balancer. • SNMP - Identity Info: Identifies an SNMP device.
SNMP - F5 BIG-IP - Identity	• SNMP - F5 BIG-IP - Identity - Serial: Retrieves the BIG-IP chassis serial number, which is globally unique for this vendor. • SNMP - Identity Info: Identifies an SNMP device.
SNMP - Identity	SNMP - Identity Info: Identifies printer CIs
SNMP - Load Balancer - Identity	• SNMP - F5 BIG-IP - Identity - Serial: Gets BIG-IP chassis serial numbers, which are globally unique for this vendor • SNMP - Identity Info: Identifies SNMP devices
SNMP - Netscaler - Identity	• SNMP - Netscaler - Identity - Serial: Retrieves the Netscaler chassis serial number, which is globally unique for this vendor. • SNMP - Identity Info: Identifies SNMP devices
Solaris - ADM	• Solaris - Active Processes: Gets active running processes • Unix - Active Connections: Gets active connections information
Solaris - Identity	• Solaris - Network: Determines network interfaces, IPs, and MACs • Solaris - Serial Number: Gets serial numbers for Solaris devices †
Unix - ADM	• Unix - Active Processes: Gets active running processes • Unix - Active Connections: Gets active connections information
UNIX - Classify	• ESX - OS: Identifies ESX machines • UNIX - OS: Runs after the ESX - OS probe to determine the operating system for SSH devices not identified as ESX

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Multiprobe	Includes These Probes
Windows - ADM	• Windows - Active Processes: Gets active running processes • Windows - Active Connections: Retrieves active connections information
Windows - Identity	• Windows - Network: Determines network interfaces, IPs, and MACs • Windows - OS/Hardware Information: Probes a Windows machine for WMI information
Windows - Storage 2008	• Windows - Storage 2008 - PS: Retrieves disk and file system information for Windows 2008 and earlier. • Windows - Storage 2008 - WMI: Retrieves disk and file system information for Windows 2008 and earlier.
Windows - Storage 2012	• Windows - Storage 2012 - WMI: Retrieves disk and file system information for Windows 2012 and later. • Windows - Storage 2012 - PS: Retrieves disk and file system information for Windows 2012 and later.

Note: † This probe requires the installation of a command line tool from Oracle called SNEEP. To download and install this tool, log in to the Oracle website. After this tool is installed, the **Solaris - Serial Number** probe runs automatically when Discovery detects a Solaris device. For Fujitsu PRIMEPOWER devices, you must run this probe with root credentials.

Add a simple probe to a multiprobe

You can add simple probes to your multiprobe in the *Includes Probes* related list.

Before you begin
Role required: admin

About this task
You can add simple probes of any type to a multiprobe and mix probe types if necessary.

Warning: Do not add a multiprobe to the *Includes Probes* related list.

Procedure

1. Navigate to **All** > **Discovery Definitions** > **Probes**.
2. Open a multiprobe from the list.
   You can see the MultiProbe designation in the **ECC queue topic** column.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux identity multiprobe

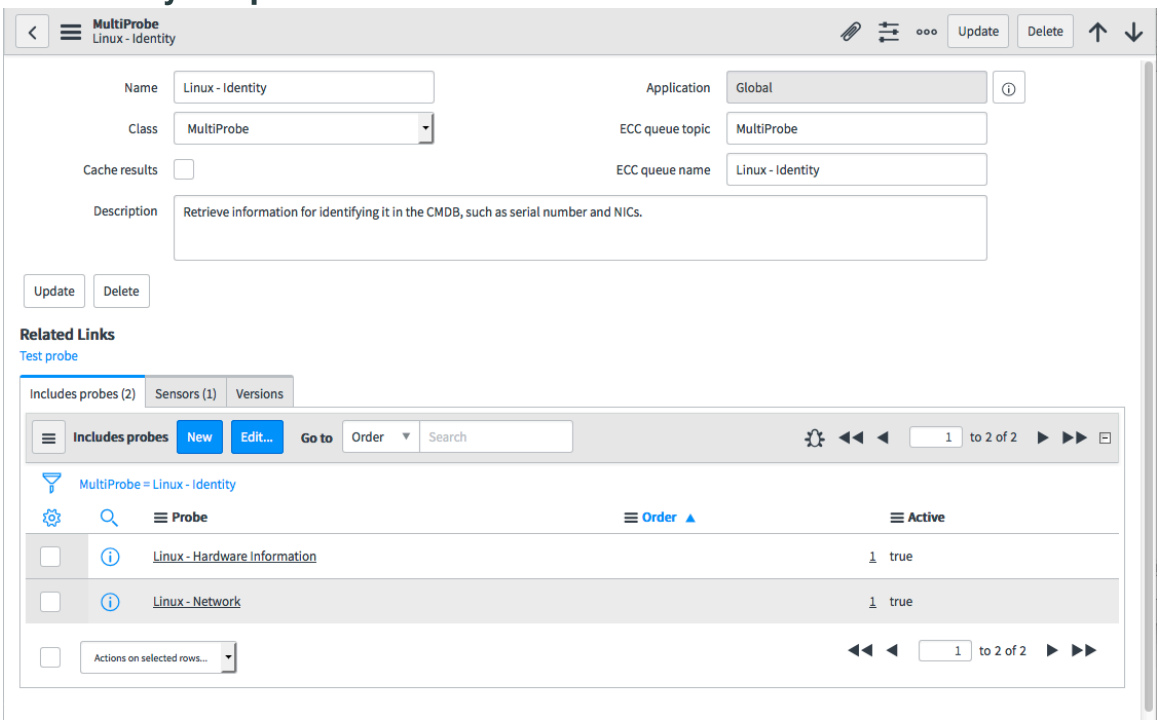

3. In the MultiProbe record, click **New** in the **Includes probes** related list.
4. Select a simple probe from the collection list in the left column and move it into the included list..

Adding a simple probe to a multiprobe

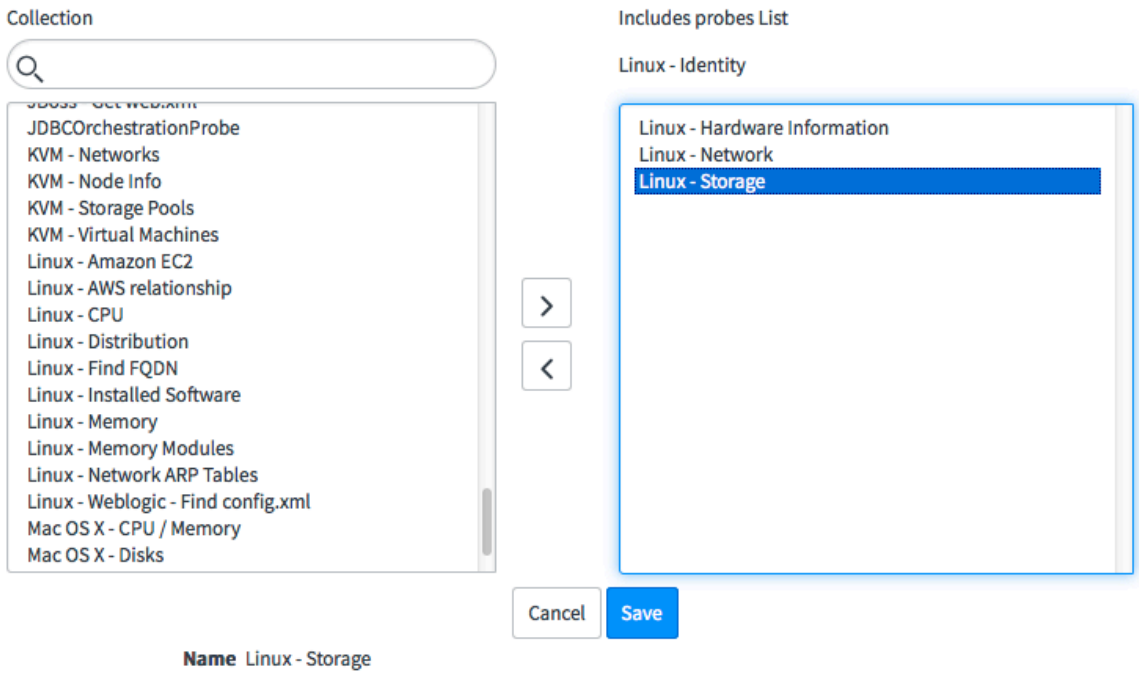

5. Click **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a Discovery multiprobe

You can use the default multiprobes and multisensors provided in the platform or create your own.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Definition** > **Probes**.
2. Click **New**.
3. Complete the form using the following settings:
   - Probe type: Multiprobe.
   - ECC queue topic: MultiProbe.
4. In the *Includes probes* related list, add the probes you want to include in the multiprobe.
5. Click **Save**.

Create a Discovery multisensor

Create a multiensor to process data returned from a multiprobe.

Before you begin

Role required: admin

About this task

Discovery Multisensor and Script

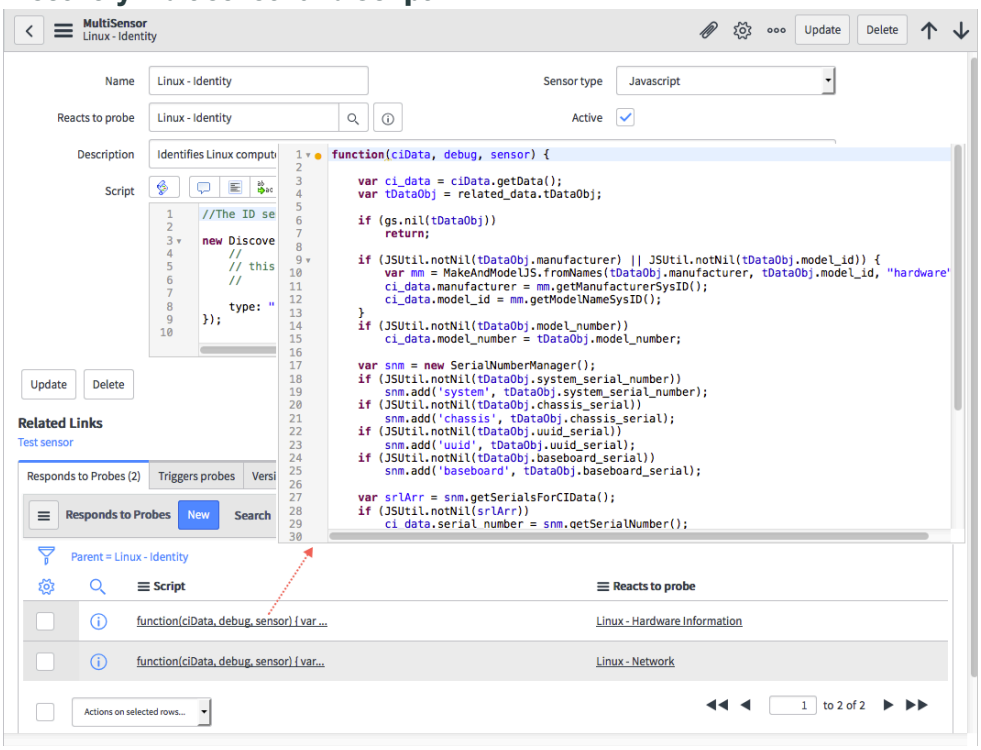

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Definition** > **Sensors**.
2. Click **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Complete the form using the following settings:
   - *Sensor type*: MultiSensor.
   - *ECC queue topic*: MultiSensor.
   - *Reacts to probe*: Select the probe whose payload this sensor must process.
   - *Sensor type*: Select one of these options.
     - *Javascript*: Returned data from the probe is processed in the sensor itself, outside the application, and is visible to the user. This is the most common sensor type.
     - *XML*: The XML data from the probe is broken into pieces. Some pieces can be used to launch other probes that the original sensor needs to complete all the necessary information about a device.
4. Right-click the form header, and click **Save**.
5. In the *Responds to Probes* related list, click **New**.
6. Complete the form using the following settings:
   - *Reacts to probe*: Select the probe within the multiprobe that this sensor reacts to. The sensor must be linked to the probe by function, such as network identification or serial number.
   - *Script*: Enter a script to run before the script in the multisensor. This script defines how the data returned from each probe should be processed.
7. Click **Save**.

Custom Discovery probe to populate a CI from a text file

This custom Discovery probe reads a text file from a Windows computer and populates a configuration item (CI) in the CMDB with the values from the file.

Before you begin
Role required: admin

About this task
In this example the user wanted to read files created by *BGinfo*.

Note: When you have completed the probe and sensor, place the probe in the appropriate Windows classifier at **Discovery Definition** > **CI Classification** > **Windows**.

Procedure
1. Navigate to **Discovery Definition** > **Probes**.
2. Select **New**.
3. Complete the following fields:
   - *Name*: Unique and descriptive name for the probe
   - *Probe type*: Select **Probe**.
   - *Description*: Describe the function of this probe.
   - *Used by Discovery*: Select this check box
   - *ECC queue topic*: This is name of the probe the MID Server is to run. In this example, we use *WMIRunner*.
   - *ECC queue name*: In this example, we use the descriptive name *WMI: BGInfo files*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Right-click in the header bar and select **Save** from the context menu.
5. Select the Probe Parameters tab in the Probe form, and then select **New**.
6. Enter `WMI_GetFiles.js` as the Name of this parameter.
7. Copy the script below into the *Script* field and edit as needed.
8. Select **Submit**.

```
//
 // Use ServiceNow WMIAPI to gather stats
 //
 var CMD_RETRIES = 3;
 var scanner = getScanner();
 if (scanner) {
   var output = "";
   for(var i = 0; i < CMD_RETRIES; i++) {
      output =
 scanner.winExec("%SystemRoot%\\system32\\cmd.exe /C type
 \\\"C:\\Information Systems\\BgInfo\\*.txt\\\"");
      if (output)
         break;
   }
   scanner.appendToRoot("output", output);
 }
```

9. Navigate to **Discovery Definition** > **Sensors**.
10. Select **New**.

    Complete the following fields:

    - *Name*: Use the same name as the matching probe. In this example, we use *Windows - Get BGInfo files*.
    - *Reacts to probe*: The name of the probe created in the previous procedure: *Windows - Get BGInfo files*
    - *Sensor type*: Select the type of sensor to create - in this example *Sensor*.
    - *Description*: Describes the function of this sensor.
    - *Script*: Copy the script below into the *Script* field and edit as needed.
    - *Sensor type*: Determines how the answer from the probe is processed - in this example *JavaScript*.

11. Select **Submit**.

```
new DiscoverySensor({ data: {}, process: function(result) {

        this.parseOutput(result.output);
        this.update(this.data);
    },

 parseOutput: function(output) {
        var currentFile;
        var files = {};

 if (output.startsWith("<wmi")) {
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
            var bgout = new XMLHelper(output).toObject();
            if (!bgout)
                return;

output = bgout.output;
        }

var lines = output.split(/\n/);

for(var i = 0; i < lines.length; i++) {
            var line = lines[i];
            if (line.startsWith("C:\\Information
Systems\\BgInfo\\")) {
                currentFile = line.substr(30);
                files[currentFile] = "";
            } else if (currentFile) {
                var newLine = line.trim();
                if (newLine)
                    files[currentFile] += (files[currentFile]?
"\n " : "") + newLine;
            }
        }

this.data['u_jack_id'] = files['JackID.txt'];
this.data['warranty_expiration'] = files['Warranty.txt'];
this.data['po_number'] = files['Ponum.txt'];},
type: "DiscoverySensor"

});
```

Patterns and horizontal discovery

A pattern is a series of operations that tell Discovery which CIs to find on your network and what credentials to use. Patterns also define what tables to populate in the CMDB.

This topic assumes you understand the phases of horizontal discovery. If you need to review the horizontal discovery process for probes and for patterns, see the following:

- Horizontal discovery process flow with probes and sensors
- Horizontal discovery process flow with patterns

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

How Discovery uses patterns

A pattern performs the same function as a probe: it identifies and explores a target CI. Discovery uses patterns only during the last two phases of discovery: Identification and Exploration.

When you kick off horizontal discovery with patterns, the Scanning and Classification phases run as they would if you were not using patterns (only using probes and sensors). After the Classification stage completes, Discovery looks at the trigger probes on the classifier to see which probe to launch. When the *Horizontal Pattern* probe is specified as a trigger probe, Discovery launches both the *Horizontal Pattern* probe and the pattern that it specifies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: Discovery now enforces code signing for Horizontal Pattern probes, probe parameters, and sensors when code signing is enabled on the instance. Unsigned or tampered payloads are blocked, and signatures are validated before execution. For more information about code signing, see Discovery on Code Signing instances.

Differences between probes and patterns

The main differences between using patterns for horizontal discovery and using only probes for horizontal discovery are outlined in this table:

Topic	Horizontal discovery with probes	Horizontal discovery with patterns
Discovery phases	Probes run during all four discovery phases.	Patterns run only during the identification and exploration phases. Discovery still uses the Shazzam probe and port probes to kick off discovery.
Finding new CIs	Discovery offers several out-of-box probes and sensors. To find new CIs or change CMDB data, you must create probes and sensors or customize the existing ones. To do this, you need knowledge of Java and JavaScript, and you need to understand how probes work together with other probes and with sensors.	Discovery is offering more patterns in every major release and on the ServiceNow App Store. Use the pattern designer to construct steps and enter values. It provides an interface for creating new patterns or customizing existing ones.
Feature support	Supports all standard network, and CI discovery. Certain applications are not supported without the use of patterns. See Detailed information on products discovered by ITOM Visibility for a full list.	In addition to supporting all standard discovery, patterns support: • Cloud discovery • configuration file tracking • CI deletion strategies
The ECC queue and troubleshooting	Discovery creates Multiple ECC Queue records depending on how many probes are launched. To troubleshoot any issues with horizontal discovery, you must analyze several ECC queue records to see the data that Discovery retrieved.	Discovery creates only one additional ECC queue record for the pattern. To troubleshoot any issues with horizontal discovery during identification and exploration, you only have to analyze one ECC queue record for the pattern.
Performance	Depending on the target CI type, Discovery might trigger multiple probes during the identification phase and additional probes during the discovery phase. Horizontal discovery with probes can take	Discovery triggers one pattern during the identification an exploration phases. All operations that Discovery needs to run are in the pattern and Discovery can execute them in order until the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Topic	Horizontal discovery with probes	Horizontal discovery with patterns
	longer because of the overhead in triggering and processing many probes.	pattern is finished. Therefore, pattern discovery can lead to faster performance.

Patterns for top-down and for horizontal discovery

Both Discovery and Service Mapping can use the same pattern for horizontal and top-down discovery. But they are edited differently. See Create or customize patterns for all steps. If you take a pattern that was exclusively used for top-down discovery and you want to use it for horizontal discovery, you have to make a few modifications. See Use a pattern for horizontal discovery for instructions.

Pattern Orchestrator

Pattern Orchestrator is a process that allows a discovery to trigger multiple patterns related to a collection of data in a hierarchical manner. This is to help address issues with large datasets (or large payloads) that can cause slow discoveries and OOM (out-of-memory) issues on a MID Server.

The Pattern Orchestrator is accessed in the **Pattern Orchestrator** tab in the **Pattern Designer**. The tab contains three new tables: **Trigger Rules**, **Orchestrator Input**, **Orchestrator Output**.

The **Trigger Rules** table specifies the parent pattern with the **Parent Pattern** field. The **Batch Size** field specifies how many items should be triggered from the parent pattern to the child pattern. **Batch Size** has the options **Count**, **All**, and **Use Parent Context Only**.

- **Count** must specify a number for the **Batch Size Count**. The child pattern is triggered each time the parent pattern sends the specified number of CIs.
- **All** uses every CI sent from the parent pattern and triggers the child one time.
- **Use Parent Context Only** relies on the parent pattern's output settings to dictate which child pattern should trigger. This is useful when two patterns have the same parent, but you need one to trigger before the other.

The **Orchestrator Input** uses the **CI Type** field to accept CI types from its parent pattern or parent's parent pattern. The **Saved As** field saves the CI type's name to be used in the pattern. Service accounts are a special case and must be named service_account. The **CI Attributes** field uses any outputs from the parent pattern to select which attributes to use.

The **Orchestrator Output** controls which data will be sent to subsequent patterns. It has the fields **CI Class Type**, which is the type of CI being discovered, and **Variable Names** which specifies the variables to be passed on as inputs to child patterns. Any variables not specified in

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Variable Names are not sent to child patterns. The variable **source_native_key** can be manually added by searching for it in the **Variable Names** field.

Source Native Key

In some situations, one pattern may get partial information about a CI, while another pattern gets partial info about the same CI. Instead of creating redundant records or discarding partial information, the information is stored in the CMDB IRE Partial Payloads [cmdb_ire_partial_payloads_list] table. After enough patterns gather partial information, it can correlate that data to provide complete information on the CI.

Pagination

Another way to reduce stress on the MID Server's memory is pagination. Pagination breaks a large payload up into smaller, ordered slices called pages. Pagination limits the maximum size of items received and generates a token. Once that page is processed, the token is passed to the next page, which is then processed, until all the pages are completed.

Pagination is controlled from the Custom Operations table [sa_custom_operation]. The field **Is pagination capable** is hidden by default on the form, and can be added to the form with the **Personalize Form** controls. Once that is set to true, the associated pattern in the pattern designer has the field **Paginated**. Setting the **Paginated** field to true runs that step in the pattern with pagination.

Add the Horizontal Pattern probe to a classifier

To use a pattern for the identification and exploration phases of horizontal discovery, add the Horizontal Pattern probe to the relevant classifiers. The classifiers must match the CI types you are trying to discover.

Before you begin

Role required: discovery_admin

About this task

⚠ **Warning:** This procedure is only intended for custom classifier records or those that do not include a pattern by default. However, if you have already populated your CMDB with data using probes from this classifier and then switch to use patterns, Discovery may create duplicate CIs. Therefore, only follow this procedure if you have never run Discovery with probes for this classifier type.

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **{classifier type}**.
2. Open the classifier record.
3. Click the **Triggers probe** related list.
4. Deactivate the existing identification and exploration probes.
5. Click **Edit**, and add the **Horizontal Pattern** probe.
   The probe appears in the related list.
6. From the related list view, double click the field under the **Pattern** column and add the pattern you want to associate with the classification.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

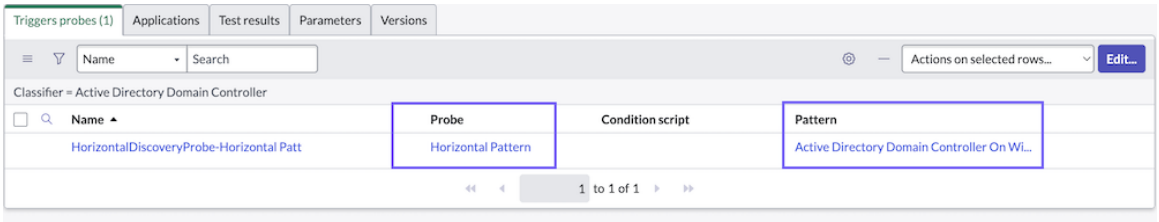

7. Remove or deactivate the other probes from the **Triggers probe** related list.

> **Note:** If you delete a pattern, the Horizontal Pattern probe is not automatically removed from the classifier. You must select another pattern for the Horizontal Probe, or you can switch back to using identification and exploration probes specific to the classifier.
> If you use the Horizontal Probe without a pattern specified, discovery stops after the classification stage.

Use a pattern for horizontal discovery

If you want to use a new pattern, or if you already have a pattern that you were using for top-down discovery, you can use the pattern for horizontal discovery with a few modifications to the relevant classifier.

Before you begin

Role required: discovery_admin

About this task

> **Warning:** This procedure is only intended for custom classifier records or those that do not include a pattern by default. However, if you have already populated your CMDB with data using probes from this classifier and then switch to use patterns, Discovery may create duplicate CIs. Therefore, only follow this procedure if you have never run Discovery with probes for this classifier type.

Procedure

1. Verify that Discovery can use the pattern:

   a. Navigate to **Pattern Designer** > **Discovery Patterns**.

   b. Open the pattern.
   Application patterns that Service Mapping uses are indicated as type **1-Application**.

   c. In the **Identification Sections** on the **Basic** tab, verify that there is at least one section that allows for an entry point type of **TCP** or **All**.
   If not, create one. See the Identification steps for creating a new pattern.

   d. Save the pattern.

2. On the instance, create or modify the classification for the CI type you want to discover.
   Configure the classifier as follows:

   a. Navigate to **Discovery Definition** > **CI Classification** > **{classification type}**.

   b. Open the relevant classifier.

   c. Configure the classifier as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Relationship type**: Select **Runs on::Runs** (for process classifiers only)
- **Condition**: Configure the same condition you defined in the pattern.
- **Triggers probes** Related list: Add the Horizontal Pattern probe, and then add the pattern you are using to the **Pattern** column.

See Create a Discovery CI classification for a description of the other fields on the classifier.

What to do next

Run the pattern in Debug mode to test it. When you are sure the pattern works, you can run discovery by setting up a discovery schedule or running an on-demand discovery. See Schedule a horizontal discovery for more information.

Discovery resource utilization

Standard transactions on Windows and UNIX generate various amounts of network traffic, depending on what is being discovered.

These tables show the bandwidth consumption for each data flow segment of a typical discovery using probes and patterns over different operating systems.

Discovery Bandwidth Over Different Operating Systems (in Megabytes)

Device Type	MID > Instance	Instance > MID	MID > Target	Target > MID	Total
Windows 2016	0.104966	0.101271	0.77739	2.364353	3.34798
Windows 2012	0.126327	0.07928	1.177146	3.70751	5.089804
Windows 2008	0.141816	0.104674	1.032673	3.594784	4.873947
Windows 10	0.091466	0.075601	0.642313	2.221103	3.030483
Linux CentOS	0.164232	0.111376	0.148742	0.690117	1.114467
Mac OSX	0.103707	0.068302	0.021681	0.461365	0.655055
HP-UX	0.120358	0.106676	0.042669	0.101149	0.370852
Solaris	0.130551	0.099414	0.060243	0.346605	0.636813
Cisco UCS Switch	0.029655	0.027465	0.094918	0.097444	0.249492
F5 Load Balancer	0.043935	0.03689	0.017179	0.012132	0.110136
A10 Load Balancer	0.046631	0.032266	0.018313	0.03182	0.12903
EMC Storage	0.4776	0.373828	1.215954	4.741926	6.809308

The following table shows the bandwidth comparison between an initial discovery for three-tier applications and for each subsequent discovery. Bandwidth is broken up into the three tiers: UI (Apache), application (Websphere), and data (Oracle). This measures the total data transfer for each discovery run once for a device class.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Bandwidth for Three Tier Applications (in Megabytes)

Device Type	MID > Instance	Instance > MID	MID > Target	Target > MID	Total
Three-tier application - Initial discovery	0.712829	0.678862	7.084678	9.430181	17.90655
F5 Load Balancer				0.017179	0.012132
Apache on Linux				0.540161	1.107108
Websphere on Linux				0.729403	1.165112
Oracle on Windows				5.797935	7.145829
Three-tier application - subsequent discovery	0.150882	0.107409	2.536535	0.560122	3.354948
F5 load balancer				0.001347	0.012132
Apache on Linux				0.136366	0.79392
Websphere on Linux				0.341042	0.11365
Oracle on Windows				2.05778	0.354948

This table shows discovery of different OS types using patterns. This measures, in megabytes, the total amount of data created and the total amount of data in subsequent scans for each device.

Discovery Bandwidth for Patterns over Different Operating System Types (in Megabytes)

Device		MID > Instance	Instance > MID	MID > Target	Target > MID	Total
Linux	Create	0.39	0.486	0.098	0.273	1.247
	Update	0.382	0.499	0.093	0.264	1.238
Windows Server	Create	0.289	0.316	5.628	8.508	14.741
	Update	0.273	0.306	5.621	8.458	14.658
Solaris	Create	1.222	1.4	0.383	0.917	3.922
	Update	1.24	1.42	.399	.675	3.734
HP-UX	Create	0.176	0.222	0.063	0.13	0.591

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Bandwidth for Patterns over Different Operating System Types (in Megabytes) (continued)

Device		MID > Instance	Instance > MID	MID > Target	Target > MID	Total
	Update	0.178	0.247	0.062	0.128	0.615
Citrix Netscaler	Create	0.424	1.919	0.019	0.042	2.404
	Update	0.355	0.619	0.016	0.041	1.031
F5	Create	0.087	0.135	0.026	0.047	0.295
	Update	0.132	0.171	0.026	0.047	0.376
L3 Switch	Create	0.172	0.125	0.282	0.478	1.057
	Update	0.178	0.126	0.282	0.479	1.065

Note: These measurements were taken with base operating configurations. Your local system results may vary.

CPU Usage examples

Examples from CPU Usage will vary among the matrix of thousands of combinations of Operating Systems, chip sets and specific loads for each system.

Your mix of these variables will determine your unique level of CPU consumption.

You can identify unique builds using internal templates and discover them by watching the performance impact, or lack of performance impact that your Discovery tool has on your system.

Discovery monitoring and issue resolution

Learn how to monitor the progress of your discoveries and how to configure the system to aggregate performance metrics that are important to you. Find descriptions of the error messages you see, and possible steps you can take to solve problems. The Now Support Knowledge Base on Hi contains several articles to help you troubleshoot discovery issues.

Discovery Admin Workspace Diagnostics

Prioritize and address Discovery errors and anomalies from a centralized view in Discovery Admin Workspace. Use the Error Framework to identify root causes and take remediation actions directly from error records. For more information, see Error Framework in Discovery Admin Workspace.

Discovery start and classification issue resolution

Resolve common causes when Discovery doesn't start or a device isn't classified, using Error Framework refined codes and manual checks for causes not yet covered by a code.

Discovery Home Page

Open the Discovery Manager to create or edit schedules for IP-based and Cloud Discovery. View the CIs and cloud resources that your schedules discover and then track any errors that occurred.

Discovery Troubleshooting Guide

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use this article in the Now SupportKnowledge Base if you're experiencing symptoms, such as the inability for Discovery to classify a discovered host, or the Shazzam probe not finding certain types of devices.

Troubleshoot Discovery Performance, Cancellation, and Timeout

Use this article on the Now Support Knowledge Base if your Discovery schedule runs slowly or is canceled before it completes. Find common causes for performance degradation and steps to address them, including guidance on the **Cancel discovery if longer than schedule** setting.

KB0657528: Credentials and Permissions troubleshooting on Discovery, Service Mapping, and Orchestration

Discovery issues may be caused by incorrectly configured or missing credentials. Find out how to identify and configure a missing credential or permission.

Discovery error messages

View the list of common error messages and possible solutions.

Discovery performance metrics

View performance metrics on probe/pattern and sensor processing times. You can use the roll-up data to monitor the performance of specific discoveries or to compare performance between versions after an upgrade.

Daily Slow Patterns

Monitor day-over-day trend and anomaly signals on platform performance patterns such as slow queries, scripts, and transactions.

Discovery Home page

The Discovery Home page provides a summary of discoveries that were triggered by cloud and non-cloud schedules. You can view any errors that occurred during a Discovery and find remediation suggestions. Use the Home page view to examine the cloud resources discovered for the service accounts you selected in a cloud schedule.

To open the Discovery Home page, navigate to **Discovery** > **Home**. The Home page shows the Discovery totals for each of the following categories. Select a tile to view details for that category.

Tip: Discovery Admin Workspace provides an updated experience for monitoring discoveries, managing schedules, and resolving errors. For more information, see Discovery Admin Workspace Home.

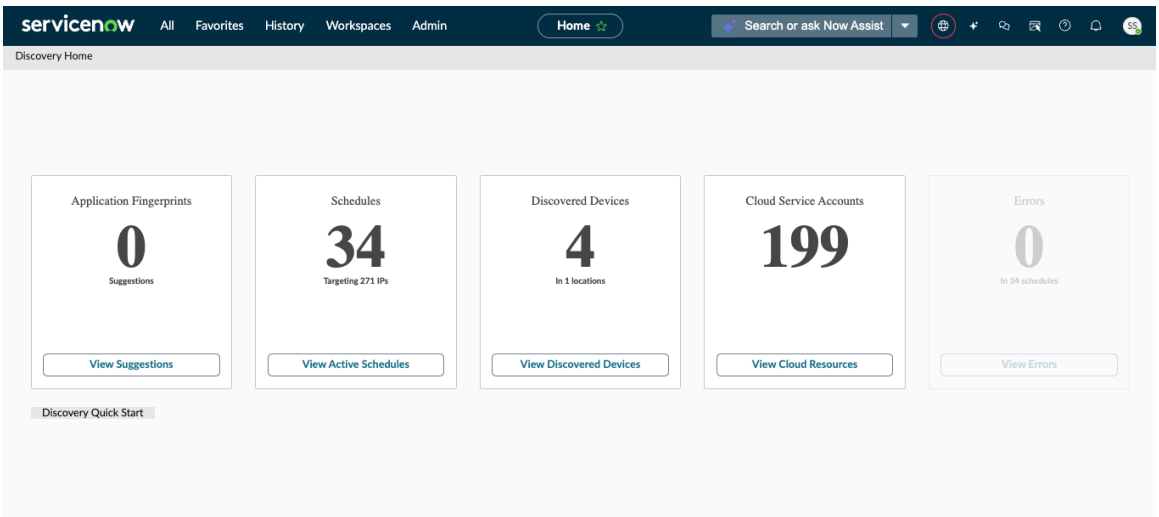

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Item	Description
Schedules	The number of Discovery schedules that are configured to find cloud resources and virtual machines (IP addresses). Select **View Active Schedules** to open the configuration item (CI) schedule view, where you can see specific details about each CI schedule separately. View the results of your cloud or IP address schedules and create schedules.
Discovered Devices	The total number of device CIs discovered. Devices include all the CIs in tables that extend the Hardware [cmdb_ci_hardware] table, such as Computer [cmdb_ci_computer] and Cisco UCS Equipment [cmdb_ci_ucs_equipment].
Cloud Resources	The total number of resources (virtual machines) discovered for the service accounts selected. See Enable discovery of virtual machines for details about setting up Discovery for your cloud resources. Select **View Cloud Resources** to open a list of discovered cloud resources, by type.
Errors	The total number of unresolved errors that occurred during discovery. Select **View Errors** to drill down to specific types of errors.

Viewing all schedules

The schedule view displays details for each schedule that ran for the selected filter. By default, the list shows all of the active Discovery schedules in alphabetical order, regardless of their current running status.

Note: You can also monitor and manage IP and cloud-based schedules in Discovery Admin Workspace Schedules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

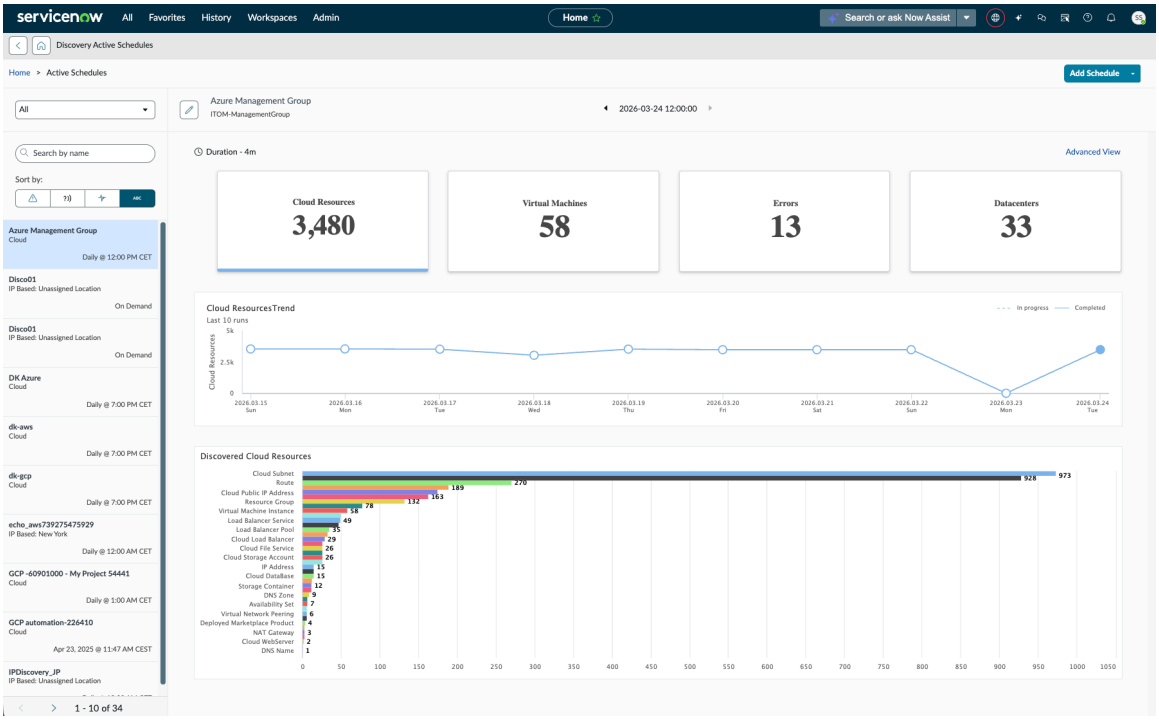

Use these controls to filter Discovery results by schedules:

Schedule filters

You have the choice of these filters:

- **All**: Results of all Discovery schedules that ran. This is the default filter.
- **Cloud Schedules**: Results of all Cloud Discovery schedules that ran. When you select this filter, a second filter appears, allowing you to display cloud schedules by **Service Account**.

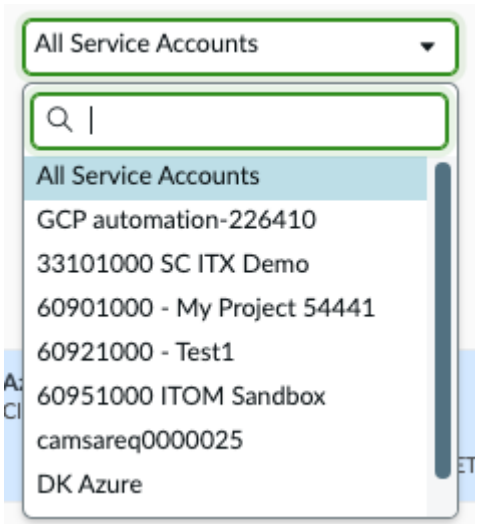

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **IP Based Schedules**: Results of all Discovery schedules that ran for Configuration Items (CI), using IP ranges. When you select this filter, a second filter appears, allowing you to display the CI schedules by a specific location or all locations.

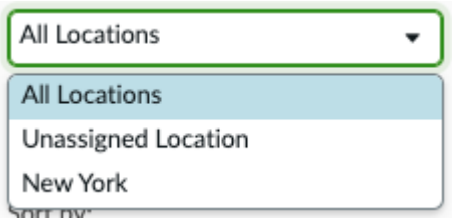

Schedule sorter

From the schedule selector in the left pane, select a sorting criteria, from left to right:

- Errors for each schedule listed.
- Unidentified IP addresses.
- Recent activity, which lists the schedules that are currently running first.
- Alphabetical listing of schedules.

Schedule list

List of available Discovery schedules. Select a schedule to display its details. Select the pencil icon at the top of the page to edit the schedule.

Date/time indicator

Date and time the selected schedule ran.

Schedule topics

Tiles that show the Discovery result categories for the selected schedule. Select a tile to see more detail.

- **Cloud resources**: Displays the total number of CIs updated or created for this filter.
- **Virtual Machines**: Displays the total number of virtual machines discovered with the selected schedule.
- **Errors**: Displays the total number of errors that occurred from running the selected schedule.
- **Datacenters**: Displays the number of datacenters selected during the creation of the Discovery schedule.

Cloud resources trend chart

Trend graph for the selected schedule. Each circle in the graph represents a specific Discovery run.

Place the cursor over any circle to see a summary of that discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

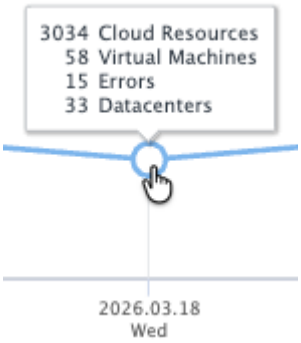

Select the circle to filter the values you see in the Discovery topics. The results populate the appropriate tiles above the graph.

Errors by category

Category tiles for the errors encountered during the Discovery for the selected schedule. Select a tile to view the errors in that category and to see recommended actions for remediation. See Discovery error messages for additional details.

Discovered cloud resources

Breakdown of the cloud resources discovered for the selected schedule, presented in a horizontal bar chart.

Add schedules

Types of schedules to add.

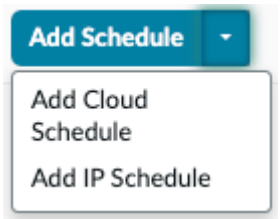

Advanced View

Displays the Discovery status record for the selected schedule. Use this link to examine the results of the schedule in the ECC Queue and the Discovery log.

Viewing IP-based schedules

To display results from IP range discoveries, select **IP Based Schedules** in the primary filter. Use the secondary field that appears to filter the search results by **Location**. The tiles show the result categories for the selected IP-based schedule. Select a tile to see more detail.

- **Discovered devices**: Total count of virtual machines discovered by the IP-based schedule, including both targeted and unidentified IP addresses.
- **Errors**: Displays the total number of errors encountered running the selected IP-based schedule.
- **Unidentified IPs**: Count of unidentified CIs discovered by the IP-based scheduled.
- **Targeted IPs**: Count of CIs found that were configured in the IP-based schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

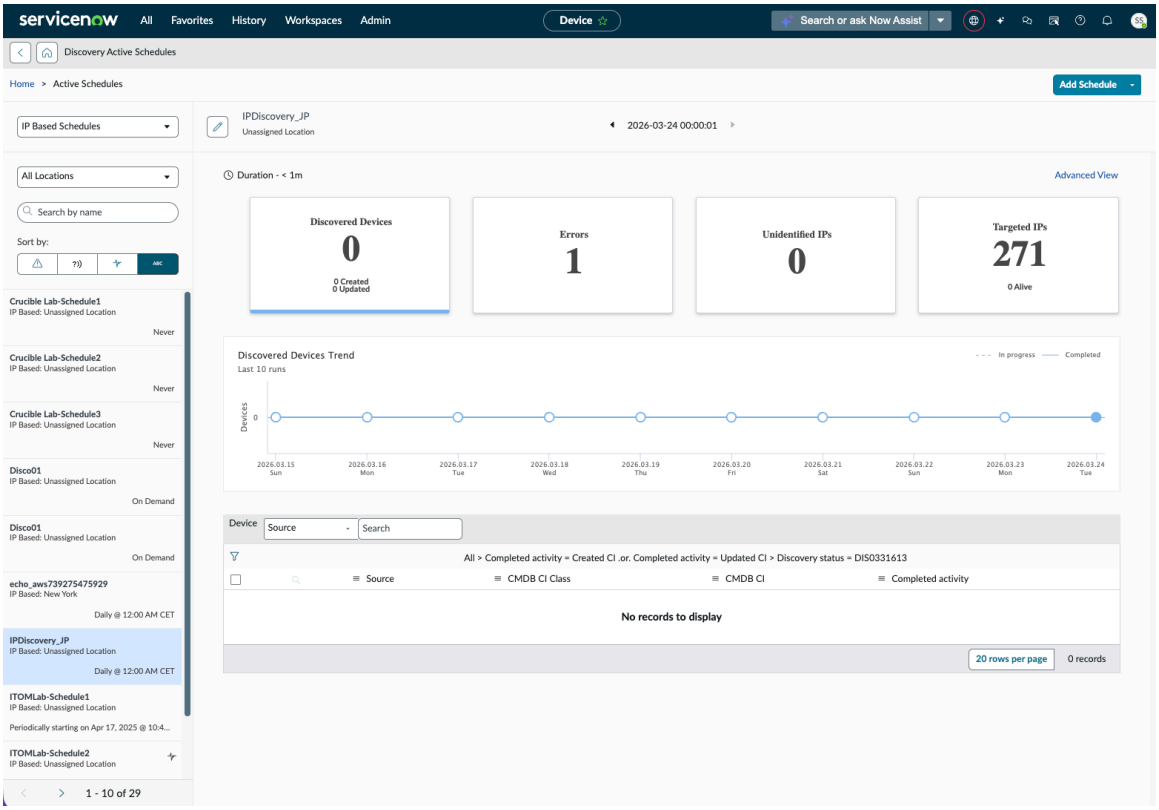

Resolve CI Discovery schedule errors

View and resolve Discovery errors from the ServiceNow Home page. Get suggestions for fixing errors and filter by all schedules or a single schedule.

Before you begin

Role required: discovery_admin

Note: You can also investigate and resolve Discovery errors in Discovery Admin Workspace. For more information, see Discovery Admin Workspace Error Details.

Procedure

1. Navigate to **All** > **Discovery** > **Home**.
   The summary page for all schedules appears.

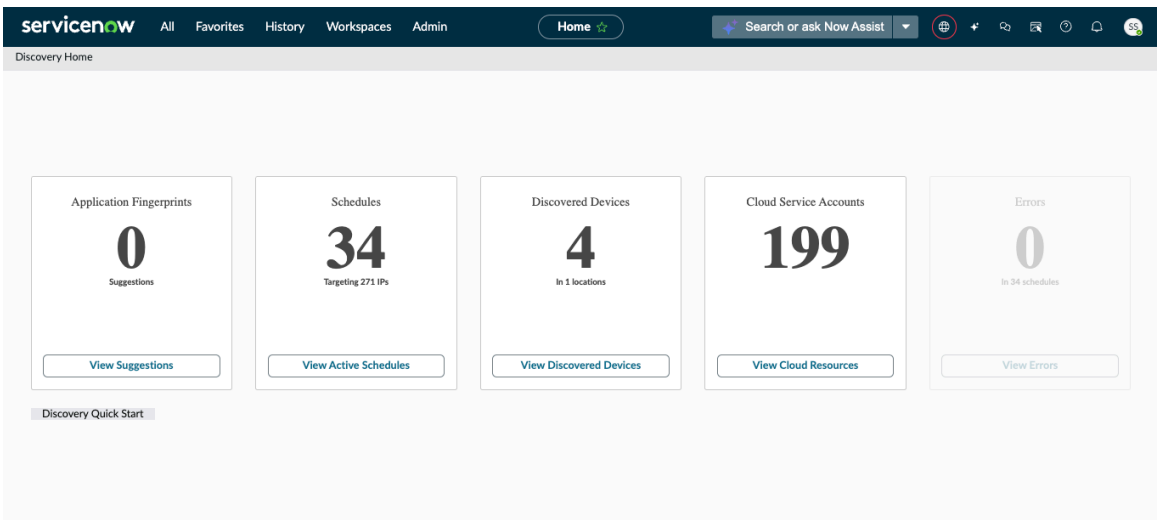

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. View the errors, using one of these methods:
   - Select the **Errors** tile to see the errors by category for all configuration item (CI) discovery schedules.
   - Select the **Schedules** tile to display the errors for a single schedule.
     a. In the summary screen, select the schedule that you want to investigate for errors.

        If you have numerous schedules, you can filter by location, using the location drop-down list, or filter by the schedule name in the **Search by name** field.

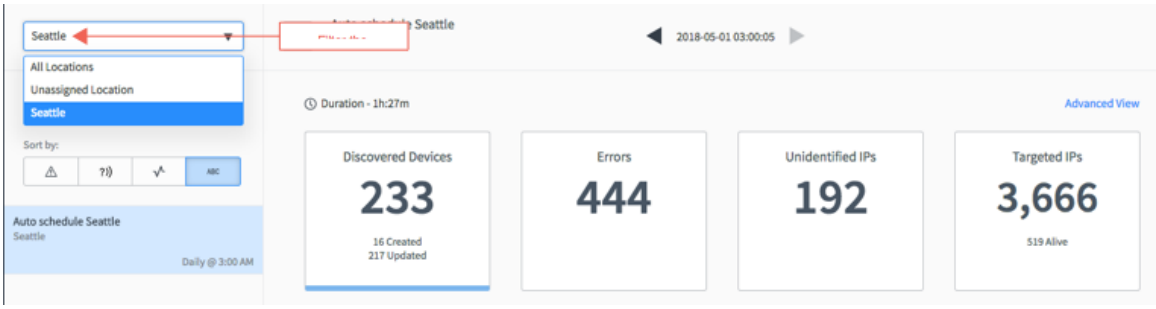

     b. After you select the schedule, select the **Errors** tile.
     c. Scroll down to view the errors by category.

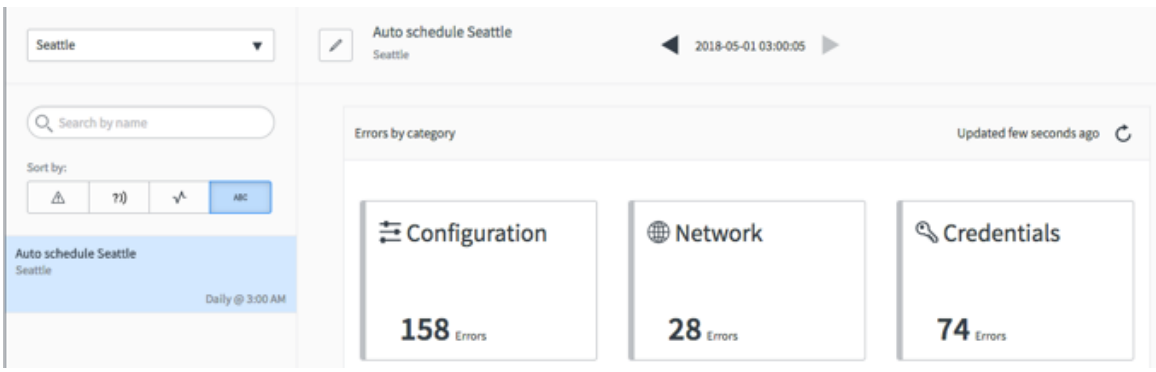

3. Select a category tile to display the specific error codes that occurred in that category.
4. Select the tile for the error code you want to investigate.

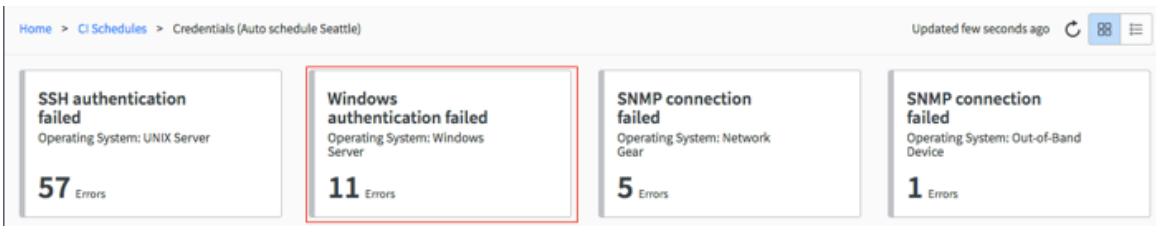

The list displays all occurrences of that error code, either for a single CI schedule or for all schedules, depending on the path you took to view the errors. The list shows the IP address of each instance that experienced the error and the error status **Active error**. The **Recommended Actions** pane on the right contains the recommended action for the selected error. The **ACTION ON ALL** pane contains operations you can execute on all the errors in the list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

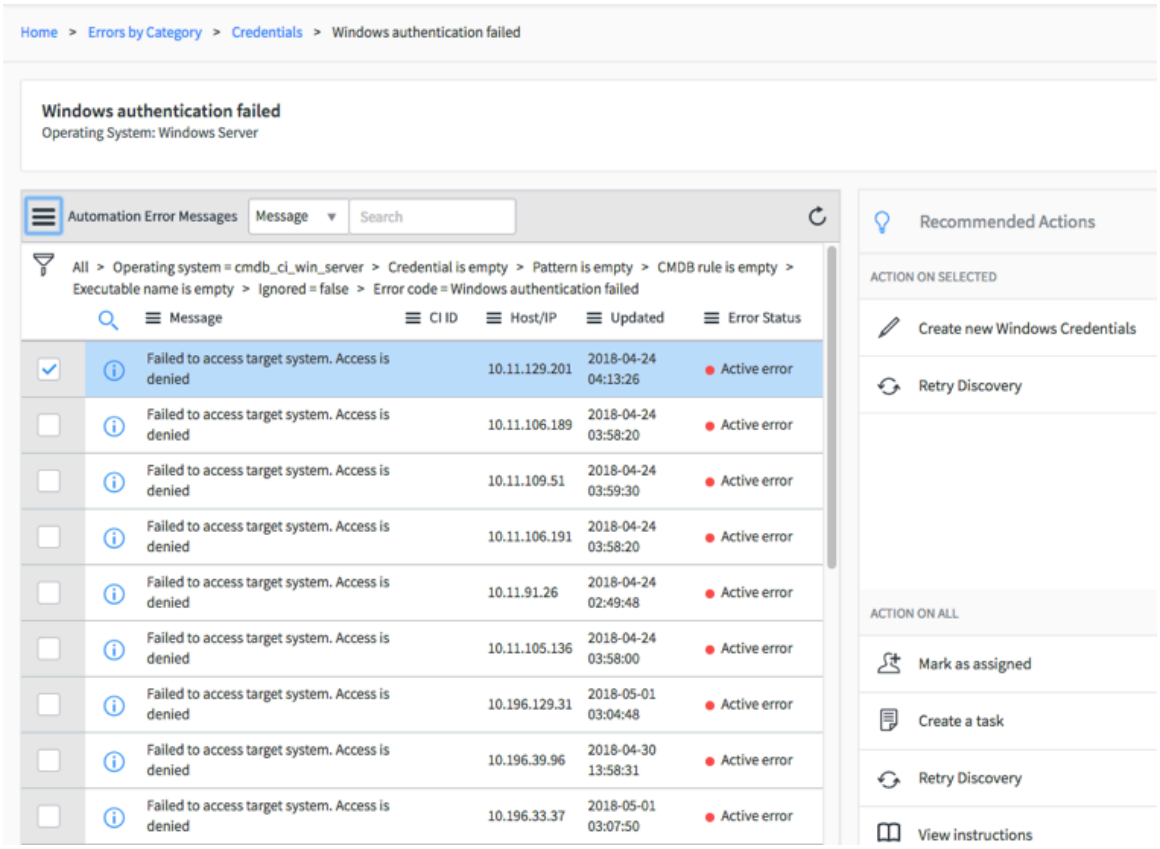

5. Select **View instructions** to see details and remediation suggestions for the error code.

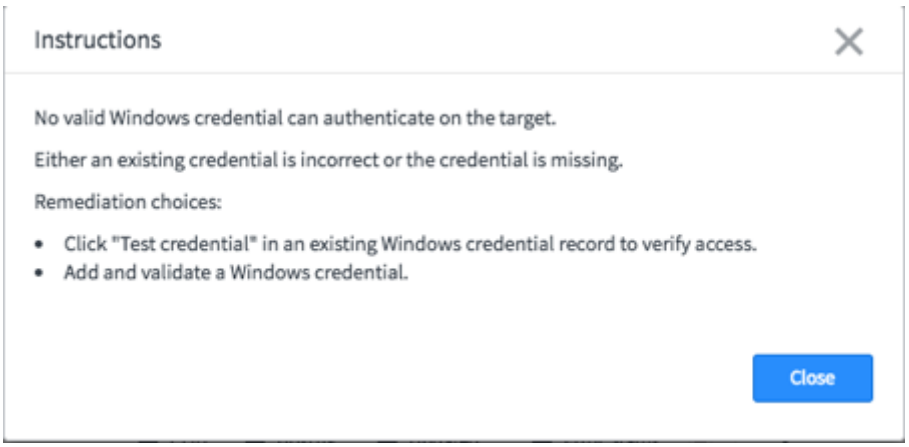

6. To begin work, either select the check box next to an error to resolve that specific error or select an action for **ALL** to resolve that error code for all instances of that error.
   In this example, the system recommends creating new Windows credentials to resolve this error code.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

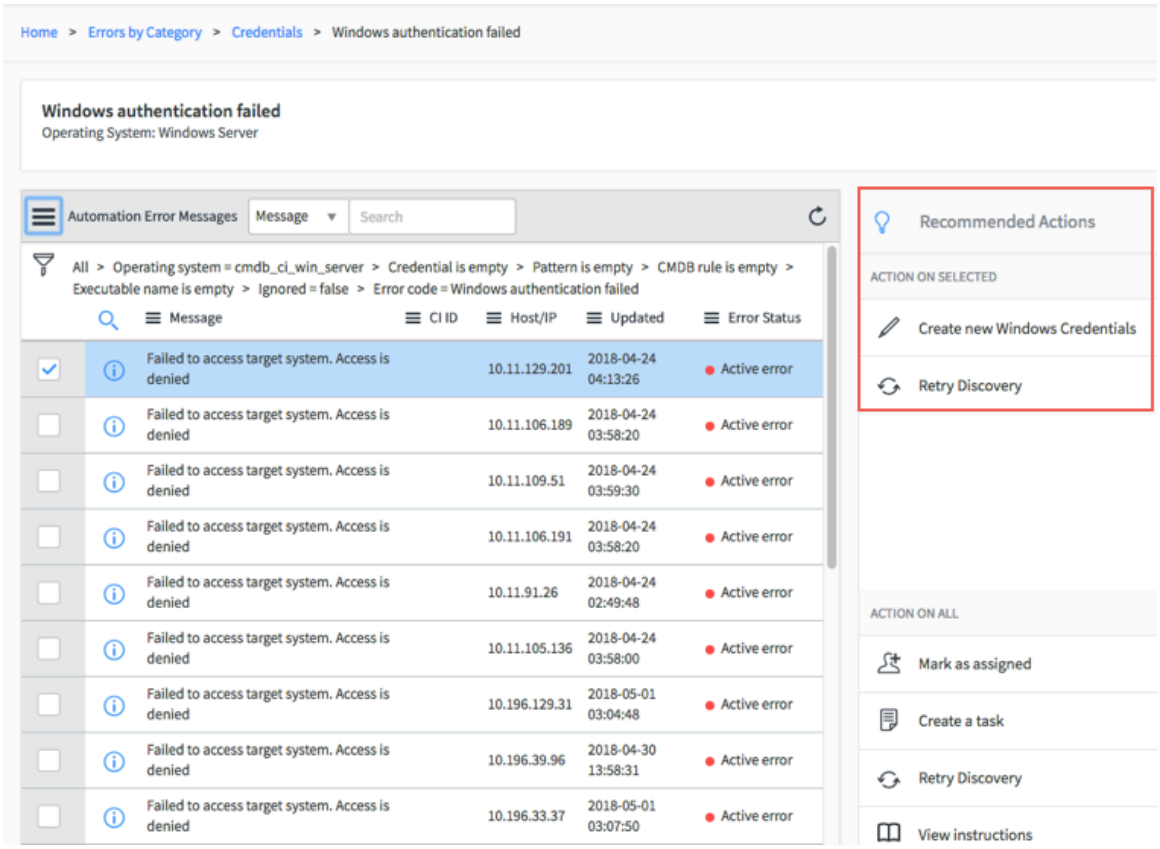

7. Complete the suggested remediation tasks before attempting to retry Discovery.
   In this example, the **Create new Windows Credentials** link opens a form for creating and testing new Windows credentials.

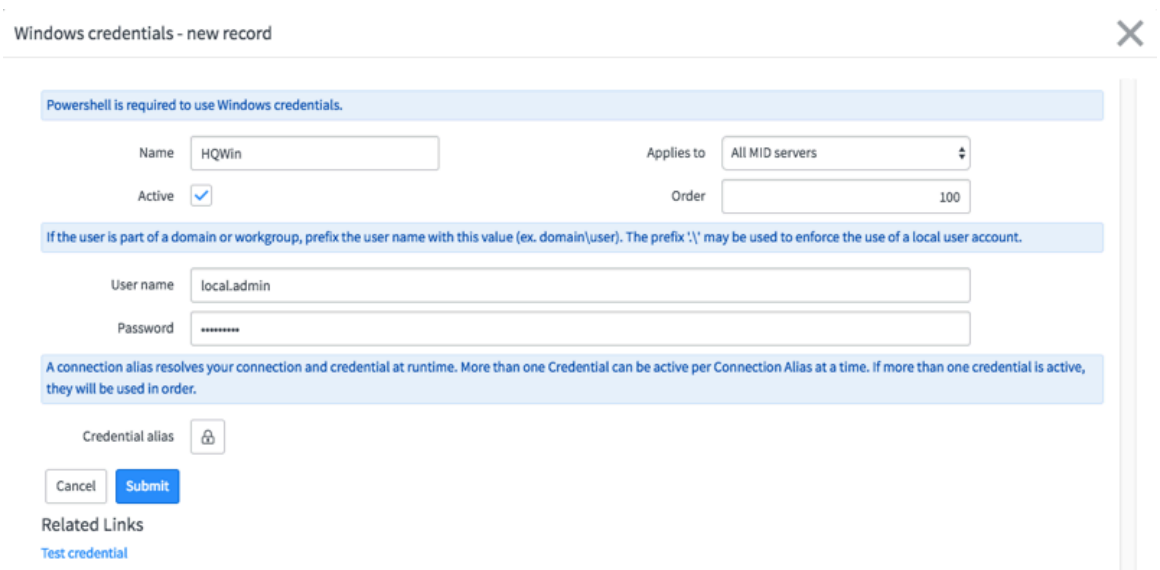

8. After all the recommended tasks are complete, select **Retry Discovery**.
   A possible course of action might be to retry discovery for **ALL** instances after creating the credentials. The new credentials might solve the issue for all the instances experiencing this error code.
9. Check the **Error Status** column to see the results of your remediation efforts.
   Single retries complete quickly. Retries for large numbers of IP addresses can take several minutes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

An error can have one of these statuses:

- **Active error**: Unresolved error. This is the default status of all new errors.
- **Resolved**: Resolved error. The recommended actions show Error successfully resolved.
- **Assigned**: Task assigned to resolve this error. Select **Create a task** to assign a user to a task for all active errors. Assigned errors are considered **Active errors**, and actions can still be performed on them.
- **Pending Discovery**: Waiting for Discovery to start after you execute the **Retry Discovery** action on all errors.
- **In Discovery**: Discovery is currently active.

Related topics

Discovery status

Review the logs for horizontal discovery

The ECC queue for Discovery

Create Discovery schedules for cloud resources

Cloud Discovery home page in Cloud Discovery Workspace

The Cloud Discovery home page displays a summary of the discoveries triggered through a discovery schedule. You can view the count of CIs discovered and errors encountered over time. You can also select any dashboard report to view additional information.

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

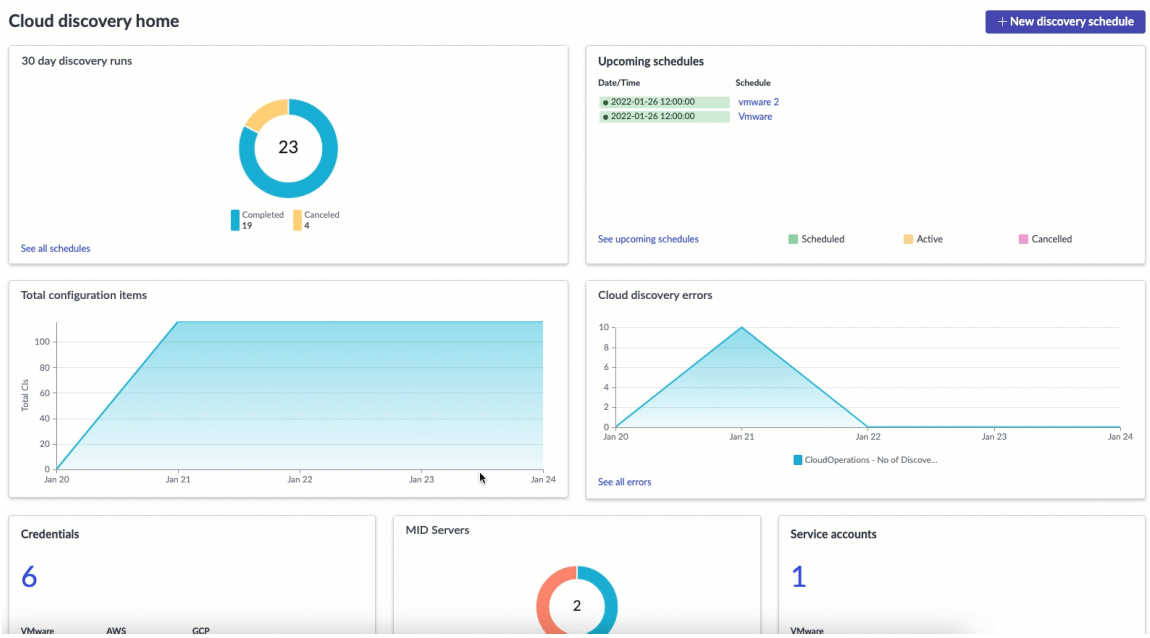

Required ServiceNow AI Platform roles

sn_cloud_ops_ws.cloud_ops_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Access the Cloud Discovery home page

To open the dashboard, navigate to **Cloud Discovery Workspace** > **Cloud Discovery** > **Dashboard**.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
sn_cloud_ops_ws.cloud_ops_admin	The sn_cloud_ops_ws.cloud_ops_admin can review the following information in the Cloud Discovery dashboard: • 30 day discovery runs • Upcoming schedules • Total configuration items • Cloud discovery errors • Credentials • MID Servers • Service accounts • Cloud VMs by datacenter

Indicators

Credentials

This indicator displays the total count of credentials defined in the system and their breakdown according to cloud provider.

Select the number mentioned under any credential type to see a filtered list of matching credentials.

Service accounts

This indicator displays the total count of service accounts defined in the system and their breakdown according to cloud provider.

Select the number mentioned under any cloud provider to see a filtered list of matching service accounts.

Reports

Title	Type	Source table	Description
30 day discovery runs	Donut chart 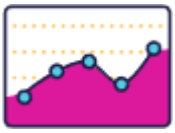	discovery_status	This report displays the count and the state of the discovery runs over the last 30 days. Cloud Discovery displays the following states in the report:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
			• Active • Canceled • Completed Select the report to view the list of all the discoveries.
Upcoming schedules	List	discovery_schedule	This report displays information about up to eight upcoming scheduled discoveries. Select any schedule to view its details.
Total configuration items	Line and area chart 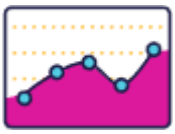	cmdb_ci	This report displays the number of CIs discovered over time. Select the report to view the list of all the configuration items discovered so far. You can select any CI record to view more information about the CI and its multi-source data if available.
Cloud discovery errors	Line and area chart 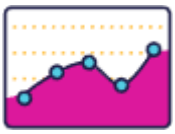	discovery_log	This report displays the number of Cloud Discovery errors over time. Select the report to view the list of all Cloud Discovery errors in the Discovery log. Alternatively, you can access this information, by navigating to **Diagnostics** > **Discovery log**.
Credentials	Indicator	discovery_credential	This report displays the total count of credentials defined in the system and their breakdown per cloud provider. Select the number mentioned under any credential type to see a filtered list of matching credentials.
MID servers	Donut chart	ecc_agent	This report displays the total count of MID Servers and their status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
Service accounts	Indicator	cmdb_ci_cloud_service_account	This report displays the total count of service accounts defined in the system and their breakdown per cloud provider. Select the number mentioned under any cloud provider to see a filtered list of matching service accounts.
Cloud VMs by datacenter	Horizontal bar chart 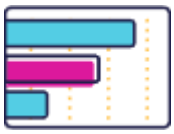	sn_disco_cd_vm_analytics	This report displays the count of VMs discovered per datacenter. Cloud Discovery displays this report for Microsoft Azure, Amazon Web Services (AWS), and vCenter datacenters only.

Cloud Discovery schedules dashboard in Cloud Discovery Workspace

The Cloud Discovery schedules dashboard displays information about the selected discovery run.

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

Required ServiceNow AI Platform roles

sn_cloud_ops_ws.cloud_ops_admin

Access the Cloud Discovery schedules dashboard

To open the dashboard, navigate to **Cloud Discovery Workspace** > **Cloud Discovery** > **Schedules**.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
sn_cloud_ops_ws.cloud_ops_admin	• View the resource count from the recent discovery runs. • View the time taken for the recent discovery runs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

User	Dashboard use
	• View the list of newly discovered CIs • View the list of CIs missing from the previous discovery run. • View the count of CIs discovered during the selected discovery run. • View the count of discovery errors that occurred during the selected discovery run.

Indicators

Total Resources

The indicator displays the total count of Configuration Items (CI) discovered during the selected discovery run.

Newly discovered

The indicator displays the count of new CIs discovered during the selected discovery run.

Absent from last

The indicator displays the count of CIs missing during the selected discovery run. That is, CIs were discovered during the previous discovery run, but are missing during the selected discovery run.

Errors

The indicator displays the count of discovery errors that occurred during the selected discovery run.

Duration

The indicator displays the time taken to complete the selected discovery run.

Note: Cloud Discovery calculates the **Newly discovered** and **Absent from last** indicators only for the pattern-based discoveries.

Reports

Title	Type	Source table	Description
Resource counts from recent discovery runs	Bar chart	discovery_cloud_results	This report displays the count of CIs discovered during the last 10 runs of the selected discovery schedule.
Time taken for recent discovery runs	Bar chart	discovery_status	This report displays the time taken to complete the last 10 runs of the selected discovery schedule. This report covers both completed as

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
			well as canceled discovery runs.

Filters

Name	Type	Description
Schedules filter	Choice-based filter	By default, the list shows all the Cloud Discovery schedules in alphabetical order, regardless of their current running status. Use the schedule filter to filter the Cloud Discovery schedules. You can filter the schedules by cloud provider, schedule name, or both. You can also sort the filtered list by any one of the following criteria: • **Alphabetically**: Sort the schedule names alphabetically. • **Errors**: Sort by the count of errors encountered during the discovery. • **Discovery in progress**: Sort by status of the discovery. After filtering the schedule, select the discovery run from the Report drop-down list to visualize its data in the schedules dashboard displayed on the **Results** tab. **Note:** Only the last 10 discovery runs appear in the Report drop-down list.

Discovery performance metrics

This Discovery enhancement collects performance metrics on probe/pattern and sensor processing times and then aggregates that data over time. You can use the roll-up data to monitor the performance of specific discoveries or to compare performance between versions after an upgrade.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metrics

Discovery provides these individual performance metrics:

- Probe and pattern processing time.
- Sensor processing time.
- Identification and Reconciliation Engine (IRE) processing time for Discovery patterns. This processing time is already included in the sensor processing time, but is isolated here to provided more insight into the identification and reconciliation of pattern payloads.

Discovery can aggregate individual metrics for these attributes:

- Builds/versions
- Discovery status
- Target IP address

How metric aggregations are triggered

Metric roll-ups are initiated as follows:

- **Aggregated by build**: Implemented by the **Aggregate Discovery Probe And Sensor Metrics By Build** scheduled job. This job runs at 0200, local time.
- **Aggregated by status**: Implemented by the **Rollup Probe/Sensor Metrics by Status** Script Action, which is triggered by the **discovery.complete** or **discovery.cancelled** registered events.
- **Aggregated by target**: Implemented by the **Rollup Probe/Sensor Metrics by Target** Script Action that is triggered by the **discovery.device.complete** registered event.

Note: If Discovery execution is cancelled before it completes, Discovery cannot update the IP target metric aggregation table. This is because the *`discovery.device.complete`* event that triggers aggregation does not run. IP target data for an interrupted Discovery is collected when subsequent discoveries run successfully on the target. Cancelling Discovery execution does not affect the aggregation of other metrics, which are triggered differently.

Tables

Discovery performance metrics data is stored in these tables:

Table	Description
Probe and Sensor Metrics (Individual) [discovery_perf_metric_probe_sensor]	Stores the individual performance metrics for probes/patterns, sensors, and IRE processing times.
Probe and Sensor Metrics (Aggregate) [discovery_perf_metric_probe_sensor_rollup]	This is the base table for the metric aggregations and does not store data itself. The three roll-up tables, by-build, by-status, and by-target all extend this table.
Probe and Sensor Metrics (Aggregated by Build) [discovery_perf_metric_probe_sensor_rollup_by_build]	Stores the aggregated performance metrics for probes/patterns, sensors, and IRE by build and version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Probe and Sensor Metrics (Aggregated by Status) [discovery_perf_metric_probe_sensor_rollup_by_status]	Stores the aggregated performance metrics for probes/patterns, sensors, and IRE by discovery status.
Probe and Sensor Metrics (Aggregated by Target) [discovery_perf_metric_probe_sensor_rollup_by_target]	Stores the aggregated performance metrics for probes/patterns, sensors, and IRE by IP targets.

Discovery properties

Performance metrics properties control whether or not aggregation occurs, but not what data is included in the aggregation. Status and IP target data is collected as follows:

- Rollups for status always contain new data. Discovery continuously collects data on all probes and sensors during the discovery execution for that discovery status and stores it in the Probe and Sensor Metrics (Individual) [discovery_perf_metric_probe_sensor] table. Aggregation rolls up all probe and sensor data for that particular status after *`discovery.cancel`* and *`discovery.complete`* events are fired for that status, but only if the aggregation property for status roll-ups is enabled.
- Discovery continuously collects data on IP targets and stores it in the Probe and Sensor Metrics (Individual) [discovery_perf_metric_probe_sensor] table. Aggregation rolls up all existing IP target data after the *`glide.discovery.perf.metrics.rollup_by_target`* property is enabled and creates records in the Probe and Sensor Metrics (Aggregated by Target) [discovery_perf_metric_probe_sensor_rollup_by_target] table.

These properties control the gathering of probe and sensor metrics:

Property	Description
glide.discovery.perf.metrics.enable_collection	Enables collection of performance metrics. • **Type**: true | false • **Default**: true
glide.discovery.perf.metrics.rollup_by_build	Enables aggregation of individual discovery performance metrics into a roll-up table that groups by build. • **Type**: true | false • **Default**: true
glide.discovery.perf.metrics.rollup_by_status	Enables aggregation of individual discovery performance metrics into a roll-up table that groups by discovery status. • **Type**: true | false • **Default**: false
glide.discovery.perf.metrics.rollup_by_target	Enables aggregation of individual discovery performance metrics into a roll-up table that groups by target IP address. By default, Discovery continuously collects individual

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	IP address metrics, even when aggregation by target IP is disabled. When you enable IP target aggregation, Discovery includes all target metrics in the roll-up table. • **Type**: true | false • **Default**: false

View Discovery performance metrics for probes, sensors, and patterns

By default, Discovery tracks the performance of individual probes, sensors, and patterns by measuring the processing time. When patterns are used, Discovery measures the Identification and Reconciliation Engine (IRE) processing time.

Before you begin

Role required: discovery_admin, discovery_operator, or admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Performance Metrics** > **Probe/Sensor (Individual)**.
2. Sort the list by Discovery status to see the list of probes and patterns that ran in a specific Discovery.
3. You can view the metrics for each probe, sensor, or pattern in the list or open the record. All probe and sensor metrics data are read-only.

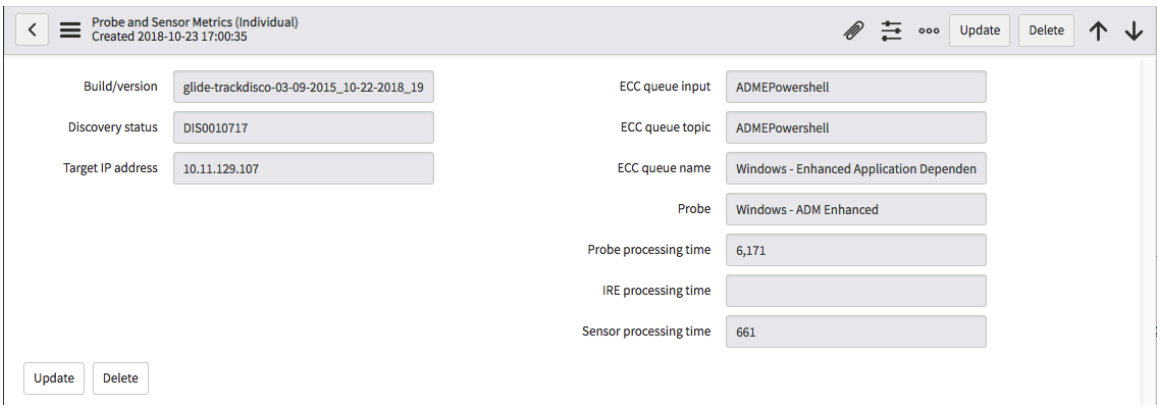

The Probe and Sensor Metrics (Individual) form provides these fields:

Field label	Field name	Description
Build/version	build_version	Build on which the Discovery was run.
Discovery status	discovery_status	ID number of the Discovery status from which the metrics were collected.
Target IP address	target_ip	IP address of the target for this Discovery.
ECC queue input	ecc_queue_input	Identifies a particular ECC input record in the ECC queue table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field label	Field name	Description
ECC queue topic *	ecc_queue_topic	Identifies the Java class in the MID Server that executes the probe.
ECC queue name *	ecc_queue_name	Identifies the probe/pattern evaluated for performance in this aggregation.
Probe *	probe	Name of the probe used for this Discovery.
Probe processing time	probe_time	Interaction time of the MID Server with the target, including construction of the payload that is sent back to the instance. The time is in milliseconds.
IRE processing time	ire_time	Time required to process the pattern payload on the instance by the Identification and Reconciliation Engine (IRE). IRE time is useful because it shows the part of the sensor time used by the pattern. The time is in milliseconds.
Sensor processing time	sensor_time	Time it took the sensor to process the payload on the instance for a Discovery. The time is in milliseconds.

* Used to uniquely identify a probe/pattern and accompanying sensor when gathering metrics for a Discovery.

View Discovery performance metrics aggregated by build

Use the roll-up by build data to ensure that the processing times for Discovery components remain consistent for discoveries in a 24 hour period. View aggregate build data before and after an upgrade to compare the performance of the old and new versions. All aggregated performance data is read-only.

Before you begin

Role required: discovery_admin, discovery_operator, or admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Performance Metrics** > **Probe/Sensor (Rollup-By-Build)**.
2. Sort the list by **Build/version**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

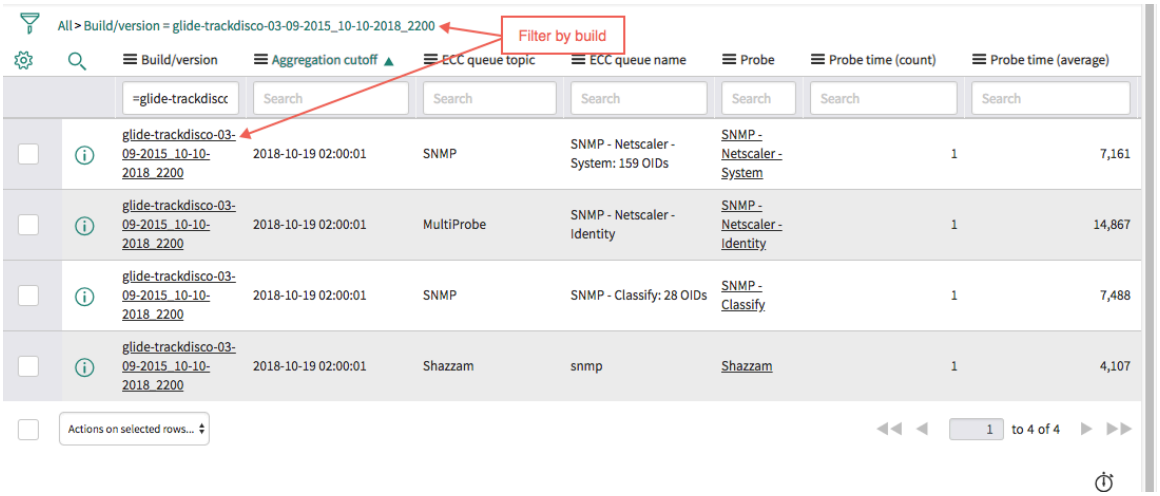

3. Filter by a specific build to see the aggregated processing times for the probes and patterns that performed a Discovery on that build.
4. Open a record to see the probe/pattern statistics for the selected build.
   The form displays additional fields not visible on the list. Roll-up calculations are over a 24 hour period, beginning every night at 0200.
5. See the table of performance aggregation data for descriptions of additional metrics displayed on the form for roll-ups by build.

View Discovery performance metrics aggregated by status

Use the roll-up by status data to ensure that the processing times for probes/patterns and sensors remain consistent for a specific Discovery. All aggregated performance data is read-only.

Before you begin

Role required: discovery_admin, discovery_operator, or admin

About this task

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Performance Metrics** > **Probe/Sensor (Rollup-By-Status)**.
2. Sort the list by **Discovery status** to see the aggregated processing times for the probes and patterns that ran during a specific Discovery.

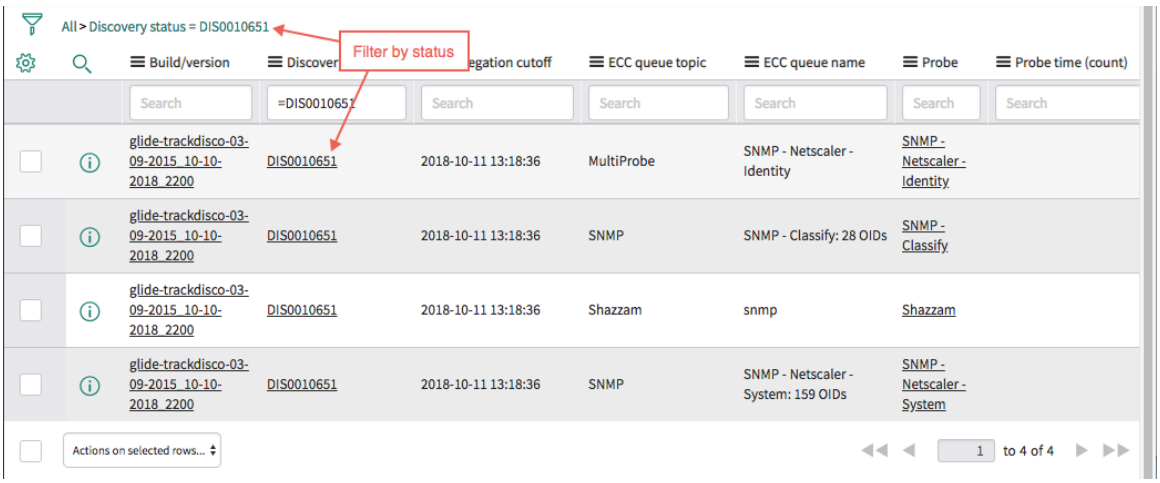

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Filter by a specific status to display metrics available for probes and patterns that ran in that status.
4. Open a record to see the probe/pattern statistics for the selected status.
   The form displays additional fields not visible on the list. Roll-ups are only created for a completed or cancelled status.
5. See the table of performance aggregation data for descriptions of additional metrics displayed on the form for roll-ups by status.

View Discovery performance metrics aggregated by IP address

Use the roll-up by target data to ensure that the processing times for probes/patterns and sensors remain consistent for each Discovery of a specific IP address. All aggregated performance data is read-only.

Before you begin

Role required: discovery_admin, discovery_operator, or admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Performance Metrics** > **Probe/Sensor (Rollup-By-Target)**.
2. Sort the list by **Target IP address**.

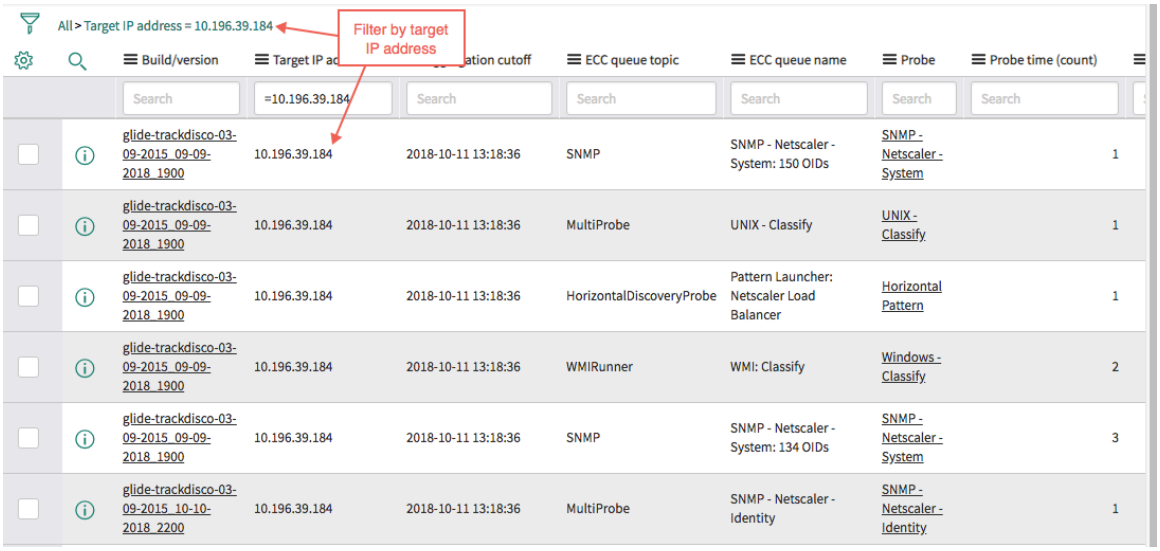

3. Filter by a specific IP address to see the aggregated processing times for the probes and patterns that performed the Discovery of that IP address.
4. Open a record to see the statistics about the selected probe/pattern for the specific IP address.
   The form displays additional fields not visible on the list. Roll-ups are performed after IP Discovery has completed successfully.
5. See the table of performance aggregation data for descriptions of additional metrics displayed on the form for roll-ups by target.

Aggregated data for Discovery performance metrics

Discovery performance metrics can accumulate data for probes, patterns, and sensors each time Discovery runs. Discovery calculates processing times and increments the number of times a component runs for each roll-up profile: status, target, or build. All aggregated performance data is read-only.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample roll-up form

This is an example of an aggregation record for probe and sensor metrics. The metrics fields shown here are used for each aggregation.

Sample roll-up by Status form

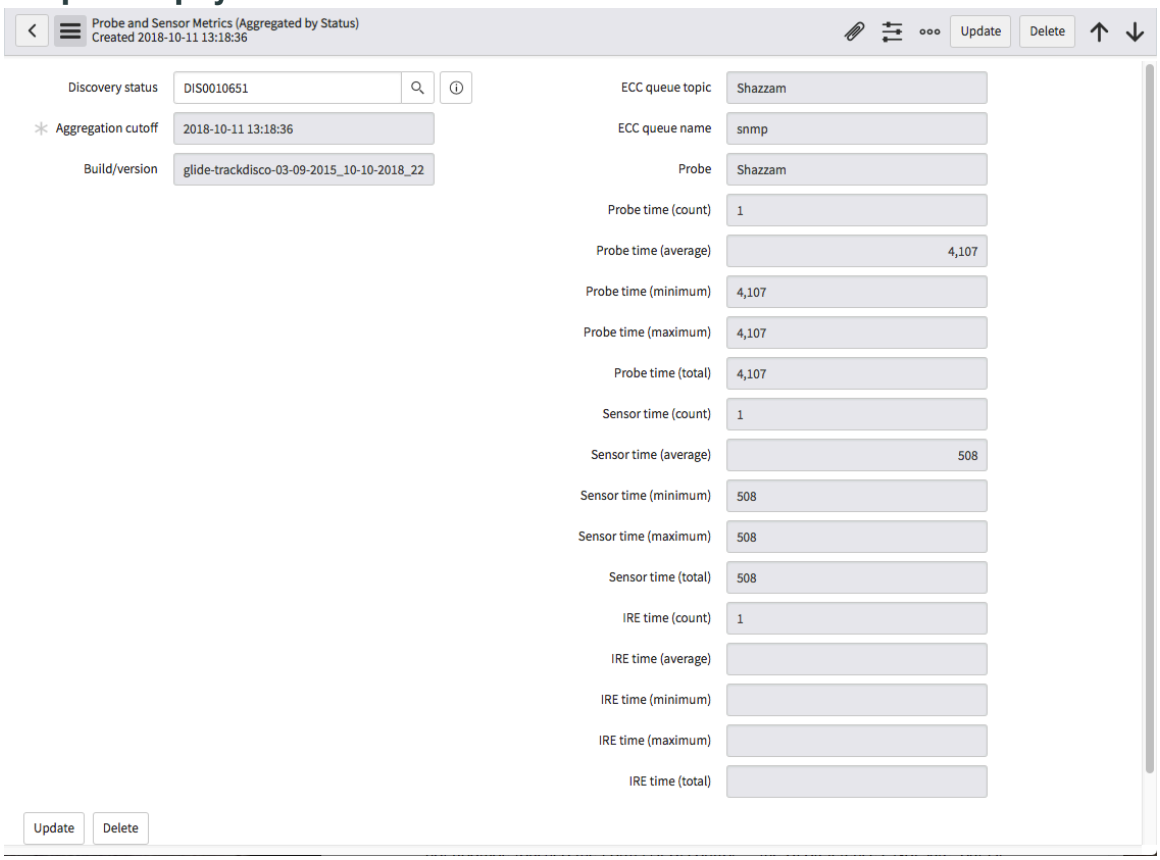

Performance Framework aggregated data

Except where noted, these fields are common to all aggregation records.

Field label	Field name	Description
Build/version	build_version	Name of the build running on the instance. This name identifies the version, patch level, and release date of the ServiceNow platform.
Discovery status	discovery_status	ID number of the Discovery status record for this aggregation. **Note:** This field only appears on the form for roll-ups by status.
Target IP address	target_ip_address	IP address of the target for this Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field label	Field name	Description
		Note: This field only appears on the form for roll-ups by target.
Aggregation cutoff	aggregation_cutoff	The cutoff time varies, depending on the roll-up profile. • **By-build**: Occurs daily at 02:00, by default. • **By-status**: Closing time of the last aggregation for that Discovery status, which might have run the last time **discovery.complete** or **discovery.cancelled** ran for that status. • **By-target**: Closing time of the last aggregation for that target IP address, which might have run the last time **discovery.device.complete** ran for that IP address.
ECC queue topic *	ecc_queue_topic	Identifies the Java class in the MID Server that executes the probe.
ECC queue name *	ecc_queue_name	Identifies the probe/pattern evaluated for performance in this aggregation.
Probe *	probe	Name of the probe used for this Discovery.
Probe time (count)	probe_time_count	Number of times a probe ran for a given roll-up profile.
Probe time (average)	probe_time_average	Average time a probe took to gather data on the target and format the payload for a given roll-up profile.
Probe time (minimum)	probe_time_min	Minimum time a probe took to gather data on the target and format the payload for a given roll-up profile.
Probe time (maximum)	probe_time_max	Maximum time a probe took to gather data on the target and format the payload for a given roll-up profile.
Probe time (total)	probe_time_total	Total time used by a probe to gather data on the target and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field label	Field name	Description
		format the payload for a given roll-up profile.
Sensor time (count)	sensor_time_count	Number of times a sensor processed payloads for a given roll-up profile.
Sensor time (average)	sensor_time_average	Average time a sensor took to process payloads on the instance for a given roll-up profile.
Sensor time (minimum)	sensor_time_min	Minimum time a sensor took to process a payload on the instance for a given roll-up profile.
Sensor time (maximum)	sensor_time_max	Maximum time a sensor took to process a payload on the instance for a given roll-up profile.
Sensor time (total)	sensor_time_total	Total time used by a sensor to process payloads on the instance for a given roll-up profile.
IRE time (count)	ire_time_count	Number of times a pattern's payload was processed by the Identification and Reconciliation Engine (IRE) for a given roll-up profile.
IRE time (average)	ire_time_average	Average time used for IRE processing of a pattern's payload for a given roll-up profile.
IRE time (minimum)	ire_time_min	Minimum time used for IRE processing of a pattern's payload for a given roll-up profile.
IRE time (maximum)	ire_time_max	Maximum time used for IRE processing of a pattern's payload for a given roll-up profile.
IRE time (total)	ire_time_total	Total time used for IRE processing of a pattern's payload for a given roll-up profile.

* Together, these values uniquely identify a probe/sensor pair (a "probe execution") that is used for a Discovery.

Discovery error messages

Review common Discovery error messages and their causes to identify and resolve issues that occur during discovery runs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Common error messages

Input/output errors

- CONNECTION_FAILED
- CONNECTION_TIMEOUT
- CONNECTION_REFUSED
- CONNECTION_CIPHER_UNSUPPORTED
- CONNECTION_PROTOCOL_UNSUPPORTED
- CONNECTION_LOST
- INVALID_CONNECTION_PARAMETER
- FILE_NOT_FOUND
- IO_ERROR

Authentication errors

- NO_CREDENTIALS
- AUTHENTICATION_FAILED
- INVALID_CREDENTIALS_TYPE

SSH errors

- SSH_SHELL_UNSUPPORTED
- SSH_SCRIPT_TRANSFER_FAILED
- SSH_INVALID_SESSION

Windows errors

POWER_SHELL_VERSION_UNSUPPORTED

Command errors

- COMMAND_TIMEOUT
- INVALID_COMMAND
- COMMAND_FILTER_ERROR
- COMMAND_PARSER_ERROR
- COMMAND_KEY_MAPPER_ERROR
- COMMAND_VALUE_NORMALIZER_ERROR
- TERMINAL_SESSION_LOST_ERROR
- TERMINAL_SESSION_ENVIRONMENT_MODIFICATION_ERROR

Generic component errors

- INVALID_COMPONENT
- INVALID_COMPONENT_PARAMETER
- INVALID_OPERATION

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

General error messages

Message: org.xml.sax.SAXParseException: The entity name must immediately follow the '&' in the entity reference

A version of this message occurs if you have used special characters in a password that is saved in an XML file.

Message: Identified, ignored extra IP

This message can appear at the end classification phase of Discovery if a targeted IP address belongs to a device that is being discovered at the same time.

For example, a Windows server has two NIC cards with two IP addresses. Discovery targets both IP addresses within the same Discovery schedule. This message is generated to note that second IP address is ignored because we don't want to update the same CI twice within the same Discovery run.

This message is a warning and is expected. No action is needed.

Message: Authentication failures

The discovery process could not discover the CI because the discovery application could not authenticate. To resolve, add the credentials of that machine in to the discovery credentials table.

Message: Identified, not updating CI, now finished

No match on any of the CI Identifiers.

Message: The impersonation of user failed

This message originates in the Powershell. Check that the domain is specified along with the user name in the credentials.

Message: Connection failed to WMI service. Error: Permission denied

This message originates in WMI. Check that the MID Server service is running with the correct credentials and has access to the target device.

To check this, run the following command from the command prompt on the MID Server host:

```
wmic /node:target /user:user /password:password path
 win32_operatingsystem
target = IP address of target device
user = user account used by the mid server service
password = password used by the mid server service
```

Message: Connection failed to WMI service. Error: The remote server machine does not exist or is unavailable

This message originates in WMI. Check that the MID Server service account has access to the targeted machine. Check if a domain admin account is used as the MID Server service account. Check if any existing firewalls are open to the connection. To check this, run the following command from the command prompt on the MID Server host.

Execute for runner_type=WMI:

```
wmic /node:"<target>" /user:"<user>" /
password:"<password>" path win32_operatingsystem
From within a Powershell console on the mid server
 host, execute for runner_type=Powershell:
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
gwmi win32_operatingsystem -computer <ip> -credential
 '<username>'
```

Message: Provider is not capable of the attempted operation

WMI repository was corrupted. After following the "WMI: Repository Corruption, or Not?" article on the Microsoft TechNet site, the problem was corrected.

Message: The result file can't be fetched because it doesn't exist

PRB581515 - Powershell does not work when customer has locked down write rights to admin share.

Message: Please run sneep as root to ensure correct serial number from fserial data source

The Oracle **Sneep** command line tool must be installed for the Solaris - Serial Number probe to work correctly. There is a known limitation with Fujitsu PRIMEPOWER devices. To work around this limitation, run the Solaris discovery with root credentials.

Discovery sensor error messages

Message: The multisensor will not process because its major version = X while probe_name responding script's major version = Y

This error message indicates that there is a major version mismatch.

Message: Payload length exceeds limit of 5000000

This error message indicates that the Discovery probe results exceed the limit. The MID Server will discard the results and return a warning message. To resolve, set the value of the *mid.discovery.max_payload_size* MID Server property to any negative number (for example, -1) to disable the payload limit and allow payloads of any size to be sent to the instance.

Message: The multisensor will not process because its responding script's major version = X while its referenced probe probe_name major version = Y

This error message indicates that there is a major version mismatch.

Message: sensor_name sensor's major version = X while its related probe's major version = Y

This error message indicates that there is a major version mismatch in the probe and sensor versions and the sensor will stop processing until this condition is resolved.

Message: sensor_name multisensor's minor version = X while probe_name responding script's minor version = Y

This error message indicates that there is a minor version mismatch.

Message: sensor_name multisensor's responding script's minor version = X while its referenced probe probe_name minor version = Y

This error message indicates that there is a minor version mismatch.

Message: sensor_name sensor's minor version = X while its related probe's minor version = Y

This error message indicates that there is a minor version mismatch in the probe and sensor versions. Processing will continue, but you may want to resolve this condition.

Message: Sensor error when processing . . . : No sensors defined

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Every active probe looks for a corresponding sensor to process the data that is collected by the probe. This error message indicates that the corresponding sensor for the probe is missing or inactive.

Message: Sensor error when processing . . . : typeError: . . .

This error message indicates that the sensor has one of the core error constructors in JavaScript.

Message: Sensor error when processing ...: Exception while processing sensor: XML Parse Problem: ![CDATA]

This XML parsing error occurs when the discovery sensor fails to parse the input ECC queue payload. The payload is not formatted according to XML specifications. Possible solutions include:

- Escape problematic characters. For example, change **<** to **&lt** and **&** to **&amp**.
- Escape entire blocks of text with CDATA sections, or put encoding declarations at the start of the text.
- Remove any whitespace characters (space, tabs, newlines) before the XML declarations.

Message: Sensor error when processing Linux - Network ARP Tables: Exception while running probe post processing script: No XML data

The system displays the No XML data error when the XML sensor processor fails to find the expected XML data in the probe output. During sensor processing, Discovery attempts to retrieve the probe results but finds the probe output empty. Verify that the probe returns the correct output for the sensor to process.

Message: Sensor error when processing Shazzam: Exception while running probe post processing script: Probe not found: null

The Probe not found error occurs when the sensor processor fails to find the probe record supposedly associated with the sensor. During sensor processing, Discovery gets the probe record from the probe cache and stores it for later reference. The "Probe not found" error occurs either when the sys_id of the probe cannot be found or there is an issue with the probe cache.

See Find the cause of a "Probe not found" error for more information.

Find the cause of a sensor error message

Use this procedure to identify a Discovery sensor error.

Before you begin

Role required: sn_cloud_ops_ws.cloud_ops_admin or admin

About this task

This process requires that you identify the following:

- The error.
- The name of the sensor where the error occurred.
- The datestamp of the error.
- The associated stack trace of the error.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Depending on the application you are using, do one of the following:
   - To view the **Discovery Log** in Cloud Discovery Workspace, navigate to **Cloud Discovery Workspace** > **Cloud Discovery** > **Discovery Log**.
   - To view the **Discovery Log** in Discovery, navigate to **All** > **Discovery** > **Discovery Log**.
2. To find the error, filter on the **Level name** of **Error**.
   If you are using Discovery, you can also view the error and the date stamp in the Error Report. The report is available on the Discovery dashboard.
3. Note the **Created** field of the error record.
   This timestamp is used to find the proper stack trace.
4. Note the **Short Message** field of the error record.
   This field shows the name of the sensor where the error has occurred.
5. To find the stack trace, navigate to **System Logs** > **Errors**.
6. Search for the date and time that matches the value in the **Created** field of the error record in the Discovery Log.
7. The **Message** field of that entry contains the full stack trace of the error.
   The stack trace contains the error message, the sys_id of the script, and the line number where the error has occurred.
8. After determining the error details, you can fix the JavaScript file.

Find the cause of a "No Sensor Defined" error message

Every active probe looks for a corresponding sensor to process the data that is collected by the probe. The "No Sensors Defined" message indicates that the corresponding sensor for the probe is missing or inactive.

Before you begin
Role required: sn_cloud_ops_ws.cloud_ops_admin or admin

About this task
To find the cause of this error, identify the **Short Message** field of the error record. This field shows the name of the probe for which the sensor is missing.

Procedure

1. Depending on the application you're using, do one of the following:
   - To identify the cause of the error using Cloud Discovery Workspace, navigate to **Cloud Discovery Workspace** > **Cloud Discovery** > **Diagnostics** > **Discovery Log**.
   - To identify the cause of the error using Discovery, navigate to **All** > **Discovery** > **Output and Artifacts** > **Discovery Log**.
2. Filter the **Level** column by **Error**.
   If you are using Discovery, you can also view the Error Report on the Discovery dashboard.
3. Note the **Short Message** field of the error record.
   This field shows the name of the sensor where the error has occurred.

Fix the cause of a sensor error message

To fix a Discovery sensor error message, you must fix the JavaScript file containing the code that generated the error.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Verify the following:

- You have access to the Discovery Definition modules.
- You have completed the steps in Find the cause of a sensor error message.

Role required: admin

About this task
Use the sensor name and stack trace you identified in Find the cause of a sensor error message to locate the failing code. The stack trace provides the line number and function name where the error occurred.

Note: The exact fix depends on the specific sensor and payload. This topic describes where and how to fix the error, not specific code changes. For errors with a refined code, you can view the root cause and remediation steps directly on the Discovery Admin Workspace Error Details page.

Procedure
1. Navigate to the Discovery Definition module where the error occurred.

   For example, navigate to **Discovery Definition** > **Script Includes**.
2. Search for the sensor name in the **Name** field.
3. Select the link to open the record.
4. Locate the failing function or line using the stack trace, then correct the JavaScript.
   For example, if the error indicates it "Cannot read property X of undefined", add a defensive check around the payload fields.
5. Select **Update**.

What to do next
Rerun Discovery to verify the sensor no longer generates the error. If the error persists after applying the fix, contact ServiceNow Support and provide the full stack trace, the sensor name, and a description of the changes you made.

Find the cause of a "Probe not found" error

Verify the correct sys_id of the probe associated with the *probe* parameter in the payload of the ECC Queue input record.

Before you begin
Role required: discovery_admin or admin

Procedure
1. Record the sys_id of the probe as it appears in the ECC Queue payload.

ECC Queue payload sample

```
▼<parameters>
   <parameter name="port" value="22"/>
   <parameter name="probe" value="10e0eebd0a0a0b4f61f46a5027df7fb6"/>
   <parameter name="priority" value="0"/>
   <parameter name="use_class" value="discovery_classy_unix"/>
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Navigate to **Discovery Definition** > **Probes** and open the record for the probe you want to inspect.
3. Right-click in the record's header bar.
4. Select **Copy sys_id** from the context menu.
   Follow browser instructions to copy the sys_id if browser security measures restrict this function.
5. Compare this value with the sys_id of the probe from the payload.
   If the sys_id of the probe record does not match the value in the payload, try to determine the cause of the incorrect value.

Discover missing compute resources

Use the Discovery Admin Workspace to identify compute resources that weren't discovered recently. Run a quick discovery to update the resource's data. If a resource isn't discoverable, ignore the resource and exclude it from the report. This procedure helps you maximize the use of your resources.

Before you begin

Role required: discovery_admin or admin

About this task

By default, the report provides the following information:

- Resources discovered by the ServiceNow® platform and ServiceWatch Suite.
- Resources that weren't discovered since the previous discovery ran 14 days previously or that have no discovery date in the **Most recent discovery** field.
- Resources that don't have an **ignore from report** tag.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** > **Missing compute resources from Discovery**.
2. Run a quick discovery for a resource.

   a. In the list, select the resource.

   b. Select **Quick discovery**.

   c. Select a MID Server to run discovery.
   The **Enter IP List** is automatically populated.

   d. In the **MID Selection Method** and **MID Server** fields, select the preferred options.
   For more information, see MID Server selection.

   e. Select **Run discovery**.

3. If the resource's discovery failed, ignore the resource and exclude it from the report.

   If you successfully ignore a resource, a green banner appears to indicate that the resource is ignored from the report. When the resource is ignored, it's removed from the **Missing compute resources** report.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the list, select the resource.

b. Select **Ignore from report**.

4. **Optional:** If you want to enable the discovery for a previously ignored resource, remove the **ignore from the report** tag.

a. Navigate to **All** and type in the filter `cmdb_ci_server.list`.

b. Personalize the list to include the **Tags** column, using the settings button.

c. Search for the resource name.

d. In the **Tags** field, select the removal option.

Result

If the quick discovery was successful, the instance redirects to the **Discovery status** dashboard. When the resource is discovered, it's removed from the **Missing compute resources** report and is listed in the Server [cmdb_ci_server] table.

Run Tuning checks

Tuning checks are scans aimed at optimizing configurations, grouped into suites with each specific check detailed below.

Before you begin

Role required: discovery_admin or admin

About this task

The configured checks are listed under the following tabs:

- Discovery suite
- MID suite
- Self Monitoring suite

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Tuning**.
2. Select **Checks**.
3. To manually run a check for Discovery suite, MID suite, or Self Monitoring suite, select **Run now**.

Note:

The instance scan can only run one at a time, so wait for the in-progress scan to finish before re-running the suite.

4. To access more information about an issue, select the check name.

Note:

- To understand the issue, select the **Show check description** link.
- To know more details on the issue, select the **Documentation on the check** link.
- To view the definition page, select the **Check definition page** link. To edit a check, see Edit an existing check.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Edit an existing check

By editing the definition check, you can change description and script.

Before you begin

Role required: admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Tuning**.
2. Select the check name.
3. Select the **Check definition page** link.
4. To edit the check, select **Switch scope**.
5. Update all the required details and select **Save**.

Note: If the Active option is enabled:

- Discovery suite runs once a week on Sunday at 8 AM.
- MID suite runs once a week on Sunday at 10 AM.
- Self monitoring suite runs every 15 mins.

Resolve Cloud Discovery errors in Cloud Discovery Workspace

View the Cloud Discovery errors that occurred during the discovery and resolve them. You can view the errors for all the Cloud Discovery runs or a single Cloud Discovery run.

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

Before you begin

Role required: sn_cloud_ops_ws.cloud_ops_admin

Procedure

1. Navigate to **Cloud Discovery Workspace** > **Cloud Discovery**.
2. Select **Errors**.
   The **Status** tab of the Schedules module displays the list of errors for the selected discovery run.
3. Select a category tile to display the specific error codes that occurred in that category.
4. Select the error code tile that you want to investigate.
   The list displays all occurrences of that error code for all the Cloud Discovery schedules. The list shows the error message, Configuration Item (CI) ID, IP address of each instance that experienced the error, and the error status.
5. Perform one of the following actions:
   - Select the specific errors, and then select the appropriate action from the **Selected errors** section.
   - Select appropriate action from the **All errors** section.

     Cloud Discovery applies the selected action to all errors of the list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Available actions

Action	Description
Create task	You can create an error task and assign it to an user for resolution. Assigned errors are considered **Active errors**, and actions can still be performed on them. Cloud Discovery displays the error task number on the error code tile.
Retry discovery	Cloud Discovery retries an IP-based discovery. When you select **Retry discovery**, Cloud Discovery performs this action on all the active errors that are available in the list. You can retry the discovery, only if the ServiceNow® IP-based Discovery plugin (com.snc.discovery.ip_based) is installed and activated in the instance.
Ignore	Cloud Discovery ignores the error.

6. Check the **Error Status** column to see the results of the remediation.
   Single retries complete quickly. Retries for large numbers of IP addresses can take several minutes.

   An error can have one of the following statuses:

   - **Active error**: Unresolved error. This is the default status of all new errors.
   - **Resolved**: Resolved error.
   - **Assigned**: Task assigned to resolve this error.
   - **Pending Discovery**: Waiting for Discovery to start after you execute the **Retry discovery** action on all errors.
   - **In Discovery**: Discovery is currently active.

Discovery start and classification issue resolution

Use the Discovery Admin Workspace Diagnostics page to identify and resolve the cause when Discovery fails to start or classify a device.

Most start and classification failures have one of a few causes: the MID Server can't reach the target, the configured account lacks the required rights on the target, or software on the target host blocks Discovery scripts from running. The Diagnostics page identifies most of these causes with a refined code that includes the root cause and remediation steps.

Requirements

The **Diagnostics** page, refined codes, and AI Insights described here require the Error Framework. Confirm the following:

- The ServiceNow AI Platform is running the Brazil release, or the Australia release beginning with Patch 3.
- Discovery Admin Workspace, beginning with v1.17.0, is installed.
- The *sn_disco_workspace.enable_error_framework* system property is set to `true` (default behavior).
- You have the discovery_admin role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To use AI Insights, install the plugins listed in the following table.

Plugin name	Plugin ID	Availability
ServiceNow Otto for Error Framework	com.sn_ef_gen_ai	Australia Patch 3
ServiceNow Otto	sn_genai_platform	Required by ServiceNow Otto for Error Framework

Issue resolution process

The **Diagnostics** page displays Discovery errors that have been refined and assigned a specific error code. Each refined code identifies the error type and includes the root cause and recommended remediation steps. Start and classification failures fall into two groups. When a cause can be refined, the **Diagnostics** page provides guidance to help resolve it. When a cause cannot be refined, no error code is available and you must verify the cause manually. Work through the causes in the order that follows.

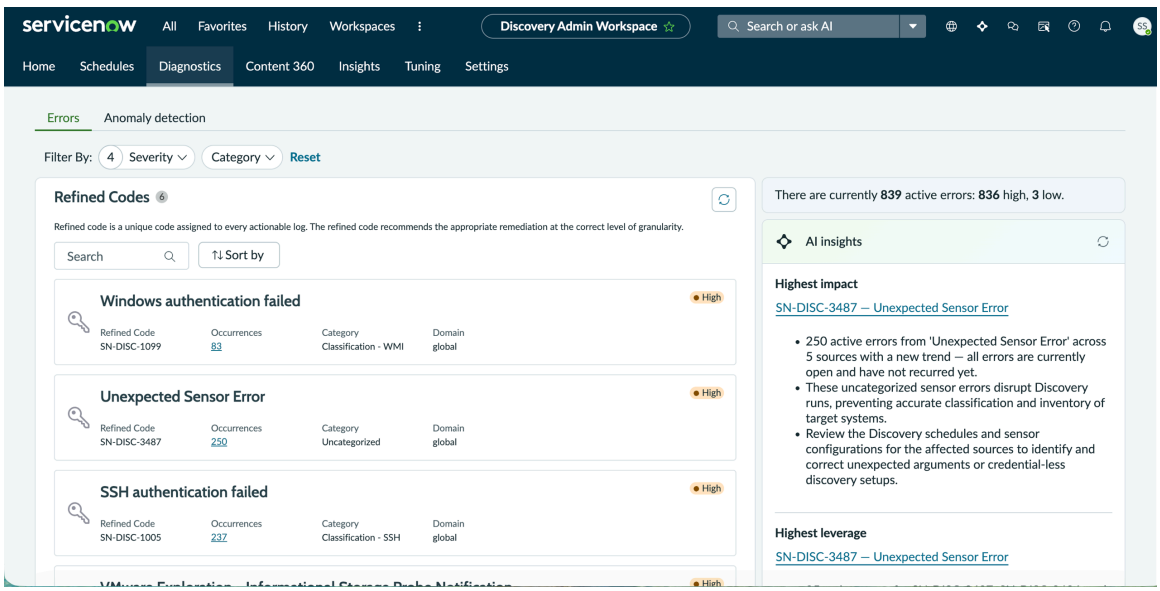

Review the Errors tab

Navigate to **Workspaces** > **Discovery Admin Workspace** > **Diagnostics** and review the **Errors** tab. Locate the error by refined code, then open it to view the root cause and remediation steps. Errors that stop Discovery for a device or schedule display with **Critical** severity.

The following refined codes commonly correspond to start and classification failures.

Refined code	Cause
SN-DISC-5602	MID Server isn't running on Windows.
SN-DISC-1007	SSH connection failed.
SN-DISC-3486	Probe skipped because a previous probe failed.
SN-DISC-3888	Discovery ECC Queue is missing or empty.
SN-DISC-6004	The target's SSL certificate isn't trusted.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Refined code	Cause
	Note: Requires Discovery Admin Workspace v1.20.0.

Verify connectivity and credentials

If Discovery still fails after you follow the remediation steps for a refined code, verify connectivity and credentials.

Confirm that the MID Server can reach the target on the ports that the protocol requires, and that no firewall blocks that path. For Windows targets, confirm that inbound traffic is enabled on the SMB port 445.

Confirm that the account Discovery uses has the rights on the target that the protocol requires. For Windows targets, confirm that the account can create and read files in the administrative share used during classification.

Check host configuration

The following conditions don't produce a refined code. If the earlier refined codes don't resolve the failure, confirm each one manually.

PowerShell execution policy

The execution policy on the target host must allow Discovery to run PowerShell commands, including any policy applied through Group Policy. If a restrictive policy is blocking scripts, update the policy or add an exemption for Discovery.

Endpoint security

Confirm that endpoint protection software on the target isn't blocking the Discovery PowerShell script. If script control is enabled, add an exclusion for the Discovery script and add the ServiceNow publisher to the trusted publisher list.

Legacy Diagnostics experience

The refined codes and AI Insights described here require the Error Framework. If your instance runs on an earlier version of Discovery Admin Workspace, the **Diagnostics** page shows the legacy experience instead. In that case, review the Diagnostics tab to review Discovery errors by error code and severity, then apply the same environmental checks.

Related topics

Error Framework in Discovery Admin Workspace

Discovery Admin Workspace Error Details

Discovery Platform Analytics Solutions

Platform Analytics Solutions contain preconfigured dashboards. These dashboards contain actionable data visualizations that help you improve your business processes and practices.

Platform Analytics Solutions

Platform Analytics data visualizations use Performance Analytics indicator data to show you data over time, helping you analyze your business processes and identify areas of improvement. With Platform Analytics Solutions, you can get value from Performance Analytics for your application with minimal setup. You can always create your own objects as well.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To enable the solution plugin for Discovery, an admin can navigate to **System Definitions** > **Plugins** and activate the Performance Analytics - Content Pack - Discovery plugin.

Related topics

Activate your Performance Analytics subscription

Discovery for data-center virtualization

Discovery and Service Mapping Patterns identify and classify information about virtual machines discovered in a data-center.

Data collected

Label	Table Name	Field Name
CPUs	cmdb_ci_vm_instance	cpus
Disks	cmdb_ci_vm_instance	disks
Disks size	cmdb_ci_vm_instance	disks_size
Image path	cmdb_ci_vmware_instance	image_path
Memory	cmdb_ci_vm_instance	memory
Network adapters	cmdb_ci_vm_instance	nics
State	cmdb_ci_vm_instance	state
Template	cmdb_ci_vmware_instance	template
vCenter Reference	cmdb_ci_vmware_instance	vcenter_ref
vCenter Reference UUID	cmdb_ci_vmware_instance	vcenter_uuid

Discovery for VMware virtualization

A Discovery schedule for VMware virtualization discovers vCenter and ESX hosts and individual ESXi hosts that manage VMs and related components without a vCenter.

The following options are available for discovering VMware vCenter data:

- Discovery runs the VMware - vCenter Datacenters probe when it identifies a VMware vCenter process running on a Windows machine or detects activity with the vmapp port probe.
- Classic Orchestration runs any of the vCenter probes from a workflow.

Use standalone ESXi discovery to discover VMware data without a vCenter.

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

VMware Workstation

In the basic VMware system, the VMware Workstation runs on a Windows or Linux host machine, but not managed directly thorugh vCenter.

This system clones instances from templates, but can't be automated. The relationships between VMware components for this type of installation are in the following diagram:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware Workstation relationships

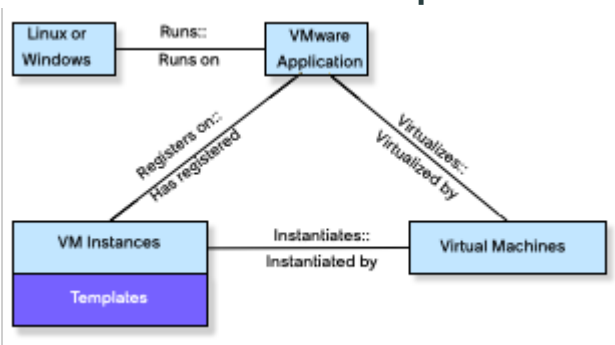

VMware Workstation Component relationships

Component	Relationships
Windows or Linux Server	Runs the VMware application
VMware application	• Runs on a Windows or Linux host machine • Has registered VM instances • Virtualizes virtual machines
VM Instances (including images and templates)	• Registers on the VMware application • Instantiated by individual virtual machines
Virtual machines	• Instantiates VM instances • Virtualized by VMware application

Discovery reference

Reference topics provide additional information about Discovery lists and forms.

Cloud Discovery reference

Reference topics provide additional information about the lists, forms, and cloud settings used in Cloud Discovery.

Cloud Discovery service account form reference

The Cloud Discovery service account form displays detailed information about the selected cloud provider. The Cloud Discovery service account form is available in the Cloud Discovery Workspace.

Amazon Web Services (AWS) service account fields

Field	Description
Service Account	Name of the service account whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Name of the service account whose resources you want to discover.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon Web Services (AWS) service account fields (continued)

Field	Description
Account ID	12-digit Identity and Access Management (IAM) account number.
Credentials	Name of the credentials that you have created for AWS discovery. Leave the field empty if the account is a member account of an AWS organization and you've configured the associated management account with a credential. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).
URL	URL of the datacenter. Enter this field only for the AWS GovCloud (US) and AWS China region accounts.
Access role name	Name of the IAM role created for the trusting account.
Select account for access	Name of the trusted account. Configure this field only for accounts that don't use permanent AWS credentials, and rely on IAM roles for access.

Microsoft Azure service account fields

Field	Description
Service Account	Name of the service account whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Microsoft Azure subscription name.
Subscription ID/Management group ID	Microsoft Azure subscription ID.
Credentials	Name of the credentials that you have created for Microsoft Azure discovery. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).
URL	URL of the datacenter. Configuring this field is required for Azure GovCloud (US) accounts. Add the required URL.
Pull Events	Option to enable the ServiceNow AI Platform to collect resource change information from the Microsoft Azure cloud. The ServiceNow AI Platform uses the resource change information to update the Configuration Management Database (CMDB). For more information on the supported Azure resource types, see Azure change processing.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Azure service account fields (continued)

Field	Description
	If you select the Pull Events check-box for a management group, the ServiceNow AI Platform enables Azure change processing for all the subscriptions in the management group.

Google Cloud Platform service account fields

Field	Description
Service Account	Project ID of the Google Cloud Platform (GCP) project whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Project ID of the Google Cloud Platform (GCP) project.
Project ID/Folder ID	Project ID of the Google project.
Credentials	Name of the credentials that you have created for GCP discovery. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).
Pull Events	Option to enable the ServiceNow® event-driven discovery to collect events for the resource types hosted in the Google Project LDCs.

IBM service account fields

Field	Description
Service Account	Name of the service account whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Name of the service account.
Credentials	Name of the credentials that you have created for IBM discovery. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).
Account ID	IBM software user number that the instance should use to access the IBM Cloud Console.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Cloud Infrastructure service account fields

Field	Description
Service Account	Name of the service account whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Name of the service account.
Account ID	OCI Compartment id.
Credentials	Name of the credentials that you have created for Oracle Cloud Infrastructure (OCI) discovery. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).
Datacenter URL	Datacenter value. For example, `https://$service.ap-mumbai-1.oraclecloud.com`

OpenStack service account fields

Field	Description
Service Account	Name of the service account whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Name of the service account.
OpenStack account ID	Project ID of the OpenStack Project to be discovered.
Account credentials	Name of the credentials that you have created for OpenStack discovery. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).
Cloud URL	OpenStack identity service endpoint. For example: https://<openstack dns>:<port>.

Red Hat Virtualization (RHV) service account fields

Field	Description
Service account	Name of the service account whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Name of the service account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Red Hat Virtualization (RHV) service account fields (continued)

Field	Description
Manager URL	The account ID to which this credential belongs. Specify the API url of the RHV-M instance, including the port. For example, https://my-manager-instance.com:443.
Manager credentials	Name of the credentials that you have created for RHV discovery. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).

VMware service account fields

Field	Description
Service Account	Name of the service account whose resources you want to discover. This field appears only when the Existing service account option is selected.
Name	Name of the service account.
UUID	Universal Unique Identifier (UUID) of the instance. To access the UUID, do the following: **1.** Open the Managed Object Browser (MOB) for the vCenter. **2.** Select **Content** > **About**. **3.** Paste the value of the instance UUID property in this field.
Credentials	Name of the credentials that you have created for VMware discovery. To add a credential in the ServiceNow AI Platform, select the Add new credentials icon ( ) and fill in the fields. If required, you can select an existing credential and then edit it by selecting the Edit credentials icon ( ).
URL	Full URL of the vCenter.

Edit settings form reference

The Edit settings form displays information about the Amazon Web Services (AWS) settings used to record and deliver the AWS events to the listening endpoint.

Edit settings form

Field	Description
General settings	
Resource types to record	Resource types for which you want to record the configuration change. Also specify whether you want to record the configuration change for global resources such as AWS AMI resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Edit settings form (continued)

Field	Description
	For more information, see the Selecting Which Resources AWS Config Records topic in the AWS documentation.
Data retention period	(Optional) Retention policy for the recorded events.
AWS Config role	An AWS role that can access the following services and resources: • Simple Notification Service (SNS) • AWS Config service • Resource types for which you want to track the configuration change
Delivery method	
AWS S3 bucket	An AWS S3 bucket to store the recorded events. Create an AWS S3 bucket or select an existing AWS S3 bucket to which you have access.
Stream configuration changes and notifications to an Amazon SNS topic	Option to publish the event notifications to the SNS topic.
Amazon SNS topic	Amazon SNS topic to which the AWS Config service should publish the events. 1. Select the **Choose a topic from your account** option. 2. Select the Amazon SNS topic from the SNS topic name drop-down list.

Command Validation Tasks form for pattern-based discovery

Command Validation Tasks form displays the fields that you must fill when you create or modify a task for testing a pattern command.

For more information about testing pattern commands, see Validate commands used in pattern-based discovery and View command validation task results in pattern-based discovery.

Field	Description
Number	The task number is automatically created by the system.
MID Server Selection Method	Select the relevant method: • **Specific MID Server** - To define the MID Server for this test. • **Automatic MID Selection** - To use a MID Server that's assigned to the IP range or IP network that you want to use for the test. Select this method only if you have

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	configured a MID Server with IP ranges or as a default MID Server for Discovery. • **Try All Eligible MID Servers** - To use any MID Server available for the the IP range or IP network that you want to use for the test. Select this method only if you have configured a MID Server with IP ranges or as a default MID Server for Discovery.
MID Server	The MID Server name. This field appears only if **Specific MID Server** is selected from **MID Server Selection Method**.
Validate All Pattern Commands	The check box for validating all commands used by a specific pattern simultaneously.
OS Class	The CI class that represents the servers on which you run this validation task.
Command Type	The command type that depends on the kind of the server: • **Shell Command** • **WMI Query** • **SNMP Query** (Tabular and Scalar) • **SNMP Walk** • **HTTP GET** (for Basic Authentication credentials, not Token) This field is available only when **Validate AI Patterns Commands** isn't selected.
IP Addresses	Any combination of IP v4 Addresses, IP Ranges, or IP networks. For example: • `10.196.39.12, 10.196.39.12` - Comma-separated IP addresses for a list of standalone addresses. • `1.1.1.1-1.1.1.100` - Hyphen-separated IP addresses for an IP range. • `1.1.1.1/28` - Slash notation for an IP network. **Note:** To use IP range or IP network, ensure you have installed the ParseIPRanges script include. For more information, refer to KB1123625: Command Validation Tool.
Commands	The commands for which you want to test credentials. Typically, it's a command that

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	discovery process failed to execute. For example, **/api/v1/pods?limit=50** or **sudo dmidecode**. You can select only commands from the list of available commands. This field is available only when **Validate All Patterns Commands** isn't selected.
Pattern	The name of the pattern for which you want to validate the commands. This field is available only when **Validate All Patterns Commands** isn't selected.

Related topics

- Command List form for pattern-based discovery
- Discovery commands for probes and patterns

Command List form for pattern-based discovery

The **Command List** form contains all commands used in discovery by infrastructure or application type patterns.

For information on viewing a list of commands used by a pattern, see View Patterns commands through the Discovery Patterns module and View Patterns commands through the Command List module.

Command List properties

Attribute	Description
Applicable OS Class	The command is intended to run on this OS class.
CI Type	The pattern that contains this command is assigned to this CI type. **Note:** Some ServiceNow applications refer to CI types as CI classes.
Command	The discovery command.
Command Type	Type of command: • **Shell Command** • **WMI Query** • **SNMP Query** (Tabular and Scalar) • **SNMP Walk** • **HTTP GET** (for Basic Authentication credentials, not Token)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command List properties (continued)

Attribute	Description
Is Failover Step	This command is used in a failover step inside a pattern. Possible values are true or false.
Outdated	This field is for internal use. Possible values are true or false.
Pattern	The command is part of this discovery pattern.
Step name	Name of the step that contains this command.
Domain	The command is validated for discovering CIs in this domain. This field is applicable for domain-separated environment.

Related topics

Command Validation Tasks form for pattern-based discovery

Discovery commands for probes and patterns

Discovery log details

Discovery logs contain information on the horizontal discovery process based on probes and patterns.

Discovery log related list

The Discovery Log related list contains information about the horizontal discovery based on probes.

The Discovery Log provides the following information:

Field	Input value
Created	Timestamp of the Discovery activity. Each timestamp defines the approximate time of the activity. Several Discovery events may occur in random order within a second.
Level	Classifies the activity into one of the following levels for general sorting: • Error • Information • Warning
Short Message	Informative message detailing the outcome of the activity or the Discovery progress. Look here for the result of a classify probe or for authentication failure.
ECC queue input	The related input record from the ECC queue for this discovery. You can also view these records from the **ECC Queue** related list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Input value
CI	Names a device for which a matching CI was found in the CMDB. Select the link to view the CI record for the device.
Source	Names the particular activity, such as the Shazzam probe or a UNIX classify probe.
Device	Lists the IP address of the CI discovered. All devices identified by IP address appear in the log, even if they refused all invitations to communicate. Any port activity from a device places it into the log, even if all subsequent efforts to identify it fail. Select the IP address of the device to view the events associated with discovering that device.

Horizontal Discovery Log

The Horizontal Discovery Log window contains information on pattern-based horizontal discovery.

Discovery Log Item	Description
Pre Pattern Execution	The actions that were run before the pattern was launched.
Selecting Pattern for Execution	The pattern that was run for discovery.
{Identification Section name}	Displays the results of the operations in the pattern. Expand the name to see each operation.
Pre Payload Processing Scripts	The results of scripts that were run before the payload was received.
Payload Processing	Details about the payload and how it was processed. Find errors that might have been encountered during various activities, such as the running of identification rules, updates to the CMDB, and so on.
Post Payload Processing Scripts	The results of scripts that were run after the payload was received.

Discovery Schedule form reference

The Discovery Schedule form provides details about the configuration items (CIs) targeted for discovery, the associated MID Server, and the schedule settings for when and how the discovery runs.

For instructions on this procedure, see Schedule a horizontal discovery.

Discovery Schedule Form

Field	Description
Name	Enter a unique, descriptive name for your schedule.
Discover	Select one of the following scan types:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule Form (continued)

Field	Description
	• **Configuration items**: Uses Discovery identifiers to match devices with CIs in the CMDB and update the CMDB appropriately. Perform a simple discovery by selecting a specific MID Server to scan for all protocols (SSH, WMI, and SNMP). Or, perform advanced discoveries with discovery behaviors. When you select a behavior, the *MID Server* field isn't available. **Note:** An IPv6 via address list scan is enabled on Discovery schedule for Configuration Items. • **IP addresses**: Scans devices without the use of credentials. These scans discover all the active IP addresses in the specified range and create device history records, but don't update the CMDB. IP address scans also show multiple IP addresses that are running on a single device. Identify devices by class and by type, such as Windows computers and Cisco network gear. The *Max range size* Shazzam probe property determines the maximum number of IP addresses Shazzam scans. See Configure Shazzam probe for details. • **Networks**: Discovers IP networks (routers and switches). Results from this search are used to populate the IP Network [cmdb_ci_ip_network] table in **Discovery > IP Networks** with a list of IP addresses and network masks. Network scans update routers and layer 3 switches in the CMDB. • **Service**: Discovers services for the Service Mapping application. See Schedule a top-down discovery by Service Mapping for instructions. • **Serverless**: Finds CIs without needing to run discovery on a host, or CIs on a proxy host that is already in the CMDB. See Serverless Discovery for more information. • **Cloud application**: Discovers only the cloud resources for the patterns that you specify. See Discovery for cloud environment for instructions. • **Cloud resources**: Discovers resources for one of the supported cloud providers. This option only appears when you run Discovery on a cloud service account. You can't select it from a new Discovery schedule. • **Certificates**. Discovers certificates based on URLs. Selecting this option adds the Certificate Discovery Type field: URL Certificate Discovery. See Run Certificate Discovery via individual URL scans for more information.
MID Server selection method	Select the method that Discovery uses to select a MID Server: • **Auto-Select MID Server**: Enable Discovery to select the MID Server automatically based on the discovery IP ranges you configure. To find a matching MID Server, you must configure MID Servers to use: ○ The Discovery application, or ALL applications. This setting authorizes the MID Server access from Discovery. ○ The IP range that includes the ranges you configure on the discovery schedule. **Note:** MID Server auto-select isn't supported with IPv6. • **Specific MID Cluster**: Use a preconfigured cluster of MID Servers. Select the cluster. You aren't required to specify one member of the cluster. The MID Server can't be part of multiple clusters, such as one that supports load balancing and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule Form (continued)

Field	Description

	one that supports failover. You can add any cluster regardless of the application that the MID Servers are assigned to. When you select the cluster, the **Discovery** application is automatically added when it doesn't exist for the MID Servers in the cluster.
	• **Specific MID Server**: Use only one MID Server. If that MID Server is part of a cluster, only that MID Server is used. The cluster isn't used. You can add any MID Server regardless of the application it's assigned to. The **Discovery** application is automatically added when it isn't already assigned for the MID Server you select. You can assign a specific MID Server for all types of Discover scans except **Service**.
	• **Use Behavior**: Use a behavior when a single schedule requires the use of multiple MID Servers to perform any of the following activities:
	  ○ Scans requiring multiple Windows credentials.
	  ○ A schedule that must execute two or more particular protocols (SNMP, SSH, or WMI) using more than one MID Server.
	  ○ Load balancing for large discoveries where a single MID Server would be inadequate.
	  ○ Scanning multiple domains.
	Note: The discovery schedule enforces domain separation. The MID Servers that are available for selection are limited to the same domain of the user who is configuring the schedule.
	See MID Server selection sequence for Discovery schedules for additional information.

MID Server	Select the MID Server to use for this schedule. This field is available if *MID Server selection method* is set to *Specific MID Server*, or if you discover IP addresses, networks, or web services. To verify that the MID Server you selected is up and validated, look at the MID Server dashboard.
MID Server Cluster	Select the MID Server cluster to use for this schedule. This field is available if *MID Server selection method* is set to *Specific MID Cluster*.
Behavior	Select a behavior configured for the MID Servers in your network. This field is available only if *MID Server selection method* is set to *Use Behavior*.
Active	Select the check box to enable this schedule. If you clear the check box, the schedule is inactive, but you can still run a discovery manually from this form, using the configured values.
Location	Choose a location to assign to the CIs that the schedule discovers. If this field is empty, then no location is assigned.
Max run time	Set a time limit for running this schedule. When the configured time elapses, the remaining tasks for the discovery are canceled, even if the scan isn't complete. Use

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule Form (continued)

Field	Description
	this field to limit system load to a desirable time window. If no value is entered in this field, this schedule runs until complete.
Max Idle time	Set a time limit for inactivity during discovery to automatically cancel stalled processes and log the event. A discovery is considered idle only when no input records are being processed and nothing is left in the processing state. Idle time is measured from the most recent change to an ECC Queue record. The configured idle time is fixed at the start of the discovery and won't reflect any edits until the next run. If no value is entered in this field, this schedule runs until complete or canceled. **Note:** By default, the **Max Idle time** field doesn't display in the Discovery Schedule form. To include this parameter, open the Discovery Schedules [discovery_schedule] table, add the column, and configure the Max Idle time directly in the table. For information about customizing columns in tables, see Personalize a list.
Run and related fields	Determines the run schedule of the discovery. Configure the frequency in the **Run** field and the other fields that appear to specify an exact time. **Note:** The run time uses the system time zone. If you add the optional **Run as tz** field, it has no effect on the actual runtime.
Log state changes	Select this check box to create a log entry every time the state changes during a discovery, such as a device going from *Active* to *Classifying*. View the discovery states from the Discovery Devices related list on the Discovery Status form. The *Completed activity* and *Current activity* fields display the states.
Shazzam batch size	Enter the number of IP addresses that each Shazzam probe can scan. Dividing the IP addresses into batches improves performance by enabling classification for each batch to begin after the batch completes, rather than after all IP addresses have been scanned. The probes run sequentially. For example, the value is set to 1000 and a discovery scans 10,000 IP addresses using a single MID Server. It creates 10 Shazzam probes with each probe scanning 1000 IP addresses. By default, the batch size is 1000. A UI policy enforces a minimum batch size of 256 because batch sizes below 256 IP addresses don't benefit from clustering. The policy converts any value below 256 to a value of zero. The value for this field can't exceed the value defined in the maximum range size property for the Shazzam probe.
Discovery Agent	Select the agent to perform this discovery: • **AWS MID**: Use AWS SSM to perform discovery. • **MID**: Use the MID Server to perform discovery. By default, the Discovery Agent is set to MID. This field is available only when the system property `glide.discovery.enable_ssm` is enabled, the **Discover** type is set to **Configuration Items**, and the **Run after** field is configured to follow a Cloud Resources schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule Form (continued)

Field	Description
Shazzam cluster support	Select the check box to distribute Shazzam processing among multiple MID Servers in a cluster and improve performance. This setting works with the Shazzam batch size. For example, a schedule is created to scan 100,000 IP addresses, with 10 MID Servers assigned to do the work. Each MID Server is assigned to scan 10,000 IP addresses. If the Shazzam batch size is set to 5,000 IP addresses per probe, the schedule runs two Shazzam probes per MID Server (10,000 IP addresses/5,000 per batch). These probes are run in sequence and not concurrently.
Use SNMP Version	Use this field to designate the SNMP version to use for this discovery. The default value is **ALL**. You can change this value to **v1**, **v2c**, or **v3**.
URL Certificate Batch Size	Define the number of URLs to discover per batch during discovery. Leave the batch size as it is unless recommended to change.
Quick ranges	Define IP addresses and address ranges to scan by entering IP addresses in multiple formats (network, range, or list) in a single, comma-delimited string. For more information, see Create a Quick IP range for a Discovery schedule.
Discover now	Use this link to immediately start this Discovery.
Related lists	
IP Ranges	This related list defines the ranges of IP addresses to scan with this schedule. If you're using a simple CI scan (no behaviors), use this related list to define the IP addresses to discover. **Note:** To improve security, limit the range of discovery targets to exclude unnecessary networks and devices.
Discovery Range Sets	This related list defines each range set in a schedule to scan by one or more Shazzam probes.
Discovery Status	This related list displays the results of current and past schedule runs.
Global IP Exclusion	This related list displays the global IP exclusion records that apply to this schedule. Active global IP exclusions prevent Discovery from scanning the specified IP addresses across all schedules that discover configuration items, IP addresses, or networks. Global IP exclusion records are managed from **Global IP Exclusion** > **IP Exclusion**. For more information, see Use Global Excludes List for IP addresses and ranges.
Discovery Attributes	This related list defines the attributes to apply to CIs discovered within this schedule. Use attributes to set CI field values such as location or asset tag at the schedule, range set, or IP address range level. For more information, see Discovery generic attributes.
Certificate URLs	This related list displays the URLs that are discovered using this schedule. You can add or delete URLs from this list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a new discovery schedule form reference

The Create a new discovery schedule form displays information related to the schedule and the target cloud provider. The Create a new discovery schedule form is available in the Cloud Discovery Workspace.

For instructions on this procedure, see Create a discovery schedule in Cloud Discovery Workspace.

Choose the provider form

Field	Description
Discovery schedule name	Name of the discovery schedule.
Cloud provider	Name of the cloud provider that hosts the resources you want to discover.
MID selection types	MID Server or MID cluster that you want to use to discover the cloud resources. The following options are available: • Auto-Select MID Server: Discovery automatically selects an appropriate MID Server. • Specific MID Cluster: Discovery uses the MID cluster specified in the MID cluster field. • Specific MID Server: Discovery uses the MID Server specified in the MID Server field. If the selected MID Server is a part of a MID cluster, the jobs are equally distributed among all the MID Servers of that cluster. This field appears only when one of the following values is selected from the cloud provider field: • AWS • Azure • Google • Alibaba Cloud
MID clusters	MID clusters you want to use with the Cloud Discovery schedule. This field appears only when the Specific MID Cluster value is selected from the MID selection types field.
Credential-less discovery	Option to use credential-less discovery. To view and filter the MID Servers that can assume an IAM role to discover the cloud resources, move the **Credential-less discovery** toggle to the right.
MID Servers	MID Servers you want to use with use with the Cloud Discovery schedule. This field appears only when the Specific MID Server value is selected from the MID selection types field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Admin Workspace reference

Reference topics provide additional information about the lists, forms, and settings used in Discovery Admin Workspace.

Discovery Admin Workspace support tools

The **Support tools** tab provides tools to investigate and troubleshoot your instance.

Available tools

Tool	Description
CAPI to patterns	Migrates your Cloud Discovery from Cloud API REST Calls (CAPI) to Patterns.
Probes to pattern migration	Migrates older probes to new, flexible, and efficient patterns.
Retrieve logs	Retrieves diagnostic logs to help identify ITOM Visibility and Discovery issues.
Detect Frequently discovered CIs	Displays which CIs are discovered more often than others.
Pattern to customization check	Determines if a customization might be breaking a pattern.
IRE payload explorer	Identifies the underlying key component in identification and reconciliation (IRE) across different data sources.
Discovery customization report	Detects probe, sensor, classifier, and pattern code that have been modified from base system settings.
Skipped files from store updates	Examines all file names from pattern plugins that were skipped from store updates.
Command validation task	Runs and validates commands on a range of IP addresses.
Missing IPs in schedules	Identifies which Discovery schedules are missing IPs and adds them to the schedule.
Missing SNMP OID data	Checks for missing SNMP device OIDs from the SNMP data that cause the devices to be misclassified.
SaCmdManager	Searches for information that might be needed when creating Discovery patterns.

Application fingerprint suggestion form

When viewing suggestions based on application fingerprints, verify the information on the Application fingerprints suggestion form.

Application fingerprint suggestion form data

Field	Description
Discovery Pattern name	The name of the pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application fingerprint suggestion form data (continued)

Field	Description
Suggested CMDB CI Class	The name of the CMDB CI class. This field is pre-populated, but as the administrator you can modify it if necessary.
Suggested Classifier Rule	The name of the process excluded from discovery using the Ignore action.
Accuracy Percent	A percentage, based on a validation algorithm, that checks the quality of the prediction. Pick suggestions with higher accuracy. You can't modify this field.
Suggested Group Name	The name ITOM Visibility generates for the suggestion. This field is pre-populated, but you can modify it if necessary.
Process Count	The number of discovered processes that this suggestion is based on.
Extended Group Name	The top 10 key words that provide more information about the process. Predictive Intelligence extracts these key words from the process. This field is pre-populated and you can't modify it.

Schedule details data visualizations

When you select a schedule from the Schedules page in Discovery Admin Workspace, the schedule details page opens. The **Overview** tab of that page displays data visualizations showing key metrics for the schedule, such as total runs, CI counts, errors, and run duration.

Note: For certificate discovery schedules, the Total CIs and CIs discovered count visualizations are replaced by the Total Certificates Discovered and Certificates discovered count visualizations, respectively. Certificate discovery is available in Discovery Admin Workspace starting with v1.20.0 and requires the Brazil, Australia, or Zurich release starting with Patch 8.

Overview tab data visualizations

Report title	Type	Description
Total runs	Indicator	Total number of times this Discovery schedule ran in a certain time period.
Timeline of events during schedule run	Timeline	Events that occurred during the schedule runs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview tab data visualizations (continued)

Report title	Type	Description
Total CIs	Indicator	Number of CIs detected on the last run.
Total Certificates Discovered	Indicator	Number of certificates detected on the last run.
CIs discovered count	Line chart	Trends of discovered CI attributes for this Discovery schedule across the statuses listed in Run History. The anomaly threshold trend reflects the most severe anomaly for this schedule. **Note:** The anomaly threshold only displays when anomaly detection is enabled. For more information, see Discovery Admin Workspace Settings.
Certificates discovered count	Line chart	Trends of discovered certificates for this Discovery schedule across the statuses listed in Run History.
Errors	Indicator	Total number of errors detected on the last run.
Errors count	Line chart	Error trends for this Discovery schedule across the statuses listed in Run History. The anomaly threshold trend reflects the most severe anomaly for this schedule. **Important:** The anomaly threshold only displays when anomaly detection is enabled. For more information, see Discovery Admin Workspace Settings.
Duration in mins	Indicator	Duration in minutes it took for the last Discovery to run.
Discovery run time	Line chart	Run time trends for this Discovery schedule across the statuses listed in Run History. The anomaly threshold trend reflects the most severe anomaly for this schedule. **Important:** The anomaly threshold only displays when anomaly detection is enabled. For more information, see Discovery Admin Workspace Settings.

Queue Registration form reference

The Queue Registration form displays the fields available when creating or modifying a queue.

Queue Registration form

Field	Description
Queue	Name of the queue.
Application	Scope of the application.
Automatic Job Scheduling	Option to enable automatic job scheduling. When selected, the queue switches to

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Queue Registration form (continued)

Field	Description
	automatic job scheduling and the system handles all job configurations.
Event Processing Order	Method used to process events. Options include: • Parallel: Process multiple events simultaneously. • Sequential: Process only one event at a time. The events are interdependent and can only process after the previous event completes. **Note:** After you configure the processing order and submit the queue, it can't be edited. Instead, you must delete the queue and create another queue.
Job configuration type	Type of job configuration. Options include: • Constant number: Creates the required number of jobs. • Scale with node: Creates multiple jobs based on the number of nodes. Select this option if you don't know the number of available nodes. **Note:** This field is automatically set to Constant number when the Event Processing Order field is set to Sequential.
Scale factor	Total number of jobs, which is calculated based on the value of the **Job configuration type** field.
Poll interval	How frequently jobs should poll and process events in this queue.
Description	Description of the queue.

Quick start test for IP Discovery

Quick start test for IP discovery validate configuration items (CI) and their relationships.

IP Discovery quick start test

Test	Description
TestIPDiscovery	IP discovery is performed on an IP input by the user. After discovery, the CI and relationships are created in the cmdb_ci and rel_ci tables. This test validates the CIs and relationships

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Discovery quick start test (continued)

Test	Description
	created and output and discovery errors during the discovery run.

To learn more about IP Discovery, see Schedule a horizontal discovery.

Related topics

Quick start tests

Tracked Configuration file form

The Tracked Configuration file form displays details about a specific configuration file.

Tracked Configuration file fields

Field	Description
File Name	Name and file path of the file.
File Size (B)	File size in bytes.
Last Modified	Date the file was last modified.
File Content	Contents of the file, if you selected the option to save the contents on the file tracking definition.

Related topics

Compare versions of CI configuration files

Firewall Audits and Reporting

Use the Firewall Audits and Reporting application to explore and create an inventory of your firewall security policies, devices, device groups, and manager information. Additionally, you can initiate requests for new firewall rules through the Service Catalog application and conduct audits of firewall security policies within specified time frames.

Get started

Explore Firewall Audits and Reporting	Manage Firewall Audits and Reporting	Firewall Audits and Reporting reference
Leverage Firewall Audits and Reporting to create inventory	Learn how to set up Firewall inventory	Familiarize yourself with dashboards and report calibration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Troubleshoot and get help

- What is firewall audit?
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Firewall Audits and Reporting

With Firewall Audits and Reporting, you have the capability to explore and conduct an inventory of your firewall security policies, devices, device groups, and manager information. Additionally, you can utilize it to submit requests for new firewall rules and audit security policies within a specified time frame.

Firewall Audits and Reporting overview

A firewall device acts as a network security system, monitoring and controlling traffic based on specific policies. It establishes a protective barrier between trusted internal and untrusted external networks, incorporating multiple security policies to safeguard against threats. Ongoing maintenance and audits are crucial as security policies evolve, preventing potential loopholes. The firewall audit process streamlines rule tracking and updating, ensuring alignment with company security policies. Firewall vendors, including Palo Alto Panorama and Fortinet FortiManager, provide centralized managers for efficient control over devices and policies.

Firewall Audits and Reporting workflow

The Firewall Audits and Reporting application enables the ServiceNow Discovery process to discover firewalls from supported vendors (Palo Alto Panorama and Fortinet FortiManager). It discovers CMDB CIs for the firewall devices, firewall manager, firewall device groups, and firewall policies using serverless patterns. Firewall policy audit tasks are generated from the firewall managers or devices. You can also request new firewall security policies through the Service Catalog and archive older firewall rule requests, audit requests, and audit tasks to improve system performance. For more information, see Visibility to Firewall inventory. To see various reports to track discovered policies and audit tasks, see Firewall Admin Workspace dashboard.

Firewall Audits and Reporting benefits

Benefit	Feature	Users
Up-to-date inventory of firewall security policies, devices, device groups, and manager information by persistently storing data in the CMDB through regular queries to a firewall manager.	Visibility to Firewall inventory	Firewall Admin [sn_disco_firewall.firewall_admin]
Provision for requesting a new firewall rule using the Service Catalog.	Firewall rule requests	Firewall Requester [sn_disco_firewall.firewall_requester]
Provision for auditing firewall security policies for a specific time period.	Firewall Admin Workspace dashboard	Firewall User [sn_disco_firewall.firewall_user]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Firewall Audits and Reporting roles and responsibilities

Specific users and roles to optimize the monitoring and tracking of requests for new firewall rules.

The configuration settings are activated or deactivated depending on the assigned roles.

Role	Responsibilities
Firewall Admin [sn_disco_firewall.firewall_admin]	Responsible for editing Configuration Item (CI) attributes not populated by Discovery, including "assigned_to" and others. Firewall Admins possess the capability to view the Firewall managers list with specific details for each manager. Furthermore, they have the authority to view and edit the dashboard, initiate audits from the firewall manager details page, create, modify, and view policies, edit audit tasks, and request policy changes. **Note:** The Firewall Admin role includes the Firewall User role. This role is not automatically granted with the discovery_admin role and needs to be explicitly assigned to the user.
Firewall User [sn_disco_firewall.firewall_user]	Responsible for viewing firewall dashboards and policies in a Read-only capacity. **Note:** This role includes the Firewall Requester role.
Firewall Requester [sn_disco_firewall.firewall_requester]	Responsible for requesting new firewall rules through the Service Catalog.

AI in Firewall Audits and Reporting

AI provides an agentic workflow in Firewall Audits and Reporting that helps firewall administrators request new firewall rules using natural language.

Firewall Audits and Reporting includes an AI-powered agentic workflow. This feature requires ServiceNow Otto for IT Operations Management (ITOM) to be installed.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI features in Firewall Audits and Reporting

AI feature	Description	Use case	Resources
Firewall Management Task Creation agentic workflow	Uses AI to accept natural-language firewall rule requests, extract required parameters, run a risk analysis, and create a rule task. The risk level determines whether the task is created automatically or requires confirmation.	A firewall administrator wants to request a new firewall rule by describing it in plain language rather than filling out a form manually.	• Firewall rule requests using agentic workflows • Request firewall rules using agentic workflow

Configuring Firewall Audits and Reporting

Set up the process to gain insight into discovering and compiling an inventory of firewall security policies, devices, device groups, and manager information.

Configuration overview

Before using Firewall Audits and Reporting, it is crucial to set up the necessary plugins, ensure activation, upgrade your instance to the latest version, and download the Firewall Audits and Reporting application from ServiceNow Store. For additional details, see Get started with Firewall Audits and Reporting.

After the prerequisites are met, you can proceed with the following tasks.

Firewall Audits and Reporting tasks

Task	Description
Discover firewall policies	Discover firewall policies that includes uncovering firewall devices, policies, and owner groups, ensuring a centralized view of the footprint.
Customize Firewall Audits and Reporting	Customize Firewall Audits and Reporting by modifying Discovery properties to tailor specific aspects.
Archive firewall rule audit, rule requests, and audit tasks	Enhance system performance by archiving firewall rule requests, audit requests, and audit tasks older than a designated time period.
Initiate audit request	Initiate audits against a specified firewall manager or device, ensuring adherence to their organization's security policies.

Get started with Firewall Audits and Reporting

Prior to diving into Firewall Audits and Reporting functionality, meet the necessary requirements by installing and activating the plugin, upgrading your instance, and obtaining the application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: admin

About this task

Note: The ServiceNow Store regularly releases new applications and updates to ServiceNow applications. If you already have an application, you can download the latest version to enhance your existing experience with ServiceNow products. Features vary by release. The version number indicates which content and features are available in each release.

Procedure

1. Ensure the following plugins are installed and activated.
   - ITOM Visibility [com.snc.itom.vis.license] plugin
   - Discovery [com.snc.discovery] plugin
   - Configuration Management for Scoped Apps (CMDB) [com.snc.cmdb.scoped] plugin
2. Confirm that your system has been upgraded to Brazil or a later version.
   During the upgrade process, the Firewall Audits and Reporting [com.sn_disco_firewall] plugin is automatically installed if the ITOM Visibility [com.snc.itom.vis.license] and Discovery [com.snc.discovery] plugins are already installed.
3. Download the Firewall Audits and Reporting application from ServiceNow Store.

Visibility to Firewall inventory

The Firewall Audits and Reporting application helps you to discover and take inventory of your firewall security policies, devices, device groups, and manager information.

The ServiceNow Platform helps organizations gain visibility to IT estate with the centralized Configuration Management Database (CMDB) while also automating workflows for the request and change management processes. The Firewall Audits and Reporting application enables security teams to discover all firewalls along with their policies, firmware versions, and other hardware attributes. The Service Portal from ITSM helps customers to digitize workflows when requesting new firewall rules. Combined with ITSM incident management and change management workflows, this solution enables transparency in the fulfillment process. To make it more simple for teams, policy ownership is tied to the system user and groups in the ServiceNow Platform. A centralized dashboard provides visibility to the firewall estate and comprehensive visibility to the inflow of request records landing in the Firewall admin's fulfillment queue.

Currently, ServiceNow works with Palo Alto Networks Panorama. In future release, other firewall providers will be added. The integration uses XML APIs and SNMP to connect the Firewall Audits and Reporting application to firewall inventory and firewall policies.

Discover firewall policies

As a member of a security team, you can discover firewall devices, policies, and owner groups from supported vendors (Palo Alto Panorama and Fortinet FortiManager), allowing a central view of the footprint. This data is updated in the ServiceNow CMDB. Set up a schedule to discover your firewall policies to help you keep track of your company's valuable information.

Before you begin
Role required: discovery_admin, firewall_admin

About this task
Administrators in charge of Discovery can establish a recurring schedule for firewall policy discovery. This schedule utilizes the serverless pattern, connecting with the firewall manager to discover and update information for Configuration Items (CIs) in the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For Fortinet FortiManager-specific requirements, including MID Server sizing and API endpoint information, see Fortinet FortiManager discovery requirements.

Firewall Configuration Items and tables by vendor

Vendor	CMDB Configuration Items	Non-CMDB Tables
Palo Alto Panorama	• Panorama Firewall Manager [cmdb_ci_firewall_manager_panorama] • Palo Alto Firewall Devices [cmdb_ci_firewall_device_palo_alto] • Panorama Firewall Device Group [cmdb_ci_firewall_device_group_panorama] • Panorama Firewall Security Policies [cmdb_ci_firewall_sec_policy_panorama]	• Panorama Address Objects • Panorama Address Group Objects • Panorama Service Objects • Panorama Service Group Objects • Policy Object M2M References
Fortinet FortiManager	• FortiManager Network Manager [cmdb_ci_fortimanager_network_manager] • Fortinet Firewall ADOM [cmdb_ci_firewall_device_group_fortinet] • Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet] • Fortinet Firewall Policy Package [cmdb_ci_fortinet_firewall_policy_pkg] • Fortinet Firewall Policy [cmdb_ci_fortinet_firewall_policy]	• Fortinet Address Objects [sn_disco_firewall_fortinet_address_object] • Fortinet Address Group Objects [sn_disco_firewall_fortinet_addrgrp_object] • Fortinet Service Objects [sn_disco_firewall_fortinet_service_object] • Fortinet Service Group Objects [sn_disco_firewall_fortinet_servicegrp_object] • Fortinet VIP Objects [sn_disco_firewall_fortinet_vip_object] • Fortinet VIP Group Objects [sn_disco_firewall_fortinet_vipgrp_object] • Fortinet ISDB Objects [sn_disco_firewall_fortinet_isdb_object] • Firewall Policy Object M2M Mappings [sn_disco_firewall_policy_object_m2m]

Procedure

1. Create a new credential of type **API Key Credentials**.
   Enter the API key for your firewall vendor in the credential record.
2. Create a credential alias for the API key credential created in the previous step.
   You provide this credential alias in the discovery schedule configuration. For more information, see Credential aliases for Discovery.
3. To create a Discovery schedule, perform the following steps.
   a. Select **Discover: Serverless**.
   b. Select the appropriate **MID Server**.
   c. Right-click the header and select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information on Discovery schedule, see Schedule a horizontal discovery.

4. From the tab at the bottom of the screen, select the Serverless Execution pattern and then select **New**.
5. In the Serverless Execution pattern, perform the following steps.
   a. Enter a name.
   b. Select the pattern for your vendor:
      - For Palo Alto Panorama: Select **PaloAlto - Firewall Manager**.
      - For Fortinet FortiManager: Select **Fortinet FortiManager (LP)**.
   c. Select **Run Child Patterns**.
   d. Select **Submit**.
6. Navigate to **Discovery Pattern Launcher Parameters** and set the following parameters.
   - credentialAlias: Provide the credential alias name created in step 2.
   - For Palo Alto Panorama:
     - url: Set to the base URL of your Panorama device (https://<PANORAMA_HOST>/api).
     - trustInsecureHosts: Set to **true** to turn off hostname verification and enable self-signed certificates to be accepted as trusted.
   - For Fortinet FortiManager: Set the ipAddress parameter to the IP address of your FortiManager appliance.
7. Right-click the header and select **Save**.

Customize Firewall Audits and Reporting

Configure default policy owner and approval groups for firewall rule change requests to reduce manual data entry.

Before you begin

Role required: discovery_admin or admin

About this task

When users submit firewall rule change requests, each request must be assigned to a policy owner group and an approval group. Without default values, administrators must manually populate these fields for every request. Configuring default groups streamlines the request process by automatically populating these fields.

Set these defaults during initial Firewall Audits and Reporting setup or when your organization's security team structure changes. For more information on properties, see Discovery properties.

Procedure

1. Navigate to **All** > **Discovery Definition** > **Properties**.
2. Configure the firewall rule change management groups.
   a. In the sn_disco_firewall.default.rule.task.policy.owner.group property, enter the name of the group responsible for implementing firewall rule changes.
   b. In the sn_disco_firewall.default.rule.task.approval.group property, enter the name of the group responsible for approving firewall rule changes.

   These groups are automatically assigned as the policy owners for new firewall rule requests.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

When users submit new firewall rule change requests, the policy owner group and approval group fields automatically populate with the configured default values.

Archive firewall rule audit, rule requests, and audit tasks

Archive firewall rule requests, audit requests, and audit tasks that are older than a specific time period to enhance system performance. At a later time, you can delete them from the archive table altogether to reduce the size of that table.

Before you begin

Ensure that the Data Archiver [com.glide.auxdb] plugin is activated and enabled.

Role required: firewall_admin

Procedure

1. Activate one or more of the **Firewall Archival Rules**.
   The following archival rules come with the Firewall Audits and Reporting application, but are inactive by default. You must activate any rules if you want to use them.
   - Archive Firewall Audit Request
   - Archival Firewall Audit Request Security Policy M2M
   - Archive Firewall Audit Task
   - Archival Firewall Audit Task Security Policy M2M
   - Archive Firewall Rule Task
   - Archival Firewall Rule Task Security Policy M2M
2. Navigate to **Archive Rules** > **Select a Firewall Archive rule**.
3. Select the **Active** check box.
   The Data Archiver [com.glide.auxdb] plugin moves data that is no longer needed from the primary tables to a set of archive tables daily. It checks for and archives requests and tasks based on the archival rule activated and chosen. The rule archives data from the primary table when the following two conditions are met:
   - State is Closed Complete
   - Update before Last 12 months

Firewall audits

Firewall Audits and Reporting provides a flexible auditing framework to track your firewall policy ownership and the necessity of a firewall policy. You can trigger a random audit to measure your security hygiene on the firewall policy and ownership, as well as perform proactive audits on a regular basis.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Firewall audit request workflow

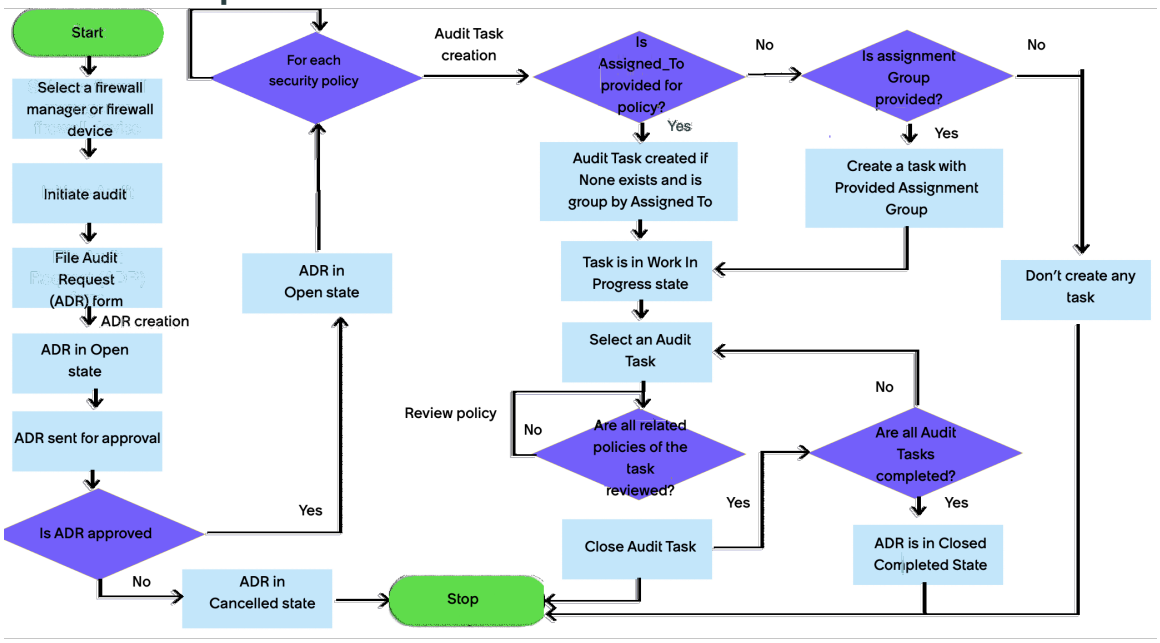

Initiate audit request

Initiate audits against a specified firewall manager or device to ensure proper configuration in alignment with the security policies of your organization.

Before you begin

Role required: firewall_admin

Procedure

1. To list all of the discovered firewall managers in your network, navigate to **All** > **Firewall Audits and Reporting** > **Records** > **Devices or Managers**.
2. Select a firewall manager CI record or device record from the list.
3. Select **Initiate Audit Request**.
4. Fill in the mandatory fields on the form.
5. Select **Submit**.

What to do next

The procedure initiates an audit for the firewall manager, validates it, and submits the audit task. Firewall audit tasks are generated based on the assigned user for the firewall security policy. If the policy lacks a designated user, no audit task is created for that security policy, and it is placed in the *Excluded Policy List* for audit requests.

This requests an audit for the firewall manager, validates, and submits. Firewall audit task are created on the basis of an assigned user for the firewall security policy. If the policy is not assigned to a user, then no audit task is created for that security policy and it is added into the *Excluded Policy List* of audit requests.

- If *Assigned To* is provided, the audit task is created and grouped by the assigned user.
- If *Assigned To* is not provided, and the *sn_disco_firewall.default.rule.task.policy.owner.group* discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

property is not set, the policy is added to *Excluded Policies* list, and no audit task is generated.

- If *Assigned To* is not provided, but the *sn_disco_firewall.default.rule.task.policy.owner.group* discovery property is configured, the audit task is created and grouped by the **Assignment Group**.

To verify if the audit task was requested, navigate to **Audits** > **Audits Request**. You should see your task in the list. Creation of audit tasks is a background job using sub-flows. You may need to wait for audit tasks to get created depending on the number of policies.

The *Assigned To* person audits each policy in the related list. They choose the action for each policy and mark the task as *Close Complete*.

Using Firewall Audits and Reporting

Firewall Audits and Reporting is a powerful application for requesting new firewall rules, streamlining the management of IP addresses.

Firewall rule requests using agentic workflows

The Firewall Management Task Creation agentic workflow provides a path to request one or more firewall rules through natural language prompts in the ServiceNow Otto panel.

Multiple rule configuration support

You can request multiple firewall rule configurations in a single conversation with the agentic workflow. The workflow parses each rule individually, evaluates conformance for each one, and creates one parent firewall task with individual configuration tasks for all rules. This approach reduces the number of tasks to track and simplifies the approval process.

Risk analysis and conformance checking

The workflow reads your natural-language request, extracts the required parameters—source IP, destination IP, protocol, traffic direction, action, and conformance type—and prompts you for any values you did not provide. Before creating a rule task, the workflow runs a risk analysis based on the request data and the specified conformance framework. When you request multiple rules, the workflow provides individual conformance feedback for each rule, such as "Out of 3 rules, 2 are non-conforming, 1 is conforming." The risk analysis returns one of three levels:

- **Low:** The workflow creates the rule task and attaches the risk analysis.
- **Medium or High:** The workflow reports the risk level in the ServiceNow Otto panel and asks whether you want to continue. You can choose to proceed with all rules or select specific ones. If you confirm, the workflow creates the rule task and attaches the risk analysis. If you decline, the workflow skips task creation. The created task includes the AI assessment so the approver can evaluate the risk.

Note: The agentic workflow provides conformance checking and risk analysis. Rule verification and creation on firewall vendor servers is handled through API calls as part of the Firewall Audits and Reporting application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Workflow process

Firewall rule request workflow

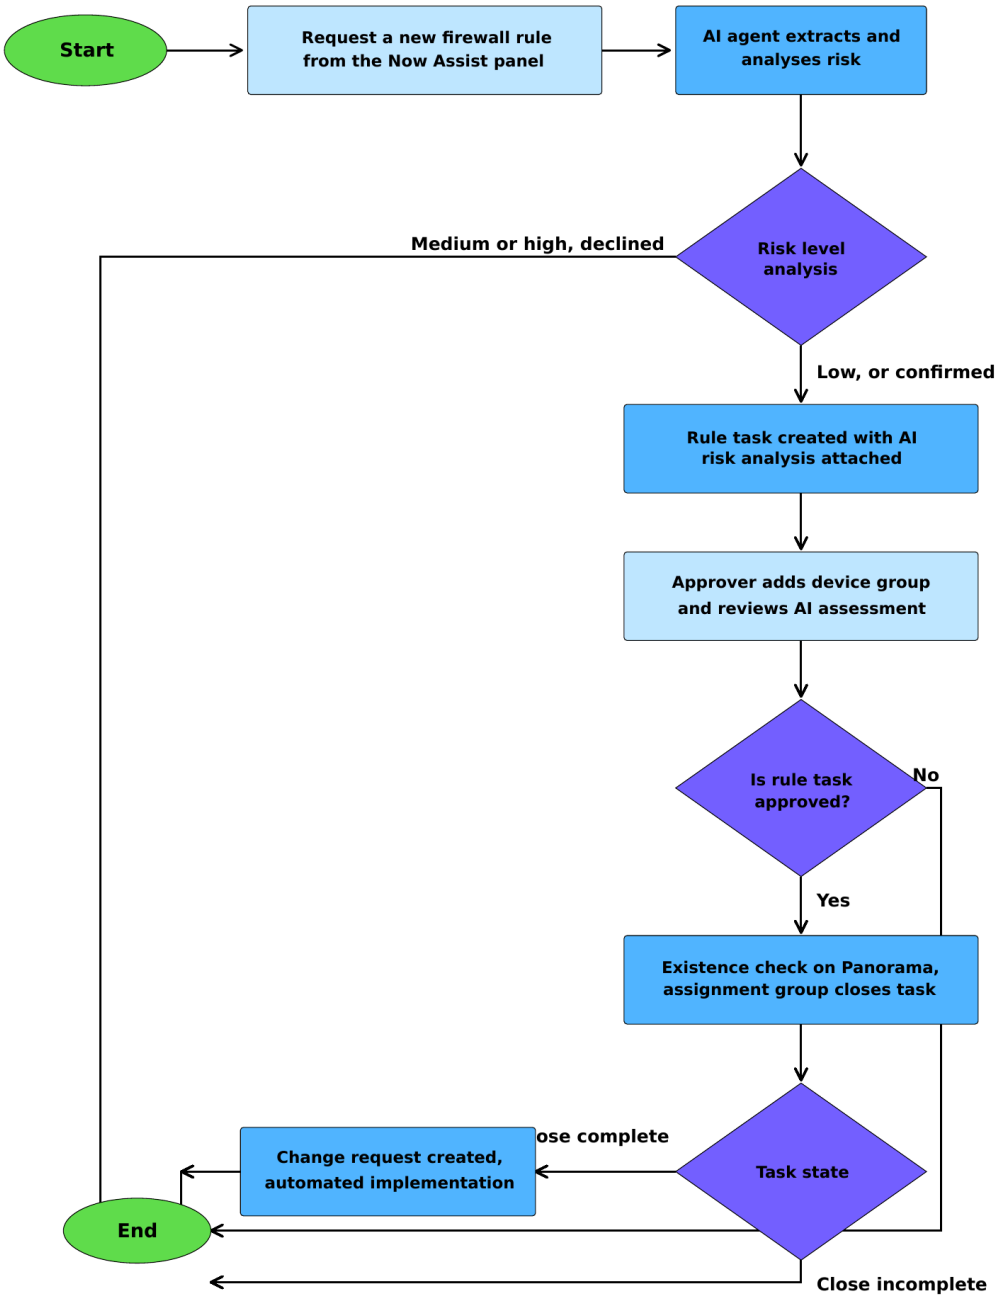

Request firewall rules using agentic workflow

Use the Firewall Management Task Creation agentic workflow to request one or more firewall rules through natural language prompts in the ServiceNow Otto panel.

Before you begin

- Install and configure the Firewall Audits and Reporting application.
- Install the AI Agents for Discovery plugin. This plugin is part of the AI Agent Bundle and requires a separate subscription.
- Discover the firewall managers and devices for your vendor, and verify that Discovery has populated the appropriate vendor-specific tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: firewall_admin

About this task

You can request multiple firewall rule configurations in a single conversation. The workflow creates one parent firewall task with individual configuration tasks for each rule. For information about how the workflow evaluates risk and determines whether to create a rule task, see Firewall rule requests using agentic workflows.

Note: Starting with version 1.12.0 of Firewall Audits and Reporting, all new firewall rule tasks are created in vendor-specific tables and display the task type based on the selected vendor.

ServiceNow Otto agentic workflow

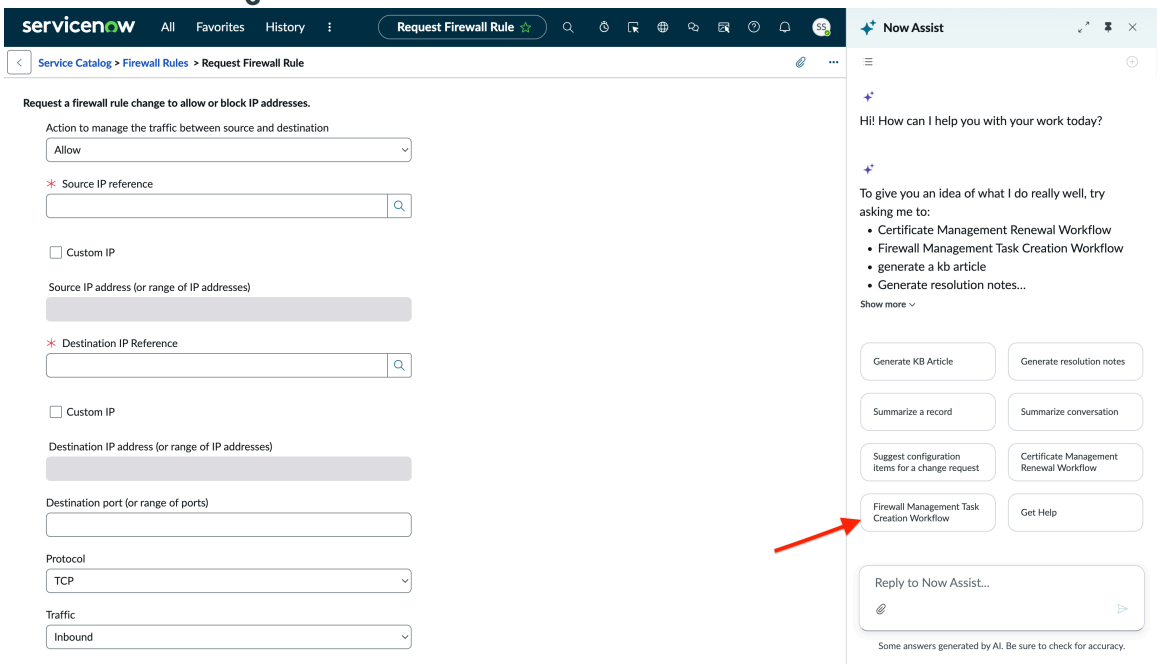

Procedure

1. Open the ServiceNow Otto panel.
   The ServiceNow Otto panel is available from the header bar of your instance. Select the ServiceNow Otto icon to open the chat window.
2. In the ServiceNow Otto chat, type `Create a firewall rule` or a similar phrase to trigger the Firewall Management Task Creation agentic workflow.
3. Enter your firewall rule request in natural language, including the source IP, destination IP, protocol, traffic direction, and action.
   You can specify multiple rule configurations in the same conversation. The workflow parses each rule individually and treats them as separate configuration tasks under one parent firewall task.
4. Respond to the workflow prompts for any missing information.
5. Review the workflow response in the ServiceNow Otto chat.
   The response includes the rule task number, the extracted request data, and the risk-analysis summary for each rule configuration. When you request multiple rules, the workflow provides conformance feedback for each rule individually, such as "Out of 3 rules, 2 are non-conforming, 1 is conforming." If any rule has a medium or high risk level, the workflow asks whether you want to continue. You can choose to proceed with all rules or select specific ones. Select **Yes** to create the rule task or **No** to cancel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. **Optional:** Verify the rule task by navigating to **Rule Requests** > **Rule Requests Task**.
   Your request appears in the list with the task type set to the selected vendor and the AI risk analysis attached as work notes. Open the parent task to view all individual rule configuration tasks.

Approve firewall rule requests

Review AI-generated firewall rule requests, evaluate risk analysis, and approve or reject requests with device group assignment.

Before you begin

Role required: none. Approval access is granted to members of the approval group specified on the rule task. The admin user can edit the approver list on the rule task.

About this task

A firewall rule task can contain multiple rule configurations, each represented as an individual configuration task under the parent task. Requests created from the agentic workflow include an AI assessment that summarizes each rule request and indicates whether the request is good to approve. Approvers can use this assessment instead of manually evaluating each parameter. Before approval, the workflow posts a chat message indicating that the device group can't be automatically determined. A device group is a logical bundle of devices on which the rule is created. The approver, who is familiar with the firewall infrastructure, must specify the device group on which to apply the rule.

Procedure

1. Navigate to **All** > **Self Service** > **My Approvals**.
2. Open the parent rule task and review the individual rule configuration tasks.
   Each configuration task includes the AI assessment in the work notes. The status **Good to Approve** indicates that the AI risk analysis returned low risk for that specific rule.
3. In the chat prompt, specify the device group on which to create the rule.
   For example, `test1`.
4. Approve or reject each configuration task.
   - Select the green checkmark to approve the configuration task. The workflow calls the Panorama REST API to check whether the rule already exists. If the rule exists, the configuration task is marked as Closed Incomplete with a work note indicating the rule was not recreated. If the rule doesn't exist, the workflow proceeds to implementation.

     Note: A background subflow creates a change request after all configuration tasks are processed and the parent rule task reaches the **Close Complete** state. The change request is created only if the change management plugin is activated.

   - Select the red X to reject the configuration task. The request is returned to the requester for revision.

Result

If the configuration tasks are approved:

- The Firewall Audits and Reporting application verifies whether each requested rule already exists on Panorama.
- For rules that don't exist, the assignment group works on the request and marks the configuration tasks **Close Complete**.
- A change request is created with an implementation plan that contains the source IP, destination IP, traffic flow, action, port, protocol, and device group for all approved rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Implement firewall rules on Panorama

Automatically implement approved firewall rules on Panorama servers through the change management process.

Before you begin

Role required: none. Implementation access is granted to members of the assignment group on the change request.

About this task

When a firewall rule task contains multiple rule configurations, each configuration is validated individually as a subtask. If a rule already exists on the Panorama server, the configuration task is marked as Closed Incomplete with a work note indicating the rule was not recreated. If a rule does not exist, the configuration task proceeds to the approval and implementation process.

The instance calls a REST API to create and commit the rule, so the assignment group does not have to log in to Panorama manually. When the request is created from the Service Catalog instead of the agentic workflow, the implementation step is not automated. In that case, a member of the assignment group must log in to Panorama, create the rule manually, and mark the change request as **Review**.

Procedure

1. Navigate to the change request created from the parent rule task.
   The change request includes all approved rule configurations from the parent task.
2. Progress the change request through the standard change management states to **Implement**.
   Update the **State** field on the change request form in sequence: **Assess**, **Authorize**, **Scheduled**, then **Implement**. Each state may require review and approval by the relevant stakeholders before you can advance to the next.
   When the state reaches **Implement**, the workflow calls the Panorama REST API to create and commit each approved rule on the device group specified during approval.
3. Review the rules on the change request.
   The work notes show whether each rule was created and committed successfully, along with the rule names. For rules that already existed on the Panorama server, the work notes indicate that the rule was not recreated.
4. If all rules were processed as expected, close the change request.

Result

All approved rules are created and committed on the Panorama server for the specified device group. The work notes on the change request confirm the result for each rule and include the rule names.

Note: Create saves the rule on the Panorama server without applying it to devices. Commit applies the rule to all relevant devices. The automation performs both actions in a single API call for each rule.

Firewall rule requests

Use Service Catalog to request new firewall policies and rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Firewall rule request workflow

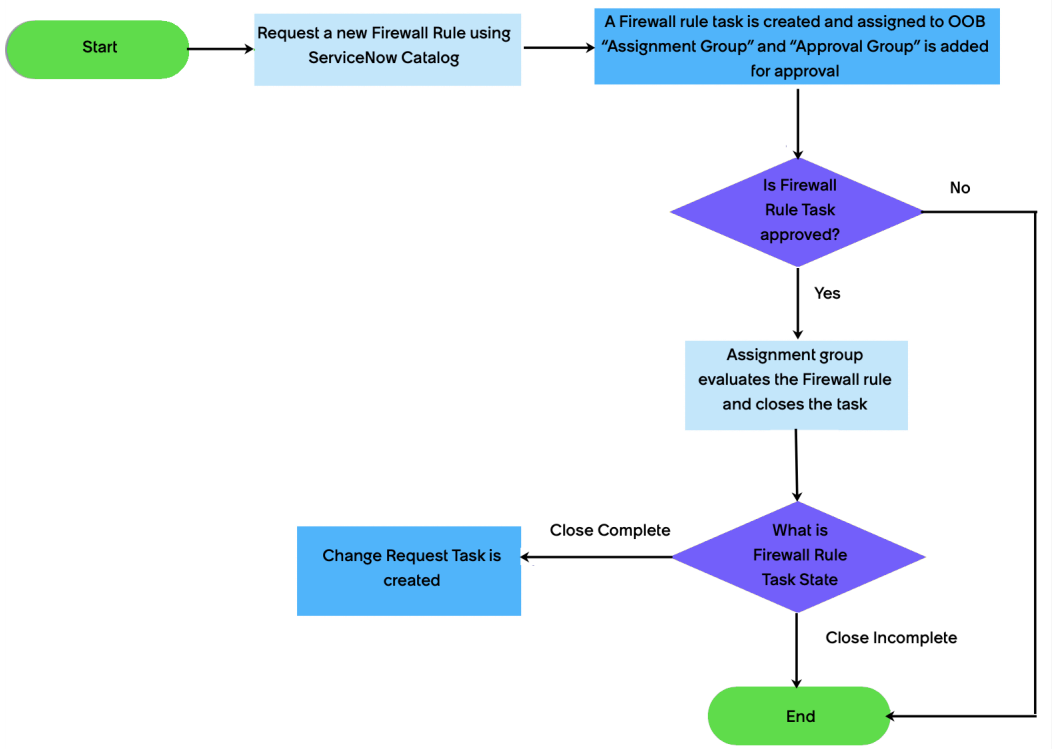

Request a firewall rule

Request one or more firewall rules using Service Catalog to manage various IP addresses and enhance network security and accommodate evolving business requirements.

Before you begin

Verify that the Firewall Audits and Reporting catalog is enabled.

Role required: firewall_admin

About this task

You can request multiple firewall rule configurations in a single request. The system creates one parent firewall task with individual configuration tasks for each rule. Administrators initiate tasks, which are automatically directed to the risk team for assessment and approval. Following approval, firewall admins smoothly implement changes, all orchestrated through automated workflows.

Note: Starting with version 1.12.0 of Firewall Audits and Reporting, all new firewall rule tasks are created in vendor-specific tables and display the task type based on the selected vendor. Open requests created before version 1.12.0 are read-only. To proceed with those requests, resubmit them using the current catalog form.

Procedure

1. Navigate to **All** > **Service Catalog** > **Firewall Rules**.
2. Select **Request Firewall Rule**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

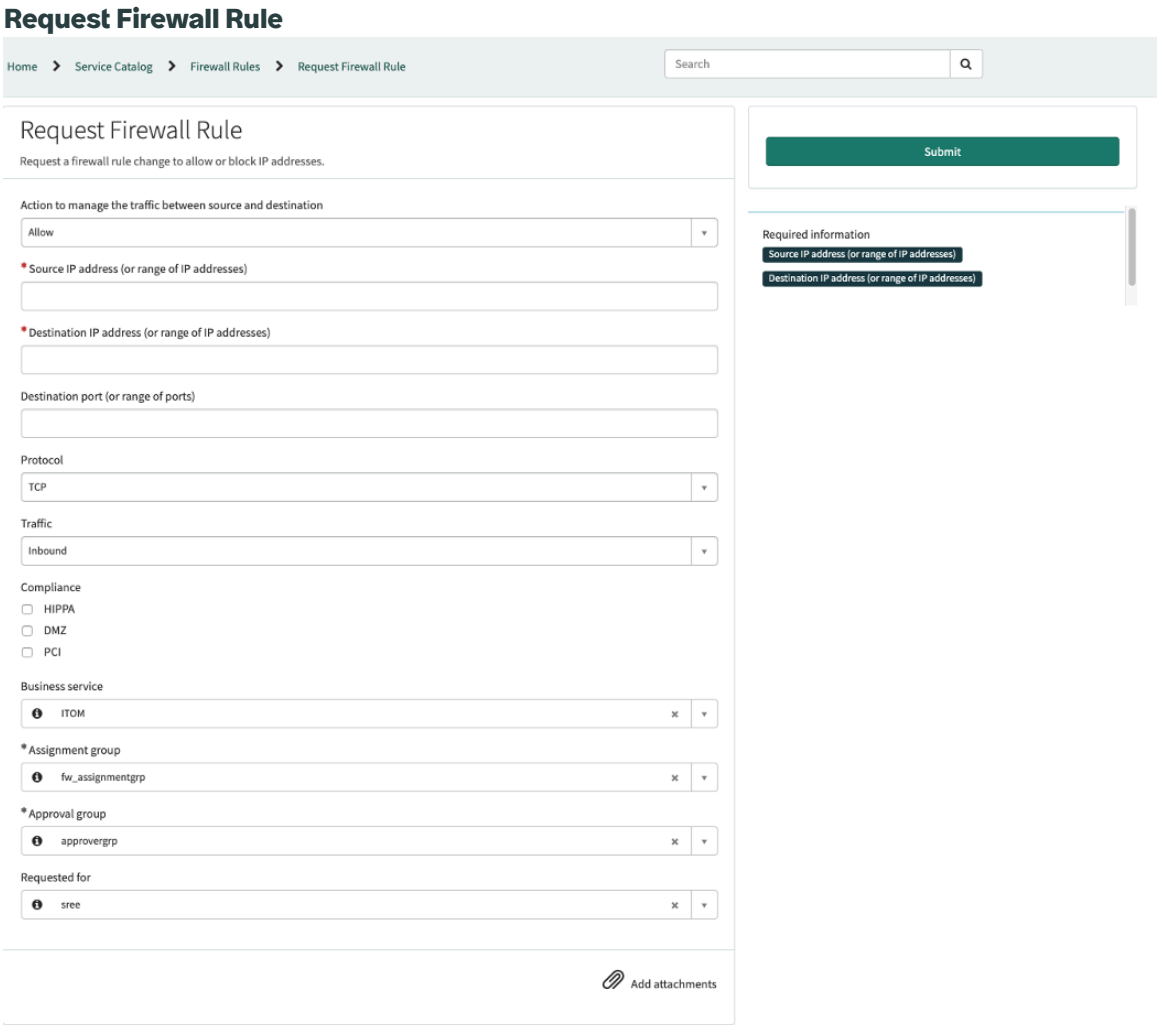

3. Enter the appropriate information for the following mandatory fields.

   - Source IP address
   - Destination IP address
   - Assignment Group

     Must have the *sn_disco_firewall.firewall_user* role.

   - Approval Group

     Must have the *approver_user* role.

4. Enter or select any details that is required.
5. **Optional:** To add additional rule configurations, select **Add Rule Config** and enter the details for each additional rule.
   Each rule configuration you add creates a separate configuration task under the same parent firewall task. All configurations share the common fields such as Assignment Group and Approval Group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Select **Submit**.
   The system creates one parent firewall rule task with individual configuration tasks for each rule you specified.

What to do next

To verify the new rule task, navigate to **Rule Requests** > **Rule Requests Task**. Your request appears in the list with the task type set to the selected vendor. Open the parent task to view all individual rule configuration tasks.

Approve firewall requests

Approval of firewall requests gives you controlled access and compliance. Members of the approver group can review and approve firewall audits and new firewall requests.

Before you begin

Role required: Members of the specified approver group **approval_group** specified in the rule task. The admin user can edit the approvers list in the Rule Request Task.

Procedure

1. Navigate to **All** > **Self Service** > **My Approvals**.
2. Select the green checkmark to approve.

Result

- The Assignment group works on the request and marks it as *Close Complete*.
- Once the assignment_group marks the request *Close Complete*, if the change request plugin is activated, a background sub-flow creates a change request.

  Note: The change request is created only if the rule task is *Approved* and in *Close Complete* state.

The Firewall rule task security policy M2M corresponds to the related list *Security policies* in *Rule task*. Firewall administrators can add *description* or *tag* fields in a security policy on a Panorama device. They can also add firewall rule task numbers or change request numbers while creating or modifying security policies on Panorama. When the next discovery runs, the M2M table populates the mapping between:

- Firewall rule task and firewall security policy
- Firewall security policy and business service if the business service is provided during the Firewall rule task request

Firewall Audits and Reporting reference

Reference topics provide additional information about Firewall Audits and Reporting components, including dashboards and report calibration.

Firewall Admin Workspace dashboard

The unified dashboard enables security and risk teams to oversee tasks, change requests, and the entire firewall inventory life cycle.

Firewall Insights tab

The **Insights** tab offers a snapshot of firewall policies and associated devices, showcasing trends in policies without owners over time. Track new and unaudited policies, and view the distribution of device manufacturers and models within the network.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

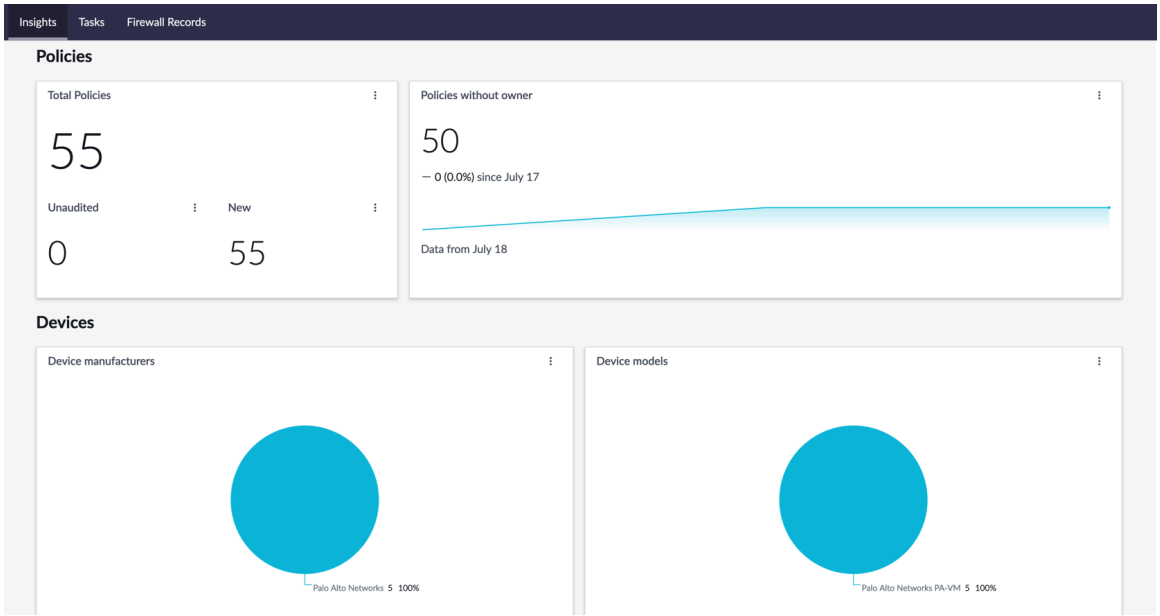

Requests and Audits tab

Within the **Tasks** tab, explore the **Requests** and **Audits** sub-tabs. The **Requests** tab provides breakdowns of request tasks by state and ownership, along with fulfillment trends over time. Similarly, the **Audits** tab displays outstanding audits by state and ownership, along with response trends over time.

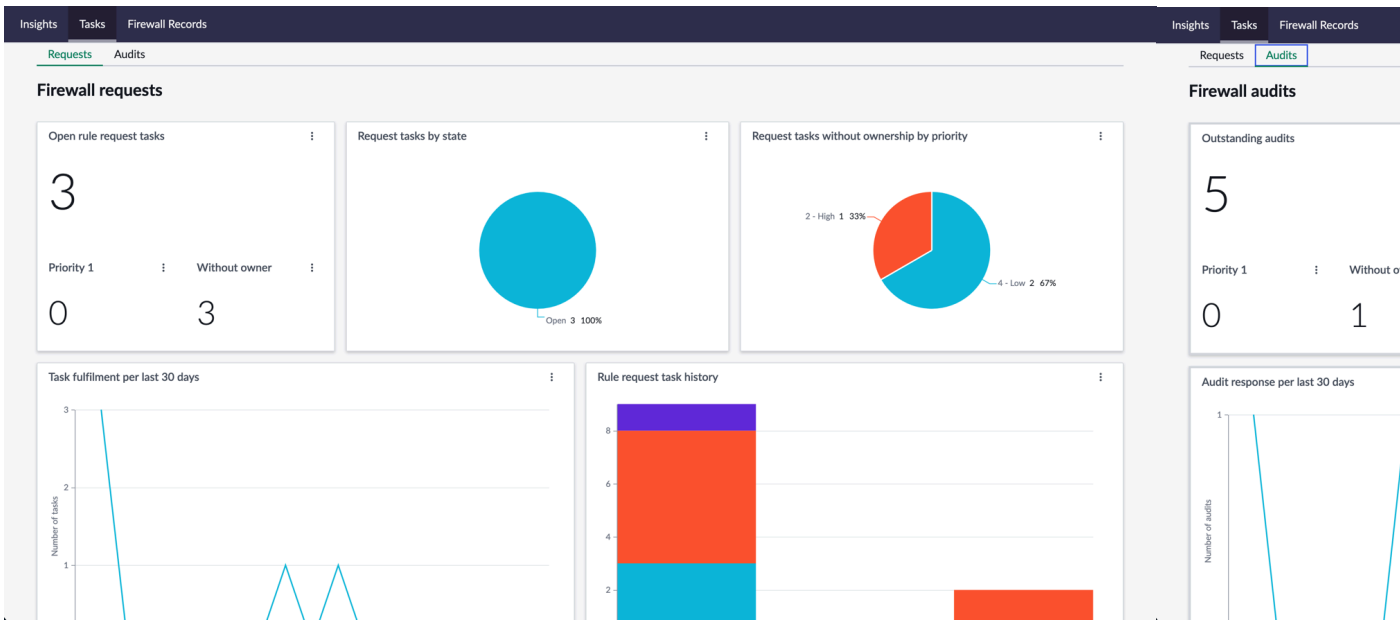

Firewall Records tab

The **Firewall Records** tab offers consolidated access to Records, Rules Requests, and Audits tables. Users can filter, edit, and export individual or groups of records. The **Devices** and **Security Policy** sub-tabs feature a dependency viewer for selected records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

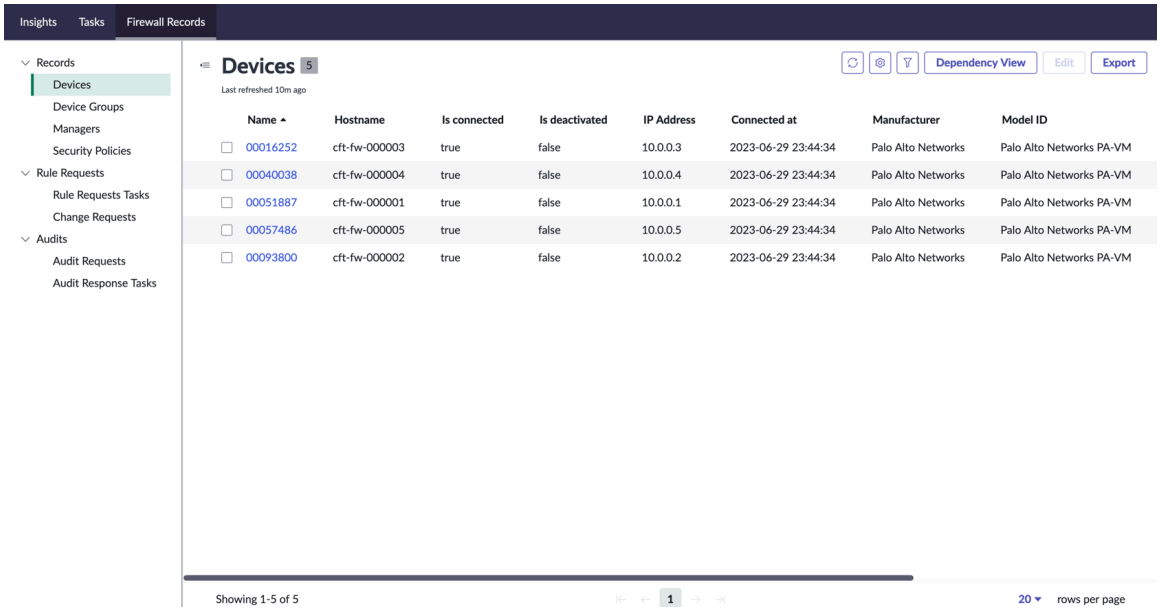

End user and roles

End user and goal	Required role
Risk teams	firewall_admin or firewall_user
Security	firewall_admin or firewall_user

Use case

To update the dashboard manually, navigate to **Performance Analytics** > **Data Collector** > **Jobs** and select **Firewall Scheduled Job**. Open the record and select **Execute Now** to update the dashboard elements with the latest information.

Firewall Audits and Reporting calibration

Explore further details on Firewall Audits and Reporting calibration by reviewing the provided material.

Firewall Audits and Reporting pattern

Palo Alto Firewall Manager

This serverless pattern utilizes the large payload feature for persisting security policies. Discovery admins can establish a schedule using this pattern. When **Run Child Patterns** is enabled, the following child patterns are automatically executed in the specified order:

- PaloAlto - Firewall Manager
- PaloAlto Firewall Devices
- PaloAlto - Firewall Device groups
- PaloAlto - Panorama Security policies

Fortinet FortiManager discovery requirements

Technical requirements and configuration details for discovering Fortinet FortiManager firewalls using the Firewall Audits and Reporting application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

Before discovering Fortinet FortiManager firewalls, verify that the Firewall Audits and Reporting application is downloaded from ServiceNow Store and installed.

Credentials and permissions

Discovery authenticates to FortiManager using an API key. Create a dedicated read-only user account on the FortiManager appliance and generate an API key for that user.

The read-only user account requires the following minimum permissions:

FortiManager user permissions

Permission area	Required access
System status	Read
ADOM management	Read
Device manager	Read
Policy & Objects	Read
Policy Package	Read
Global objects	Read

Configure the FortiManager API key in ServiceNow as follows:

1. Navigate to **Discovery** > **Credentials** and create a new API Key credential record with the FortiManager API key value.
2. Create a Credential Alias and associate it with the credential record.
3. In the FortiManager serverless discovery schedule configuration, input the credential alias.

MID Server configuration

The MID Server must have HTTPS (TCP 443) network access to the FortiManager management IP address. No agent is installed on FortiManager. All discovery is API-driven from the MID Server over HTTPS using the FortiManager JSON-RPC API.

MID Server sizing

FortiManager environments can be large, with thousands of firewall objects and policies. The following sizing is recommended:

- **JVM heap size:** Minimum 4 GB RAM per MID Server. Set *`wrapper.java.maxmemory`* to `4096` in the MID Server configuration.
- **MID Server cluster:** Deploy a cluster of 3–4 MID Servers each configured with 4 GB JVM heap to distribute the discovery load across concurrent pattern executions. You can increase this allocation if you encounter memory issues even after following the above sizing recommendations.

⚠ **Warning:** Running FortiManager discovery on an undersized MID Server (less than 4 GB JVM heap) may result in out-of-memory errors or incomplete discovery, particularly for large deployments with many ADOMs and policy packages.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

API endpoints

All requests use POST to `https://<fortimanager-host>/jsonrpc` with a JSON body. Authentication uses the API key credential alias configured in ServiceNow Discovery.

Key API endpoints used during discovery:

FortiManager API endpoints

Resource	API URL (in params.url)
FortiManager system status	/sys/status
Firewall policies	/pm/config/adom/{{adom}}/pkg/{{policy-package}}/firewall/policy
Global address objects	/pm/config/global/obj/firewall/address
Global address groups	/pm/config/global/obj/firewall/addrgrp
Global service objects	/pm/config/global/obj/firewall/service/custom
Global service groups	/pm/config/global/obj/firewall/service/group
Global VIP objects	/pm/config/global/obj/firewall/vip
Global VIP groups	/pm/config/global/obj/firewall/vipgrp
Global curated ISDB objects	/pm/config/global/obj/_fdsdb/internet-service
Global custom ISDB objects	/pm/config/global/obj/firewall/internet-service-custom
ADOM list	/dvmdb/adom
Devices in ADOM	/dvmdb/adom/{{adom}}/device
Policy packages in ADOM	/pm/pkg/adom/{{adom}}
ADOM address objects	/pm/config/adom/{{adom}}/obj/firewall/address
ADOM address groups	/pm/config/adom/{{adom}}/obj/firewall/addrgrp
ADOM service objects	/pm/config/adom/{{adom}}/obj/firewall/service/custom
ADOM service groups	/pm/config/adom/{{adom}}/obj/firewall/service/group
ADOM VIP objects	/pm/config/adom/{{adom}}/obj/firewall/vip
ADOM VIP groups	/pm/config/adom/{{adom}}/obj/firewall/vipgrp
ADOM custom ISDB objects	/pm/config/adom/{{adom}}/obj/firewall/internet-service-custom

CMDB CI attributes populated

The FortiManager discovery pattern populates comprehensive attributes for all discovered CMDB Configuration Items. The following tables detail all attributes populated for each CI type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

FortiManager Network Manager attributes [cmdb_ci_fortimanager_network_manager]

Attribute	Description
Name	Display name of the FortiManager appliance
Serial Number	Hardware serial number used as the primary identifier
IP Address	Management IP address
Is Virtual	Whether the appliance is a virtual deployment
Platform Type	Hardware platform model type
Time Zone	Configured time zone of the appliance
High Availability Mode	HA mode: standalone, active-passive, or active-active
Hostname	Configured system hostname
ADOM Count	Number of Administrative Domains managed
Managed Device Count	Total number of managed FortiGate devices
ADOM Enabled	Whether the ADOM feature is enabled on this appliance
BIOS Version	BIOS firmware version
Build Number	Firmware build number
FIPS Mode Enabled	Whether FIPS compliance mode is active
Max ADOM	Maximum number of ADOMs supported by the license
Offline Mode Enabled	Whether the appliance is operating in offline mode
System Part Number	Hardware part number

Note: Discovery populates firmware version information in the Firmware Installation [cmdb_firmware_install_list] table instead of directly on the FortiManager CI.

Fortinet Firewall ADOM attributes [cmdb_ci_firewall_device_group_fortinet]

Attribute	Description
Name	Display name of the Administrative Domain
UUID	Globally unique identifier
Device Count	Number of managed devices in this ADOM
Object ID	Unique path to this resource: /dvmdb/adom/{name}
ADOM Type	Product type restriction (for example, fos for FortiOS/ FortiGate)
Target OS Version	FortiOS version that devices in this ADOM are expected to run
Policy Package Count	Number of policy packages within this ADOM
FortiManager	Reference to the parent FortiManager CI
Description	Administrator-assigned description

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fortinet Firewall Device attributes [cmdb_ci_firewall_device_fortinet]

Attribute	Description
Name	Device display name in FortiManager
Serial Number	Hardware serial number used as primary identifier
IP Address	Device management IP address
Manufacturer	Device manufacturer
Model ID	Reference to hardware model record [cmdb_model]
Platform Family	Operating system type (for example, fos)
High Availability Role	HA cluster role: standalone, primary, or secondary
FortiManager	Reference to the managing FortiManager CI
VDOM Enabled	Whether virtual domains are enabled (derived from maxvdom > 1)
Operational Status	Device operational state integer

Note: Discovery populates firmware version information in the Firmware Installation [cmdb_firmware_install_list] table instead of directly on the device CI.

Fortinet Firewall Policy Package attributes [cmdb_ci_fortinet_firewall_policy_pkg]

Attribute	Description
Name	Policy package name
Object ID	Unique path: /pm/pkg/adom/{adom}/{name}
Policy Count	Number of policies in this package
Parent ADOM	Reference to the owning ADOM CI
FortiManager	Reference to the root FortiManager CI
Package Type	pkg (active package) or folder (organizational container)

Fortinet Firewall Policy attributes [cmdb_ci_fortinet_firewall_policy]

Attribute	Description
Name	Policy display name. Fallback: "Policy {policyid}" when name is empty
UUID	Primary identifier auto-generated by FortiOS
Policy ID	Integer policy ID within the package
Policy Enabled	Whether the policy is active
Last Hit	Timestamp of last traffic match
Object ID	Unique path: /pm/config/adom/{adom}/pkg/{pkg}/firewall/policy/{policyid}
Sequence Number	Rule evaluation order within the package

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fortinet Firewall Policy attributes [cmdb_ci_fortinet_firewall_policy] (continued)

Attribute	Description
Policy Package	Reference to the parent Policy Package CI
Parent ADOM	Reference to the ADOM CI this policy belongs to
FortiManager	Reference to the root FortiManager CI
Action	Policy action: accept, deny, or ipsec
Source Interface	Ingress interface name(s), comma-joined
Destination Interface	Egress interface name(s), comma-joined
Source Address	Source address/group object names, comma-joined
Destination Address	Destination address/group/VIP object names, comma-joined
Source Services	Service/service-group object names, comma-joined
Policy Schedule	Schedule object name (for example, always)
ISDB Src Enabled	Whether Internet Service is used at the source
ISDB Dest Enabled	Whether Internet Service is used at the destination
NAT Enabled	Whether source NAT is applied
NAT IP	NAT IP pool name when ippool is enabled
Traffic Logging Mode	Log level: disable, utm, or all
Traffic Session Logging Enabled	Whether session start events are logged
UTM/Security Profile Enabled	True when any UTM profile (AV, IPS, SSL) is applied
AV Profile	Antivirus profile name applied to this rule
IPS Sensor	IPS sensor profile name applied to this rule
SSL/SSH Profile	Deep inspection profile name
Description	Free-text description and change audit trail

Non-CMDB firewall object attributes populated

The following tables detail attributes populated for non-CMDB firewall object tables. All tables inherit common base attributes: Object ID, UUID, Name, Description, Discovery Source, Parent ADOM, and FortiManager reference.

Address Object attributes [sn_disco_firewall_fortinet_address_object]

Attribute	Description
Address Type	Type of address: ipmask, iprange, fqdn, wildcard, geography, or dynamic
Subnet	IP address and mask (present when type = ipmask)
Start IP	Start of IP range (present when type = iprange)
End IP	End of IP range (present when type = iprange)
Fully Qualified Domain Name	FQDN string (present when type = fqdn)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Address Object attributes [sn_disco_firewall_fortinet_address_object] (continued)

Attribute	Description
Geography	ISO country code(s) (present when type = geography)
Associated Interface	Interface this address is bound to. "any" indicates all interfaces

Address Group Object attributes [sn_disco_firewall_fortinet_addrgrp_object]

Attribute	Description
Member Names	Comma-joined list of member address/group names
Member Count	Number of members in the group
Exclude Members	Comma-joined exclude-member list when exclusion mode is active

Service Object attributes [sn_disco_firewall_fortinet_service_object]

Attribute	Description
Protocol	Layer-4 protocol type: ICMP, IP, TCP/UDP/SCTP, or ICMP6
TCP Port Range	Space-separated TCP ports/ranges (present when protocol includes TCP)
UDP Port Range	Space-separated UDP ports/ranges
SCTP Port Range	Space-separated SCTP ports/ranges
ICMP Type	ICMP type number (present when protocol = ICMP)
Service Category	Comma-separated category names this service belongs to

Service Group Object attributes [sn_disco_firewall_fortinet_servicegrp_object]

Attribute	Description
Member Names	Comma-joined list of service member names
Member Count	Number of service members in the group
Proxy Enabled	Whether this is a proxy-mode service group

VIP Object attributes [sn_disco_firewall_fortinet_vip_object]

Attribute	Description
VIP Type	Type of VIP: static-nat, server-load-balance, fqdn, or dns-translation
External IP	Public-facing IP that clients connect to (before NAT)
External Interface	Interface where inbound traffic arrives (for example, wan1, any)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VIP Object attributes [sn_disco_firewall_fortinet_vip_object] (continued)

Attribute	Description
Mapped Address	Internal real server IP or FQDN that traffic is forwarded to after DNAT
Port Forward Enabled	Whether port translation is applied in addition to IP translation
Active	Whether this VIP is active and processing traffic
NAT44 Enabled	Whether IPv4-to-IPv4 NAT is enabled

VIP Group Object attributes [sn_disco_firewall_fortinet_vipgrp_object]

Attribute	Description
External Interface	External interface all member VIPs are associated with
Member Names	Comma-joined list of member VIP object names
Member Count	Number of VIP members in the group

ISDB Object attributes [sn_disco_firewall_fortinet_isdb_object]

Attribute	Description
Reputation Score	Threat reputation score for this internet service entry

Policy Object M2M Mapping attributes [sn_disco_firewall_policy_object_m2m]

Attribute	Description
Firewall Policy	Reference to the Firewall Policy CI [cmdb_ci_firewall_sec_policy]. Cascade delete enabled.
Firewall Object	Reference to the base firewall object record [sn_disco_firewall_object]. Cascade delete enabled.
Reference Type	Which policy field referenced this object: srcaddr, dstaddr, or service
Discovery Source	Source system that populated this record

Relationships created

Discovery establishes the following CMDB relationships between discovered Configuration Items:

FortiManager CI relationships

Parent CI	Relationship Type	Child CI
FortiManager Network Manager	Manages::Managed By	Fortinet Firewall ADOM

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

FortiManager CI relationships (continued)

Parent CI	Relationship Type	Child CI
Fortinet Firewall ADOM	Members::Member Of	Fortinet Firewall Device
Fortinet Firewall ADOM	Contains::Contained By	Fortinet Firewall Policy Package
Fortinet Firewall Policy Package	Used By::Uses	Fortinet Firewall Device
Fortinet Firewall Policy Package	Contains::Contained By	Fortinet Firewall Policy

Discovery pattern

The FortiManager discovery uses a serverless Large Payload Pattern (LP). The pattern name is **Fortinet FortiManager (LP)**.

Discovery consists of one serverless root pattern that triggers multiple independent Large Payload Patterns for global objects, ADOMs, devices, policy packages, and firewall policies.

Quick start test for Firewall Audits and Reporting

Quick start test for Firewall Audits and Reporting to validate configuration items (CI) and their relationships.

Firewall Audits and Reporting Quick Start Test

Test	Description
Firewall Discovery	Firewall Discovery is performed on a firewall endpoint. Discovery brings the firewall manager, policies and devices, and persists them to the cmdb_ci tables. This test validates the firewall manager, policies, and devices data in the cmdb_ci tables and output and discovery errors during the discovery run.

To learn more, see Firewall Audits and Reporting.

Related topics

Quick start tests

Certificate Inventory and Management

With Certificate Inventory and Management, you can discover, conduct an inventory, and proactively manage all TLS certificates. Certificate Inventory and Management supports IPv6, providing comprehensive coverage for your certificate management needs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore	Manage	Reference
Certificate Inventory and Management	Certificate Inventory and Management	Certificate Inventory and Management
Leverage Certificate Inventory and Management to create and manage certificates	Learn how to set up Certificate Inventory and Management	Familiarize yourself with dashboards, tables, patterns, and terms

Troubleshoot and get help

- What is certificate management?
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Certificate Inventory and Management

Certificate Inventory and Management serves as a centralized and automated solution for handling the complexities of certificate management. It enhances security, ensures compliance, and provides a streamlined approach to managing the life cycle of digital certificates within an organization.

Certificate Inventory and Management overview

Certificate Inventory and Management is a centralized and automated solution designed to streamline the tracking and management of digital certificates in an IT infrastructure. By automating the discovery, inventory, and monitoring processes, Certificate Inventory and Management enables organizations to maintain a comprehensive record of certificates, including essential details like expiration dates and issuers.

The system facilitates real-time tracking, automates routine tasks such as renewals, and supports customization to align with specific organizational needs. With integration capabilities into existing CI-based Discovery schedules and support for IPv6, Certificate Inventory and Management offers flexibility and efficiency in managing the entire life cycle of certificates. Role-based access control ensures secure and tailored usage, while alerts and notifications keep administrators informed about critical events, contributing to enhanced security, compliance, and operational efficiency in certificate management.

Domain separation is supported in Certificate Inventory and Management, certificate records and certificate tasks can be separated by domain. With domain separation, you can control the visibility of certificate data across domains.

Certificate Inventory and Management workflow

Certificate Inventory and Management streamlines the process of scanning for certificates by automatically detecting them on specific ports through your established Configuration Item (CI)-

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

based Discovery schedules. Additionally, Certificate Inventory and Management empowers users with the flexibility to generate new Discovery schedules tailored for scanning individual URLs. This dynamic approach ensures a robust and adaptable method for discovering and managing certificates, contributing to a secure and well-organized IT infrastructure.

Process flow diagram

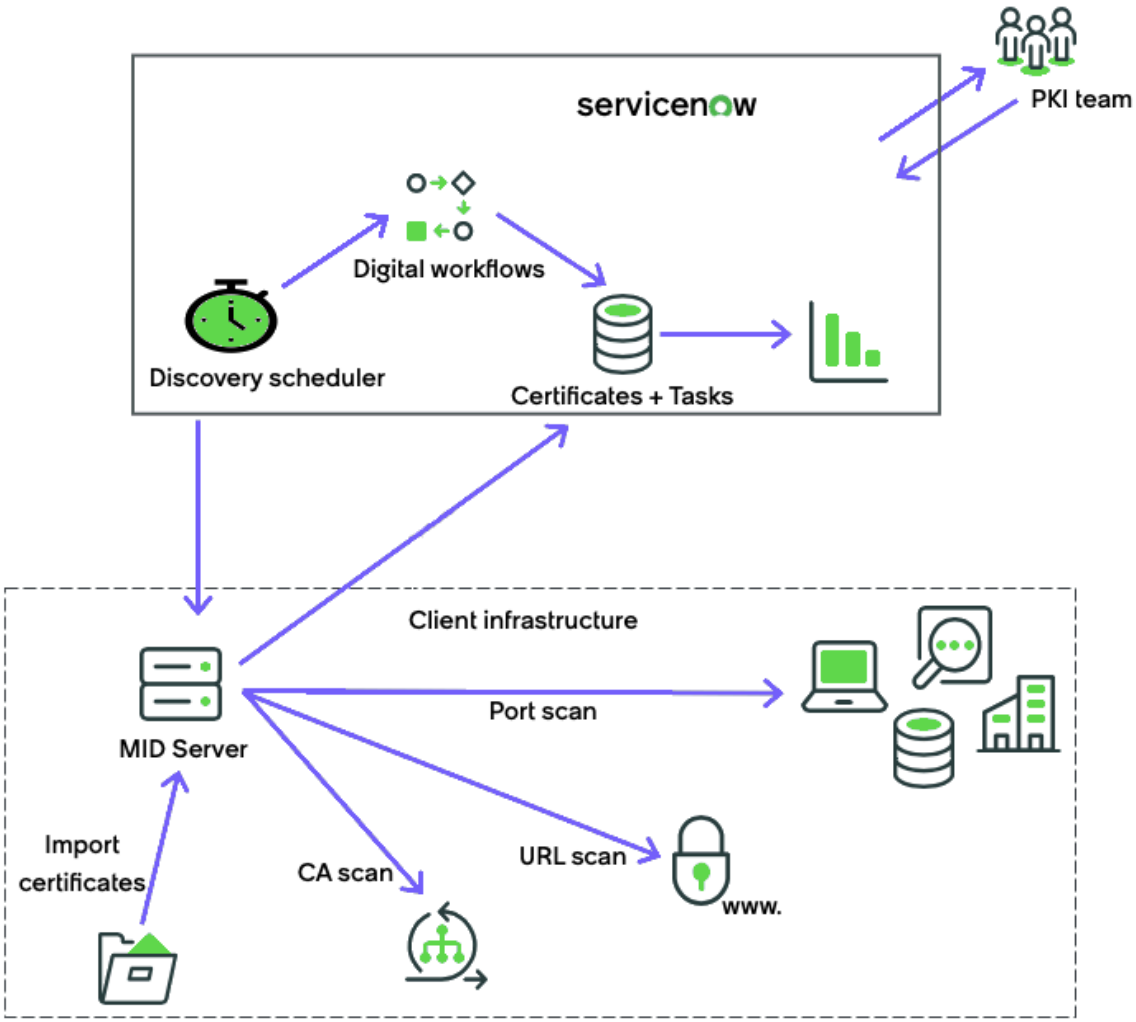

For more information on how Certificate Inventory and Management works, see Certificate Inventory and Management process flow.

Certificate Inventory and Management benefits

Benefit	Feature	Users
Management of TLS certificates by maintaining a centralized inventory, prioritizing certificates, and leveraging renewal workflows to prevent costly outages and ensure a secure environment.	Configuring Certificate Inventory and Management	Certificate Administrator [sn_disco_certmgmt.pki_admin]
Automated features to streamline tracking, preventing manual oversight of a large volume of certificates. Embrace a proactive approach by staying informed about impending	Automated certificate management with ACME	Certificate User [sn_disco_certmgmt.pki_user]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefit	Feature	Users
expirations, contributing to a well-managed and secure certificate environment.		
A secure and compliant certificate environment is maintained by overseeing, approving, validating, and authorizing necessary actions related to certificate tasks and flows. This comprehensive approach ensures that certificate renewal tasks and incidents for expired certificates are appropriately addressed, contributing to a robust security framework.	Approve certificate tasks	Certificate Approver [pki_approver]

Certificate Inventory and Management process flow

Certificate Inventory and Management allows Discovery to automatically scan for certificates on specific ports through your existing CI-based Discovery schedules. In addition, you can create new Discovery schedules to scan individual URLs.

How Certificate Inventory and Management works

In Version 1.1.7 Certificate Inventory and Management, you can also scan for certificate authorities (CA) and import certificates from files.

With the Brazil release, you can download Version 1.2.0 Certificate Inventory and Management. You can discover certificates used in the HTTP(s) Endpoint [cmdb_ci_endpoint_http] table and bulk import a list of SSL Certificates. You can create events and alerts for expiring and expired certificates and be notified in Slack of expiring and expired certificates.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Process flow diagram

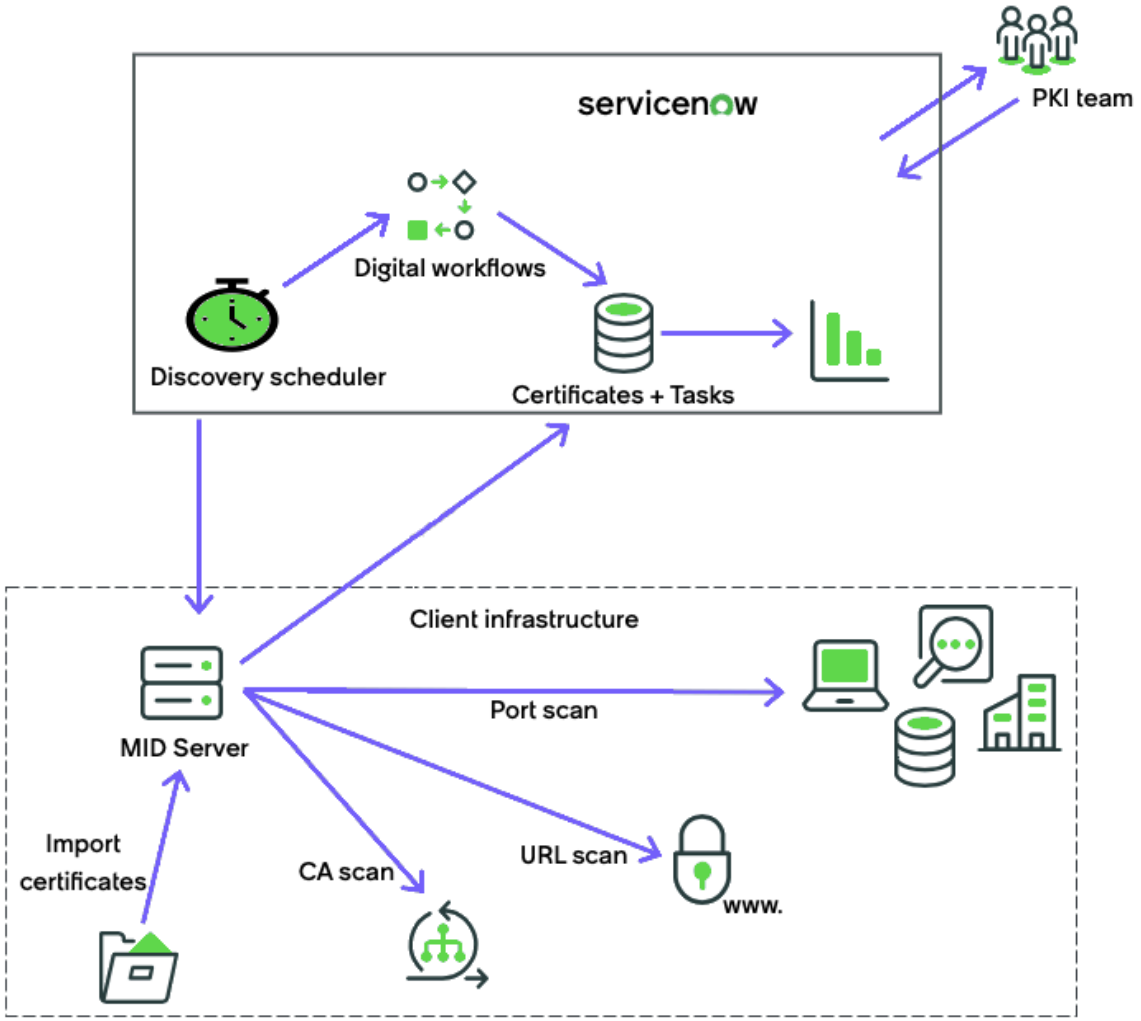

The Certificate Inventory and Management process flow is segmented into two distinct phases:

- Pre-discovery phase: The initial segment of the process varies based on the source of certificate discovery. This phase focuses on actions required for identifying certificates from different sources within the IT infrastructure.
- Post-discovery phase: This phase involves subsequent actions taken after the certificates have been successfully discovered. It includes processes such as cataloging, tracking, and managing certificates to ensure their secure and efficient integration into the system.

Pre-discovery phase

The pre-discovery phase involves preparatory steps, such as defining scanning parameters and configuring credential details, for a smooth initiation of the certificate discovery process.

Discovery via Ports

The port probe (tls_ssl_certs) automatically scans 14 default preauthorized ports.

- Typical ports for SSL: 443, 8443, 9443, 636 (ldaps), 993 (imaps), 995 (popssl), 989, 990
- StartTLS ports: 25 (smtp), 110, 143, 389, 21, 587 (smtp)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

As part of the CI Discovery process during Shazzam, the MID Server uses scanners to gather certificate chain information from the IP port number and capturing diverse attributes, including the certificate hierarchy. The MID Server then transforms these certificates into an XML payload, sharing it with the instance. The Shazzam sensor, in turn, detects the ECC queue entry and inserts a new record into the Discovered Certificate (sn_disco_certmgmt_certificate_history) table.

The following fields are pulled from the XML payload and verified in java code from the Shazzam TLS port probe for discovered certificates: certificate id, revocation_status, subject, issuer, sans/, is_self_signed, is_ca, valid_from, valid_to, signature_algorithm, fingerprint_algorithm, key_size, serial_number, and version.

Discovery via URL

The Certificate URL (sn_disco_certmgmt_cert_url) table holds a list of URLs to target for certificate discovery. Each record also has an optional reference to the Unique Certificate (cmdb_ci_certificate) table, to see what certificate is related to the given URL definition. The necessary parameters from the Discovery Schedule are combined to create and initialize the Discovery status. The CertificateDiscoveryFromURLScan probe discovers the certificate chain for each of the URLs in the batch and outputs an XML payload that contains the certificate chain for each certificate. It also adds a new record into the Discovered Certificate (sn_disco_certmgmt_certificate_history) table.

Discovery via Import Certificates (Version 1.1.7 Certificate Inventory and Management)

The Import certificates are discovered through the Import SSL Certificate pattern, which relies on the following parameters.

- Host name/IP where the certificates are hosted
- Folder where certificates are located
- TLS_keepOriginalCertificate: Setting this parameter to true may lead to increased payload size, potentially causing out-of-memory issues.
- Mid_temp_folder: The temporary folder on the MID Server where files will be temporarily copied.

Note: Auto-select MID Server option is NOT supported for Windows and Linux mid combinations. If the MID Server is used for storing the original certificates files, then Host name/IP should be set to blank or localhost and the particular MID Server should be selected for the Discovery schedule.

Discovery via CA authority (Version 1.1.7 Certificate Inventory and Management)

After the Certificate Inventory and Management credential is set up with either GoDaddy, DigiCert, Entrust, or Sectigo Certificate Authority and the Discovery schedule runs, the specific CA pattern makes REST API calls to (GoDaddy, DigiCert, Entrust, or Sectigo), collects certificate information, retrieves the list of certificates, and stores it in the cmdb_ci_certificate, certificate_domain, and sys_attachment tables.

ca_api_url and ca_api_version are optional parameters. If these parameters are left empty inside pattern parameters, default values will be used. The default values include:

- DigiCert - Certificate Management (ca_api_version = v2, ca_api_url = https://www.digicert.com/services/)
- Entrust - Certificate Management (ca_api_version = v2, ca_api_url = https://api.entrust.net/enterprise/)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- GoDaddy - Certificate Management (ca_api_version = v1, ca_api_url = https://api.godaddy.com/)
- Sectigo - Certificate Management (ca_api_version = v1, ca_api_url = https://cert-manager.com/api/ssl/)

The arguments for GoDaddy, DigiCert, Entrust, and Sectigo patterns are as follows. Starting from Version 1.2.0, you have the capability to scan Sectigo and Entrust Certificate Authorities (CAs).

- Start_offset: The offset position for reading certificates from CA authorities, with a default value of 0.
- Limit: The number of certificates to be read from the start_offset, with a default value of 1500.
- CredentialAlias: The name of the Credential Alias or Tag linked to the CA credentials, added in the serverless execution pattern configuration.

  If the TLS_keepOriginalCertificate parameter is set to true, the certificate file is attached to the Certificate CI. This may increase payload size, potentially causing out-of-memory issues.

- IncludeCertStatus: A parameter for specifying additional certificate states to discover, in addition to the defaults.

Certificate states by certificate authorities

Certificate Authority	Default States Discovered
Sectigo	○ Issued ○ Expired
DigiCert and GoDaddy	○ Active ○ Expired ○ Revoked ○ Canceled
Entrust	○ Active ○ Expired ○ Revoked

You can include multiple certificate statuses by separating each with commas.

Note: The State field in the Unique Certificate (cmdb_ci_certificate) table denotes the life cycle state of the certificate, not the raw state returned by the CA though API. Certificates discovered with a CA status of issued, valid, expired, or canceled, are all denoted with the State value Issued. The State field doesn't include an expired value. To identify certificates that have expired, use the Valid to value rather than the State value.

After the pre-discovery phase is complete, move on to the post-discovery phase.

Post-discovery phase

Following the discovery phase, the system manages TLS certificates, offering flexibility for both manual and automated request options, catering to various certificate-related tasks and processes.

The TLS certificate chains, signed by the certificate authority and root, populate the Unique Certificate [cmdb_ci_certificate] and Installed Certificate

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

[sn_disco_certmgmt_cmdb_installed_certificate] tables. Subsequently, a scheduled job reviews the Unique Certificate table for expiring and expired certificates, initiating the creation and assignment of certificate tasks and incidents.

You have the option to manually request new certificates and renew existing ones. For more information, see Manual flow for certificate requests.

In Version 1.3.8, the capability to automate requests for new certificates, renewals, or revoking certificates has been introduced. For more information, see Automated certificate management for TLS certificates.

To optimize system performance, a table cleaner automatically removes old certificate records from these two tables after a specified number of days:

- Discovered Certificate [sn_disco_certmgmt_certificate_history] table: older than 30 days
- Installed Certificate [sn_disco_certmgmt_cmdb_installed_certificate] table: older than 90 days

Note: You can toggle various behaviors related to Certificate Inventory and Management depending on your needs, using specific certificate properties as shown in Discovery properties and System properties.

The system scans for certificates expired or archived for more than six months. If a replacement certificate exists, the system removes outstanding tasks for these obsolete certificates.

The same six-month grace period applies for expired and archived certificates and certificate tasks in the *failed* state.

Certificate Inventory and Management roles and responsibilities

Dedicated users with specialized roles are assigned to optimize the monitoring and tracking of requests for new and renewing certificates.

The Certificate Inventory and Management dashboard and certificate configuration settings are activated or deactivated depending on the assigned roles.

Role	Responsibilities
Certificate administrator [sn_disco_certmgmt.pki_admin]	The headless user (an account not tied to a specific user) is included in the base system and equipped with the *sn_disco_certmgmt.pki_admin* role. It serves as the caller for automatically generated renewal certificate tasks and incidents. To customize this user, adjust the user ID using the Discovery property: *glide.discovery.certs.cert_admin_user_id* instead of leaving it as the default headless user. Responsible for changing non-standard attributes in the original certificate record, this role can modify attributes like state, status, assigned to, assignment group, renewal tracking, and service type. The certificate's inherent attributes remain unaltered. The default state for discovered certificates is installed, but this role can manually adjust it to other states such as issued, installed, revoked, and retired. Additionally, users with this role have the capability to view diverse dashboards and possess read/write access to certificates and certificate tasks associated with certificate Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role	Responsibilities
	Note: • The sn_disco_certmgmt.pki_admin role contains the sn_disco_certmgmt.pki_user role. • Any user with the pki_admin role can be added to the approver field in the certificate request form. Users with pki_admin role should be assigned with the approval_admin role manually to approve or reject requests from the certificate task form. This approver does not need to login to approve the task. In a certificate task form or related list, system admin or a user with approval_admin role or pki_admin role can add multiple approvers or edit the existing one. • The approver_user role is required if the user wants to approve a request by logging in. Without the approver_user role, a user doesn't have the option to navigate to **My Approvals** where the requests can be approved.
Certificate user [sn_disco_certmgmt.pki_user]	Responsible for overseeing certificate discovery, this role is granted the ability to access diverse dashboards and has read/write permissions for certificates and associated certificate tasks.
Certificate approver [sn_disco_certmgmt.pki_approver]	Responsible for certificate requests, a user with this role (normal user) can initiate certificate requests through the Service Catalog form.
Certificate requester [sn_disco_certmgmt.certificate_requester]	Responsible for submitting certificate requests, this role is granted the ability to request and renew certificates from the Service Catalog.

Certificate Management workspace

The Certificate Management workspace provides centralized visibility into your organization's certificates so you can make data-driven decisions. For example, you can avoid outages by looking at the numbers that are soon to expire.

Roles Required

- Certificate user [sn_disco_certmgmt.pki_user]
- Certificate approver [sn_disco_certmgmt.pki_approver]
- Certificate admin [sn_disco_certmgmt.pki_admin]

Accessing the Certificate Management workspace

To open the workspace, navigate to **Workspaces** > **Certificate Management**.

The Certificate Management workspace has five tabs:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate Management workspace tabs

Tab	Description
Inventory	Displays key insights about all certificates including certificates that are active, that have expired, that are expiring soon, and that have been revoked. Select each widget to view its list. This is the default tab of the workspace. **Note:** The data on this tab is refreshed automatically every day.
Tasks	Displays tasks to track the upcoming tasks that will expire soon. These tasks are organized into four categories: Renewals, Requests, Expirations, and Automation trends.
Certificate records	Displays a filterable list view of all certificates in the system. Use this tab to search, inspect, and manage individual certificate records. You can switch between record types and access the configuration options.
Downloads	Displays the files that are required to set up the ServiceNow external issuer (`sn-external-issuer`) in Kubernetes.
Settings	Displays the settings to manage the notifications of this workspace. For more information on setting up these notifications, see Receive certificate notifications via Microsoft Teams and Receive certificate notifications via email.

AI in Certificate Inventory and Management

AI Agents for Discovery provides an AI agent in Certificate Inventory and Management that helps PKI administrators automatically renew TLS certificates through a conversational prompt.

Certificate Inventory and Management includes an AI-powered certificate renewal agent. This feature requires ServiceNow Otto for IT Operations Management (ITOM) to be installed.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

AI features in Certificate Inventory and Management

AI feature	Description	Use case	Resources
Certificate renewal AI agent	Uses AI to automatically renew a single TLS certificate through a conversational prompt. The agent guides the user through the renewal	A PKI administrator wants to renew a certificate without	Renew certificates

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI features in Certificate Inventory and Management (continued)

AI feature	Description	Use case	Resources
	process and provides a link to the task record when complete.	manually navigating the renewal workflow.	using the AI agent

Configuring Certificate Inventory and Management

Streamline the process of gaining insights into certificate and inventory Management, covering key aspects like certificate discovery, bulk upload functionality, maintenance of certificate chain relationships, configuration of Credential Identifiers, customization of the system, creation of new/renewal requests, archiving tasks, automation of TLS certificate management, integration with Event Management, Slack notifications for alerts, and leveraging cert-Manager for efficient certificate generation.

Configuration overview

Before using Certificate Inventory and Management, it's crucial to set up the necessary plugins, ensure activation, upgrade your instance to the latest version, and download the Certificate Inventory and Management application from ServiceNow Store. For additional details, see Get started with Certificate Inventory and Management.

After the prerequisites are met, you can proceed with the following tasks.

Certificate Inventory and Management tasks

Task	Description
Visibility to TLS certificates	Establish an efficient certificate inventory management system by configuring certificate discovery, enabling bulk certificate uploads, managing chain relationships, configuring secure access, and customizing the system to meet specific organizational needs.
Manual flow for certificate requests	Enable the creation of new certificate requests or certificate renewals within the system, adhering to defined policies. Implement a process to archive completed certificate tasks, maintaining a historical record for auditing and compliance purposes.
Automated certificate management for TLS certificates	Set up automation processes to manage TLS certificates, including renewal, updates, and revocation as needed.
Certificate alerts and notifications	Establish integration between Event Management systems and Certificate Inventory Management for real-time monitoring and response to certificate-related events. Configure notifications to be sent to Slack channels, keeping stakeholders informed about critical certificate events and changes.
Certificate generation through Kubernetes cert-manager	Utilize Cert-Manager or a similar tool to automate the certificate generation process, ensuring efficient and secure issuance based on predefined policies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started with Certificate Inventory and Management

Install required plugins, upgrade your instance, and download the Certificate Inventory and Management application from the ServiceNow Store.

Before you begin

Role required: admin

About this task

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Note: The ServiceNow Store regularly releases new applications and updates to ServiceNow applications. If you already have an application, you can download the latest version to enhance your existing experience with ServiceNow products. Features vary by release. The version number indicates which content and features are available in each release.

Procedure

1. Ensure the following plugins are installed and activated.
   - ITOM Visibility [com.snc.itom.vis.license] plugin
   - Discovery [com.snc.discovery] plugin
   - Configuration Management for Scoped Apps (CMDB) [com.snc.cmdb.scoped] plugin
2. Confirm that your system has been upgraded to Brazil or a later version.
   During the upgrade process, the Certificate Inventory and Management [com.sn_disco_certmgmt] plugin is automatically installed if the ITOM Visibility [com.snc.itom.vis.license] and Discovery [com.snc.discovery] plugins are already installed.
3. Download the Certificate Inventory and Management application from ServiceNow Store.

Result

You are ready to get started with Certificate Inventory and Management.

Note: The MID Server needs access to the following endpoints based on your use case:

- Entrust: Discover Certificate URL: `https://api.entrust.net/enterprise/`
- Entrust Download root Certificate URL: `https://web.entrust.com/root-certificates/`
- DigiCert: `https://www.digicert.com/services/`
- Sectigo: `https://cert-manager.com/api/ssl/`
- GoDaddy: `https://api.godaddy.com/`

Visibility to TLS certificates

The Certificate Inventory and Management application allows Discovery to automatically scan for certificates on specific ports through your existing CI-based Discovery schedules. In addition, you can create Discovery schedules to scan for specific URLs.

The ServiceNow Store regularly releases new applications and updates to ServiceNow applications. If you already have an application, you can download the latest version to enhance your existing experience with ServiceNow products. Features vary by release. The version number indicates which content and features are available in each release.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In Certificate Inventory and Management, you can add a list of imported certificates to Run Certificate Discovery via certificate file import, and scan for certificates from your Certificate Authority (CA) such as GoDaddy and DigiCert. You can also scan Sectigo and Entrust CAs.

The existing certificate authority patterns for DigiCert to collect the following fields as part of the CA Trust Certificate discovery. These fields are required to create an automated flow (request, renew, or revoke) for certificates discovered by Digicert CA Discovery and are stored in the Certificate Extensions [sn_disco_certmgmt_certificate_extension] table.

- Certificate Id
- Order id
- Thumbprint
- Serial Number
- Certificate Status

Using the Certificate Inventory and Management, you can Run Certificate Discovery via port scans. You can also Run Certificate Discovery via individual URL scans.

To Import Certificates or Discovery CA Trust with more than 1500 certificates, create the discovery schedule with more than one serverless patterns configured. Each pattern execution supports a maximum of 1500 certificates discovery.

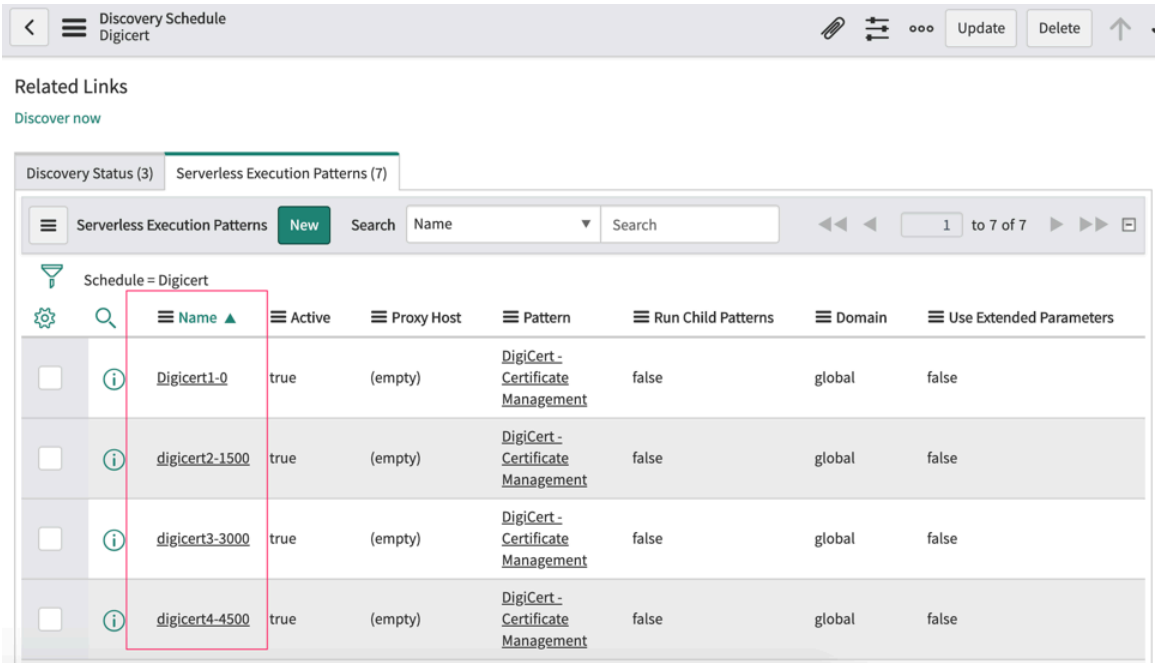

To discover all the certificates, the limit (defaults to 1500) and start_offset (defaults to 0), must be configured accordingly. For example, to fetch up to 6,000 certificates, add four serverless patterns with start_offset 0, 1500, 3000, and 4500. Start_offset and limit parameters are configured as shown.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

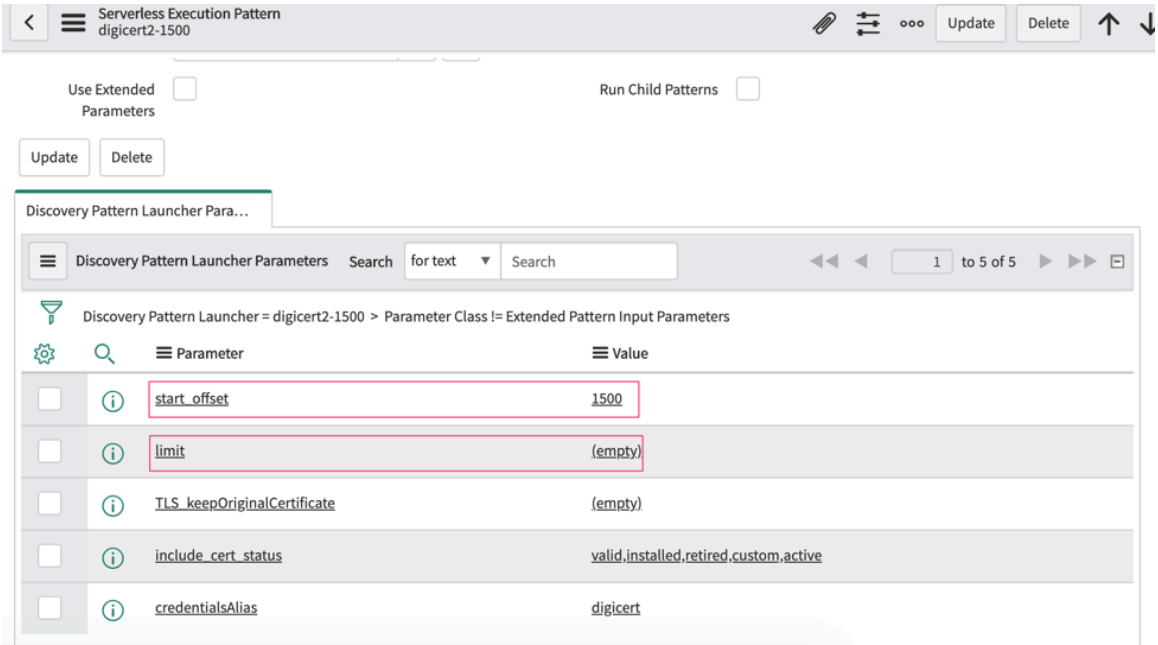

Run Certificate Discovery via port scans

When the TLS port probe [tls_ssl_certs] is enabled, Discovery automatically scans 14 pre-authorized ports as part of your existing CI Discovery schedules.

Before you begin

Role required: discovery_admin or admin

Procedure

1. Activate the TLS port probe [tls_ssl_certs].
   a. Navigate to **Discovery Definition** > **Port Probes**.
   b. Open **tls_ssl_certs**.
   c. To enable the probe, select the **Active** check box.

      By default, the check box for any new installation is unchecked.
   d. Save your changes.
2. Add IP service to help configure the TLS port probe.
   a. Navigate to **Discovery Definition** > **IP Services**.
   b. Create a new IP service with a port.
3. Configure the TLS port probe.

   Edit the Port Probe definition to add up to 138 additional ports or remove existing ones.
   a. Navigate to **Discovery Definition** > **Port Probes**.
   b. Open **tls_ssl_certs**.
   c. Unlock the **Triggered by services** field by selecting the lock icon next to it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Remove any ports from the list or add additional ones from the search area.

e. Save your changes.

Result
Your existing Discovery schedules should then automatically scan for any certificates on the specified ports.

Run Certificate Discovery via port scans in Discovery Admin Workspace

Configure a certificate port probe in the Discovery Admin Workspace, so Discovery collects certificates from devices with secure ports open during your existing IP-based Discovery schedules.

Before you begin
Confirm the following:

- Discovery Admin Workspace v1.20.0 is installed.
- Certificate Inventory and Management v4.2.4 is installed.
- The ServiceNow AI Platform is running on the Brazil, Australia, or Zurich release starting with Patch 8.

Role required: discovery_admin

About this task
Port scan discovery reuses your existing IP-based Discovery schedules. Whenever Discovery finds a device with one of the listed ports open, it reads and collects the certificate served on that port. Use this method when certificates are served on secure ports across your network.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Certificate discovery** and select **Continue**.
4. Select **Discover via port scans**, then select **Continue**.
   The Certificate Port Probe page displays with a list of ports to scan for certificates.
5. Select the **Active** check box to enable the port probe.
6. Add ports to scan for certificates.
   - To add an existing IP service, select **Add** or **Add** > **From existing IP services**. In the Add from existing Ports dialog, select the check box next to each port to add, then select **Save**.
   - To add a new IP service, select **Add** > **New IP services**. In the Add a new IP service dialog, complete the IP Service fields and select **Add**.

   To remove a port, select it and then select **Remove**.
7. Select **Save**.
   Discovery collects certificates from devices that have one of the listed ports open during your IP-based Discovery schedules.

Run Certificate Discovery via individual URL scans

To initiate certificate discovery through URL scans, you must manually include individual URLs and configure a new certificate Discovery schedule.

Before you begin
Role required: discovery_admin or admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Only the certificates that are available on the server during URL scans can be discovered. To confirm the available certificates, use the following command:

```
openssl s_client -showcerts -connect <URL>:<PORT> </dev/null
```

Procedure

1. Navigate to **All** > **Certificate Management** > **Certificate Discovery Source URLS**.
2. To add individual URLs to the table, select **New**.
   Ensure accurate results by entering URLs in the following format: scheme://host:port. The port is optional, and defaults are used if not provided. For example: https://www.servicenow.com or https://servicenow.com:443, ldaps://myldap.com or ldaps://myldap.com:636.
3. Create a Discovery schedule with the following fields.
   For more information on setting up your Discovery schedules, see Schedule a horizontal discovery.

   a. Select **Discovery**: **Certificates**.

   b. Select **Certificate Discovery Type**: **URL Certificate Discovery**.

      Keep the batch size as is, unless there is a specific recommendation to change it.

4. To add or delete other URLs, from the **Certificate URLs** tab, select **Edit**.
5. Select **Submit**.
6. Select the checkbox to include URLs from the HTTP(s) Endpoint [cmdb_ci_endpoint_http] table in the discovery process.

Result

When your Discovery schedule runs, it automatically scans for any certificates on the specified URLs and fetches all URLs from the cmdb_ci_endpoint_http table. It then creates a mapping between the URL and the schedule in the *sn_disco_certmgmt_cert_url_sched_m2m*.

With Service Mapping enabled, by default, it creates a relationship between the HTTP endpoint and application when it creates an entry in cmdb_ci_endpoint_http. For example, the Amazon application is automatically connected to amazon.com.

The relationship is: cmdb_ci_endpoint_http[parent] --> [Implement End Point To::Implement End Point From] --> cmdb_ci_appl[child].

If the above relationship exists, the URL certificate discovery creates an additional relationship between the certificate and application. This relationship is: cmdb_ci_appl[parent] --> [Uses::Used by] --> cmdb_ci_certificate[child].

Note: URL discovery schedules do not generate server configuration items (CIs).

Run Certificate Discovery via URL scans in Discovery Admin Workspace

Create a certificate Discovery schedule in the Discovery Admin Workspace to collect certificates from specific URLs rather than from ports found during IP-based Discovery.

Before you begin

Confirm the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Discovery Admin Workspace v1.20.0 is installed.
- Certificate Inventory and Management v4.2.4 is installed.
- The ServiceNow AI Platform is running on the Brazil, Australia, or Zurich release starting with Patch 8.

Role required: discovery_admin

About this task

URL scan discovery visits each URL that you list and collects the certificate served at that URL. Use this method when you need to collect certificates from specific endpoints rather than from ports found during IP-based Discovery.

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Schedules**.
2. Select **New Discovery** from the header of any tab on the Schedules page.
3. Select **Certificate discovery** and select **Continue**.
4. Select **Discover via URL scans**, then select **Continue**.
5. Enter a name for the Discovery schedule.
6. Choose your MID Server.

Option	Description
Use a cluster of MID servers	Select an existing MID Server cluster from the drop-down list. Clusters provide failover protection and load balancing between MID Servers. See for more information.
Automatically select a MID server	An available MID Server is automatically selected when the Discovery schedule runs. See Automatic MID Server selection for more information.
Search for a MID server	Select an existing MID Server from the drop-down list. Only MID Servers that are validated and up are displayed.

7. Select **Next**.
8. Choose the URLs to scan for certificates.
   a. From **Available URLs**, select the check box for each URL to scan.
   b. **Optional:** To add a URL that isn't listed, select **Add URL** and enter the URL. Select **Add more** to add multiple URLs. Then select **Add URLs**.
9. Confirm your selections appear under **Selected URLs**.
10. Select **Next**.
11. Configure the Discovery schedule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Run on demand	The schedule only runs when triggered manually. Trigger the schedule by selecting **Finish and run** at the end of this setup, or by navigating to the schedule in the Schedules table and selecting **Discover Now**.
Run at a scheduled time	The schedule runs at a scheduled date and time. Use the fields to define when the schedule runs.
Run after series	The schedule runs after another existing Discovery schedule completes, staggering or chaining them together. Selecting an existing schedule displays a relationship map of all associated schedules.
Cancel discovery if longer than maximum runtime	If the schedule exceeds the maximum run time, it's canceled. After the option is toggled on, configure the runtime threshold.
Finish and run	After you select this option, all the information you provided is validated. Then, a Discovery schedule is created in the background, a Discovery status is created, and the schedule is run. You're redirected to the Status Details page for the schedule.
Finish	After you select this option, all the information you provided is validated. A Discovery schedule is created and you're redirected to its entry in the Discovery Schedules [discovery_schedule] table. You can edit the schedule information or run it by selecting **Discover Now**.

- If you selected **Finish and Run**, Discovery visits each selected URL and collects its certificate. You're redirected to the Status details page for the schedule.
- If you selected **Finish**, you're redirected to the Schedule details page for the schedule.

Run IP-Based Certificate Discovery

Enable the Transfer Layer Security (TLS) port probe [tls_ssl_certs] and scan for certificates on an IP address or multiple IP addresses.

Before you begin

Know the IP address, range of IP addresses, or list of IP addresses you want to perform Certificate Discovery on.

Role required: Certificate administrator, discovery_admin, or admin

Procedure

1. Activate the TLS port probe [tls_ssl_certs].
   a. Navigate to **All** > **Discovery Definition** > **Port Probes**.
   b. Open **tls_ssl_certs**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. To enable the probe, select the **Active** check box.

   By default, the check box for any new installation is cleared.

d. Confirm the **Triggered by services** has the services you need.

   By default, the following services trigger the tls_ssl_certs probe: HTTPS, tomcat-ssl, IBM Websphere SSL, ldaps, IMAPS, pop3s, ftps-data, ftps, smtp, pop3, imap, ldap, ftp, submission.

e. Select **Update**.

2. Create an IP-Based Certificate Discovery Schedule.

   a. Navigate to **All** > **Discovery** > **Discovery Schedules**.

   b. Select **New**.

   c. Fill out each text field with its corresponding value.
   For more information on the fields and values, see IP-Based Discovery Schedule form and fields table.

   d. Select and hold (or right-click) above **Discovery Schedule** and select **Save**.
   By selecting **Save**, additional configuration options are available.

   e. Select the **Use SNMP** text field.

   f. In the new tabs that appear, select **Discovery IP Ranges**.

   g. Select **New**.

   h. Fill each field with its corresponding value.
   For more information on the fields and values, see Setting Your IP Addresses form and fields table.

   i. Select **Submit**.
   The port **tls_ssl_certs** will look for Certificates available inside your selected IPs at the next time interval (daily, weekly, monthly, etc.).

3. Create an IP-Based Certificate **Quick Discovery** for immediate discovery (optional).

   a. Navigate to **All** > **Discovery** > **Discovery Schedules**.

   b. Select **Quick Discovery**.

   c. Enter your Target IP into the **Target IP** field.

   ⚠ **Warning:** Make sure that your IP address is accessible to the MID Server you're using. If it doesn't, you have to set Credential aliases for Discovery.

   d. Select **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The probe tls_ssl_certs searches for Certificates inside your selected IP immediately, and reports the certificates it finds.

Use bulk certificate upload

In Certificate Inventory and Management version 1.2.0 and later, you can import up to 5000 SSL certificates in bulk using a single .xlsx file.

Before you begin
Role required: pki_user or pki_admin

The default admin is assigned both the pki_user and pki_admin roles. The users with the following roles have access to bulk certificate upload.

- sn_disco_certmgmt.pki_user
- sn_disco_certmgmt.pki_admin

Procedure

1. Navigate to **All** > **Certificate Management** > **Bulk Upload Certificates**.
2. From the Bulk Upload page, select **Download Template File** (.xlsx) link to download the `sample_cmdb_ci_certificate.xlsx` file.
3. Open the downloaded `sample_cmdb_ci_certificate.xlsx` file.
4. Refer to the values in the sample rows as a guide and add the SSL certificate's properties into the corresponding columns.

   Each entry in the Excel sheet corresponds to a CI certificate. For a description of each field in the `sample_cmdb_ci_certificate.xlsx` field, see .

   Note:
   - For maintaining the certificate's chain relationship, every certificate must have the fingerprint of its root certificate in the **root_issuer** column and the fingerprint of its issuer in the **issuer** column.
   - For self-signed certificates, both **root_issuer** and **issuer** must have the value of the certificate's fingerprint.

5. Delete the sample rows and then save the `sample_cmdb_ci_certificate.xlsx` file.
6. Navigate back to the Bulk Upload Certificates form.
7. Select **Browse File** and choose the updated `sample_cmdb_ci_certificate.xlsx` file.
8. Select **Upload**.

Result
If any errors occur while uploading, a warning message and link to the error logs appears. Only users with the admin role can view these logs. Errors can occur when a mandatory field is left empty or a valid_from epoch value is greater than that of valid_to.

After the file is successfully uploaded, you can see a success message displayed with a *`View All`* redirect link. You can view a list of the certificates that were added to the Unique Certificate [cmdb_ci_certificate] table.

Run Certificate Discovery via certificate file import

In Certificate Inventory and Management, you can discover certificates by importing certificate files into the system using pattern-based Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that **TLS_keepOriginalCertificate** is set to *False*.
- Ensure the certificate is in one of the following formats: .cert, .pem, .txt, .der, and .cer.
- Role required: discovery_admin or admin

About this task

The import certificates are discovered using the Import SSL Certificate pattern, which relies on the following parameters:

- Server: The host name or IP address where the certificates are hosted. If importing from the MID Server, specify localhost and leave temp_certificate_folder empty.
- server_certificates_folder: The folder path on the server containing the certificates.
- TLS_keepOriginalCertificate parameter: Set to false to avoid potential out-of-memory issues due to increased payload size.
- temp_certificates_folder: Temporary folder on the MID Server where the files are copied.

Procedure

1. Create a Discovery schedule with the following fields.
   For more information on setting up your Discovery schedules, see Schedule a horizontal discovery.
   a. Select **Discovery**: **Import Certificates**
   b. Select **Certificate Discovery Type**: **Import Certificates**.
   c. Select **MID Server selection method**: **Specific MID Server**.
   d. Select **MID Server**: Select the preferred MID Server.
   e. Add the credentials for the Windows or Linux host machine where the certificates are located.
   f. Fill out any other necessary fields.
   g. Select **Save**.
2. To add the Import SSL Certificate pattern, from the **Serverless Execution pattern** tab, select **New**.

Result

When your Discovery schedule runs, it then automatically scans your files.

Connect an intermediate certificate to its root certificate

Connect intermediate certificates to recently imported certificates or to root certificates located outside the servers to complete the certificate chain of trust.

Before you begin

If your root certificate is outside your server, you must first Discover root certificates hosted outside your server. If you need to import certificate files, Run Certificate Discovery via certificate file import.

> **Note:** If your root certificates are already in your servers, they're discovered and connected to the certificate chain using the standard Discovery probes deployed to your servers.

Role required: Admin, discovery_admin, Certificate administrator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Your intermediate certificate and your end-entity certificate are on your server. Discovery scans and establishes a certificate chain of trust for these certificates. If your root certificate is stored outside your server, perform a special Discovery on it and connect it to your intermediate certificate. Failure to connect might result in intermediate certificates being identified as the root certificate.

Procedure

1. Navigate to **System Properties** > **All Properties**.
2. Select **sn_disco_certmgmt.find_and_attach_root_certificate**.
3. Set the Value to **true** in the System Property record.
4. Save the record by selecting **Update**.
5. Rerun the Certificate discovery
   You can rerun the discovery using a port probe or a URL based Discovery.

Result

Your intermediate certificate identifies its root certificate, completing Discovery of your certificate chain of trust.

Maintaining certificate chain relationships

Maintaining certificate chain relationships via certificate import verifies the integrity and security of digital certificates, validating their authenticity in a system.

When you run certificate discovery by importing certificate files, the file extension and the order of the certificates within the file determine whether the system can build the chain relationships between the server, intermediate, and root certificates. This behavior applies only to certificates discovered through certificate file import. For information about this discovery method, see Run Certificate Discovery via certificate file import.

Only files with the `.txt` extension preserve certificate chain relationships. Certificates in a `.txt` chain file must appear in this order: Server certificate, intermediate certificate, and root certificate.The following outcomes apply when a file doesn't meet these requirements:

File contents	Result
Two or more certificates in a format, such as `.cert` or `.pem`	The system processes only the first certificate. It doesn't process the remaining certificates and doesn't build chain relationships.
A `.txt` file with one certificate	The system treats the certificate as a server certificate and doesn't build chain relationships.
A `.txt` file with two certificates	The system treats the first certificate as the server certificate and the second as the root certificate. The chain contains no intermediate certificate.

Note: Certificate chain relationships depend on the most recent URL or IP discovery run. If you import a file that doesn't contain chain relationships, the system removes any existing chain relationships for a certificate with the same certificate fingerprint.

Configure IDs for Certificate Management credentials

Ensure unique identification and effective management of credentials by configuring credential identifier for Certificate Inventory and Management.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: admin

About this task
External storage credential support enables TLS certificate discovery from CA types such as GoDaddy, DigiCert, and Sectigo. Set up the credential identifier in the instance for the respective Certificate Management credential type.

Procedure
1. Navigate to **All** > **Discovery** > **Credentials**.
2. Select **New**.
3. Select a credential type labeled as **Certificate Management Credential**.
4. In the Credential form, select the **External Credential Store** check box.
5. Fill in the **Credential Alias** and **Credential ID** fields.

Note:

Ensure that the Credential ID and Arg_ID in the Credential Resolver file match.

Based on the Certificate Authority you choose, the appropriate fields will appear requesting the specific information required by each CA. For additional information, refer to KB article for Digicert, Entrust, Let's encrypt, Microsoft CA.

6. Select **Submit**.

Run certificate discovery via certificate authority query

Run certificate discovery via certificate authority (CA) query to identify and import TLS certificates from specific CAs. Discover TLS certificates from CAs with Certificate Inventory and Management, using Patterns for diverse certificate authority vendors.

Before you begin
Role required: pki_admin or discovery_admin

About this task
In Certificate discovery, the CA pattern uses specific API elements. The user added to the instance credentials needs permissions for these queries. For more information on the Certificate Authorities (CA) pattern and the associated API elements, see Certificate authorities pattern API elements and permissions.

Procedure
1. Create a credential for the CA.
   a. Navigate to **All** > **Discovery** > **Credentials**.
   b. On the Credentials form, select **New**.
   c. Select **Certificate Management Credentials**.
   d. Select a CA in the CA type field.
   The available options are: DigiCert, GoDaddy, Entrust, Sectigo, Let's Encrypt, EJBCA, CyberArk Certificate Manager SaaS, Sectigo Universal, and Sectigo Public.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. Fill in the fields that are specific to the CA type you selected.

Field	Description
Name	Descriptive name for the credential.
Credential alias	Create or add a credential alias for the credential. The alias must be assigned to the corresponding credential record for the association to let Discovery scan the CA and allow automated certificate requests authenticate to it. For more information, see Credential aliases for Discovery.
API Key	API key generated from the CA for DigiCert, GoDaddy, and CyberArk Certificate Manager SaaS CA types.
Secret Key	API secret key generated from GoDaddy CA.
User name	User name associated with this credential for Entrust,
Authentication Key	Authentication key associated with this credential for Sectigo
Password	Password associated with this credential, it is stored in the database in encrypted form. This field is visible only for Entrust and Sectigo CA types.
Customer URI	Customer URI generated from the CA for Secitgo CA type.
Keystore	JKS or PKCS12 keystore to use when this protocol is selected for Entrust and EJBCA CA types.
Private Key	Private key associated with the CA for Let's Encrypt, EJBCA, Sectigo Universal, and Sectigo Public CA types.
Key Type	Encryption algorithm type for the private key. The available options are: None, RSA, and ECDSA. This field is visible only for Let's Encrypt, EJBCA, Sectigo Universal, and Sectigo Public CA types.
Key ID	Key ID for the ACME account for Sectigo Universal and Sectigo Public CA types.
Contacts	Email IDs to use when creating ACME account. The IDs must be comma separated. For example, `username1@domain.com,username2@domain.com`. This field is visible only for Let's Encrypt and EJBCA CA types.
MAC Key	MAC Key ID for the ACME account for Sectigo Universal and Sectigo Public CA types.
Active	Option to make the credential active.
ACME	Option to enable to Automated Certificate Management Environment (ACME) protocol. For more information, see Automated certificate management for TLS certificates. This check box is available for DigiCert, Entrust, Let's Encrypt, EJBCA, Sectigo Universal, and Sectigo Public CA types.
External Account Binding	Option to enable external account binding for EJBCA CA.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Create a Discovery schedule.
   For more information on setting up your Discovery schedules, see Schedule a horizontal discovery.

   a. Navigate to **All** > **Discovery** > **Discovery Schedules**.

   b. Select **New**.

   c. Select **Certificates** in the **Discovery** field.

   d. Select **CA Trust Discovery** in the **Certificate Discovery Type** field.

   e. Select **Auto-Select MID Server** or **Specific MID Server** in the **MID Server selection method** field.

   f. Fill out any other necessary fields.
   For more information, see Discovery Schedule form reference.

   g. Select **Submit**.

3. Open the Discovery schedule.
4. To add the CA pattern you need, select the **Serverless Execution pattern** tab and then select **New**.
   If you enable the **Include cert status** option, you can specify multiple certificate statuses by separating them with commas.
5. Select **Submit**.

Result
When your Discovery schedule runs, it automatically scans your files.

Discover root certificates hosted outside your server

Collect information about root certificates stored outside your server. Create a specialized Discovery schedule.

Before you begin
Role required: Admin, discovery_admin, Certificate administrator

About this task
To discover root certificates stored outside your server, create a special Discovery schedule to target them.

Procedure
1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. From the Discovery Schedule page, select **New**.
3. On the Discovery Schedule form, fill in the fields.
   For a description of the field values, see Discovery schedule form table.
4. Save this record by selecting **Update**.
5. Select the **Serverless Execution Patterns** tab.
6. Create a Serverless Execution Pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select **New**.

b. In the New Record page, select the search icon ( ) in the **Pattern** field.

c. Fill in the **Name** field.

d. Find and select **Import SSL Certificate**.

e. Select **Submit**.
You're redirected back to the Discovery Schedule page created in step 4.

7. Select the **Serverless Execution Patterns** tab.
8. Select the **Name** of the new pattern that you just made.
9. Select the **(empty)** value field, and add a value for each parameter.
For guidance on what values to set, see Run Certificate Discovery via certificate file import.
10. Collect information on your root certificate by making a unique certificate file.
    a. Navigate to the **Discovery Schedule** record that you created.
    b. Select the link **Discover now** to begin Discovery that collects your root certificate.
11. Navigate to **All** > **Unique Certificates**.
12. After the Discovery is complete, find your root certificate to confirm that it's there.

Result
Your Discovery collects information on your root certificate located outside a server.

What to do next
Connect an intermediate certificate to its root certificate to complete your certificate chain.

Customize Certificate Inventory and Management

Customize various aspects within Certificate Inventory and Management using Discovery properties, enabling you to tailor the solution to your specific needs and enhance the overall management of TLS certificates.

Before you begin
Role required: discovery_admin or admin

Procedure

To edit Discovery properties, navigate to **All** > **Discovery Definition** > **Properties**.

- glide.discovery.certs.cert_admin_user_id
- glide.discovery.certs.days_before_expiration_to_create_renewal_task
- glide.discovery.certs.enable_incident_creation_for_expired_certificates
- glide.discovery.certs.enable_renewal_task_creation_for_discovered_certificates
- glide.discovery.certs.slack_channel_id

Manual flow for certificate requests

Certificate renewal requests and incidents are automatically created when certificates are about to expire or have expired. You can also manually create requests using the Service Catalog.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In the manual flow, the request and its approval are recorded on the ServiceNow instance. However, the certificate itself is obtained outside the instance, directly with your certificate authority (CA). Use the manual flow when your CA is not covered in the automated flow. For the CAs that are supported, see Automated certificate management for TLS certificates.

A scheduled job examines the Unique Certificates [cmdb_ci_certificate] table on a daily basis and triggers renewal certificate tasks when a certificate is either expired or nearing expiration.

Things to consider for renewal of certificates:

- Tasks for certificate renewals are automatically generated 60 days prior to expiration. To modify the timeframe, update the Discovery property: *glide.discovery.certs.days_before_expiration_to_create_renewal_task*.

  Note: It's enough to configure the *glide.discovery.certs.days_before_expiration_to_create_renewal_task system property* for the automatic creation of manual certificate tasks. You don't need to custom configure the *CertificateNotifier* or other prefilters.

- Certificate tasks are automatically assigned a Priority 3 - Moderate. Modify the priority by opening the task if necessary.
- If a certificate task already exists for the current certificate, no additional tasks are created.
- For renewal certificate tasks and incidents, various fields are automatically pre-populated based on the certificate CI.

  Note: The fields **assigned_to** and **requested_for** in the renewal task are derived from the **assigned_to** field in the targeted certificate. Additionally, the **assignment_group** is carried over from the certificate to the renewal task.

- If a certificate expires and isn't replaced, a new incident is created. However, only one incident is generated for each expired certificate.
- To view the certificate incidents list, navigate to **Service Desk** > **Incidents**.

Create certificate requests

Create certificate requests manually through Service Catalog for approval when your certificate authority (CA) is not covered in the automated flow.

Before you begin

Ensure that the Certificate Management catalog is enabled.

Role required: pki_admin

Procedure

1. Navigate to **All** > **Service Catalog** > **Certificate Management**.
2. Select **Request New Certificate**.
3. Provide the necessary information for the following mandatory fields.
   - **Certificate Signing Request**
   - **Certificate Authority**
   - **Certificate Type**
   - **Validity period**
   - **Approver**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Enter or select additional details on the form as needed.
5. Select **Submit**.

Result
The Certificate task is generated and submitted for approval. The approval field is view-only for pki_user and editable for pki_admin.

Create renewal certificate requests

Manually initiate renewal requests for certificates using the Service Catalog for approval when your certificate authority (CA) is not covered in the automated flow.

Before you begin
Ensure that the Certificate Management catalog is enabled.

Role required: pki_admin

Procedure
1. Navigate to **All** > **Service Catalog** > **Certificate Management**.
2. Select **Renew Certificate**.
3. Select a Certificate CI.
4. Provide the necessary information for the following mandatory fields.
   - **Certificate Signing Request**
   - **Requested for**
   - **Approver**
5. Enter or select additional details on the form as needed.
6. Select **Submit**.

Result
The Certificate task is generated and submitted for approval. The approval field is view-only for pki_user and editable for pki_admin.

Note: The notification process for certificate renewal exclusively supports manual flow and does not involve automated certificate management. For CSR information, you must use automated certificate management.

When monitoring the status of a certificate order, if there's a configuration-related issue, the task may get stuck as a work in progress. The PKI Admin is responsible for resolving such issues. For example, if the MID Server is down during the scheduled job tracking the certificate order status, the system waits for the next scheduled job instead of logging an error. If the MID Server is brought back up before the subsequent job, the certificate retrieval proceeds accordingly.

Archive certificate tasks

From Certificate Inventory and Management version 1.1.7, the Data Archiver [com.glide.auxdb] plugin performs daily operations to transfer data that is no longer required from primary tables to a designated set of archive tables.

Before you begin
Role required: pki_user or pki_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Certificate Inventory and Management actively checks and archives certificate tasks that surpass one year in age. To facilitate this, the corresponding plugin must be enabled and activated for the scheduled archival tasks to execute. It's important to note that the Certificate Task Archive rule, which is included with the Certificate Inventory and Management application, is inactive by default.

Procedure

1. Activate **Certificate Task Rule**.
2. Navigate to **Archive Rules** > **Certificate Task Archive**.
3. Select the **Active** check box.

Result

After activation, this rule archives data from the Certificate Task table [sn_disco_certmgmt_certificate_task] under the following conditions:

- State is Closed Complete
- Last update occurred more than twelve months ago

Automated certificate management for TLS certificates

From Certificate Inventory and Management version 1.3.8, you can automate the request flow for new certificates, renewals, and revoking certificates.

Certificate Inventory and Management automatically fetches certificates from Certificate Authorities (CAs) without requiring manual intervention from the PKI team. Starting in Version 2.1.0, this feature supports DigiCert and Entrust CA Gateway for seamless automatic fulfillment flows, with the limitation that only OV DigiCert certificates can be requested. Version 2.3.2 introduces support for the Microsoft CA. For automated flows with DigiCert or Entrust CA Gateway in Certificate Inventory and Management, you must have permissions to request, renew, and revoke certificates.

For provider-specific requirements, see the following documentation:

- Microsoft CA: For information about configuring and administering a Microsoft certification authority, go to Microsoft Learn and search for the "Active Directory Certificate Services documentation" article.
- DigiCert: For information about API access and certificate order management, go to the DigiCert developer portal and search for the "CertCentral APIs" documentation.
- Entrust: Refer to your Entrust CA Gateway documentation.

For automated flows with DigiCert or Entrust CA Gateway in Certificate Inventory and Management, you must have permissions to request, renew, and revoke certificates.

The Microsoft CA user requires the following permissions:

Microsoft Gateway user permissions

Permission	Action
CredSSP on CA, intermediate server, and MID Server	Set up CredSSP on CA, intermediate server, and MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Gateway user permissions (continued)

Permission	Action
	For CredSSP configuration steps, see the Now Support Knowledge Base documented in the KB article KB1632624.
Membership in Enterprise Admins	Ensure the user holds membership in the Enterprise Admins group.
Security Group Inclusion for Template	Ensure the user is included in the Security Group of the template.
Specific Permissions in CA	Grant the user permissions: Read, Issue and Manage Certificates, Manage CA, and Request Certificates in the CA.

Using automated flow for certificate management

Certificate and Management streamlines your TLS certificate processes, offering benefits such as improved efficiency and enhanced security. Automating certificate management ensures timely renewal of certificates, which minimizes the risk of expired certificates.

Before you begin

To use the Microsoft Certificate Authority automated flow, you must install the **ServiceNow IntegrationHub Action Step - PowerShell** plugin and must have an Integration Hub subscription. See Integration Hub usage and subscription for more information.

Role required: pki_admin or admin

Procedure

1. Set the system property *sn_disco_certmgmt.cert_task_default_approval_group* to the default Approval Group name.
   If the certificate request moves into manual mode, the approval group name is the default group used. For instance, the default group is used if there's no matching policy or more than two matching policies. You can add more than one approval group, separated by commas. The first group on the list, which belongs to the task domain, is used for approval. If no domain-specific group is found, the first name in the global domain list is used.
2. To set the validity period of the certificate order, update the system property *sn_disco_certmgmt.default_cert_order_validity_period*.
   The default is 730 days (2 years).
3. Add the IP of the Microsoft CA Server.
   - Add the *ca_host_ip* field of the routing policy.
   - Add the IP of an intermediate server in *ca_host_ip* field of the routing policy.

   The intermediate server can be any Windows server in the same domain as the Microsoft CA Server. It must have access to the certutil and certreq commands available on Powershell. When an intermediate server is used, the MID Server executes a Powershell script on the intermediate server using Invoke-Command. This command uses a Remote Procedure Call (RPC) to run the certutil and certreq commands on the CA Server.
4. Create the certificate credential and map it to the credential alias.
   Each credential should map using a unique credential alias. For more information, see Credential alias for Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Confirm the Certificate and Certificate URL information are in the Certificate Authority [sn_disco_certmgmt_ca] and Certificate Authority API URL [sn_disco_certmgmt_ca_api_url] tables.
   The default URL for DigiCert and Entrust CA Gateway provide all validation type URLs. You can also add additional URLs.
6. Set the task priority.

   The change requests priority and type are mapped based on the priority of the task. Change requests have the same priority as a task priority, except a change request doesn't have P5 so in this case it's mapped to P4.

   To change the type of change requests, the change management property *com.snc.change_management.change_model.type_compatibility* must be set to true. The default is False.

   a. If needed, set the task and change the system property *sn_disco_certmgmt.default_cert_task_priority* to configure New and Renew task priorities.
      The priority defaults to P3. The possible values are 1, 2, 3, 4, 5. If the value is 1, the priority sets to P1, and so on. If any invalid value is provided, the priority resets to the default of P3.

   b. If needed, set the task and change the system property *sn_disco_certmgmt.default_revoke_cert_task_priority* to configure Revoke task priorities.
      The priority defaults to P1. The possible values are 1, 2, 3, 4, 5. If the value is 1, the priority sets to P1 and so on. If any invalid value is provided, the priority resets to the default of P1.

7. **Optional:** Install the Integration hub plugin [com.glide.hub.integrations].

   (Optional) The [com.glide.hub.integrations] plugin isn't required for requesting the DigiCert or Entrust CA Gateway Certificate and tracking the certificate order status. However, if you want to debug the certificate subflow actions or add your own customization flow for DigiCert or Entrust CA Gateway, install this plugin.

Set up routing policy for automated certificate management

Set up a routing policy to automate your Certificate Inventory and Management. Creating a policy based on Certificate Authority (CA), environment, and other features confirm efficient TLS certificate management.

Before you begin

Role required: pki_admin or admin

About this task

Use the routing policy to determine which CA must be contacted for certificate operations. It contains the CA, CA URL, Credential, Approval Group, Assignment Group, and CSR attributes. The routing policy triggers the flow for requesting certificates for specific CAs.

A certificate request is considered a duplicate if there is another certificate task with the same domain name that is still in progress. Duplicate certificate requests aren't allowed. However, you can override this setting by checking the **Allow duplicate requests** check box.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- Approvals are only supported in the Fulfiller approval experience at this time. See a table of fields that go into the routing policies at, Certificate routing policy form for EJBCA ACME.
- When the CA's certificates needs to be picked up by Discovery, the same credential alias must be set as the `CredentialAlias` argument on the corresponding CA discovery pattern. Both associations use the identical alias value, one enables authentication for discovery, the other for automated certificate requests.

Procedure

1. Navigate to **All** > **Certificate Management** > **Certificate Routing Policies**.
2. Select **New** and fill in the required fields on the form.
   While requests for new certificates and certificate renewal can be automated, many PKI teams prefer human validation before fulfillment. If so, select the **Approval required** check box.

   Note:

   Organization, Organizational Unit, Locality, State, Country, and Email accept comma-separated values. * will be considered as any. Subject common name and Subject alternative name are supported with RegEx. The RegEx format has the following restrictions:

   - It should not contain commas.
   - It should not start and end with a forward slash (/) and * matches any.
   - For more inforamtion about the fields and values on a routing policy form, see .

3. The system matches the following CSR attributes with the entries in the Routing Policy [sn_disco_certmgmt_routing_policy] table:
   - Organization
   - Organizational Unit
   - Locality
   - State
   - Country
   - Email
   - Environment
   - Certificate Purpose(internal/external)
   - Subject common name
   - Subject alternative name

   Note: Entrust CA Gateway routing policies also include these fields: Certificate Authority Identifier, Certificate Profile, and Certificate Format. For Microsoft CA also use these fields: Certificate Authority, CA template name, CA Host IP, Credential, and CSR attributes. For DigiCert, the routing policy also requires a Certificate Authority API URL field to handle automated processes and revocation flows.

4. The following options may occur.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
If a single routing policy matches	Verify the following conditions: ○ Validate the subject common name using the RegEx pattern provided in the Routing Policy table, domain name, or *. ○ Check for certificate request validity period not greater than the maximum validity peri od in the Routing Policy table. ○ Check for duplicate Certificate Request is allowed flag in the Routing Policy table.
If multiple routing policies are eligible	The task is assigned to the default approver group.
If there is no routing policy found	The task is assigned to the default approver group.
If single policy matches and approval need ed flag is true	The task is assigned to the task approval group defined in the routing policy.

Result

The approval group is assigned to the routing policy and contains the role: pki_approver and at least one of the active group members available in that group. If the routing policy requires manual approval, then approval is requested from those in the approval group.

What to do next

The following knowledge base articles guide you through the process of producing the credentials required and configuring routing policies for different certificate authorities:

For Digicert, see [Digicert] Configure automated certificate management for TLS certificates [KB2166364].

For Entrust, see [Entrust] Configure automated certificate management for TLS certificates [KB2173533].

For Let's Encrypt, see [Let's Encrypt - ACME] Configure automated certificate management for TLS certificates [KB2197962].

For Microsoft CA, see [Microsoft CA] Configure automated certificate management for TLS certificates [KB2198094].

Routing policy details

Routing policies define specific criteria, such as CA and environment, for automated certificate management in TLS certificates.

There are three different scenarios for how routing policies work for automated certificate management for TLS certificates.

- Scenario 1: A single routing policy matches and manual approval is required. The approver can verify the task details and approve the certificate request. Once approved, this triggers the automated flow.
- Scenario 2: No routing policy matches OR there are multiple routing policies that match. The approver can verify the task details and select Choose Routing Policy & Approve. The approver

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

needs to manually select the routing policy. The updated routing policy then displays. The approver selects the routing policy which then triggers the automated flow.

- Scenario 3: A single routing policy matches, manual approval is not required, and Approval Required flag is inactive. This triggers the automatic flow.

Note:

If the approver has rejected the certificate request task, the task is marked as Failed. A single approval is enough for triggering the automated flow progress.

Request certificates using automated flow

Request a new certificate and automatically retrieve the certificates for an application using automated certificate management. Certificate Inventory and Management Version 2.1.0 supports requesting certificates from DigiCert and Entrust CA Gateway. Version 2.3.2 also supports Microsoft CA.

Before you begin

- Verify the certificate management catalog is enabled and that a routing policy is created.
- For DigiCert only: Get a domain validated API key from DigiCert. If you submit a certificate request with a new domain that is not validated by DigiCert, the request status is displayed as Pending. When the automated flow can't fetch the certificate information, the request status is marked as Failed.
- Role required: Certificate requester [sn_disco_certmgmt.certificate_requester], PKI admin [sn_disco_certmgmt.pki_admin], PKI user [sn_disco_certmgmt.pki_user], or admin.

Note: Approvals are only supported in the Fulfiller approval experience at this time.

Procedure

1. Navigate to **All** > **Service Catalog** > **Certificate Management**.
2. Select **Request New Certificate - Automated flow**.
3. Provide details for the mandatory fields: **CSR** and **Validity Period**.
4. Fill in or choose additional information on the form, and then click **Submit** to proceed with the order.

   The Routing Policy [sn_disco_certmgmt_routing_policy] table helps to fetch the CA routing policy ID. If a single routing policy ID is not returned, then the approver must select the routing policy and approve the task. For more information, see Approve certificate tasks.

   This creates a certificate task which then triggers the automated flow. In the routing policy, if Approval Required field was checked, the task requires approval before the automated flow begins.

Result

- After the request is submitted, the automated flow makes the request to the CA to get the certificate.

  Note: Powershell Step is used for Microsoft CA. This requires the plugin: com.glide.hub.action_step.powershell.

- A record is created in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table once the certificate is successfully fetched.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Every 30 minutes, the following scheduled job runs and checks for status: DigiCert – Track Certificate Order Status.

  Note: There are no scheduled jobs for Entrust CA Gateway and Microsoft CA.

- If the certificate is available, it is attached to the certificate task.
- The certificate task is then marked Completed and a change request is created.
- If multiple tasks are created for the same CSR and Allow Duplication is not checked in the routing policy, the task fails.
- If the validity period exceeds the matched routing policy validity period, the task fails.

Automatically generate a Certificate Signing Request

You can generate a Certificate Signing Request (CSR) on the instance by providing the required information on the catalog form.

Creating a CSR manually can be a cumbersome and error-prone process. Automatic CSR generation can be enabled on the **Request New Certificate** and **Request New Ceriticate (Automated)** forms. This feature is disabled by default for security.

To enable automatic CSR generation, set the system property **sn_disco_certmgmt.enable_csr_generation** to true. This property also controls whether the PKCS12 certificate format is available, because a PKCS#12 key store requires a private key that is generated with the CSR. See Discovery properties for more information.

Once enabled, a **Want to generate CSR?** checkbox is displayed in the certificate request form. Selecting the checkbox enables the required input fields to generate the CSR. Provide the necessary information and select **Generate**. The input data is validated and the CSR is formated accordingly.

The CSR is generated in the **Certificate Signing Request (CSR)** field, and a private key is generated in the **Private Key** field. Select the **Download** link to save the private key for later use. By default, the private key size is 2048 but can be changed using the system property **sn_disco_certmgmt.private_key_size**. See Discovery properties for more information.

When selecting **Submit**, the form provides a warning to download and save the private key. For security, ServiceNow does not store private keys, so the private key must be saved before leaving the form or else it is lost.

Approve certificate tasks

Sometimes you need to manually approve a certificate task to ensure the validity and security of TLS certificates.

Before you begin

Role required: admin

About this task

If the approval is requested for a routing policy, select the routing policy created. Approvals are only supported in the Fulfiller approval experience at this time.

Procedure

1. Navigate to **All** > **Self-Service** > **My Approval**.
2. Open the automatically created approval task.
3. Check the new certificate details and **Choose Routing Policy & Approve**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

1. The automated flow for requesting a new certificate is initiated.
2. You can open the certificate task by navigating to **Certificate Management** > **New Certificate Task**.
3. Ensure that the automated flow is triggered.
4. If the CA approves the certificate request, both the Order ID and Certificate ID are retrieved for the newly requested certificate. If the CA does not approve, only the Order ID is fetched.

Note: For Entrust CA Gateway, Certificate Serial Number and Enrollment Id are fetched.

Renew certificates using automated certificate management

Request renewal for a certificate and automatically retrieve the certificate for an application. This maintains secure and uninterrupted services as you extend the validity period of the certificate, preventing potential service disruptions due to expired certificates.

Before you begin

Set up routing policy for automated certificate management to ensure the Certificate Management catalog is enabled.

To renew an existing certificate, make sure the certificate has its extension details populated in the [sn_disco_certmgmt_certificate_extension] table. Currently, these details populate automatically by the Discovery DigiCert CA Pattern or for any certificates created through the automated flow. However, for Entrust certificates, this only works if you Request certificates using automated flow.

Role required: PKI Admin, Admin, Certificate Owner, or Users that are part of the Certificate Owner Group.

The Certificate Owner and Certificate Owner group contains the Certificate Requester role (minimal role).

Note: Approvals are only supported in the Fulfiller approval experience at this time.

About this task

To renew an existing certificate, request a new certificate with the same details as the original certificate. Currently, there is no Renew API available for Entrust CA Gateway and Microsoft CA certificates. During a renewal request, a new certificate is generated internally with the same attributes as the selected certificate.

For renewing an existing certificate, CSR is mandatory. The requester can use an existing CSR if available or use a new CSR. If you want to use an existing CSR, use the same CSR to request a new certificate from the CA. If fields are entered using vault and java APIs, the CSR is generated.

Procedure

1. Navigate to **All** > **Service Catalog** > **Certificate Management**.
2. Click **Renew Certificate – Automated flow**.
3. Provide details for the mandatory fields: **CSR** and **Validity Period**.
4. Fill in or choose additional information on the form, and then click **Submit** to proceed with the order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

1. The Routing Policy [sn_disco_certmgmt_routing_policy] table helps fetch the CA routing policy ID.
   - If a single routing policy isn't matched, the approver must select the CA and trigger the flow.
   - If the CSR contains a different domain name than the issued certificate domain name, the task requests approval.
   - If a single routing policy is matched but renewal certificate information isn't available in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table, the task requires approval.
   - Certificates can't be renewed if the Certificate Authority and Order id or thumbprint details are missing for the certificate in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table. Discover the certificate via Certificate Authority query to populate the required details in the Certificate Extension table. After Discovery, select the routing policy and approve the task.
2. This creates a task for the ordered certificate which triggers the flow to request renewal certificates.
3. Once the request is submitted, the automated flow makes the request to the CA to get the certificate.

   Note: Powershell Step is used for Microsoft CA. This requires the plugin: com.glide.hub.action_step.powershell.

4. The Order id is stored in the Certificate Task [sn_disco_certmgmt_certificate_task] table and Certificate Extension [sn_disco_certmgmt_certificate_extension] table.

   Note: For Entrust CA Gateway, Certificate Serial Number and Enrollment Id are fetched. The Serial Number is stored in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table.

5. Every 30 minutes, the “DigiCert – Track Certificate Order Status” scheduled job runs and checks for status.

Note:

The system fetches details about the selected certificate from the Certificate Extension [sn_disco_certmgmt_certificate_extension] table and makes the request to the CA to renew the certificate. Certificates can't be renewed if the Certificate Authority, Order id, or thumbprint are missing from this table. If additional details are missing for renewing the certificate, the system logs a message and suggests what to do. In this case, you must discover the certificate using CA-based discovery. For more information, see Run Certificate Discovery via Certificate Authority query to populate these details to the Certificate Extension table.

Revoke certificates using automated certificate management

Revoke certificate for an application. Revocation doesn't require approval if *`order Id`* and *`certificate Id`* are present in the Certificate Extension table. If *`order Id`* and *`certificate Id`* aren't present in the Certificate Extension table, then you need approval.

Before you begin

Set up routing policy for automated certificate management to ensure the Certificate Management catalog is enabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To revoke an existing certificate, make sure the certificate has its extension details populated in the [sn_disco_certmgmt_certificate_extension] table. You don't need to include an endpoint URL. The system will automatically hardcode the revoke URLs for the discovery. Currently, these details populate automatically by the Discovery DigiCert CA Pattern or for any certificates created through the automated flow. However, for Entrust certificates, this only works if you Request certificates using automated flow.

Role required: pki_admin or admin

Note: Currently, approvals are only supported in the Fulfiller approval experience.

Procedure

1. Navigate to **All** > **Service Catalog** > **Certificate Management**.
2. Select **Revoke Certificate – Automated flow**.
3. Enter the Issued Certificate which you want to revoke (for instance: www.undefined.com). Multiple Certificates can be selected for revoking.
4. Provide an appropriate reason for revoking the certificate.
   For Microsoft CA the reason should be an integer. If any other value is given, a default value of 0 will be used which means unspecified.
5. To place the revoke order, select **Submit**.
6. In the confirmation pop-up, select **OK**.

Result

1. A task is automatically created when you request a revocation.
   - If *order Id* and *certificate Id* are present in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table, revocation does not require approval.
   - If *order Id* and *certificate Id* aren't present in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table, then the task requests approval.
   - If the serial number for Entrust CA Gateway isn't present in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table, then the task requests approval.
2. Once the PKI team provides approval, the mapping between the certificate and CA occurs based on the Routing policy selected.
3. This triggers the revocation operation for the CA selected which uses the CA APIs.
4. Details are stored in the Certificate Extension table.
5. Every 30 minutes, the following scheduled job runs and checks for status: DigiCert – Track Certificate Order Status.

   Note: There are no scheduled jobs for Entrust CA Gateway and Microsoft CA.

6. The status of the certificate is marked as revoked.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

Certificates can't be revoked if Certificate Authority or Certificate Id details are missing in the Certificate Extension [sn_disco_certmgmt_certificate_extension] table. For Entrust CA Gateway, certificates can't be revoked if the Serial number is missing. Discover the certificate via Certificate Authority query to populate the required details in the Certificate Extension table. After that, Discovery selects the routing policy and approves the task.

Revoke certificate API request. If "skip_approval" is true, the revoke process is completed faster. If "skip_approval" is false, the revoke process is completed when the DigiCert or Entrust CA Gateway admin has approved or rejected the revoke request. To skip the approval step, the API key must have admin privileges.

Certificate generation through Kubernetes cert-manager

Request a certificate through Kubernetes cert-manager using the ServiceNow External Issuer (sn-external-issuer) and save the certificate and its related information securely within the Kubernetes cluster as a secret. In Kubernetes, a secret is an object that allows you to store and manage sensitive information, such as passwords, API keys, and certificates.

For information on building an external issuer, see Building and Deploying External Issuer For Certificate Management [KB1435392].

Deployment Requirements

- Deploy cert-manager in your Kubernetes environment. Update **manager.yaml** with Instance URL, Certificate Owner Group, Certificate Owner, Environment, and Renewal Tracking.
- Deploy the ServiceNow External Issuer (sn-external-issuer) in your Kubernetes environment. Create a Kubernetes secret *`clusterissuer-servicenow-credentials`* with the instance username and password, ensuring the user has the necessary roles to request a certificate.
- The ServiceNow External Issuer supports amd64 architecture along with the latest version of Kubernetes, 1.33.3.

ServiceNow External Issuer (sn-external-issuer)

- External issuers expand cert-manager functionality to issue certificates through non-core APIs and services.
- The ServiceNow External Issuer is a ServiceNow-specific implementation of an External Issuer.
- When a new certificate task is created, its Certificate Request UID and Certificate Task Sys Id are stored in the local JSON cache and the Certificate Request UID to Task Map table on the instance.
- The ServiceNow External Issuer polls the instance to monitor the status of the certificate task.
- If the certificate task is in the **Work in progress** state, its Certificate Request UID and Certificate Task Sys Id are added to the External Issuer UID Map table on the instance and the local JSON cache. During this time, Cert-manager automatically attempts to request the certificate.
- On receiving a certificate request, Cert-manager checks for a matching task in the local JSON cache. If found, it polls the same task; otherwise, it queries the instance for records from the External Issuer UID Map table and populates the local JSON cache.
- After the task is marked as complete and the certificate is generated, the ServiceNow External Issuer sends another request to the instance, downloads the certificate attachment, and updates the certificate resource and corresponding secret in Kubernetes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Deploying the ServiceNow External Issuer in Kubernetes

Deploying the ServiceNow External Issuer in Kubernetes involves the following steps:

1. From the ServiceNow instance download page, obtain the Helm Chart or YAML zip package.
2. Customize the `manager.yaml` or `values.yaml` files as needed for your specific use case. These files may include essential information such as the Instance URL and Certificate Owner Group.
3. Create a Kubernetes secret named `clusterissuer-servicenow-credentials` with the instance username and password.

   Example command:

   - Create a Kubernetes secret named `clusterissuer-servicenow-credentials` with the instance username and password. Example command:

     ```
     kubectl create secret generic
      clusterissuer-servicenow-credentials
                  --from-literal=user=<user_name>
      --from-literal=password=<password> -n
                system
     ```

   - Ensure that the user has the necessary roles to request certificates.
4. Execute the following commands for deployment.

   ```
   kubectl create ns system
   kubectl apply -f crd
   kubectl apply -f rbac
   kubectl apply -f issuers
   kubectl apply -f manager/manager.yaml
   ```

5. (Optional) Customize any additional configurations in the files to suit your specific requirements.
6. Ensure that the deployment is successful and the ServiceNow External Issuer is up and running.

Request a certificate flow

After deployment, submit a certificate resource with the following information in a file named `certificate_clusterissuer.yaml`.

- issuerRef : clusterissuer-servicenow
- issuer : issuer-servicenow
- kind : ClusterIssuer
- issuerRef : servicenow-issuer.servicenow.com

Here's a sample Certificate Resource:

```
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
name: certificate-by-clusterissuer
spec:
commonName: certificate-by-clusterissuer.servicenow.com
secretName: certificate-by-clusterissuer
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
dnsNames:
- servicenow.com
- foo.servicenow.com
issuerRef:
name: clusterissuer-servicenow
group: servicenow-issuer.servicenow.com
kind: ClusterIssuer
```

Apply the certificate resource using `kubectl apply -f certificate_clusterissuer.yaml`

Certificate alerts and notifications

From version 1.2.0 onward, Certificate Inventory and Management integrates with Event Management, providing a unified solution for streamlined operations. This integration enables the generation of events and notifications, with the added capability of sending alerts to Slack. Certificate Inventory and Management also supports sending certificate notifications via email and Microsoft Teams.

Manage events and alerts with Event Management

In Certificate Inventory and Management version 1.2.0 and later, Event Management can create events and alerts for both expiring and expired certificates, enhancing the system's monitoring capabilities.

Before you begin

Ensure the Event Management Store application is installed.

Role required: admin

About this task

The integration leverages the Workflow Studio feature to initiate events. It utilizes a pre-existing daily Discovery schedule designed to identify expired or expiring certificates, activating a subflow. Upon subflow activation, it verifies the installation of the Event Management Store application. If the application is installed, the subflow generates an event for expiring or expired certificates.

Procedure

1. To view all the generated alerts or events, navigate to **All** > **Event Management** > **All Alerts** or **All** > **Event Management** > **All Events**.
2. To exclude specific Certificate events and prevent alert triggers, navigate to **Event Management** > **Rules** > **Event Rules** and perform the following steps.
   a. Select the **Certificate Expired/Expiring Events** rule.
   b. Select the **Event Filter** tab.
   c. Select **Ignore events that match this filter**.
   d. Add any other necessary details as needed.
   e. Select **Save**.

Receive certificate notifications on Slack

In Version 1.2.0 of Certificate Inventory and Management, set up Slack notifications to stay informed about expiring and expired certificates. Receive detailed alerts with convenient URL links for additional information.

Before you begin

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

A scheduled job in Certificate Inventory and Management regularly checks the expiry dates of certificates stored in the valid_to field of the cmdb_ci_certificate table. When certificates are either expired or nearing expiration, the job initiates the subflow Process Slack Message, which subsequently sends notifications to the designated Slack channel.

Procedure

1. Set up Slack spoke.
2. Create a Slack bot and add it to a Slack channel.
3. Configure default Slack credentials.
4. Configure custom Slack credentials.

   If your Slack spoke version is 1.2.0, you can bypass the next step concerning the retrieval of the Channel ID.

5. Get the ID for your Slack channel.
   a. Go to **Slack**.
   b. Right-click the Channel.
   c. Select **Additional options** > **Copy link**.
   d. From the provided link, copy the string following <archive/>.
6. Configure the Slack Channel ID property.
   a. Navigate to **Discovery** > **Properties**.
   b. Enter value in the Slack Channel ID property *glide.discovery.certs.slack_channel_id*.
7. Select **Save**.

Result

Notifications are directed to the Slack channel specified in Discovery properties. If a certificate is assigned to a user, and their configured email matches the Slack email, the user is tagged in the Slack message. The Slack scheduled job runs weekly, and you can adjust its frequency by modifying the Certificate Notification scheduled job settings.

Receive certificate notifications via email

Configure certificate notifications to be delivered via email to relevant recipients.

Before you begin

Role required: Certificate administrator (glide.discovery.certs.cert_admin_user_id) or PKI Admin (pki_admin).

About this task

Certificate Inventory and Management sends email notifications for certificate life-cycle events such as certificate expiration, task creation, renewal completion, and auto-renewal failures.

Procedure

1. Configure the email address for users.
   a. Navigate to **All** > **Organization** > **Users**.
   b. Select a relevant user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. In the **Email** field, enter the email address.

d. Select **Update**.

2. Enable email sending configuration.

a. Navigate to **All** > **System Properties** > **Email Properties**.

b. In the **Outbound Email Configuration** section, under **Email sending enabled**, select the **Yes/No** check box.

c. Select **Save**.

3. Customize which certificate events trigger email notifications.

a. Navigate to **Workspaces** > **Certificate Management**.

b. Select **Settings** in the menu bar and then select the **Notification** tab.
The available notification types in the **Certificate notification policy** [sn_disco_certmgmt_notification_policy] table are:

- Auto-Renewal Failed notification
- Certificate Completed Failed notification
- Certificate Expired notification
- Certificate Expiring notification
- Certificate Renewal Task Created notification
- Certificate Task Creation notification

c. Select the Edit form icon for the relevant notification type.

d. Verify the **Email enabled** check box is selected.

e. Select the **Email recipients** field and then select the email recipient you want to notify.

f. Select the **Email recipient Groups** field and then select the email recipient group you want to notify.

g. In **Email recipient fields**, enter comma-separated field names from the certificate task record to notify users or groups assigned to those fields.
For example, to send notifications to the user and group assigned to the certificate record, you would enter `assigned_to,assignment_group`.

h. Select **Update**.

Result

Certificate notifications are sent to relevant recipients as configured for each notification type.

Renew certificates via email

Certificate renewal notifications are sent to users configured in the **Certificate notification policy** table via email 60 days before a certificate expires or after a certificate expires.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- You must have certificate CSR information to paste or in a .csr or .pem format. If you don't have one, you can instruct the system to generate a new one or use a previous one.
- Role required: pki_admin or admin

Procedure

1. Open the email and select Renew through email.
   Opening the link triggers a new email. It's the request to renew the certificate.
2. Open the renewal email.
3. Add the CSR in one of the following ways:
   - Attach a CSR file in .csr or .pem format.
   - Paste the CSR directly.
   - To generate a new CSR, replace the text in the subject line called <Renewal Option> with the text `AUTO - CSR`.

     Note: This generates a new CSR and the private key is stored securely in your vault.

   - To use the previous CSR, replace the text in the subject line called <Renewal Option> with the text `RP - CSR`.

     Note: The system reuses the CSR from the previous certificate request.
4. Select **Send**.
   Your manager receives an email notification for final approval or rejection of the certificate renewal.
5. To approve the certificate renewal request, select **Click here to approve** in the email.

Result
The certificate renewal is complete.

Receive certificate notifications via Microsoft Teams

Configure certificate notifications to be delivered to a Microsoft Teams channel so that you can receive alerts and initiate certificate renewal workflows directly from Microsoft Teams.

Before you begin
Verify that the Microsoft Teams spoke is installed and configured. For more information, see Set up the .

Role required: pki_admin or admin.

About this task
The Microsoft Teams integration uses the ServiceNow Microsoft Teams spoke to deliver notifications. A scheduled job in Certificate Inventory and Management regularly checks the expiry dates of certificates stored in the *valid_to* field of the Certificate [cmdb_ci_certificate] table. When certificates expire or are nearing expiration, the system triggers a subflow that sends notifications to the configured Microsoft Teams channel.

Procedure

1. Navigate to **Workspaces** > **Certificate Management**.
2. Select **Settings** from the menu bar and then select the **MS Teams Channel** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Enter the details of the Microsoft Teams channel.

Field	Description
Channel alias	Alias for the channel.
Microsoft teams channel link	URL to the Microsoft Teams channel.
Microsoft team name	Name of the team the Microsoft Teams channel is associated with.
Microsoft channel name	Name of the Microsoft Teams channel.

4. Select **Add**.
   The Microsoft Teams channel is added to the Certificate Microsoft Teams channel [sn_disco_certmgmt_microsoft_teams_channel] table.
5. Configure the type of notifications the Microsoft Teams channel receives.

   a. Select the **Notification** tab.
      The available notification types in the **Certificate notification policy** [sn_disco_certmgmt_notification_policy] table are:
      - Auto-Renewal Failed notification
      - Certificate Completed Failed notification
      - Certificate Expired notification
      - Certificate Expiring notification
      - Certificate Renewal Task Created notification
      - Certificate Task Creation notification

   b. Select the Edit form icon for the relevant notification type.

   c. Enable Microsoft Teams channel notifications for the notification type by selecting the **MS Teams enabled** check box.

   d. Select the **MS Teams channel** field and then select the relevant channel you want to notify.

   e. Select **Update**.

Result

- Certificate notifications are sent to the relevant Microsoft Teams channel as configured for each notification type.
- When renewing certificates from Microsoft Teams, you can choose to attach a CSR, generate a new CSR (the private keys are stored securely in your vault), or reuse the previous CSR.

Automated certificate management with ACME

The Automated Certificate Management Environment (ACME) is a communication protocol that automates the interaction between a certificate authority (CA) and a server. It streamlines the processes of requesting, renewing, and revoking SSL/TLS certificates.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACME overview

ACME uses JSON-formatted messages transmitted over a secure HTTPS connection. This communication enables for automated certificate life-cycle management, reducing manual intervention and the risk of errors.

ACME is a widely adopted standard used by public key infrastructure admins, line of business owners, certificate management users, admins, account owners, and team managers.

Certificate Inventory and Management supports the following ACME CAs: DigiCert, Entrust, Let's Encrypt, EJBCA, Sectigo Universal, and Sectigo Public. The ACME framework is extensible to any ACME-compatible certificate authority. Admins can add a CA by creating a record in the Certificate Authority [sn_disco_certmgmt_ca] table. For more information, see Add ACME-compatible certificate authorities.

ACME benefits

During requesting, renewing, or revoking SSL/TLS certificates, ACME offers significant benefits:

- It automates tasks associated with certificate management, reducing the administrative burden for IT teams.
- It helps verify that the certificates are valid and up-to-date.
- It eliminates the need for manual certificate management tasks, reducing overall costs.

Add ACME-compatible certificate authorities

Add a certificate authority (CA) that supports the ACME protocol to enable automated certificate request, renewal, and revocation flows.

Before you begin

Role required: admin

About this task

Certificate Inventory and Management supports the following ACME CAs: DigiCert, Entrust, Let's Encrypt, EJBCA, Sectigo Universal, and Sectigo Public. If your organization uses a different CA that is compatible with the ACME protocol, you can add it by creating a record in the Certificate Authority [sn_disco_certmgmt_ca] table. No script changes are required.

Note: Errors can occur when using custom ACME-compatible CAs. Contact ServiceNow Support for assistance.

Procedure

1. Navigate to **All** > **Certificate Management** > **Certificate Automate Flows** > **Certificate Authorities**.
2. Select **New**.
3. Fill in the fields on the **Certificate Authority** form.

Field	Description
Name	Name of the CA.
Base API URL	Base URL of the CA's ACME directory endpoint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
ACME	Select this check box to enable the ACME flow for this CA.

4. Select **Submit**.

What to do next

- Create a certificate management credential for the new CA, see Create credentials for ACME certificate authority.
- Set up a routing policy for the new CA, see Set up routing policies for ACME.

Create credentials for ACME certificate authority

Create credentials so Certificate Inventory and Management can communicate with your ACME certificate authority (CA) for automated certificate life-cycle management.

Before you begin

Role required: pki_admin or admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Select **New**.
   a. From the credentials list, select **Certificate Management Credentials**.
   b. On the form, fill in the fields.

Certificate management credentials form

Field	Description
Name	Descriptive name of the credential.
Credential alias	Credential alias linked to the credential.
CA Type	Type of the CA. The available options are: ▪ DigiCert ▪ Entrust ▪ Let's Encrypt ▪ EJBCA ▪ Sectigo Universal ▪ Sectigo Public **Note:** If you added an ACME-compatible CA, it is listed here as a CA type. For more information, see Add ACME-compatible certificate authorities.
ACME	Option to enable ACME flow the CA. For Let's Encrypt, EJBCA, Sectigo Universal,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	and Sectigo Public CAs, this check box is selected by default.
Private Key	Create any private key or generate an account using the ACME CA. Provide a private key with the key type RSA or ECDSA.
Key Type	Encryption algorithm for the private key (for example, RSA or ECDSA).
Key ID	Used for account binding and is provided by the CA.
MAC Key	Used for account binding and is provided by the CA.
Active	Option to make the credential active. This check box is selected by default.

3. Select **Submit**.

Set up routing policies for ACME

Set up routing policies to establish automated certificate management based on factors such as certificate authority (CA) and environment for efficient SSL/TLS certificate management.

Before you begin

Role required: pki_admin, flow_designer, action_designer, or admin

About this task

The routing policy decides which CA must be contacted for certificate operations. It contains the CA, CA URL, Credential, Approval Group, Assignment Group, and CSR attributes. The routing policy triggers the flow for requesting certificates for specific CAs.

Duplicate certificate requests aren't allowed. However, you can override this setting by selecting the Allow duplicate requests check box. A certificate request is considered duplicate if there's another certificate task with the same domain name that is still in progress.

Procedure

1. Navigate to **All** > **Certificate Management** > **Certificate Routing Policies**.
2. Select **New**.
3. On the **Certificate Routing Policy** form, fill in the fields.

   For a description of the field values, see Certificate Routing Policy form for ACME.
4. Select **Submit**.

Request certificates using ACME manual DNS challenge flow

Request a new certificate and automatically retrieve the certificates for an application using ACME manual flow of DNS challenge.

Before you begin

- The Certificate Management catalog should be enabled.
- A routing policy without any DNS challenge action must exist.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: Certificate requester, PKI admin, PKI user, or admin

Note:

Certificate requester is a user who does not have the PKI admin or PKI user role.

Procedure

1. Access the certificate request automated flow.
   a. Navigate to **All** > **Service Catalog**.
   b. Select **Certificate Management**.
   c. Select **Automated Flow**.
2. Select **Request New Certificate (Automated)**.
3. On the form, fill in the fields.

New certificate

Field	Description
Certificate Purpose	Indicates whether the request is for an internal or external certificate. For CAs (for example, Let's Encrypt), select **External**.
Certificate Signing Request (CSR)	CSR containing the certificate information.
Validity Period for Certificate (In Days)	Number of days the certificate is valid. For Let's Encrypt, the maximum validity period is 90 days.
Certificate Owner Group	Group for which the certificate tasks will be generated.
Certificate Owner	Name or role of the person who will own the certificate.

The following CSR attributes are matched and auto-populated based on the certificate information from CSR:

- Subject Common Name
- Subject Alternative Name
- Organization
- Organizational Unit
- Locality/City
- Province
- Country
- Email Address

4. Select **Submit**.
   Once the request is submitted, a task is created and an activity is assigned for you to complete the DNS challenge and mark it complete.
5. On the **New certificate task** page, in the **DNS Task** field, select the record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. On the **DNS Task** page, add a DNS TXT record for the attached DNS challenge.
   a. In the **DNS Challenges** pane, copy the DNS value.
   b. On the web browser, go to the domain and add the DNS value as a TXT record. For example, the domain can be **godaddy.com** > **thedisconow.com**.
   c. Fill in the other mandatory fields.
   d. Select **Save**.

   Note: Most DNS updates take effect within an hour but it could take up to 48 hours to update.

7. On the **DNS Task** page, change the status of the record to **Completed**.
8. **Optional:** Check whether the DNS record has propagated successfully. To check if it's propagated successfully, use the *dig* command.
9. Select **Save**.

Result

- Once the DNS challenge is completed, the automated flow makes the request to the CA to get the certificate.
- The certificate is attached to the New certificate task.
- The certificate task status changes to Completed.

Renew certificates using ACME manual DNS challenge flow

Request to renew certificate and automatically retrieve the certificates for an application using ACME manual flow of DNS challenge.

Before you begin

- The Certificate Management catalog should be enabled.
- A routing policy without any DNS challenge action must exist.

Role required: Certificate requester, PKI admin, PKI user, or admin

Note:

Certificate requester is a user who does not have the PKI admin or PKI user role.

Procedure

1. Access the certificate request automated flow.
   a. Navigate to **All** > **Service Catalog**.
   b. Select **Certificate Management**.
   c. Select **Automated Flow**.
2. Select **Renew Certificate (Automated)**.
3. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Renew certificate

Field	Description
Issued Certificate	Certificate to renew.
Certificate Purpose	Request internal or external certificate. For CAs (for example, Let's Encrypt), select **External**.
Certificate Signing Request (CSR)	CSR containing certificate information.
Validity Period for Certificate (In Days)	Number of days the certificate is valid. For Let's Encrypt, the maximum validity period is 90 days.
Certificate Owner Group	Group for which the certificate tasks will be generated.
Certificate Owner	Name or role of the person who will own the certificate.

The following CSR attributes are matched and auto-populated based on the certificate information from CSR:

- Subject Common Name
- Subject Alternative Name
- Organization
- Organizational Unit
- Locality/City
- Province
- Country
- Email Address

4. Select **Submit**.
   Once the request is submitted, a task is created and an activity is assigned for you to complete the DNS challenge and mark it complete.
5. On the **New certificate task** page, in the **DNS Task** field, select the record.
6. On the **DNS Task** page, add a DNS TXT record for the attached DNS challenge.

   a. In the **DNS Challenges** pane, copy the DNS value.

   b. On the web browser, go to the domain and add the DNS value as a TXT record.
   For example, the domain can be **godaddy.com** > **thedisconow.com**.

   c. Fill in the other mandatory fields.

   d. Select **Save**.

   Note: Most DNS updates take effect within an hour but it could take up to 48 hours to update.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. On the **DNS Task** page, change the status of the record to **Completed**.
8. **Optional:** Check whether the DNS record has propagated successfully.
   To check if it's propagated successfully, use the *dig* command.
9. Select **Save**.

Result

- Once the DNS challenge is completed, the automated flow makes the request to the CA to get the certificate.
- The certificate is attached to the New certificate task.
- The certificate task status changes to Completed.

Request certificates using ACME automated DNS challenge flow

Request a new certificate and automatically retrieve the certificates for an application using an Automated Certificate Management Environment (ACME) automated flow of DNS challenge.

Before you begin

- Ensure that a credential has been set up. For more information, see Create credentials for ACME certificate authority
- The Certificate Management catalog has been enabled.
- A routing policy with a DNS challenge action exists. For more information, see Set up routing policies for ACME.
- Role required: Certificate requester [sn_disco_certmgmt.certificate_requester], certificate administrator [sn_disco_certmgmt.pki_admin], certificate user [sn_disco_certmgmt.pki_user], flow_designer, action_designer, or admin.

Procedure

1. Navigate to **All** > **Self-Service** > **Service Catalog**.
2. Access the form for requesting a new certificate.
   a. Select **Certificate Management**.
   b. Select **Automated Flow**.
   c. Select **Request New Certificate (Automated)**.
3. For the **Select how to manage your certificate** value, select **ServiceNow**.
4. On the form, fill in the fields.

New certificate

Field	Description
Certificate Purpose	Determines whether the Certificate is internal or external. ○ Internal: Requires a completed form ○ External: Requires importing the CSR and private key from an external source
Environment	Environment where the certificate must be deployed or installed. The available options are Development, Disaster recovery, Production, or Sub-Production.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Certificate Format	Format in which the certificate is generated. The available options are: PEM, DER, and PKCS12. The default value is PEM. **Note:** The format must match the Certificate Format configured on the routing policy that handles your request. If the formats don't match with the routing policy, you're notified that manual intervention is required
Validity Period for Certificate (In Days)	The number of days the certificate is valid. This field appears only when External is selected from the Certificate Purpose field.
Certificate Signing Request (CSR)	Details about the CSR containing certificate information.
Validity Period for Certificate (In Days)	The number of days the certificate is valid. This field appears only when External is selected in the Certificate Purpose field.
Certificate Owner Group	Owner group for the certificate.
Certificate Owner	Name or role of the person who owns the certificate.

The following CSR attributes are matched and auto-populated based on the certificate information from CSR:

- Subject Common Name
- Subject Alternative Name
- Organization
- Organizational Unit
- Locality/City
- Province
- Country
- Email Address

5. Select **Submit**.
   After the request is submitted, a task is created for completing the DNS challenge. The task is completed automatically.

Result

- The certificate is attached to the certificate task in the format you selected.
  - For PEM, the server, intermediate, and root certificates are each attached as separate PEM files.
  - For DER, the same three certificates are attached in DER format. If conversion to DER fails, the certificates are attached in PEM format and a warning is recorded in the work notes of the certificate task. The certificate has already been issued by the certificate authority, so no certificate data is lost.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

  - For PKCS12, the server certificate is attached as a key store that contains the server, intermediate, and root certificates together with the private key, encrypted with the password you stored in your external vault. The intermediate and root certificates are also attached separately in PEM format, because no private key exists for them.

- If a value is selected in the DNS Challenge Action field in the routing policy, the system creates the required DNS record automatically through the configured DNS provider. After DNS record propagation is complete (two minutes by default), the DNS challenge is completed. The automated flow sends a request to the CA to get the certificate. Admins can change this duration by modifying the `sn_disco_certmgmt.wait_time_for_dns_record_propagation` system property.
- If no value is selected in the DNS Challenge Action field and the domain is already validated with the CA, the DNS challenge task completes without additional action.
- If no value is selected in the DNS Challenge Action field and the domain is not validated, you must manually add the required DNS TXT record to complete the validation. The certificate is attached to a certificate task. The request certificate task status changes to Completed.

Renew certificates using ACME automated DNS challenge flow

Request to renew certificates and automatically retrieve the certificates for an application using the Automated Certificate Management Environment (ACME) automated flow of DNS challenge.

Before you begin

- Ensure that a credential has been set up. For more information, see Create credentials for ACME certificate authority.

  Note: The GoDaddy credential is provided with the base system inside the credential page.

- The Certificate Management catalog has been enabled.
- A routing policy with a DNS challenge action exists. For more information, see Set up routing policies for ACME.
- Role required: Certificate requester [sn_disco_certmgmt.certificate_requester], certificate administrator [sn_disco_certmgmt.pki_admin], certificate user [sn_disco_certmgmt.pki_user], flow_designer, action_designer, or admin.

Procedure

1. Navigate to **All** > **Self-Service** > **Service Catalog**.
2. Access the form for renewing a certificate.
   a. Select **Certificate Management**.
   b. Select **Automated Flow**.
   c. Select **Renew Certificate (Automated)**.
3. In the **Issued Certificate** field, select the Lookup using list icon to find and select the certificate.
   The date the certificate expires is displayed in the Certificate Expires On field.
4. Review and update the details of the certificate as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Renew certificate

Field	Description
Certificate Purpose	Request internal or external certificate.
Certificate Format	Format in which the certificate is generated. The available options are: PEM, DER, and PKCS12. The default value is PEM. **Note:** The format must match the Certificate Format configured on the routing policy that handles your request. If the formats don't match with the routing policy, you're notified that manual intervention is required
Validity Period for Certificate (In Days)	The number of days the certificate is valid. This field appears only when External is selected from the Certificate Purpose field.
Certificate Signing Request (CSR)	CSR containing certificate information.
Validity Period for Certificate (In Days)	The number of days the certificate is valid. This field appears only when External is selected in the Certificate Purpose field.
Certificate Owner Group	Owner group for the certificate.
Certificate Owner	Name or role of the person who owns the certificate.

The following CSR attributes are matched and auto-populated based on the certificate information from CSR:

- Subject Common Name
- Subject Alternative Name
- Organization
- Organizational Unit
- Locality/City
- Province
- Country
- Email Address

5. Select **Submit**.
   Once the request is submitted, a task is created for completing the DNS challenge. The task is completed automatically.

Result

- The certificate is attached to the certificate task in the format you selected.
  - For PEM, the server, intermediate, and root certificates are each attached as separate PEM files.
  - For DER, the same three certificates are attached in DER format. If conversion to DER fails, the certificates are attached in PEM format and a warning is recorded in the work notes of

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

the certificate task. The certificate has already been issued by the certificate authority, so no certificate data is lost.

- For PKCS12, the server certificate is attached as a key store that contains the server, intermediate, and root certificates together with the private key, encrypted with the password you stored in your external vault. The intermediate and root certificates are also attached separately in PEM format, because no private key exists for them.

- If a value is selected in the DNS Challenge Action field in the routing policy, the system creates the required DNS record automatically through the configured DNS provider. After DNS record propagation is complete (two minutes by default), the DNS challenge is completed. The automated flow sends a request to the CA to get the certificate. Admins can change this duration by modifying the *sn_disco_certmgmt.wait_time_for_dns_record_propagation* system property.
- If no value is selected in the DNS Challenge Action field and the domain is already validated with the CA, the DNS challenge task completes without additional action.
- If no value is selected in the DNS Challenge Action field and the domain is not validated, you must manually add the required DNS TXT record to complete the validation. The certificate is attached to a certificate task. The request certificate task status changes to Completed.

Revoke certificates using ACME automated flow

Request a revoke certificate for an application.

Before you begin

- Ensure that a credential has been set up. For more information, see Create credentials for ACME certificate authority.
- The Certificate Management catalog has been enabled.
- A routing policy with a DNS challenge action exists. For more information, see Set up routing policies for ACME.
- Role required: certificate administrator [sn_disco_certmgmt.pki_admin] or admin.

Procedure

1. Navigate to **All** > **Self-Service** > **Service Catalog**.
2. Access the form for renewing a certificate.
   a. Select **Certificate Management**.
   b. Select **Automated Flow**.
   c. Select **Revoke Certificate (Automated)**.
3. Unlock the **Issued Certificate** field.
4. Select the Lookup using list icon (🔍) and select the certificate you want to revoke. You can select more than one certificate.
5. Provide an appropriate reason for revoking the certificate.
6. Place the revoke order by selecting **Submit**.
7. In the confirmation dialog box, select **OK**.

Result

A task is automatically created that triggers the revocation operation. After the operation is complete, the status of the certificate changes to Revoked.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACME integration with KeyFactor EJBCA for automated flows

Automate the flow of requesting, renewing, and revoking your certificates by integrating Keyfactor EJBCA with the Automated Certificate Management Environment (ACME).

Keyfactor EJBCA is a certificate authority that issues your certificates. ACME are a set of protocols and rules that give you a secure environment to use an automated flow of managing certificates.

By configuring your routing policy fields, you can ensure that the content in your Certificate Signing Request (CSR) aligns with the correct routing policy. This streamlines the process of requesting, renewing, and revoking your certificates.

EJBCA has two types of Credentials. One includes External Account Binding (EAB), and the other doesn't.

In EJBCA, automated certificate workflows start when you create routing policies for EJBCA ACME Certificates. For every routing policy, there are required fields where you have to give information.

Your platform has routing policies where you fill in all the fields of the routing policy. Your platform aligns that information to each CSR you create to request, renew, and revoke certificates.

For more information, see Create routing policies for EJBCA ACME certificates.

Perform the following task to use EJBCA ACME to automate your certificate life cycle:

1. Create a Connection & Credential alias
2. Configure base API URL for EJBCA ACME
3. Validate EJBCA ACME base API URL

Configure base API URL for EJBCA ACME

Configure your base API URL for EJBCA ACME to your organization's root URL address.

Before you begin

Role required: discovery_admin, Public Key infrastructure (PKI) admin, admin

Procedure

1. Navigate to **All** > **Certificate Authorities**.
2. In the table of certificate authorities, select **EJBCA ACME**.
3. Edit the text in one of the following ways:
   - Select **edit** text.
   - Double-click (or use the keyboard shortcut)" the **Base URL** field.

   Note: If a warning message appears, you might be in the "Global" application. Select the link that appears to edit the record.
4. Paste your own API URL over the `[ejbca_url]` text inside the URL.
5. Make sure that the ACME option is selected.
6. Select **Save**.

Result

The base API URL for your EJBCA ACME environment is configured to your systems.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create routing policies for EJBCA ACME certificates

Automate your EJBCA ACME workflows by creating a routing policy that aligns with your Certificate Signing Requests (CSRs) to request, renew, and revoke certificates.

Before you begin
Role required: discovery_admin, Public Key infrastructure (PKI) admin, admin

About this task
Streamline your certificate life-cycle workflow by filling out a routing policy form. This form configures your routing policy fields to make sure your Certificate Signing Request (CSR) aligns with the correct Certificate Authority (CA) routing policy.

Procedure
1. Navigate to **All** > **Certificate Management** > **Certificate Routing Policies**.
2. Select **New**.
3. On the Certificate Routing Policy form, fill in the fields.
   For a description of the field values, see Certificate routing policy form for EJBCA ACME.

   ⚠ **Warning:** For your automated workflow to work, you must fill in every field in your routing policy form. If you don't, your CSR requires manual steps. You can see all the fields and values at the Certificate routing policy form for EJBCA ACME.
4. Select **Submit**.

What to do next
Navigate to **All** > **Certificate Management** > **Certificate Routing Policies**.

By locating your routing policy on the Certificate Routing Policies page, you validate your EJBCA ACME Certificate routing policy.

Validate EJBCA ACME base API URL

Validate that your base API URL for EJBCA ACME has been updated to your organization's root URL address.

Before you begin
Role required: discovery_admin, Public Key infrastructure (PKI) admin, admin

Procedure
1. Navigate to **Workspaces** > **Certificate Management**.
2. Select the **Certificate Records** tab.
3. Select **Certificate Authority API URL**.
4. Select **EJBCA ACME** in the table.

Result
The UI is redirected to the certificate field and its Base URL.

What to do next
Review the URL and confirm that it's not the standard default settings.
To connect with EJBCA, you must change the default URL to your specific environment (env) URL. If you still see the default setting, reconfigure your base API URL. For more information, see Configure base API URL for EJBCA ACME.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate management with CyberArk Certificate Manager SaaS

The ServiceNow Certificate Inventory and Management application has been integrated with CyberArk Certificate Manager SaaS for automated certificate life-cycle management, providing centralized certificate provisioning, renewal, and revocation capabilities.

Integration overview

The integration connects Certificate Inventory and Management with CyberArk Certificate Manager SaaS to automate certificate operations so you can use the certificate management capabilities of CyberArk while maintaining centralized visibility and control through Certificate Inventory and Management.

Automated certificate requests, renewals, and revocations are handled through routing policies that direct certificate operations to the appropriate CyberArk certificate authority. Certificate life-cycle events are tracked and managed within Certificate Inventory and Management.

Note: Automated certificate renewal for CyberArk is managed by the CyberArk Certificate Manager SaaS platform. Certificate Inventory and Management does not trigger auto-renewal for CyberArk-managed certificates from your instance.

Key benefits

- Centralized certificate life-cycle management: Request, renew, and revoke certificates managed by CyberArk directly from Certificate Inventory and Management, without switching between platforms for certificate operations.
- Automated certificate provisioning: Automate certificate request and renewal workflows through routing policies, reducing manual effort and the risk of expired certificates causing service outages.
- Secure private key handling: CyberArk manages private key storage, so private keys aren't transferred to or stored in Certificate Inventory and Management.
- Change management governance: Certificate operations automatically generate change requests in Certificate Inventory and Management, providing audit trails and compliance documentation for certificate life-cycle events.
- Centralized visibility: Monitor all CyberArk managed certificates alongside certificates from other providers in the **Certificate Management** workspace.

Certificate management workflow

Certificate operations follow this process:

1. Configure CyberArk credentials in Certificate Inventory and Management. For more information, see Configure CyberArk Certificate Manager SaaS credentials.
2. Create routing policies to direct certificate requests to CyberArk. For more information, see Create routing policies for CyberArk Certificate Manager SaaS.
3. Submit certificate requests and monitor certificate life-cycle events through automated flows:
   - Request certificates through CyberArk Certificate Manager SaaS
   - Renew certificates through CyberArk Certificate Manager SaaS
   - Revoke certificates through CyberArk Certificate Manager SaaS

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure CyberArk Certificate Manager SaaS credentials

Configure authentication credentials so Certificate Inventory and Management can communicate with CyberArk Certificate Manager SaaS for automated certificate life-cycle management.

Before you begin

Access to CyberArk Certificate Manager SaaS with appropriate permissions for certificate operations is required.

Role required: pki_admin or admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Select **New**.
   a. From the credentials list, select **Certificate Management Credentials**.
   b. On the form, fill in the fields.

Field	Description
Name	Descriptive name for the credential.
CA Type	This value should be **CyberArk Venafi SaaS**.
API Key	API key generated from CyberArk.
Active	Option to make the credential active.

3. Select **Submit** to save the credential configuration.
4. **Optional:** Validate that the credential is valid and reachable.
   a. Navigate to **All** > **Connections & Credentials** > **Credentials**
   b. Open the credential you created.
   c. Select **Validate Credential**.
   If the validation result is unsuccessful, verify the **API Key** value and retry.
5. **Optional:** Assign a credential alias to the credential.
   For more information, see Credential aliases for Discovery.
   a. In the **Credential alias** field, select the Unlock Credential alias icon.
   b. Select the Lookup using list icon to find and select the credential alias.
   c. Select the Lock Credential alias icon.
   d. Select **Update**.

Create routing policies for CyberArk Certificate Manager SaaS

Create routing policies to direct certificate requests to CyberArk Certificate Manager SaaS by matching attributes such as organization, domain, and certificate type to the Routing Policy [sn_disco_certmgmt_routing_policy] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Configure CyberArk credentials. For more information, see Configure CyberArk Certificate Manager SaaS credentials.

Role required: pki_admin or admin.

Procedure

1. Navigate to **All** > **Certificate Management** > **Certificate Routing Policies**.
2. Select **New**.
3. On the Certificate Routing Policy form, fill in the fields.
   For a description of the field values, see Certificate Routing Policy form.
4. If used by your organization as filtering criteria, set values for the **Organization**, **Organizational Unit**, **Locality**, **Country**, **State**, or **Email Address** fields.

   Note: Use an asterisk (*) for wildcard matching and RegEx patterns for complex matching requirements.

5. Select **Submit**.

Result

The routing policy is created and directs matching certificate requests to CyberArk. Certificate requests matching the policy criteria are automatically routed to CyberArk for processing. For more information, see Routing policy details.

Request certificates through CyberArk Certificate Manager SaaS

Submit a certificate request managed through CyberArk Certificate Manager SaaS using configured routing policies.

Before you begin

- CyberArk credentials must have been configured. For more information, see Configure CyberArk Certificate Manager SaaS credentials.
- Routing policies must have been configured. For more information, see Create routing policies for CyberArk Certificate Manager SaaS.

Role required: pki_admin, pki_user, or admin.

Procedure

1. Navigate to **All** > **Self-Service** > **Service Catalog**.
2. Access the form for requesting a new certificate.
   a. Select **Certificate Management**.
   b. Select **Automated Flow**.
   c. Select **Request New Certificate (Automated)**.
3. For the **Select how to manage your certificate** value, select **CyberArk Certificate Manager SaaS**.
4. On the Request New Certificate (Automated) form, fill in the fields.
   For a description of the field values, see Certificate request form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: A configured routing policy is applied automatically based on certificate attributes. For more information, see Create routing policies for CyberArk Certificate Manager SaaS

5. Select **Submit**.

Result
The certificate request is routed to CyberArk for processing. The request status is updated in the relevant certificate task as it progresses through the workflow.

Renew certificates through CyberArk Certificate Manager SaaS

Renew an existing certificate through the CyberArk Certificate Manager SaaS to extend its validity period before expiration.

Before you begin
Role required: pki_admin, pki_user, or admin

About this task
The renewal process uses the same certificate attributes and routing policy as the original certificate request.

Procedure
1. Navigate to **All** > **Self-Service** > **Service Catalog**.
2. Access the form for renewing a certificate.
   a. Select **Certificate Management**.
   b. Select **Automated Flow**.
   c. Select **Renew Certificate (Automated)**.
3. In the **Issued Certificate** field, select the Lookup using list icon to find and select the certificate.
   The date the certificate expires is displayed in the **Certificate Expires On** field.
4. Review and update the details of the certificate as needed.
   For a description of the field values, see Certificate request form.
5. Select **Submit**.

Result
The renewal request is routed to CyberArk Certificate Manager SaaS. The original certificate remains valid until the renewed certificate is issued.

Revoke certificates through CyberArk Certificate Manager SaaS

Revoke a certificate through the CyberArk Certificate Manager SaaS to invalidate it before its scheduled expiration date.

Before you begin
The certificate must be originally issued through CyberArk and be currently valid.

Role required: pki_admin or admin

About this task
Common reasons for revocation include key compromise, change in certificate details, or decommissioning of the associated service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Self-Service** > **Service Catalog**.
2. Access the form for renewing a certificate.
   a. Select **Certificate Management**.
   b. Select **Automated Flow**.
   c. Select **Revoke Certificate (Automated)**.
3. For the **Select how to manage your certificate** value, select **CyberArk Certificate Manager SaaS**.
4. Select the certificate that must be revoked.
   a. In the **Issued Certificate** field, select the Lock icon .
   b. Select the Lookup using list icon to find and select the certificate.
   c. On the Unique Certificates form, select the certificate to be revoked.
5. In the **Reason for Revoking Certificate** field, provide one of the revocation reasons for CyberArk Certificate Manager SaaS.
   You can also display these reasons by expanding **More information**.
   - SUPERSEDED: The certificate has been replaced
   - KEY_COMPROMISE: The private key has been compromised
   - AFFILIATION_CHANGED: The certificate holder's affiliation has changed
   - CESSATION_OF_OPERATION: The service is no longer operational
   - UNSPECIFIED: Any other reason.
6. Select **Submit**.

Result

The revocation request is sent to CyberArk for processing. The certificate status is updated in the relevant certificate task to reflect the revocation. The certificate is invalidated and can no longer be used for authentication.

Automated certificate renewal

Detect certificates that are about to expire and renew them automatically before expiration to keep your digital systems secure and continuously available.

Expired certificates are a serious problem for organizations.

Many organizations face unexpected outages due to their Transport Layer Security (TLS) certificates expiring without their knowledge.

Manually tracking your digital TLS certificates is a time consuming, tedious job that can cause certificates to expire unnoticed. Overlooking expiring certificates renders sites insecure and dangerous to users.

Automatic certificate renewal agents let you find certificates approaching expiration and automatically renew them before they expire.

You can also set a certificate to renew automatically when requesting a new certificate.

Certificate Inventory and Management supports HashiCorp, Azure KeyVault, and CyberArk PVWA for external private key storage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate renewal AI agent

Find certificates that are about to expire and renew them by describing what you want in natural language, or renew a specific certificate directly from its record.

The certificate renewal AI agent is part of AI Agents for Discovery. You interact with it from the AI panel in Certificate Management workspace. You can use the agent in two ways:

- From a certificate record. When you open the AI panel while viewing a certificate, the agent picks up that certificate as context and offers to renew it.
- From a prompt. Describe the certificates you want to renew, and the agent returns a list of matching certificates for you to confirm before renewal begins. In either case, the agent creates a certificate renewal task and gives you a link to the task record so that you can track progress.

Before you use the agent, configure your system for automatic certificate renewal:

1. Configure MID Server for automatic certificate renewal
2. Add required applications and capabilities to MID Server

Renew certificates using the AI agent

Renew a specific certificate by opening it and asking the certificate renewal AI agent to renew it.

Before you begin

Configure your system for certificate renewal:

1. Configure MID Server for automatic certificate renewal
2. Add required applications and capabilities to MID Server

Role required: sn_disco_certmgmt.pki_admin

Procedure

1. Navigate to **Workspaces** > **Certificate Management**.
2. Open the certificate that you want to renew.
3. Select the ServiceNow Otto icon to open the AI panel.
   The agent identifies the certificate you're viewing and asks whether you want to renew it. Confirm that the certificate name in the agent's prompt matches the certificate you intend to renew before you continue.
4. Confirm the renewal.

Result

The agent creates a certificate renewal task and returns a link to the task record. Select the link to track the status of the renewal.

Find and renew expiring certificates using the AI agent

Describe the certificates that you want to renew, and the certificate renewal AI agent finds them and renews them after you confirm.

Before you begin

Configure your system for certificate renewal:

1. Configure MID Server for automatic certificate renewal
2. Add required applications and capabilities to MID Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: sn_disco_certmgmt.pki_admin

About this task

You can describe certificates by any field on the Certificate [cmdb_ci_certificate] table using the field label as it appears in the record. For example, expiration date, assigned user, description, or thumbprint.

Procedure

1. Navigate to **Workspaces** > **Certificate Management**.
2. Select the ServiceNow Otto icon to open the AI panel.
3. Describe the certificates that you want to renew.
   For example,`Renew all certificates expiring in the next 5 days`, `Renew the certificates assigned to me that expire before 1 August 2026`, or `Show me certificates expiring between 1 May 2026 and 1 August 2026`.
4. Review the list of certificates that the agent returns.
5. Confirm the renewal.

Result

The agent creates a renewal task for each certificate and returns links to the task records. If no certificates match your description, the agent tells you that none were found.

Configuring automated certificate renewal

Configure your MID Server and system properties to automatically renew certificates before they expire.

To automatically renew your certificates before they expire, use the automated certificate renewal workflow or the certificate renewal AI agent.

In both cases, you must first configure your system for automatic renewal.

Configuring for automated certificate renewal involves the following steps:

1. Configure MID Server for automatic certificate renewal
2. Add required applications and capabilities to MID Server
3. Configure system properties for auto-renewal. This step is optional if you use the certificate renewal AI agent.

Configure MID Server for automatic certificate renewal

Collect information about root certificates stored outside your server. Create a specialized Discovery schedule.

Before you begin

Role required: pki_admin or admin

About this task

Configure your MID Server to renew certificates automatically by setting the configuration parameters in your MID Server.

For information about version compatibility and troubleshooting, see the Renewal of TLS certificates using AI Agents for Discovery knowledge article [KB2470998] in the Now Support Knowledge Base. The Certificate Inventory and Management on Yokohama Patch 8 or later supports the certificate renewal agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Discovery** > **MID Servers**.
2. Select the MID Server that you want to configure.
3. Select the **Configuration Parameters** tab.
4. Add the parameters for your external vault provider.
   For each parameter, select **New**, select the **Parameter name** and **Value**, and then select **Submit**.

HashiCorp Vault MID Server parameters

Parameter name	Value
ext.vault.hashicorp.address	Address of your external HashiCorp vault. The default value is `http://127.0.0.1:8200`.
ext.vault.hashicorp.path	File path in your HashiCorp vault.

Azure Key Vault MID Server parameters

Parameter name	Value
ext.vault.azure.keyvault.name	Name of your Azure Key Vault instance.
ext.vault.azure.client.id	Application (client) ID of your Azure AD registered application.
ext.vault.azure.tenant.id	Directory (tenant) ID of your Azure AD tenant.
ext.vault.azure.client.secret	Client secret for your Azure AD registered application.

Note: The system caches the Azure authentication token in memory and refreshes it before expiry to maintain secure access without unnecessary token requests.

CyberArk PVWA MID Server parameters

Parameter name	Value
ext.vault.cyberark.pvwa.url	Base URL of your CyberArk PVWA instance.
ext.vault.cyberark.pvwa.username	Username used to authenticate to CyberArk PVWA.
ext.vault.cyberark.pvwa.safe_name	Name of the CyberArk safe where private keys are stored.
ext.vault.cyberark.pvwa.platform_id	Platform ID configured in CyberArk PVWA for the stored accounts.

5. Add the secure credential for your external vault provider to your MID Server configuration file.
   a. Navigate to the IP address in the **IP address** field of your MID Server record.
   b. Navigate to the **MID Server installed folder** where you installed your MID Server.
   c. Select the `agent/config.xml` file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Add the parameter for your provider to your `config.xml` file.

Provider	Parameter	Value
HashiCorp Vault	ext.vault.hashicorp.token	Your HashiCorp vault token.
Azure Key Vault	ext.vault.azure.client.secret	Client secret for your Azure AD registered application.
CyberArk PVWA	ext.vault.cyberark.pvwa.password	Password for your CyberArk PVWA user.

e. Insert the following code, replacing the parameter name and value for your provider.

```
<parameter name="Parameter" secure="true" value="<Value>"/>
```

For example,

```
<parameter name="ext.vault.hashicorp.token" secure="true"
 value="<Your token value>"/>
```

f. Restart your MID Server.

Result
Your MID Server is configured for automatic certificate renewal.

What to do next
To complete the process of configuring automatic certificate renewal, you must complete the required steps to Add required applications and capabilities to MID Server and Configure system properties for auto-renewal.

Add required applications and capabilities to MID Server

Enable your MID Server to auto-renew your certificates by adding the Certificate Inventory and Management and GenerateCSR applications.

Before you begin
Check that you have completed the task Configure MID Server for automatic certificate renewal.

Role required: pki_admin or admin

About this task
Move the applications Certificate Inventory and Management and GenerateCSR from your collection list to your supported applications list.

Procedure

1. Navigate to **All** > **Mid Servers**.
2. Select the MID Server that you want to configure.

   Note: The MID Server must be the same MID Server that you configured for automated certificate renewal.

3. Modify the MID Server
   a. Select **Supported Applications**.
   b. Select **Edit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Move **Certificate Inventory and Management** to the **Supported Applications List**.

d. Select **Save**.
You're redirected to the MID Server page.

4. Add MID Server Capabilities

a. Select **Capabilities**.

b. Select **Edit**.

c. Move **GenerateCSR** from the **Collections List** to the **Supported Applications List**.

d. Select **Save**.
You're redirected to the MID Server page.

Result
Your MID Server supports the Certificate Inventory and Management and GenerateCSR applications, fulfilling a requirement to enable automatic certificate renewal.

What to do next
Configure system properties for auto-renewal, to complete the configuration for automatic certificate renewal.

Configure system properties for auto-renewal

Enable the auto-renewal options in your System Properties to configure your system to renew automatically Transport Layer Security (TLS) certificates before they expire.

Before you begin
Check that you have completed the tasks, Configure MID Server for automatic certificate renewal, and Add required applications and capabilities to MID Server.

Role required: pki_admin or admin

About this task
Configuring your system properties to enable automatic certificate renewal is required for using the automatic renewal workflow and the AI certificate renewal agent.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for *auto_renew*.
3. Select**sn_disco_certmgmt.enable_auto_renewal_options_in_service_catalog**
4. Set the **Value** field to **true**.

Result
Your system is configured to renew automatically certificates before they expire.

What to do next
You can Set certificates to renew automatically or use the Certificate renewal AI agent.

Define maximum retries for auto-renewal

Use the system property sn_disco_certmgmt.auto_renewal_max_retries to set the maximum retries your system enables for automatic certificate renewal.

Before you begin
Role required: pki_admin or admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

If an automatic certificate renewal fails, the system will reattempt to renew the next day until it exhausts the maximum number of retries. Set this property to configure the maximum number of retries your system makes to renew your certificate.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for *auto_renew*.
3. Select**sn_disco_certmgmt.auto_renewal_max_retries**
4. Set the **Value** field to the maximum retries that you want.

Set certificates to renew automatically

Configure automatic certificate renewal to avoid service interruptions from expired certificates. Set the number of days before expiration when the system should automatically renew certificates.

Before you begin

Complete the following tasks to configure your system to renew your certificates automatically:

1. Configure MID Server for automatic certificate renewal
2. Add required applications and capabilities to MID Server
3. Configure system properties for auto-renewal

Role required: pki_admin or admin

Procedure

1. Navigate to **Service Catalog** > **Certificate Management** > **Automated Flow** > **Renew Certificate (Automated)**.
2. Select the **Issued Certificate** that you would like to renew by selecting the search icon .
3. Set the **Renew Automatically** field to **Yes**.

   Note: By selecting **Yes**, a new field appears called **How many days before expiry does this certificate need to be renewed?**.

4. Select the number of days in the **How many days before expiry does this certificate need to be renewed?** field.
5. Select **Submit**.

Result

The certificate you chose is set to auto-renew the number of days before expiration that you request.

Request certificates with auto-renewal

Set your certificate to auto-renew when you first request it.

Before you begin

Complete the following tasks to configure your system to renew your certificates automatically:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Configure MID Server for automatic certificate renewal
2. Add required applications and capabilities to MID Server
3. Configure system properties for auto-renewal

Role required: pki_admin or admin

Procedure

1. Navigate to **Service Catalog** > **Certificate Management** > **Automated Flow** > **Request New Certificate (Automated)**.
2. Select the check box **Want to Generate CSR?**
3. On the "request a new certificate" form, fill in the fields.
   For a description of the field values, see ,.
4. Select **Generate**.
5. Set the **Validity Period for Certificate (in Days)** field to validity period of the certificate.
6. Set the **Renew Automatically** field to **Yes**.

   Note: By selecting **Yes**, a new field appears called **How many days before expiry does this certificate need to be renewed?**.

7. Select the number of days before the certificate you want to renew in the **How many days before expiry does this certificate need to be renewed?** field.
8. Select **Submit**.

Result

Your new certificate is issued with instructions to renew automatically before it expires.

Validate auto-renew status of certificate

Confirm that you set your certificate to renew automatically before it expires.

Before you begin

Complete one of the following tasks before you validate your certificate:

- Set certificates to renew automatically
- Request certificates with auto-renewal

Role required: pki_admin or admin

Procedure

1. Navigate to **All** > **Unique Certificates**.
2. Select the Certificate that you want to validate.
3. Select the **Certificates metadata** tab.
4. Check that the **Auto renew enabled** field is set to **true**.

   Note: The pki_admin can select this field and edit it.

Certificate attestation for certificate owners

Attestation jobs run in regular intervals, sending emails to certificate owners, giving them a chance to review certificate ownership.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can configure the system to send certificate ownership attestation reviews by email at regular intervals. Certificate owners can attest to or reject ownership of certificates via these emails. Alternatively, you can reassign certificate ownership to align with user changes.

Related topics

- Configure certificate attestation reviews
- Attest or reject certificate ownership via email
- Reassign certificate ownership

Configure certificate attestation reviews

Send out emails to all certificate owners confirming their ownership or confirming they're no longer certificate owners at set intervals.

Before you begin

Role required: pki_admin or admin

Procedure

1. Navigate to **System Definition** > **Scheduled Jobs**.
2. Select **Certificate Ownership Attestation Policy Processor**.
3. Configure the fields to set when and how often the job runs.

   Note: The default settings are to run the job every month.

4. Select **Update** or **Execute Now**.

   Note: You can test your attestation settings by executing them immediately to verify emails are sent to the configured users.

Result

Email notifications are sent out to all certificate users based on the time and time intervals you set. As each user receives their notification of certificate ownership, they can attest or reject their ownership.

Attest or reject certificate ownership via email

Attest or reject ownership of a certificate from your certification attestation email notification.

Before you begin

- Verify that your system is configured to send certificate attestation notifications. For more information, see Configure certificate attestation reviews
- Role required: pki_admin or admin

Procedure

1. Select **Review CIs** in your email notification.
2. Select each certificate by selecting the check box next to it.
3. Select **Attest**.
4. Select all certificates you want to reject.
5. Select **Reject**
6. In the text box that appears, explain why you are rejecting the certificate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Select **Submit**.
8. In the Submit Attestation confirmation window, select **Submit**.

Result
For the certificates that you attest ownership, you receive further notifications at monthly intervals, unless otherwise configured. For the certificates that you reject ownership, you no longer receive certificate attestation notifications.

Reassign certificate ownership

View a list of all certificates and their respective owners, and reassign ownership for certificates.

Before you begin
Role required: pki_admin or admin

Procedure
1. Navigate to **Workspaces** > **Certificate Management**.
2. Select **Certificate records**.
3. Select the certificate you want to reassign.
4. Select the Search for Record icon in the **Assigned to** field.
5. Select a new user to assign the certificate.
6. Select **Save**.

Result
The certificate is reassigned to a new owner. The new owner receives email notifications to attest or reject ownership of the certificate.

Certificate Inventory and Management reference

Reference topics provide additional information about Certificate Inventory and Management dashboard and components, including tables, patterns, and terms.

Certificate Inventory and Management tables

This framework supports security, compliance, and streamlined operations. Certificate Inventory and Management tables provide a centralized system to track and manage digital certificates. They capture key details, including discovered certificates, installation locations, historical data, and associated tasks such as renewals and requests.

Certificate Inventory and Management tables

Table	Description
Unique Certificate [cmdb_ci_certificate]	The Fingerprint column is unique for every server certificate. **Note:** You can also view the entire certificate chain using the related list in the Unique Certificate table. The State field records the life cycle state of the certificate and doesn't include an Expired value. Expired certificates are denoted as Issued. Use the Valid to field to identify certificates that have expired.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate Inventory and Management tables (continued)

Table	Description
Installed Certificate [sn_disco_certmgmt_cmdb_installed_certificate]	Discovery populates this table with all discovered certificates and shows all locations where certificates are installed. It does not have a reference to the server CI and does not go through IRE. Relationships with the managed certificate are stored in the CI Relationship [cmdb_rel_ci] table. The relationships can be for servers, applications, or business services.
Discovered Certificate [sn_disco_certmgmt_certificate_history]	Discovered certificates are stored in this table first as part of the Shazzam and URL discovery sensor and shows the XML payload of all certificates with that address.
Certificate Domain [certificate_domain]	Contains one or more domains for the subject alternative name.
Certificate Task [sn_disco_certmgmt_certificate_task]	Contains all certificate renewal tasks and new certificate request tasks.
Certificate URL [sn_disco_certmgmt_cert_url]	Contains URLs targeted for certificate discovery.
Certificate Management Credential [sn_disco_certmgmt_certificate_management_credential] (Version 1.1.7 Certificate Inventory and Management)	Stores the CA types including GoDaddy and DigiCert. Others are populated based on the *certificate_authority* script. Only discovery_admin can create other CA credentials.
Scheduled Certificate URL [sn_discovery_cert_url_sched_m2m]	Links the URLs to discover [discovery_cert_url] to a particular Discovery schedule [discovery_schedule].
Certificate Authority [sn_disco_certmgmt_ca] (Version 1.3.8 Certificate Inventory and Management)	Contains the Certificate Authority name and base URL of the REST API.
Certificate Authority API URL [sn_disco_certmgmt_ca_api_url] (Version 1.3.8 Certificate Inventory and Management)	Contains the reference of certificate authority, end point URL of REST API, and certificate validation type.
Routing Policy [sn_disco_certmgmt_routing_policy]	Determines which CA to contact for certificate operations. This table contains the CA, CA URL, Credential, Approval Group, Assignment Group, and CSR attributes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate Inventory and Management tables (continued)

Table	Description
(Version 1.3.8 Certificate Inventory and Management)	
Automated Certificate Task [sn_disco_certmgmt_task] (Version 1.3.8 Certificate Inventory and Management)	Contains all automated certificate request tasks (new, renew, and revoke).
New Certificate Tasks [sn_disco_certmgmt_new_task] (Version 1.3.8 Certificate Inventory and Management)	Extends from [sn_disco_certmgmt_task] and contains new certificate request tasks.
Renew Certificate Tasks [sn_disco_certmgmt_renew_task] (Version 1.3.8 Certificate Inventory and Management)	Extends from [sn_disco_certmgmt_task] and contains renew certificate request tasks.
Revoke Certificate Tasks [sn_disco_certmgmt_revoke_task] (Version 1.3.8 Certificate Inventory and Management)	Extends from [sn_disco_certmgmt_task] and contains revoke certificate request tasks.
Certificate Extensions [sn_disco_certmgmt_certificate_extension] (Version 1.3.8 Certificate Inventory and Management)	Stores information for all server certificates.
Microsoft Teams Channel [sn_disco_certmgmt_microsoft_teams_channel] (Version 4.2.0 Certificate Inventory and Management)	Stores the Microsoft Teams channel configuration.
Certificate Notification Policy sn_disco_certmgmt_notification_policy (Version 4.2.0 Certificate Inventory and Management)	Stores the Microsoft Teams channel notification settings for certificate events.

Certificate Routing Policy form

The Certificate Routing Policy form enables you to configure routing policies for CyberArk Certificate Manager SaaS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	Descriptive name for the routing policy
Certificate Authority	This value should be CyberArk Certificate Manager SaaS.
Environment	Environment where the certificate is deployed or installed.
Assignment Group	Group to which certificate tasks created for this routing policy are assigned to automatically.
Vault Type	This should be set to CyberArk PVWA.
Application Name	Unique application name from the CyberArk portal. This field is case-sensitive.
Issuing Template Alias	API alias of the issuing template from the CyberArk portal. This field is case-sensitive.
Credential Alias	Credential alias associated with the certificate authority.
Certificate Purpose	Whether the certificate request is for internal or external use.
Certificate Format	Format in which the certificate is generated. The available options are: PEM, DER, and PKCS12. The default value is PEM.
PKCS12 Password Vault Reference	Reference to the PKCS#12 key store password stored in your external vault. The format depends on your vault type: • For HashiCorp Vault, enter the full path to the secret • For Azure Key Vault, enter the secret name • For CyberArk PVWA, enter the account name This field is required when the Certificate Format field is set to PKCS12.
PKCS12 Password Vault Key	Name of the key within the vault secret that holds the PKCS#12 key store password. This field is required when the Certificate Format field is set to PKCS12 and the Vault Type field is set to HashiCorp Vault.
Is Active	Option to determine whether the routing policy is active.
Allow Duplicate Request	Option to allow duplicate requests with the same Certificate Signing Request (CSR).
Approval Required	Option to require approval before the automated flow begins.
Task Approval Group	If Approval Required is selected, the task approval group with the pki_approver role to provide approval.
Mid Server	Specific MID Server that handles all requests matching this routing policy.
Subject Common Name	Domain name secured by the certificate.

Certificate Routing Policy form for ACME

Fill in the Certificate Routing Policy form to set up the routing policy for ACME.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate routing policy

Field	Description
Name	Name of the routing policy.
Certificate Authority	Certificate Authority (CA) used to create, renew, or revoke certificates. The available options are: • DigiCert • Entrust • Let's Encrypt • EJBCA • Sectigo Public ACME • Sectigo Universal ACME **Note:** If you added an ACME-compatible CA, it is listed here as a CA type. For more information, see Add ACME-compatible certificate authorities.
Certificate Authority API URL	For DigiCert CA, the API URL to handle automated processes and revocation flows.
Assignment Group	Group to which a manual certificate task is assigned.
DNS Challenge Action	Flow designer option to resolve DNS challenges automatically (for example, base script - **ACME DNS Challenge - GoDaddy**). **Note:** You can also create a new action similar to the base script to support any other DNS provider. The flow_designer role or action_designer role is required for this action. If **None** is selected, then the new or renew certificates are done using ACME manual flow of DNS challenge.
Credential Alias	Credential alias linked to the CA credential.
Vault Type	External vault provider to use for private key storage with this routing policy. When set, the system stores private keys in the specified vault during automated certificate operations. The available options are: • HashiCorp Vault • Azure Key Vault • CyberArk PVWA

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate routing policy (continued)

Field	Description
Certificate Purpose	Indicates whether the certificate request is for an internal or external purpose.
Certificate Format	Format in which the certificate is generated. The available options are: PEM, DER, and PKCS12. The default value is PEM.
PKCS12 Password Vault Reference	Reference to the PKCS#12 key store password stored in your external vault. The format depends on your vault type: • For HashiCorp Vault, enter the full path to the secret • For Azure Key Vault, enter the secret name • For CyberArk PVWA, enter the account name This field is required when the Certificate Format field is set to PKCS12.
PKCS12 Password Vault Key	Name of the key within the vault secret that holds the PKCS#12 key store password. This field is required when the Certificate Format field is set to PKCS12 and the Vault Type field is set to HashiCorp Vault.
Allow validity override	For DigiCert ACME, Sectigo Public ACME, and Sectigo Universal ACME CAs, option to include a custom certificate validity period in the ACME order sent to the CA. When selected, the system sends the requested validity period (notAfter) to the CA. When cleared, the system omits the validity period and the CA applies its default validity. This check box is selected by default. **Note:** Clear this check box if the CA profile rejects client-requested validity.
Task Approval Group	Group to which the task approval is assigned.
DNS Task Assignment Group	Group to which the DNS challenge for this task is assigned.
Mid Server	MID Server for certificate actions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate routing policy (continued)

Field	Description
Note: The **Organization**, **Organizational Unit**, **Locality**, **State**, **Country**, and **Email Address** fields accept comma-separated values. The **Subject Common Name** and **Subject Alternative Name** fields take regex expressions. The regex format has the following restrictions: • Shouldn't contain commas. • Shouldn't start or end with a forward slash (/).	

Certificate routing policy form for EJBCA ACME

To automate the processes of your certificate life cycle, you must fill out a routing policy form that populates your Certificate Signing Requests. This table shows you the required fields and values.

Fields for a Certificate Routing Policy form

Field	Description
Name	Name of your routing policy.
Certificate Authority	Certificate authority. To see a list of supported Certificate Authorities, select the search icon. Choose **EJBCA ACME**. **Note:** The certificate authority determines if additional fields are necessary for your routing policy. Fill in this field first, to see what information you need for this form.
Environment	The environment that you want your certificate for. Options are: • Development • Disaster recovery • Production • Sub-Production .
Assignment Group	Assignment group. You can select an assignment group or make a new one.
DNS Challenge Action	Domain name system challenge action.
Credential Alias	Credential alias. Your options are based on the aliases that you create in the Certificate Management Credentials form.
Vault Type	External vault provider to use for private key storage with this routing policy. When set, the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields for a Certificate Routing Policy form (continued)

Field	Description
	system stores private keys in the specified vault during automated certificate operations. The available options are: • HashiCorp Vault • Azure Key Vault • CyberArk PVWA
Certification Purpose	Certification Purpose. Options are: • Internal • External
Certificate Format	Format in which the certificate is generated. The available options are: PEM, DER, and PKCS12. The default value is PEM.
PKCS12 Password Vault Reference	Reference to the PKCS#12 key store password stored in your external vault. The format depends on your vault type: • For HashiCorp Vault, enter the full path to the secret • For Azure Key Vault, enter the secret name • For CyberArk PVWA, enter the account name This field is required when the Certificate Format field is set to PKCS12.
PKCS12 Password Vault Key	Name of the key within the vault secret that holds the PKCS#12 key store password. This field is required when the Certificate Format field is set to PKCS12 and the Vault Type field is set to HashiCorp Vault.
Is Active	Option to activate the policy.
Allow Duplicate Request	Option to allow multiple requests to receive multiple certificates.
Approval Required	Option to require approval.
Task Approval Group	Task approval group. This field appears only when the **Approval Required** field is activated.
DNS Task Assignment Group	DNS Task Assignment Group. Select **Certificate Inventory and Management**.
Domain	Domain of the policy. This field is automatically set to *global*.
MID Server	MID Server you want to use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields for a Certificate Routing Policy form (continued)

Field	Description
Subject Common Name	Specific entity or domain name that the certificate is issued to. Enter a name or `*`.
Subject Alternative Name	Domain or subdomain included in the Subject Common Name. Enter an alternative name or `*`.
Organization	Organization making the certificate signing request for the given Subject Common Name. Enter the organization or `*`.
Locality	Locality (city) of the organization making the certificate signing request for the given Subject Common Name. Enter the locality or `*`.
Country	Country of the organization making the certificate signing request for the given Subject Common Name. Enter the country or `*`.
Organizational Unit	Organizational unit making the certificate signing request for the given Subject Common Name. Enter the unit or `*`.
State	State of the organization making the certificate signing request for the given Subject Common Name. Enter the state or `*`.
Email Address	Email address of the administrator in the organization making the certificate signing request for the given Subject Common Name. Enter an email address or `*`.

Certificate Inventory and Management patterns

Efficiently manage your Certificate Inventory with patterns to streamline the management of digital certificates. Patterns reduce the risk of vulnerabilities by bolstering security, compliance, and efficient life cycle control.

Certificate Inventory and Management patterns and descriptions

Pattern	Description
GoDaddy – Certificate Management pattern	Part of the Certificate Inventory and Management application. Allows you to import and scan for certificates from your GoDaddy certificate authority (CA).
DigiCert – Certificate Management pattern	Part of the Certificate Inventory and Management application. Allows you to import and scan for certificates from your DigiCert CA.
Entrust - Certificate Management pattern	Part of the Certificate Inventory and Management application. Allows you to import and scan for certificates from your Entrust CA.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate Inventory and Management patterns and descriptions (continued)

Pattern	Description
Sectigo - Certificate Management pattern	Part of the Certificate Inventory and Management application. Allows you to import and scan for certificates from your Sectigo CA.
Import SSL Certificate pattern	Allows you to discover certificates by importing certificates from files. Run Certificate Discovery via certificate file import.
Microsoft CA - Certificate Management pattern	Discovers certificates issued by your Microsoft CA server. For more information, see Microsoft Certificate Authority (CA) certificates discovery.
Amazon AWS - Certificate Manager pattern and **Amazon AWS - Collect Certificates Tags** shared library. Certificate Inventory and Management, at least 3.4.0. Discovery and Service Mapping Patterns, at least 1.12.0.	Discovers certificates as part of AWS Cloud discovery. AWS Certificate Manager discovery.
Google Cloud Platform (GCP) - Certificate Manager Certificates pattern and **Google Cloud Platform (GCP) - Classic Certificates** pattern. Certificate Inventory and Management, at least 3.4.0. Discovery and Service Mapping Patterns, at least 1.12.0.	Discovers certificates as part of GCP Cloud discovery. GCP Certificate Manager discovery.
Azure – Key Vault Certificates pattern. Certificate Inventory and Management, at least 3.4.0. Discovery and Service Mapping Patterns, at least 1.12.0.	Discovers certificates as part of Azure Cloud discovery. Azure Key Vault certificate discovery.
Java KeyStore and Windows Certificates	Discovers certificates in the Java KeyStore or Windows Certificate Store. Java KeyStore and Windows Certificate Store discovery.

AWS Certificate Manager discovery

Cloud Discovery uses Patterns to discover certificate data that the Amazon AWS Cloud Certificate Manager (ACM) manages. Discovering this data requires installing and updating Discovery and Service Mapping Patterns and Certificate Inventory and Management.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the configuration of an AWS account

For more information, see Discovery for cloud environment

Verify the installation of the plugins

Certificate Inventory and Management (sn_disco_certmgmt) at least 3.4.0

Discovery and Service Mapping Patterns (sn_itom_pattern) at least 1.12.0

Verify that Cloud Discovery has permissions to discover AWS

Run the following commands with AWS CLI to check the permissions and policies attached to your active IAM role:

- List Certificates

```
AWS_PAGER='' aws acm list-certificates --region
<region>
```

- Describe Certificates

```
AWS_PAGER='' awsacm describe-certificate --certificate-
arnarn:aws:acm:<region>:<accout_id>:certificate/
<certificate_id>
```

- Get Certificates

```
AWS_PAGER='' awsacm get-certificate --certificate-
arnarn:aws:acm:<region>:<account_id>:certificate/
<certificate_id>
```

- Get tags

```
AWS_PAGER='' awsresourcegroupstaggingapi get-
resources --tags-per-page 100 --resource-type-filters
'acm:certificate'
```

Set the configuration of the system property

When the system property *sn_itom_pattern.issuer_certificate_search_by_idn* is set to **false**, Discovery uses the certificate fingerprints to find issuers and root issuers.

Verify the MID Server requirements

The MID Server must have either ALL capability or AWS capability.

Verify the configuration of Cloud Discovery schedule

For more information, see Create a discovery schedule in Cloud Discovery Workspace

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected by Discovery during horizontal discovery

The **Amazon AWS - Certificates Manager** pattern and **Amazon AWS - Collect Certificates Tags** shared library support the discovery of the following table and fields.

Field	Description
Unique Certificates [cmdb_ci_certificate]	
Name	The host name/domain associated with the certificate. For example, **.service-now.com*
Fingerprint	The hash value of the certificate. For example, *d708c8583c78c176d5df1a4f01aac746294 e390a03038f280b0d8f5*
Fingerprint algorithm	The algorithm that's being used to hash the certificate. For example, Discovery calculates fingerp SHA-256 algorithm, so the value that is populated is: *SHA-256*
Serial Number	The serial number of the certificate. For example, *70 d8 c9 52 77 1c 2d 54 97 00 0e 76 b5 e8 c1 73*
Subject common name	The host name/domain associated with the certificate. For example, **.service-now.com*
Subject distinguished name	The distinguished name of the entity that the certificate is issued to. The subject distinguished name consists of the following • Common name (CN). • Organization (O)- The organization that owns the domain that the certificate is issued to. • organizational unit (OU)- The organizational unit that owns the domain that the certificate is iss
Issuer common name	The common name of the certificate issuer. For example, Entrust Certification Authority.
Issuer distinguished name	The distinguished name of the certificate issuer. The issuer distinguished name consists of the following

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• Common Name (CN)- the authority that issued the certificate. For example, *Entrust Certi Authority*. • Organization (O)- The organization that issued the certificate. For example, *"Entrust", I* • Organizational Unit (OU)- The unit that has the legal rights to issue the certificate.
Renewal tracking	Indicates whether to create any priority 1 or priority 3 tasks for the expiring certificates. Discovery sets **Renewal tracking** to *priority3* when the system property *glide.discovery.certs.enable_renewal_task_creation_for_discovered* is set to **true**.
Valid From	The certificate is valid from this date (UTC). For example, *2023-09-25 10:43:03*
Valid To	The expiry date of the certificate (UTC). For example, 2024-09-24 10:43:03
Subject organization	The organization (O) that the certificate is issued to.
Subject organizational unit	The organizational unit (OU) that the certificate is issued to.
Subject country	The country (C) of the organization that the certificate is issued to. Populated in a two-letter coun
Subject state	The region, state (ST), or province of the organization that the certificate is issued to. Populated w code.
Subject locality	The city, location (L) of the organization that the certificate is issued to.
Subject email	The email address of the organization that the certificate is issued to.
Issuer	A reference to the entity that signed and issued the certificate. The reference is available if the is a part of the same payload.
Root Issuer	A reference to the root certificate. The reference is available if the issued certificate is a part of th
Subject alternative name	The name of the certificate domain record.

CI relationships

The **Amazon AWS - Certificates Manager** pattern and **Amazon AWS - Collect Certificates Tags** shared library support the discovery of the following relationships:

CI	Relationship	CI
Unique Certificate [cmdb_ci_certificate]	Hosts::Hosted on	AWS Datacenter [cmdb_ci_aws_datacenter]
Unique Certificate [cmdb_ci_certificate]	Hosts::Hosted on	Cloud Service Account [cmdb_ci_cloud_service_account]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Key Value [cmdb_key_value]	Reference	Unique Certificate [cmdb_ci_certificate]

GCP Certificate Manager discovery

Cloud Discovery uses Patterns to discover certificate data that the GCP Certificate Manager manages. Discovering this data requires installing and updating Discovery and Service Mapping Patterns and Certificate Inventory and Management.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the configuration of an GCP account

For more information, see Discovery for cloud environment

Verify the installation of the plugins

Certificate Inventory and Management (sn_disco_certmgmt), at least 3.4.0

Discovery and Service Mapping Patterns (sn_itom_pattern), at least 1.12.0

Verify the configuration of the system property

When the system property *sn_itom_pattern.issuer_certificate_search_by_idn* is set to **False**, Discovery uses the certificate fingerprints to find issuers and root issuers.

Verify the MID Server requirements

The MID Server must have either **ALL** capability or **GCP** capability.

Verify the configuration of Cloud Discovery schedule

For more information, see Create a discovery schedule in Cloud Discovery Workspace

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected by Discovery during horizontal discovery

The **Google Cloud Platform (GCP) - Certificate Manager Certificates** pattern and **Google Cloud Platform (GCP) - Classic Certificates** pattern support the discovery of the following table and fields.

Field	Description
Unique Certificates [cmdb_ci_certificate]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	The host name/domain associated with the certificate. For example, **.service-now.com*
Serial Number	The serial number of the certificate. For example, *70 d8 c9 52 77 1c 2d 54 97 00 0e 76 b5 e8 c1 73*
Subject common name	The host name/domain associated with the certificate. For example, **.service-now.com*
Subject distinguished name	The distinguished name of the entity that the certificate is issued to. The subject distinguished name consists of the following • Common name (CN). • Organization (O)- The organization that owns the domain that the certificate is issued to. • organizational unit (OU)- The organizational unit that owns the domain that the certificate is iss
Issuer common name	The common name of the certificate issuer. For example, Entrust Certification Authority.
Issuer distinguished name	The distinguished name of the certificate issuer. The issuer distinguished name consists of the following • Common Name (CN)- the authority that issued the certificate. For example, *Entrust Certi Authority*. • Organization (O)- The organization that issued the certificate. For example, *"Entrust", I* • Organizational Unit (OU)- The unit that has the legal rights to issue the certificate.
Renewal tracking	Indicates whether to create any priority 1 or priority 3 tasks for the expiring certificates. Discovery sets **Renewal tracking** to *priority3* when the system property *glide.discovery.certs.enable_renewal_task_creation_for_discovered* is set to **true**.
Valid From	The certificate is valid from this date (UTC). For example, *2023-09-25 10:43:03*
Valid To	The expiry date of the certificate (UTC). For example, 2024-09-24 10:43:03
Subject organization	The organization (O) that the certificate is issued to.
Subject organizational unit	The organizational unit (OU) that the certificate is issued to.
Subject country	The country (C) of the organization that the certificate is issued to. Populated in a two-letter coun
Subject state	The region, state (ST), or province of the organization that the certificate is issued to. Populated w code.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Subject locality	The city, location (L) of the organization that the certificate is issued to.
Subject email	The email address of the organization that the certificate is issued to.
Issuer	A reference to the entity that signed and issued the certificate. The reference is available if the is a part of the same payload.
Root Issuer	A reference to the root certificate. The reference is available if the issued certificate is a part of th
Subject alternative name	The name of the certificate domain record.

CI relationships

The **Google Cloud Platform (GCP) - Certificate Manager Certificates** pattern and **Google Cloud Platform (GCP) - Classic Certificates** pattern support the discovery of the following relationships:

CI	Relationship	CI
Unique Certificate [cmdb_ci_certificate]	Hosts::Hosted on	Google Datacenter [cmdb_ci_google_datacenter]
Unique Certificate [cmdb_ci_certificate]	Hosts::Hosted on	Cloud Service Account [cmdb_ci_cloud_service_account]
Key Value [cmdb_key_value]	Reference	Unique Certificate [cmdb_ci_certificate]

Azure Key Vault certificate discovery

Cloud Discovery uses Patterns to discover Azure Key Vault certificates. Discovering this data requires installing and updating Discovery and Service Mapping Patterns and Certificate Inventory and Management.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the configuration of an Azure account

For more information, see Discovery for cloud environment

Verify the installation of the plugins

Certificate Inventory and Management (sn_disco_certmgmt), at least 3.4.0

Discovery and Service Mapping Patterns (sn_itom_pattern), at least 1.12.0

Verify the configuration of the system property

Verify that the *sn_itom_pattern.issuer_certificate_search_by_idn* set to **True** to configure Discovery to search issuers by issuer distinguished name only.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery can't calculate fingerprint for each certificate in the chain because Key Value Certificates PFX records don't contain certificate chains.

Verify the MID Server requirements

The MID Server must have either **ALL** capability or **Azure** capability.

Verify the configuration of Cloud Discovery schedule

For more information, see Create a discovery schedule in Cloud Discovery Workspace

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected by Discovery during horizontal discovery

The **Azure – Key Vault Certificates** pattern support the discovery of the following table and fields.

Field	Description
Unique Certificates [cmdb_ci_certificate]	
Name	The host name/domain associated with the certificate. For example, *`*.service-now.com`*
Serial Number	The serial number of the certificate. For example, *`70 d8 c9 52 77 1c 2d 54 97 00 0e 76 b5 e8 c1 73`*
Subject common name	The host name/domain associated with the certificate. For example, *`*.service-now.com`*
Subject distinguished name	The distinguished name of the entity that the certificate is issued to. The subject distinguished name consists of the following • Common name (CN). • Organization (O)- The organization that owns the domain that the certificate is issued to. • organizational unit (OU)- The organizational unit that owns the domain that the certificate is iss
Issuer common name	The common name of the certificate issuer. For example, Entrust Certification Authority.
Issuer distinguished name	The distinguished name of the certificate issuer. The issuer distinguished name consists of the following

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• Common Name (CN)- the authority that issued the certificate. For example, *Entrust Certi Authority*. • Organization (O)- The organization that issued the certificate. For example, *"Entrust", I* • Organizational Unit (OU)- The unit that has the legal rights to issue the certificate.
Renewal tracking	Indicates whether to create any priority 1 or priority 3 tasks for the expiring certificates. Discovery sets **Renewal tracking** to *priority3* when the system property *glide.discovery.certs.enable_renewal_task_creation_for_discovered* is set to **true**.
Valid From	The certificate is valid from this date (UTC). For example, *2023-09-25 10:43:03*
Valid To	The expiry date of the certificate (UTC). For example, 2024-09-24 10:43:03
Subject organization	The organization (O) that the certificate is issued to.
Subject organizational unit	The organizational unit (OU) that the certificate is issued to.
Subject country	The country (C) of the organization that the certificate is issued to. Populated in a two-letter coun
Subject state	The region, state (ST), or province of the organization that the certificate is issued to. Populated w code.
Subject locality	The city, location (L) of the organization that the certificate is issued to.
Subject email	The email address of the organization that the certificate is issued to.
Issuer	A reference to the entity that signed and issued the certificate. The reference is available if the is a part of the same payload.
Root Issuer	A reference to the root certificate. The reference is available if the issued certificate is a part of th
Subject alternative name	The name of the certificate domain record.

CI relationships

The **Azure – Key Vault Certificates** pattern support the discovery of the following relationships:

CI	Relationship	CI
Unique Certificate [cmdb_ci_certificate]	Hosts::Hosted on	Azure Datacenter [cmdb_ci_azure_datacenter]
Unique Certificate [cmdb_ci_certificate]	Hosts::Hosted on	Cloud Service Account [cmdb_ci_cloud_service_account]
Key Value [cmdb_key_value]	Reference	Unique Certificate [cmdb_ci_certificate]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Java KeyStore and Windows Certificate Store discovery

Discovery uses the **Collect Certificates** extension section of the **Linux Server** and **Windows OS – Servers** patterns to discover certificates stored in the Java KeyStore or Windows Certificate Store. Discovering the certificate information requires installing and updating Discovery and Service Mapping Patterns and Certificate Inventory and Management.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the following plugins are up to date

- Discovery and Service Mapping Patterns (sn_itom_pattern), at least 1.12.0
- Certificate Inventory and Management (sn_disco_certmgmt), at least 3.4.0

Verify that the pattern is enabled

The **Collect Certificates** pattern extension is disabled by default. On the **Discovery Patterns** list, verify that the **Collect Certificates** pattern extension has the **Active** field set to **true**.

Verify the configuration of the Java KeyStore credentials

In the **Discovery Credentials** module, select **jks credentials**. Make sure you configured the Discovery IP Affinity. In the **Type** field, insert `java_keystore`. For more information, see Get started with credentials

Verify the configuration of the Discovery schedule

For more information, see Running discoveries in your network

Data collected by Discovery during horizontal discovery

The **Collect Certificates** extension section of the **Linux Server** and **Windows OS – Servers** patterns support discovery of the following tables and fields.

Unique Certificates [cmdb_ci_certificate] table

Field	Description
Fingerprint [fingerprint]	The certificate fingerprint number.
Serial Number [serial_number]	The serial number of the certificate.
Signature Algorithm [signature_algorithm]	The signature algorithm of the certificate.
Fingerprint Algorithm [fingerprint_algorithm]	The fingerprint algorithm of the certificate.
Issuer distinguished name [issuer_distinguished_name]	The distinguished name of the certificate issuer. The issuer distinguished name consists of the following

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Unique Certificates [cmdb_ci_certificate] table (continued)

Field	Description
	• Common Name (CN)- the authority that issued the certificate. For example, *Entrust Certification Authority*. • Organization (O)- The organization that issued the certificate. For example, *"Entrust", Inc*. • Organizational Unit (OU)- The unit that has the legal rights to issue the certificate.
Subject distinguished name [subject_distinguished_name]	The distinguished name of the entity that the certificate is issued to. The subject distinguished name consists of the following • Common name (CN). • Organization (O)- The organization that owns the domain that the certificate is issued to. • organizational unit (OU)- The organizational unit that owns the domain that the certificate is issued to.
Subject common name [subject_common_name]	The host name or domain associated with the certificate.
Subject country [subject_country]	The country (C) of the organization that the certificate is issued to. Populated in a two-letter country code.
Subject State [subject_state]	The region, state (ST), or province of the organization that the certificate is issued for. This data is populated with two-letter code.
Subject locality [subject_locality]	The city, location (L) of the organization that the certificate is issued for.
Subject organizational unit [subject_organizational_unit]	The organizational unit (OU) that the certificate is issued for.
Subject organization [subject_organization]	The organization (O) that the certificate is issued for.
Subject email [subject_email]	The email address of the organization that the certificate is issued for.
Valid From [valid_from]	The certificate is valid from this date (Displayed in your time zone.). For example, *2023-09-25 10:43:03*
Valid To [valid_to]	The expiry date of the certificate (Displayed in your time zone.). For example, 2024-09-24 10:43:03
State [state]	When Discovery collects the certificate the field indicates **installed**
Subject Alternative Name [subject_alternative_name]	A list of references to alternative host names used by the server.
Issuer [issuer]	In this field, there's a reference to the certificate record. The issuer is the entity that signed and issued the certificate. The

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Unique Certificates [cmdb_ci_certificate] table (continued)

Field	Description
	reference is available if the issued certificate is a part of the same payload.

Installed Certificate [sn_disco_certmgmt_cmdb_installed_certificate] table

Field	Description
Discovery Method [discovery_method]	If a pattern discovered this installed certificate, the field indicates **Pattern**.
Source [source]	The IP address of the host server on which discovery is run.
Server [server]	The Java-based web server or Windows based web server that store the certificate in java KeyStores on the host machine.
Root Issuer [root_issuer]	A reference to the root certificate record. The reference is available if the issued certificate is a part of the same payload.

Certificate Inventory and Management terms

The Certificate Inventory and Management glossary comprises a set of concise definitions and explanations for terms related to the tracking, storage, and management of digital certificates within an IT environment.

Certificate Authority (CA)

Entity that issues a digital certificate.

Certificate Fingerprint

Carries the unique information of the certificate, meta data.

Discovery Method - CA (Version 1.1.7 Certificate Inventory and Management)

TLS certificates belonging to this category of discovery method are discovered from a Certificate Authority such as GoDaddy or DigiCert.

Discovery Method - Certificate Import (Version 1.1.7 Certificate Inventory and Management)

TLS certificates belonging to this category of discovery method are discovered by importing the files into the system using pattern-based Discovery.

Discovery Method - Port

TLS certificates belonging to this category of discovery method are discovered from remote servers listening ports.

Discovery Method - URL

TLS certificates belonging to this category of discovery method are discovered from URLs.

External Certificate

Indicates that the digital certificate is signed by either root or intermediate CA.

Internal certificate or self-signed

Indicates that the digital certificate is signed by the same entity whose identity it certifies. For example, internal or self-signed certificates aren't signed by root or intermediate CAs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Transport Layer Security (TLS)

Successor protocol to SSL. TLS is an improved version of SSL. It works in much the same way as the SSL, using encryption to protect the transfer of data and information.

Trust manager

Validates the TLS certificate.

Trust store

Repository to store CA certificates.

Unique vs installed certificates

Unique: Each unique instance or fingerprint of certificate.

Installed: Shows where each discovered certificate is installed.

Certificate authorities pattern API elements and permissions

Explore the Certificate Authorities (CA) pattern API elements and permissions for comprehensive insights into the associated functionalities and authorization requirements.

API elements and permissions

Certificate Authorities	API element
GoDaddy	• * /v1/certificates • * /v1/certificates/{certificate_id}/download
DigiCert	• * /services/v2/order/certificate • * /services/v2/certificate/{certificate_id}/chain
Entrust	• * / v2/certificates • * / v2/certificates/{certificate_id}
Sectigo	• * /cert-manager/api/ssl/v1 • * /cert-manager/api/ssl/v1/collect/{certificate_id}/pem

The patterns available with Certificate Inventory and Management Version 1.1.7 are:

- GoDaddy
- DigitCert

The patterns available with Certificate Inventory and Management Version 1.2.0 are:

- Entrust
- Sectigo

Setting Your IP Addresses form and fields table

A table of the fields and values to establish the IP address, range of IP addresses, or list of IP addresses for Certificate Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Setting Your IP Addresses Form and Fields

Field	Description
Name	Enter a name for your IP Discovery search.
Active	Select the check box.
Type - IP address list	Enter each IP address, separated by a comma.
Type - IP address range	If you select this option for Type, two fields appear: **Starting IP** and **Ending IP**.
Type - IP Network	The Network IP field appears. Enter the IP address of the network. For more information, see Discovery IP address configuration.

IP-Based Discovery Schedule form and fields table

A table of the fields and values to set up an IP-based Certificate Discovery schedule.

IP-Based Discovery Schedule form and fields Table

Field	Description
Name	Enter a name for your IP Based Discovery Schedule.
Discover	Select **Configuration Items**.
MID Server selection method	The default value is **Specific MID Server**.
MID Server	Select a MID Server. ⚠ **Warning:** Warning: Make sure that your IP address is accessible to the MID Server you're using. If it doesn't, you have to Credential aliases for Discovery.
Created by Quick Start	Check box not selected
Application	Certificate Inventory and Management
Active	Check box selected
Location	Optional value
Run	Set the frequency that you want the Discovery to run. Daily, Weekly, Monthly, and so on
Credential alias	Optional value

Quick start test for Certificate Inventory and Management

Quick start test for certificate discovery for GoDaddy, Entrust, Digicert, Setigo.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate Inventory and Management - quick start test

Test	Description
Certificate Discovery	Certificate discovery for GoDaddy, Entrust, Digicert, and Setigo. Discovery brings the certificates and persist them in cmdb_ci_certificate table. This test validates the certificates and CIs in the cmdb_ci_certificate tables.

To learn more about Certificate Inventory and Management, see Certificate Inventory and Management.

Related topics

Quick start tests

Discovery schedule form table

To perform Discovery on root certificates stored outside your server, you must set a special Discovery schedule that locates these certificates. This table guides you through the form required for this task.

Set a NewDiscovery Schedule for Certificates

Field	Value
Name	Name for your schedule.
Discover	**Certificates**. **Note:** Once you fill the Discover field, the **Certificate Discovery Type** field appears.
Certificate Discovery Type	**Import Certificates**.
MID Server selection method	**Specific MID Server**.
MID Server	The MID Server you want to perform this discovery.
Credential Alias	Your credential alias. If you must set one up, Create a Connection & Credential alias.

Certificate request form

The Request New Certificate (Automated) and Renew Certificate (Automated) forms enable you to submit or update a Certificate Signing Request (CSR) to be submitted to a certificate authority.

Field	Description
Certificate Purpose	Determines whether the Certificate is internal or external.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• Internal: Requires a completed form • External: Requires importing the CSR and private key from an external source
Environment	Environment where the certificate must be deployed or installed. The available options are Development, Disaster recovery, Production, or Sub-Production.
Validity Period for Certificate (In Days)	The number of days the certificate is valid. This field appears only when External is selected from the Certificate Purpose field.
Choose how to generate CSR	Determines which CSR is used to generate the certificate. The available values are: • Use external CSR • Generate CSR in CyberArk: Available only when CyberArk Certificate Manager SaaS is selected as the **Select how to manage your certificate** value when requesting or revoking a certificate.
Certificate Signing Request (CSR)	CSR sent to the external Certificate Authority (CA) to request the certificate.
Subject Common Name	Entity or domain name the certificate is issued to.
Subject Alternative Name	Domain or subdomain included in the Subject Common Name.
Organization	Organization submitting the CSR.
Organizational Unit	Organizational unit submitting the CSR.
Locality/City	Locality (city) of the organization submitting the CSR.
Province	State or province of the organization submitting the CSR.
Country	Country of the organization submitting the CSR.
Email Address	Email address of the administrator in the organization submitting the CSR.
Key Algorithm	Cryptographic algorithm for generating the certificate key pair, based on algorithms supported by CyberArk. The available options are RSA 1024, RSA 2048, RSA 3072, RSA 4096, EC P256, EC P384, EC P521, and EC ED25519.
Application Services	Application services for the certificate.
Servers	Servers the certificate can be hosted on.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Application Names	Specific application the certificate is issued for.
Application Servers	Application servers the certificate is hosted on.
Certificate Owner Group	Certificate owner group for the certificate.
Certificate Owner	Specific entity that the certificate is issued to.
Renewal Tracking	Determines whether to track certificate renewal automatically. The available options are Create priority 1 tasks, Create prioirty 3 tasks, and Do not create renewal tasks.
Renew Automatically	Option to automatically renew the certificate.
How many days before expiry does the certificate need to be renewed?	Number of days before certificate expiration to trigger renewal.

Bulk certificate upload template fields

The bulk certificate upload Excel template contains fields that map to the fields in the Unique Certificate [cmdb_ci_certificate] table. You can import up to 5000 SSL certificates in bulk using the Excel file.

Bulk certificate upload template

Field	Required	Description
valid_to	Yes	Expiry date of the certificate. Must be greater than the *valid_from* value.
root_issuer	Yes	Reference of the root certificate in the certificate chain. For a self-signed certificate, enter the certificate's own fingerprint.
subject_common_name	Yes	Host name or domain associated with the certificate. For example, `www.servicenow.com`.
valid_from	Yes	Date from which the certificate is valid.
state	Yes	State of the certificate. Enter `issued`, `installed`, `revoked`, `retired`, or `other`.
fingerprint	Yes	Hash value of the certificate. Enter the value as space-separated hexadecimal pairs.
fingerprint_algorithm	Yes	Algorithm used to hash the certificate. For example, `SHA-256`.
issuer	Yes	Reference of the certificate that signed and issued the certificate. For a self-signed certificate, enter the certificate's own fingerprint.
issuer_common_name	Yes	Common name of the certificate issuer. For example, `Entrust Certification Authority - L1K`.
issuer_distinguished_name	Yes	Distinguished name of the certificate issuer. Includes the common name (CN) of the issuing authority, the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Bulk certificate upload template (continued)

Field	Required	Description
		organization (O) that issued the certificate, and the organizational unit (OU) with the legal rights to issue it.
serial_number	No	Serial number of the certificate. Enter the value as space-separated hexadecimal pairs.
signature_algorithm	Yes	Signature algorithm of the certificate. For example, `SHA256WITHRSA`.
subject_alternative_name	No	Alternative host names covered by the certificate. Enter a comma-separated list.
subject_country	No	Country (C) of the organization that the certificate is issued to, as a two-letter country code.
subject_distinguished_name	Yes	Distinguished name of the entity that the certificate is issued to. Consists of the common name (CN), the organization (O) that owns the domain, and the organizational unit (OU) that owns the domain.
subject_email	No	Email address of the organization that the certificate is issued to.
subject_locality	No	City or location (L) of the organization that the certificate is issued to. For example, `San Diego`.
subject_organization	No	Organization (O) that the certificate is issued to. For example, `ServiceNow, Inc.`
subject_organizational_unit	No	Organizational unit (OU) that the certificate is issued to.
subject_state	No	Region, state (ST), or province of the organization that the certificate is issued to. For example, `California`.
version	No	X.509 version of the certificate. For example, `3`.
key_size	Yes	Size of the certificate key, in bits. For example, `2048`.
is_ca	No	Whether the certificate is a certificate authority certificate. Enter `TRUE` or `FALSE`.

Service Mapping

The ServiceNow® Service Mapping application discovers all application services in your organization and builds a comprehensive map of all devices, applications, and configuration profiles used in these application services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping added value

Service Mapping is available as a separate subscription and requires activation by Service Mapping personnel.

Gain control with service-aware operations

Service Mapping enables IT departments of companies, organizations, and cloud companies providing platform as a service to create a service-aware view of infrastructure.

Reach time to value quickly

Shorten the mapping process and improve accuracy with out-of-the-box options. Service Mapping deploys different methods for collecting information about configuration items (CIs) and organizing them into application services. The available mapping methods are: pattern-based, tag-based, traffic-based, and discovery based on Predictive Intelligence.

Deliver on your multicloud strategy

Map services across dynamic cloud-based providers with a high degree of accuracy.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Explore Learn about the various methods of mapping application services.	Configure Plan and configure your implementation.	
Use 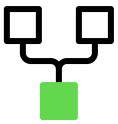 Review mapping application services for analysis and maintenance.	Reference Get details about components like roles, tables, and properties.	

Exploring Service Mapping

Service Mapping discovers all service instances in your organization and builds a comprehensive map of all devices, applications, and configuration profiles used in these service instances.

Who uses Service Mapping

Service Mapping enables IT departments of companies, organizations, and cloud companies providing platform as a service to create a service-aware view of infrastructure.

How do you use Service Mapping

A service instance is a set of interconnected applications and hosts that are configured to offer a service to the organization. Service instances can be internal, like an organization email system or customer-facing, like an organization website. For example, creating financial reports through a web-based application requires a computer, web server, application server, databases, middleware, and network infrastructure. These applications and hosts are all configured to offer the service of financial reporting.

Typically, IT departments create and maintain an inventory that treats devices and applications as standalone, independent objects. Connections between the devices and applications are not included. This is usually referred to as horizontal discovery. This method does not address the biggest challenge for IT departments, which is understanding the connection and dependencies between each object.

Service Mapping maps dependencies, based on a connection between devices and applications. This method is referred to as top-down mapping. The top-down mapping helps you immediately see the impact of a problematic object on the rest of the service instance operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison of horizontal and top-down mapping results

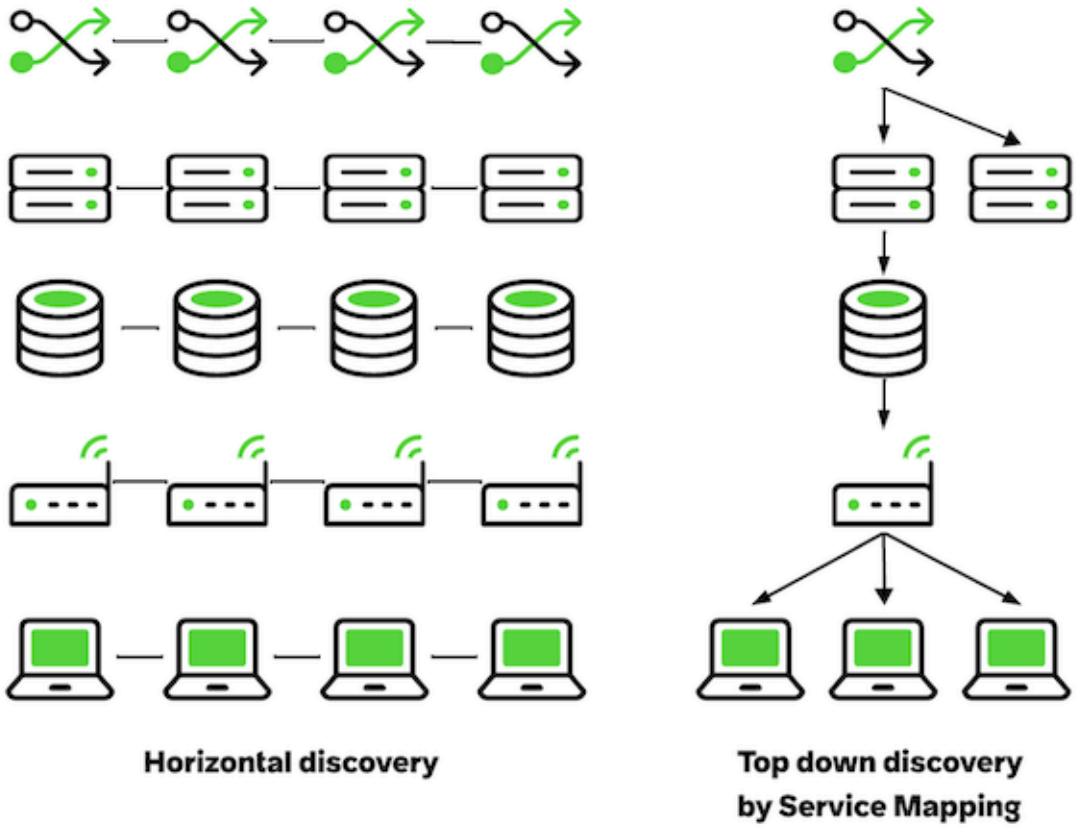

Service instance maps show infrastructure objects and semantic connections between them. Service Mapping regenerates service instance maps regularly, to keep them updated and relevant. Any faulty objects are shown along with the devices and applications they affect, providing a visual clue of the state of the service instance.

How does Service Mapping work

Service Mapping can deploy different methods for creating application services. The main method of Service Mapping discovering and mapping devices and applications is using patterns. A pattern is a sequence of operations whose purpose is to detect attributes of devices and applications and their outbound connections. If your organization uses tags for asset management, you can use these tags to map service instances. Service Mapping can also discover devices and applications by following traffic connections between them. This method is referred to as traffic-based mapping. For more information, see Choose the right method for discovering and mapping services.

What to know before you begin

You can use Service Mapping only if the Discovery product is activated and set up.

You must define users and configure credentials to allow Service Mapping and Discovery access to applications and devices inside your organization network. For more information, see Configuring Service Mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping on the ServiceNow AI Platform

Service Mapping relies on Discovery and the MID Server to discover devices and applications. Service Mapping uses results of horizontal discovery performed by Discovery. The MID Server facilitates communication between Service Mapping and devices and applications it discovers.

Data collected and organized by Service Mapping is visible in Event Management, Dependency Views, and Application Portfolio Management (APM). With Event Management, you can view events to take actions for recovering your organization application services. Dependency Views shows relationships between devices and applications in the context of application services they belong to. When integrated with APM, Service Mapping provides information about components making up a business application and helps APM users to monitor business application performance.

Service Mapping supports domain separation. If your ServiceNow AI Platform uses domain separation, administrators and users can only see and manage application services belonging to their own domain.

To view a list of Service Mapping plugin dependencies, see Plugins or applications installed with ITOM Visibility.

Related topics

- Application services
- Event Management
- Dependency Views

Choose the right method for discovering and mapping services

Service Mapping deploys different methods for collecting information about configuration items (CIs) and organizing them into application services. The available mapping methods are: pattern-based, tag-based, traffic-based, and discovery based on Predictive Intelligence. Learn about the mapping methods to use the ones that best suit the needs of your organization.

Mapping method	Description	Advantages	Disadvantages
Automated Service Suggestions	Predictive Intelligence evaluates connections between application fingerprints, CIs, and processes, and ranks their relevancy. Service Mapping uses this information to create connections based on connection rules. It also generates connection suggestions for servers and load balancers for you to decide which connections to add or remove from the service instances.	Service Mapping automatically adds CI connections and CIs they lead to based on connection rules. You can modify the default connection rules or create your own to enhance this discovery method. Discovery based on Predictive Intelligence is effective for mapping multiple application services.	When using discovery based on Predictive Intelligence to identify TCP connections and processes, deploy it in large on-premise deployments. Predictive Intelligence is a machine learning solution and needs training. To achieve best results using this discovery method, you must run discovery repeatedly. Service Mapping may erroneously add or remove dynamic connections.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Mapping method	Description	Advantages	Disadvantages
Pattern-based discovery	Service Mapping uses patterns to discover and map CIs. A pattern is a sequence of commands designed to detect attributes of a CI and its outbound connections. A typical Service Mapping pattern consists of separate algorithms for identifying CIs and finding CI connections.	Pattern-based discovery creates precise and complete service instances that reliably represent the service-aware view of your organization's IT infrastructure.	Pattern-based mapping requires configuring credentials, users, and user permissions to let Service Mapping access applications inside your organization private network. This process may take time and effort.
Tag-based discovery	If your organization uses tags for asset management, you can use these tags to map service instances. Discovery and Cloud Provisioning and Governance discover tags assigned to CIs, and populate the CMDB with this data. Service Mapping uses the tag-related data from the CMDB to map services.	You can effectively use tags to map multiple application services. Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.	Service instances mapped using tags might not include relevant CIs if those CIs do not have the correct tags.
Traffic-based discovery	Service Mapping can discover and map configuration items (CIs) following their traffic-based connections. This method is referred to as traffic-based mapping and complements pattern-based mapping. Depending on your configuration, the	Using traffic-based discovery is like casting a finer net, allowing Service Mapping to find even those CIs that it failed to discover using patterns.	While using traffic-based discovery creates a more inclusive map, it may also result in mapping many redundant CIs that do not influence the application service operation. It may clutter an application service with irrelevant CIs. Use traffic-based discovery at the initial stages of discovering

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Mapping method	Description	Advantages	Disadvantages
	behavior of traffic-based discovery varies.		an application service and disable it once you completed discovery and fine-tuned the application service.

Multi-source service mapping

Multi-source service mapping combines data from multiple discovery methods to create a single, comprehensive service map that provides complete visibility into your organization's hybrid IT infrastructure.

What is multi-source service mapping

Multi-source service mapping enables you to create unified services. You can create unified services combining mapped and unmapped services from different data sources in your CMDB. Instead of managing separate service maps created through different discovery methods, you can add them into a single composite service that represents your organization's interconnected service architecture.

This feature addresses the challenge of managing complex IT environments where a single service instance might be discovered through multiple methods, each providing partial visibility. For example, some components can be discovered through pattern-based top-down discovery, while others are more effectively identified through tag-based Service Mapping. Multi-source service mapping brings together data from all available sources to create a complete picture.

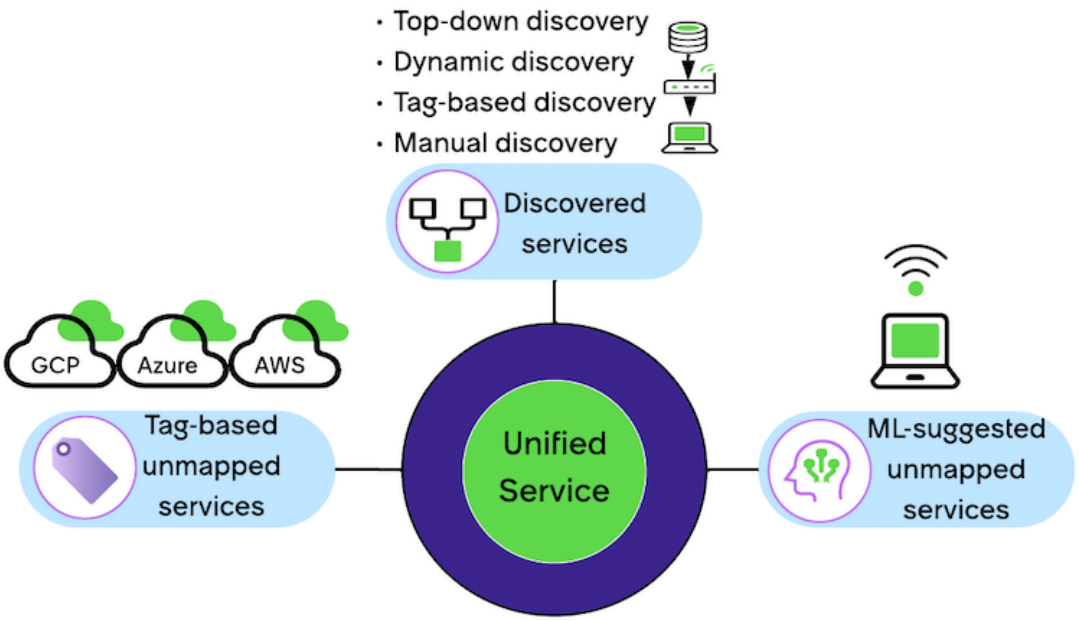

Supported data sources

Multi-source service mapping can combine data from the following sources:

Discovered maps

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Horizontal or top-down discovery that identifies configuration items (CI) and their connections using Discovery and Service Mapping Patterns. These methods involves running discovery using a MID Server.

Tag-based maps

Services mapped using cloud-native tags and metadata from your asset management system. This method is effective for cloud environments where tags are used for organization and categorization.

ML-powered maps

- Machine learning-based discovery that evaluates connections between application fingerprints, CIs, and processes. This method provides additional data through connections and processes discovered during monitoring.
- Traffic-based mapping is based on network traffic connections between devices and applications. This method helps identify communication patterns and dependencies based on actual network behavior.

Manual service maps

Manually selecting specific CIs without automation. By converting manual services to a dynamic application services, you can synchronize the manually created CI relationships into a Service Map view. These maps are maintained by administrators and are best fit for static services. For more information, see: Convert legacy manual services to dynamic application services

Key benefits

Multi-source service mapping offers several advantages for organizations managing complex, hybrid IT environments:

Complete visibility

Your organization can gain broader and more detailed visibility of application dependencies across your entire infrastructure, including on-premises, cloud, and hybrid environments. Different discovery methods excel at finding different types of components, and combining them ensures that nothing is missed.

Reduced duplication issues

Your organization can eliminate redundant service definitions by consolidating multiple partial views into a single authoritative service map. The same CIs often appear across different service maps, and unification removes this redundancy. By including unmapped services, the service maps reflect a single source of truth.

Improved incident response

Service-level actions can improve Mean Time to Resolve (MTTR) and business outcomes compared to CI-level fixes, as teams work from a unified view rather than fragmented information.

Business context and reporting

Enriching unified services with business context details such as business applications, parent services, and service offerings enables better visibility into how technical services support business operations. This improves reporting and analytics.

Use cases

Multi-source service mapping is valuable in the following scenarios:

Incident Management

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When a single server appears in multiple separate service maps, its failure can trigger multiple teams to work on the same issue. A unified service map provides a complete dependency picture, reducing Mean Time to Resolve (MTTR) and helping to prevent duplicate work across teams.

Change Management

When a database schema change is reflected in a single service map but not another, it increases the risk of unplanned outages. Multi-source mapping can promote visibility to all dependencies, enabling accurate analysis of impact across the entire service.

Reporting and Compliance

A unified view provides more accurate statistics and reports for Service Level Agreement (SLA) compliance than manually correlating data from multiple separate service maps.

Multi-source service-mapping workflows

The multi-source service-mapping process follows these key phases:

Creating a unified service

Service Mapping administrators can access this flow through the Service Mapping workspace. They can search for services based on properties related to the services, such as server tag, or UR and choose the services to unify. They can choose mapped and unmapped services from various types. For example, they can unify a top-down service with a tag-based service. Unmapped services can be added to the unified service and converted to mapped services along the process. A ranking algorithm suggests the top recommendations of services to merge, according to the criteria provided.

Note: The feature supports unification of 10 services maximum, by default. Service Mapping administrators can change the limit using the sn_sm_scoped_app.sm.unified_services.max_children_per_service system property.

After creating the unified service, the feature automatically transforms the selected unmapped services into non-operational service instances. This happens in order to prevent automatic discovery and impact on ITOM or ITSM processes. The Service Mapping administrator can choose to convert the services to operational using the **Mapped application services** list.

When the process of creating the unified service is complete, the unified service map is displayed in the "Unified map" format through the CMDB Workspace. This shows the complete consolidated view.

Editing the unified service

The Service Mapping administrator can add additional services from various types or remove a service if needed. The feature doesn't support unification of unified services, only single services.

Deleting the unified service

The Service Mapping administrator can delete the unified service if needed. Child services remain non-operational by default.

Getting Started

To use multi-source service mapping, your ServiceNow instance must have the following setting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- At least Australia family version installed.
- Service Mapping Plus installed and updated to the latest version.
- At least two existing services to unify, or at least two sources of data that are actively updated. If there are no existing services on your instance, choose a method to create the services. For more information, see: Choose the right method for discovering and mapping services.

Service Mapping with Agent Client Collector

Agent Client Collector enables the discovery and mapping of application services without requiring you to set up credentials for the MID Server.

Using Agent Client Collector, you can deploy both top-down and machine learning-based Service Mapping and benefit from a simplified approach to mapping application services.

For information on configuring Service Mapping to work with Agent Client Collector, see Configure Agent Client Collector for use with Service Mapping.

CMDB-based mapping

Using data stored in the Configuration Management Database (CMDB), you can run top-down discovery and map your organization's application services without having to access a MID Server.

Top-down discovery begins with an IP address and retrieves the associated host. It then tries to locate an available MID Server. If a MID Server isn't available and CMDB-based mapping is active, it is initiated. This feature uses incoming traffic already registered in the CMDB to pinpoint the target process running on an identified IP address. It then searches for an application linked to that process. The application is then mapped as a node on the service map.

If an application isn't identified, CMDB-based mapping uses Application Dependency Mapping (ADM) logic to classify and create an application linked to the target process within the CMDB.

If ADM is unable to classify the application, it creates a simple application in the Application [cmdb_ci_appl] table in the CMDB. For more information, see Application Dependency Mapping (ADM) for Discovery.

For information on configuring Service Mapping for CMDB-based mapping, see Configure Service Mapping to work with the Configuration Management Database (CMDB).

Automated Service Suggestions

Automated Service Suggestions simplifies the processing of mapping application services or adding to an existing service by providing application service candidates using machine learning. Mapped services provide key information to help you make data-driven decisions, prioritize services, and respond to situations.

Using machine learning (ML), Service Mapping quickly identifies potential application fingerprints based on discovered processes. It evaluates the relationship between these fingerprints and configuration items (CIs) and creates meaningful connections. These connections are used to create an application service map, a visual representation of a service.

Automated Service Suggestions further simplifies the mapping process by using the ML information to calculate entry points, efficiently identifying where to begin mapping an application service. It also generates a list of application service candidates. Each candidate includes an entry point as well as a set of previously identified resources to include in the service map. You can then select a candidate to create a new application service or to add to an existing service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Fingerprints

Service Fingerprints complements Automated Service Suggestions by identifying unique, classified components of identified application service candidates. Presenting these components facilitates the selection of the most suitable suggested application service. You can view components of an application in your actual environment, including details such as product, product category, publisher, and description. Use this information to search for specific components you want to discover or to make informed decisions about which candidates to convert into application services.

To use Service Fingerprints, you must download and install ITOM Content Service from the ServiceNow® Store and opt in to ITOM Content Service. For more information, see ITOM Content Service.

Domain separation with Automated Service Suggestions

Automated Service Suggestions supports domain separation starting with version 1.9.0. No support is currently offered for on-premise instances. Reach out to product management for further information.

Related topics

Exploring Service Mapping

Service Mapping Workspace

Application service readiness dashboard in configurable workspace

Map application services based on Automated Service Suggestions

Discovery based on Predictive Intelligence

Service Mapping uses data processed by Predictive Intelligence to generate suggestions for traffic-based connections.

Connection suggestions generated using Predictive Intelligence

Predictive Intelligence evaluates connections between application fingerprints, CIs, and processes, and ranks their relevancy. Service Mapping uses this information to create connections based on connection rules. It also generates connection suggestions for servers and load balancers for you to decide which connections to add or remove from the service instances. The confidence level indicates the likelihood of this connection being part of the service instances.

- **High** - Commonly a connection to an internal component that is only used by this specific component or application.
- **Medium** - Commonly a connection to a middleware component used by multiple services, such as IBM MQ or Oracle WebLogic Server.
- **Low** - Commonly a connection to an application deployed on numerous servers in the organization, typically for monitoring purposes, such as Tivoli Enterprise Monitoring Agents or MID Server.
- **Very low** - Commonly a connection to a central application used by the entire organization, for example, Microsoft Active Directory or Okta.

If Predictive Intelligence is enabled, Service Mapping automatically adds connections to service instances based on connection rules. These suggestions are generated using traffic-related data from the Configuration Management Database (CMDB) and the analysis of application fingerprints, CIs, and processes by Predictive Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection rules based on connection suggestions

Connection rules enhance the connection suggestions generated using Predictive Intelligence. Each rule contains a condition that describes an aspect of a traffic-based connection. The fields available in conditions come from the Connection Suggestion [sa_ml_connection_suggestion] table, which is only populated when you run top-down discovery. You can base your rules directly on the connection suggestions. For example, create a rule for any connection with confidence level High. Alternatively, if you know which CIs must be part of the service instances, you can choose fields that point at these CIs.

When Service Mapping adds a connection based on connection rules, it performs the following actions:

1. Associating the connection rule with the connection suggestion.
2. Changing the connection suggestion *`Decision`* attribute to `Added By Rule`.
3. During service rediscovery, adding one level of the actual CI connections to relevant application services.

Every time application services are rediscovered, Service Mapping adds only one level of CI connections. You may want to change the frequency in your discovery schedules to complete mapping of the application services using connection rules.

Service Mapping evaluates only active connection rules and only for connection suggestions that have the Undecided decision attribute. The system applies the rules using the following algorithm:

1. Service Mapping checks local rules for a CI connection. If multiple rules are relevant for the same CI connection, Service Mapping uses the Order attribute for the sequence: It evaluates the rules from least to greatest Order value. By default, Service Mapping checks only 20 local rules.
2. If a connection suggestion matches the connection rule condition, Service Mapping adds this connection to the relevant application service.
3. If Service Mapping finds no match for local rules, the system starts evaluating and applying global rules in ascending order using the Order attribute. By default, the system checks only 20 global rules.
4. If a connection suggestion matches the rule condition, the system adds it to all application services. If a rule is configured to exclude suggested connections from certain services, the system does not apply this rule to these services.
5. Every time Service Mapping rediscovers application services of the manual or discovered type, it validates automatically added traffic-based connections against the active connection rules. That way the system ensures that application services contain only relevant connections based on the current and active connection rules.
6. If a rule is no longer valid, Service Mapping changes the decision attribute for its corresponding connection suggestion to Undecided. Service Mapping also removes the association between the connection rule and the connection suggestion. Unless Service Mapping can apply a different connection rule to CI connections that became irrelevant, Service Mapping removes them from the application services.

If necessary, you can modify the number of the local or global connection rules that the system uses in the algorithm, as described in Properties installed with Service Mapping.

Enabling discovery based on Predictive Intelligence

Discovery based on Predictive Intelligence works at the product level. The *`sa_ml.connection_suggestions.active`* property controls this feature. If traffic-

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

based discovery was enabled in your deployment prior to the Quebec release and you used it to discover at least one application service, discovery based on Predictive Intelligence is disabled by default.

Related topics

- Add CIs to application services using connection rules
- Add CIs to an application service using connection rules
- Manage CI connections in bulk using suggestions
- Manage CI connections for a single service using suggestions

Pattern-based discovery in Service Mapping

Pattern-based discovery is the main method of Service Mapping collecting data about devices and applications used in application services. After Service Mapping collects data, it then creates a map of application services and stores the collected data in the CMDB.

ServiceNow applications refer to devices and applications that comprise a service instance as configuration items (CIs).

Service Mapping uses patterns to discover and map CIs. A pattern is a sequence of commands designed to detect attributes of a CI and its outbound connections. A typical Service Mapping pattern consists of two types of algorithms for identifying CIs and finding CI connections.

The starting point of any discovery process is an entry point. An entry point is a point where clients access a service instance. For example, to map your electronic mailing service instance, define an IP address or host name of the email server as an entry point. The discovery and mapping process begins from Discovery performing the horizontal discovery to identify the host. Once the host discovery is complete, Service Mapping starts the top-down discovery to find and map applications running on this host.

Service Mapping uses MID Servers to communicate with CIs in your organization. MID Servers are located inside your organization network and Service Mapping can communicate with them without traversing firewalls.

Note: Currently, pattern operations do not support multi-languages. If values returned from pattern operations are not in English, the returned data cannot be parsed properly and the pattern discovery will fail.

The discovery and mapping process consists of the following interactions:

1. An administrator defines an service instance with an entry point for an application CI.

   Service Mapping creates a record for the new service instance in the Mapped Application Service [cmdb_ci_service_discovered] table.

2. The device hosting the application is identified.
   a. Service Mapping checks if the device hosting this application CI exists in the CMDB.
   b. If the device hosting this application CI does not exist, Service Mapping triggers Discovery to detect host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

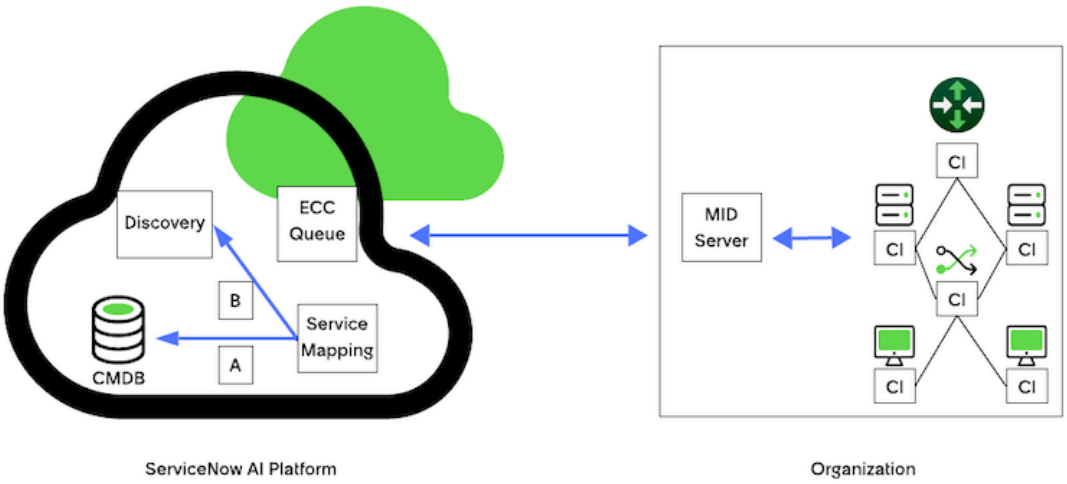

c. Discovery creates the first set of probes for port discovery, referred to as Shazzam probes, and places them as a discovery request in the External Communication Channel (ECC) queue.

d. The MID Server checks the ECC queue and retrieves the discovery request assigned to it.

e. The MID Server runs the probes against the host and discovers open ports.

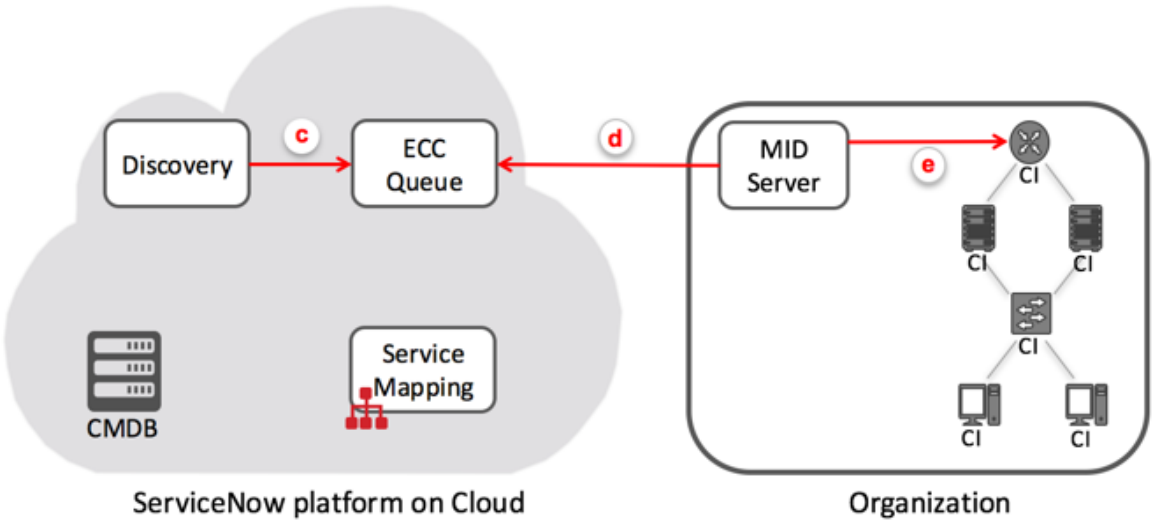

f. The MID Server passes information on the host ports to the ECC queue.

g. Discovery checks the ECC queue and receives information on the host ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

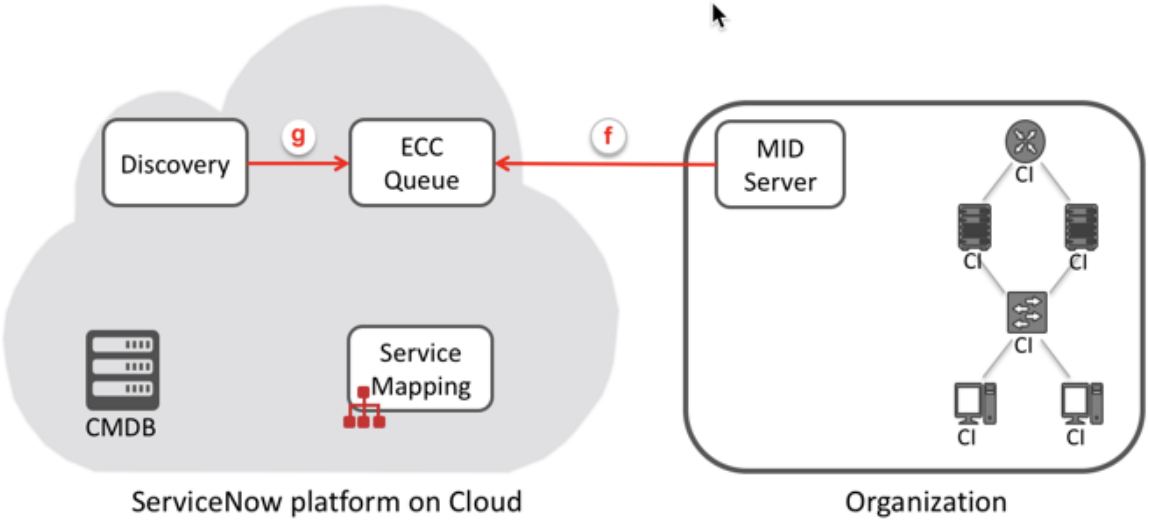

h. These steps are repeated for other types of probes: classification, identification, and exploration.

i. Discovery adds the host to the CMDB.

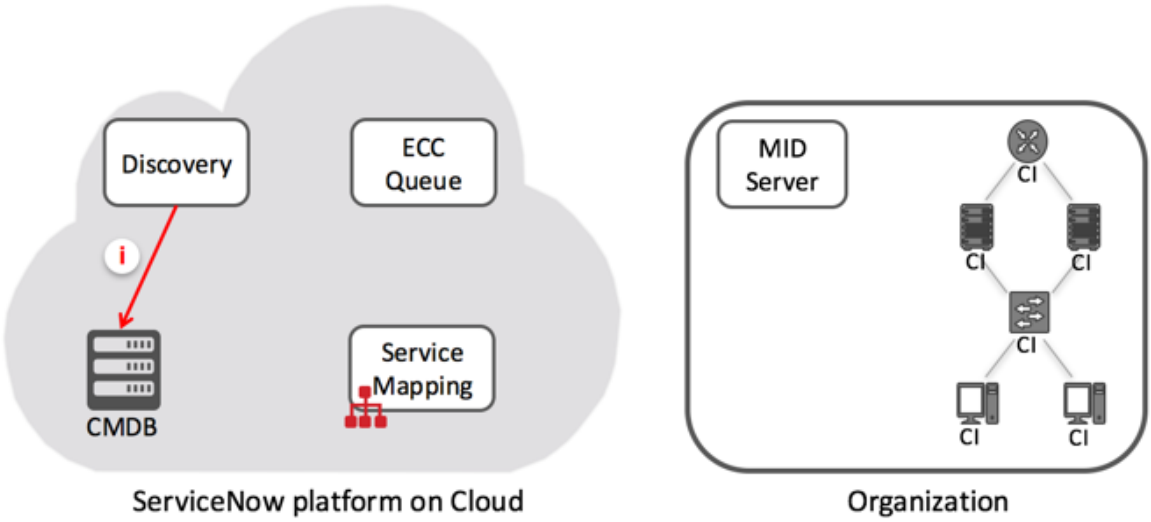

j. During the host discovery using probes, Service Mapping checks the ECC queue if this process is complete. When the host discovery is complete, Service Mapping checks whether this host exists in the CMDB.

Note: For the detailed description of the horizontal discovery flow, refer to Horizontal discovery process flow with probes and sensors.

3. Once the host is found in the CMDB, Service Mapping discovers the application running on this host.

a. Service Mapping creates an application discovery request for the IP address of the entry point. It then writes the request in the ECC queue and assigns a MID Server to the request.

b. The MID Server checks the ECC queue and retrieves the discovery request assigned to it.

c. The MID Server starts running identification sections of the patterns associated with the classifier to find the match for the entry point. When the identification section matches the entry point, the pattern discovers a CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

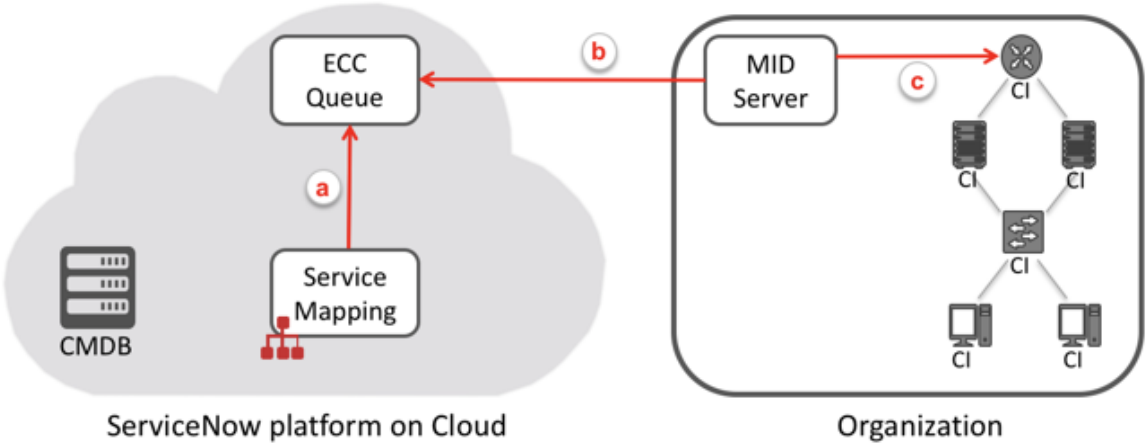

d. The MID Server starts running connectivity sections of the pattern to find outgoing connections of the newly discovered CI.

e. The MID Server passes information on the discovered CI, its attributes, and connections to the ECC queue.

f. Service Mapping checks the ECC queue and receives information on the newly discovered CI.

Whenever Service Mapping checks the ECC queue and receives information on a discovered CI, it checks these tables for any data on outbound connections related to the CI: the cmdb_tcp and sa_flow_connection tables. If these two tables contain unique data that patterns did not discover, Service Mapping enriches the information about the CI connections and adds them to the map.

g. Service Mapping writes the information into the CMDB and adds this CI to the service instance map.

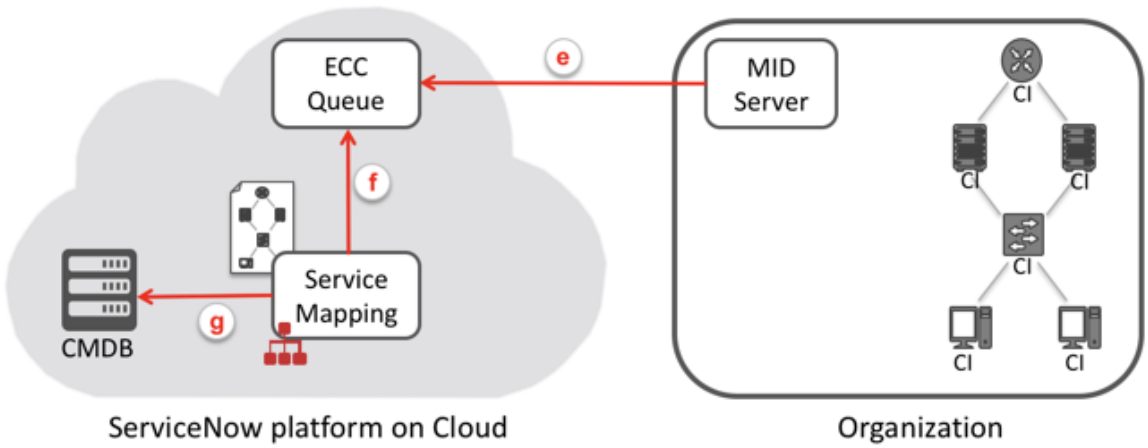

h. Service Mapping creates the discovery requests for all applications to which the newly discovered CI connects. Mapping is complete after Service Mapping maps a CI that either does not have any outbound connections or is marked as a boundary. A boundary makes Service Mapping stop discovery from this point and not follow outgoing connections.

4. The system regularly runs the service recomputation to query the CMDB for the latest CI changes, add data for network and storage paths, and apply CI impact rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- The ECC queue for Discovery
- Horizontal discovery process flow with probes and sensors
- Traffic-based discovery in Service Mapping

Tag-based discovery in Service Mapping

If your organization uses tags, or establishes them in a local ServiceNow instance, you can use these tags to map application services.

A tag is a label that consists of a key-value pair. Your organization might use tags to categorize its assets and to enhance query and reporting capabilities. Discovery and Cloud Provisioning and Governance can discover tags used by all major cloud providers and container ecosystems. Once the tags are discovered, Service Mapping can create application services based on these tags. Typically organizations use tagging in virtualized, hyper-converged, or multi-cloud infrastructures.

Example of using tagging in organization infrastructure

Segment	Technology providers	What is tagged
Cloud	Amazon Web Services (AWS), Azure, Google Cloud Platform (GCP), IBM Cloud Platform	IaaS, PaaS, FaaS, CaaS resources
Hyper-converged	Nutanix	Virtual machines
Containers	Kubernetes, OpenShift Container Platform, AWS ECS	Containers
Virtualization	VMware vCenter	Virtual machines

Unlike other mapping methods, tag-based mapping does not require configuring credentials or providing users with elevated rights. You can effectively use tags to map multiple service instances.

Tag-based discovery stages

The tag-based discovery and mapping process consists of the following stages:

1. Discovery and Cloud Provisioning and Governance discover tags for cloud and resource configuration items (CIs) and then populate data for discovered tags into the Key Value [cmdb_key_value] table.
2. The administrator creates CI tag categories and defines tag keys that the categories contain. Tag categories contain tags with similar use. For example, if your organization has production and staging environments, you could create "production" and "staging" tag values.
3. The administrator creates a tag-based service family and selects tag categories to use for its mapping. CIs that have more than one tags assigned to them, can be part of multiple services.
    - Only CIs with discovered tag keys for the selected tag categories become part of application services.
    - The administrator can narrow the criteria by defining specific tag values in addition to tag categories.
    - Service Mapping uses only CIs that have the matching values to create application services.
    - CIs that have more than one tag assigned to them can be a part of multiple services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: When creating a service family for mapping services based on tags in domain-separated environment, you must pick the relevant leaf domain.

4. Service Mapping queries the CMDB for CIs with tag values that match the tag definitions for this tag-based service family.
5. Service Mapping creates service candidates based on the defined tags.

   Service Mapping populates tag-based services based on this service family only with CIs belonging to the leaf domain to which the service family belongs.
6. The administrator selects the service candidates to map and starts the mapping process.
7. Service Mapping adds new application services to the Tag-Based Application Service [cmdb_ci_service_by_tags] table.

   Important: Service instances mapped using tags might not include relevant CIs if those CIs do not have the correct tags.
8. Service Mapping maps application services by creating connections between tagged CIs based on CI relationships. The Traversal Rules for Application Services [svc_traversal_rules] table contains information used for creating tag-based application services.

   Note: Service Mapping includes CIs that are part of these relationships even if these CIs do not have tags assigned to them.

Creating tag-based services

You can create single tag-based application services using the Common Service Data Model (CSDM) flow as described in populate application services using tags.

For information about how to create multiple tag-based services, see Map application services using tags with classic Service Mapping.

Related topics

Tag-based discovery configuration

Tag Governance

Tag-based discovery for the Service Mapping Workspace

If your organization uses tags for asset management, you can use the Service Mapping workspace to manage these tags and create tag-based application services.

A tag is a label that consists of a key-value pair, used to categorize and organize assets within your organization. Tags enhance query and reporting capabilities, making it easier to manage and track resources. Unlike other mapping methods, tag-based mapping does not require configuring credentials or providing users with elevated rights. You can effectively use tags to map multiple service instances. Discovery and Cloud Provisioning and Governance can automatically detect tags from all major cloud providers and container ecosystems. Once these tags are discovered, Service Mapping can leverage them to create and manage application services. This approach is useful for resource management in virtualized, hyper-converged, and multi-cloud environments.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information about the traditional tag-based discovery process, see Tag-based discovery in Service Mapping. Tag-based discovery for the Service Mapping workspace operates differently and consists of the following stages:

Note:

Access to tag-based service mapping in the Service Mapping workspace requires the installation of Service Mapping Plus version 1.16.3. For more information, see Install Service Mapping Plus.

1. Discovery and Cloud Provisioning and Governance discover tags for cloud and resource configuration items (CIs) and then populate data for discovered tags into the Key Value [cmdb_key_value] table.
2. The administrator creates CI tag categories and defines tag keys that the categories contain. Tag categories contain tags with similar use. For example, if your organization has production and staging environments, you could create "production" and "staging" tag values.

   Note: Starting with version 1.7.0, Tag Governance offers five predefined categories that you can use without modification. Administrators can also modify these categories or add new ones.

3. The process begins with the selection of a tag category from the available categories or a category that the administrator creates.
   - Only CIs with discovered tag keys for the selected tag categories become part of application services.
   - The administrator can narrow the criteria by defining specific tag values in addition to tag categories.
   - Service Mapping uses only CIs that have the matching values to create application services with the service family being created automatically.
4. Service Mapping queries the CMDB for CIs with tag values that match the tag definitions for the selected categories.
5. Service Mapping creates service candidates based on the defined tags.
6. The administrator selects the service candidates to map, starts the mapping process, and names the tag-based service family. This record is saved so the administrator can return at any time to monitor or continue the process later.
7. Service Mapping adds new application services to the Tag-Based Application Service [cmdb_ci_service_by_tags] table.

   Important: Service instances mapped using tags might not include relevant CIs if those CIs do not have the correct tags.

8. The tag-based service content is calculated on demand either when the map is first viewed or initiated manually.

   Service Mapping maps application services by creating connections between tagged CIs based on CI relationships. The Traversal Rules for Application Services [svc_traversal_rules] table contains information used for creating tag-based application services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Manual CI Exclusions/Inclusions [svc_manual_ci_exclusions_inclusions] table can filter out CI classes, helping to prevent them from being added by tags or traversal rules.

Note: Service Mapping includes CIs that are part of these relationships even if these CIs do not have tags assigned to them.

Related topics

Tag-based mapping in the Service Mapping Workspace

Map your application services using tags in the Service Mapping Workspace

Traffic-based discovery in Service Mapping

Service Mapping can discover and map configuration items (CIs) following their traffic-based connections. This method is referred to as traffic-based mapping and complements pattern-based mapping.

What is traffic-based mapping

Traffic-based mapping can be seen as an enhancement to pattern-based mapping. While pattern-based mapping relies on predefined rules and patterns to identify connections, traffic-based mapping analyzes network traffic to automatically discover connections between CIs in your environment. This approach allows Service Mapping to find those CIs not discovered using patterns alone.

Depending on your configuration, the behavior of traffic-based discovery can vary. If Predictive Intelligence is enabled, Service Mapping automatically adds connections to service instances based on connection rules. These suggestions are generated using traffic-related data from the Configuration Management Database (CMDB) and the analysis of application fingerprints, CIs, and processes by Predictive Intelligence.

If discovery based on Predictive Intelligence isn't enabled, Service Mapping automatically adds traffic-based connections using data from the CMDB to the service instances. However, to keep your service instances well-organized, you might need to manually remove connections to irrelevant CIs.

Typically, if the connection suggestions feature is turned off, you use traffic-based discovery at the initial stages of discovering application services. To maintain a clean, organized service map, you can opt to turn off traffic-based discovery after you have completed the discovery and fine-tuning of your application services. For more information, see Fine-tune application services to implement owner requests.

Connections created only through traffic-based discovery aren’t removed. However, if pattern-based discovery runs after traffic-based discovery has run, it might create duplicate connections between CIs. To avoid duplicates, the system removes any connections that were created by traffic-based discovery if they’re also created by pattern-based discovery.

Traffic-related data from the CMDB

The system uses commands and network flow logs to collect traffic-related data and saves it in the CMDB tables. Service Mapping retrieves this data from the tables to detect CI inbound and outbound connections.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables containing data collected using traffic-based methods

Table	Source	Used by Service Mapping to
Flow Connector [sa_flow_connection]	Netflow and VPC logs	Discover dependencies, add connections during top-down discovery.
Flow Services IP/Port and Statistics [sa_flow_service]	Netflow and VPC logs	Discover all services listening on ports. In a base system, Service Mapping does not use data from this table.
Flow Server Communication [sa_flow_server_comm]	Netflow and VPC logs	Discover services communicating to other services. In a base system, Service Mapping does not use data from this table.
TCP Connection [cmdb_tcp]	netstat and lsof commands	Discover connections during top-down discovery.

In base systems, which are the default or standard configurations, traffic-based discovery relies solely on TCP-related data collected using the *netstat*, *ss*, and *lsof* commands. Discovery based on Netflow and VPC logs requires additional configuration. To enhance traffic-based discovery, you can configure Service Mapping to collect data from Netflow and VPC logs. Service Mapping also accesses the TCP connection data collected by improved Application Dependency Mapping (ADM), which Discovery performs as part of horizontal discovery. See Application Dependency Mapping (ADM) for Discovery to learn more about this process.

Enabling traffic-based discovery in the system

Traffic-based discovery using commands is available in Service Mapping for use at any level. You can enable it at different levels ranging from the most global to the most specific. For more information, see Enable traffic-based discovery for an application service.

Product level

By default, traffic-based discovery in Service Mapping is turned off. The **Traffic based discovery** (*sa.traffic_based_discovery.active*) property controls traffic-based discovery at this level. Enabling traffic-based discovery at the product level facilitates the discovery of all CIs in your environment. However, more specific rules at the service instance, CI type, or specific CI levels can override these settings.

> **Important:** You can't enable traffic-based discovery at other levels unless it's enabled at the product level.

The connection suggestions feature works at the product level. The *sa_ml.connection_suggestions.active* property controls this feature. If you enabled traffic-based discovery before the Quebec release and used it to discover at least one application service, the connection suggestions feature is turned off by default.

Service instance level

You can enable traffic-based discovery for an individual service instance. In this case, Service Mapping uses traffic-based discovery for all CIs that make up this service instance. However, if traffic-based discovery is turned off for some CI types

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

or specific CIs, those CIs aren't discovered using this method, even if the service instance is set to use traffic-based discovery.

For example, you enable traffic-based discovery for a specific service instance that includes a Tomcat server, a MySQL database, and a web application. Service Mapping uses traffic-based discovery to discover the Tomcat server, the MySQL database, and the web application.

CI type level

You can create a discovery rule to include or exclude a CI type from traffic-based discovery. This rule takes precedence over the setting that you choose for an individual service instance.

For example, you enable traffic-based discovery for a specific service instance that includes a Tomcat server, a MySQL database, and a web application. You create a CI type rule that excludes all Tomcat servers from traffic-based discovery. Service Mapping uses traffic-based discovery to discover the MySQL database and the web application since there are no rules to exclude them.

Specific CI level

You can create a discovery rule to include or exclude a specific CI type from traffic-based discovery. This rule takes precedence over the setting that you choose for an individual service instance or any rules for CI types.

For example, you enable traffic-based discovery for a specific service instance that includes a Tomcat server, a MySQL database, and a web application. You create a CI type rule that excludes all Tomcat servers from traffic-based discovery. However, you also create a discovery rule enabling traffic-based discovery for a specific Tomcat server. Service Mapping uses traffic-based discovery to discover the MySQL database, the web application, and that specific Tomcat despite the CI type rule excluding all Tomcat servers.

See Enable traffic-based discovery for CI types or specific CIs to learn more about creating rules for CI types or specific CIs.

Related topics

- Pattern-based discovery in Service Mapping
- Manage CI connections in bulk using suggestions
- Manage CI connections for a single service using suggestions
- Remove CIs not belonging to application services
- Data collection and discovery using Netflow
- Data collection and discovery using VPC Flow Logs

Service Mapping for containerized environments using KVA

Kubernetes Visibility Agent (KVA) and Service Mapping discover and visualize application dependencies across Kubernetes clusters and related resources, providing complete visibility into containerized environments.

Service Mapping for containerized environments provides automated discovery and visualization of application dependencies across your entire technology stack. Service Mapping creates comprehensive topology maps that show how services connect and depend on each other. This applies whether your applications run entirely within Kubernetes clusters or span multiple environments including cloud databases, API gateways, and traditional servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping and KVA combined capabilities transform complex architectures into clear, actionable service maps. Teams can troubleshoot faster, assess change impact more accurately, and optimize application performance across environments.

Key capabilities

Seamless entry point discovery and mapping

Start with any application endpoint such as a URL or hostname and port, and Service Mapping automatically identifies the corresponding Kubernetes resource serving that endpoint. KVA recognizes multiple Kubernetes exposure methods including load balancers, ingress controllers, NodePort services, Gateway API resources, and OpenShift routes.

Complete Kubernetes topology-mapping

Discover and visualize traffic flows between pods within your Kubernetes clusters. Service maps display workloads including deployments, stateful sets, daemon sets, jobs, and cron jobs. You can choose to view either workloads or Kubernetes services on your maps, providing flexibility in how you analyze your container infrastructure.

External dependency tracking

Automatically detect connections from Kubernetes pods to external resources including cloud databases, third-party APIs, and traditional infrastructure. KVA resolves DNS connections and Kubernetes ExternalName services to be able to map dependencies beyond cluster boundaries.

Bidirectional discovery

Service maps work regardless of whether your entry point leads directly to a Kubernetes resource or to an external component with downstream connections into clusters. This flexibility promotes comprehensive visibility across any architecture.

Use Cases

Accelerated incident resolution

When incidents occur, Service Mapping provides immediate visibility into all dependencies, helping teams quickly identify root causes and understand the impact across Kubernetes and traditional infrastructure.

Change impact analysis

Before changing Kubernetes deployments, APIs, or infrastructure components, visualize all upstream and downstream dependencies to understand potential impacts and plan maintenance accordingly.

Application rationalization

Gain complete visibility into which services support critical business applications, enabling data-driven decisions, resource allocation, and technical debt reduction.

Benefits

- Reduced mean time to resolution through comprehensive visibility into service architectures
- Minimized risk from changes by reviewing dependencies before modifying Kubernetes or infrastructure components
- Improved operational efficiency by automating discovery instead of manually documenting complex service relationships

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Better decision-making through accurate, up-to-date topology information that reflects actual application behavior
- Simplified compliance and audit processes with automated documentation of service dependencies and data flows

Getting Started

To begin discovering and mapping Kubernetes services, install Kubernetes Visibility Agent (KVA) latest version and make sure Service Mapping requirements are met. For more information, see:

- Configuring Kubernetes Visibility Agent
- Configuring Service Mapping

Create hybrid service maps

Enter an application entry point into Service Mapping. KVA automatically identifies whether this entry point connects to Kubernetes resources and initiates comprehensive topology discovery.

Service Mapping Lightweight Service Model

The Lightweight Service Model is an optimized architecture for storing and determining service topology. It maintains current service data without storing historical snapshots, reducing storage overhead and improving performance for large configuration item inventories.

What is the Lightweight Service Model

The Lightweight Service Model is an alternative architecture for Dynamic service instances and Tag-Based application services in Service Mapping. Unlike the traditional model, which maintains a complete history of service topology snapshots, the Lightweight Service Model stores only the current state of service topology.

Comparison of traditional and Lightweight Service Model

Traditional model	Lightweight Service Model
Stores complete history of service topology snapshots	Stores only current service topology data
Historical records are available for all past updates	No historical snapshots retained; timeline shows change points but not full history
Slower update cycles due to historical data management	Faster update cycles
Can query past service topology states and change history	Can query current state only, with direct access to service topology data

Benefits of the Lightweight Service Model

Service Mapping AI capabilities enable creating services at scale. While scaling up the mapping, the Lightweight Service Model becomes crucial for scaling up the performance.

When a service operates under the Lightweight Service Model, service topology updates are significantly faster because the system doesn't maintain historical snapshot data. In addition, storage requirements are reduced because only current service topology data is stored.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Converting a service to Lightweight Service Model

Conversion to the Lightweight Service Model is permanent and one-way. Once you convert a service to Lightweight Service Model, this change can't be reversed. Historical snapshots that were stored under the traditional model will be deleted and can't be restored.

When you convert a service to Lightweight, the system optimizes how it stores and manages data across the service hierarchy tables. The service types that can be converted to Lightweight are Tag-Based and Dynamic.

Lightweight Service Model related link

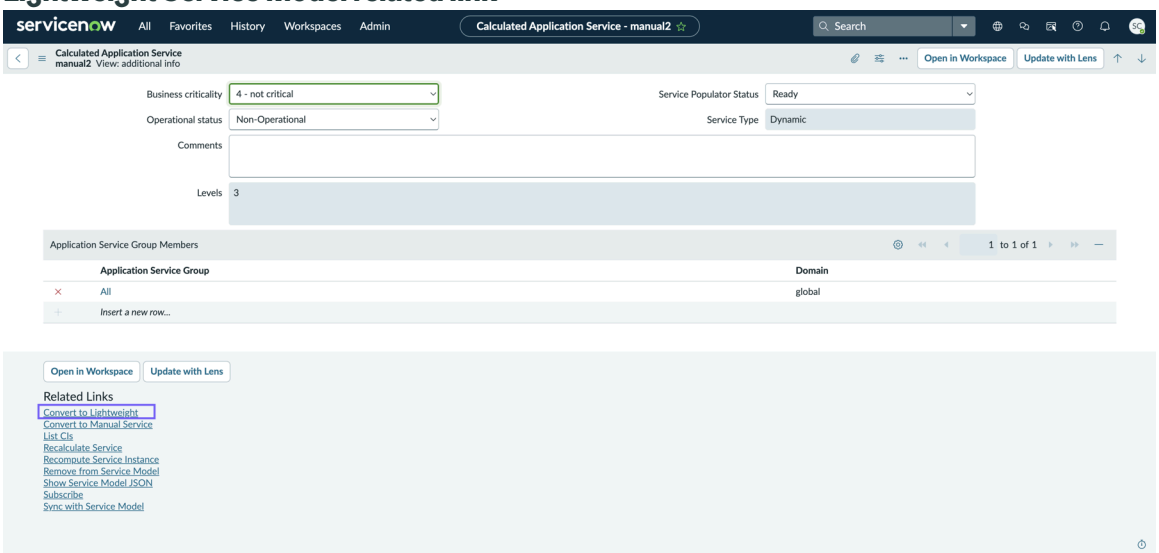

Use the Lightweight Service Model related link to convert a service. For the full procedure, see Convert a service instance to Lightweight Service Model.

Related topics

Convert a service instance to Lightweight Service Model

Discovery of application services on cloud using Service Mapping

Service Mapping in cloud environments provides critical visibility into application dependencies and connections. By identifying how different application components interact within IaaS and PaaS environments, your organization can gain better insight into its application services and improve overall service management.

IaaS and PaaS environments

To navigate the discovery process effectively, you must first understand the fundamental differences between Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) environments.

- **IaaS** - IaaS gives organizations virtual computing resources over the internet, enabling them to manage and deploy their own apps and operating systems. In IaaS, Service Mapping uses a similar discovery process as in private networks, identifying hosts and applications through patterns.
- **PaaS** - PaaS offers a complete platform for building, launching, and running applications, removing the need for organizations to manage hardware or servers directly. In PaaS, Service Mapping must use cloud platform tools and service accounts to identify devices and applications.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Performing Discovery in IaaS and PaaS Environments

The discovery process in IaaS environments mirrors that used for private networks.

1. Locating hosts - Service Mapping starts by using Discovery to locate hosts.
2. Identifying applications and connections - Patterns are used to identify applications and the connections between those applications and the discovered hosts.
3. Top-down discovery - Service Mapping performs top-down discovery to map the relationships and connections between devices and applications. It then creates a comprehensive map of the application service.

In contrast to IaaS environments, PaaS environments require a distinct discovery process.

1. Using cloud platform tools and service accounts - The process begins with a dedicated service account configured for the cloud environment hosting the application service. Service Mapping and Discovery use this service account to access devices and applications. To learn more about setting up an account, see Set up a cloud service account.
2. Horizontal discovery - Discovery performs horizontal discovery in two steps:
   a. Probes are used to help identify infrastructure devices. See Horizontal discovery process flow with probes and sensors for detailed information about the four phases of discovery using probes.
   b. Patterns help find the applications and services running on those devices. See Pattern-based discovery in Service Mapping and Horizontal discovery process flow with patterns to learn more about discovery using patterns.
3. Top-down discovery - Service Mapping then performs top-down discovery, creating a comprehensive map of the application service.

Discovery methods for cloud environments

Service Mapping/Discovery Approach	IaaS (Virtual Machines, Servers)	PaaS (Functions, Databases, Queues)	Key requirement difference
Top-down service mapping	Supported (Requires OS credentials or deep agent access)	Not supported (or rarely supported) by cloud metadata discovery alone	Access to the OS layer to trace live connections
Tag-based mapping	Supported (Using collected tags from the cloud resources and containers)	Supported and highly applicable (Uses cloud-native tagging/metadata)	Availability of organized metadata (tags)
Machine Learning (ML)-based mapping	Supported (Requires capturing running processes and network traffic, thus requires deep discovery)	Supported, but requires deep discovery (Only available when using methods that access the underlying OS/container runtime)	Access to application process data

Consider an e-commerce application hosted on a cloud provider. In an IaaS environment, Service Mapping identifies virtual machines, databases, and their connections, providing a comprehensive view of the application's infrastructure. In a PaaS environment, it uses cloud platform tools to map application components like managed databases, revealing critical

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

dependencies that impact performance and reliability. This visibility enables better resource management, remediation, and optimization of the application service.

Related topics

Discovery for cloud environment

Service Mapping flow

Learn about high-level tasks users having different roles perform in Service Mapping.

In Service Mapping, users collaborate to map, review, and monitor an organization's application services.

Note: For a complete list of Service Mapping roles, see Components installed with Service Mapping.

A typical Service Mapping workflow has the following stages:

1. The administrator performs basic obligatory configurations to set up Service Mapping. See Configuring Service Mapping for more details.
2. The administrator maps organization application services in bulk. In addition, the administrator may map some application services individually. See Map multiple application services suggested by classic Service Mapping for more information about mapping services in bulk.
3. The administrator fixes application service errors in bulk, as described in Fix application service errors in bulk.
4. The administrator reviews the results of the initial mapping, resolving errors in individual application services. See Fix errors in individual application service maps for more details.
5. The administrator sends fixed application services to the application service owner for review. See Send application service maps for review for more information.
6. The service instance owner checks that the service instance maps are complete and all major components comprising it are correctly represented. If necessary, the owner leaves comments, referred to as reject messages, on service instance maps for the Service Mapping administrator to implement. See Review application service maps.
7. The administrator uses the feedback from the application service owner to fine-tune the application service maps, and then resends them to the owner for review. For more information about fine-tuning an application service map, see Fine-tune application services to implement owner requests.
8. If the revised service instance maps are satisfactory, the service instance owner approves them. If not, the owner requests further fixes, which the administrator must address. See Review application service maps.
9. The administrator finalizes the application services by configuring access and setting advanced attributes, such as criticality. To learn more, see Application service completion.
10. After the service instance definition is complete, the Service Mapping user can view application service maps. To learn more about maintaining the application services using Service Mapping, see Application service analysis and maintenance using classic Service Mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping flow

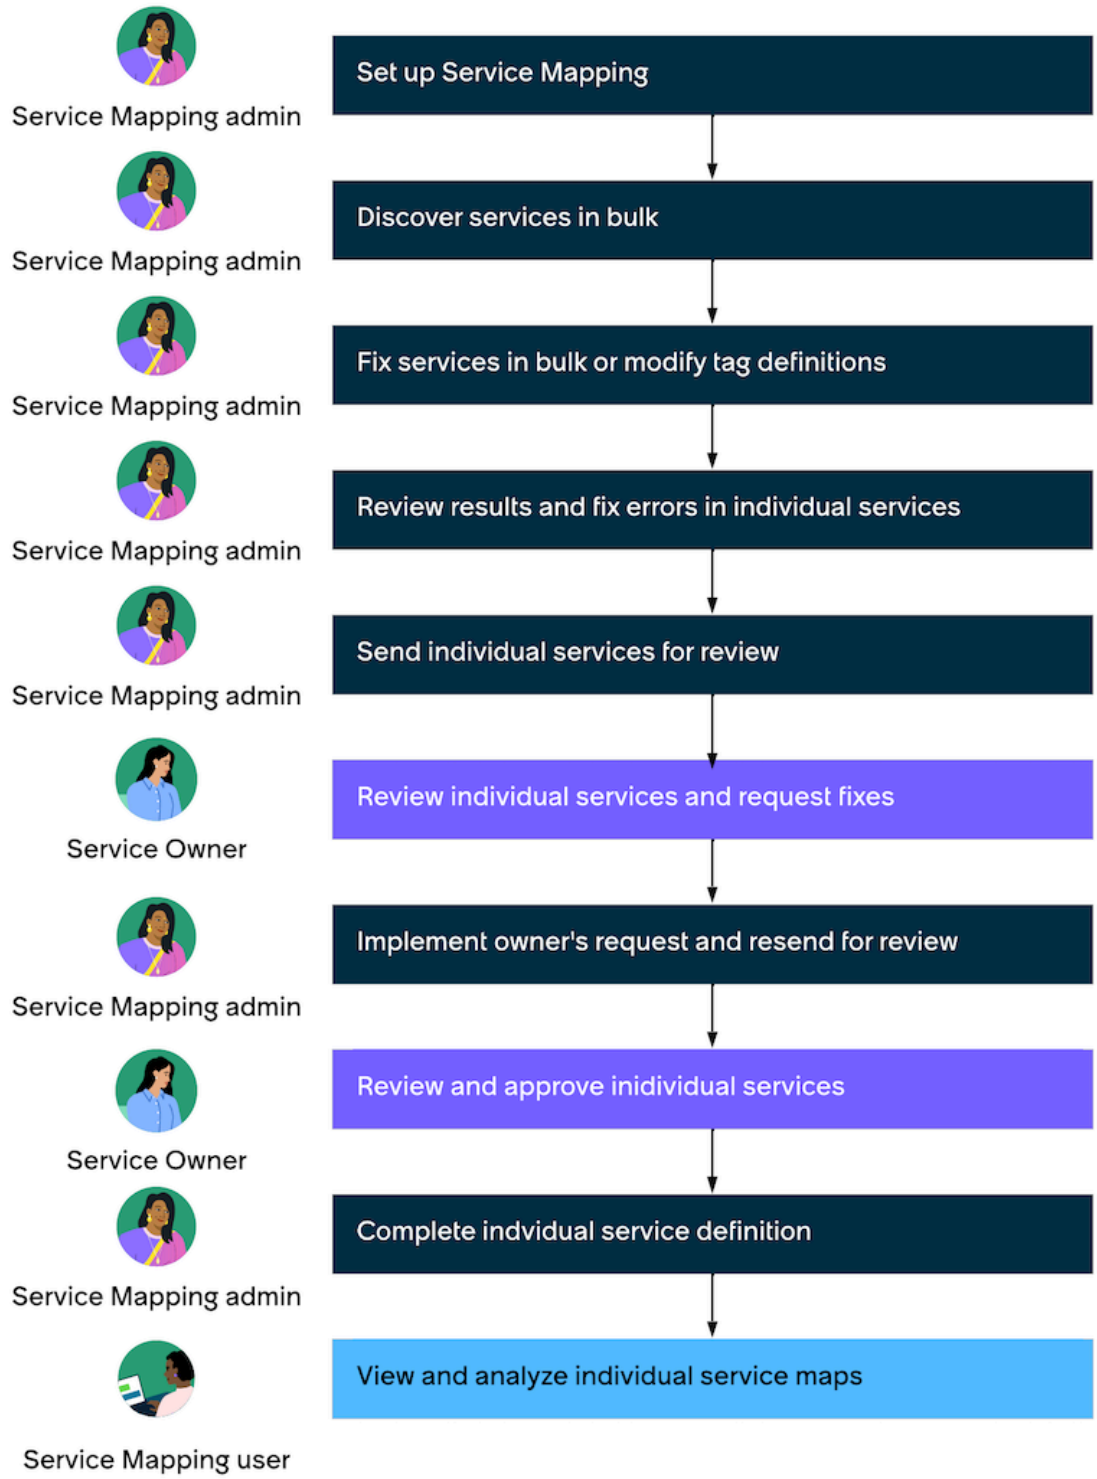

In addition to these tasks, there may be a need to customize patterns that Service Mapping uses to discover devices and applications forming application services. Users must have the pd_admin role to customize patterns. In the base system, the service_mapping_admin role contains the pd_role. Customizing patterns requires basic knowledge of programming.

Related topics

Control user access to application services

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag-based mapping in the Service Mapping Workspace

The Tag-based service mapping dashboard offers an intuitive way to manage tag-based services. Use the widgets, data visualizations, and list navigation to review and update your services, or create new tag-based services in a few steps. The dashboard enhances the Service Mapping Workspace, offering an efficient and user-friendly interface for organizing and optimizing your service mappings.

Prerequisites

Access to tag-based service mapping in the Service Mapping workspace requires the installation of Service Mapping Plus version 1.16.3. For more information, see Install Service Mapping Plus.

Required role

service_mapping_admin

Access

To open the dashboard, navigate to **Workspaces** > **Service Mapping**, then select the Tag-based service mapping ( ) icon from the navigation pane.

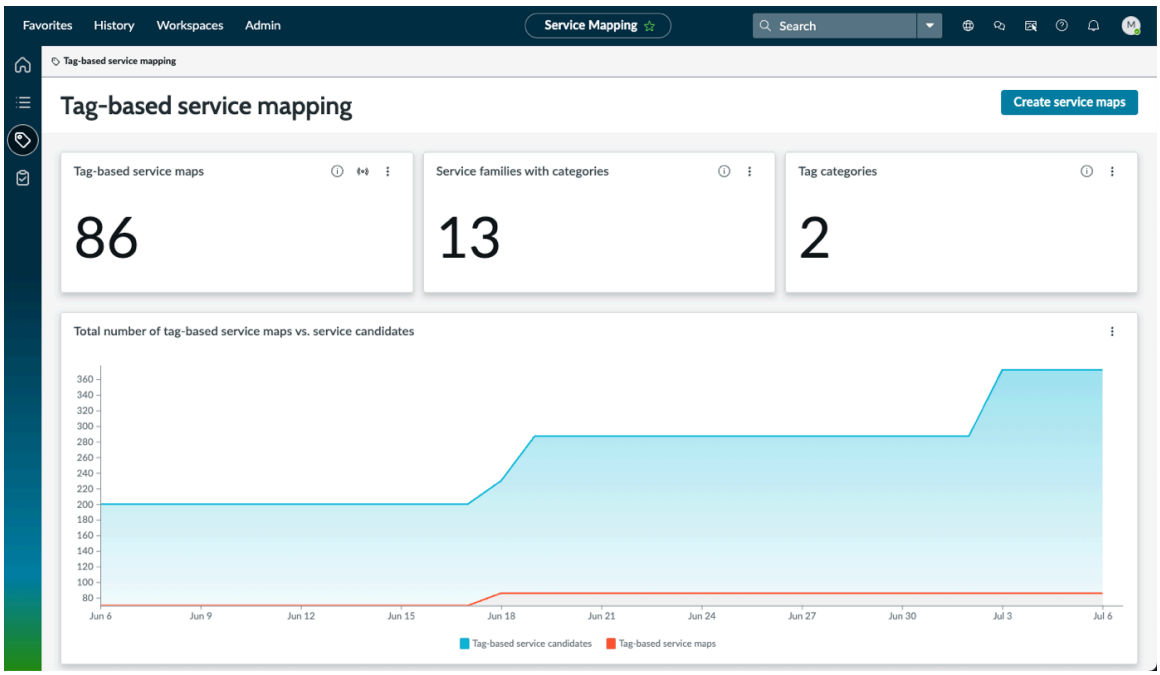

Key metrics

Selecting a widget in this section redirects you to the list for more detailed information.

Tag-based service maps

The number of service maps created using tag-based service mapping.

Service families with categories

The number of service families created to store service candidates and service maps.

Tag categories

The number of tag categories, which group tags that have a related use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Report

Title	Type	Description
Total number of tag-based service maps vs. service candidates	Two-variable line graph	A chart that shows the tag-based service candidates and the tag-based service maps created on a given day over a 30-day period. Use this information to evaluate the actual number of services mapped as compared to the number of service candidates created.

List navigation

The List navigation pane displays tag-based service mapping information. 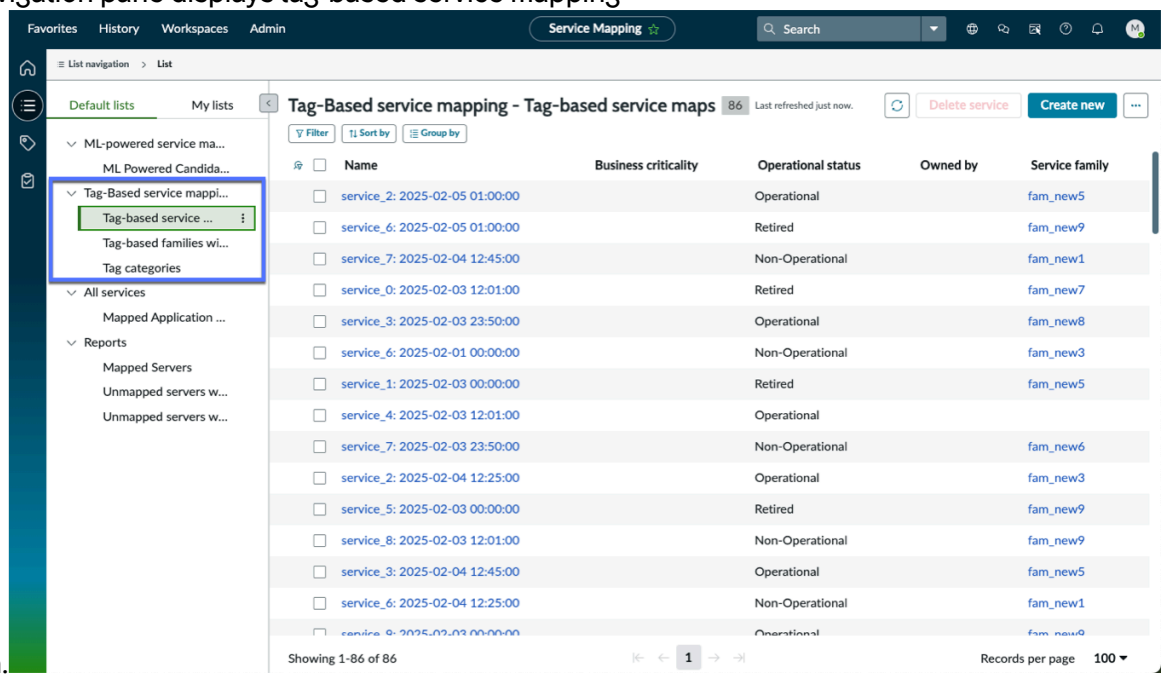

Display tag-based information and access tag-based actions through the following pages:

Tag-based service maps

- View or delete existing service maps.
- Initiate the flow to create a tag-based service.

Tag-based families with categories

- View service families that have at least one selected category.
- Select a service family to view detailed information including mapped services and service candidates.
- Initiate the flow to create a tag-based service.

Tag categories

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- View or delete existing tag categories.
- Select a tag category to view detailed information, including tag keys.
- Initiate the flow to create a tag category.

Domain separation

When creating a service family for mapping services based on tags in domain-separated environment, you must pick the relevant leaf domain. Service Mapping populates tag-based services based on this service family only with CIs belonging to the leaf domain to which the service family belongs. If you have domain separation without Performance Analytics domain support, you can't see data visualization for the total number of tag-based service maps versus service candidates. You're then redirected to install the Performance Analytics Domain Support plugin.

For more information about domain separation and Service Mapping, see Domain separation and Service Mapping.

About the tag-based application service form

The tag-based application service form provides a centralized location for viewing detailed information about a selected tag-based service, including its configuration and its associated configuration items.

The tag-based application service form offers a comprehensive view of a tag-based application service. With actionable buttons and multiple tabs, you can manage and troubleshoot your tag-based application services and their associated configuration items (CIs).

Access

To access the tag-based application service form:

1. Navigate to **Workspaces** > *Service Mapping*>.
2. Select the **Tag-based** *Service Mapping* icon from the navigation pane.
3. Select the **Tag-based service maps** tile.
4. Select the desired tag-based service to view the tag-based application service form.

Details tab

The **Details** tab displays basic information about the selected tag-based application service, including the service family and tag categories associated with the service. Use this tab to update the basic information about an application service, such as the description, the criticality, or the operational status.

Service CIs tab

The **Service CIs** tab lists the configuration items (CIs) that are part of the selected tag-based application service. The *svc_by_tags.max_ci_limit* property controls number of CIs listed in this tab. Currently, the list is limited to the first 1,000 CIs that match the criteria for the tag-based service.

Note: This list filters out manual endpoints, ensuring that only relevant CIs are displayed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tagged CIs tab

The **Tagged CIs** tab lists all CIs that meet the criteria for a tag-based application service, up to 4,000. The *sa.it_service.list_cis_max_count* property controls this limit. This list includes CIs belonging to the application service (also listed on the **Service CIs** tab) and up to 3,000 additional CIs. This tab is useful for building new tag-based application services or troubleshooting existing services with a preview of relevant CIs. However, no CIs beyond the 4000 limit are displayed.

Traversal rules in Service Mapping

Traversal rules identify and map relationships between configuration items to create comprehensive application service maps. They connect relevant configuration items based on predefined relationships, promoting accurate service mapping and visualization.

Traversal rules facilitate comprehensive service mapping by identifying related configuration items (CIs) through specific relationships that exist in the Configuration Management Database (CMDB). These rules support the mapping of related CIs even when they aren't directly tagged, providing a more accurate service map.

Components of traversal rules

When configuring traversal rules, it's essential to add them in the correct order, as tag-based discovery uses these rules to identify CIs that can be mapped.

The Service Mapping administrator (service_mapping_admin) manages traversal rules in the Traversal Rules for Application Services [svc_traversal_rules] table. When configuring these rules, the following key components are considered:

- **Parent Type:** The CI class that serves as the starting point for mapping related configuration items that belong to the Child class as specified by the traversal rule.
- **Relationship:** The specific association or link between the Parent CI and Child CI that determines how they're related in the service mapping context. The direction of the traversal depends on the **Is Reverse** setting.
- **Child Type:** The CI class that contains the related configuration items being mapped from the Parent CI class as specified by the traversal rule.
- **Priority Level:** The level that specifies the order in which traversal rules are applied, with lower numbers taking precedence over higher numbers.
- **Is Reverse:** A flag that determines the direction of traversal for the relationship between Parent CI and Child CI. When selected, the mapping process starts from the Child CI instead of the Parent CI.

> **Note:** To promote accurate mapping, configure traversal rules with the correct priority levels, as tag-based discovery relies on these rules to identify mappable CIs.

Configuring traversal rules

Service Mapping includes predefined traversal rules, simplifying the initial configuration. These rules provide a starting point for your tag-based service mapping setup. For details on the specific predefined rules, see Predefined traversal rules for application services.

In addition to these predefined traversal rules, you can configure custom traversal rules to meet your organization's specific needs. For detailed instructions on how to configure traversal rules, including information on the required fields and settings, see Traversal Rules for Application Services form completion.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, you can only have five traversal rules participate in tag-based discovery. However, you can adjust this limit by configuring the *sys_property svc_traversal.rules.active* system property to enable more active traversal rules.

AI in Service Mapping

AI-powered features in Service Mapping help administrators automate service map creation, connect business applications to discovered services, and query live service topology using natural language.

Service Mapping includes several AI-powered features. All features require ServiceNow Otto for IT Operations Management (ITOM) to be installed.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

MCP servers

The Service Mapping MCP server exposes instance data to external AI clients through the Model Context Protocol, giving those clients structured, ACL-enforced access to live application service data.

MCP servers in Service Mapping

AI feature	Description	Use case	Resources
Service Mapping MCP server	Exposes five read-only tools through the CMDB MCP Server. The tools give AI clients such as Claude structured access to live application service data, including full service topology, server impact graphs, and unmapped CIs. Authentication uses OAuth 2.0 with JWT tokens and enforces the same ACLs as standard instance API calls.	A Service Mapping administrator or operator wants to query service topology, identify mapping gaps, or assess server impact. They do so by using natural language in an AI client, without navigating the instance UI.	• Service Mapping MCP tools • Activate the CMDB MCP Server for Service Mapping tools • Connect Claude Desktop to the Service Mapping MCP Server • Service Mapping MCP tools reference

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI agents and agentic workflows

Service Mapping AI agents and agentic workflows automate service map creation, business application mapping, and change impact analysis.

AI agents and agentic workflows in Service Mapping

AI feature	Description	Use case	Resources
Service Mapping AI Agent	Autonomously creates service maps from ML-powered Application Service Candidates (ASCs). The agent evaluates candidates, filters noise such as monitoring clients and operating system processes, and creates the service topology in the CMDB.	A Service Mapping administrator wants to process a large volume of ML-powered candidates and create service maps automatically, without reviewing each candidate individually.	• AI Agents for Service Mapping • Activate AI Agents for Service Mapping
Business App Mapping AI Agent	Automatically creates CSDM Consumes::Consumed by relationships between Business Applications and discovered Application Services by invoking the Business App Identification skill. **Note:** The Business App Mapping AI Agent is available through the Feature Preview Program. For more information, see Feature Preview Program.	A Service Mapping administrator wants to connect discovered application services to their business application context in the CSDM without manually mapping each relationship.	• AI Agents for Service Mapping • Activate AI Agents for Service Mapping
Analyze potential impact agentic workflow	Analyzes how a change request might affect relevant servers and services. The agent verifies prerequisites, retrieves the change request, selects up to 10 affected servers, identifies matches with suggested services, and displays an impact analysis in the ServiceNow Otto panel. The analysis is also saved to the change request work notes. Uses the Service Mapping Candidate and Service Mapping Candidates Impact skills.	An operator or change manager wants to assess the risk of a change request before approving it, without manually reviewing all affected CIs.	• Analyze potential impact agentic workflow • Assess a change request with the Analyze potential impact workflow • Activate the Service Mapping Candidate skill • Activate the Service

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI agents and agentic workflows in Service Mapping (continued)

AI feature	Description	Use case	Resources
			Mapping Candidates Impact skill

AI Agents for Service Mapping

The Service Mapping AI agents automate the creation and maintenance of service maps in the Configuration Management Database (CMDB), reducing manual effort for Service Mapping administrators.

AI Agents for Service Mapping automate key parts of the Service Mapping workflow. After activation, all agents run autonomously in the background on a recurring schedule, processing all relevant records without requiring administrator input for each one.

Available AI Agents

The ITOM AI Agents for Service Mapping application [sn_sm_gen_ai] provides the following AI Agents:

Service Mapping AI Agent

Automatically creates ML-based service maps from Application Service Candidates (ASCs) and persists the full service topology in the CMDB. The Agent roles:

- Evaluates ML-powered candidates, focuses on those with a high-confidence name suggestion source such as Service Fingerprints (SFPs).
- Filters out noise (such as monitoring clients, security clients, and operating system processes),
- Creates the service topology.

Created service maps are set to non-operational by default so that the Service Mapping administrator can review them before making them operational.

Business App Mapping AI Agent

Tag-Based Service Map Creation AI Agent

Automatically creates service maps from tag-based candidates generated by cloud infrastructure metadata. This Agent analyzes cloud tags (such as Application and Environment) applied to virtual machines and determines whether they represent real, coherent business services. The Agent performs two critical evaluations:

- It signals whether the tag represents a meaningful business service using both the tag name and running process evidence.
- It verifies that all virtual machines (VMs) grouped under the tag actually belong to the same application service rather than multiple unrelated systems that share a tag. Created service maps are set to non-operational by default.

Why these AI Agents are needed

Without these AI Agents, Service Mapping administrators have to manually review each ML-powered candidate or tag-based candidate and decide whether to create a service map. Similarly, connecting discovered application services to their business application context in the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CSDM required manual one-to-one or one-to-many mapping. Using the AI Agents makes these processes significantly less time-consuming and enables processing at a much larger scale.

The following examples illustrate the challenges these AI Agents solve.

Consider a bank that has a payment application defined as a business application. Under it, Service Mapping discovers many application services:

- Online payment processing.
- Cash payment handling.
- Net banking.
- Mobile payment systems.

Business stakeholders are only concerned with the payment application at the top level. They need it connected to all the relevant discovered services so they can monitor health and impact. Manually maintaining these connections as new application services are created is impractical. The Business App Mapping AI Agent solves this by iterating over all business applications and automatically finding and connecting the matching discovered services. It works top-down, from the smaller set of business applications to the larger set of discovered application services.

In a typical scenario, an organization runs hundreds of servers and generates a large number of ML-powered candidates daily. A Service Mapping administrator who previously had to manually evaluate each candidate and create a service map can instead activate the Service Mapping AI Agent. The agent processes all high-confidence candidates automatically, creates service maps, and sets them to non-operational for administrator review.

Organizations using cloud infrastructure often apply metadata tags to virtual machines for organization and governance. However, manual service creation from these tags is time-consuming and error-prone, especially at scale. The Tag-Based Service Map Creation AI Agent solves this by:

- Automatically analyzing tagged infrastructure, detecting when multiple unrelated systems are incorrectly grouped under a single tag.
- Creating service maps only when the tag represents a real, coherent business application.

How AI Agents work

All agents use a ReAct (Reasoning and Acting) strategy. The agent reasons about each step before deciding which tool to invoke. After activation, a script runs every 15 minutes. On each run, the agent processes all relevant records in a single batch — for example, all business applications in the system. No per-record input from an administrator is required.

Service Mapping AI Agent process:

1. The agent retrieves full ML candidate data from the Application Service Candidate (ASC) record, including the candidate name, process name, and server count.
2. The agent performs a noise check. If the candidate is flagged as noise, such as monitoring clients, security clients, or operating system processes, the agent skips it.
3. The agent retrieves all running processes associated with the candidate.
4. The agent analyzes the data and reasons about what the actual service is, prioritizing candidates with a high-confidence name suggestion source such as Service Fingerprints (SFPs).
5. The agent creates the service topology in the CMDB and sets the service to non-operational. If a topology for the same candidate already exists, the agent skips creation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag-Based Service Map Creation AI Agent process:

1. The agent retrieves a tag-based candidate record from [svc_by_tags_candidates], including the cloud tags (such as Application and Environment) and the set of member virtual machines.
2. The Agent validates the candidate through prerequisite checks. It confirms the candidate is not hidden, ensures at least 2 configuration items (CI) exist to form a topology, and verifies no existing service already maps it.
3. The agent performs a coherence evaluation by analyzing the running processes and process groups (AFP groups) across all member VMs. It determines whether all VMs running under this tag belong to the same application service or whether multiple unrelated systems are incorrectly grouped.
4. If the candidate is incoherent (containing multiple distinct business systems), the agent records the outcome and skips service creation, flagging the candidate for administrator review.
5. If the candidate is coherent or uncertain, the agent evaluates the signal strength. The agent does so by combining the cloud tag name (for example, "order-management") with running process evidence to determine whether a real business service exists.
6. If the signal is strong enough, the agent derives a service name using this priority order:
   a. Service Fingerprint (SFP) name if the candidate's CIs match an Application Service Fingerprint in the Discovery Content Library.
   b. Strong cloud tag value if no service fingerprint exists.
   c. AFP process-derived name if the tag signal is weak and no fingerprints exist.
   The agent then creates the service topology and sets the service to non-operational. If a topology for the same candidate already exists, the agent skips creation.

Activation and monitoring

All agents are inactive by default. A user who has the Service Mapping admin role activates them from the Service Mapping home page. For steps to activate the AI Agents, see Activate AI Agents for Service Mapping.

The Service Mapping AI Agent and Tag-Based Service Map Creation AI Agent are activated together through a single workspace activation button. The Business App Mapping AI Agent is activated independently.

The Business App Mapping AI Agent is available through the Feature Preview Program. For more information, see Feature Preview Program.

Note: The Service Mapping AI Agent and Tag-Based Service Map Creation AI Agent operate independently despite sharing the same orchestrator and writing the results to a shared run record table [sn_sm_gen_ai_svc_create_run]. They process different input sources and have no data dependencies.

No configuration in AI Agent Studio is required for Service Mapping administrators. AI Agent Studio is used by platform administrators or users who have the AI admin role to manage agent definitions and token usage.

After activation, administrators can monitor activity and review results in the **AI Activity** section of the Service Mapping list navigation. For more information, see Service Mapping AI Agent activity list tables.

Service maps created by any AI Agent are non-operational by default and can be made operational from the application service record in the Mapped Application Services list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Agent processing quota

Each AI Agent can process up to 100 records by default. When an AI Agent reaches this limit, it stops processing and its status is set to **Paused**.

To raise the quota and reactivate the agents, first configure the following properties:

1. For the Service Mapping AI Agent:
   sn_sm_gen_ai.agent_invocation_limit.ServiceMapCreationAISpecialist
2. For the Tag-Based Service Map Creation AI Agent:
   sn_sm_gen_ai.agent_invocation_limit.TagBasedServiceMapCreationAISpecialist

You can change the value according to your needs, with consideration of your AI resources. After raising the quota, reactivate the Service Mapping AI Agents in the Service Mapping workspace.

Note: Raising the quota results in a higher AI resource consumption.

Service Mapping MCP tools

The Service Mapping tools expose live application service data and enable AI clients to query service topology, identify mapping gaps, and create application services in natural language.

The Service Mapping MCP tools provide query and create processes for investigating, visualizing, and building service topology, by implementing the Model Context Protocol (MCP) on the ServiceNow AI Platform. The MCP gives AI clients structured, secured, tool-based access to your data in the ServiceNow® instance.

Once a license is obtained, the admin can set up the CMDB MCP Server. Then users can connect Claude Desktop, and use it to query in natural language and create services. Claude selects the appropriate tool and retrieves live data or creates records directly from the ServiceNow® instance over OAuth 2.0 and JWT authentication. The results are presented in several visualizations or confirmations.

Benefits

Reduced time to insight

Getting a complete picture of an application service through a natural-language query in Claude.

AI access to live Service Mapping data

Without the MCP tools, Claude has no programmatic access to live Service Mapping state. They must rely on static snapshots or user-provided context, leading to stale or inaccurate outputs. The MCP tools give Claude access to current data directly from the instance at query time.

Visibility into traffic connections outside the CMDB

CIs connected only by observed TCP traffic, with no matching Configuration Management Database (CMDB) relationship, are invisible to operators until something breaks. The get_unmapped_topology tool surfaces these CIs so admins can prioritize mapping work proactively.

Service creation at scale

The create_top_down_application_service tool enables bulk onboarding of applications without manual efforts. An admin can prompt "create services using these entry points" and the tool creates application service records, detects entry point types (HTTP vs TCP), and triggers discovery. For detailed information on prompts and outputs, see Service Mapping MCP tools reference.

Secure, role-controlled access

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The MCP tools enforce the same ACLs and role permissions that govern standard ServiceNow REST API calls. Each request is executed under the authenticated user's session using caller-scoped data access (GlideRecordSecure). OAuth 2.0 with JWT tokens is used to authenticate the AI client connection.

All six tools require the sm_mcp_user role.

No additional scripting required

The six Service Mapping tools ship as part of the CMDB MCP Server scoped application. Once the server is activated and the AI client is connected, the tools are available immediately. All tools return structured, consistent responses and degrade gracefully when CMDB data is incomplete, returning partial results with a warning flag rather than a hard failure.

How the Service Mapping MCP tools work

The Service Mapping MCP tools are built on the following technical stack:

Six scripted REST API tools

Five tools map to read-only scripted REST API endpoints under /api/sn_sm_gen_ai/ on the ServiceNow instance. One tool, create_top_down_service, maps to a write-enabled endpoint that creates cmdb_ci_service_discovered records and triggers asynchronous Service Mapping discovery. All tools are registered as REST API type tools in the MCP Server Console.

MCP Framework integration

The server uses MCP Framework version 1.4.2 (sn_mcp v1.4.2). The MCP Discovery Layer at `/sncapps/mcp-server/mcp/...` receives tool invocation requests from external clients and routes them to the scripted REST API.

OAuth 2.0 and JWT authentication

External clients authenticate using OAuth 2.0 Authorization Code grant with JWT token format. ACL filtering and role-based scope enforcement are applied at the scripted REST API layer. The OAuth inbound integration must use the service_mapping_mcp_auth_scope authorization scope, limited to the MCP Tools API and CMDB Mcp Api. Using a different scope results in a successful OAuth connection but no tools visible in the AI client.

Caller-scoped data access

Business logic is executed by the Service Mapping MCP tools, ensuring data is returned only for CIs and services the authenticated user is permitted to access.

Write operations (service creation) are performed under the authenticated user's permissions. A user with the sm_mcp_user role can call the create_top_down_application_service tool.

The data sources are CMDB Services tables, CMDB relationships, TCP Traffic, and service record creation tables.

Asynchronous discovery on write

When create_top_down_service creates a new service record, Service Mapping discovery job is initiated from the provided entry point. Discovery runs asynchronously in the background and does not block the tool response. The tool returns immediately to the AI client with the new service system ID. Claude can continue with topology queries or other operations without waiting for discovery to complete.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scale limits

The Service Mapping MCP tools enforce the following scale limits to maintain performance.

Edge maximum

2,000 edges per topology response.

Traversal depth maximum

4 levels of relationship depth per query.

Response time target

Under 5 seconds per tool call. The get_all_application_service_names tool targets under 3 seconds, enforced through pagination with a page size capped at 800 entries.

For application services that approach these limits, request summary data rather than full topology to stay within the bounds. For example, ask for member count and edge count only, rather than the full topology.

Available tools

The CMDB MCP Server provides six tools: five read-only query tools and one write tool for service creation. All tools are consumable by Claude via the MCP protocol without additional transformation.

Access to each tool is controlled by the same ACLs that apply to the corresponding ServiceNow REST API. If the authenticated user does not have the required role, the tool returns an authorization error. For detailed role requirements, see Configure roles for the Service Mapping MCP tools.

For detailed input and output specifications and example queries, see Service Mapping MCP tools reference.

get_all_application_service_names

Returns a catalog of all application service names. Use this tool to browse the service catalog or locate a specific service before querying its topology.

The following example shows get_all_application_service_names returning a categorized portfolio view of all mapped services in an instance, broken down by service type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

get_all_application_service_names output

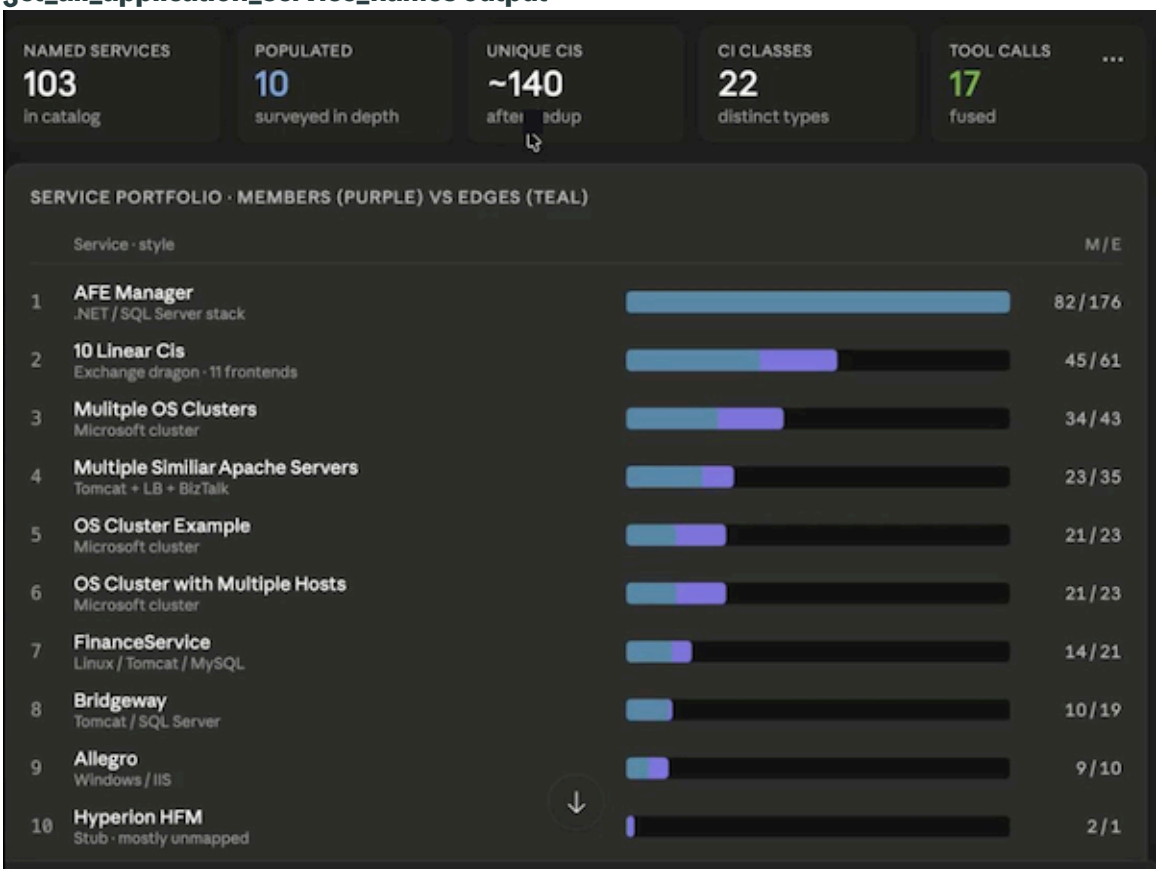

get_all_application_service_for_server

Returns all application services that include a specified server as a member CI. Use this tool to assess which services are at risk when a server is degraded or undergoing maintenance.

The following example shows get_all_application_service_for_server identifying every named service at risk from a single Windows server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

get_all_application_service_for_server output

get_application_service_topology

Returns the full topology of a specified application service, including all member CIs and the CMDB relationships (edges) connecting them, with CI class, relationship type, and direction. Reflects the current state of the service map, not a cached version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

get_application_service_topology output

get_server_impact_graph

Given a server CI, returns the observed connections using TCP traffic edges and CMDB relationships. Use this tool to walk the graph outward from a host and view how far the cascade goes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

get_server_impact_graph output

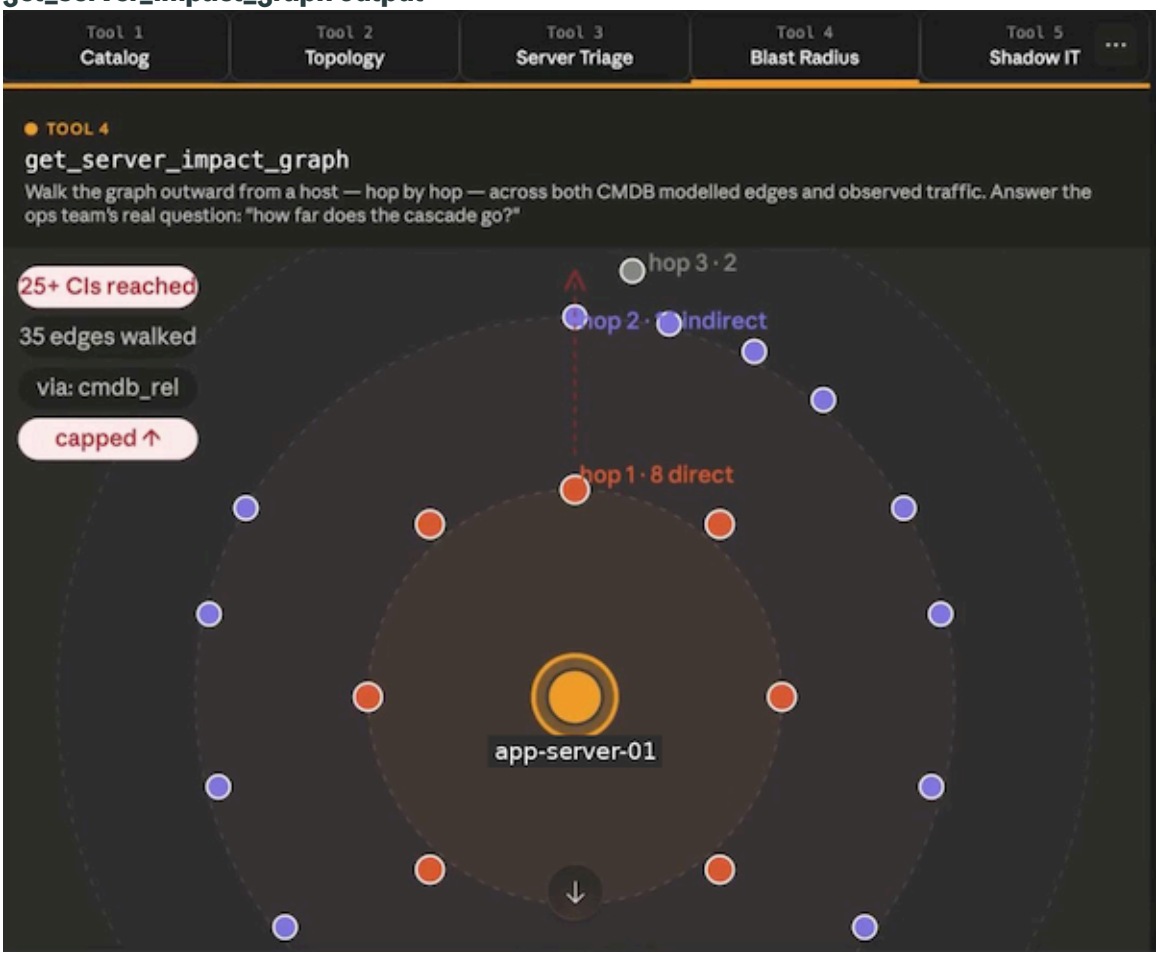

get_unmapped_topology

Starting from a single server or application CI, returns the CIs reachable from it by observed TCP traffic only, without traversing CMDB relationships. Use this tool alongside get_server_impact_graph to see what a CMDB-only map is missing.

The following example shows get_server_impact_graph and get_unmapped_topology side by side for the same server, illustrating the difference between CMDB-modeled edges (left) and traffic-only connections not already captured in the CMDB (right).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

get_unmapped_topology output

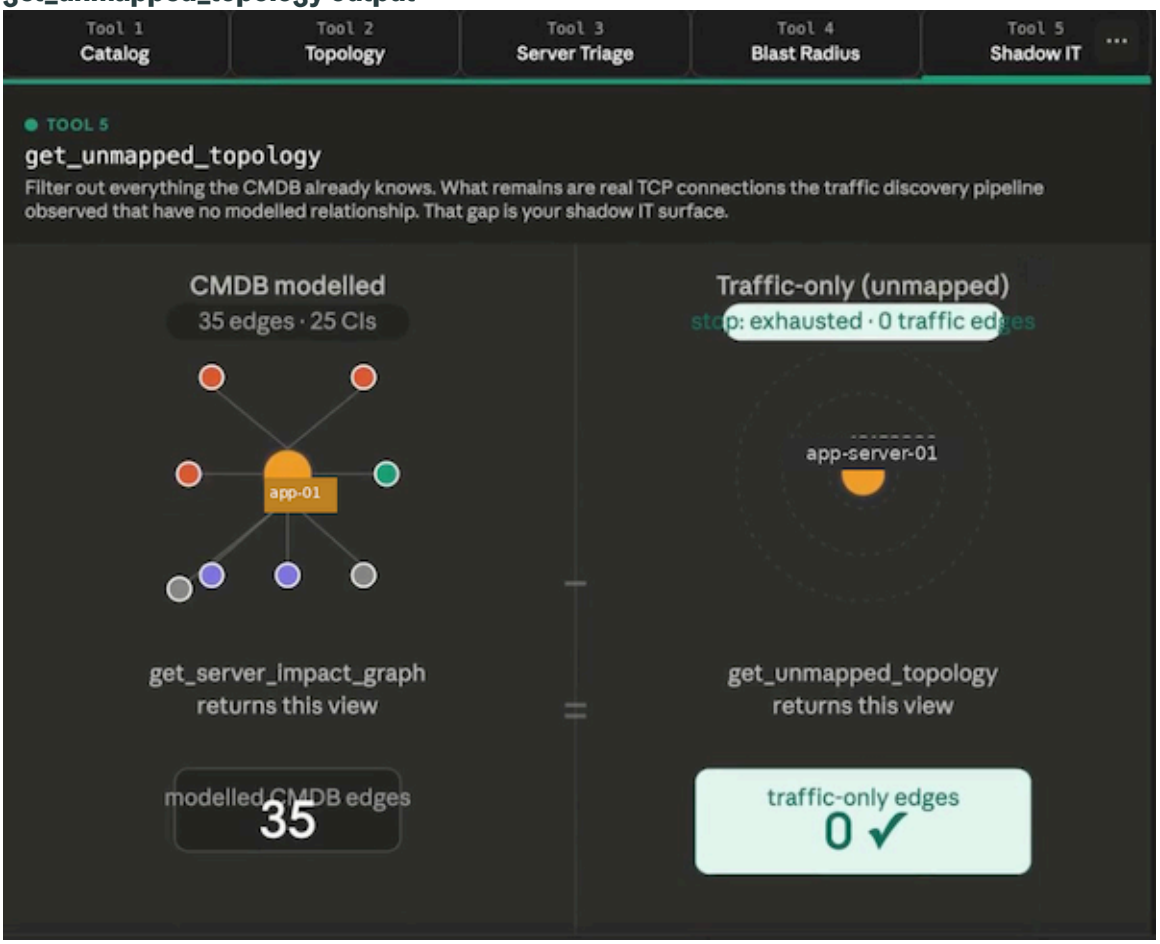

create_top_down_service

Creates a new application service from one or more entry points (HTTP URLs or TCP host-and-port pairs). Automatically detects each entry point's type and validates inputs. Use this tool to onboard applications, stand up services on demand, or automate service creation workflows.

Requires the sm_mcp_user role. For detailed information, see Service Mapping MCP tools reference

Setting up the Service Mapping MCP tools

Setting up the Service Mapping MCP tools involves sequential tasks performed by a system administrator, followed by a connection step performed by each Service Mapping user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping MCP tools setup flow

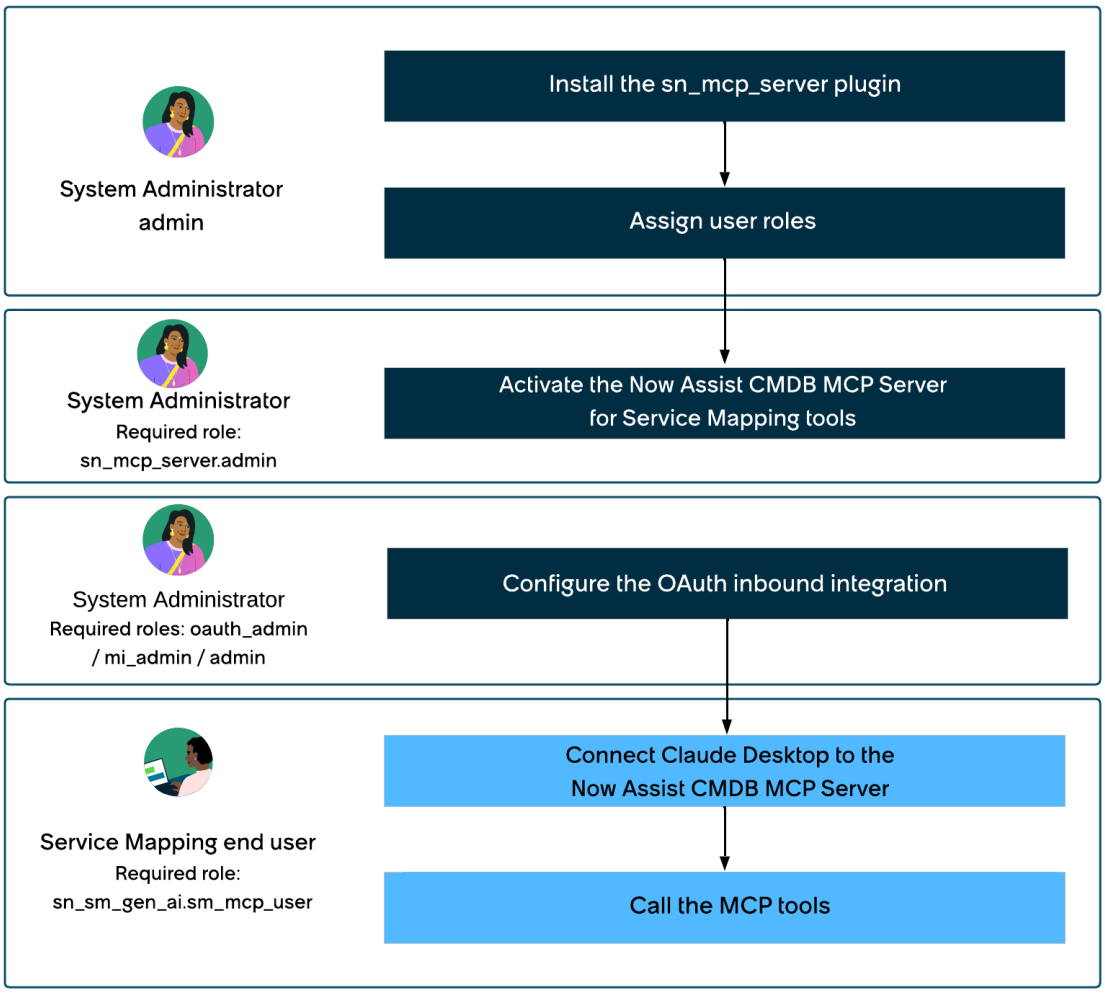

Configure roles for the Service Mapping MCP tools

Assign the required roles to users so they can connect to the CMDB MCP Server and call the Service Mapping MCP tools.

Before you begin

Before assigning roles, confirm the following requirements are met.

- You have the latest version of MCP Platform Manager plugin activated.
- You have the CMDB MCP Server [sn_cmdb_mcp_server], version 1.1.1, application installed.

Role required: admin

About this task

For information about the Service Mapping tools, see Service Mapping MCP tools.

The sn_sm_gen_ai.sm_mcp_user role grants access to all six Service Mapping MCP tools, including the create_top_down_application_service write tool.

sn_sm_gen_ai.sm_mcp_user automatically includes service_mapping_user and the MCP platform role sn_mcp_server.viewer. No manual role containment configuration is needed.

Note: If you configured role containment manually in a previous release, no action is needed. The shipped containment does not conflict with existing records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following table describes the roles involved and the access each one grants.

Role	Type	Granted rights
sn_sm_gen_ai.sm_mcp_admin	MCP admin role	Contains sn_sm_gen_ai.sm_mcp_user and service_mapping_admin. Grants no access to the Service Mapping MCP tools beyond what sn_sm_gen_ai.sm_mcp_user already grants; service_mapping_admin covers general Service Mapping administration, unrelated to the MCP tools.
sn_sm_gen_ai.sm_mcp_user	MCP access role	Access to all six Service Mapping MCP tools, including create_top_down_application_service, enforced by the REST endpoint ACL. Contains service_mapping_user and sn_mcp_server.viewer.
service_mapping_admin	Standard Service Mapping role	Administrative access to Service Mapping configuration. Included automatically under sn_sm_gen_ai.sm_mcp_admin.
service_mapping_user	Standard Service Mapping role	Read access to application service maps and topology data. Included automatically under sn_sm_gen_ai.sm_mcp_user.
sn_mcp_server.viewer	MCP platform role	Grants the ability to discover and invoke tools on an MCP server. Included automatically under sn_sm_gen_ai.sm_mcp_user.

Note: Role checks for Service Mapping tools and CMDB tools are independent, even though both tool groups run under the same MCP server. A user with only a Service Mapping role gets an authorization error on every CMDB tool, and a user with only a CMDB role gets an authorization error on every Service Mapping tool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The role containment hierarchy

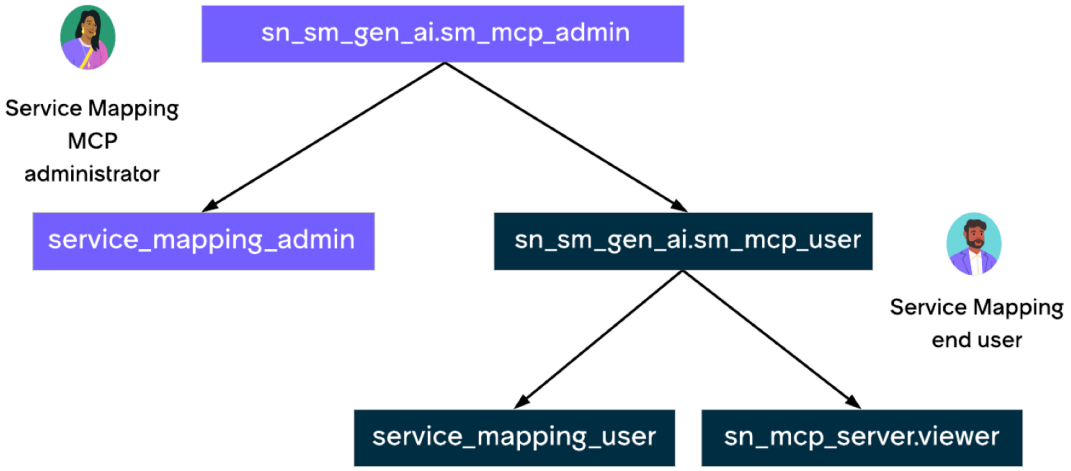

Procedure

1. Navigate to **All** > **User Administration** > **Users** and open the record of a user who needs access to the Service Mapping MCP tools.
2. Scroll to the **Roles** related list and select **Edit**.
3. Add the **sn_sm_gen_ai.sm_mcp_user** role.
   This role grants access to all six Service Mapping MCP tools, including create_top_down_application_service. It automatically includes service_mapping_user and the MCP platform role sn_mcp_server.viewer; no additional roles need to be assigned.
4. Select **Save**.

Result

The user is assigned the sn_sm_gen_ai.sm_mcp_user role and can call all six Service Mapping MCP tools, including create_top_down_application_service. The role automatically includes service_mapping_user and the MCP platform role sn_mcp_server.viewer.

What to do next

Activate the CMDB MCP Server for Service Mapping tools

Activate the CMDB MCP Server for Service Mapping tools

Activate the CMDB MCP Server and configure the OAuth inbound integration so that external AI clients can connect to your ServiceNow® instance and query application service data.

Before you begin

Before activating the CMDB MCP Server, confirm the following requirements are met.

- You have MCP Platform Manager version 1.4 or later installed.
- You have the CMDB MCP Server [sn_cmdb_mcp_server], version 1.1.1 or later, application installed.
- You have the roles required as described in Configure roles for the Service Mapping MCP tools.

Important: The Service Mapping tools are added to the CMDB MCP Server only if com.sn.itom.sm.gen.ai was active when the CMDB MCP Server application was installed. If com.sn.itom.sm.gen.ai was activated after the application was installed, reinstall the CMDB MCP Server [sn_cmdb_mcp_server] application to add the Service Mapping tool records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin assigned with the sn_mcp_server.admin role, and the oauth_admin (or mi_admin) role

The role containment chain

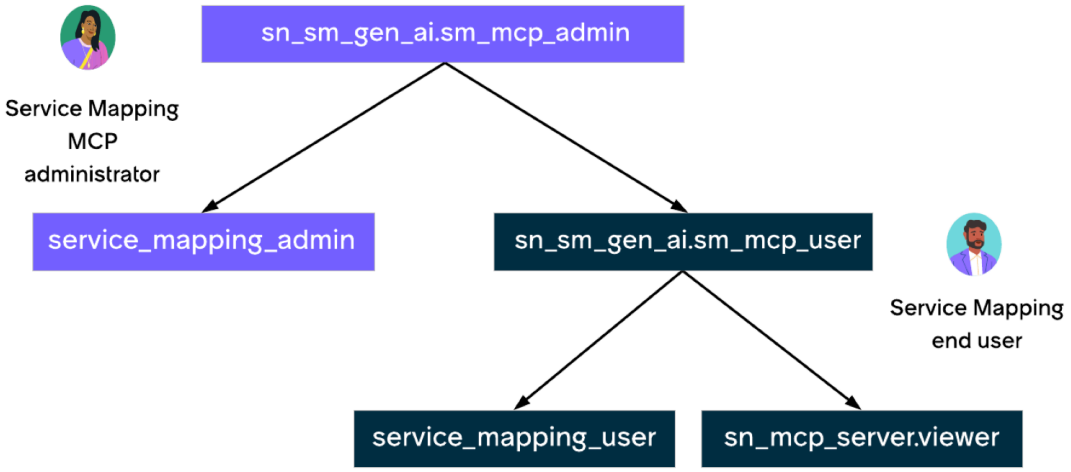

About this task

The Service Mapping MCP tools are delivered as part of the CMDB MCP Server. Activating it exposes five read-only tools and one write tool that AI clients can use to query live application service data. The OAuth inbound integration for this server is auto-generated when the server is installed. You must locate this integration, verify its configuration, and add the redirect URL for the AI client before users can authenticate.

For detailed information, see Service Mapping MCP tools

Procedure

1. On your ServiceNow instance, navigate to **All** > **MCP Server Console**.
2. From the **Configuration** tab, select **Servers**.
3. Select the **CMDB MCP Server** tile.

   Note: If the server does not appear, change the application scope to **ITOM AI Agents for Service Mapping**. The MCP Server Console screen opens with all fields populated, including the server URL and short description.

4. From the **Deactivate** menu, select **Activate**.
   The server status is set to **Active**. The individual tools are active by default.

5. Configure the OAuth inbound integration so that AI clients can authenticate with your instance.

   Important: Do not create a new OAuth client for this integration. The integration named sn_cmdb_mcp_server.cmdb_mcp_server is auto-generated when the server is installed. Creating a custom client requires manual auth scope configuration, and misconfiguration is silent — the connection succeeds but no tools appear in the AI client.

   a. Set the scope to **Global**.

   b. Navigate to **All** > **Machine Identity Console** and select the **Inbound integrations** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Locate and open the integration named **sn_cmdb_mcp_server.cmdb_mcp_server**.

> **Note:** The inbound integrations list may also contain integrations whose names match other server names, including underscores. These are monitoring integrations created by AI Control Tower and should not be used for client connections.

d. Verify that the **Auth scope** field is set to **service_mapping_mcp_auth_scope**, limited to **MCP Tools API** and **CMDB Mcp Api**.

> **Important:** The auth scope must be service_mapping_mcp_auth_scope. Using a different scope, such as useraccount, results in a successful OAuth connection but no tools appear in the AI client.

e. In the **Redirect URLs** field, add the redirect URL for the AI client.
For Claude Desktop, use: `https://claude.ai/api/mcp/auth_callback`.

f. Select **Save**.

6. Note the **Client ID** and **Client Secret** displayed on the integration record, and provide them to users along with the server URL so they can configure their AI client connections.
The server URL follows this format: `https://<instance>.service-now.com/sncapps/mcp-server/mcp/sn_cmdb_gen_ai_cmdb_mcp_server`

Result
The CMDB MCP Server is active and the OAuth inbound integration is configured. Users can connect an AI client and authenticate using the Client ID, Client Secret, and server URL.

What to do next
Share the following details with users so they can connect their AI clients: the server URL, the OAuth Client ID, and the OAuth Client Secret. For instructions on connecting Claude Desktop, see Connect Claude Desktop to the Service Mapping MCP Server.

Connect Claude Desktop to the Service Mapping MCP Server

Add the Service Mapping MCP Server as a custom connector in Claude Desktop so you can query application service data from your ServiceNow® instance in natural language.

Before you begin
Before connecting Claude Desktop to the Service Mapping MCP Server, confirm the following requirements are met.

- Claude Desktop is installed on your computer.
- You have a Claude account with permission to add custom MCP connectors. If your organization manages Claude through an Enterprise or Team plan, verify that custom connectors are enabled in your organization's settings.
- The Service Mapping MCP Server is active, and you have the following details:
  - MCP Server URL
  - OAuth Client ID
  - OAuth Client Secret

  For more information, see: Activate the CMDB MCP Server for Service Mapping tools
- Verify that you have the required role configuration. For more information, see Configure roles for the Service Mapping MCP tools.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: sn_sm_gen_ai.sm_mcp_user

The role containment chain

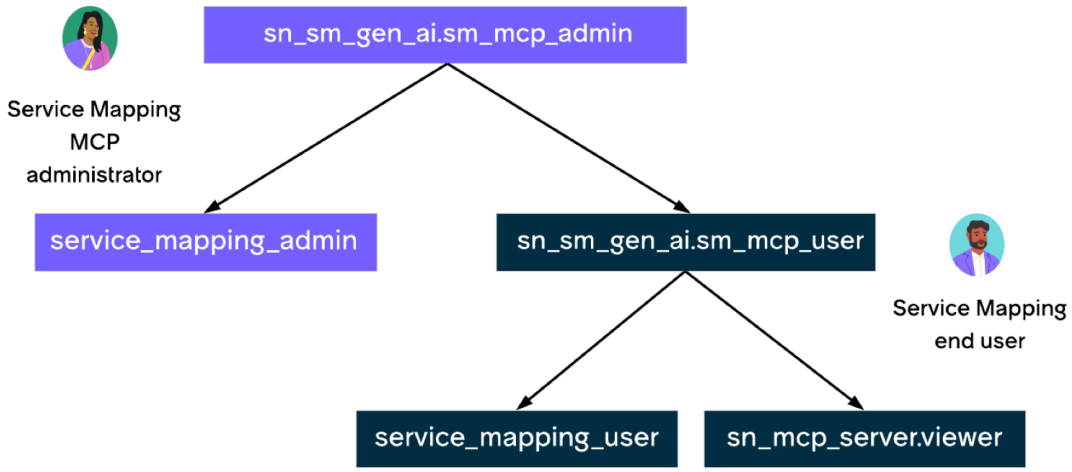

About this task

Claude Desktop supports custom MCP connectors that let Claude query external data sources on your behalf. After you add the Service Mapping MCP Server as a connector and authorize it with your ServiceNow credentials, Claude can call the Service Mapping MCP tools according to the prompt.

For detailed information, see Service Mapping MCP tools

Procedure

1. Open Claude Desktop on your computer.
2. Open **Settings**.
3. Select the **Connections** tab.
4. Select **Add custom connector**.
5. Fill in the connection details on the form.
   a. In the **Name** field, enter a name for the connector, for example `ServiceNow Service Mapping`.
   b. In the **Remote MCP server URL** field, enter the MCP Server URL.
   c. Select **Advanced settings**.
   d. In the **OAuth Client ID** field, enter the Client ID.
   e. In the **OAuth Client Secret** field, enter the Client Secret.
6. Select **Add**.
7. Complete the authorization in the browser.
   a. If prompted, enter your ServiceNow® user name and password and select **Log In**.
   b. Review the requested permission scope and select **Allow**.
   The browser displays: Claude connected successfully. Close the browser window and return to Claude Desktop.
8. Verify the connection.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Open a new chat in Claude Desktop.

b. Enter a test prompt, for example: `Use the get_all_application_service_names tool to list all application services.`
For more information about the prompts, see Service Mapping MCP tools reference.
Claude returns a list of application service names from your ServiceNow® instance.

Result
Claude Desktop is connected to the Service Mapping MCP Server. In any new chat, you can ask questions about your application services in natural language and Claude will query your ServiceNow® instance using the available tools.

What to do next
For the full list of available tools, input parameters, example queries, and usage guidelines, see Service Mapping MCP tools reference.

Service Mapping MCP tools reference

Details on the six Service Mapping MCP tools, including their inputs, outputs, and example natural-language queries for use with Claude, and service creation workflows.

The CMDB MCP Server exposes six tools that an MCP-compatible AI client can invoke to retrieve application service data from a ServiceNow instance and create application services. Five tools are read-only and do not create, update, or delete records. One tool, create_top_down_service, creates application service records and initiates Service Mapping discovery.

All six tools require the sm_mcp_user role.

The following five tools retrieve application service data without modifying records:

get_all_application_service_names

Returns a list of all application service names in the instance.

Input

An optional filter parameter to limit results by mapping type: pattern-based, tag-based, or calculated. If no filter is provided, all service types are returned. Optional *`limit`* and *`offset`* parameters control pagination. *`limit`* is capped at 800 entries per page.

Output

A paginated list of entries. Each entry includes:

- Service name
- System ID
- Service type

The response is paginated or bounded to prevent oversized payloads. The response includes a *`has_more`* flag and a *`next_offset`* value for retrieving subsequent pages.

Example queries

- "Use the ServiceNow Service Mapping tool, get_all_application_service_names, to list all application services."
- "List all tag-based application services in ServiceNow."

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

get_all_application_service_for_server

Returns all application services that include a specified server as a member CI.

Input

Server CI name or System ID.

Output

A list of application services. Each entry includes:

- Service name
- System ID
- Service type
- Mapping status

If no services are found for the specified server, an empty list is returned. This is not treated as an error. The list is capped at 150 entries; if the cap is reached, the response includes a flag noting the result set was capped. If the server name matches more than one CI, the tool resolves to the first match and the response includes a warning recommending a system ID for a deterministic result.

Example queries

- "Use the ServiceNow Service Mapping tool, get_all_application_service_for_server, to find which services contain server emse-10152008.servicenow.com"
- "Which application services include server haproxy-s?"

get_application_service_topology

Returns the full topology of a specified application service, including all member CIs and the CMDB relationships (edges) connecting them.

Input

Application service name or System ID.

Output

The complete topology of the service. Each CI entry includes:

- CI name
- CI class
- Relationship type
- Direction: upstream or downstream

The topology reflects the current state of the service map, not a cached version.

Example queries

- "Use the ServiceNow Service Mapping tool, get_application_service_topology, to get the topology for the Inclusion service. Just show me the member count and edge count, don't visualize."
- "Show me the topology of the Payroll service."

Usage note

Application services with a large number of member CIs and edges (more than 50 members) return a significant volume of data. When querying topology for large

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

services, ask Claude for a summary or specific counts to avoid reaching message-length limits on the Claude free tier.

get_server_impact_graph

Given a server CI, returns all CIs related to it via CMDB relationships and all CIs with observed TCP traffic connections to or from that server.

Input

server_names: Required. One or more server CI names or system IDs, comma-separated. *max_depth*: Optional. Maximum relationship depth to traverse, from 1 to 4. Defaults to 2. *servers_limit_per_start*: Optional. Maximum number of servers to traverse per starting point, from 1 to 25. *sources*: Optional. A comma-separated subset of traffic, cmdb_rel to limit which connection types are returned. Defaults to both.

Output

A set of related CIs. The response distinguishes between:

- CIs related via CMDB relationships. Uses [cmdb_rel_ci] or equivalent.
- CIs with observed TCP traffic connections to or from the specified server

This distinction enables AI-assisted identification of candidates for service map inclusion.

Example queries

- "Use the ServiceNow Service Mapping tool, get_server_impact_graph, to get impact analysis for server haproxy-s."
- "What CIs are related to server db-cluster-02 via CMDB or traffic?"

get_unmapped_topology

Starting from a single server or application CI, returns the CIs reachable from it by observed TCP traffic only. This tool doesn't traverse CMDB relationships. Use it together with get_server_impact_graph, which returns the fused CMDB-relationship-and-traffic view, to see what a CMDB-only map is missing.

Input

ci: Required. The name or system ID of the starting server or application CI. *max_depth*: Optional. Maximum traversal depth, from 1 to 4. Defaults to 2. *servers_limit_per_start*: Optional. Maximum number of servers to traverse per starting point, from 1 to 25.

Output

An adjacency graph of the CIs reached by traffic-only traversal. The response includes:

- An adjacency list and CI names for the traversed nodes
- The starting server or application CI
- A stop reason for the traversal, for example reaching the maximum depth or exhausting all connections
- The count of CIs visited
- The count of traffic edges found
- A fixed reminder that CMDB relationships are excluded and to use get_server_impact_graph for the fused view

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the input name matches more than one CI, the response also includes a warning noting that the tool resolved to the first match and recommending a system ID for a deterministic result. If the specified CI doesn't exist, or exists but isn't a supported class, the tool returns an error identifying which case applies.

Example queries

- "Use the ServiceNow Service Mapping tool, get_unmapped_topology, to show me what's connected to server db-cluster-02 by traffic only."
- "Use the ServiceNow Service Mapping tool, get_unmapped_topology, starting from application server app-01."

create_top_down_application_service

Creates an application service record from one or more entry points (HTTP URLs or TCP host-and-port pairs). The user provides a unique service name and at least one entry point that isn't already used by another service. The user can also provide optional metadata details. Each entry point's type is automatically detected from its format.

Role requirement

Requires the sm_mcp_user role, the same endpoint-level ACL that applies to all six tools.

Input

Required:

- name (string) — A unique name for the application service.
- entry_points (array) — One or more entry points. Each entry point is either a URL string (e.g., "https://payments.acme.com") or an object with host and port for a TCP endpoint (e.g., {host: "10.10.5.20", port: 6379}). At least one entry point is required.

Optional:

- short_description (string) — A brief description of the service purpose.
- business_criticality (string or integer) — Service criticality level, set on the business_criticality field as provided.
- service_classification (string) — Service classification category (e.g., "Business Service", "Application Service"). Maps to the service_classification field.
- owned_by (string) — Owner user name or system ID.
- support_group (string) — Support group name or system ID for operational support assignment.
- operational_status (string) — Operational status (e.g., "Operational", "Pilot", "Planned"). Defaults to Operational if not provided.

Output

On success:

- success (boolean) — true
- service_sys_id (string) — The system ID of the newly created application service record.
- http_status (integer) — 201 (Created).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

On failure:

- success (boolean) — false
- error_code (string) — One of missing_name, invalid_entry_points, duplicate_name, duplicate_entry_point, service_insert_failed, endpoint_insert_failed, or endpoint_link_failed.
- error_message (string) — Human-readable error description.
- http_status (integer) — 400 for missing_name or invalid_entry_points, 409 for duplicate_name or duplicate_entry_point, 500 for the remaining error codes.

Behavior

- Entry point format is validated before record creation. If any entry point format is invalid, no record is created and an error is returned.
- A service name or entry point already used by another service is rejected rather than merged into the existing record.
- Optional metadata fields are set directly on their corresponding CMDB service record fields, with no value translation. Provide values in the format the target field accepts.
- If the service record is created but a later step, such as linking an entry point, fails, the created service record is rolled back rather than left partially created.
- The created service record is immediately visible in ServiceNow and available for topology queries via get_application_service_topology.

Example queries

- "Create a new application service called 'Payment Gateway' that monitors the endpoint at https://payments.acme.com"
- "I need to onboard our Redis cache into Service Mapping. Create an application service called 'Cache Layer' pointing at host 10.10.5.20 on port 6379"
- "Set up a new top-down service called 'Order Processing API' with two entry points: the frontend at https://orders.acme.com and the backend TCP service at host 10.10.5.100 port 8443"
- "Create a new application service for our customer-facing checkout flow: Name: Checkout Service, Entry point: https://checkout.acme.com, Description: Handles all checkout and payment flows for the region, Business criticality: High, Classification: Business Service"
- "We need to map our new authentication service. Create it in Service Mapping with the entry point https://auth.acme.com, then once it's created show me the topology so I can confirm it exists"

Usage notes

- To give one service multiple entry points, pass them all in a single entry_points array rather than calling the tool once per entry point.
- For bulk service creation across separate services, invoke create_top_down_application_service multiple times with different service names. Each invocation is independent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The service created by this tool can be queried immediately with get_application_service_topology.
- If you provide all optional metadata fields at creation time, all service properties will be populated in a single call, reducing the number of follow-up edits needed.

Analyze potential impact agentic workflow

The Analyze potential impact agentic workflow analyzes how a change request might impact servers and services. This analysis helps you make informed decisions about the next steps regarding the change request.

Analyze potential impact agentic workflow overview

Using the Analyze potential impact agentic workflow, the Analyze potential impact agent produces an analysis of potentially impacted relevant servers and services.

Note: Relevant servers refer to those servers that are currently operational and in use, excluding any non-operational or retired servers.

AI agent used in the Analyze potential impact agentic workflow

AI Agent and its role

AI Agent	AI Agent role
Analyze Potential Impact Agent	Analyzes the potential impact of a change on relevant servers and services and generate an impact analysis.

Generating the impact analysis

The Analyze potential impact agentic workflow uses the agent and its related skills to evaluate the potential impact of a change request. When you select the Analyze potential impact agent in the ServiceNow Otto panel, the agent initiates the workflow as follows:

- Prerequisite verification: The agent verifies that all prerequisites have been met.
- Change request identification: If you have a change request open, the agent retrieves the change request number from the current active page. Otherwise, it prompts you to provide the change request number.
- Server selection: The agent selects up to 10 affected servers from the configuration items (CIs) in the change request.
- Match identification: The agent identifies matches between servers and suggested services.
- Impact analysis: Eventually, the agent prioritizes and displays up to 10 impacted relevant servers, giving priority to servers that are part of services, and up to 3 impacted services. Additionally, the agent provides you with a summary about each service and how it might be impacted by related servers. This information is displayed in the ServiceNow Otto panel.
- Saving the impact analysis: The analysis generated by the Analyze potential impact agentic workflow is saved to the change request work notes.

To learn more about using the workflow, see Assess a change request with the Analyze potential impact workflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Analyze potential impact agentic workflow sample

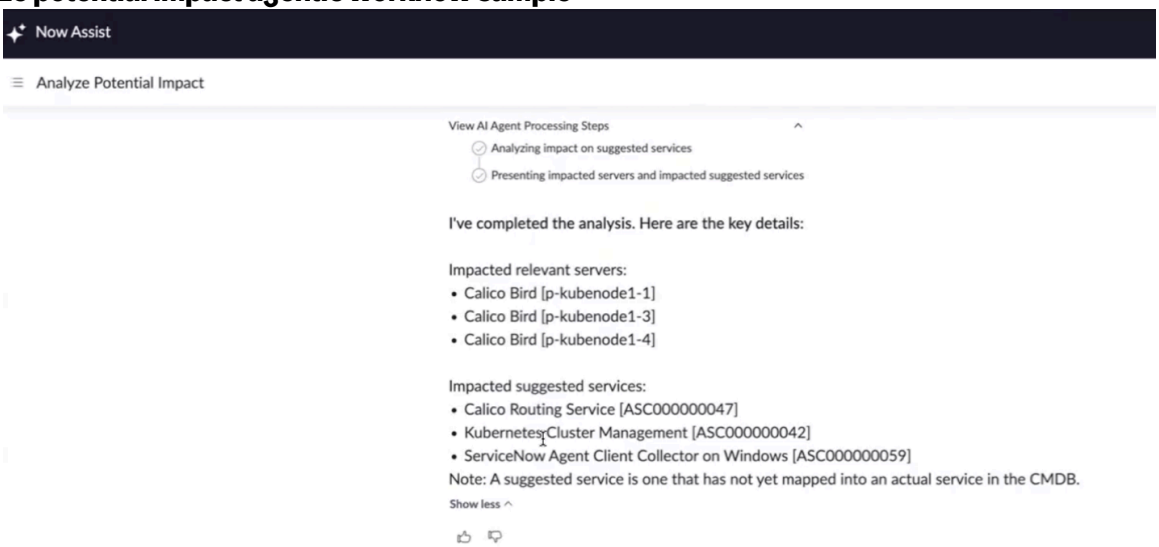

report

Domain separation

Currently, the Analyze potential impact agentic workflow doesn't support instances with domain separation. For optimal results, conduct all conversations in English.

Service Mapping Skills

Service Mapping skills are intelligent capabilities that enhance the Analyze potential impact agentic workflow. These skills leverage generative AI to identify, classify, and assess the impact of changes to your IT infrastructure.

Service Mapping provides two primary Gen AI skills:

- Service Mapping Candidate
- Service Mapping Candidates Impact

These Service Mapping skills are active by default. For information on activating inactive skills, see:

- Activate the Service Mapping Candidate skill
- Activate the Service Mapping Candidates Impact skill

Service Mapping Candidate skill

This skill automatically identifies and names processes and application service candidates by analyzing process characteristics, commands, and parameters. It provides detailed descriptions and categorization to help you understand which services are running in your environment.

The skill uses a two-stage process:

1. In the process classification stage, the agent uses the skill to analyze individual processes within an application service candidate. The agent uses the skill to extract:
   - Publisher: The company or organization that released the product. For example: "Project Calico".
   - Product: The specific product name. For example: "Calico".
   - Description: Detailed explanation of what the process does. For example: "BIRD (BIRD Internet Routing Daemon) process running as part of Calico networking solution for

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes and container orchestration. This process handles BGP routing functionality with remote control socket enabled, running as a daemon with specific Calico configuration for pod-to-pod networking and network policy enforcement."

- Service Hints: Keywords that help identify the service type. For example: "bird,calico,BGP,routing,daemon,networking,kubernetes,container,policy,socket".

2. In the service information generation stage, the agent combines process details with load balancer information to generate:
   - Service Name: A concise, accurate name for the service candidate. For example: "Calico BGP Routing Service".
   - Service Description: A comprehensive description of the service's purpose. For example: "BIRD Internet Routing Daemon process running as part of Project Calico's networking solution for Kubernetes container orchestration. This service handles BGP routing functionality for pod-to-pod networking and network policy enforcement, operating as a daemon with remote control socket capabilities. The component is connected through Calico's container networking infrastructure to provide routing services in Kubernetes environments".
   - G2 Category: Industry-standard categorization from G2.com. For example: "Container Networking".

Service Mapping Candidates Impact skill

This skill generates comprehensive impact assessments when changes are made to your infrastructure. It analyzes service connections and dependencies to predict how a change to a single component might affect other services and servers. In addition, it provides a summary of the impact analysis.

The skill analyzes:

- Connection Topology: How servers within a service candidate are connected.
- Affected Servers: Which servers are directly impacted by a change.

Based on the candidate number, and the analysis, the output summary is provided per a service candidate. For example:

- Service candidate: "bird [ASC000000015]".
- Impact: "Network routing disruption affecting pod-to-pod communication across Kubernetes cluster nodes due to Calico Bird BGP routing daemon failure on p-kubenode1-2, potentially causing connectivity issues for workloads on connected nodes p-kubenode1-3, p-kubenode1-4, and p-kubenode1-5".

Important: This agentic workflow is turned on by default. For more information, see AI agents, skills, and agentic workflows on by default.

Assess a change request with the Analyze potential impact workflow

Use the Analyze potential impact agentic workflow to assess the effects of a change request. This workflow highlights the potential impacted servers and suggested services to help you identify risks, benefits, and make any necessary adjustments.

Before you begin

Using the Analyze potential impact agentic workflow requires the installation of the following:

- Service Mapping Plus version 1.16.3.
- ServiceNow Otto for IT Operations Management (ITOM) version 2.3.0.
- ServiceNow Otto for IT Service Management (ITSM) version 11.1.3.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: sn_change_write

Role masking

Required role:

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with your applications have specific roles that come included with the application. If you select *Users with specific roles* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

About this task

The agent verifies the prerequisites at the beginning of the agentic workflow. If any prerequisites aren't met, the workflow stops, and you're informed so that you can acquire the prerequisites. For more information about the agentic workflow, see Analyze potential impact agentic workflow.

Procedure

Either initiate the Analyze potential impact agentic workflow using a change request number or from an open change request.

Situation	Steps
Change request number is known	**a.** From the header menu, select the Service Now Otto icon to open the ServiceNow Ot to panel. **b.** Select **Analyze Potential Impact**. **c.** When prompted, enter the change request number.
Open change request	**a.** Navigate to **Workspaces** > **Service Operations Workspace** > **.** **b.** Select the list icon. **c.** Navigate to **Changes** > **All**. **d.** Select the applicable change request from the list. **e.** From the header menu, select the Now Assist icon to open the Now Assist panel. **f.** Select **Analyze Potential Impact**.

The Analyze potential impact agentic workflow is initiated and the agent starts with the verification of the prerequisites.

Result

In the panel, the Analyze potential impact agentic workflow generates an analysis of the potentially impacted relevant servers and suggested services.

The impact analysis is updated in the work notes in the change request details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuring Service Mapping

You get started with Service Mapping by configuring roles, credentials, and MID Server connections.

Setting up Service Mapping is the first stage in the Service Mapping workflow.

Service Mapping workflow

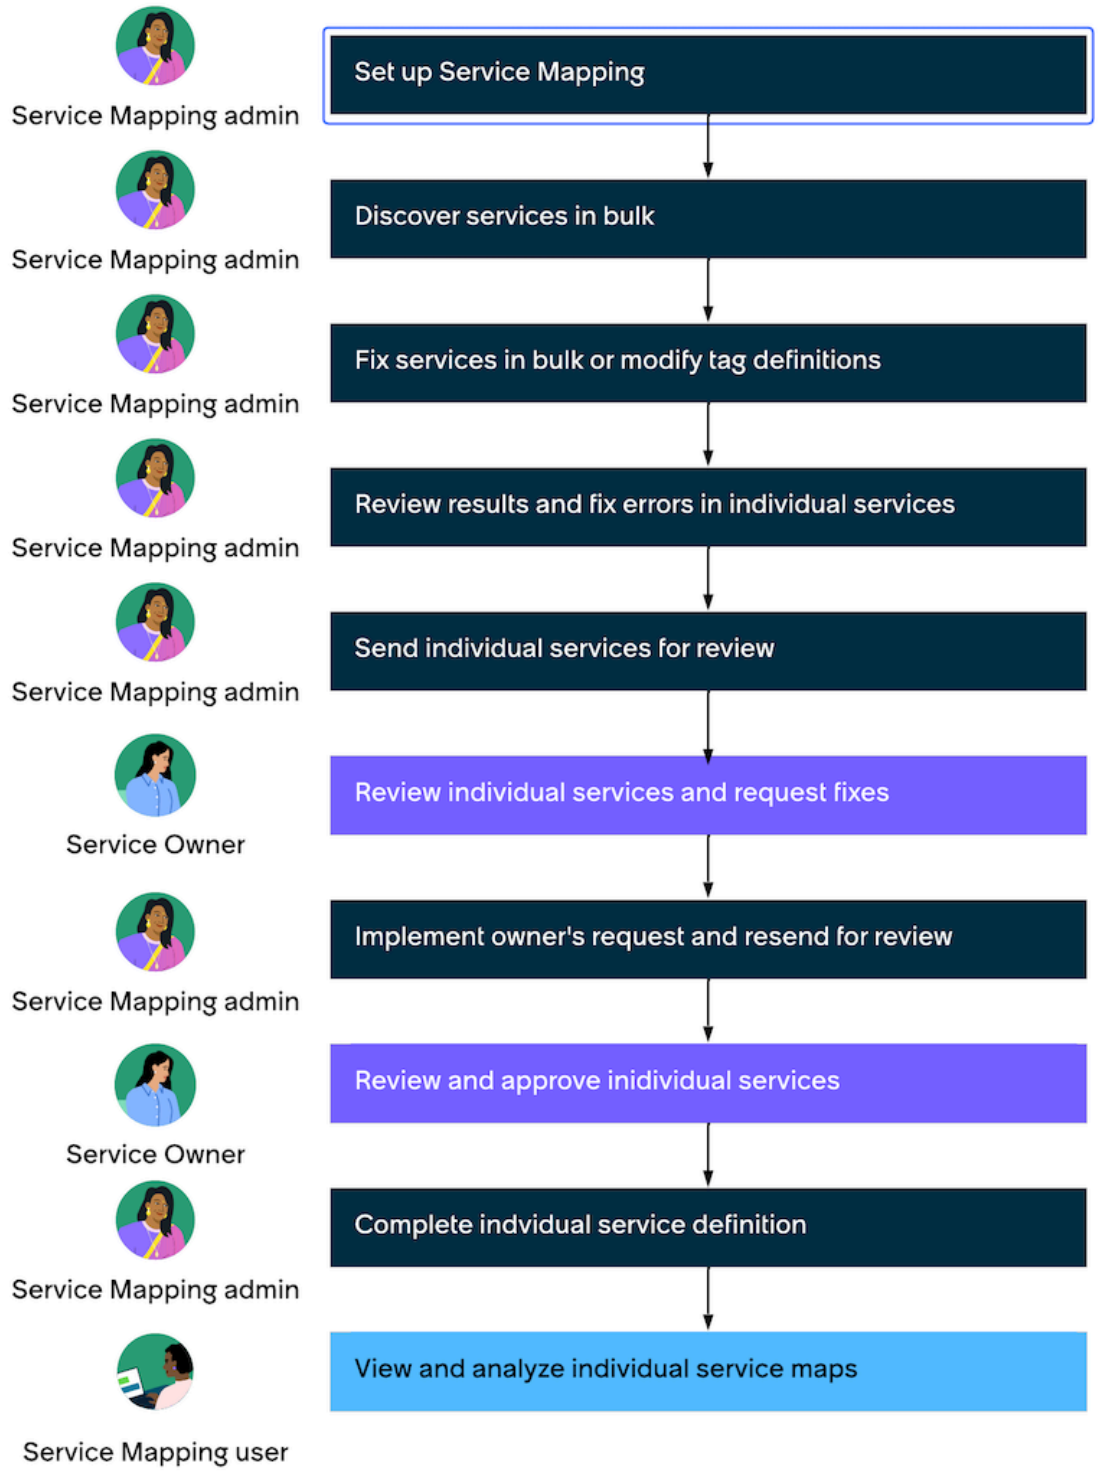

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping is part of the ServiceNow AI Platform and deploys some of its platform-wide mechanisms and features. At the same time, there are some configurations that are specific to Service Mapping only.

> ⚠ **Warning:** The Service Mapping setup for the Advance Release is different from the standard Service Mapping setup.

Perform the following tasks in the exact order they are listed below:

1. Request Service Mapping.
2. Install Service Mapping Plus.
3. Install and configure MID Server. MID Servers, which are located in the enterprise private network, facilitate communication between servers on the network and some ServiceNow applications, such as Service Mapping, and Discovery. For more information, see MID Server configuration for Service Mapping.
4. Verify that Discovery is set up and runs horizontal discovery as expected.
5. Configure credentials required for host discovery.
6. Configure credentials required for Service Mapping to access applications inside your organization private network. See Prerequisites for performing top-down discovery using Service Mapping.
7. If your organization network contains load balancers running on a Linux host, create a discovery behavior for load balancers.
8. Grant the following Service Mapping roles to relevant users:

   service_mapping_admin

   Sets up the Service Mapping application. Maps, fixes, and maintains service instances. Also performs advanced configuration and customization of the product. Assign this role to application administrators.

   service_mapping_user

   Views maps for operational service instances to plan change or migration, as well as analyze the continuity and availability of services. Assign this role to application users.

   sm_app_owner

   Provides information necessary for successful mapping of a service instance. Once a service is mapped, this user reviews the results and either approves it or suggests changes. Assign the sm_app_owner role to users who own service instances and are familiar with the infrastructure and applications that make up the services.

9. Verify that Service Mapping is set up properly.
10. If your organization has a ServiceNow deployment with customized Discovery or CMDB attributes, perform additional configuration described in KB0647574: Preparing customized ServiceNow deployments to work with Service Mapping.

In addition to the obligatory setup configurations described here, you may need to perform additional configurations after you complete the initial mapping process. For more information, see Advanced Service Mapping configuration.

Request Service Mapping

Service Mapping is available under ITOM Visibility subscription and requires installation and updates on the ServiceNow Store. Visit the ServiceNow Store to view all the available apps, and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes. For details, see Request Service Mapping.

Before you begin

Role required: none

To purchase a subscription, contact your ServiceNow account manager. The account manager can arrange to have the plugin activated on your organization's production and subproduction instances, generally within a few days.

If you don't have an account manager, decide to delay activation after purchase, or want to evaluate the product on a subproduction instance without charge, follow these steps.

Role required: admin

About this task

The following plugins are activated automatically when the Service Mapping plugin (com.snc.service-mapping) is activated:

- Discovery (com.snc.discovery)

  Note: If you're activating the Discovery plugin for the first time in the Brazil release, install the Discovery (com.snc.discovery) plugin first. Then install the Discovery and Service Mapping Patterns (com.sn_itom_pattern) plugin. You can then upgrade to Brazil.

- Pattern Designer (com.snc.pattern.designer)
- Cloud Provisioning and Governance Core (com.snc.cloud.core)
- Performance Analytics – Content Pack – Service Mapping (com.snc.service-mapping.pa.content)
- Event Management and Service Mapping Core (com.snc.service-watch)

  Note: The Event Management and Service Mapping Core (com.snc.service-watch) plugin is different from the Event Management plugin (com.glideapp.itom.snac).

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Related topics

Credentials required for host discovery

Install Service Mapping Plus

You can install the Service Mapping Plus application (sn_sm_scoped_app) if you have the admin role. The application installs related ServiceNow® Store applications and plugins if they are not already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

The following items are installed with Service Mapping Plus:

- Scheduled jobs
- Tables

For more information, see Components installed with Service Mapping Plus.

To view a list of Service Mapping plugin dependencies, see Plugins or applications installed with ITOM Visibility.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Service Mapping Plus application (sn_sm_scoped_app) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

MID Server configuration for Service Mapping

Configure Service Mapping and MID Servers to work together.

What is the MID Server

The Management, Instrumentation, and Discovery (MID) Server is a Java application that runs as a Windows service or UNIX daemon on a server. MID Servers, which are located in the enterprise private network, facilitate communication between servers on the network and some ServiceNow applications, such as Service Mapping, and Discovery.

Note: If the Service Mapping (com.sn_itom_pattern) plugin is activated on the instance and a language plugin (for example, Spanish) is also activated, make sure that the MID Server language is also set to English/None on the user record.

MID Server selection criteria

Most environments require multiple MID Servers, with Service Mapping using the relevant MID Server for discovery. MID Servers have the following selection criteria that Service Mapping can use to choose the relevant MID Server:

- Application — defines what application a MID Server works with. Set it to Service Mapping to reserve this MID Server exclusively to Service Mapping discovery requests. Alternatively, set it to ALL to allow any ServiceNow application to use this MID Server.
- Capability — defines the network capability. For Service Mapping, set this parameter to ALL or any combination of SSH, WMI, SNMP, and Cloud Provisioning and Governance . See Configure MID Server capabilities.
- IP range — limits operation of this MID Server to this IP range. Service Mapping does not choose this MID Server for a discovery request whose endpoint is outside this IP range. See Configure an IP address range for the MID Server.

  Note: For information on how to bypass MID Server configuration based on IP range, see Fine-tune Service Mapping with MID affinity and IP reuse.

Service Mapping selects a MID Server using the following algorithm:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Service Mapping chooses the MID Server whose selection criteria best match the parameters of the discovery request.
- If there are no MID Servers with matching selection criteria, Service Mapping chooses the default MID Server.
- If there are no MID Servers with matching selection criteria or default MID Server, Service Mapping cannot start the discovery process.

While by default Service Mapping uses this algorithm in all deployments upgraded from Istanbul or Jakarta, it can support both new and legacy algorithms for selecting a MID Server. For more information, see Choose MID Server selection algorithm.

Default MID Servers for Service Mapping

In addition to selection criteria, you can configure one of the MID Servers as the default server that Service Mapping uses. If there are no MID Servers with matching application, capability, or IP range, Service Mapping uses the default MID Server. See Configure a default MID Server for each application.

Using PowerShell for discovery

MID Servers can use PowerShell to directly communicate with Windows servers using both WMI and WinRM protocols. For Windows services using the WinRM protocol, the PowerShell process establishes a secure PSSession (PowerShell Remoting session) that stays open until the MID Server finishes querying a Windows server. For Windows servers using the WMI protocol, the PowerShell process sends every PowerShell command with credentials.

If you do not configure MID Servers to use PowerShell and PowerShell Remoting, MID Servers use WMI.

ServiceNow now supports PowerShell 3.0 up to 5.1.

MID Server credential-less discovery with Nmap

If the MID Server does not have sufficient credentials to access a device or application, it can run Network Mapper (Nmap) commands to collect basic information without using credentials. Credential-less discovery with Nmap requires additional configuration as described in Install and uninstall Nmap on a MID Server.

ServiceNow applications refer to devices and applications that comprise a service instance as configuration items (CIs).

Placing MID Servers

The number of MID Servers that you require and where you place them depends on your organization needs. If you want to map devices and applications inside your private network, place the MID Servers inside the private network. If you want to map devices and applications located in the DMZ, place the MID Servers both in the DMZ and inside the private network.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Placing MID Servers inside private network

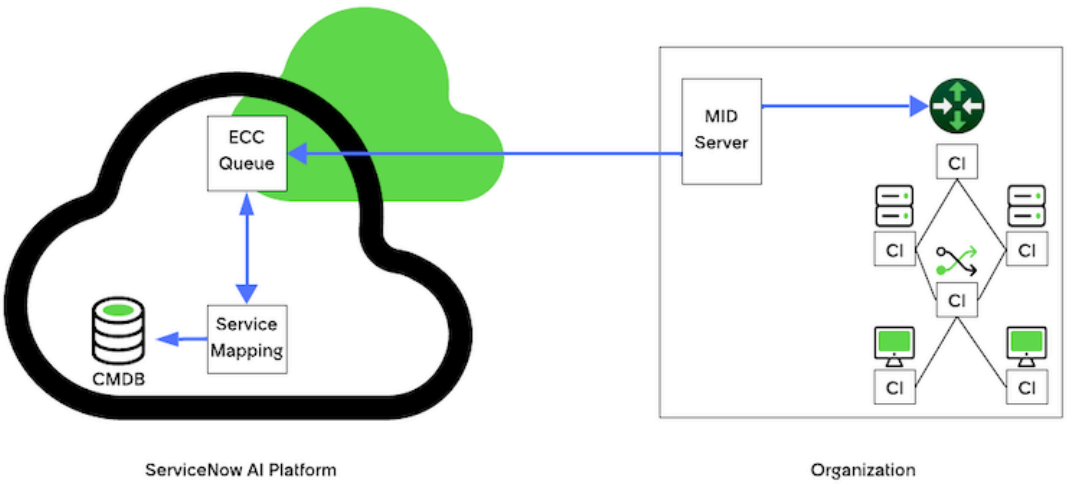

In deployments where domain separation is enabled and domains are configured to form a hierarchy, MID Servers must be placed in the lowest domain level, a "leaf domain".

Placement of MID Servers in domain-seprated environments

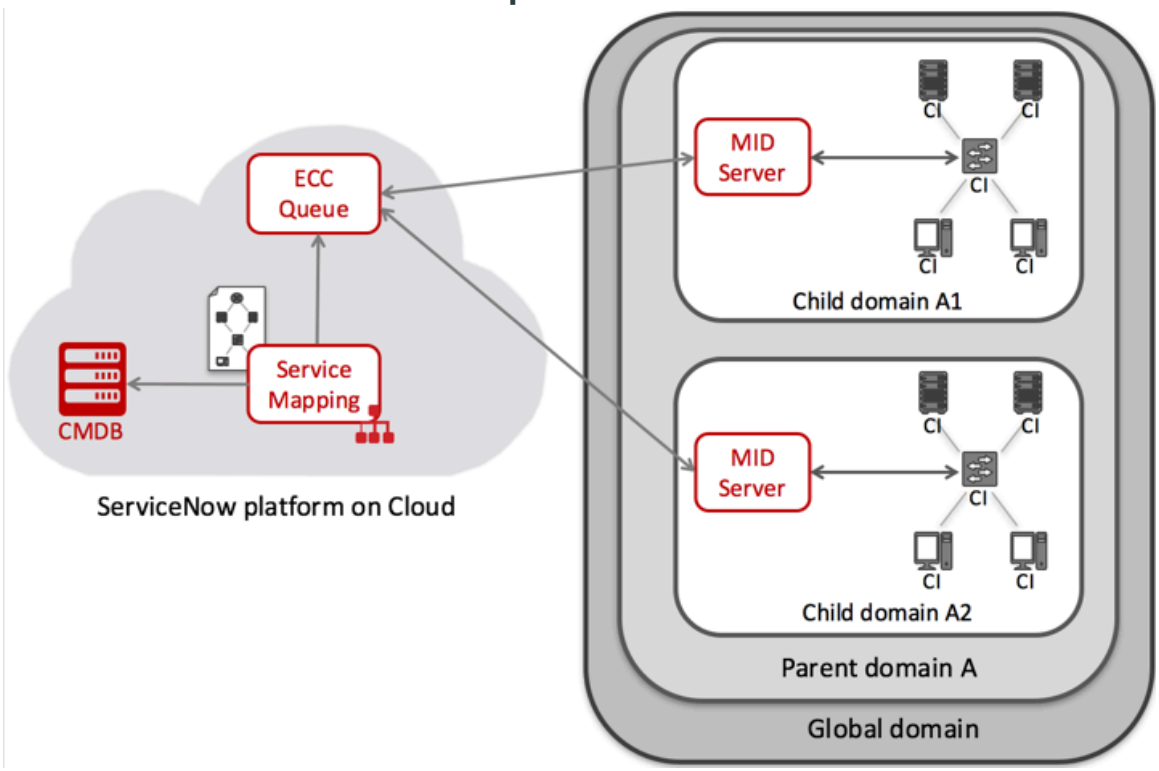

Once MID Servers MID Servers are installed, configure them to work with Service Mapping for the best discovery results.

Related topics

- Pattern-based discovery in Service Mapping
- MID Server connection prerequisites
- Credential-less discovery with Nmap

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credential-less Discovery with Nmap

Credential-less Application Discovery

Multiple MID Server deployments

MID Server installation

MID Server properties used by Service Mapping

Choose MID Server selection algorithm

Service Mapping supports the new and the legacy algorithms for selecting a MID Server for a discovery request. Depending on your organization needs, you can choose which algorithm to enable.

Before you begin

Role required: admin or service_mapping_admin

About this task

Service Mapping selects a MID Server using the following algorithm:

- Service Mapping chooses the MID Server whose selection criteria best match the parameters of the discovery request.
- If there are no MID Servers with matching selection criteria, Service Mapping chooses the default MID Server.
- If there are no MID Servers with matching selection criteria or default MID Server, Service Mapping cannot start the discovery process.

By default, Service Mapping uses this algorithm in all fresh installs and deployments upgraded from Istanbul or Jakarta.

In upgraded deployments, Service Mapping selects a MID Server using a legacy algorithm:

- Service Mapping chooses the MID Server whose application criteria is set to *`ServiceMapping`* or to *`ALL`*.
- By default, Service Mapping selects the MID Server whose IP range matches the IP in the discovery request.
- If Service Mapping cannot find a MID Server with the matching IP range and matching application, it assigns the discovery request to the default MID Server.
- If there is no default MID Server and none of the MID Servers have IP ranges configured for them, Service Mapping uses a random MID Server.
- If there is no default MID Server and no IP range matches the IP in the discovery request, Service Mapping cannot start the discovery process.

By default, Service Mapping uses the legacy algorithm in all deployments upgraded to Istanbul and earlier.

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Properties**.
2. To enable the new algorithm for deployments upgraded to Istanbul and earlier:
   a. Clear the **Enable Legacy MID selection algorithm** check box.
   b. Click **Save**.
   c. Configure MID Server selection criteria as described in MID Server configuration for Service Mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. To enable the legacy algorithm for the Brazil deployment:
   a. Select the **Enable Legacy MID selection algorithm** check box.
   b. Click **Save**.
   c. Configure MID Server selection criteria as described in MID Server configuration for Service Mapping in upgraded deployments.

MID Server configuration for Service Mapping in upgraded deployments

For ServiceNow deployments upgraded from earlier versions to Istanbul, Service Mapping uses the legacy algorithm to choose a MID Server for a discovery request.

MID Servers have the following selection criteria that Service Mapping in upgraded deployments uses to choose a MID Server for a discovery request:

- Application — defines what application a MID Server works with. Set it to Service Mapping to reserve this MID Server exclusively to Service Mapping discovery requests. Alternatively, set it to ALL to allow any ServiceNow application to use this MID Server.

  In earlier releases, you could configure MID Servers to work only with Service Mapping by using the *Capability* parameter. If you had MID Servers with this configuration, during the upgrade to Istanbul the application parameter is automatically set to *ServiceMapping* for them.

- IP range — limits operation of this MID Server to this IP range. Service Mapping does not choose this MID Server for a discovery request whose endpoint is outside this IP range.

In addition to the IP range, you can configure one of the MID Servers as the default server that Service Mapping uses. For operational information, see Configure a default MID Server for Service Mapping for upgraded deployments.

In upgraded deployments, Service Mapping selects a MID Server using a legacy algorithm:

- Service Mapping chooses the MID Server whose application criteria is set to *ServiceMapping* or to *ALL*.
- By default, Service Mapping selects the MID Server whose IP range matches the IP in the discovery request.
- If Service Mapping cannot find a MID Server with the matching IP range and matching application, it assigns the discovery request to the default MID Server.
- If there is no default MID Server and none of the MID Servers have IP ranges configured for them, Service Mapping uses a random MID Server.
- If there is no default MID Server and no IP range matches the IP in the discovery request, Service Mapping cannot start the discovery process.

The new algorithm for selecting MID Servers uses an extra selection criteria: capability. It allows Service Mapping to take a network protocol into consideration. For information about updated selection criteria and the new MID Server selection algorithm, see Choose MID Server selection algorithm.

While by default Service Mapping uses the legacy algorithm in upgraded deployments, it can support both new and legacy algorithms for selecting a MID Server. If necessary, you can enable the new MID Server selection algorithm for your upgraded deployment. For more information, see Choose MID Server selection algorithm.

Related topics

Configure an IP address range for the MID Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure a default MID Server for Service Mapping for upgraded deployments

In deployments upgraded to Istanbul or earlier, Service Mapping uses the default MID Server when it cannot find a MID Server with the matching IP range. Configuring a default MID Server improves the discovery process.

Before you begin

Role required: service_mapping_admin
Perform this procedure only for deployments upgraded to Istanbul or earlier. For fresh installs, perform Configure a default MID Server for each application.

Ensure that you know the name of the MID Server you want to configure as the default MID Server for Service Mapping.

About this task

In upgraded deployments, Service Mapping selects a MID Server using a legacy algorithm:

- Service Mapping chooses the MID Server whose application criteria is set to *`ServiceMapping`* or to *`ALL`*.
- By default, Service Mapping selects the MID Server whose IP range matches the IP in the discovery request.
- If Service Mapping cannot find a MID Server with the matching IP range and matching application, it assigns the discovery request to the default MID Server.
- If there is no default MID Server and none of the MID Servers have IP ranges configured for them, Service Mapping uses a random MID Server.
- If there is no default MID Server and no IP range matches the IP in the discovery request, Service Mapping cannot start the discovery process.

Procedure

1. Navigate to System Properties by entering `sys_properties.list` in the navigation filter.
2. Click **New**.
3. Configure values for the new property as follows:

Property	Description
Name	Enter `mid.server.sm_default`.
Description	Enter a free text description of the use of this MID Server.
Type	Select **string**.
Value	Enter the name of the MID Server you want to configure as the default for Service Mapping.

4. Click **Update**.

Related topics

MID Server configuration for Service Mapping

Activate AI Agents for Service Mapping

Activate the Service Mapping AI Agent and the Business App Mapping AI Agent from the Service Mapping home page to start automated service map creation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Before you begin, verify the following:

- ITOM AI agents for Service Mapping [sn_sm_gen_ai, version 1.3.1 or later] is installed as part of ServiceNow Otto for ITOM.
- Service Mapping Plus version 1.20.0 or later is installed. The AI Agent UI is delivered through Service Mapping Plus.

Role required: service_mapping_admin

About this task

The AI Agents are inactive by default. You can activate one or both AI Agents. After activation, the agents run automatically in the background without further input. You can deactivate an agent at any time.

Note: If the ITOM AI Agents for Service Mapping application is not installed, the **Support your team with an AI Agent** section does not appear on the home page. The Business App Mapping AI Agent is available through the Feature Preview Program. For more information, see Feature Preview Program.

Procedure

1. Navigate to **Workspaces** > **Service Mapping**.
2. On the Service Mapping home page, locate the **Support your team with an AI Agent** section and expand the section.
3. On the card for the AI Agent you want to activate, select **Activate**.
   The button changes to **Deactivate** and an **Active** badge appears on the card. The AI Agent begins running in the background.
4. *Repeat the previous step for the second AI Agent.*

   The activated AI Agents run automatically. You can monitor their activity in the list navigation under **AI Activity**.

What to do next

To deactivate an AI Agent, select **Deactivate** on the respective card. The button returns to **Activate**.

Activate the Service Mapping Candidate skill

The Service Mapping Candidate skill provides ServiceNow Otto the ability to classify app service candidates and generate a description for them. The skill is active by default. If needed, administrators can activate or deactivate the skill.

About this task

The Service Mapping agentic workflow, Analyze potential impact, uses two ServiceNow Otto skills:

- Service Mapping Candidate
- Service Mapping Candidates Impact

The Service Mapping Candidate skill provides names for processes and application service candidates. The skill gets details about the processes within the application service candidate, such as publisher, product, and description. Then, it gets details about the application service candidate, such as name, and description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Before activating the ServiceNow Otto skills, you must install the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

You must configure the following settings:

- ITOM pro plus SKU
- ServiceNow Otto for ITOM
- Service Mapping Plus (minimum version 1.16.3)
- ServiceNow Otto for ITSM
- ServiceNow Otto®

Important: This skill is now turned on by default. The skill will be automatically available to appropriate role users for the application. This change simply activates the skill and does not touch the roles that are needed to use the skill. The new default behavior works as follows:

New customers

When you install an AI product, designated skills will turn on automatically.

Existing customers who are upgrading

Any previously unconfigured skill will turn on automatically (the skill was never turned on, then off again).

There is no change to skills that are currently enabled and customized.

Previously configured skills that were turned on, then off, will remain inactive.

Role required: admin

Procedure

1. Navigate to **All** > *AI Admin Hub* > **Skills**.
2. On the navigation panel, select **ITOM**.
3. On the Service Mapping Candidate tile, select **Activate**.

Activate the Service Mapping Candidates Impact skill

The Service Mapping Candidates Impact skill analyzes connections and effects on servers. It enables generating an impact summary, using ServiceNow Otto. The skill is active by default. If needed, administrators can activate or deactivate the skill.

About this task

The Service Mapping agentic workflow, Analyze potential impact, uses two ServiceNow Otto skills:

- Service Mapping Candidate
- Service Mapping Candidates Impact

The Service Mapping Candidates Impact skill enables the agent to create a summary of impacted components, following a change request.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Before activating the ServiceNow Otto skills, you must install the ServiceNow Otto for IT Operations Management (ITOM) plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

You must configure the following setting:

- ITOM pro plus SKU
- ServiceNow Otto for ITOM
- Service Mapping Plus (minimum version 1.16.3)
- ServiceNow Otto for ITSM
- ServiceNow Otto®

Important: This skill is now turned on by default. The skill will be automatically available to appropriate role users for the application. This change simply activates the skill and does not touch the roles that are needed to use the skill. The new default behavior works as follows:

New customers

When you install an AI product, designated skills will turn on automatically.

Existing customers who are upgrading

Any previously unconfigured skill will turn on automatically (the skill was never turned on, then off again).

There is no change to skills that are currently enabled and customized.

Previously configured skills that were turned on, then off, will remain inactive.

Role required: admin

Procedure

1. Navigate to **All** > *AI Admin Hub* > **Skills**.
2. On the navigation panel, select **ITOM**.
3. On the Service Mapping Candidates Impact tile, select **Activate**.

Prerequisites for performing top-down discovery using Service Mapping

Learn about credentials, users, and user permissions you must configure to let Service Mapping access and discover applications inside your organization private network.

Configure the relevant requirements to enable successful top-down discovery:

Host credentials

Service Mapping uses credentials of the hosts to access applications running on this host. Typically, host credentials are enough for top-down discovery. For example, to discover the HP Service Manager Application Server, configure only credentials for accessing the server hosting this application, SSH credentials, or Windows credentials.

Applicative credentials

To discover some applications, Service Mapping needs separate credentials from credentials of the host on which such applications run. This type of credentials is referred to in ServiceNow as applicative credentials. For example, to discover the ABAP SAP Central Services (ASCS), configure applicative credential in addition to the host credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Elevated rights for Unix-based hosts

For Unix-based hosts, configure on MID Servers, you must provide a user with sudo-level credentials on all Unix-based hosts in your organization. For example, if the server hosting ASCS is Unix-based, provide a user with elevated rights to be able to run discovery commands against this Unix-based server.

SNMP community credentials

Configure SNMP community credentials to allow Service Mapping and Discovery to query network devices using the SNMP protocol. For example, configure SNMP community credentials for A10 Networks Thunder ADC.

Users or user permissions

You may need to provide additional users or user permissions to run certain commands or access certain directories. For example, for ASCS, provide a user with permissions to run SOAP on RFC read table function.

The ServiceNow AI Platform provides a centralized way of configuring credentials for Discovery, Service Mapping, and other applications. It allows you to configure credentials for hosts and applications only once. If you have already configured host-related credentials for another application, you do not need to do it again for Service Mapping.

You can assign credentials to specific MID Servers or keep credential generic so all MID Servers can use them. MID Servers retrieve commands and credentials from the ServiceNow instance and apply them to discover CIs inside organization private network.

Assigning credentials to the MID Server that uses them to access CIs

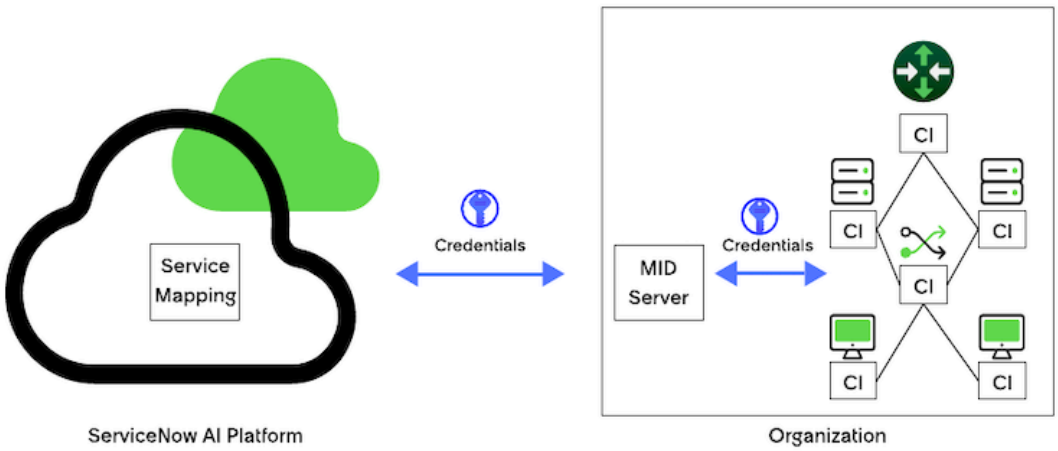

Use the following guidelines to decide for which MID Server to create credentials:

- If all CIs, belonging to the same CI type, share a credential, you do not need to specify a MID Server for it. In that case, it is used on all MID Servers by default.

  Note: Specifying the MID Server when configuring credentials is optional.

- If CIs of the same type and on the same network have different credentials, configure these credentials and define order in which Service Mapping uses them when trying to connect to these CIs.

In addition to generic credentials, you configure on MID Servers, you must configure sudo-level credentials on all Unix-based hosts in your organization.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the MID Server does not have sufficient credentials to access a device or application, it can run Network Mapper (Nmap) commands to collect basic information without using credentials. Credential-less discovery with Nmap requires additional configuration as described in Install and uninstall Nmap on a MID Server.

Review the list of credentials necessary to discover hosts and applications supported in the base system.

Application	CI type	Configure this
AIX OS	AIX Server [cmdb_ci_aix_server]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.

IBM z/OS	IBM zOS server Server [cmdb_ci_ibm_zos_server]	• Run the USS service on the port 22 on the z/OS server.
		• On the ServiceNow AI Platform, configure SSH credentials.
		• Provide the user with the following permissions:
		  ○ To write into the /tmp directory.
		  ○ To see all user processes.
		  ○ To run REXX scripts.

		On the ServiceNow AI Platform, configure SSH credentials.
		Provide the user with the following permissions: • To write into the /tmp directory. • To see all user processes. • To run REXX scripts.
HP-UX OS	HPUX Server [cmdb_ci_hpux_server]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		On the ServiceNow AI Platform, configure SSH credentials.
Linux OS	Linux Server [cmdb_ci_linux_server]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		On the ServiceNow AI Platform, configure SSH credentials.
Solaris OS	Solaris Server [cmdb_ci_solaris_server]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
Windows OS	Windows Server [cmdb_ci_windows_server]	On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		Provide all necessary credentials and permissions as described in Credentials required for host discovery.
A10 Load Balancing	A10 Load Balancer [cmdb_ci_lb_A10]	Configure applicative credentials.
		Configure SNMP credentials.
Windows Active Directory Domain Services	Active Directory Domain Controller [cmdb_ci_ad_controller]	On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
AWS DynamoDB	DynamoDB Global Table [cmdb_ci_dynamodb_global_table]	• On the AWS Console, provide a user with the following permissions: AmazonDynamoDBReadOnly Access and AWSLambdaInvocation-DynamoDB. • On the ServiceNow AI Platform®, configure AWS credentials, using a secret key and an access key.
Amazon Web Services (AWS) Elastic Load Balancing	Load Balancer Service [cmdb_ci_lb_service]	Configure Amazon Web Services credentials.
		On the ServiceNow AI Platform, create a service account. Set the Account ID to the Amazon ELB account ID as it appears in the AWS Management Console.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Amazon Web Services (AWS) API Gateway		Configure Amazon Web Services credentials.
		On the ServiceNow AI Platform, create a service account. Set the Account ID to the Amazon account ID to which API Gateway belongs. Use the Account ID as it appears in the AWS Management Console.
		Provide a user with permissions to run the apigateway:GET function for just the GET action in API Gateway.
		Ensure to place the MID Server on the virtual machine that is part of the service account that you are planning to discover.
Amazon Web Services (AWS) Elastic Compute Cloud (EC2)	Windows Server [cmdb_ci_windows_server] or Linux Server [cmdb_ci_linux_server]	(If hosted on a Unix server), on the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		On the ServiceNow AI Platform, create a service account. Set the Account ID to the Amazon account ID to which API Gateway belongs. Use the Account ID as it appears in the AWS Management Console.
		Ensure to place the MID Server inside the Amazon cloud.
Amazon Web Services (AWS) Relational Database Service (RDS)	Cloud DataBase [cmdb_ci_cloud_database]	Configure Amazon Web Services credentials.
		On the ServiceNow AI Platform, create a service account. Set the Account ID to the Amazon account ID to which RDS belongs. Use the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		Account ID as it appears in the AWS Management Console.
		Ensure to place the MID Server inside the Amazon cloud.
Amazon Web Services (AWS) Lambda	Cloud Function [cmdb_ci_cloud_function]	Configure Amazon Web Services credentials.
		On the ServiceNow AI Platform, create a service account. Set the Account ID to the Amazon account IT to which AWS Lambda belongs. Use the Account ID as it appears in the AWS Management Console.
		Ensure to place the MID Server on the virtual machine that is part of the service account that you are planning to discover.
		Provide a user with permissions to run the ListFunctions API with the Resource: arn:aws:lambda:region:account-id: function:function-name. Provide a user with permissions to run the GetFunction API with the Resource: '*'.
Apache Tomcat	Tomcat [cmdb_ci_app_server_tomcat]	Give the user elevated rights to execute Service Mapping commands requiring a privileged user.
		(If hosted on a Unix server), on the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Apache Tomcat WAR	Tomcat WAR [cmdb_ci_app_server_tomcat_war]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on UNIX), on the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
Apache Web Server	Apache Web Server [cmdb_ci_apache_web_server]	(If hosted on UNIX) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		On the server hosting the Apache Web Server, give the user permission to execute the following command: `$APACHEHOME/bin/httpd -V`.
		On the server hosting the Apache Web Server, verify that the $APACHEHOME/bin/envvar file is configured with the correct native library path.
		(If hosted on UNIX), on the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
App TNS Service	[cmdb_ci_appl_ora_tns]	Give the user elevated rights to run Service Mapping commands requiring a privileged user.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on Linux or UNIX), on the ServiceNow AI Platform, configure SSH credentials.

F5 BIG-IP Global Traffic Manager (GTM) and F5 BIG-IP Local Traffic Manager (LTM)	Load Balancer Service [cmdb_ci_lb_service]	• Verify that the following applications are up to date:
		  ○ Discovery and Service Mapping Patterns
		  ○ Visibility Content
		  ○ CMDB CI Class Models
		• To successfully discover the load balancer pool members, set the **display service names** option (*bigpipe.displayservicenames*) of the load balancer to **false** on the F5 load balancer.
		• For F5 Load Balancer pattern:
		  ○ Configure SNMP credentials on the ServiceNow AI Platform.
		  ○ (Optional) If there are iRules or SNMP community credentials aren't enough for discovering outgoing connections, configure SSH credentials.
		Note: If you don't want to use SSH credentials, you can use the REST API to create a connection to F5 BIG-IP devices.
		Service Mapping uses the SSH credentials to retrieve connections that aren't from CMDB. Discovering connections using the SSH protocol is a failover mechanism for the SNMP-based discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		• For F5 Load Balancer SSH pattern: ○ Configure SSH credentials. ○ Verify permissions for the following SSH commands. In newer versions of F5 BIG-IP, only A-type GTM pools are supported. ▪ `list auth partition all` ▪ `list sys global-settings` ▪ `show sys hardware` ▪ `list sys db failover.state` ▪ `list net self one-line` ▪ `show sys mac-address | grep interface` ▪ `show net interface field-fmt` ▪ `show net self | grep -e 'Mac Address' 'IP:'` ▪ `list ltm pool | grep -e 'pool' 'mode'` ▪ `list ltm pool /<partition_name>/* | grep -e 'pool' 'mode'` ▪ `show ltm pool | grep -e 'Ltm::' 'Total Connections'` ▪ `show ltm pool /<partition_name>/* |grep -e 'Ltm::' 'Total Connections'` ▪ `list sys db bigpipe.displayservicenames | grep value`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		▪ `show /net vlan | grep -e 'Interface Name' 'Mac' 'Tag'` ▪ `show ltm virtual detail | grep -e 'Server:' 'State' 'Pool:' 'Destination' 'Total Connections'` ▪ `show ltm virtual detail /<partition_name>/* | grep -e 'Server:' 'State' 'Pool:' 'Destination' 'Total Connections'` ▪ `show gtm wideip | grep -e 'WideIp' 'State'` ▪ `list wideip a /<partition_name>/* | grep -e 'WideIp' 'State'` ▪ `show gtm wideip all | grep -e 'WideIp' 'State'` ▪ `list gtm wideip` ▪ `list gtm wideip all` ▪ `list gtm pool one-line` ▪ `list gtm pool a one-line` ▪ `list gtm pool /<partition_name>/* one-line` ▪ `list gtm pool a /<partition_name>/* one-line` ▪ `list gtm server one-line` • For F5REST pattern:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		○ For horizontal discovery using Discovery, provide read-only permissions to run the following APIs: ▪ `"https://" + $ipAddress + "/mgmt/tm/ltm/pool"` ▪ `"https://" + $ipAddress + "/mgmt/tm/sys/global-settings"` ▪ `"https://" + $ipAddress + "/mgmt/tm/sys/hardware"` ▪ `"https://" + $ipAddress + "/mgmt/tm/sys/failover"` ▪ `"https://" + $ipAddress + "/mgmt/tm/net/self"` ▪ `"https://" + $ipAddress + "/mgmt/tm/ltm/virtual?$select=destination,mask,vla` ▪ `"https://" + $ipAddress + "/mgmt/tm/net/interface"` ▪ `"https://" + $ipAddress + "/mgmt/tm/net/interface/stats"` ▪ `"https://" + $ipAddress + "/mgmt/tm/net/trunk"` ▪ `"https://" + $ipAddress + "/mgmt/tm/net/vlan"` ▪ `"https://" + $ipAddress + "/mgmt/tm/net/`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		vlan/"+<vlan id>+"/interfaces" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip" ▪ "https://" + $ipAddress + "/mgmt/tm/cm/traffic-group/" ▪ "https://" + $ipAddress + "/mgmt/tm/cm/device" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/server" ▪ For BIG-IP starting with version 12: ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/a" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/aaaa" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/cname" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/mx" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/naptr" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/srv" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/a"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/aaaa" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/cname" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/mx" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/naptr" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/srv" ○ Create basic authentication credentials. • For F5 REST OAuth 2.0 authentication pattern: ○ Verify read-only permission for the following APIs: ▪ "https://" + $ipAddress + "/mgmt/shared/authn/login" ▪ "https://" + $ipAddress + "/mgmt/tm/ltm/pool" ▪ "https://" + $ipAddress + "/mgmt/tm/sys/global-settings" ▪ "https://" + $ipAddress + "/mgmt/tm/sys/hardware" ▪ "https://" + $ipAddress + "/mgmt/tm/sys/failover"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		▪ "https://" + $ipAddress + "/mgmt/tm/net/self" ▪ "https://" + $ipAddress + "/mgmt/tm/ltm/virtual?$select=destination,mask,vla ▪ "https://" + $ipAddress + "/mgmt/tm/net/interface" ▪ "https://" + $ipAddress + "/mgmt/tm/net/interface/stats" ▪ "https://" + $ipAddress + "/mgmt/tm/net/trunk" ▪ "https://" + $ip_address + "/mgmt/tm/net/vlan" ▪ "https://" + $ipAddress + "/mgmt/tm/net/vlan/"+<vlan id>+"/interfaces" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip" ▪ "https://" + $ipAddress + "/mgmt/tm/cm/traffic-group/" ▪ "https://" + $ipAddress + "/mgmt/tm/cm/device" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/server" ▪ For BIG-IP starting with version 12:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/a" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/aaaa" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/cname" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/mx" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/naptr" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/pool/srv" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/a" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/aaaa" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/cname" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/mx" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/naptr" ▪ "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/srv"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		◦ Create basic authentication credentials. ◦ Create a credential alias for the basic authentication credential. ◦ Create a serverless discovery schedule. • For top-down discovery using Service Mapping, provide permissions to run the following APIs: ◦ `https://"get_attr{"managementIP"}"/mgmt/tm/ltm/"get_attr{"irule"}` ◦ `https://"get_attr{"managementIP"}"/mgmt/tm/ltm/virtual/~Common~"get_attr{"vip_name"}` ◦ `https://"get_attr{"computer_system.managementI mgmt/tm/ltm/pool/"get_attr{"pool_name"}"/members?` • (Optional) For top-down discovery with Service Mapping, configure applicative credentials. • (Optional) For top-down discovery with Service Mapping, provide a user with permissions to run the following commands: ◦ bigpipe commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 9) ◦ bigpipe and Traffic Management Shell (TMSH) commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 10)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		○ Traffic Management Shell (TMSH) commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 11) ○ Traffic Management Shell (TMSH) advanced commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 10, 11, and 12) ○ The *`show cm traffic-group -get command`* for discovering F5 BIG-IP Device Service Clustering • (Optional) For top-down discovery with Service Mapping, give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
BMC CTRL-M Enterprise Manager	Control-M [cmdb_ci_appl_controlm]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on a Linux or UNIX server) (Optional) on the ServiceNow AI Platform, configure SSH credentials.
BMC CTRL-M Gateway	Control-M [cmdb_ci_appl_controlm]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on a Linux or UNIX server) (Optional) on the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		ServiceNow AI Platform, configure cc.
BMC IT Asset Management (ITAM)	ITAM Asset Center [cmdb_ci_appl_itam]	On the ServiceNow AI Platform, configure a Windows OS user with permissions to access the ITAM registry as described in Windows credentials.
CA Identity Manager Provisioning Server	CA Identity Manager Provisioning Server [cmdb_ci_appl_ca_id_man]	On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Cadebill Interconnect	inter connect [cmdb_ci_inter_connect]	On the ServiceNow AI Platform, configure Windows credentials.
Cadebill Interconnect Instance	Interconnect Instance [cmdb_ci_interconnect_instance]	On the ServiceNow AI Platform, configure Windows credentials.
Cisco Application Control Engine ACE (on Cisco CSM)	ACE [cmdb_ci_lb_ace]	On the ServiceNow AI Platform, configure applicative credentials.
		On the ServiceNow AI Platform, configure SSH credentials.
		On the ServiceNow AI Platform, configure SNMP community credentials as described in SNMP credentials.
Cisco CallManager	Cisco CallManager [cmdb_ci_appl_cisco_call_man]	(If hosted on a Linux or UNIX server) (Optional) on the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
Cisco Global Site Selector (GSS)	Load Balancer Service [cmdb_ci_lb_service]	Configure SNMP community credentials as described in SNMP credentials.
		On the ServiceNow AI Platform, configure SSH credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		Provide a user who has permissions to run query on the domain list, DNS rule output, and VIP answer to get the outgoing connections of GSS.
Citrix Delivery Controller	Delivery Controller [cmdb_ci_appl_delivery_controler]	• Configure Windows administrative credentials to access the host machine. • Make sure that PowerShell for Discovery is installed on your MID Servers. The patterns used for Citrix discovery require the ability to run PowerShell commands. • For OS users, make sure you can execute "Get-XDSite." The Citrix Studio Admin should provide permissions. • (Optional) For top-down discovery with Service Mapping, on Citrix Delivery Controller, install the Citrix.Broker module. • (Optional) For top-down discovery with Service Mapping, a user with permission to run PowerShell commands. • (Optional) (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Citrix Usage Collector	Citrix Collector [cmdb_ci_appl_citrix_collector]	• Configure applicative credentials. • Configure a user with permissions to run PowerShell commands on the Citrix server. • (If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this

		described in Windows credentials.
		
		Note: To use this credential later for XenAPP, set the **Credential alias** attribute to cmdb_ci_appl_citrix_xenapp. To use this credential later for Citrix, set the **Credential alias** attribute to cmdb_ci_appl_citrix_collector.
		
		• To discover Epic, provide permission to run a `farm` command.
		• Provide a user with permissions to read and query the Citrix repository:
		  ○ For versions 4.5 and later, a PowerShell script from the XenApp server.
		  ○ For versions earlier than 4.5, a Microsoft Visual Basic Script (VBScript) from the Citrix Presentation Server
		• (If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		
		Note: To use this credential later for XenAPP, set the *`Credential alias`* attribute to cmdb_ci_appl_citrix_xenapp. To use this credential later for Citrix, set the *`Credential alias`* attribute to cmdb_ci_appl_citrix_collector.

Citrix Presentation Server/ Citrix XenApp	Citrix XenAPP or Presentation Server [cmdb_ci_appl_citrix_xenapp]	• On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Note: Citrix XenApp is the new name for Presentation Server beginning with the XenApp 5 release.		described in Windows credentials. • (Optional) To discover using the WinRM protocol, see Configure WinRM trusted hosts on MID Server
Citrix Netscaler	Citrix Netscaler [cmdb_ci_lb_netscaler]	• Configure SNMP credentials. The SNMP community string requires read-only access to the Citrix NetScaler MIB (`1.3.6.1.4.1.5951`). **Note:** You also have the option of using SSH credentials. The SSH user requires permission to run NetScaler `show` commands only. However, if the payload size of the data that's returned to the instance is large, the discovery process can fail. In this case, use SNMP credentials instead. • (Optional) When using the NS.CONF pattern, provide a user with read-only access to the *NS.conf* configuration file if you don't want to provide SNMP or SSH credentials. • Use basic authentication credentials for discovery via REST. Configure the credential as a credential alias. The NITRO API user requires read-only access to the following resources: ◦ `lbvserver` ◦ `csvserver` ◦ `service` ◦ `servicegroup` ◦ `interface` ◦ `lbvserver_binding`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		◦ `csvserver_binding` ◦ `nspartition` ◦ `nshardware` ◦ `nshostname` ◦ `nsconfig` ◦ `hanode` Service Mapping uses the SSH credentials to retrieve connections that aren't from CMDB. Discovering connections using the SSH protocol is a failover mechanism for the SNMP-based discovery.
Citrix Netscaler Global Server Load Balancer	Load Balancer Service [cmdb_ci_lb_service]	• Configure SNMP credentials. The SNMP community string requires read-only access to the Citrix NetScaler MIB (`1.3.6.1.4.1.5951`). **Note:** You also have the option of using SSH credentials. The SSH user requires permission to run NetScaler `show` commands only. However, if the payload size of the data that's returned to the instance is large, the discovery process can fail. In this case, use SNMP credentials instead. • (Optional) When using the NS.CONF pattern, provide a user with read-only access to the *`NS.conf`* configuration file if you don't want to provide SNMP or SSH credentials. • Use basic authentication credentials for discovery via REST. Configure the credential as a credential alias. The NITRO API user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		requires read-only access to the following resources: ◦ `lbvserver` ◦ `csvserver` ◦ `service` ◦ `servicegroup` ◦ `interface` ◦ `lbvserver_binding` ◦ `csvserver_binding` ◦ `nspartition` ◦ `nshardware` ◦ `nshostname` ◦ `nsconfig` ◦ `hanode`
Connect APK	Connect-It Service [cmdb_ci_appl_connectit]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on a Linux or Unix server) on the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
EMC Documentum	Documentum DocBase [cmdb_ci_appl_doc_docbroker]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Linux or Unix server) (Optional) on the ServiceNow AI Platform, configure SSH credentials.
Epic	epicd app server [cmdb_ci_epicd_app_server] epic agent [cmdb_ci_epic_agent] epic cache [cmdb_ci_epic_cache]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Linux or Unix server) (Optional) on the ServiceNow AI Platform, configure SSH credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
	epicd [cmdb_ci_endpoint_epicd]	
FormEngine	Oracle Forms Engine [cmdb_ci_appl_ora_forms]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Unix server) on the ServiceNow AI Platform, configure SSH credentials.
Google Cloud Platform	Cloud Service Account [cmdb_ci_cloud_service_account]	• Provide a read-only user with permission to use the following API elements. • On the ServiceNow AI Platform, configure Google API credentials.
HAProxy	HAProxy Load Balancer [cmdb_ci_lb_haproxy]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Linux or Unix server) (Optional) on the ServiceNow AI Platform, configure SSH credentials.
HP Operations Manager	HP Operations Manager [cmdb_ci_appl_hp_operations]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
HP Service Manager Application Server	HP Service Manager [cmdb_ci_appl_hp_service]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Linux or Unix server) (Optional) on the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
HP Service Manager Index	HP SM Index Server [cmdb_ci_appl_hp_index]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Linux or Unix server) On the ServiceNow AI Platform, configure SSH credentials.
HP Service Manager Knowledge Base	HP SM KnowledgeBase [cmdb_ci_appl_hp_sm_kb]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
HP Quality Center	HP Quality Center [cmdb_ci_appl_hp_qc]	On the ServiceNow AI Platform, configure Windows credentials.
HP UCMDB	HP uCMDB [cmdb_ci_app_server_hp_ucmdb]	On the ServiceNow AI Platform, configure Windows credentials.
IBM Cloud Platform	Cloud Service Account [cmdb_ci_cloud_service_account]	On the ServiceNow AI Platform, configure the IBM Credentials for accessing the IBM account.
IBM WebSphere Message Broker	IBM WebSphere Message Broker [cmdb_ci_appl_ibm_wmb]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Unix server) On the ServiceNow AI Platform, configure SSH credentials for a user with permissions to run the following commands: • "source " + $directory + "/mqsiprofile" • "echo $MQSI_VERSION " • $directory + "/mqsireportbroker " + $name • $directory + "/mqsibrowse " + $name + " -t BROKERRESOURCES " Set the *Credential alias* attribute to [cmdb_ci_appl_ibm_wmb]. Configuring this attribute

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		allows you later to use the credential for this CI in patterns.
IBM CTRL-M Server	Control-M [cmdb_ci_appl_controlm]	(If hosted on Unix) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials.
IBM CICS Transaction Gateway CTG	IBM CTG [cmdb_ci_appl_ibm_ctg]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
IBM Customer Information Control System (CICS)	IBM CICS [cmdb_ci_appl_ibm_cics]	(If hosted on Unix) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
IBM DB2	DB2 Instance [cmdb_ci_db_db2_instance] **Note:** Starting with CMDB CI Class Models version 1.69.0, the label has been renamed to **DB2 Subsystem**.	• (If hosted on Linux) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user. • (If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials. Set the tag attribute to [cmdb_ci_db_db2_instance]. Configuring this attribute allows you later to use the credential for this CI in patterns. (Optional) To discover connections to IBM DB2, configure SSH credentials for a user with permissions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		to run the following commands: ○ `db2 list tablespaces` ○ "db2cmd /c /w /i db2 -tvf "+$ script_name • (If hosted on z/OS) Provide the user with the following permissions: ○ To write into the /tmp directory. ○ To see all user processes. ○ To run REXX scripts. ○ To read into the db2 folder. • (If hosted on z/OS) On the ServiceNow AI Platform, configure SSH credentials.
IBM J2EE EAR	[cmdb_ci_app_server_ws_ear]	(If hosted on Linux only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials.
IBM PowerHA SystemMirror for AIX discovery	• Cluster Resource Group [cmdb_ci_cluster_resource_group] • UNIX Cluster Resource [cmdb_ci_unix_cluster_resource] • UNIX Cluster Node [cmdb_ci_unix_cluster_node] • AIX Server [cmdb_ci_aix_server] • UNIX Cluster [cmdb_ci_unix_cluster] • Cluster Virtual IP [cmdb_ci_cluster_vip]	Configure a UNIX user with permissions to run all commands defined in the official IBM documentation.
IBM Tivoli Access Manager WebSEAL or IBM Security	Webseal [cmdb_ci_app_server_webseal]	(For ISAM appliance only) Configure applicative

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Access Manager (ISAM) appliance		credentials for the user with permissions to run the following REST commands: • /reverseproxy – for retrieving data about the reverse proxy and instances. • /wga/reverseproxy/"+ ${junctionstable[].instance}+"/ junctions – for getting junctions per each instance. • /wga/reverseproxy/" + $instance_name + "/junctions?junctions_id=" + $junction_name – for getting forwarding URLs for the entry point junction.
		(For ISAM appliance only) Ensure that the system can use the cURL commands for discovery: • On the MID Server used for ISAM appliance discovery, define the PATH variable to contain the cURL package installation for the relevant MID Server users. • Configure permissions to run the `curl.exe` file for the relevant MID Server user.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
IBM Virtualization and Hardware Management Console (HMC)	• HMC server • IBM frame • Logical partition (LPAR) instance An LPAR is virtual machine equivalent of IBM. • LPAR resource	• Verify that the following applications are up to date: ○ Discovery and Service Mapping Patterns ○ CMDB CI Class Models • Verify you have permissions for the following HMC commands: ○ `lssyscfg` ○ `lshmc` ○ `lshwres` ○ `lssysconn -r all -F type_model_serial_num:ipaddr` • Verify you have permissions for the following AIX (LPAR) command: `lparstat -i`. **Note:** The `lparstat -i` command is in addition to the commands required for AIX server discovery. For more information on AIX server discovery, see AIX server discovery. For a list of the AIX commands, see Service Mapping commands requiring a privileged user. • Configure SSH credentials on the ServiceNow AI Platform. For more information, see SSH credentials.
IBM WebSphere Application Server	IBM Websphere [cmdb_ci_app_server_websphere]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user. (If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		(If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on z/OS) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on z/OS) Provide the user with the following permissions: • To write into the /tmp directory. • To see all user processes. • To run REXX scripts. • To run commands using oeconsol. • To run the command that brings info on the WebSphere Application Server task: `"oeconsol 'D OMVS,A=ALL' | grep "+$taskname`. • To read content of the WAS installation folder and the permission to read was.env.file where the installation variables are situated.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
IBM WebSphere DataPower	Data Power [cmdb_ci_app_server_datapower]	**SNMP credentials** SNMP credentials are necessary for using the following SNMP queries: • For horizontal discovery of the Data Power Hosting Server [cmdb_ci_datapower_server]: ○ 1.3.6.1.2.1.1.1 ○ 1.3.6.1.2.1.1.5

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		○ 1.3.6.1.2.1.4.20.1 ○ 1.3.6.1.2.1.2.2.1 ○ 1.3.6.1.4.1.14685.3.1.196.1 ○ 1.3.6.1.4.1.14685.3.1.11.1 ○ 1.3.6.1.4.1.14685.3.1.112.1 • For top-down discovery of the Data Power [cmdb_ci_datapower_server]: ○ 1.3.6.1.2.1.1.1 ○ 1.3.6.1.2.1.1.5 ○ 1.3.6.1.4.1.14685.3.1.11.1 **Applicative credentials** Together with SNMP credentials, applicative credentials are required to discover DataPower application domains, otherwise known as logical CIs. In addition, SNMP and applicative credentials are used for top-down discovery of applicative relationships (corresponding application DataPower domain). **User with permissions to run SOAP** Permissions to run SOAP queries are necessary for top-down discovery of the DataPower Web Service and MQ connection.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
IBM WebSphere Message Broker Flow	WMB Flow [cmdb_ci_appl_wmb]	(If hosted on a Unix server) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		Provide an IBM WebSphere Message Broker OS user with permissions to run the WebSphere Message Broker service.
		(If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials for a user with permissions to run the following commands: • "source " + $directory + "/mqsiprofile” • "echo $MQSI_VERSION “ • $directory + "/mqsireportbroker " + $name • $directory + "/mqsibrowse " + $name + " -t BROKERRESOURCES “ Set the *Credential alias* attribute to [cmdb_ci_appl_ibm_wmb]. Configuring this attribute allows you later to use the credential for this CI in patterns.
IBM WebSphere Message Broker	IBM WebSphere Message Broker [cmdb_ci_appl_ibm_wmb]	(If hosted on Unix) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials. Set the *Credential alias* attribute to [cmdb_ci_appl_ibm_wmb]. Configuring this attribute allows you later to use the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		credential for this CI in patterns.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		Provide an IBM WebSphere Message Broker OS user with permissions to run the WebSphere Message Broker service.
IBM MQ (formerly IBM WebSphere MQ)	IBM MQ Manager (formerly IBM WebSphere MQ) [cmdb_ci_appl_ibm_wmq]	(If hosted on Unix) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		Configure applicative credentials for a user with permissions to run the following commands: • (If hosted on a Windows server) "echo display QMSTATUS ALL | runmqsc -e " + $queue_manager – for getting the queue manager (QM) status. • (If hosted on a Windows server) "echo display QMGR ALL | runmqsc -e " + $queue_manager – for getting the complete information about QM. • (If hosted on a Unix server) "echo \" display QMSTATUS ALL \" | runmqsc " + $queue_manager – for getting the QM status. • (If hosted on a Unix server) "echo \" display QMGR ALL \" | runmqsc " + $queue_manager – for getting the complete information about QM. • (If hosted on a Unix server) "echo \" display queue(" + $entry_point.queue + ") \" | runmqsc " +

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		$queue_manager – for getting the information about the EP queue.
		(If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) configure SSH credentials. While configuring OS credentials, set the *Credential alias* attribute to cmdb_ci_appl_ibm_wmq. Configuring this attribute allows you later to use the credential for this CI in patterns. (Optional) To find IBM MQ connections, configure SSH credentials for a user with permissions to run the following commands: • "echo \" display queue("+ $entry_point.queue +") \" | runmqsc "+ $queue_manager • "echo \" display QMSTATUS ALL \" | runmqsc "+ $queue_manager • echo \" display QMGR ALL \" | runmqsc "+ $queue_manager
		(If hosted on z/OS) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on z/OS) Provide the user with the following permissions: • To write into the /tmp directory. • To see all user processes. • To run REXX scripts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		• To run commands using oeconsol. • To run MQ commands.
IBM MQ Queue (formerly IBM WebSphere MQ Queue)	IBM MQ Queue (formerly IBM WebSphere MQ Queue) [cmdb_ci_appl_ibm_wmq_queue]	(If hosted on a Unix server) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		Configure applicative credentials for a user with permissions to run the following commands: • dspmq – for getting information about the queue manager. • "echo \" display queue(" + $entry_point.queue + ") \" | runmqsc " + $queue_manager – for getting the queue information. • "echo \" display clusqmgr(*) conname \" | runmqsc " + $queue_manager – for getting all cluster queue manager (QM) lines. • "echo \" display qstatus(" + $entry_point.queue + ") TYPE(HANDLE) OPENTYPE(INPUT) CONNAME PID \" | runmqsc " + $queue_manager – for getting the queue status information.
		(If hosted on Unix server) on the ServiceNow AI Platform, configure SSH credentials for a user with permissions to run the following commands: • dspmq echo \" display queue("+ $entry_point.queue +") \" | runmqsc "+ $queue_manager • (Optional) "echo \" display qcluster("+$rqmname

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		+") \" | runmqsc "+ $queue_manager • (Optional) "echo \" dischl(*) CHLTYPE(SDR) CONNAME XMITQ WHERE(XMITQ EQ "+ $xmit_name+")\" \| runmqsc "+$queue_manager • (Optional) "echo \" display clusqmgr(*) conname \" | runmqsc "+ $queue_manager • (Optional) "echo \" display qstatus("+ $entry_point.queue +") TYPE(HANDLE) OPENTYPE(INPUT) CONNAME PID \" \| runmqsc "+$queue_manager While configuring OS credentials, set the *Credential alias* attribute to cmdb_ci_appl_ibm_wmq. Configuring this attribute allows you later to use the credential for this CI in patterns.
		(If hosted on Windows), on the ServiceNow AI Platform, configure Windows credentials for a user with permissions to run the following commands: • dspmq echo \" display queue("+ $entry_point.queue +") \" | runmqsc "+ $queue_manager • (Optional) "echo \" display qcluster("+$rqmname +") \" | runmqsc "+ $queue_manager • (Optional) "echo \" dischl(*) CHLTYPE(SDR) CONNAME XMITQ WHERE(XMITQ EQ "+ $xmit_name+")\" \| runmqsc "+$queue_manager

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		• (Optional) "echo \" display clusqmgr(*) conname \" | runmqsc "+ $queue_manager • (Optional) "echo \" display qstatus("+ $entry_point.queue +") TYPE(HANDLE) OPENTYPE(INPUT) CONNAME PID \" \| runmqsc "+$queue_manager While configuring OS credentials, set the *Credential alias* attribute to cmdb_ci_appl_ibm_wmq. Configuring this attribute allows you later to use the credential for this CI in patterns..
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
IBM WebSphere Portal	Websphere Portal [cmdb_ci_appl_websphere_portal]	(If hosted on Linux only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Interconnect Web Service	Inter connect [cmdb_ci_inter_connect]	(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Jboss	JBoss [cmdb_ci_app_server_jboss]	Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Unix server) on the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Jboss Module	Jboss module [cmdb_ci_app_server_jb_module]	(If hosted on a Unix server) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
JRun	Jrun [cmdb_ci_app_server_jrun]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
JRun WAR Inc	Jrun WAR [cmdb_ci_app_server_jrun_war]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
Microsoft Azure Web Sites	Cloud WebServer [cmdb_ci_cloud_webserver]	On the ServiceNow AI Platform, add an Azure service principal and set up a cloud account for Azure. Set the account Id to the Azure Service Principal as it appears at the Microsoft Azure Console.
Microsoft Azure Load Balancer	Load Balancer Service [cmdb_ci_lb_service]	On the ServiceNow AI Platform, configure add an Azure service principal and set up a cloud account for Azure. Set the account Id to the Azure Service Principal as it appears at the Microsoft Azure Console.
Microsoft Azure SQL Database	Cloud Database Rule [cmdb_ci_cloud_database]	On the ServiceNow AI Platform, add an Azure service principal and set up a cloud account for Azure. Set the account Id to the Azure Service Principal as it appears at the Microsoft Azure Console.
Microsoft Dynamic CRM	Dynamic CRM Component [cmdb_ci_appl_ms_dynamic_crm]	On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Microsoft BizTalk Orchestration	BizTalk Orchestration [cmdb_ci_appl_biztalk_orch]	On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Microsoft BizTalk Server	BizTalk [cmdb_ci_appl_biztalk]	On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
Microsoft Exchange BackEnd Server	ExchangeBackEndServer [cmdb_ci_exchange_backend]	On the ServiceNow AI Platform, configure Windows credentials.
		Configure credentials required for host discovery for a Windows server.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Microsoft Exchange FrontEnd Server	ExchangeFrontEndServer [cmdb_ci_exchange_frontend]	On the ServiceNow AI Platform, configure Windows credentials.
		Configure credentials required for host discovery for a Windows server.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Microsoft FAST Search Server	Fast Search [cmdb_ci_appl_fastsearch]	Create a user with the permission to read the following files: • `install_directory +"\etc\VERSION.xml` • and `install_directory +"\etc\CSConfig.xml`
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Microsoft Exchange Hub Transport Server Microsoft Exchange Mailbox Microsoft Exchange CAS	ExchangeHub [cmdb_ci_exchange_hub] Exchange MailBox [cmdb_ci_exchange_mailbox] Exchange Client Access Server [cmdb_ci_exchange_cas]	• On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials. Use the Microsoft Exchange Server administrator's user name and password as Windows credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		For example, user name: local\administrator and password: Are1300. **Note:** If you configure Windows credentials for a user different from the Microsoft Exchange Server administrator, Service Mapping uses traffic-based discovery that results in mapping all Microsoft Exchange clients. • (For Exchange CAS 2010 and 2013, 2016) On the ServiceNow AI Platform, configure applicative credentials, setting the CI type to Exchange Mailbox. **Warning:** Do not confuse the Exchange Mailbox CI type with the Exchange Mailbox Server CI type. Service Mapping uses these applicative credentials to access the of/Servers/Servers.slab folder on the Key Exchange Admin Center (EAC). Service Mapping also uses the applicative credentials to run the following PowerShell commands against the MID Servers: ○ Get-ExchangeServer ○ Get-ExchangeServer ○ Get-MailboxServer ○ Get-ClusteredMailBoxServerStatus ○ Get-ExchangeServer ○ Get-StorageGroup ○ Get-MailboxDatabase • (For Microsoft Exchange CAS 2010 and 2016) On the Windows Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		hosting Microsoft Exchange CAS, install the Microsoft.Exchange.Management.PowerShel module. • (For Microsoft Exchange CAS 2007) On the Windows Server hosting Microsoft Exchange CAS, install the Microsoft.Exchange.Management.PowerShel module. • (For Microsoft Exchange CAS 2007, 2010) On the ServiceNow AI Platform, upload the rctrlx tool for running PowerShell commands as described in Upload the rctrlx.exe file to MID Servers. • (For Microsoft Exchange CAS 2007 and 2010) Provide an Exchange Mailbox OS user with the rights to run Exchange services on Windows and has permission to query Exchange repository: On the ServiceNow AI Platform, configure Windows credentials setting the *Credential alias* attribute to [cmdb_ci_exchange_mailbox]. Configuring this attribute allows you later to use the credential for this CI in patterns. • (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server. • Verify that you know the entry point necessary for discovering the Microsoft Exchange Server. Use the URL for the Outlook web access in the following format:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		`http(s)://<hostname>:<port>/owa` For example, `http(s)://myhost:<443>/owa`
Microsoft Identity Integration Server	IIFP [cmdb_ci_directory_iifp]	On the ServiceNow AI Platform, configure Windows credentials.
Microsoft Internet Information Services (IIS)	Microsoft iis Web Server [cmdb_ci_microsoft_iis_web_server]	On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Microsoft Internet Information Services (IIS) Virtual Directory	IIS Virtual Directory [cmdb_ci_iisdirectory]	On the ServiceNow AI Platform, configure the following credentials: • Windows credentials • applicative credentials
		Verify read access to the following files: • `web.config`: Possible locations `$entry_point.path + "Web.config"` or `$entry_point.path + "\Web.config"` • `machine.config`: possible location `"dir \"C:\Windows\Microsoft.NET\"`
		• (Optional) To find ADO.NET connections, set the *`Credential alias`* attribute to [cmdb_ci_iisdirectory]. Configuring this attribute allows you later to use the credential for this CI in patterns. • (Optional) For discovering using the WinRM protocol,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		see Configure WinRM trusted hosts on MID Server.
Microsoft Message Queuing (MSMQ)	MSMQ [cmdb_ci_appl_msmq]	On the ServiceNow AI Platform, configure Windows credentials.
Microsoft .NET Framework	.NET Application [cmdb_ci_appl_dot_net]	On the ServiceNow AI Platform, configure Windows credentials. (Optional) To find ADO.NET connections, set the *Credential alias* attribute to [cmdb_ci_iisdirectory]. Configuring this attribute allows you later to use the credential for this CI in patterns.
		Provide an IIS Virtual Directory OS user with the rights to run the IIS service on Windows.
		Configure applicative credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Microsoft SharePoint	SharePoint [cmdb_ci_appl_sharepoint]	On the ServiceNow AI Platform, configure the following credentials: • Windows credentials • Set the *Credential alias* attribute to [cmdb_ci_appl_sharepoint]. Configuring this attribute allows you later to use the credential for this CI in patterns. • Configure WMI Query access to the remote Sharepoint server. • (Optional) To find Windows SharePoint connections, configure Windows credentials for a user with

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		permissions to run http get request to the SharePoint Central Administration URL.
		Provide a SharePoint OS user with permissions: • To run SharePoint services on Windows • To log into the SharePoint Admin page
		To use two separate credentials with different permissions, configure credentials as covered in Prerequisites for performing top-down discovery / Service Mapping discovery to Sharepoint system [KB0813083] article in the Now Support Knowledge Base.
Microsoft SQL Database	MySQLClusterDataNode [cmdb_ci_db_mysql_clusternode]	Provide the following user or users: • For discovering Microsoft Biztalk using SQL queries, provide a user and credentials for Microsoft BizTalk. • For discovering licensing information on Microsoft SQL Server, provide a user and credentials for accessing Microsoft SQL Database. Refer to the product documentation for operational information.
		Credentials for host discovery as described in Windows credentials.
		Configure applicative credentials.
MongoDB	Mongo Config Server [cmdb_ci_appl_mongo_config_server]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		Provide a user with permissions to run Mongo commands described in Service Mapping commands not requiring a privileged user.
MySQL Cluster MGM Node	MySQLClusterMGMNode [cmdb_ci_db_mysql_clustermgmtnode]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		On the ServiceNow AI Platform, configure Windows credentials.
Nginx	Nginx Web Server [cmdb_ci_nginx_web_server]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
OpenText Documentum	Documentum DocBase [cmdb_ci_appl_doc_docbase]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
Oracle Application Server	Oracle iAS [cmdb_ci_app_server_ora_ias]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on Unix) Configure applicative credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Oracle Advanced Queue Queue	Advanced Queue Queue [cmdb_ci_appl_ora_queue]	(If hosted on Unix) Configure applicative credentials for Service Mapping.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
Oracle Clusterware	UNIX Cluster [cmdb_ci_unix_cluster], UNIX Cluster Node [cmdb_ci_unix_cluster_node], UNIX Cluster Resource [cmdb_ci_unix_cluster_resource], Cluster Virtual IPs [cmdb_ci_cluster_vip]	See Oracle database discovery.
Oracle Concurrent Server	Oracle Concurrent Server [cmdb_ci_appl_ora_conc]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle Database	Oracle Instance [cmdb_ci_db_ora_instance]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		(If hosted on Unix) Configure applicative credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

Oracle Database 12c	Oracle PDB Instance [cmdb_ci_db_ora_pdb_instance]	• Configure Applicative credentials or Oracle Wallet authentication
		• For Windows, configure Windows credentials
		• For Unix, configure SSH credentials
		• For the Oracle database, configure Oracle database discovery
		• Provide a user with permissions to run queries on the following tables:
		  ○ V$DATAFILE
		  ○ V$TEMPFILE
		  ○ V$LOG
		  ○ DBA_FREE_SPACE
		  ○ V$DATABASE
		  ○ V$PDBS
		  ○ V$PARAMETER

Oracle Database Advanced Queuing	Advanced Queue Queue [cmdb_ci_appl_ora_queue]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		Provide an Oracle instance user with permission to run the sqlplus command as described in Service Mapping commands not requiring a privileged user.
Oracle Discoverer Engine	Oracle Discoverer Engine [cmdb_ci_appl_ora_disc]	(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle Discoverer UI	Oracle Discoverer UI [cmdb_ci_appl_ora_disc_ui]	(If hosted on Unix) Configure applicative credentials for Service Mapping.
		On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on Unix only) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle E-Business Suite	Oracle ESB [cmdb_ci_appl_ora_ebs]	(If hosted on Unix) Configure applicative credentials.
		On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		see Configure WinRM trusted hosts on MID Server.
Oracle Form UI	Oracle Forms UI [cmdb_ci_appl_ora_forms_ui]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle Fulfillment Server	Oracle Fulfillment Server [cmdb_ci_appl_ora_fs]	(If hosted on Unix) Configure applicative credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle GlassFish Server	GlassFish [cmdb_ci_appl_glassfish] GlassFish WAR [cmdb_ci_appl_glassfish_war]	• (If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials. • (If hosted on a Linux or UNIX server) (Optional) on the ServiceNow AI Platform, configure SSH credentials. • (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle HTTP Server	Oracle HTTP Server [cmdb_ci_appl_ora_http]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		see Configure WinRM trusted hosts on MID Server.
Oracle iAS Web Module	Oracle iAS Web module [cmdb_ci_app_server_ora_ias_m]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
Oracle Metric Client	Oracle Metric Client [cmdb_ci_appl_ora_metric_client]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
Oracle Metric Server	Oracle Metric Server [cmdb_ci_appl_ora_metric_svr]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle MySQL Server	MySQL Instance [cmdb_ci_db_mysql_instance]	Provide the following user or users: • For discovering Microsoft Biztalk using SQL queries, provide a user and credentials for Microsoft BizTalk. • For discovering licensing information on Microsoft SQL Server, provide a user and credentials for accessing Microsoft SQL Database.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		Configure applicative credentials.
		(If hosted on Linux) Give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(If hosted on a Linux server) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle Net Listener	Oracle Database Listener [cmdb_ci_db_ora_listener]	(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
Oracle Notification Server	Oracle Notification Server [cmdb_ci_appl_ora_notif_svr]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
Oracle OACORE Server	Oracle OACORE Server [cmdb_ci_appl_ora_oacore]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
Oracle OAFM Server	Oracle OAFM Server [cmdb_ci_appl_ora_oafm]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on Unix) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Oracle PeopleSoft	Peoplesoft Application Server [cmdb_ci_appl_peoplesoft]	(If hosted on Unix) On the ServiceNow AI Platform, configure SSH credentials.
		Provide a user with permissions to run commands described in Service Mapping commands not requiring a privileged user.
Oracle Process Manager	Oracle Process Manager [cmdb_ci_appl_ora_pm]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle Report Server	Oracle Report Server [cmdb_ci_appl_ora_report]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle Tnslsnr Engine	Oracle TNS Listener Engine [cmdb_ci_appl_ora_tnslsnr]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on Linux) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Oracle Tuxedo and Oracle Tuxedo Portal	Tuxedo [cmdb_ci_appl_tuxedo] and Tuxedo Portal [cmdb_ci_appl_tuxedo_portal]	• On the ServiceNow AI Platform, configure SSH credentials. • (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server. • (If hosted on Unix) Provide a user with the elevated rights for running commands described in Service Mapping commands requiring a privileged user.
Oracle WebLogic Module	Oracle iAS Web module [cmdb_ci_app_server_ora_ias_m]	(If hosted on Unix) Configure applicative credentials.
Oracle WebLogic Server (version 10.3)	Weblogic [cmdb_ci_app_server_weblogic]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Oracle WebLogic On-demand Router Load Balancer	Websphere ODR LB [cmdb_ci_app_server_ws_odr]	(If hosted on Unix) Configure applicative credentials for Service Mapping.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
Radware Alteon NG	Alteon Load Balancer [cmdb_ci_lb_alteon]	On the ServiceNow AI Platform, configure SNMP credentials.
Radware AppDirector	Radware Load Balancer [cmdb_ci_lb_radware]	On the ServiceNow AI Platform, configure SNMP credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Red Hat Cluster Suite	Cluster Virtual IPs [cmdb_ci_cluster_vip] UNIX Cluster [cmdb_ci_unix_cluster] UNIX Cluster Node [cmdb_ci_unix_cluster_node] UNIX Cluster Resource [cmdb_ci_unix_cluster_resource]	On the ServiceNow AI Platform, configure SSH credentials.
		Provide a user with permissions to read these files: • `/etc/cluster/cluster.conf` • `/etc/hosts`
		Provide a user with elevated rights for running this command: `clustat -x`.
PostgreSQL Database	PostgreSQL Instance [cmdb_ci_db_postgresql_instance]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(For hosted on Unix only) Give this user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
Pivotal RabbitMQ	RabbitMQ Cluster [cmdb_ci_appl_rabbitmq_cluster]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
SAP Business Object XI Schedule	SAP BO BOXISchedule Router [cmdb_ci_appl_sap_bo_schedule]	On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, seew Configure WinRM trusted hosts on MID Server.
SAP Crystal Management Server	SAP Business Objects CMS server [cmdb_ci_appl_sap_bo]	On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
SAP Central Instance/ SAP Central Services (SCS)/ SAP Development Infrastructure/ SAP Evaluated Receipt Settlement (ERS)/ SAP Java Connector/ SAP NetWeaver Dialog Instance/ ABAP SAP Central Services (ASCS)	SAP CI Application [cmdb_ci_appl_sap_ci]/ SAP SCS Application [cmdb_ci_appl_sap_scs]/ SAP DI Application [cmdb_ci_appl_sap_di]/ SAP ERS Application [cmdb_ci_appl_sap_ers]/ SAP JC Application [cmdb_ci_appl_sap_jc]/ SAP DI Application [cmdb_ci_appl_sap_di]/ SAP ASCS Application [cmdb_ci_appl_sap_ascs]	• (If hosted on a UNIX server) On the ServiceNow AI Platform, configure applicative credentials and SSH credentials. The SSH credentials are used to create connections from SAP Central Instance (CI) to other applications and devices. In addition, provide the user elevated rights to be able to run commands requiring a privileged user: Service Mapping commands requiring a privileged user. • (If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials to create connections from SAP CI to other applications and devices. • (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server. • (Optional for top-down discovery by Service Mapping) Provide a user with permissions to run SOAP on RFC read table function. • (Optional for top-down discovery by

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		Service Mapping) If your organization deploys SAP Digital Business Services using web services and you want to map all connected services, perform the following actions: ○ Configure applicative credentials for a user with permissions to access the following tables: CVERS, DBCONS, and RFCDES. ○ Ensure that you have the web service URL. For information on how to create web services, see https://blogs.sap.com/2012/10/24/how-to-create-web-services-abap-based/.
SAP HANA Database	SAP Hana Db [cmdb_ci_appl_sap_hana_db]	(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
SAP Sybase DB	Sybase Instance [cmdb_ci_db_syb_instance]	(If hosted on Unix) Configure applicative credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(For hosted on Unix only) Give this user elevated rights to be able to run Service Mapping commands requiring a privileged user.
		(Optional) For discovering using the WinRM protocol,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		see Configure WinRM trusted hosts on MID Server.
SQL Server Analysis Services (SSAS)	SQL Server Analysis Services [cmdb_ci_db_mssql_analysis]	On the ServiceNow AI Platform, configure Windows credentials.
SQL Server Integration Services (SSIS) Job	SQL Server Integration Services Job [cmdb_ci_db_mssql_int_job]	On the ServiceNow AI Platform, configure Windows credentials.
SQL Server Integration Services (SSIS)	SQL Server Integration Services [cmdb_ci_db_mssql_integration]	On the ServiceNow AI Platform, configure Windows credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
SQL Server Reporting Services (SSRS)	SQL Server Reporting Services [cmdb_ci_db_mssql_reporting]	Provide an SSRS OS user with permissions to run the SSRS Service.
		Configure credentials for a user with permissions to run http get request to the entry point URL.
		(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials. Set the *Credential alias* attribute to [cmdb_ci_db_mssql_reporting]. Configuring this attribute allows you later to use the credential for this CI in patterns.
		(Optional) To discover connections to report datasources, configure credentials for a user with the permission to run http post request to the entry point URL.
		Configure applicative credentials.
		(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
Sun iPlanet Web Server	Iplanet Web Server [cmdb_ci_iplanet_web_server]	Configure applicative credentials.
		On the ServiceNow AI Platform, configure SSH credentials.
Sun Directory	Sun LDAP Server [cmdb_ci_sun_ldap_dir_server]	Configure applicative credentials.
		On the ServiceNow AI Platform, configure SSH credentials.
Sun JES	JES [cmdb_ci_email_server_jes]	Configure applicative credentials.
		On the ServiceNow AI Platform, configure SSH credentials.
Symantec Enterprise Vault	Enterprise Vault [cmdb_ci_email_server_ent_vault]	On the ServiceNow AI Platform, configure Windows credentials.
TIBCO ActiveMatrix BusinessWorks/ TIBCO ActiveMatrix BusinessWorks Process/ TIBCO EMS Queue/ TIBCO Enterprise Message Service (EMS)/	ActiveMatrix Business Works [cmdb_ci_appl_tibco_matrix]/ ActiveMatrix Business Works Process [cmdb_ci_appl_tibco_matrix_proc]/ EMS Queue [cmdb_ci_appl_tibco_queue]/ Tibco Enterprise Message Service [cmdb_ci_appl_tibco_message]	• (If hosted on Unix) Configure applicative credentials for the Tibco EMS. • (If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials. • (If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials. • (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
TIBCO ActiveMatrix Adapter	Tibco Adapter [cmdb_ci_appl_tibco_adapter]	(If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
		(If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
Veritas Enterprise Vault	Enterprise Vault [cmdb_ci_email_server_ent_vault]	(If hosted on a Windows server) On the ServiceNow AI

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	CI type	Configure this
		Platform, configure Windows credentials.

Resolving credential-related issues

If there are discovery issues caused by incorrectly configured or missing credentials, resolve them as explained in the KB article: Credentials & Permissions troubleshooting on Discovery, Service Mapping, Orchestration

Credentials required for host discovery

There are credentials and permissions that Service Mapping requires for discovering hosts.

Prior to discovering applications, Service Mapping accesses hosts that applications run on and discovers them. Make sure that you provide all necessary credentials to allow host discovery.

The ServiceNow AI Platform provides a centralized way of configuring credentials for Discovery, Service Mapping, and other applications. It allows you to configure credentials for hosts and applications only once. If you have already configured host-related credentials for another application, you do not need to do it again for Service Mapping.

Credentials required for discovery of hosts running on UNIX or Linux

Discovery, Orchestration, and Integration Hub explore UNIX and Linux devices by using SSH credentials to execute commands over Secure Shell (SSH).

To provide sufficient permissions, configure one of the following UNIX and Linux credentials:

- Non-root user and password and using the 'sudo' utility to run selected commands as root
- Root user and password

For information on commands requiring sudo-level rights, see Service Mapping commands requiring a privileged user and Commands that require root privileges for Discovery and Orchestration.

To access Unix-based hosts with non-root credentials, provide the read access to the following files and directories:

- /etc/*release
- /etc/bashrc
- /etc/profile
- /proc/cpuinfo
- /proc/vmware/sched/ncpus
- /var/log/dmesg
- APD directory

Credentials required for discovery of hosts running on Windows servers

To provide sufficient permissions, configure one of the following Windows credentials:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- A domain user with local administrator access on the target Windows hosts.
- A domain administrator.

Note: You may need domain administrator credentials only in some cases. For example, when discovering domain controllers.

Configure MID Server to use Windows credentials.

(Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

Resolving credential-related issues

If there are discovery issues caused by incorrectly configured or missing credentials, resolve them as explained in the KB article: Credentials & Permissions troubleshooting on Discovery, Service Mapping, Orchestration

Service Mapping commands requiring a privileged user

Service Mapping uses commands requiring elevated rights to discover and map Unix-based hosts in your organization. In addition to configuring necessary credentials, configure servers in your organization to allow Service Mapping to run these commands with elevated rights.

You do not run these commands directly. Service Mapping uses commands requiring elevated rights as part of the following processes:

- host detection
- process identification on port
- discovering CIs using patterns

Some of these commands do not require elevated rights, unless directories that Service Mapping must access are protected. For more information, see Service Mapping commands not requiring a privileged user.

Operating system commands requiring elevated rights

AIX commands requiring elevated rights

Command	Parameter	Description
`cat`	`file-name`	Shows the file content.

`ls`	`-F-1`
	`-1HF`
	`-w 1`
	`-1`
		Lists the directory content.

`procwdx`	`Process_id`	Gets working directory of a process.
`rmsock`	`Socketname tcpcp`	Finds process listening on a specific port.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AIX commands requiring elevated rights (continued)

Command	Parameter	Description
`lsof`	`-Pnl +M -i`	Shows files or connections associated with the process.
`lstat`	Various options	Fetches information about a link.
`ps`	`eww`	Fetches environment variables for the process on AIX.

HP-UX commands requiring elevated rights

Command	Parameter	Description
`cat`	`file-name`	Shows the file content.

`ls`	`-F-1`
	`-1HF`
	`-w 1`
	`-1`
		Lists the directory content.

`pfiles`	`Process_id`	Shows files or connections associated with the process.
`lsof`	`-Pnl +M -i`	Shows files or connections associated with the process.

Solaris commands requiring elevated rights

Command	Parameter	Description
`cat`	`File_name`	Shows the file content.

`ls`	`ls -F-1`
	`ls -1HF`
	`ls -w 1`
	`ls -1`
		Lists the directory content.

	`e-l`	Fetches file attributes used for file caching decision.
`pargs`	`-e`	Gets the executable directory.
`pargs`	`-a`	Gets the process.
`lsof`	`-Pnl +M -i`	Shows files and connections associated with the process.
`netstat`	`-anu`	Lists the open ports. Required for Solaris version 11.2 or later.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solaris commands requiring elevated rights (continued)

Command	Parameter	Description
Ifconfig	*Ifconfig -a*	Shows interface information (need sudo to get the MAC addresses).
pwdx	*Process_id*	Gets the process information.
pargs	*-e process_id*	Gets the process information.
ps	*-eo user, pid, ppid, comm, args*	Gets the process list.
inetadm	*-l or without params*	Handles the case of application using the inet daemon.

Linux commands requiring elevated rights

Commands	Parameter	Description
cat	*File_name* *scsi* *tr '||n'* */proc* */proc/cpuinfo* *sys* *sys/hypervisor* *sys/hypervisor/compilation* *sys/hypervisor/compilation/compiled by*	Shows information about file content.
cut	*-d* *-f1*	Shows the entire output of a file.
dmidecode	*-t 17* *-s bios-vendor* *-t 1* *-t 2* *-t 3*	Shows information about hardware of the Linux server in a readable format.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux commands requiring elevated rights (continued)

Commands	Parameter	Description
lshw		Shows information about hardware of the Linux server in a readable format.
ls	*ls -F-1* *ls -1HF* *ls -w 1* *ls -1* *ls -l*	Lists directory content.
netstat	*-ltnup* *-ltnp* *-ntup* *-an*	Shows the open network connections.
Isof **Note:** Used in some cases as an alternative to netstat.	*-Pnl +M -i*	Shows files and connections associated with the process.
ss	*-ltnup* *-ntup* *-an*	Shows the open network connections.
stat	*--format="%Y"* *-L file_name*	Shows the modification time of the file.
fdisk	*-l*	Displays all disk partitions

Application commands requiring elevated rights

Service Mapping uses some of these commands in patterns.

Apache Web Server (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Web Server (on UNIX) (continued)

Command	Parameter	Description
ss	Various options	Lists open ports.
ls	Various options	Lists files and folders in the specified folder.

Apache Tomcat (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.
netstat	Various options	Lists the open ports.
ss	Various options	Lists open ports.
find	*-name*	This command is used only for creating the web services connections. Finds specific strings in files and folders.

Apache Tomcat WAR (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.
version.sh/ version.bat	-	Gets the Tomcat version.
netstat	Various options	Lists the open ports.
ss	Various options	Lists open ports.
find	*-name*	This command is used in connection sections of the pattern for discovering Apache Tomcat WAR connections. Finds specific strings in files and folders.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

App TNS Service (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
grep	Various options	Extracts strings from the output.

BIG-IP Global Traffic Manager (GTM) F5 (on F5 BIG-IP)

Command	Parameter	Description
ping	*-a -c 1 [url[1].host]*	Gets the host IP.

BMC CTRL-M Enterprise Manager (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

BMC CTRL-M Gateway (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
find	Various options	This command is used only to create connections to Control M Server. Finds file and folder paths.

Citrix Netscaler (on Citrix Netscaler)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
show	Various options	Retrieves Netscaler IP address.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

EMC Documentum Broker (on UNIX)

Command	Parameter	Description
find	Various options	Finds connections to Docbase.

FormEngine (on UNIX)

Command	Parameter	Description
grep	Various options	(On Unix only) Extracts strings from the output.
ls	Various options	(On Unix only) Lists files and folders in the specified folder.
cat	-	(On Unix only) Shows the file content.

IBM Customer Information Control System (CICS) (on UNIX)

Command	Parameter	Description
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
grep	Various options	Extracts strings from the output.

IBM CTRL-M Server (on UNIX)

Command	Parameter	Description
grep	Various options	(For Unix only) Extracts strings from the output.
find	Various options	(For Unix only) Finds file and folder paths.

IBM CICS Transaction Gateway CTG (on UNIX or Windows)

Command	Parameter	Description
grep	Various options	(For Unix only) This command is used only to create CICS connections. Extracts strings from the output.
cat	-	(For Unix only) This command is used only to create CICS connections. Displays the file content in the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM CICS Transaction Gateway CTG (on UNIX or Windows) (continued)

Command	Parameter	Description
find	Various options	(For Unix only) This command is used only to create CICS connections. Finds specific strings in files and folders.

IBM DB2 (on Linux)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
cat	-	Displays the file content in the output.

IBM J2EE EAR (on Linux)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
find	*-name*	Finds files and folders in the specified folder.

IBM WebSphere Application Server (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.
find	*-name* *-type*	This command is used only for creating the Web Services connections. Finds files and folders for the specific name.

IBM WebSphere Portal (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM WebSphere Portal (on UNIX) (continued)

Command	Parameter	Description
cat	-	Shows the file content.
netstat	Various options	Lists the open ports.
ss	Various options	Lists open ports.

IBM WebSphere Message Broker Flow (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.
ps	*-ef*	Gets the process information.
netstat	Various options	Lists the open ports.
ss	Various options	Lists open ports.

IBM WebSphere Message Broker (on UNIX)

Command	Parameter	Description
ps	*-ef*	Gets the process information.
netstat	Various options	Lists the open ports.
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

IBM MQ (on UNIX) — formerly IBM WebSphere MQ

Command	Parameter	Description
grep	Various options	Extracts strings from the output.

IBM MQ Queue (on UNIX) — formerly IBM WebSphere MQ Queue

Command	Parameter	Description
grep	Various options	Extracts strings from the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Jboss Module (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

Jboss (on UNIX and Windows)

Command	Parameter	Description
ls	Various options	(On Unix only) Lists files and folders in the specified folder.
grep	Various options	Extracts strings from the output.
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
find	Various options	(For Windows only) Finds specific strings in files and folders.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

JRun (on UNIX)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.

JRun WAR Inc (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.
grep	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Exchange Client Access Server (CAS), Microsoft Exchange Hub Transport Server, Microsoft Exchange Mailbox (on Windows)

Command	Parameter	Description
Get-ExchangeServer	*-status -Identity "+$hostname+"| export-clixml $env:TEMP \exchange_pwrshell_output.xml*	Extracts the list of Microsoft Exchange hosts and role
Get-ExchangeServer	*format-table -autosize -HideTableHeaders Name,Fqdn,IsMailboxServer*	Extracts the list of Microsoft Exchange hosts and role in table format
Get-MailboxServer	*format-list*	Extracts the list of Microsoft Exchange Mailbox servers.
Get-ClusteredMailBoxServerStatus	*format-table -Property OperationalMachines*	Extracts information about Microsoft Exchange Mailbox clusters.
Get-ExchangeServer	*format-table -autosize -HideTableHeaders Name,Fqdn,IsHubTransportServer*	Extracts information on Microsoft Exchange Hub Transport servers
Get-StorageGroup	*-server "+$hname +" | select SystemFolderPath | Export-Csv out.csv -notype;cat out.csv*	Extracts information about storage and exports it into a file.
Get-MailboxDatabase	*| select LogFolderPath | format-table -autosize -hideTableHeaders; Get-MailboxDatabase | select EdbFilePath| format-table -autosize -hideTableHeaders;Get-MailboxDatabase | select TemporaryDataFolderPath| format-table -autosize -hideTableHeaders*	Extracts information about the Microsoft Exchange Mailbox database.

MySQL Server (on Linux)

Command	Parameter	Description
grep	Various options	(On Unix only) Extracts strings from the output.
netstat	Various options	Lists open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Server (on Linux) (continued)

Command	Parameter	Description
ss	Various options	Lists open ports.
ps	*--pid=[process.pid] --no-headers -o " %U : %p : %a"*	Gets the userid parameter value.

MySQL Cluster MGM Node (on Linux)

Command	Parameter	Description
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
ps	*--pid=[process.pid] --no-headers -o " %U : %p : %a"*	(On Unix only) Gets the userid parameter value.

Nginx (on UNIX)

Command	Parameter	Description
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
ps	*--pid=[process.pid] --no-headers -o " %U : %p : %a"*	Gets the userid parameter value.
ls	*[IncludeTabletmp[].files]*	This command is used if necessary to create HTTP connections. Lists files and folders in the specified folder.
cat	*[IncludeTabletmp[].files]*	This command is used if necessary to create HTTP connections. Shows the file content.

Oracle E-Business Suite (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.
grep	Various options	Extracts strings from the output.
export	-	Sets variables.
echo	Various options	Prints strings in the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle E-Business Suite (on UNIX) (continued)

Command	Parameter	Description
sqlplus	Various options	Creates the connection to the Oracle instance.
awk	Various options	Manipulates the output.
netstat	Various options	Gets open ports.

Oracle Advanced Queue Queue (on UNIX)

Command	Parameter	Description
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

Oracle Concurrent Server (on UNIX)

Command	Parameter	Description
grep	Various options	(For Unix only) Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
awk	Various options	(For Unix only) Manipulates the output.
ls	Various options	(For Unix only) Lists files and folders in the specified folder.
cat	-	(For Unix only) Displays the file content in the output.
find	Various options	Finds specific strings in files and folders.

Oracle Discoverer Engine (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.
grep	Various options	Extracts strings from the output.
cat	-	Shows the file content.

Oracle Discoverer UI (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Discoverer UI (on UNIX) (continued)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
cat	-	Shows the file content.

Oracle Form UI (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle Fulfillment Server (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle HTTP Server (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle Metric Client (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle OACORE Server (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle OAFM Server (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle Database (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.
ps	*-ef*	Gets the process information.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

Oracle iAS Web Module (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	(For Unix only) Lists files and folders in the specified folder.
cat	-	(For Unix only) Shows the file content.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Process Manager (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle Report Server (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle TnsIsnr Engine (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle WebLogic Module (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.
find	*-name* *-type*	This command is used only for creating the Web Services connections. Finds files and folders for the specific name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle WebLogic Server (version 10.3) (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle WebLogic On-demand Router Load Balancer (on UNIX)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
ls	Various options	Lists files and folders in the specified folder.
cat	-	Shows the file content.

Oracle Tuxedo (on UNIX)

Command	Parameter	Description
grep	Various	Extracts strings from the output
ls	Various	Lists files and folders on the given folder.
cat	-	Displays the file content in the output.
find	*-name*	This command is used only to create web service connections. Finds files and folders for specific name.
netstat	Various	Gets open ports.
ss	Various options	Lists open ports.
ps	*-ef*	(For Unix only) Gets process attributes.

Oracle Tuxedo Portal (on UNIX)

Command	Parameter	Description
grep	Various	Extracts strings from the output
ls	Various	Lists files and folders on the given folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Tuxedo Portal (on UNIX) (continued)

Command	Parameter	Description
cat	-	Displays the file content in the output.

PostgreSQL Database (on UNIX)

Command	Parameter	Description
ls	*[IncludeTabletmp[].files]*	This command is used only to create the HTTP connections. Lists files and folders in the specified folder.

Rabbit MQ (on UNIX)

Command	Parameter	Description
ps - ef	*ps -ef | grep "+ $process.parentProcessId +" |egrep -v -e grep -e beam*	Gets the parent process.

Red Hat Cluster Suite (on Linux)

Command	Parameter	Description
clustat	*-x*	Displays the cluster configuration and the status in XML format.

Commands requiring a privileged user for discovering SAP applications: ABAP, SAP Central Services (ASCS), SAP Control Instance, SAP Netweaver Dialog Instance, SAP Evaluated Receipt Settlement (ERS), SAP Java Cluster (SJC), SAP Central Services (SCS) (on UNIX)

Command	Parameter	Description
pargs	*-e*	Gets the executable directory.
ls	Various options	Lists files and folders in the specified folder.
ping	*-c 1*	Verifies that the host is answering.
netstat	Various options	Lists open ports.
Isof **Note:** Used in some cases as an alternative to netstat.	*-Pnl +M -i*	Shows files and connections associated with the process.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Commands requiring a privileged user for discovering SAP applications: ABAP, SAP Central Services (ASCS), SAP Control Instance, SAP Netweaver Dialog Instance, SAP Evaluated Receipt Settlement (ERS), SAP Java Cluster (SJC), SAP Central Services (SCS) (on UNIX) (continued)

Command	Parameter	Description
CVERS	-	Retrieves the version of installed SAP modules.
DBCONS	-	Retrieves the connection strings to SAP DB.
RFCDES	-	Retrieves the connection string to systems integrated with SAP.
sapcontrol	*-function ABAPGetComponentList*	Retrieves the sysid for SAP applications and components.

Sun Directory (on Solaris)

Command	Parameter	Description
netstat	Various options	Lists the open ports.

Sun iPlanet Web Server (on Solaris)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.
netstat	Various options	Lists the open ports.

Sun JES (on Solaris)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.
grep	Various options	Extracts strings from the output.

Sybase (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.

TIBCO ActiveMatrix Adapter (on UNIX)

Command	Parameter	Description
cat	-	Shows the file content.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

TIBCO ActiveMatrix Adapter (on UNIX) (continued)

Command	Parameter	Description
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

TIBCO ActiveMatrix BusinessWorks (on UNIX)

Command	Parameter	Description
hostname	-	Gets the hostname.
ls	Various options	Lists files and folders in the specified folder.
grep	Various options	Extracts strings from the output.
cat	-	Shows the file content.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

TIBCO ActiveMatrix BusinessWorks Process (on UNIX)

Command	Parameter	Description
hostname	-	Gets the hostname.
ls	Various options	Lists files and folders in the specified folder.
grep	Various options	Extracts strings from the output.
cat	-	Shows the file content.
findstr (for Windows)	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
cut	Various options	Splits the output line.

TIBCO EMS Queue (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

TIBCO Enterprise Message Service (EMS) (on UNIX)

Command	Parameter	Description
ls	Various options	Lists files and folders in the specified folder.

Related topics

Service Mapping commands not requiring a privileged user

Service Mapping commands not requiring a privileged user

Most of commands utilized by Service Mapping for discovery and mapping do not require elevated rights.

Review this list of commands to understand how Service Mapping uses them and to make sure that the virtual security of your company is not compromised.

You do not run these commands directly. Service Mapping uses commands requiring elevated rights as part of the following processes:

- host detection
- process identification on port
- discovering CIs using patterns

Operating systems

AIX commands

Command	Parameter	Description
awk	various options	Parses the string.
istat	*file_name*	Gets last modification time of file.
nslookup	*hostname*	Resolve DNS host name.
ping	*-c count ip/host*	Pings the host.
arp -an	various options	Shows ARP table.
grep	String to search	Finds string in previous command output.
traceroute	*-n ip_address*	Shows layer 3 network route to target host.
netstat	*-an* *-Aan*	Shows open network connections.
ps	*-eo user, pid, ppid, comm, args*	Gets the process list.
find	Directory and parameters to search	Searches for files by name or type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HP-UX commands

Command	Parameter	Description
ping	*-c count ip/host*	Pings the host.
arp -an		Shows the ARP table.
grep	String to search	Finds string in previous command output.
awk	various options	Parses the string.
traceroute	*-n ip_address*	Shows layer 3 network route to target host.
netstat	*-an*	Shows open network connections.
ps	*-eo user, pid, ppid, comm, args*	Gets the process list.
find	Directory and parameters for search	Searches for files by name or type.

IBM / z /OS

Commands	Parameter	Description
sysvar	*SYSNAME*	Gets the system name.
hostname		Retrieves the device name.
grep	String to search	Finds the string in command output.
hostname	*-*	Gets the hostname.
netstat	*-g | awk '{if (NR>4) print \$1,\$3}'*	Gets network information.
netstat	*-r | awk '{if (NR>3) print \$1,\$2,\$5}'*	Gets gateway information.
netstat	*-h | awk '{if (NR>4) print \$1}' | grep -v '127.0.0.1'*	Gets IP array.
netstat	*-R ALL | awk '{if (NR>1) print \$0}'*	Gets interfaces.
netstat	*df -k | awk '{print \$1,\$2}' |sed -e s'/(//' | sed -e s'/)//' 2> /dev/null*	Gets the file system information.
ps	*-*	Gets the process list.
ls	Various options	Lists files and folders in the specified folder.
cat	*file-name*	Shows the file content.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM / z /OS (continued)

Commands	Parameter	Description
cut	Various options	Splits the output according to entered parameters.
uname	*l, v, rm, s*	Gets the OS type.

Linux

Commands	Parameter	Description
nslookup	*hostname*	Resolves the DNS host name.
ping	*-c count ip* *host*	Pings the host.
arp	*-an/n*	Shows the ARP table.
lp n	*-an/n*	Shows the ARP table.
grep	String to search *-v* *hypervisor* *proc* *cpuinfo* *-i 'attached scsi'* *var* *log* *dmesg* *'SCSI device'* *'Disk* *UUID*	Finds the string in previous command output.
awk	Various options	Parses the string.
netstat	*-an*	Shows open network connections.
ss	*-an*	Shows open network connections.
ps	*-eo user, pid, ppid, comm, args, wwwx*	Gets the process list.
egrep	Various options	Searches text in previous command output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux (continued)

Commands	Parameter	Description
find	Directory and parameters for search	Searches for files by name or type.
traceroute	*-n ip_address*	Shows the layer 3 network route to target host.
df	*-kP, echo*	Extracts the information about the file system, including the size in KB.
uptime	Various options	Shows the current time, how long the system has been running, how many users are currently logged on, and the system load averages for the past 1, 5, and 15 minutes.
date	Various options	Displays the date in the time zone on which Linux operating system is configured.
perl	*-e 'print time*	Opens the Perl shell and prints time.
command	*-v vxprint*	This command is used in the Debug mode.
uname	*-a* *-r*	Displays the information about the system.
ipfconfig	*-a*	Shows the configuration of the network interface.
ip	route list *-4 -o addr show* *-6 -o a* *r* String to parse	Displays information about network interface, including the routing table.
route	*-n*	Shows the IP routing table.
ip r		Shows the IP routing table.
sed	*-n '2!p'*	Executes a function using the SED stream editor.
bash	*-c*	Executes the command using the UNIX shell or command language.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux (continued)

Commands	Parameter	Description
waagent	*-version 2>/dev/null*	Displays the information extracted from the Microsoft Azure Linux agent.

Solaris

Command	Parameter	Description
nslookup	*hostname*	Resolves the DNS host name.
ping	*-c count ip/host*	Pings the host.
arp	*-an*	Shows the ARP table.
grep	String to search	Finds the string in previous command output.
awk	Various options	Parses the string.
traceroute	*-n ip_address*	Shows the layer 3 network route to target host.
netstat	*-an*	Shows the open network connections.
zoneadm	*List*	Shows the list of zones.
egrep	Various options	Searches text in previous command output.
find	Directory and parameters	Searches for files by name or type.

Windows WMI Queries

Query	Description
*Select * from Win32_ComputerSystem*	Gets the server basic information like serial number.
Select LastModified from CIM_LogicalFile Where Name=...	Gets the last modification time of a file.
Select AddressWidth from Win32_Processor	Gets the computer architecture (32 bit or 64 bit).
*Select * From Win32_Process where (ProcessId = ?)*	Gets the process list running on the computer.
Select Name from CIM_LogicalDisk	Gets the file systems on the computer, such as C: or D:.
Select FileName, Extension from CIM_Directory where Drive='drive' and Path='path'	Gets the list of files in a specific directory.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows commands

Command	Parameter	Description
get_process_info.exe	*Process_id*	Extracts information on the processes using executable created by ServiceNow.
type	*File_name*	Shows the text file content.
nslookup	*Host_name*	DNS lookup
arp	*-a*	Shows the ARP table.
netstat	*-ano*	Shows the network connections.
findstr	Various options	Finds the string in the previous command output.
dir	*/q* */s /b*	Lists files in the directories.
paping	*--nocolor -c 1 -p* *port_num* *ip_address*	Utility to run TCP ping against given host.
ping	*-n 1* *ip_address/host*	Pings the host.
tracert	*Ip_address*	Shows the layer 3 network route to the target host.

Applications

A10 load balancers

Command	Parameter	Description
show partition	Various options	Get the partition configured.
sh-run all-partitions	Various options	Get details like virtual-server, port, service group, and pool members information.
show slb virtual-server all-partitions	Various options	Get the status of the virtual server and hit count details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Commands not requiring a privileged user for discovering ABAP SAP Central Services (ASCS) (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
msg_server	*-V*	Gets the version.
ping	*-n 1*	(For Windows only) Verifies that the host is answering.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

Active Directory Domain Controller (on Windows)

Command	Parameter	Description
netstat	Various options	Gets open ports on
ss	Various options	Lists open ports.
hostname	*-*	Gets the hostname.
dir	Various options	Lists all files in the specified folder.

Apache Web Server (on UNIX)

Command	Parameter	Description
httpd	*-V*	Gets the version.
cut	Various options	Splits the output line.
opmnctl	*status -fmt %cmp32%prt32%por40%pid* and *@farm status -noheaders -fsep "|" -fmt %cmp%prt%clu %ins%por*	This command is used only for discovering the OC4J connectivity. Gets the status of the OPMN CTRL service.
sort	*-u*	This command is used only for discovering the Weblogic connectivity. Sorts the output.

Apache Web Server (on Windows)

Command	Parameter	Description
httpd.exe	*-V*	Gets the version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Web Server (on Windows) (continued)

Command	Parameter	Description
opmnctl	*status -fmt %cmp32%prt32%por40%pid* and *@farm status -noheaders -fsep "|" -fmt %cmp%prt%clu %ins%por*	This command is used only for discovering the OC4J connectivity. Gets the status of the OPMN CTRL service.

Apache Tomcat (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.
find	*-name*	Creates the web services connections. Finds specific strings in files and folders.
version.sh/ version.bat	*-*	Gets the Tomcat version.
java -cp " + $install_directory + "/lib/catalina.jar org.apache.catalina.util.ServerInfo	*-*	Gets the Tomcat version.
netstat	Various options	Lists the open ports.
find	*-name*	This command is used only for creating the web services connections. Finds specific strings in files and folders.
ss	Various options	(For Linux only) Lists the open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Tomcat WAR (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.
version.sh/ version.bat	-	Gets the Tomcat version.
netstat	Various options	(Windows only) Lists the open ports.

App TNS Service (on Windows)

Command	Parameter	Description
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
netstat	Various options	(On Windows only) Lists open ports.
grep	Various options	(On Windows only) Extracts strings from the output.

BIG-IP Global Traffic Manager (GTM) F5 (on F5 BIG-IP)

Command	Parameter	Description
show	*/gtm pool [wide_ip_pool_4_cmd[1].pool] members*	This command is used only to create connections to pool members. Gets the pool members.
list	*/gtm server [servers_from_cmd[].server] addresses*	This command is used only to create connections to pool members. Lists the alias server addresses.

BIG-IP Local Traffic Manager (LTM) F5 (on F5 BIG-IP)

Command	Parameter	Description
b	*rule [rule name] list*	This command is used only for creating connections from rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

BIG-IP Local Traffic Manager (LTM) F5 (on F5 BIG-IP) (continued)

Command	Parameter	Description
	and *virtual address [entry_point.ip_address]*	Lists existing rules on BIGPipe.
list	*ltm rule [rule name]* and *data-group wts_routing_destination_prod | grep -A 1 [ep_uri]*	This command is used only for creating connections from rules. Lists existing rules on TMSH.

BIG-IP Device Service Clustering

Command	Parameter	Description
show cm traffic-group – get	Various options	This command is used to get traffic and clustering information.

BMC CTRL-M Enterprise Manager (on UNIX or Windows)

Command	Parameter	Description
findstr	Various options	(For Windows only) Extracts strings from the output.
netstat	Various options	(For Windows only) Lists open ports.

BMC CTRL-M Gateway (on UNIX or Windows)

Command	Parameter	Description
findstr	Various options	(For Windows only) Extracts strings from the output.
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
find	Various options	This command is used only to create connections to Control M Server. Finds file and folder paths.
xargs	*grep <string>*	This command is used only to create connections to IBM CTRL-M Server. Runs the command on all lists from the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

BMC IT Asset Management (ITAM) Software

Command	Parameter	Description
tnsping	Various options	Retrieves information about Oracle connections.

CA Identity Manager Provisioning Server (for Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.

Cassandra (for UNIX)

Command	Parameter	Description
find	Various options	Finds the base init configuration file path.
locate	Various options	Finds the base init configuration file path if the 'find' command fails.
nodetool version	Various options	Retrieves the Apache Cassandra version.
nodetool describecluster	Various options	Retrieves the cluster name.
nodetool info	Various options	Retrieves the data center and rack information
nodetool ring	Various options	Retrieves the workload running for Apache Cassandra.
nodetool tablestats	Various options	Retrieves list of keyspaces available on a cluster.
dse	*-v*	Retrieves the DataStax Cassandra version.
dsetool ring	Various options	Retrieves the workload running for DataStax Cassandra.

Cisco ACE Command Line Interface (on Cisco CSM)

Command	Parameter	Description
show	*context, running*	Shows the requested information.
include	-	Extracts strings from the output.
echo	-	Displays strings to the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cisco Global Site Selector (on Cisco CSM)

Command	Parameter	Description
show	*gslb-config domain-list*	Shows the domain list.
show	*gslb-config dns rule*	Shows the DNS rule output.
show	*gslb-config answer-group*	Show the VIP answer.

Citrix Delivery Controller (on Windows)

Command	Parameter	Description
powershell	*Add-PSSnapin Citrix.Broker.Admin.V2; Get-BrokerApplication -Name [entry_point.name] | select Name,CommandLineExecutable,CommandLineArguments,Working | Format-List* and *Add-PSSnapin Citrix.Broker.Admin.V2; Get-BrokerApplication -Name [entry_point.name] | select Name,CommandLineExecutable,CommandLineArguments,Workin | Format-List* And *Add-PSSnapin Citrix.Broker.Admin.V2; Get-BrokerMachine -DesktopGroupName ' [delivery_groups[1].name]*	Gets applications managed by this Delivery Controller.

Citrix Usage Collector (on Windows)

Command	Parameter	Description
qfarm		Discovers epic icons.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Citrix XenApp (on Windows)

Command	Parameter	Description
get_xenapp_apps.ps1	*[entry_point.icon_path]*	Runs powershell commands that retrieve the Citrix icon info from the Citrix repository.
powershell	Various options	Runs powershell commands.

Citrix Presentation Server (on Windows)

Command	Parameter	Description
cscript	*//NoLogo*	Runs VB scripts without a popup box.
GetAppsInFolder.wsf	*[entry_point.icon_path]*	Runs vbscript commands that retrieve the Citrix icon info from the Citrix repository.

Citrix Netscaler (on Citrix Netscaler)

Command	Parameter	Description
show	*cs policy, cs action,cs Policy, lb vserver*	This command is used only for discovering outgoing connections. Shows the requested information.
grep	Various options	Extracts strings from the output.

Citrix Netscaler Global Server Load Balancer (on Citrix Netscaler)

Command	Parameter	Description
show	*ns.conf*	This command is used only for discovering outgoing connections if SNMP or SSH credentials are not provided. Shows the contents of the ns.conf file.

Connect-IT Service (on UNIX or Windows)

Command	Parameter	Description
findstr	Various options	(For Windows only) Extracts strings from the output.
dir	Various options	(For Windows only) Lists files and folders in the specified folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

EMC Documentum Docbase (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

FormEngine (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.

Generic application (on UNIX or Windows)

Command	Parameter	Description
findstr	Various options	(On Windows only) Extracts strings from the output.
tasklist	Various options	Lists all running tasks.
netsh	*http show servicestate*	Shows a snapshot of the HTTP service.

Generic application (on Linux)

Command	Parameter	Description
systemctl list-units --type service	Various options	Show the service status.
/bin/systemctl status "+service_name +".service	Various options	Show the service status.
service --status-all	Various options	Show the service status.
read /etc/services	Various options	Allows access to the file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HAProxy (on UNIX or Windows)

Command	Parameter	Description
ls	Various options	(On Unix only) Lists files and folders in the specified folder.
tail	Various options	(On Unix only) Displays the end of the output.

HP Service Manager Application Server

Command	Parameter	Description
tnsping	Various options	Retrieves information about Oracle connections.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

HP Service Manager Index (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.

HP Service Manager Knowledge Base (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.

IBM WebSphere Message Broker (on Windows)

Command	Parameter	Description
dmqdocbroker	*-s -c getdocbasemap* and *-t [computer_system.primaryHostname] -s -c getdocbasemap* and *-s -c getservermap [docbas]* *-t [computer_system.primaryHostname] -s -c getdocbasemap* [docbase]	Gets the docbase information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM WebSphere Message Broker (on Windows) (continued)

Command	Parameter	Description
del	Various options	Deletes the XML file containing the old binding information.
BTSTask	*ExportBindings /Destination:%TEMP%\MyBindings.xml /Database:[MgmtDbName] /server:[serverName]*	This command is used only if there are not MSSQL credentials. Extracts the binding info and places it into the XML file.

IBM CTRL-M Server (on Unix or Windows)

Command	Parameter	Description
findstr	Various options	(For Windows only) Extracts strings from the output.
dir	Various options	(For Windows only) Lists files and folders in the specified folder.

IBM CICS Transaction Gateway CTG (on UNIX or Windows)

Command	Parameter	Description
findstr	Various options	(For Windows only) This command is used only to create CICS connections. Extracts strings from the output.
dir	Various options	(For Windows only) This command is used only to create CICS connections. Lists files and folders in the specified folder.
type	-	(For Windows only) This command is used only to create CICS connections. Displays the file content in the output.

IBM DB2 (on Linux)

Command	Parameter	Description
source	*[installed_dir]db2profile*	Sets the DB2 variables.
db2	list database directory	Displays the DB2 installation parameters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM DB2 (on Linux) (continued)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
head	*-1*	Displays only the first line from the output.
db2level	*-*	Displays the DB2 version.
cat	*-*	Displays the file content in the output.
xargs	*-I {} echo list tablespace containers for {} ';'*	This command is used only to create storage connections. Runs the command on all lists from the output.

IBM DB2 (on Windows)

Command	Parameter	Description
db2cmd	*/c /w /i*	Lists the DB2 directories.
db2	list database directory	Displays the DB2 installation parameters.
find	Various options	Finds specific strings in files, folders, or standard output.
dir	Various options	Lists files and folders in the specified folder.
db2level	*-*	Displays the DB2 version.
echo	Various options	This command is used only for storage connectivity. Prints the strings in the output.
findstr	Various options	This command is used only for storage connectivity. Extracts strings from the output.
for	Various options	This command is used only for storage connectivity. Runs loops.

IBM DB2 (on IBM z/OS)

Command	Parameter	Description
netstat	Various options	Gets the DB2 name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM J2EE EAR (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.
find	Various options	Finds specific strings in files and folders.
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.

IBM Websphere Application Server (on UNIX)

Command	Parameter	Description
versionInfo.sh	*-*	Gets the Websphere Application Server version.

IBM Websphere Application Server (on IBM z/OS)

Command	Parameter	Description
"oeconsol 'D OMVS,A=ALL' | grep "+$taskname	*-*	Gets the Websphere Application Server task.
netstat	*-a | grep "+$port+" | awk '{print $1}'*	Gets the task name from the listening port.
ps -ef -o user,pid,ppid,comm | awk '{print $1,$2, $3,$4}'		Gets attributes for processes.
$install_dir+"/bin/ versionInfo.sh | grep Version | tail -1"		Gets the version from script.
$install_dir +"/properties/ version/ BASE.product"		Gets the version from version file (failover).
"ls "+$conf_dir+"// */applications/ */deployments/ */META-INF/ application.xml | xargs grep \"<context-root>/"+ $uri_search+"\""+" 1>/tmp/app.txt 2>/dev/null; cat / tmp/app.txt"*		Gets the list of URI files.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM Websphere Application Server (on IBM z/OS) (continued)

Command	Parameter	Description
$ear_directory_name +"//*/*/wsdl/ *.wsdl"*		Reads the wsdl files under the discovered inclusion on the Ear directory.

IBM Websphere Portal (on Linux)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
find	*-name* *-type*	This command is used only for creating the Web Services connections. Finds files and folders for the specific name.

IBM Websphere Portal (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
netstat	Various options	Lists the open ports.
find	Various options	Finds specific strings in files and folders.

IBM WebSphere Message Broker Flow (on UNIX)

Command	Parameter	Description
find	*-name* *-type*	This command is used only for creating the Web Services connections. Finds files and folders for the specific name.
mqsiprofile	*-*	Sets the variables for WebSphere Message Broker.
mqsireportbroker	*[broker name]*	Gets the WebSphere Message Broker information.
mqsibrowse	*[broker name] -t BROKERRESOURCES*	Displays the Message Broker information resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM WebSphere Message Broker Flow (on UNIX) (continued)

Command	Parameter	Description
echo	*$ODBCINI*	(On Unix only) This command is used only to create connections to IBM DB2. Prints strings in the output.

IBM WebSphere Message Broker Flow (on Windows)

Command	Parameter	Description
mqsiprofile	*-*	Sets the variables for WebSphere Message Broker.
mqsireportbroker	*[broker name]*	Gets the WebSphere Message Broker information.
ps	*-ef*	Gets the process information.
mqsibrowse	*[broker name] -t BROKERRESOURCES*	Displays the Message Broker information resources.
db2		Displays DB2 installation parameters: database directory and node directory.

IBM WebSphere Message Broker (on UNIX)

Command	Parameter	Description
mqsiprofile	*-*	Sets the variables for WebSphere Message Broker.
mqsireportbroker	*[broker name]*	Gets the WebSphere Message Broker information.
mqsibrowse	*[broker name] -t BROKERRESOURCES*	Displays the Message Broker information resources.
find	*-name* *-type*	This command is used only for creating the Web Services connections. Finds files and folders for the specific name.
echo	*$ODBCINI*	This command is used only to create connections to IBM DB2. Prints strings in the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM WebSphere Message Broker (on Windows)

Command	Parameter	Description
mqsiprofile	*-*	SSets the variables for WebSphere Message Broker.
mqsireportbroker	*[broker name]*	Gets the WebSphere Message Broker information.
ps	*-ef*	Gets the process information.
mqsibrowse	*[broker name] -t BROKERRESOURCES*	Displays the Message Broker information resources.

IBM MQ (on UNIX) — formerly IBM WebSphere MQ

Command	Parameter	Description
dspmq	*-*	Gets the queue manager status.
dspmqver	*-*	Gets the queue manager version.
runmqsc	*[queue manager]*	Gets the queue information.

IBM MQ (on Windows)

Command	Parameter	Description
grep	Various options	Extracts strings from the output.
dspmq	*-*	Gets the queue manager status.
dspmqver	*-*	Gets the queue manager version.
runmqsc	*[queue manager]*	Gets the queue information.

IBM MQ (on z/OS)

Command	Parameter	Description
netstat	*-a | grep "+$port*	Gets the IBM MQ name.
df	*-k | grep MQ"*	Gets the file system.
"echo "+ $netstat_info[1].mq_name +" | cut -c1-4"	*-*	Gets the queue manager name.
"oeconsol '-"+ $queue_manager+" display qmgr ALL' "	*-*	Gets the MQ logical name information.
"oeconsol '-"+ $queue_manager+" display queue ("+	*-*	Gets the MQ queue information together with

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM MQ (on z/OS) (continued)

Command	Parameter	Description
$entry_point.queue_name +") ALL' | tr \"(\" \"#\" | tr \")\" \"|\""		the information on its connections.
"oeconsol '-"+ $queue_manager+" display queue ("+ $name+")'"	-	Gets the MQ queue information.
"oeconsol '-"+ $queue_manager +" display channel(*) CONNAME where(QMNAME EQ "+ $remote_queue_mngr_name +")'"	-	Gets the channel information.
"oeconsol '-"+ $queue_manager+" display qlocal(*) '"		Gets the local queue information.

IBM MQ Queue (on UNIX) — formerly IBM WebSphere MQ Queue

Command	Parameter	Description
dspmq	-	Gets the queue manager status.
dspmqver	-	Gets the queue manager version.
runmqsc	*[queue manager]*	Gets the queue information.

IBM MQ Queue (on Windows)

Command	Parameter	Description
dspmq	-	Gets the queue manager status.
dspmqver	-	Gets the queue manager version.
runmqsc	*[queue manager]*	Gets the queue information.

Jboss Module (on Unix or Windows)

Command	Parameter	Description
dir	Various options	(On Windows only) Lists files and folders in the specified folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Jboss Module (on Unix or Windows) (continued)

Command	Parameter	Description
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(On Windows only) Finds specific strings in files and folders.
netstat	Various options	(On Windows only) Lists open ports.

Jboss (on UNIX and Windows)

Command	Parameter	Description
find	*-name*	Finds files and folders in the specified folder.
xargs	*grep <string>*	Finds strings in all files found in the output file.
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.
netstat	Various options	Lists open ports.

JRun (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.

JRun WAR Inc (on UNIX and Windows)

Command	Parameter	Description
find	*-name*	Finds files and folders in the specified folder.
xargs	*grep <string>*	Finds strings in the files found in the output file.
grep	Various options	Extracts strings from the output.
dir	Various options	(On Windows only) Lists files and folders in the specified folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

JRun WAR Inc (on UNIX and Windows) (continued)

Command	Parameter	Description
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.

Microsoft Exchange CAS (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.
hostname	-	Gets the hostname.
dir	Various options	Lists files and folders in the specified folder.
tasklist	Various options	Lists all running tasks.
del	Various options	Deletes the XML file containing the old binding information.
echo	Various options	Prints strings in the output.
powershell	*Add-PSSnapin Microsoft.Exchange.Management.PowerShell.E2010; Get-ExchangeServer -status -Identity [hostname]| export-clixml $env:TEMP\exchange_pwrshell_output.xml* and *Add-PSSnapin Microsoft.Exchange.Man...dmin; Get-ExchangeServer -status -Identity [hostname]| export-clixml $env:TEMP\exchange_pwrshell_output.xml* and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.E2010; Get-ExchangeServer| format-table -autosize*	This command is used only if you do not use TCP connections. Gets the CAS status, server roles, and Exchange setup structure. **Important:** Do not use the dollar sign ($) in credentials for Exchange CAS, because Service Mapping uses the dollar sign for pattern variables in the Powershell command.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Exchange CAS (on Windows) (continued)

Command	Parameter	Description

	`-HideTableHeaders Name,Fqdn,IsMailboxServer`
	
	and
	
	`Add-PSSnapin Microsoft.Exchange.Management.PowerShell.Admin; Get-ExchangeServer | format-table -autosize -HideTableHeaders Name,Fqdn,IsMailboxServer`
	
	and
	
	`Add-PSSnapin Microsoft.Exchange.Management.PowerShell.Admin; Get-MailboxServer | format-list`	

`find`	Various options	This command is used only if PowerShell is not operational. Finds specific strings in the files and folders.
`findstr`	Various options	Extracts strings from the output.

Microsoft Dynamic CRM (on Windows)

Command	Parameter	Description
`netstat`	Various options	Lists open ports.
`dir`	Various options	Lists files and folders in the specified folder.
`findstr`	Various options	Extracts strings from the output.
`find`	Various options	Finds specific strings in files and folders.

Microsoft BizTalk Orchestration (on Windows)

Command	Parameter	Description
`echo`	Various options	Prints strings in the output.

Microsoft BizTalk Server (on Windows)

Command	Parameter	Description
`echo`	Various options	Prints strings in the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft BizTalk Server (on Windows) (continued)

Command	Parameter	Description
del	Various options	Deletes the XML file containing the old binding information.
BTSTask	*ExportBindings /Destination:%TEMP%\MyBindings.xml /Database:[MgmtDbName] /server:[serverName]*	This command is used only if there are not MSSQL credentials. Extracts the binding information and places it into the XML file.

Microsoft Exchange BackEnd Server (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.
hostname	-	Gets the hostname.
tasklist	Various options	Lists all running tasks.

Microsoft Exchange FrontEnd Server (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.
hostname	-	Gets the hostname.
tasklist	Various options	Lists all running tasks.

Microsoft FAST Search Server (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.

Microsoft Exchange Hub Transport Server (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.
hostname	-	Gets the hostname.
dir	Various options	Lists files and folders in the specified folder.
tasklist	Various options	Lists all running tasks.
echo	Various options	Prints strings in the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Exchange Hub Transport Server (on Windows) (continued)

Command	Parameter	Description
del	Various options	Deletes the XML file containing the old binding information.
powershell	*Add-PSSnapin Microsoft.Exchange.Management.PowerShell.E2010; Get-ExchangeServer -status -Identity [entry_point.host_name] | export-clixml $env:TEMP |exchange_pwrshell_out* and *"Add-PSSnapin Microsoft.Exchange.Man Get-ExchangeServer -status -Identity [entry_point.host_name] | export-clixml $env:TEMP |exchange_pwrshell_output.xml*	This command is used only if you do not use TCP connections. Gives the CAS status, server roles, and Exchange setup structure. **Important:** Do not use the dollar sign ($) in credentials for Exchange Hub Transport Server, because Service Mapping uses the dollar sign for pattern variables in the powershell command.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Microsoft Internet Information Services (IIS) (on Windows)

Command	Parameter	Description
appcmd.exe	Various options	This command is used only for IIS version 7 and later. Retrieves information about the specified application name.

Microsoft Internet Information Services (IIS) Virtual Directory (on Windows)

Command	Parameter	Description
appcmd.exe	Various options	This command is used only for IIS version 7 and later. Retrieves information about the specified application name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Internet Information Services (IIS) Virtual Directory (on Windows) (continued)

Command	Parameter	Description
iisapp.vbs	Various options	This command is used only for IIS version 6. Retrieves information about the specified application name.
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.
netstat	Various options	Lists open ports.

Microsoft .NET commands (on Windows)

Command	Parameter	Description
ildasm.exe	text that contains process exe path	This file is necessary for performing the put file operation on the ILDisassembler alias. This file disassembles strings from the command .exe file.
.findstr	Various options	Gets the strings from the output.
dir	Various options	Lists all files in the specified folder.
connection_strings_browser.exe	*[encrypted_configs[*].name]*	This command is used only for creating database encrypted connections. This file is necessary for performing the put file operation on theConnectionStringsBrowser alias. Decrypts the database connections.

Microsoft Exchange Mailbox (on Windows)

Command	Parameter	Description
netstat	Various options	Lists open ports.
hostname	-	Gets the hostname.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Exchange Mailbox (on Windows) (continued)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
tasklist	Various options	Lists all running tasks.
del	Various options	Deletes the XML file containing old binding information.
echo	Various options	Prints strings in the output.
powershell	*Add-PSSnapin Microsoft.Exchange.Management.PowerShell.Admin; Get-ExchangeServer -status -Identity [entry_point.host_name]| export-clixml $env:TEMP |exchange_pwrshell_out* and *Add-PSSnapin Microsoft.Exchange.Man Admin; Get-ExchangeServer -status | export-clixml $env:TEMP |exchange_pwrshell_output.xml* and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.E2010; Get-ExchangeServer -status -Identity [entry_point.host_name]| export-clixml $env:TEMP |exchange_pwrshell_output.xml* and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.Admin; Get-ClusteredMailBoxServerStatus | format-table -Property OperationalMachines* and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.Admin; Get-ExchangeServer*	This command is used only if you do not use TCP connections. Gets the status, server roles, and the setup structure of Exchange Mailbox. **Important:** Do not use the dollar sign ($) in credentials for Exchange Mailbox, because Service Mapping uses the dollar sign for pattern variables in the Powershell command.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Exchange Mailbox (on Windows) (continued)

Command	Parameter	Description
	| format-table -autosize -HideTableHeaders Name,Fqdn,IsHubTransportServer and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.E2010; Get-ExchangeServer | format-table -autosize -HideTableHeaders Name,Fqdn,IsHubTransportServer* and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.Admin; Get-StorageGroup -server [hostname] | select SystemFolderPath | Export-Csv out.csv -notype;cat out.csv* and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.Admin; Get-StorageGroup | select SystemFolderPath | Export-Csv out.csv -notype;cat out.csv* and *Add-PSSnapin Microsoft.Exchange.Management.PowerShell.E2010; Get-MailboxDatabase | select LogFolderPath | format-table -autosize -hideTableHeaders; Get-MailboxDatabase | select EdbFilePath| format-table -autosize -hideTableHeaders;Get-MailboxDatabase | select TemporaryDataFolderPath| format-table*	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Exchange Mailbox (on Windows) (continued)

Command	Parameter	Description
	-autosize -hideTableHeaders	
find	Various options	This command is used only if PowerShell is not operational. Finds specific strings in the files and folders.
findstr	Various options	Extracts strings from the output.

Microsoft SharePoint (on Windows)

Command	Parameter	Description
hostname	*-*	Gets the hostname.

Microsoft SQL Database (on Windows)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
hostname	*-*	Gets the hostname.
findstr	Various options	Extracts strings from the output.
sqlcmd	*-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Q select nvcname###inboundtransporturl FROM adm_ReceiveLocation rl, bts_receiveport rp where rl.receiveportid=rp.nid* and *-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Qselect*	This command is used only to discover Microsoft BizTalk. Gets the BizTalk information from the database.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft SQL Database (on Windows) (continued)

Command	Parameter	Description
	nID,nvcName,nApplicationID from bts_receiveport and *-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Qselect nID,nvcName,nApplicationID from bts_sendport* and *-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Qselect nOrcPortID,nReceivePortID,nSendPortid from bts_orchestration_port_binding* and *-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Qselect nID,nOrchestrationID from bts_orchestration_port* and *-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Qselect nID,nvcFullName from bts_orchestration* and	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft SQL Database (on Windows) (continued)

Command	Parameter	Description
	-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Qselect ah.Name###bo.nvcNamespace.bo.nvcName###ase.Name from bts_orchestration bo, adm_Host ah,adm_server ase,adm_Server2HostMapping ash where bo.nAdminHostID = ah.ID and ase.ID = ash.ServerID and ah.ID = ash.HostID and *-Stcp: [computer_system.primaryManagementIP], [port] -d [entry_point.database] -U [username] -P '[password]' -y 0 -h-1 -Qselect nID,nSendPortID,nvcAddress from bts_sendport_transport*	

Microsoft Message Queuing (MSMQ) (on Windows)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.

Microsoft SQL Database (on Windows)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.
sqlservr.exe		Gets the version of the database from the executable file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MongoDB Database (on Windows or UNIX)

Command	Parameter	Description
mongo	--eval db.printShardingStatus()	Retrieves the status of MondoDB that is used for load balancing.
mongo	--db.Name()	Retrieves the name of the database connected to MongoDB server.

MySQL Server (on Windows or UNIX)

Command	Parameter	Description
findstr	Various options	(On Windows only) Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
ps	*--pid=[process.pid] --no-headers -o " %U : %p : %a"*	Gets the userid parameter value.
mysqld-nt/mysqld/ mysqld.exe	*-V*	Gets the version.
mysql	*--user=[username] --password=[password] --port=[jdbc_port] --skip-column-names --silent --execute="SHOW ENGINE NDB STATUS;;"* and *--user=[username] --password=[password] --port=[jdbc_port] --skip-column-names --silent --execute="SHOW SLAVE STATUS;;"* and *--user=[username] --password=[password] --port=[jdbc_port] --skip-column-names --silent --execute="select host from information_schema.processlist*	This command is used only to create cluster node connections, primary replications, and the secondary. Gets the deployment structure of the MySQL server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Server (on Windows or UNIX) (continued)

Command	Parameter	Description
	where command like '%%binlog%%';;"	
which	*mysql*	This command is used only to create cluster node connections, primary replications, and the secondary. Gets the full path of MySQL cluster nodes.

MySQL Cluster MGM Node (on Windows or UNIX)

Command	Parameter	Description
netstat	Various options	Lists open ports.
ss	Various options	Lists open ports.
ndb_mgmd	*-V* and *-e show*	Gets the version and status of the MySQL Cluster MGM node.

Nginx (on UNIX)

Command	Parameter	Description
nginx	*-v*	Gets the version.
egrep	*-v -e ^# -e "|t#"*	This command is used if necessary to create HTTP connections. Ignores special characters.

Oracle E-Business Suite (on Windows or UNIX)

Command	Parameter	Description
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.
export	-	(For Unix only) Sets variables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle E-Business Suite (on Windows or UNIX) (continued)

Command	Parameter	Description
echo	Various options	Prints strings in the output.
sqlplus	Various options	Creates the connection to the Oracle instance.
awk	Various options	(For Unix only) Manipulates the output.
netstat	Various options	(For Windows only) Gets open ports.

Oracle Advanced Queue Queue (on Windows or UNIX)

Command	Parameter	Description
export	-	(On Unix only) Sets a variable.
echo	Various options	Prints strings in the output.
sqlplus	Various options	Creates connection to the Oracle instance.
set	-	(On Windows only) Sets a variable.
netstat	Various options	(On Windows only) Lists open ports.

Oracle Concurrent Server (on Windows or UNIX)

Command	Parameter	Description
findstr	Various options	(For Windows only) Extracts strings from the output.
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
netstat	Various options	(For Windows only) Lists open ports.
awk	Various options	(For Unix only) Manipulates the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.

Oracle Discoverer Engine (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Discoverer Engine (on Windows) (continued)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle Discoverer UI (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle Form UI (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle Fulfillment Server (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle HTTP Server (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle HTTP Server (on Windows) (continued)

Command	Parameter	Description
find	Various options	Finds specific strings in files and folders.

Oracle Metric Client (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle OACORE Server (on Windows)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
dir	Various options	Lists files and folders in the specified folder.
find	Various options	Finds specified strings in the file and folders.

Oracle OAFM Server (on Windows)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
dir	Various options	Lists files and folders in the specified folder.
find	Various options	Finds specified strings in the file and folders.

Oracle Database (on UNIX)

Command	Parameter	Description
export	-	Sets the variable.
sqlplus	*-s [username]/[password]@[entry_point] @[thread]*	This command is used only for getting the advanced queue information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Database (on UNIX) (continued)

Command	Parameter	Description
		Gets the Oracle version and information about advance queue.
echo	*-e set head off feed off pages 0 line 3000||n|"select NAME || '##' || QUEUE_TABLE || '##' || QID || '##' || QUEUE_TYPE || '##' || RETENTION from DBA_QUEUES where OWNER = UPPER([entry_point.scheme]);* and *set head off feed off pages 0 line 3000||n|"select NAME || '##' || QUEUE_TABLE || '##' || QID || '##' || QUEUE_TYPE || '##' || RETENTION from DBA_QUEUES where OWNER = UPPER([entry_point.scheme]);* and *select member as NEEBULA from v|$logfile;* and *select value as NEEBULA from v\|$parameter where name='log_archive_dest'* and *select file_name as NEEBULA from dba_data_files;* and *select DB_LINK,HOST from DBA_DB_LINKS*	This command is used only for getting the advance queue information. Gets the information about advance queue, storage devices, and database links.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Database (on UNIX) (continued)

Command	Parameter	Description
lsnrctl	*version* And *services*	This command is used only for discovering the real application cluster (RAC) version. Gets the RAC version.

Oracle Database (on Windows)

Command	Parameter	Description
set	*-*	Sets the variable.
sqlplus	*-s [username]/ [password]@[entry_point.ip]*	This command is used only for getting the advance queue information. Gets the Oracle version and information about advance queue.
echo	*-e set head off feed off pages 0 line 3000||n|"select NAME || '##' || QUEUE_TABLE || '##' || QID || '##' || QUEUE_TYPE || '##' || RETENTION from DBA_QUEUES where OWNER = UPPER([entry_point.scheme]);* and *set head off feed off pages 0 line 3000||n|"select NAME || '##' || QUEUE_TABLE || '##' || QID || '##' || QUEUE_TYPE || '##' || RETENTION from DBA_QUEUES where OWNER = UPPER([entry_point.scheme]);* and *select member as NEEBULA from v| $logfile;*	This command is used only for getting the advance queue information. Gets the information about advance queue, storage devices, and database links.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Database (on Windows) (continued)

Command	Parameter	Description
	and *select value as NEEBULA from v\$parameter where name='log_archive_dest'* and *select file_name as NEEBULA from dba_data_files;* and *select DB_LINK,HOST from DBA_DB_LINKS*	
lsnrctl	*version* And *services*	This command is used only for discovering the real application cluster (RAC) version. Gets the RAC version.
netstat	Various options	Lists open ports.

Oracle iAS Web Module (on Windows)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
dir	Various options	Lists files and folders in the specified folder.
find	Various options	This command is used only to create the Web Services connections. Finds specified strings in the file and folders.
netstat	Various options	Lists open ports.

Oracle Net Listener (on UNIX)

Command	Parameter	Description
lsnrctl	*version* and	Gets the version and status of the Oracle Net Listener.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Net Listener (on UNIX) (continued)

Command	Parameter	Description
	`status`	
`export`	`-`	Sets the variable.

`crsctl`	`status res -t`
	and
	`config service -d`
		Gets the services status and configuration.

Oracle PeopleSoft (on UNIX)

Command	Parameter	Description
`psadmin`	Various options	Retrieves information from the PeopleSoft repository or server.
`tnsping`	Various options	Retrieves information about the database instance given to the tnsping command.
`sqlplus`	Various options	Connects to the database instance and runs the sql query.

Oracle Process Manager (on Windows)

Command	Parameter	Description
`findstr`	Various options	Extracts strings from the output.
`dir`	Various options	Lists files and folders in the specified folder.
`find`	Various options	Finds specified strings in the file and folders.

Oracle Report Server (on Windows)

Command	Parameter	Description
`findstr`	Various options	Extracts strings from the output.
`dir`	Various options	Lists files and folders in the specified folder.
`find`	Various options	Finds specified strings in the file and folders.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle TnsIsnr Engine (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle WebLogic Module (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle WebLogic Server (version 10.3) (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specific strings in files and folders.

Oracle WebLogic On-demand Router Load Balancer (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
find	Various options	(For Windows only) Finds specific strings in files and folders.
versionInfo.sh	-	Gets the Websphere Application Server version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Tuxedo (on Unix or Windows)

Command	Parameter	Description
findstr	Various	(For Windows only) Extracts strings from the output
dir	Various	(For Windows only) Lists all existing files in the given folder.
find	Various	(For Windows only) Finds specific strings in files and folders.
find	*-name*	This command is used only to create web service connections. Finds files and folders for specific name.
tmadmin	*-v*	Gets the Tuxedo version.
netstat	Various	Gets open ports.
ss	Various options	Lists open ports.

Oracle Tuxedo Portal (on UNIX or Windows)

Command	Parameter	Description
findstr	Various	(For Windows only) Extracts strings from the output
dir	Various	(For Windows only) Lists all existing files in the given folder.
find	Various	(For Windows only) Finds specific strings in files and folders.
find	*-name*	This command is used only to create web service connections. Finds files and folders for specific name.

PostgreSQL Database (on UNIX and Windows)

Command	Parameter	Description
postgres	*-v*	Gets the version.
findstr	Various options	(For Windows only) Extracts strings from the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PostgreSQL Database (on UNIX and Windows) (continued)

Command	Parameter	Description
dir	Various options	(For Windows only) Lists files and folders in the specified folder.

Red Hat Cluster Suite (on Linux)

Command	Parameter	Description
ps	*-ef*	Displays the currently running process.
grep	Various options	Searches plain-text data sets for lines that match the regular expression.
grep	*-v <string>*	Searches plain-text data sets for lines that do not match the regular expression.
ps	*-ef | grep corosync | grep -v grep'*	Checks if the corosync process is running on the server. If the process is running, the Red Hat Cluster is running on this server.
ifconfig	*-*	Displays the current network configuration information.
ip a	*-*	Displays the current network configuration information.
awk	Various options	Manipulates the output.
hostname	*-s*	Displays the Fully Qualified Domain Name (FQDN).
hostname	*-f*	Displays the short host name.

SAP BO BOXI ScheduleRouter (on Windows)

Command	Parameter	Description
findstr	Various options	(For Windows only) Extracts strings from the output.
type	*-*	
tnsping	*instance*	This command is used to create connections to Oracle.

SAP Business Objects CMS Server (on Windows)

Command	Parameter	Description
findstr	Various options	(For Windows only) Extracts strings from the output.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SAP Business Objects CMS Server (on Windows) (continued)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
tnsping	*instance*	This command is used to create connections to Oracle. Gets the version.

SAP Central Instance (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
disp + work	*-V*	Gets the version.
ping	*-n 1*	(For Windows only) Verifies that the host is answering.
netstat	Various options	Lists open ports.
ss	Lists open ports.	Lists open ports.

SAP NetWeaver Dialog Instance (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
disp + work	*-V*	Gets the version.
ping	*-c 1*	(For Unix only) Verifies that the host is answering.
ping	*-n 1*	(For Windows only) Verifies that the host is answering.
netstat	Various options	Lists open ports.

SAP Evaluated Receipt Settlement (ERS) (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
disp + work	*-V*	Gets the version.
ping	*-n 1*	(For Windows only) Verifies that the host is answering.
netstat	Various options	Lists open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SAP Evaluated Receipt Settlement (ERS) (on UNIX or Windows) (continued)

Command	Parameter	Description
ss	Lists open ports.	Lists open ports.

SAP HANA Database (on UNIX or Windows)

Command	Parameter	Description
hdbsql	*-v*	Gets the version.

SAP HANA Database Catalog on UNIX

Command	Parameter	Description
hdbsql	Various options	• Gets schemas info • Gets time zone • Gets disk size

SAP Java Cluster (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
jlaunch	*-V*	Gets the version.
ping	*-n 1*	(For Windows only) Verifies that the host is answering.
netstat	Various options	Lists the open ports.

SAP Central Services (SCS) (on UNIX or Windows)

Command	Parameter	Description
dir	Various options	(For Windows only) Lists files and folders in the specified folder.
msg_server	*-V*	Gets the version.
ping	*-n 1*	(For Windows only) Verifies that the host is answering.
netstat	Various options	Lists the open ports.
ss	Lists open ports.	Lists open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SQL Server Analysis Services (SSAS) (on Windows)

Command	Parameter	Description
MsMdSrv.exe	*-n*	Gets the instance name.
netstat	Various options	Lists the open ports.

SQL Server Integration Services (SSIS) Job (on Windows)

Command	Parameter	Description
MsMdSrv.exe	*-n*	Gets the instance name.
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specified strings in the file and folders.
sqlcmd	*-Stcp: [computer_system.primaryManagementIP], [port] -y 0 -h-1 -E -Qselect convert(varchar(max), convert(varbinary(max),packagedata))from msdb.dbo.sysssispackages where name = '[name]' -o [tmp_dir]/[name].xml* and *-Stcp: [computer_system.primaryManagementIP], [port] -y 0 -h-1 -E -Qselect convert(varchar(max), job_id)###[command]@@@ from msdb.dbo.sysjobsteps* and *-Stcp: [computer_system.primaryManagementIP], [port] -y 0 -h-1 -E -Qselect [name]###convert(varchar(max), job_id) from msdb.dbo.sysjobs*	This command is used only for discovering jobs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SQL Server Integration Services (SSIS) (on Windows)

Command	Parameter	Description
dir	Various options	Lists files and folders in the specified folder.
findstr	Various options	Extracts strings from the output.
find	Various options	Finds specified strings in the file and folders.
sqlcmd	*-Stcp: [computer_system.primaryManagementIP], [port] -h-1 -E -QSELECT [name]###DB FROM msdb.dbo.sysssispackages* and *-Stcp: [computer_system.managementIP], [port] -h-1 -E -QSELECT [name]###DB FROM msdb.dbo.sysssispackages* and *-Stcp: [computer_system.primaryManagementIP], [port] -h-1 -E -QSELECT [name]###JOB FROM msdb.dbo.sysjobs* and *-Stcp: [computer_system.managementIP], [port] -h-1 -E -Q SELECT [name]###JOB FROM msdb.dbo.sysjobs*	This command is used only if you need to discover jobs.

Sun Directory (on Solaris)

Command	Parameter	Description
slapd-dirserv	*-D [instance] -v*	Gets the version.
netstat	Various options	Lists the open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sun JES (on Solaris)

Command	Parameter	Description
configutil	-	Gets the JES configuration.

Sybase (on Unix or Windows)

Command	Parameter	Description
dataserverv	*-v*	Gets the version.
dir	Various options	(For Windows only) Lists files and folders in the specified folder.

Symantec Enterprise Vault (on Windows)

Command	Parameter	Description
findstr	Various options	Extracts strings from the output.
netstat	Various options	Lists open ports.
ping	*-n 1[ip]*	Verifies that the SQL server is running.

TIBCO ActiveMatrix Adapter (on UNIX or Windows)

Command	Parameter	Description
hostname	-	Gets the hostname.
"IFS=$'|n';for i in `find " + $tra_home + " -type f | grep -v /tmp | egrep \".process $\|.substvar$\"`; do result=`grep \"" + $entry_point.queue + "\" $i`; echo $i### $result@@@; done| grep -v \"###@@@\""	-	Scans the label file and brings the actual values.

TIBCO ActiveMatrix BusinessWorks (on UNIX or Windows)

Command	Parameter	Description
TibcoFilesParser.ksh	-	(For Unix only) This file is necessary for performing the put file operation. It replaces all tags in Tibco configuration files and puts them under the /tmp folder.
hostname	-	Gets the hostname.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

TIBCO ActiveMatrix BusinessWorks (on UNIX or Windows) (continued)

Command	Parameter	Description
cd	-	(On Unix only) Changes directory.
chmod	-	(On Unix only) Changes file permissions.
for	Various options	(On Unix only) Runs loops.
echo	Various options	(On Unix only) Prints strings in the output.
if	Various options	(On Unix only) Starts a condition.
wc	*-l*	(On Unix only) Counts the output lines.
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
netstat	Various options	Lists open ports.

TIBCO ActiveMatrix BusinessWorks Process (on UNIX or Windows)

Command	Parameter	Description
hostname	-	Gets the hostname.
cd	-	(On Unix only) Changes directory.
chmod	-	(On Unix only) Changes file permissions.
for	Various options	(On Unix only) Runs loops.
echo	Various options	(On Unix only) Prints strings in the output.
if	Various options	(On Unix only) Starts a condition.
wc	*-l*	(On Unix only) Counts the output lines.
dir	Various options	(On Windows only) Lists files and folders in the specified folder.
findstr	Various options	(On Windows only) Extracts strings from the output.
netstat	Various options	Lists open ports.
ss	Lists open ports.	Lists open ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

TIBCO ActiveMatrix BusinessWorks Process (on UNIX or Windows) (continued)

Command	Parameter	Description
cut	Various options	(On Unix only) Splits the output line.

TIBCO Enterprise Message Service (EMS) Queue (on UNIX or Windows)

Command	Parameter	Description
rm	*-f /tmp/ems.cmd*	Prepares the EMS script.
tibemsadmin	*-server tcp://+[computer_system.primaryIPAddress] [port] -user [username] -password [password] -script [script]*	Connects to the EMS administration server and gets the list of EMS Queue consumers.
echo	Various options	Prints strings in the output.

TIBCO Enterprise Message Service (EMS) (on UNIX and Windows)

Command	Parameter	Description
rm	*-f /tmp/ems.cmd*	Prepares the EMS script.
tibemsadmin	*-server tcp://+[computer_system.primaryIPAddress] [port] -user [username] -password [password] -script [script]*	Connects to the EMS administration server and gets the list of EMS Queue consumers.
echo	Various options	(On Windows only) Prints strings in the output.

Related topics

Service Mapping commands requiring a privileged user

SNMP-based queries

Service Mapping accesses network infrastructure devices like load balancers and routers using Simple Network Management Protocol (SNMP) v1/v2c/v3. Configure SNMP community credentials to enable this type of access.

Configure SNMP community if you use SNMP v1/v2 or SNMPv3 credentials if you use SNMP v3. For operational information, see SSH credentials.

SNMP-based queries have the format of strings of integers.

General SNMP-based queries

Make sure that the user has credentials for SNMPv3 unless read-only community string is used.

- 1.3.6.1.2.1.47.1.1.1.1.13
- 1.3.6.1.2.1.1.6

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 1.3.6.1.2.1.47.1.1.1.1
- 1.3.6.1.4.1.9.9.774.1.1.1
- 1.3.6.1.4.1.9.6

SNMP-based queries for switches and routers

- 1.3.6.1.4.1.9.9.46.1.3.1.1
- 1.3.6.1.4.1.9.9.68.1.2.1.1
- 1.3.6.1.2.1.17.1.4.1
- 1.3.6.1.2.1.17.4.3.1
- 1.3.6.1.2.1.17.7.1.2.2.1
- 1.3.6.1.2.1.17.2.15.1
- 1.3.6.1.4.1.9.9.23.1.2.1.1
- 1.3.6.1.4.1.45.1.6.13.2.1.1
- 1.0.8802.1.1.2.1.4.2.1
- 1.0.8802.1.1.2.1.4.1.1

A10 Thunder Series

- 1.3.6.1.4.1.1872.2.5.1.1.1.11
- 1.3.6.1.4.1.1872.2.5.3.1.3.2.1
- 1.3.6.1.4.1.1872.2.5.3.1.2.4.1
- 1.3.6.1.4.1.1872.2.5.1.3.1.10
- 1.3.6.1.4.1.1872.2.5.4.1.1.4.2.1
- 1.3.6.1.4.1.1872.2.5.3.1.6.3.1
- 1.3.6.1.4.1.22610.2.4.1.1.1
- 1.3.6.1.4.1.22610.2.4.1.6.2
- 1.3.6.1.4.1.22610.2.4.3.4.1.2.1
- 1.3.6.1.4.1.22610.2.4.3.2.1.2.1
- 1.3.6.1.4.1.22610.2.4.3.3.1.2.1
- 1.3.6.1.4.1.22610.2.4.3.4.3.1.1
- 1.3.6.1.4.1.22610.2.4.3.3.3.1.1

BIG-IP Local Traffic Manager (LTM) F5 (on F5 BIG-IP) and BIG-IP Global Traffic Manager (GTM) F5

- 1.3.6.1.4.1.3375.2.1.6.3
- 1.3.6.1.4.1.3375.2.1.4.1
- 1.3.6.1.4.1.3375.2.1.3.3.3
- 1.3.6.1.4.1.3375.2.1.6.2
- 1.3.6.1.4.1.3375.2.1.2.1.1.2.1
- 1.3.6.1.4.1.3375.2.2.10.1.2.1
- 1.3.6.1.4.1.3375.2.3.12.4.2.1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 1.3.6.1.4.1.3375.2.3.12.1.2.1
- 1.3.6.1.4.1.3375.2.1.4.1
- 1.3.6.1.4.1.3375.2.1.6.2
- 1.3.6.1.4.1.3375.2.1.4.2
- 1.3.6.1.4.1.3375.2.2.10.1.2.1
- 1.3.6.1.4.1.3375.2.2.5.6.2.1
- 1.3.6.1.4.1.3375.2.2.10.8.2.1
- 1.3.6.1.4.1.3375.2.2.5.6.2.1
- 1.3.6.1.4.1.3375.2.2.10.1.2.1
- 1.3.6.1.4.1.3375.2.1.4.1
- 1.3.6.1.4.1.3375.2.1.4.2
- 1.3.6.1.4.1.3375.2.1.6.2
- 1.3.6.1.4.1.3375.2.3.11.1.2.1
- 1.3.6.1.4.1.3375.2.3.12.3.2.1
- 1.3.6.1.4.1.3375.2.3.12.4.2.1
- 1.3.6.1.4.1.3375.2.2.5.6.2.1
- 1.3.6.1.4.1.3375.2.3.12.3.2.1
- 1.3.6.1.4.1.3375.2.3.12.4.2.1
- 1.3.6.1.4.1.3375.2.3.12.5.2.1
- 1.3.6.1.4.1.3375.2.3.6.7.2.1
- 1.3.6.1.4.1.3375.2.3.11.4.2.1
- 1.3.6.1.4.1.3375.2.3.3.1.2.1
- 1.3.6.1.4.1.3375.2.3.3.3.2.1
- 1.3.6.1.4.1.3375.2.3.9.1.2.1
- 1.3.6.1.4.1.3375.2.3.9.3.2.1

Cisco Content Services Switch (CSS)

- 1.3.6.1.4.1.2467.1.16.4.1
- 1.3.6.1.4.1.2467.1.31.3
- 1.3.6.1.2.1.1.2
- 1.3.6.1.4.1.2467.4.1
- 1.3.6.1.4.1.2467.4.2
- 1.3.6.1.4.1.2467.4.3
- 1.3.6.1.4.1.2467.4.4
- 1.3.6.1.2.1.1.1
- 1.3.6.1.4.1.2467.1.15.2.1
- 1.3.6.1.4.1.2467.1.16.4.1
- 1.3.6.1.4.1.2467.1.18.2.1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cisco Global Site Selector (GSS)

- 1.3.6.1.2.1.1.5 – to get the hostname.
- 1.3.6.1.2.1.1.1 – to get product name.

Citrix Netscaler

- 1.3.6.1.4.1.5951.4.1.3.6.1
- 1.3.6.1.4.1.5951.4.1.3.2.1
- 1.3.6.1.4.1.5951.4.1.3.1.1
- 1.3.6.1.4.1.5951.4.1.2.7.1
- 1.3.6.1.4.1.5951.4.1.2.2.1
- 1.3.6.1.4.1.5951.4.1.2.11.1
- 1.3.6.1.4.1.5951.4.1.2.1.1
- 1.3.6.1.4.1.5951.4.1.1.6
- 1.3.6.1.4.1.5951.4.1.1.26.1
- 1.3.6.1.4.1.5951.4.1.1.24.1
- 1.3.6.1.4.1.5951.4.1.1.14
- 1.3.6.1.2.1.55.1.8.1
- 1.3.6.1.2.1.47.1.1.1.1
- 1.3.6.1.2.1.4.20.1
- 1.3.6.1.2.1.2.2.1
- 1.3.6.1.2.1.2.2.1
- 1.3.6.1.2.1.1.5

Citrix Netscaler Global Server Load Balancer

- 1.3.6.1.4.1.5951.4.1.1.51.3.1.1
- 1.3.6.1.4.1.5951.4.1.2.1.1

HPE BladeSystem Enclosure

- 1.3.6.1.4.1.232.22.2.3.1
- 1.3.6.1.4.1.11.5.7.1.2

IBM WebSphere DataPower

- For horizontal discovery of the Data Power Hosting Server [cmdb_ci_datapower_server]:
  - 1.3.6.1.2.1.1.1
  - 1.3.6.1.2.1.1.5
  - 1.3.6.1.2.1.4.20.1
  - 1.3.6.1.2.1.2.2.1
  - 1.3.6.1.4.1.14685.3.1.196.1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 1.3.6.1.4.1.14685.3.1.11.1
- 1.3.6.1.4.1.14685.3.1.112.1

- For top-down discovery of the Data Power [cmdb_ci_datapower_server]:
  - 1.3.6.1.2.1.1.1
  - 1.3.6.1.2.1.1.5
  - 1.3.6.1.4.1.14685.3.1.11.1

Juniper Junos

1.3.6.1.4.1.2636.3.1.3

Network Appliance Filer Storage Array

- 1.3.6.1.4.1.789.1.1.2
- 1.3.6.1.4.1.789.1.1.9

Radware Alteon NG

- 1.3.6.1.4.1.1872.2.5.4.2.32.1
- 1.3.6.1.4.1.1872.2.5.1.1.1.11
- 1.3.6.1.4.1.1872.2.5.3.1.3.2.1
- 1.3.6.1.4.1.1872.2.5.3.1.2.4.1
- 1.3.6.1.4.1.1872.2.5.1.3.1.10
- 1.3.6.1.4.1.1872.2.5.4.1.1.4.2.1
- 1.3.6.1.4.1.1872.2.5.3.1.6.3.1
- 1.3.6.1.4.1.1872.2.5.4.1.1.4.2.1
- 1.3.6.1.2.1.1.1
- 1.3.6.1.4.1.1872.2.5.4.1.1.4.5.1
- 1.3.6.1.4.1.1872.2.5.4.1.1.3.3.1
- 1.3.6.1.4.1.1872.2.5.4.1.1.2.2.1
- 1.3.6.1.2.1.4.24.4.1
- 1.3.6.1.4.1.89.35.1.13.1.2
- 1.3.6.1.4.1.89.2.12
- 1.3.6.1.4.1.89.2.4
- 1.3.6.1.4.1.89.35.1.40.52.1.1
- 1.3.6.1.4.1.89.35.1.13.1
- 1.3.6.1.4.1.89.35.1.30.1
- 1.3.6.1.2.1.68.1.3.1
- 1.3.6.1.4.1.89.35.1.13.1
- 1.3.6.1.4.1.89.35.1.11.1
- 1.3.6.1.4.1.89.35.1.40.52.1.1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

ITOM Content Service reference

Create a discovery behavior for Unix-based load balancers

If your network uses load balancers running on Linux hosts, create a discovery behavior to ensure that Service Mapping and Discovery discover them correctly.

Before you begin

Make sure that Discovery is activated.

Role required: service_mapping_admin or admin

About this task

The discovery and mapping process begins with Discovery identifying the host. After the host discovery is complete, Service Mapping starts discovering and mapping applications running on this host.

The default discovery process identifies a configuration item (CI) as a Linux server if it finds an SSH port in use. It may lead to incorrect discovery of load balancers running on Linux as Linux servers. Create a discovery behavior to first discover network devices using SNMP protocol. After Discovery has discovered Unix-based load balancers correctly and created the appropriate CIs for them, you can continue with discovering and mapping application services in Service Mapping.

Discovery performs discovery through a MID Server located in the enterprise private network. A behavior determines:

- What MID Server Discovery is using.
- What probes and patterns the chosen the MID Server is running.

Procedure

1. Create a new behavior for discovering Unix-based load balancers:
   a. Navigate to **Service Mapping** > **Administration** > **Behavior**.
   b. Click **New**.
   c. Enter the name for this behavior.
   d. Right-click on the header bar and click **Save**.
   e. Click **New** on the **Discovery Functionality** tab.
   f. Enter 1 in the **Phase** field.
   g. Select **SNMP only** from the **Functionality definition** list.
   h. Select the relevant MID Servers in the **MID Servers** area.
   i. Check the **Active** check box.
   j. Click **Submit**.
2. Create a discovery schedule for the newly defined behavior:
   a. Navigate to **Discovery** > **Discovery Schedules**.
   b. Click **New**.
   c. Fill in the fields as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule Form for Discovery

Field	Description
Name	Enter a unique, descriptive name for your schedule.
Discover	Select **Configuration items**.
MID Server selection method	Select **Use behavior**.
Behavior	Select the behavior you created in step 1.
Active	Select the check box to enable this schedule. If you clear the check box, the schedule is disabled, but you can still run a discovery manually from this form, using the configured values.
Max run time	Set a time limit for running this schedule. When the configured time elapses, the remaining tasks for the discovery are canceled, even if the scan is not complete. Use this field to limit system load to a desirable time window. If no value is entered in this field, this schedule runs until complete.
Max Idle time	Set a time limit for inactivity during discovery to automatically cancel stalled processes and log the event. A discovery is only considered idle if no input records are processed and nothing remains in the processing state; idle time is measured from the most recent change to an ECC Queue record. The configured idle time is fixed at the start of the discovery and won't reflect any edits until the next run. If no value is entered in this field, this schedule runs until complete or cancelled. **Note:** By default, the **Max Idle time** field doesn't display in the Discovery Schedule form. To include this parameter, open the Discovery Schedules [discovery_schedule] table, add the column, and configure the Max Idle time directly in the table. For information about customizing columns in tables, see Personalize a list.
Run and related fields	Determines the run schedule of the discovery. Configure the frequency in the **Run** field and the other fields that appear to specify an exact time. **Note:** The run time always uses the system time zone. If you add the optional **Run as tz** field, it has no effect on the actual runtime.

3. Create an IP range for the discovery schedule:

Important: The IP range for the schedule must correlate with the IP range of the MID Server that you assigned to the behavior. The MID Server must also have the appropriate application or the **ALL** application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. On the Discovery IP Ranges tab, click **New**.

b. Enter the name for this IP range.

c. Configure the IP range by entering IP addresses in the **Starting IP** and **Ending IP** fields.

4. Click **Submit**.

Related topics

Pattern-based discovery in Service Mapping

MID Server

Verify that Service Mapping is set up properly

Before you start mapping application services, verify that you configured MID Servers to discover application services and provided all the necessary host and application credentials.

Before you begin

To check readiness of mapping based on Predictive Intelligence, install Service Mapping Plus from ServiceNow Store.

Role required: service_mapping_admin

About this task

Service Mapping performs an automatic scan of the following resources and configurations to give you the readiness status:

- (Mandatory) There are MID Servers with correctly configured IP ranges and the applications and capabilities settings. For more information, see MID Server configuration for Service Mapping.

  If this check does not recognize correctly configured MID Servers, you cannot perform the bulk discovery.

- (Optional) There are credentials for operating systems and applications as described in Prerequisites for performing top-down discovery using Service Mapping.
- (Optional) There are at least three load balancers recently discovered during horizontal discovery.
- (Optional) There are at least 100 hosts discovered by the horizontal discovery in the last 14 days.
- (Optional) The data collection using Netflow or VPC Flow Logs is configured. This configuration is necessary to enhance the traffic-based discovery in Service Mapping. For more information, see Data collection and discovery using Netflow and Data collection and discovery using VPC Flow Logs.
- (Optional) Cloud Discovery is activated and configured. This setting is necessary to discover application service based on IaaS.

If an optional resource or configuration is missing, you can still map application service, but it may result in partial discovery of some CIs and connections.

In addition, Service Mapping performs the readiness check for mapping functionality based on Predictive Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Service Mapping** > **Home**.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.
2. Verify that Service Mapping is set up and ready for discovery.
   If some obligatory configurations are missing, the error displays at the top of the window.

   a. Click the link in the error or click the **Readiness Checklist** icon ( ).

   b. In the Service Mapping Readiness Checklist window, click the link next to the setup category with the error icon ( ).
   You cannot perform discovery and mapping unless there are no errors in the Readiness Checklist window.

   c. Perform the necessary action to complete the setup.
   For example, configure missing credentials.

3. If more than one setup category has an error, return to the Readiness Checklist window and fix any other errors.
4. Verify that Service Mapping is ready for discovery based on Predictive Intelligence.

   a. Navigate to **Workspaces** > **Service Mapping**, and then select **Application Service Readiness** icon .

   b. Verify that all items in the **Prerequisites status** list are in the **Ready** state.
   If any prerequisites are not fulfilled, perform the necessary procedure, and then select **Check readiness**.

What to do next

Map business services in bulk

Related topics

Application service readiness dashboard in configurable workspace

Quick start tests for Service Mapping

After upgrades and deployments of new applications or integrations, run quick start tests to verify that Service Mapping still works. If you have customized Service Mapping, copy the quick start tests and configure them for your customizations.

Service Mapping quick start tests require activating the Service Mapping (com.snc.service-mapping) plugin and loading the demo data.

Service Mapping quick start tests

Test	Description	Release version
SM OOTB: Tests SM application visibility by roles	Validate the visibility of Service Mapping applications and modules for different roles. For example, the test verifies	Madrid

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping quick start tests (continued)

Test	Description	Release version
	that a user logged in with the service_mapping_user role cannot access the Administration module under Service Mapping.	
SM OOTB: Service map verification	Use this test template to create custom tests for verifying that the topology of the most significant services is unchanged. Configure values under Test Run Data Sets, to identify the name of the service to test and the nodes that you expect to find in this service. For the node name values, enter the node attribute exactly as it appears on the map, for example, "Apache server." If the attribute name for a node is truncated or shows that a node is a CI cluster, configure this node name value to reflect the way it appears on the map, for example, "11x ExchangeF" for a group of 11 Exchange FrontEnd servers. **Note:** You cannot use this test to verify the service content beyond the top, unexpanded level.	Orlando, updated Paris
SM OOTB: Check UI accessibility after an upgrade	Run this test to check that the Service Mapping UI is fully functional after an upgrade.	Orlando

Related topics

Run an automated test

Advanced Service Mapping configuration

Fine-tune Service Mapping collaboration with other components and modules as well as customize data display in service instance maps.

In addition to the basic setup described in Configuring Service Mapping, you can perform optional configurations to meet the needs of your organization.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Enable traffic-based discovery for CI types or specific CIs
- Tag-based discovery configuration
- Configuration file tracking
- Fine-tune tracking changes for the change history
- Data collection and discovery using Netflow
- Data collection and discovery using VPC Flow Logs
- Configure Search Assistant for Windows
- Modify display for CI attributes
- Configure support for Windows servers with non-English OS
- Upload the rctrlx.exe file to MID Servers
- Fine-tune Service Mapping with MID affinity and IP reuse

Enable traffic-based discovery for CI types or specific CIs

Service Mapping can discover and map CIs by detecting the inbound and outbound traffic that the CIs generate. Create a traffic-based discovery rule to determine which configuration items are available for traffic-based mapping.

Before you begin

If your ServiceNow instance uses domain separation and you have access to the global domain, log in to the relevant domain. The selected domain must be a domain without any child domains.

Role required: service_mapping_admin

- Ensure that the traffic-based discovery is enabled at the Service Mapping product level: navigate to **Service Mapping** > **Administration** > **Properties** and verify that the **Traffic based discovery** check box is selected.
- Enable traffic-based discovery for a specific service instance as described in Map a single application service using classic Service Mapping.

About this task

You may choose to use traffic-based discovery in addition to the pattern-based mapping.

You can create rules to configure Service Mapping to use traffic-based discovery for certain CIs. You can configure traffic-based discovery rules to monitor specific CIs or types of CIs.

For example, you enable traffic-based discovery for a specific service instance that includes a Tomcat server, a MySQL database, and a web application. You create a CI type rule that excludes all Tomcat servers from traffic-based discovery. However, you also create a discovery rule enabling traffic-based discovery for a specific Tomcat server. Service Mapping uses traffic-based discovery to discover the MySQL database, the web application, and that specific Tomcat despite the CI type rule excluding all Tomcat servers.

Depending on your configuration, the behavior of traffic-based discovery is different. If Predictive Intelligence is enabled, Service Mapping automatically adds connections to service instances based on connection rules. These suggestions are generated using traffic-related data from the Configuration Management Database (CMDB) and the analysis of application fingerprints, CIs, and processes by Predictive Intelligence.

If discovery based on Predictive Intelligence isn't enabled, Service Mapping automatically adds traffic-based connections using data from the CMDB to the service instances. However, to keep your service instances well-organized, you might need to manually remove connections to

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

irrelevant CIs. To learn more about traffic-based method, see Traffic-based discovery in Service Mapping.

If your instance uses domain separation, you can create traffic-based rules for specific domains. Rules in the base system are assigned to the global domain and apply to all domains of all levels.

When you create a rule for a specific domain, the new rule is used only for this domain and does not exist in any other domains. If you customize an existing rule in the global domain and assign it to a specific domain, you create a copy of the global rule, which is still used in all other domains except the domain that has the customized version of this rule. Likewise, if you customize a rule in the global domain, the change affects all domains except the one that uses a customized copy of this rule.

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Traffic Based Discovery**.
2. Click **New**.
3. Define the rule parameters as follows:

Field	Description
Action	Select **Enable** to add traffic-based connections to the specific CI or CI Type.
Rule Scope	Select **Specific CI** to detect traffic for one configuration item, or select **CI Type** to detect traffic for all configuration items in one of the CI-based table.
CI/CI Type	Select a specific CI, or select the table that contains the CIs for which you want to detect traffic.
Domain	Select the domain to which the rule belongs.

4. Click **Submit**.
5. To fine-tune traffic-based discovery, define advanced parameters as follows:

Parameter	Description
sa.traffic_based_discovery.conn_aging_time	Time period in hours for a Traffic Based Connection to remain active since last discovered. ○ **Type**: integer ○ **Default value**: 72 ○ **Other possible values**: any number higher than 24 ○ **Location**: **Service Mapping** > **Administration** > **Properties**.
sa.traffic_based_discovery.ignored_ports	Ports to ignore when found by traffic-based discovery. This property is available in the System Property [sys_properties] table. Change this property to define ports that Service Mapping ignores while performing traffic-based discovery. It makes discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
	more efficient since resources are not wasted on discovering irrelevant connections. ○ **Type**: string ○ **Default value**: 445, 139, 111, 2049, 860, 3260, 135, 53 ○ **Other possible values**: any relevant port numbers ○ **Location**: System Property [sys_properties] table
sa.traffic_based_discovery.max_connections	Maximum number of traffic-based connections from a single CI. This property is available in the System Property [sys_properties] table. This property helps to keep the map size reasonable by limiting the number of possible CI connections. ○ **Type**: integer ○ **Default value**: 30 ○ **Other possible values**: any number higher than 1 ○ **Location**: System Property [sys_properties] table

Tag-based discovery configuration

You can refine the default configuration to control which CIs Service Mapping includes in application services during the tag-based discovery process.

Note: Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Service Mapping provides preconfigured CI relationships which can participate in tag-based discovery. You have the ability to customize these CI relationships and refine the default configuration using one of the following methods:

- include CIs based on classes in tag-based discovery;
- exclude CIs based on their install status;
- exclude preconfigured CI relationships;
- add or edit CI relationships.

To view preconfigured CI relationships which Service Mapping includes in application services during tag-based discovery, see Preconfigured CI relationships in tag-based discovery.

To view properties installed with Service Mapping that relate to tag-based discovery, see Properties related to tag-based discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Include CIs based on classes in tag-based discovery

Use CI classes to identify which CIs Service Mapping includes in application services during tag-based discovery. By default, this list appears empty and includes all CIs belonging to all CI classes. Define CI classes and allow only CIs belonging to these classes or their extensions to participate in tag-based discovery.

Before you begin

Learn about default feature configuration in Tag-based discovery in Service Mapping.

> **Note:** Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Role required: service_mapping_admin

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Properties**.
2. Enter CI classes, separated by commas, in the allow list that controls which tagged CIs participate in tag-based mapping.

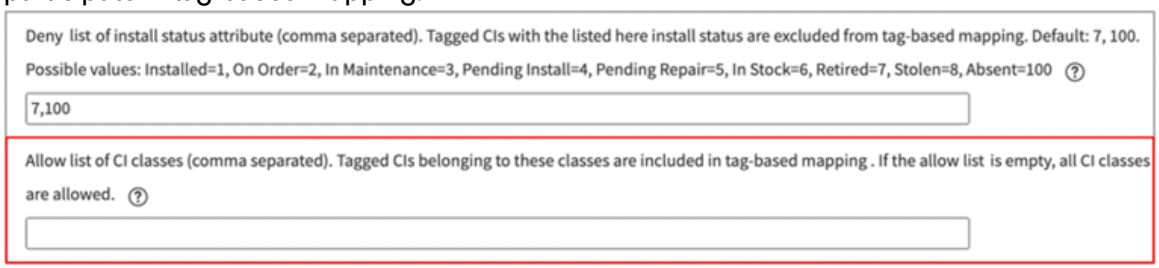

3. Click **Save**.

What to do next

If you refine the default configuration after creating appliations based on tags, see Recalculate previously created application services for the next steps.

Exclude CIs by installation status in tag-based discovery

Use the CI installation status to create CI classes and exclude CIs from application services during tag-based discovery. By default, application services based on tags exclude CIs with the Retired or Absent install status. You can expand this list of excluded CIs to include additional installation statuses such as Pending install or Stolen.

Before you begin

Learn about default feature configuration in Tag-based discovery in Service Mapping.

> **Note:** Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Role required: service_mapping_admin

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Properties**.
2. Enter values for installation statuses into the deny list that controls which CIs to exclude from tag-based mapping based on install status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use commas to separate the values.

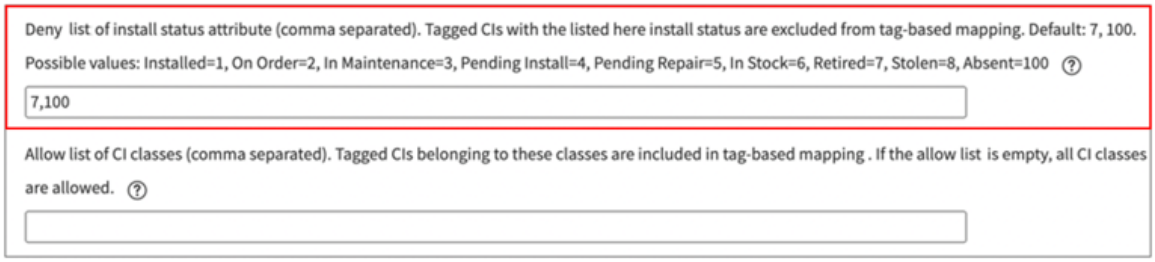

3. Click **Save**.

What to do next

If you refine the default configuration after creating appliations based on tags, see Recalculate previously created application services for the next steps.

Exclude CI relationships from tag-based discovery

Exclude preconfigured CI relationships from the tag-based discovery process to refine the application services created. Service Mapping consists of preconfigured CI relationships and includes CIs in these relationships even if they do not have assigned tags. Choose which preconfigured CI relationships to exclude and the remaining CI relationships participate in tag-based discovery.

Before you begin

Learn about default feature configuration in Tag-based discovery in Service Mapping.

Note: Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Role required: service_mapping_admin

About this task

The Traversal Rules for Application Services [svc_traversal_rules] table stores the CI relationships that participate in tag-based discovery. By default, preconfigured CI relationships used in tag-based discovery have the **true** value in the **Base System Rule** column. Use this information to exclude preconfigured CI relationships from tag-based discovery.

Important: You cannot delete or modify preconfigured CI relationships used for tag-based discovery from the Traversal Rules for Application Services [svc_traversal_rules] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Tag-based Service Traversal Rules**.

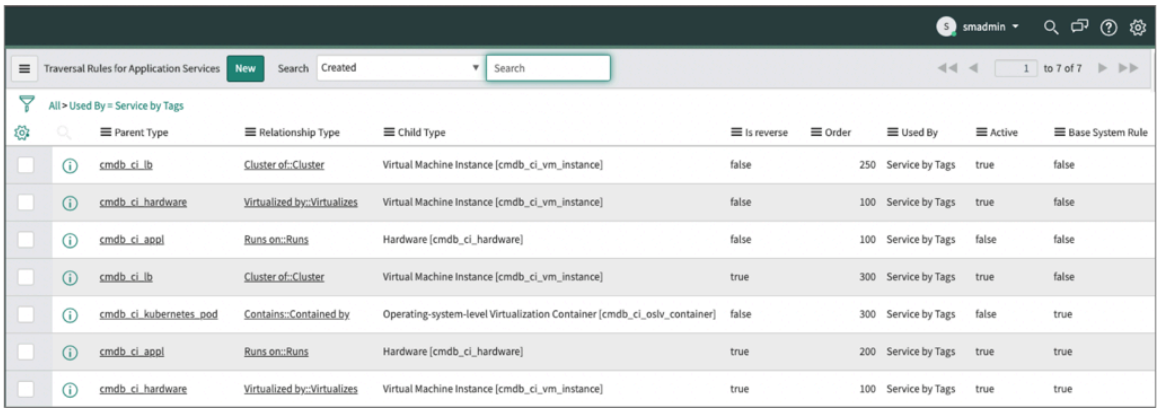

2. Click the preconfigured CI relationship you want to exclude from tag-based discovery.

   To view these CI relationships, see Preconfigured CI relationships in tag-based discovery

3. Clear the **Active** check box.
4. Click **Submit** in the **Traversal Rules for Application Services** form.

What to do next

If you refine the default configuration after creating appliations based on tags, see Recalculate previously created application services for the next steps.

Add or edit CI relationships in tag-based discovery

Activate pre-configured CI relationships that Service Mapping does not include by default or add additional relationships to use in the tag-based discovery process. For example, you can add a CI relationship between Linux servers and storage devices to discover servers hosting storage devices based on tags.

Before you begin

Learn about default feature configuration in Tag-based discovery in Service Mapping.

Note: Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Role required: service_mapping_admin

About this task

The Application Services [svc_traversal_rules] table appears filtered to show only CI relationships used for tag-based discovery. By default, preconfigured CI relationships used in tag-based discovery have the **true** value in the **Base System Rule** column. Use this information to create new CI relationships or to edit preconfigured CI relationships in tag-based discovery.

Important: You cannot delete or modify preconfigured CI relationships used for tag-based discovery from the Traversal Rules for Application Services [svc_traversal_rules] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Tag-based Service Traversal Rules**.

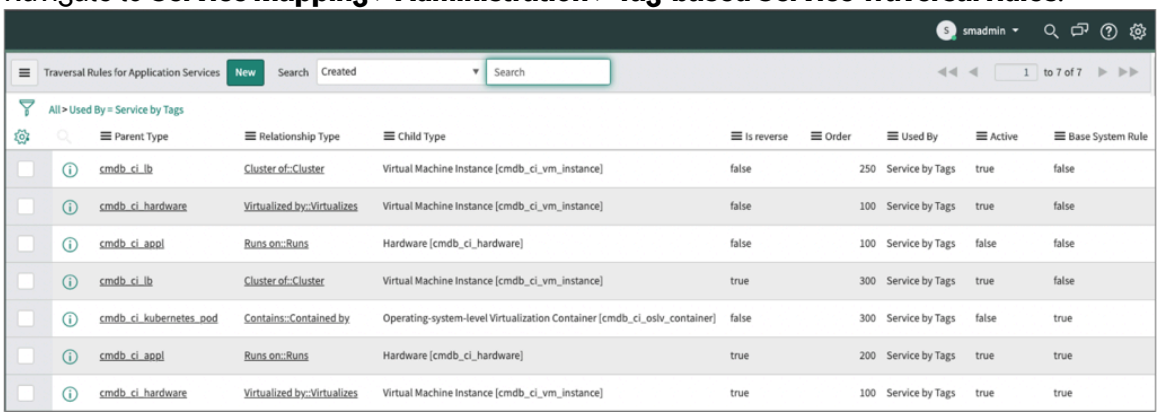

2. Click the existing CI relationship that you want to modify or click **New** to use an additional CI relationship in the tag-based discovery process.

   To view default CI relationships, see Preconfigured CI relationships in tag-based discovery.

3. Complete the **Traversal Rules for Application Services** form.
   To complete this form, see Traversal Rules for Application Services form completion.
4. Click **Submit**.

What to do next

If you refine the default configuration after creating appliations based on tags, see Recalculate previously created application services for the next steps

Recalculate previously created application services

Recalculate application services after you refine the default configuration to control which CIs to include during the tag-based discovery process.

Before you begin

Complete the steps to refine the default configuration for the tag-based discovery process. For more details, see Tag-based discovery configuration.

Note: Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Role required: service_mapping_admin

Procedure

1. Navigate to **Service Mapping** > **Services** > **Application Services**.
2. Click the relevant application service.
3. Click **Recalculate Service** under **Related Links**.
4. Repeat these steps for all relevant tag-based services.

Fine-tune tracking changes for the change history

Define CI fields for which the system reflects changes in the change history for application services. The change history view shows changes to CIs making up application services as well as changes to application services themselves.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: admin

About this task
Details about changes to a service instance and to its CIs are stored in the CMDB. Typically, these changes reflect adding or removing CIs from a service instance, upgrading or updating CIs, or modifying CI configuration files. The system gathers this data by querying CMDB tables and then creating the change history view. In deployments where Service Mapping is activated, the type of change information Service Mapping queries depends on discovery patterns that Service Mapping uses to discover CIs.

Changes to configuration files are associated with CIs to which these files belong. Maps show configuration file changes as changes to related CIs.

Procedure
1. Navigate to **All** > **Configuration** > **Exclusion List**.
   The Exclusion Lists list opens displaying records that the system uses to exclude changes from the change history view.
2. To make the system ignore changes to an extra field:
   a. Click **New**.
   b. Enter the CI class name in the **Table Name** field.
   For example, enter `cmdb_ci` for the Configuration Item [cmdb_ci] class.

   Note: The Service Mapping user interface refers to CI classes as CI types.

   c. Enter the name of the field you want to exclude in the **Field Name** field.
   For example, enter `work_notes` to exclude changes to the item work notes from the change history view.
   d. Select **Submit**.
3. To make the system include previously ignored changes in the change history:
   a. Select the check boxes next to the relevant records.
   b. Select **Delete** from the **Actions on selected rows** list.

Related topics
View the change history of application services in legacy Agent Workspace

View the change history of application services in classic Service Mapping

Data collection and discovery using Netflow
Service Mapping can perform discovery based on data collected using the Netflow protocol. Netflow is a protocol that Service Mapping can use to collect data about CIs and their connections along with Netstat and lsof commands.

Using the Netflow protocol for collecting data is one of the traffic-based discovery methods. Other methods deployed by Service Mapping are using netstat and lsof commands and the VPC Flow Logs. For more information, refer to Traffic-based discovery in Service Mapping.

In base systems, which are the default or standard configurations, traffic-based discovery relies solely on TCP-related data collected using the *`netstat`*, *`ss`*, and *`lsof`* commands. Discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

based on Netflow and VPC logs requires additional configuration. You can enrich your traffic-based discovery by configuring Service Mapping to use the Netflow protocol.

The component, which receives data in the Netflow format is the Netflow Collector. Its location depends on whether you configure data collection for testing purposes or standard operation:

For the test purposes

This setup results in half automated data collection flow, where Service Mapping imports data only if you manually copy it from the Netflow Collector. You place the Netflow Collector on a server inside your organization network. This must be a server different from the server hosting the MID Server. You configure and test this setup as described in Configure a Netflow data import for testing.

For standard operation

This setup results in fully automated data collection flow, where all involved components send, collect and analyze data automatically. You place the Netflow Collector on the same server as the MID Server inside your organization network. For instructions, see Configure data collection using Netflow.

Netflow-based discovery has the following flow:

1. The Netflow daemon runs and receives data from switches communicating with servers in the organization. The Netflow Collector writes received data from the Netflow daemon.
2. The server, hosting the Netflow collector, uses the Netflow nfdump utility to write the data into the nfdump output file. This file summarizes the raw data on all switches used for server communication.

Collecting data and writing it into the nfdump output file

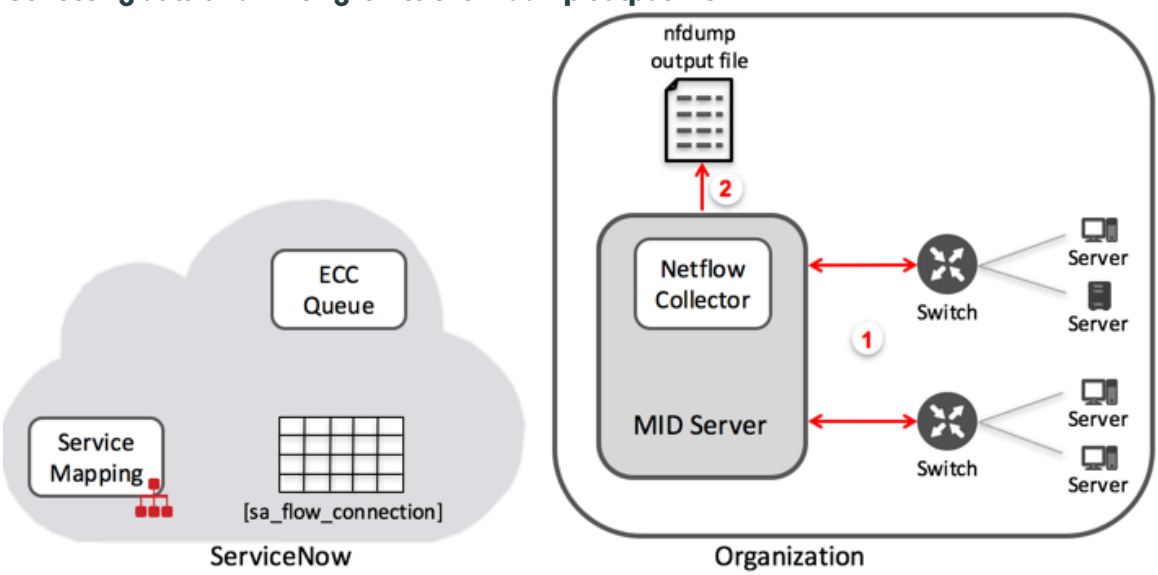

3. In testing setups, where the Netflow Collector is located not on the same server as the MID Server, you may need to convert the nfdump into the gzip format. Then you must manually copy the raw data in the nfdump output file onto the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Copying the nfdump output file onto the MID Server

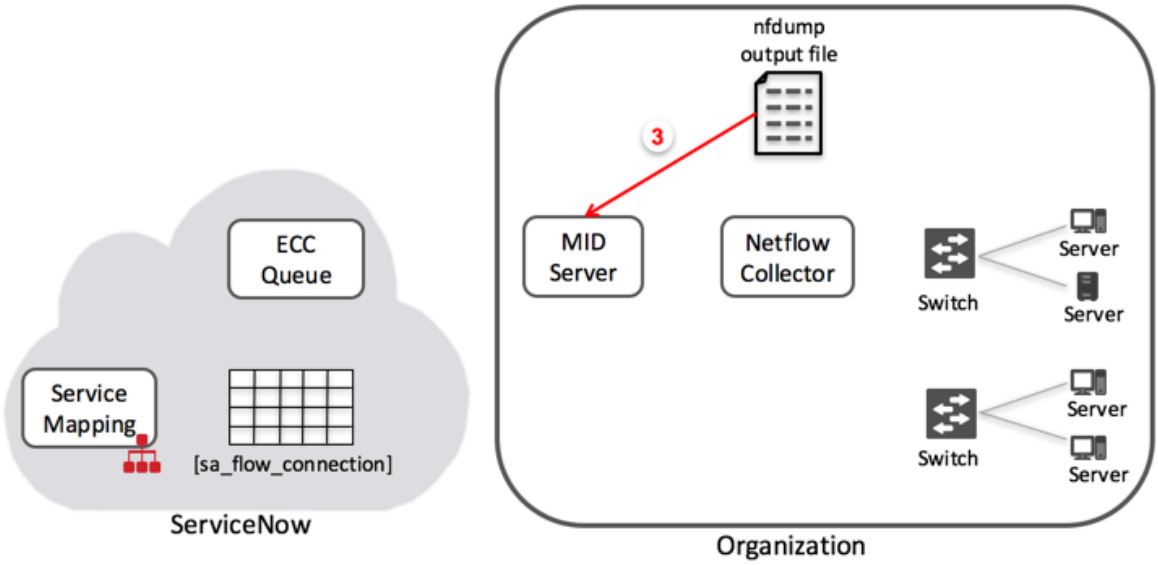

4. The MID Server processes the raw data in the nfdump output file and places the processed information onto the ECC queue.

Analyzing the raw data and placing it at the ECC Queue

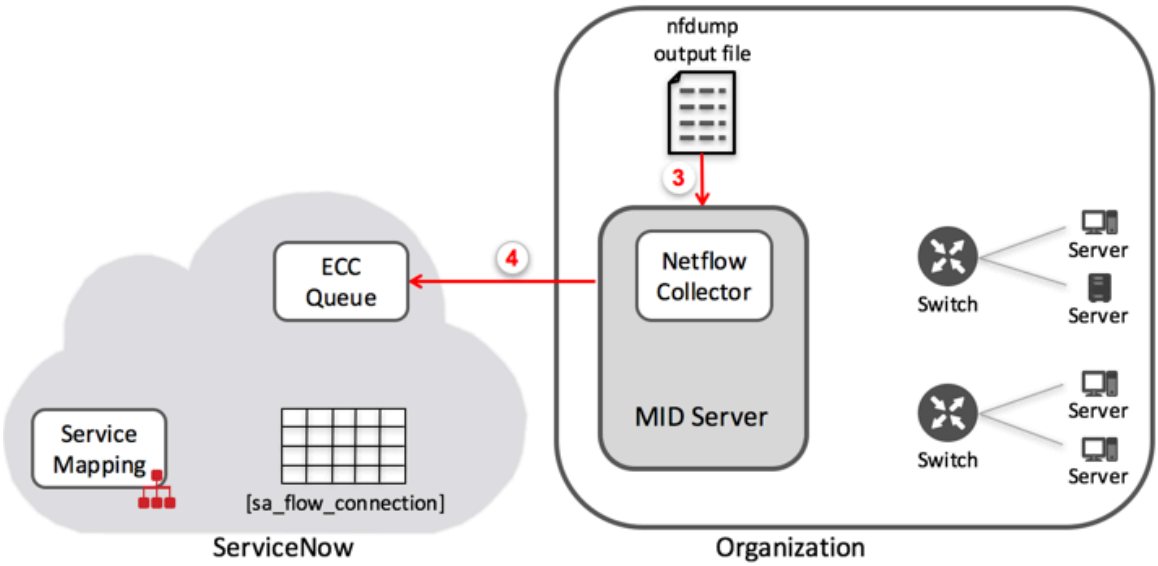

5. A sensor retrieves the processes data from the ECC queue and writes it into the Flow Connection [sa_flow_connection] table.

6. Whenever Service Mapping checks the ECC queue and receives information on a discovered CI, it checks these tables for any data on outbound connections related to the CI: the cmdb_tcp and sa_flow_connection tables. If these two tables contain unique data that patterns

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

did not discover, Service Mapping enriches the information about the CI connections and adds them to the map.

Service Mapping retrieves data from the sa_flow_connection table

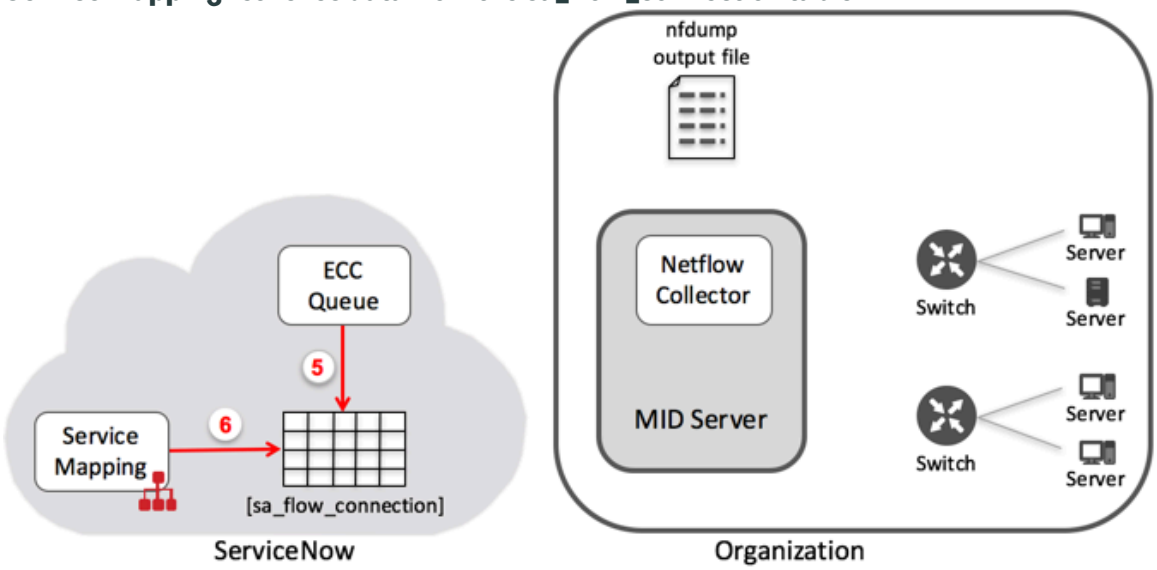

Related topics

https://support.servicenow.com/kb_view.do?sysparm_article=KB0597467

Configure a Netflow data import for testing

Configure and test Service Mapping discovery process based on data collected using the Netflow protocol.

Before you begin

Learn about Traffic-based discovery in Service Mapping.

Role required: admin or service_mapping_admin

About this task

In base systems, which are the default or standard configurations, traffic-based discovery relies solely on TCP-related data collected using the *netstat*, *ss*, and *lsof* commands. Discovery based on Netflow and VPC logs requires additional configuration. You can enrich your traffic-based discovery by configuring Service Mapping to use the Netflow protocol. For more information about the way Service Mapping to collect Netflow data, see Data collection and discovery using Netflow.

For testing purposes, install the Netflow Collector (nfdump) on a Unix server inside your organization. In this case, this Unix server should be different from the server hosting the MID Server server.

Configure the ServiceNow connector to trigger MID Server to collect the data from the flow log and processes it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Download and install the Netflow collector (nfdump) on a Unix or Ubuntu server inside your organization.
   - For a Linux server, download, compile, and install the nfdump package. You can download the nfdump package from https://sourceforge.net/projects/nfdump/.
   - For an Ubuntu server, install the nfdump package without predownloading or compiling it. Open the command-line window and run the following command:

     ```
     sudo apt-get install nfdump
     ```

   - For an Ubuntu server, if the *apt-get* command fails, predownload the nfdump package, save it locally and then install it. Open the command-line window and run the following commands:

     ```
     sudo dpkg -i nfdump_1.6.15-3_i386.deb -
     ```

     ```
     sudo apt-get -f install
     ```

     Note: The file name for the nfdump package has the following format: nfdump_<version number> .deb. In this example it is nfdump_1.6.15-3_i386.deb.

2. Configure the Netflow collector to save data for one day.

   a. Open the command-line window on the server hosting the Netflow collector.

   b. Create a cron job.
   `crontab -e`

   c. Enter the following command using the correct paths.
   `*/10 * * * * /usr/local/bin/nfexpire -e /data/nfdump -t 1d`

3. Create a file with the nfdump data.
   For example, use the following command:

   nfdump -q -m -R /data/nfdump/ -o extended -t 2016/07/06.07:00:00-2016/07/06.07:10:00 'inet and proto tcp' >> /tmp/my_file

4. If the file is very large, you can compress it using the gzip format by using the following command:
   `gzip /tmp/my_file`
5. Copy the nfdump data file to the MID Server.
6. Configure Service Mapping to receive data collected by the Netflow collector:

   a. Navigate to **Service Mapping** > **Administration** > **Flow Connectors**.

   b. Click **New**.

   c. Click **ndfdump file**.

   d. On the dfdump file page, configure parameters as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	A descriptive name for the connector.
nfdump data path	The path to a location on the MID Server to which you saved the nfdump data file in 5.
MID Server	The MID Server, onto which you copied the nfdump file.
Gzipped file	If you converted the nfdump file into the gzip format before saving it on the MID Server, set this parameter to true to unzip it.

e. Click **Submit**.

7. Verify that Service Mapping collects data using Netflow:
   a. On the **nfdump file** form, select the newly configured connector and click **Run now** to start the data collection flow and populate the Flow Connection [sa_flow_connection] table.
   b. Navigate to **System Definitions** > **Tables**.
   c. Click the Flow Connection [sa_flow_connection] table.
   d. Under **Related Links**, click **Show List**.
   e. Verify that the table contains data.

What to do next
If you are satisfied with the results of the test, configure Netflow-based data collection as described in Configure data collection using Netflow.

Configure data collection using Netflow

Enable Service Mapping to perform discovery based on data collected using the Netflow protocol. This setup results in fully automated data collection flow, where all involved components send, collect, and analyze data automatically.

Before you begin
Learn about Traffic-based discovery in Service Mapping.

Enable the following scheduled job: Flow Discovery Scheduler [sysauto_script_74c676f0dbb0220060ff742eaf9619f2].

Role required: admin or service_mapping_admin

About this task
In base systems, which are the default or standard configurations, traffic-based discovery relies solely on TCP-related data collected using the *netstat*, *ss*, and *lsof* commands. Discovery based on Netflow and VPC logs requires additional configuration. You can enrich your traffic-based discovery by configuring Service Mapping to use the Netflow protocol. For more information about the way Service Mapping uses Netflow, see Data collection and discovery using Netflow.

Configure the ServiceNow Netflow connector to trigger the MID Server to collect the data from the Netflow flows and process them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Install the nfdump package on a server hosting the MID Server in your organization:
   - For a Linux server, download, compile, and install the nfdump package. You can download the nfdump package from https://sourceforge.net/projects/nfdump/.
   - For an Ubuntu server, install the nfdump package without predownloading or compiling it. Open the command-line window and run the following command:

     ```
     sudo apt-get install nfdump
     ```

   - For an Ubuntu server, if the *apt-get* command fails, predownload the nfdump package, save it locally and then install it. Open the command-line window and run the following commands:

     ```
     sudo dpkg -i nfdump_1.6.15-3_i386.deb -

     sudo apt-get -f install
     ```

     Note: The file name for the nfdump package has the following format: nfdump_<version number> .deb. In this example it is nfdump_1.6.15-3_i386.deb.

2. Configure the Netflow collector to save the nfdump file in the required directory.

   a. Open the `/etc/init.d/nfdump` file.

   b. Modify the parameter responsible for saving this file in the required location. For example, on an Ubuntu server, specify the location using the *DEAMON_ARGS* parameter:

   ```
   DATA_BASE_DIR="/var/cache/nfdump"

   DAEMON_ARGS="-D -l $DATA_BASE_DIR -P $PIDFILE"
   ```

   For operational information, refer to https://sourceforge.net/projects/nfdump/.

3. Configure the switches to forward their nfdump files to the MID Server. The default value for the MID Server is port 9995.
4. Configure the Netflow collector to save data for one day.

   a. Open the command-line window on the server hosting the Netflow collector.

   b. Create a cron job.
   `crontab -e`

   c. Enter the following command using the correct paths.
   `*/10 * * * * /usr/local/bin/nfexpire -e /data/nfdump -t 1d`

5. Verify that the Netflow collector is configured correctly and receives the correct data from the network resources.

   a. Run the following command:

   ```
   nfdump -q -O tstart -R /data/nfdump/ -o extended
   ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. In the command output, verify that marked fields contain real data:

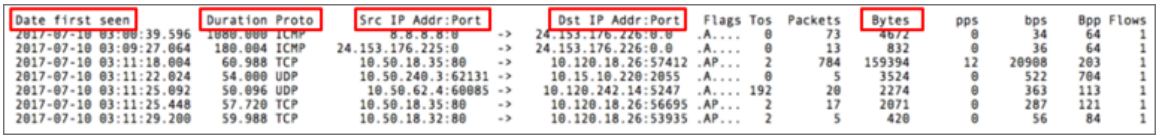

Date first seen	Duration	Proto	Src IP Addr:Port		Dst IP Addr:Port	Flags	Tos	Packets	Bytes	pps	bps	Bpp	Flows
2017-07-10 03:08:39.596	1080.000	ICMP	0.0.0.0:0	->	24.153.176.226:0.0	.A....	0	73	4672	0	34	64	1
2017-07-10 03:09:27.064	180.004	ICMP	24.153.176.225:0	->	24.153.176.226:0.0	.A....	0	13	832	0	36	64	1
2017-07-10 03:11:18.004	60.988	TCP	10.50.18.35:80	->	10.120.18.26:57412	.AP...	2	784	159394	12	20908	203	1
2017-07-10 03:11:22.024	54.000	UDP	10.50.240.3:62131	->	10.15.10.220:2055	.A....	0	5	3524	0	522	704	1
2017-07-10 03:11:25.092	50.096	UDP	10.50.62.4:60085	->	10.120.242.14:5247	.A....	192	20	2274	0	363	113	1
2017-07-10 03:11:25.448	57.720	TCP	10.50.18.35:80	->	10.120.18.26:56695	.AP...	2	17	2071	0	287	121	1
2017-07-10 03:11:29.200	59.988	TCP	10.50.18.32:80	->	10.120.18.26:53935	.AP...	2	5	420	0	56	84	1

6. Configure Service Mapping to receive data collected by the Netflow collector:

a. Navigate to **Service Mapping** > **Administration** > **Flow Connectors**.

b. Click **New**.

c. Click **nfdump install**.

d. On the nfdump install page, configure parameters as follows:

Field	Description
Name	A descriptive name for the connector.
MID Server	The MID Server on which you installed the Netflow collector.
nfdump data directory	The data directory where you configured the Netflow collector to save the nfdump files.

e. Click **Submit**.

7. Verify that Service Mapping collects data using Netflow:

a. On the **nfdump install** form, select the newly configured connector and click **Run now** to start the data collection flow and populate the Flow Connection [sa_flow_connection] table.

b. Navigate to **System Definitions** > **Tables**.

c. Click the Flow Connection [sa_flow_connection] table.

d. Under **Related Links**, click **Show List**.

e. Verify that the table contains data.

Data collection and discovery using VPC Flow Logs

Service Mapping can perform discovery based on data collected using VPC Flow Logs. Amazon VPC hosts Amazon Elastic Compute Cloud (EC2) instances that provide Amazon Web Services. VPC Flow Logs collect data on IP traffic going to and from network interfaces in the VPC.

In base systems, which are the default or standard configurations, traffic-based discovery relies solely on TCP-related data collected using the *netstat*, *ss*, and *lsof* commands. Discovery based on Netflow and VPC logs requires additional configuration. You can enrich your traffic-based discovery by configuring Service Mapping to use VPC Flow Logs.

Service Mapping discovery based on VPC Flow logs has the following flow:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Amazon EC2 instances collect their individual logs into log streams and forward them to the central flow log group.

Amazon EC2 instances collect their individual logs

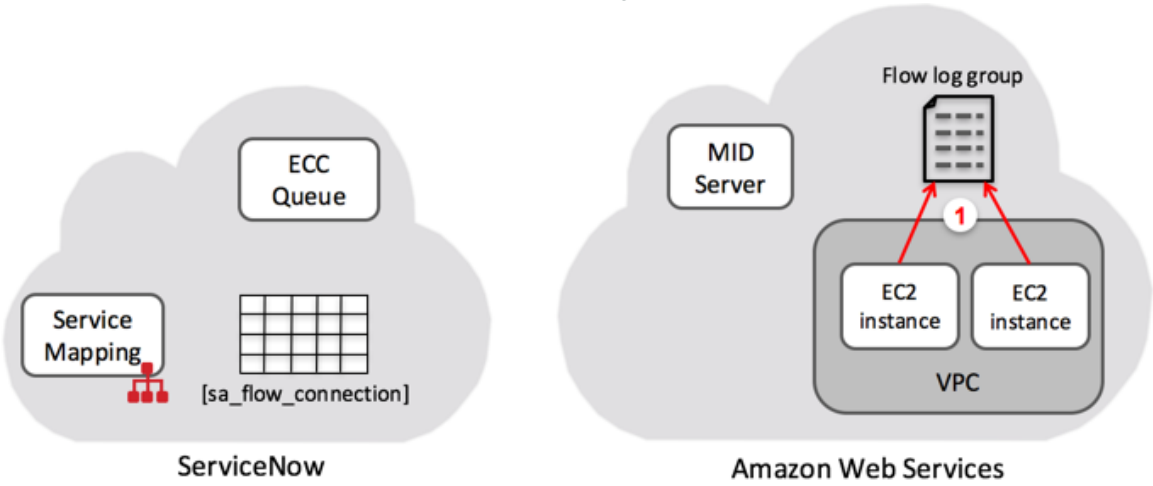

2. The ServiceNow connector triggers MID Server to collect the data from the flow log and processes it.
3. The MID Server places the processed information onto the ECC queue.

MID Server collects the data from the flow log and places it onto the ECC queue

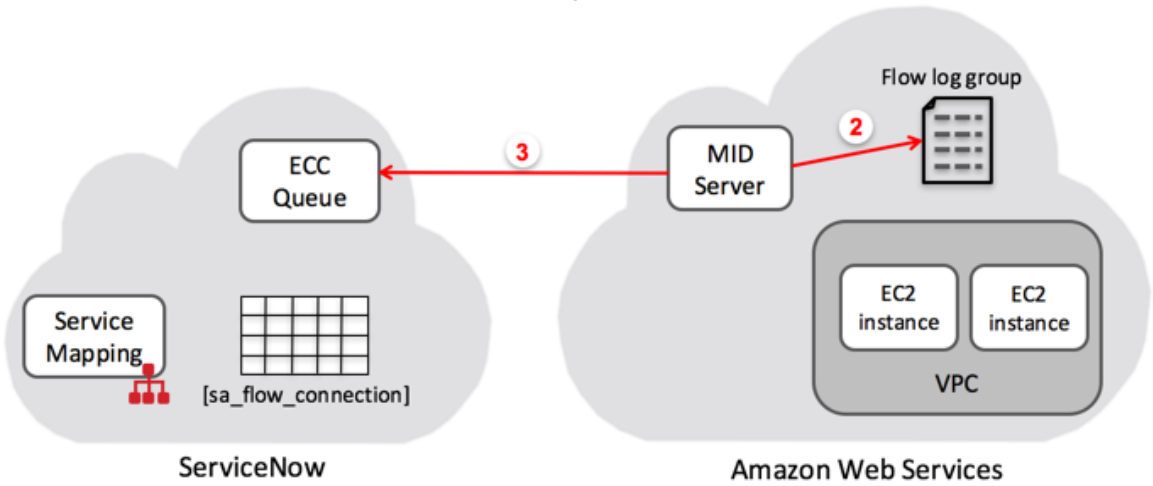

4. A sensor retrieves the processes data from the ECC queue and writes it into the Flow Connection [sa_flow_connection] table.

5. Whenever Service Mapping checks the ECC queue and receives information on a discovered CI, it checks these tables for any data on outbound connections related to the CI: the cmdb_tcp and sa_flow_connection tables. If these two tables contain unique data that patterns

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

did not discover, Service Mapping enriches the information about the CI connections and adds them to the map.

Collected data is written into the sa_flow_connection table from which Service Mapping collects the data

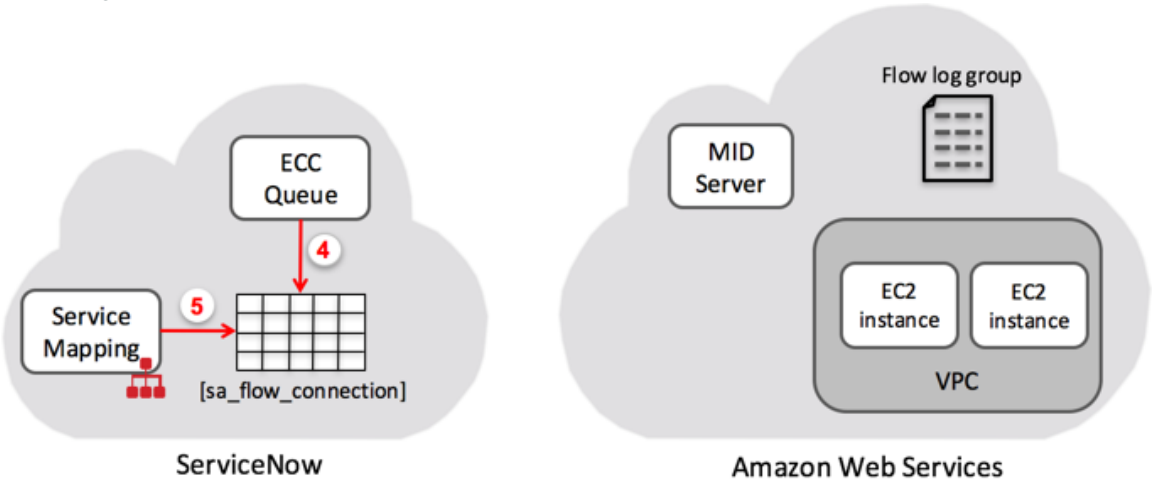

In deployments with multiple flow log groups, configure a dedicated connector that works with one MID Server for every flow log group. Multiple flow log groups my use the same AWS credentials.

Related topics

https://support.servicenow.com/kb_view.do?sysparm_article=KB0597467

Traffic-based discovery in Service Mapping

Configure data collection using VPC Flow Logs

Enable Service Mapping to perform discovery based on data collected using Virtual Private Cloud (VPC) logs. This method is relevant for organizations using Amazon Web Services (AWS).

Before you begin

Verify that credentials for the AWS account are configured in the Credentials module of the platform by navigating to **Service Mapping** > **Credentials** > **AWS Credentials**. For more information about AWS credentials, see Cloud credentials.

Role required: admin or service_mapping_admin

About this task

In base systems, which are the default or standard configurations, traffic-based discovery relies solely on TCP-related data collected using the *netstat*, *ss*, and *lsof* commands. Discovery based on Netflow and VPC logs requires additional configuration. You can enrich your traffic-based discovery by configuring Service Mapping to use VPC Flow Logs. For more information about the Service Mapping discovery flow based on VPC Flow logs, see Data collection and discovery using VPC Flow Logs.

Amazon VPC hosts Amazon Elastic Compute Cloud (EC2) instances that provide Amazon Web Services. VPC flow logs collect data on IP traffic going to and from network interfaces in the VPC.

Configure the ServiceNow connector to trigger MID Server to collect the data from the flow log and processes it. In deployments with multiple flow log groups, configure a dedicated connector that works with one MID Server for every flow log group. Multiple flow log groups my use the same AWS credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Configure VPC Flow Logs on the Amazon EC2 console as described in the official AWS documentation.
2. Configure Service Mapping to work with VPC Flow Logs:
   a. Navigate to **Service Mapping** > **Administration** > **Flow Connectors**.
   b. Click **New**.
   c. Click **AWS VPC flow logs**.
   d. On the AWS VPC flow logs page, configure the connector parameters as follows:

Field	Description
Name	A descriptive name for the connector.
Group name	The name of the central flow log group to which Amazon EC2 instances forward their log streams.
MID Server	The MID Server that Service Mapping uses to collect the data from the flow log group on AWS.
AWS Credentials	Select AWS credentials relevant for the account for which you want to collect flow logs.

   e. Click **Submit**.
3. Verify that Service Mapping collects data using VPC Flow Logs:
   a. On the **AWS VPC flow logs** form, select the newly configured connector and click **Run now** to start the data collection flow and populate the Flow Connection [sa_flow_connection] table.
   b. Navigate to **System Definitions** > **Tables**.
   c. Click the Flow Connection [sa_flow_connection] table.
   d. Under **Related Links**, click **Show List**.
   e. Verify that the table contains data.

Configure Search Assistant for Windows

The Search Assistant feature of Pattern Designer allows you to search within files or registries. Upload grep files on to your instance to enable this feature to search on Windows servers.

Before you begin

Role required: admin or service_mapping_admin

About this task

Grep is a utility that the Search Assistant feature uses to search files for matches. Grep is native to Unix-based systems, therefore you must upload it only for Windows-based configuration items (CI). Once you upload grep files onto the instance, it transfers them onto all MID Server which distribute the files onto all Windows-based CIs in your organization.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Download the following files onto the computer from the Internet.
   - egrep.exe
   - libiconv2.dll
   - libintl3.dll
   - pcre3.dll
   - regex2.dll

   For example, you can download the zip files for the **Dependencies** and **Binaries** folders at the following location: http://gnuwin32.sourceforge.net/packages/grep.htm
2. Extract the files.

   After you extract the downloaded folders, the files are located as follows:
   - grep-2.5.4-bin/bin/egrep.exe
   - grep-2.5.4-dep/bin/libiconv2.dll
   - grep-2.5.4-dep/bin/libintl3.dll
   - grep-2.5.4-dep/bin/pcre3.dll
   - grep-2.5.4-dep/bin/regex2.dll
3. Navigate to **Pattern Designer** > **Uploaded Files**.
4. Click **New**.
5. Enter the original file name in the **Name** field on the **Logical Name** tab, for example `egrep.exe`.
6. Click **OS Types** and select **Windows**.

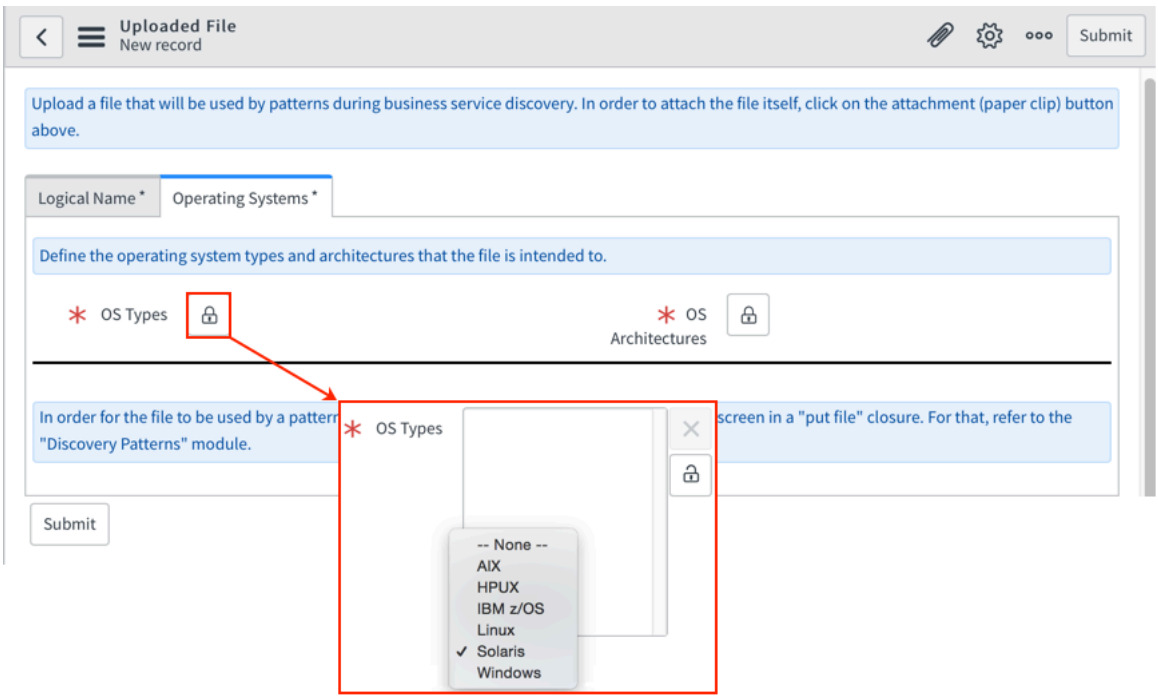

7. Click **OS Architectures** and select both 32-bit and 64-bit option.
8. Click the **Manage Attachments** icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. Click **Choose Files**.
10. In the **Downloads** folder on your computer, navigate to the file which name you entered before, for example `egrep.exe`.
11. Click **Choose**.
12. Close the Attachments window.
13. Repeat Step 3 to Step 11 to upload the rest of the files:
    - libiconv2.dll
    - libintl3.dll
    - pcre3.dll
    - regex2.dll
14. Connect to each Windows MID Server and update the PATH environmental variable with the location of the `bin` directory where `egrep.exe` is installed.

Modify display for CI attributes

You can control what configuration item (CI) attributes the system displays in the Properties pane of service instance maps.

Before you begin

Get familiar with the notion of views. For more information on views, see View Management.

Decide for which CI classes you need to modify display.

Decide whether you want to modify display for parent CI classes to see the view change in all their child CI classes. To check the parent CI class of a CI class, perform the steps below:

Note: The Service Mapping user interface refers to CI classes as CI types.

1. Navigate to **System Definition** > **Tables**.
2. Set the search field to **Name**, and enter the relevant CI class.

   The list of all CI class that match the name is displayed.
3. Check the **Extends table** column to see the parent CI class.

Role required: service_mapping_admin or app_service_admin

About this task

When working with service instance maps, the Properties pane displays attributes for CIs and connections you select in a map. Since there is a large number of attributes to display, you may want to choose which attributes are displayed, and in what order.

The attributes and their order of display is determined by a combination of the CI type definition and the view for this CI type.

CMDB stores CI class in the form of a hierarchy where some CI classes are parents to other CI classes, who automatically derive their parent's attributes in addition to attributes you configure specifically for child CI classes. CI class hierarchy is used widely for configuring CI behavior, relationships, and display.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Child CI classes derive attributes from their parents

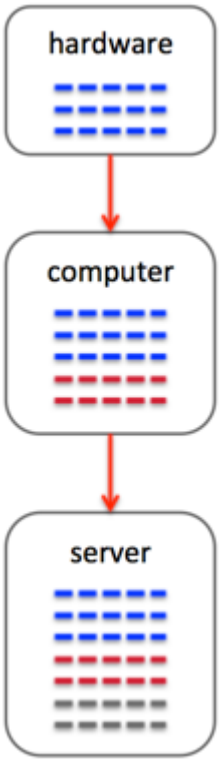

The view that controls the attribute display in service instance maps is called *sa_map_properties*. You can assign this view to new CI classes or modify the view for CI classes which already use it. Use the CI type hierarchy to configure the *sa_map_properties view*.

- The parent CI class at the top of the hierarchy is *cmdb*.
- If you do not define the view for a child CI class, the system displays attributes derived from the parent with the addition to all the attributes of this child CI class.

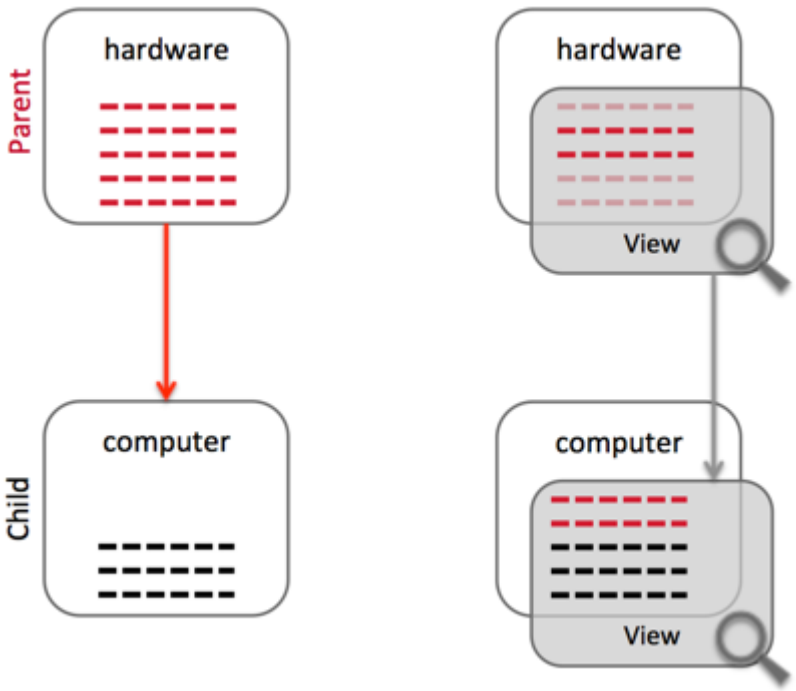

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you define a view for a child CI class, it includes only the attributes that you specifically added to the child CI class. The system does not display any of its parent attributes automatically.

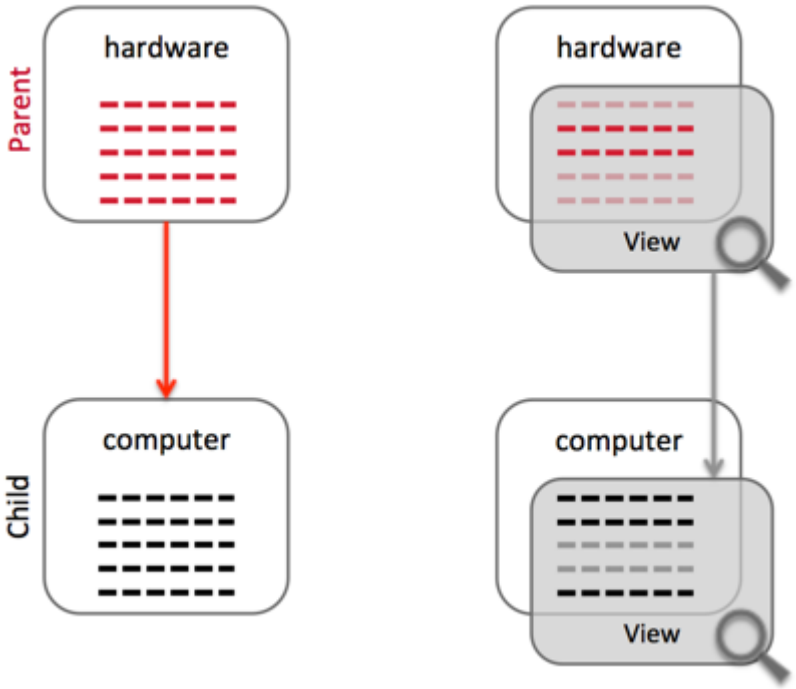

- If you do not add a parent CI class or its child CI class to the view, the system uses the ascendant CI class that was added to the view in addition to all attributes of all parent CI

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

classes in the hierarchy between this child CI classes and the parent CI class that is added to the view.

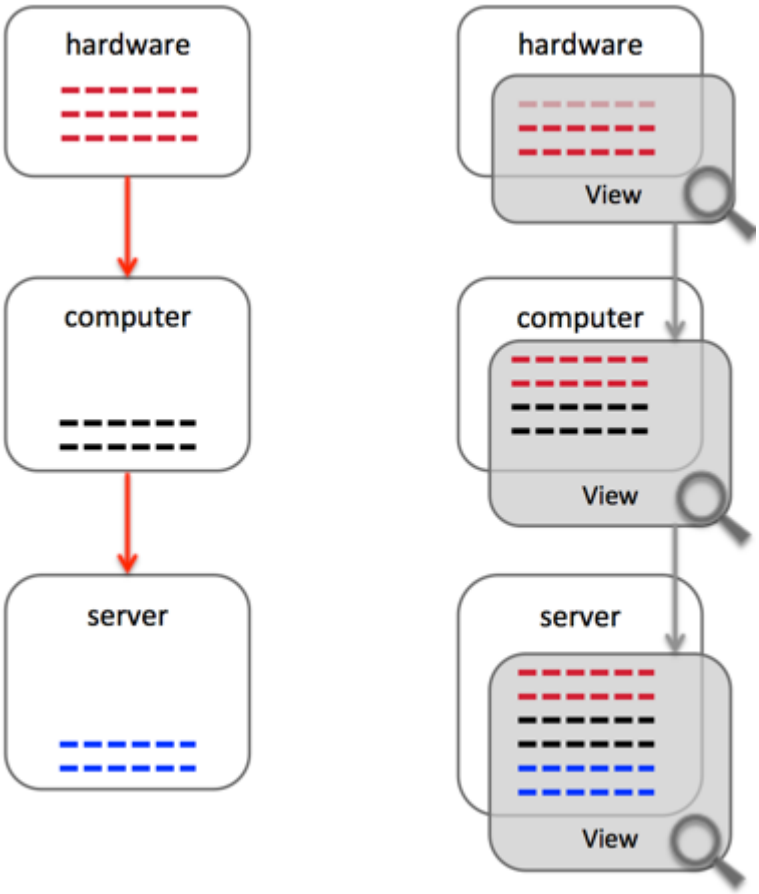

For example, if you define the *sa_map_properties* view for the CI class for Windows Servers, the Properties pane displays the same attributes for all Windows Servers: CI type (class), Class, Model ID, and name.

CIs belonging to the same class

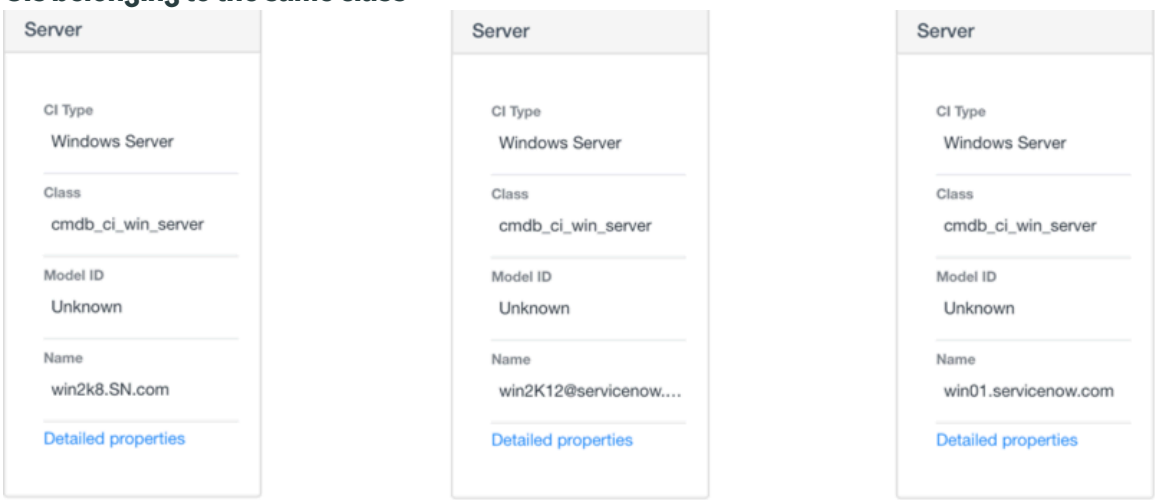

You can modify the view at any time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. In the service instance map, select the CI whose attributes you want to change.
3. In the **Properties** pane, click **Detailed Properties**.

4. In the CI Detailed Properties page, click the **Additional actions** icon and select **Configure > Form Layout**.
5. In the **View name** list under **For view and section**, define which view you want to modify:

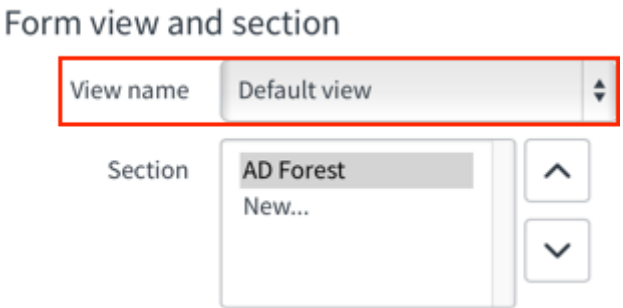

   - Select **sa_map_properties** if it appears in the list, or
   - If **sa_map_properties** does not appear in the list:
     a. Select **New**.
     b. In the **View name** field, enter `sa_map_properties`.
     c. Click **OK**.
6. Use the slushbucket to add or remove attributes from the view.
7. Click **Save**.
8. Return to the service instance map and reselect the CI for which you changed the view.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Properties** pane displays the attributes that you configured for the view.

9. If necessary, repeat the procedure for other CI classes.

Related topics

View CI attributes in an application service map

Configure support for Windows servers with non-English OS

You can configure your ServiceNow AI Platform to support Windows servers that use non-English Windows operating system (OS).

Before you begin

You can perform this procedure only on the MID Servers that serve locations using non-English operating system. Alternatively, you can perform this configuration on all MID Servers in one go, even MID Servers that are used for locations using English operating systems.

Role required: mid_server

About this task

If your organization does not deploy Service Mapping, this configuration is not necessary.

Sometimes organizations choose to use operating systems in local languages. While having a localized operating system (OS) may be user-friendly, it poses a problem when it comes to automatically discovering Windows servers running the localized OS. During the discovery and mapping process Service Mapping sends discovery commands to Windows servers on your infrastructure via a MID Server. If a Windows server returns a message in a language other than English, Service Mapping does not recognize it as a valid response and the interaction fails.

If your organization uses non-English operating system for some of the Windows servers, you must perform this procedure to make sure that Service Mapping can access all Windows servers, not exclusively Windows servers using the English operating system. This configuration allows the MID Server, that is located between Service Mapping and Windows servers, to recognize a non-English response from a Windows server and to change the language of the CI operating system into English.

This configuration affects only the user for which you perform this procedure.

Procedure

1. Navigate to **All** > **MID Servers** > **Properties** in the **Modules** pane.
2. Click **New**.
3. In the **Name** field, enter `mid.servicewatch.wmi.mui`.
4. In the **Value** field, enter `true`.
5. In the **MID server** field, select the server to which you want to apply this change.

Option	Action
To apply the change to all MID Servers in your organization	Leave the MID Server field blank.
To apply the change to a specific MID Server	Select the relevant MID Server from the list.

6. Click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Resolve pattern-related mapping errors

MID Server

Upload the rctrlx.exe file to MID Servers

Upload the rctrlx.exe file to MID Servers to enable running discovery commands on Microsoft Exchange 2007 and 2010, and Citrix XenApp.

Before you begin

Role required: admin or service_mapping_admin

About this task

During the discovery process, the MID Server runs Service Mapping and Discovery patterns on applications and devices located inside the organization network. A pattern is a sequence of commands designed to detect attributes of a CI and its outbound connections. To execute commands over PowerShell correctly, MID Server runs patterns as services using the rctrlx.exe file. The rctrlx.exe file is part of the open source ManagePC utility and is available for downloading from the Internet.

Upload an `rctrlx.exe` file for every MID Server in your deployment.

Procedure

1. Download the `rctrlx.exe` file onto your computer from the Internet.

   For example, https://github.com/leonsodhi/rctrlx/releases offers the latest versions of the `rctrlx.exe` file:

2. Place the rctrlx.exe file into the folder for Windows 32-bit servers:

   ```
   <MID Server installation directory>\agent\bin\sw_wmi\bin\32
   ```

   For example, C:\SN_MID\SN_MID_Dev\agent\bin\sw_wmi\bin\32.

3. Place the rctrlx.exe file into the folder for Windows 64-bit servers:

   ```
   <MID Server installation directory>\agent\bin\sw_wmi\bin\64
   ```

   For example, C:\SN_MID\SN_MID_Dev\agent\bin\sw_wmi\bin\64.

Related topics

Pattern-based discovery in Service Mapping

MID Server

Fine-tune Service Mapping with MID affinity and IP reuse

The MID Server needed to map an application service might be misidentified in subnetworks with overlapping IP ranges and reused IP addresses. To enable the successful discovery of resources associated with a subnetwork, Service Mapping provides several properties for MID Server identification.

The networks of large organizations are often divided into numerous subnetworks, each acting as a network of its own. Service Mapping selects the MID Server based on its IP range. Subnetworks with overlapping IP ranges and reuse of IP addresses risk misidentification of the MID Server needed to map an application service. Discovery can therefore fail, resulting in incomplete or inaccurate service maps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Two properties included with Service Mapping offer a solution to system administrators.

Note: For full functionality, this feature requires the installation of CMDB CI Class Models version 1.50.0 or later.

- MID affinity: If horizontal discovery has been configured, you can enable the *sm.topdown.reuse_mid_from_hd_schedule* system property. It identifies the subnetwork MID Server that was previously configured during horizontal discovery and reuses it for top-down discovery in Service Mapping. Configuring the MID Server using this method requires less time and results in fewer misidentifications.

  Note: For subnet-based MID affinity to take effect, an executed horizontal Discovery schedule configured to use the IP range targeted by top-down Service Mapping must exist. If no matching horizontal Discovery history exists for the IP range, affinity will not apply.

- Network location: After locations have been assigned to the horizontal discovery schedules, you can enable the *sm.network_location.detect_host_by_location* system property, which adds a **Network Location** field to the Entry point form. The top-down discovery process starts with the entry point defined on this form and can then identify both the MID Server and the host using the location provided in the horizontal discovery schedule.

Legacy MID Server selection: Verify that *sa.mapping.legacy_mid_selection* is set to false. When set to true, Service Mapping uses the legacy MID Server selection algorithm and the MID affinity properties have no effect, regardless of their configuration.

To learn more about enabling each property, see Components installed with Service Mapping

Related topics

Schedule a horizontal discovery

MID Server configuration for Service Mapping

Configure Service Mapping to work with the Configuration Management Database (CMDB)

Configure Service Mapping to use the existing data within your CMDB to run top-down discovery and map application services.

Before you begin

Verify that the necessary tables are populated with data according to the process explained in CMDB-based mapping.

- Transmission Control Protocol (TCP) Connection [cmdb_tcp] table
- Running Process [cmdb_running_process] table
- Hardware [cmdb_ci_hardware] table

Confirm that you have installed the Service Mapping Plus August 2024 version from the ServiceNow® Store.

Role required: service_mapping_admin

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. In the System Properties [sys_properties] table, search for the *sa.topdown.enable_cmdb_based_mapping* property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Set the *sa.topdown.enable_cmdb_based_mapping* system property to configure Service Mapping to use the data in your CMDB.
   a. Confirm that the **Type** field value is set to true|false.
   b. Set the **Value** field to `true`.

Configure Agent Client Collector for use with Service Mapping

Configure Agent Client Collector to discover and map application services without having to configure credentials for the MID Server.

Before you begin

Ensure that Agent Client Collector has been configured. For more information, see Application patterns for the Agent Client Collector.

Verify that agents have the necessary privileges to execute the required application pattern commands for successful completion of top-down discovery. For more information, see Privileged SSH commands for probe-based discovery and Non-privileged SSH commands during probe-based discovery.

Important: This configuration requires a MID Server that is up and running to work correctly. Service Mapping selects the MID Server to which the target agent connects, in order to run commands through the agent.

Role required: service_mapping_admin

About this task

Agent Client Collector enables the discovery and mapping of application services without requiring you to set up credentials for the MID Server. You can deploy both top-down and machine learning-based Service Mapping and benefit from a simplified approach to mapping application services.

Procedure

1. Enter `sys_properties.list` in the navigation filter.
2. In the System Properties [sys_properties] table, search for the *sa.topdown.mapping_with_acc_mode* property.
3. Set the property to configure Agent Client Collector for Service Mapping and define the MID Server selection mode for top-down discovery.
   a. Confirm that the **Type** field value is set to integer.
   b. Determine how top-down discovery should run by setting the **Value** field value.
   Possible values are:
   - `0` (the default) - Top-down discovery runs using the MID Server exclusively.
   - `1` - Top-down discovery initially attempts to run using Agent Client Collector. If no agent is available, then it defaults to using the MID Server.
   - `2` - Top-down discovery runs using Agent Client Collector exclusively. If unsuccessful, the process terminates.

Convert a service instance to Lightweight Service Model

Convert a Service Mapping service instance to Lightweight to minimize storage footprint and optimize performance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that the Service Mapping Plus store app is up to date.

Role required: service_mapping_admin

About this task

When you convert a service instance to Lightweight, the system deletes the existing historical snapshot data and reconfigures the service to operate under the Lightweight architecture. The service continues to calculate membership automatically, but only current state data is retained going forward.

Note: Converting a service instance to Lightweight is permanent and irreversible.

Supported service types for conversion:

- Dynamic
- Tag-based

For detailed information on the Lightweight Service Model, see Service Mapping Lightweight Service Model.

Procedure

1. In the navigation filter, enter: `cmdb_ci_service_calculated_list.do`.
2. In the **Calculated Application Services** list, select a service instance of type Dynamic or Tag-Based.
3. On the service record, locate the **Related Links** section and select the **Convert to Lightweight** link.

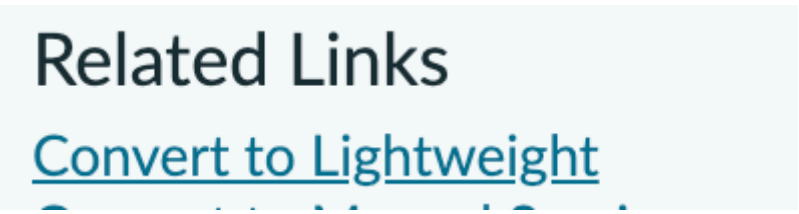

4. In the confirmation dialog, review the warning about deleting historical data.
   a. If you need to preserve historical snapshots, select **Cancel** and complete any required analysis before proceeding.
   b. If you're ready to proceed, select **Convert**.

   A notification appears: "Converting service to Lightweight. This may take a few minutes." Keep the page open during the conversion process. Don't navigate away or close the page until the conversion is complete.

Result

The service operates under the Lightweight Service Model.

Using Service Mapping

As the Service Mapping administrator, you support companies in their efforts to map and maintain critical application services so that they remain service-aware.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a service instance from various data sources

Search for mapped and unmapped services to unify. The flow is available using the Service Mapping workspace.

About this task

Multi-source service mapping

Before you begin

You must have at least Australia platform version installed.

You must have the latest version of Service Mapping Plus.

At least two existing services to merge, or at least two sources of data that are actively updated.

Role required: service_mapping_admin

Procedure

1. Navigate to **Workspaces** > **Service Mapping**.
2. Select **Create service instance**
3. Add service details.

   a. Fill in the basic details form.

   Basic details

Field	Description
Service instance name	Enter a descriptive name for this service instance. For example: Payment processing. You must enter a value to proceed.
Support group	Select the group responsible for supporting this service instance. For example: Application support. You must select a value to proceed.
Service instance owner	Select the person accountable for this service instance.
Short description	Provide a brief summary of what this service instance does. For example: Unified service containing tag-based services and ML-suggestion-based services for payment processes.

   b. *Expand the Additional fields section and customize additional fields as needed. For example, add production environment or operational status.*

   Field availability is configured by your administrator.

   c. Select **Next**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Find and select services to unify.

   To add mapped and unmapped services to your unified map, add criteria in the **Search criteria** section. You can search for services by entering different types of criteria.

   Note: A maximum of ten services can be unified by default (configurable via the sn_sm_scoped_app.sm.unified_services.max_children_per_service system property).

   a. Enter a service name to add a service for unification.

   b. Enter a server name to add a server for unification.

   c. Enter a tag key to add a resource for unification.
   You can add as many tags as you need. You can add multiple tag values per each tag key.

   d. Enter a URL or multiple URLs separated by commas.
   You can add as many URLs as needed.

   e. Select **Find matching services**.

      Services that match the criteria are listed in two ways:

      - Top recommendations: The top 50 services that match the criteria, including unmapped services.
      - All matching services: All mapped services that match the criteria.

   f. Select the services to unify using the check boxes.

      While selecting the services to unify, details and preview maps are available for unmapped and mapped services, using the info icon. A list of configuration items (CI) is displayed for review before the final step of creating the unified service. The CIs origins are mentioned in the **Parents** field.

      Note: Map preview is not supported for unmapped tag-based services.

   g. Select **Next**.

5. Enrich your service by linking it to its business context.

   Provide the business and operational details that connect this service to the rest of your organization. These details help drive visibility and reporting.

   a. Set relationships and connect this service to its CSDM components.

      All three fields are optional. Leaving any or all fields empty is valid and does not prevent you from proceeding to the next step.

      CSDM fields

Field	Description
Business application	Search for and select a business application. This field allows you to link your

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	service instance to the business application it supports. Single selection only. If you type text without selecting from the list, an error state appears (red border). The **Next** button is disabled until you select a valid option or clear the field.
Parent service	Search for and select a parent application service. Use this field to establish hierarchical relationships and build logical context for your service instance. Single selection only. If you type text without selecting from the list, an error state appears (red border). The **Next** button is disabled until you select a valid option or clear the field.
Business service offerings	Business service offerings represent specific workflows or solutions that are delivered through business services. Select one or more business service offerings. Multiple selections allowed. Selected offerings appear as removable tags.

b. Select **Next**.

6. Review the service details and select **Create**.

   When the process of creating the unified service is complete, the unified service map is displayed in the "Unified map" format. You can review the map using the CMDB Workspace or the Service Graph Workspace, showing the complete consolidated view. It can take a few minutes for the map to load.

   Note: The unified service is created as non-operational by default. All selected unmapped services convert into non-operational mapped services.

Related topics

Edit a unified service

Delete a unified service

Edit a unified service

Add or remove services to update a unified service. The flow is available using the Service Mapping workspace.

About this task

Multi-source service mapping

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

You must have at least Australia platform version installed.

You must have the latest version of Service Mapping Plus.

Create a service instance from various data sources

Role required: service_mapping_admin

Procedure

1. Navigate to **Workspaces** > **Service Mapping**.
2. Select the **Mapped application services** widget.
3. Select a service of type Unified, and then select **Edit**
4. Edit the unified service details if needed, and select **Next**.
5. To add mapped and unmapped services to your unified map, add criteria in the **Search criteria** section.

   a. Enter a service name to add a service for unification.

   b. Enter a server name to add a server for unification.

   c. Enter a URL or multiple URLs separated by commas.

   d. Enter a tag key to add a resource for unification.

   e. Select **Find matching services** to add services by the matching criteria.

   f. Review the service details, using the info icon, and select the check box to add it to the unified service.

   g. Review the "Selected services" panel and remove service instances if needed.
   The "Delete service" button is for removing the entire unified service.

   h. Select **Next** when your list of services in the "Selected services" panel is final.

6. Enrich your service by linking it to its business context.

   Provide the business and operational details that connect this service to the rest of your organization. These details help drive visibility and reporting.

   a. Set relationships and connect this service to its CSDM components.

      All three fields are optional. Leaving any or all fields empty is valid and does not prevent you from proceeding to the next step.

CSDM fields

Field	Description
Business application	Search for and select a business application. This field allows you to link your

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	service instance to the business application it supports. Single selection only. If you type text without selecting from the list, an error state appears (red border). The **Next** button is disabled until you select a valid option or clear the field.
Parent service	Search for and select a parent application service. Use this field to establish hierarchical relationships and build logical context for your service instance. Single selection only. If you type text without selecting from the list, an error state appears (red border). The **Next** button is disabled until you select a valid option or clear the field.
Business service offerings	Business service offerings represent specific workflows or solutions that are delivered through business services. Select one or more business service offerings. Multiple selections allowed. Selected offerings appear as removable tags.

b. Select **Next**.

7. Review the service details and select **Update**.

Your unified service is updated.

In the process of editing the unified service, the list of configuration items (CI) is updated. After completion, a map refresh is needed to present the updated unified service.

Related topics

Delete a unified service

Delete a unified service

Delete a unified service using the Service Mapping workspace.

About this task

Multi-source service mapping

Before you begin

You must have at least Australia platform version installed.

You must have the latest version of Service Mapping Plus.

Create a service instance from various data sources

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: service_mapping_admin

Procedure

1. Navigate to **Workspaces** > **Service Mapping**.
2. Select the **Mapped application services** widget.
3. Select the service, and then select **Edit**
   You can group by service types to find the service more efficiently. For example, if you want to remove a unified service, grouping by service type can help you to narrow the list of unified services.
4. Select **Delete service**.

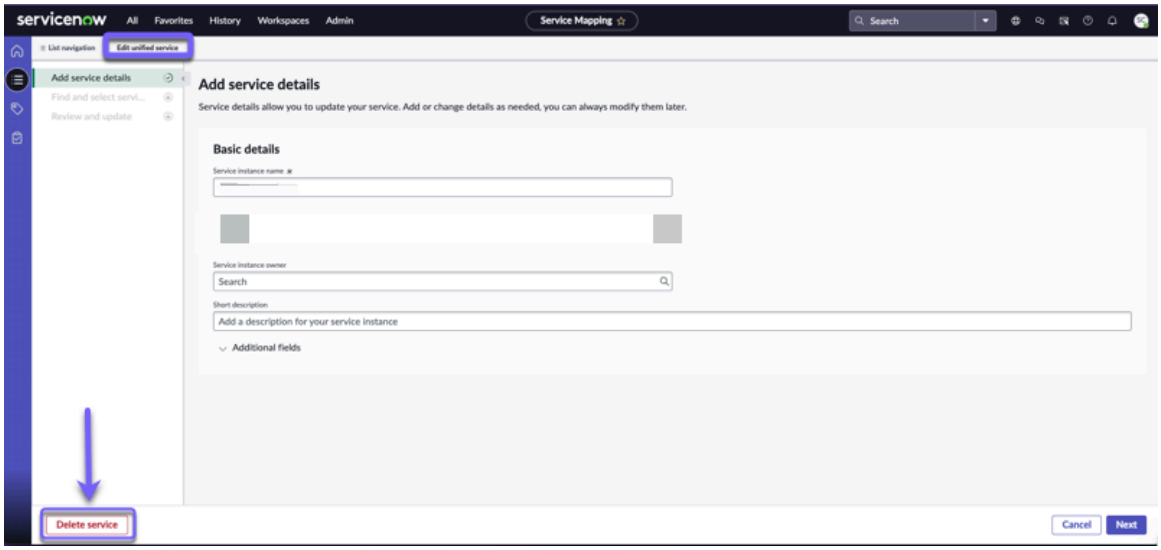

Result

The service is deleted. The UI is redirected to the home page. Child services remain non-operational.

Related topics

Edit a unified service

Create an application service for unmapped servers

Maximize the use of your organization's resources by mapping application service candidates that include unmapped servers. Use the Service Mapping workspace's unmapped servers widget to create an application service and ensure that your servers are used efficiently.

Before you begin

- Ensure that Predictive Intelligence is enabled.
- Verify Machine Learning readiness.
- Confirm that the data is trained.

For more information, see Application service readiness dashboard in configurable workspace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

Service Mapping Plus v.1.10.0 or higher:

To update your data and view application candidates for unmapped servers, see Troubleshooting steps for Mapping unmapped servers with Application service candidates [KB1585065] in the Now Support Knowledge Base.

Role required: service_mapping_admin

Procedure

1. Navigate to **Workspaces** > **Service Mapping**.
2. Select the unmapped server that you want to use to create an application service.

<table>
<tr><th>Version of Service Mapping Plus</th><th>Actions</th></tr>
<tr><td>Versions released before version 1.16.3</td><td>
a. In the Unmapped servers tile, select the Unmapped servers with Candidate ID graph.
b. Review the list of suggested candidates.
c. Organize the candidate list to prioritize your work.
<table>
<tr><th>Organization method</th><th>Action</th></tr>
<tr><td>Order</td><td>Select a column name to sort the results in ascending or descending order.</td></tr>
<tr><td>Group</td><td>Select the menu (⋮) icon to group results by the selected column.</td></tr>
</table>
d. Select the preferred candidate.
e. Select Map application service.
</td></tr>
<tr><td>Starting with version 1.16.3</td><td>
a. In the Unmapped servers tile, select the Unmapped servers with relevant candidate graph.
Note: All application service candidates have been filtered to exclude irrelevant tools, so only the most relevant candidates are presented.
b. Review the list of servers with their associated candidates.
</td></tr>
</table>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Version of Service Mapping Plus	Actions
	c. Organize the candidate list to prioritize your work. **Organization method** / **Action** Order: Select a column name to sort the results in ascending or descending order. Group: Select the menu (⋮) icon to group results by the selected column. **d.** Select the preferred server or candidate. **e.** Select **Map application service**.

Result
The selected application service candidate is converted to an application service, and the selected unmapped servers are mapped.

Map application services based on Automated Service Suggestions

Map a new application service based on automatically generated suggestions. Use the Automated Service Suggestions or Service Fingerprints features to create a new application service or add to an existing service in a few clicks.

Before you begin

- Service Mapping Plus must be installed. For more information, see Install Service Mapping Plus.
- Confirm that the application is ready to discover and map application services. See Application service readiness dashboard in configurable workspace.

Role required: service_mapping_admin

About this task
In addition to being provided with automatically generated suggestions, you can also identify and choose an application service by using Service Fingerprints suggestions that include information about the components in the application service candidates. For more information, see Automated Service Suggestions.

Procedure

1. Navigate to **Workspaces** > **Service Mapping**
2. Select **ML powered candidates**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Starting with Service Mapping Plus version 1.16.3, the filtering process has been updated to include only the most relevant candidates, excluding any ignored or irrelevant ones. This enhancement requires the installation of ITOM Content Service.

3. Select an application service either based on candidate suggestions or identify services with desired component.

Search method	Actions
Select from candidates based on a desired component	**a.** Select the Service Fingerprints tab. **b.** Review the list of Service Fingerprints sug gestions that include components. **c.** Select your preferred application service candidate.
Select the candidate using basic details	**a.** Select the Application Service Suggestions tab. **b.** Review the list of suggested application ser vices. **c.** Select your preferred candidate.

4. **Optional:** Preview the map.

   a. Select **Preview map**.

   b. Review the components of the proposed application service map.

   Note: This preview map might appear different from the candidate because it doesn't show absent connections. For more information, see the Automated Service Suggestions Troubleshooting [KB1205971] article in the ServiceNow Support Knowledge Base.

   c. Select **List navigation** to return to the candidate list.

5. Select **Map application service**.
6. Either create a new application service with the suggested candidate or add the suggested candidate to an existing application service.

Action	St
Create a new application service	**a.** Select **Cre ate new ser vice**. **b.** On the form, fill in

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	St
	the fields. **Map application service form** Name Pre-populated, unique name suggested for the new application service. Use this name. enter a unique name, or select **Other name suggestions** for additional options. For more information, see Name suggestions for application service candidates. Number Pre-populated unique

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	St
	ID for the application service.
	Description of the new application service.
	The service owner who approves of or rejects any change to the mapped application service.
	Service group application services group to which you want to add this service. The default All setting means the application service is available

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	St
Add a candidate to an existing service	to all groups. **a.** Select **Add to existing service**. **b.** Begin entering the name of the existing service and select the desired service from the list.

7. Select **Map service**.
8. **Optional:** Select **View map** to view the new or modified application service.

Related topics

Automated Service Suggestions

Map your application services using tags in the Service Mapping Workspace

Categorize and organize organization's configuration items and map them into application services using the Tag-based dashboard in the Service Mapping workspace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that you have installed Service Mapping Plus version 1.16.3, which is required to access tag-based mapping in the Service Mapping workspace. For more information, see Install Service Mapping Plus.
- Review and analyze your organization's tag usage, making a list of all tags and their purposes. Use the Key Value [cmdb_key_value] table to see the tags in the Configuration Management Database (CMDB).
- Use Tag Governance to optimize tag application in your organization. See Tag Governance for more information.
- Assign tags to configuration items (CIs) that you want to include in application services.
- Run Discovery to horizontally discover CIs. Discovery populates the CMDB with the CI data, including tags. For more information about discovery tag-based discovery for the Service Mapping workspace, see Tag-based discovery for the Service Mapping Workspace.

Role required: service_mapping_admin

Procedure

1. Navigate to **Workspaces** > **Service Mapping**, then select the Tag-based Service Mapping icon from the navigation pane.
2. Select **Create service maps**.
3. Either select an existing tag category or create a new tag category.

Tag category option	Action
Select an existing tag category	**a.** Start entering the category name in the search field. **b.** Select the tag category from the list. **c.** Select **Define the tag values** and enter val ues for the selected category.
Create a new tag category	**a.** Create a new tag category by selecting **Cre ate a new tag category**. **b.** Enter the tag category name. **c.** Enter values for the tag keys in the **Tags** field.

4. Save the tag category and values by selecting **Save**.

   Note: To save the tag categories in a service family so you can return later to complete the mapping process, select **Save and exit** at any time.

5. Search for potential services matching the tag categories by selecting **Find candidates**.
6. Review the list of tag-based service candidates by selecting **Map selected candidates** when you have chosen your preferred candidates.
7. When prompted, enter a unique name for the tag-based service family.
8. Select **Save and map**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
A tag-based application service and tag-based service family have been created.

Access the Unified Map feature from the Service Mapping Workspace

View a hierarchical map of CIs and the relationships between them by adding access to the Unified Map feature to the Service Mapping Workspace.

Before you begin
Ensure that you have the following:

- Vancouver release or later
- Service Mapping Plus, November 2023 release
- CMDB Workspace, v. 4.0.1 or later.

Role required: service_mapping_admin

About this task
Review Unified Map to learn more about the Unified Map feature.

Note: The Unified Map feature offers limited Service Mapping functionality.

Procedure
1. Navigate to **All** > **System Properties** > **All Properties**.
2. Enter *sn_sm_scoped_app.sa.unified_map.enabled* in the Search field.
3. Confirm that the value of the property is set to true.
4. Navigate to **Workspaces** > *Service Mapping*.

   Note: Access to the Unified Map is only available through the Service Mapping Workspace.

5. Select the Mapped application services widget.
6. Select the check box next to the desired application service.
7. Select the **View map** button.
   A new tab opens displaying the Unified Map for the application service in the CMDB Workspace.

Related topics
View the attributes of a CI or a relationship

View service instances on Unified Map

View historical changes for a CI

Application service mapping using classic Service Mapping

You can perform initial mapping of multiple or single application services in your organization using a method that best suits you.

You can map application services after you complete the Service Mapping setup.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping workflow

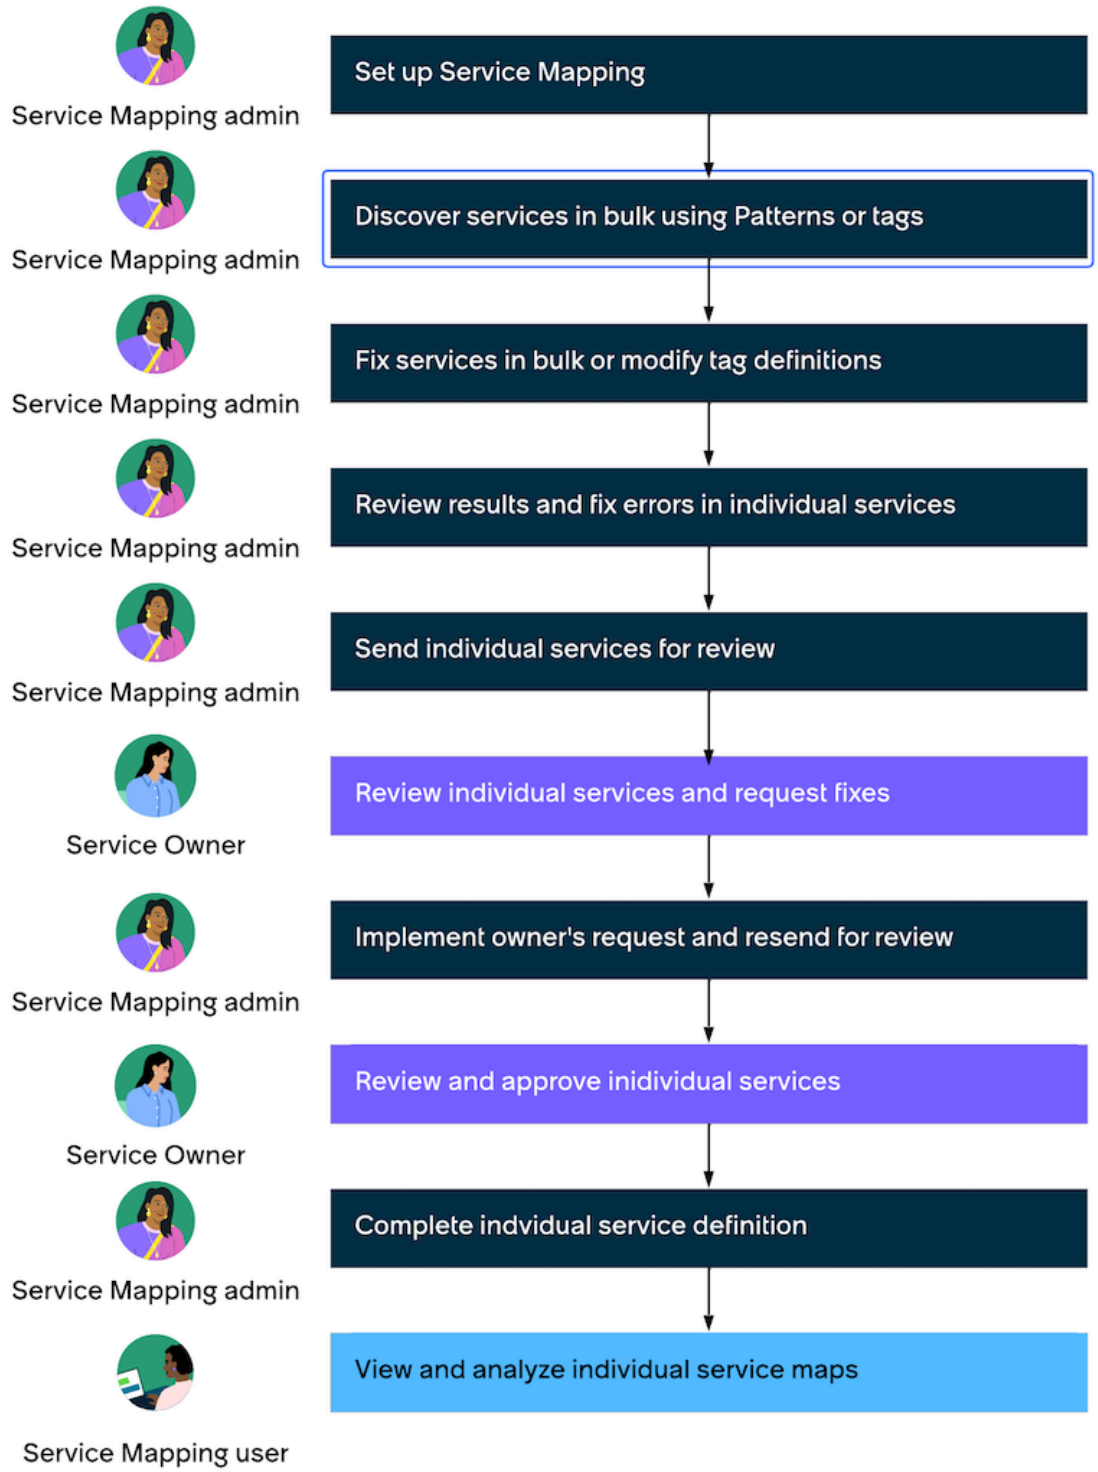

A typical organization has hundreds of application services. Mapping each of them individually can be time and effort consuming. With Service Mapping, you can map most IT services in bulk. Mapping multiple application services at once is very effective, but is not precise. You may need to individually map application services omitted during bulk mapping.

A discovery method depends on how you have been managing your application services prior to mapping them and what information about them you gathered:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Map multiple application services suggested by classic Service Mapping

  Service Mapping uses load balancers on your network to identify possible application services (candidates). You can map all or selected candidate in bulk, creating multiple application services in a single operation. This method suits your organization if you do not have much information about application services.

- Map application services using tags with classic Service Mapping

  If your organization uses tags for asset management, you can use these tags to map service instances. Unlike other mapping methods, tag-based mapping does not require configuring credentials or providing users with elevated rights. You can effectively use tags to map multiple service instances.

- Map multiple application services from a CSV file using classic Service Mapping

  This method suits you if your organization has performed cross-organization mapping and analysis and collected some information about planned service instances. If so, you can organize the collected information in a specific order and save it as a CSV file. Service Mapping extracts information from this file and creates potential service instances referred to as service candidates.

- Map a single application service using classic Service Mapping

  You can map individual, single application services one by one by defining attributes for each service instance. Use this mapping method for application services Service Mapping cannot discover during bulk mapping. You must know or plan to find out application service attributes necessary for mapping.

Map multiple application services suggested by classic Service Mapping

You can map multiple application services identified and suggested by Service Mapping in a single operation. This method suits your organization if you do not have much information about application services.

Before you begin

Role required: service_mapping_admin

About this task

You must have Service Mapping enabled to map or review service instances.

Immediately after you complete the Service Mapping setup, Service Mapping automatically creates potential application services referred to as candidates:

1. Service Mapping extracts entries directly from load balancers on your network.

2. Service Mapping converts load balancer entries into potential entry points and stores them in the Candidate Entry Points [sa_cand_entry_point] table.

   Service Mapping also creates candidate entry points based on Netflow and TCP and stores them in the same table.

   An entry point is a point where clients access a service instance. Service Mapping starts the mapping process from this point. For example, to map your electronic mailing service instance, define an IP address or host name of the email server as an entry point.

3. Service Mapping checks that potential entry points created from Netflow and VPC logs are unique and not in use by any existing application services to avoid duplication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Service Mapping creates a service instance candidate for each entry point.

   Not all candidate service instances correlate to actual service instances. Service Mapping filters out unlikely candidates by applying an algorithm that takes into account such parameters as empty IP addresses, port numbers, and the load balancer VIP hit counts less than 1000. Service Mapping stores candidates in the Candidate Entry Point [sa_cand_entry_point] table in the CMDB.

During the discovery and mapping process, Service Mapping creates a service instance from each candidate and assigns it the entry point retrieved from the candidate. While mapping in bulk is an effective and fast way to create numerous service instances, the result of the bulk mapping is not precise. There may be some false or duplicate service instances created by mistake because the data from the load balancer is raw. During the review and approval process, you review all mapped service instances and get rid of the ones that were erroneously created.

You can further filter the list of candidates by selecting specific ones to include in or exclude from discovery. Mapping only selected candidates provides the following advantages:

- Reduces the discovery time
- Minimizes the number of potential irrelevant application services

You can remove the unwanted candidates from the list using these methods either by filtering them out or deleting them from the Candidate Entry Point [sa_cand_entry_point] table. In both cases, you can bring hidden or deleted candidates back to the list.

Procedure

1. Navigate to **All** > **Service Mapping** > **Home**.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.
2. Verify that Service Mapping is set up and ready for discovery.
   If some obligatory configurations are missing, the error displays at the top of the window.

   a. Click the link in the error or click the **Readiness Checklist** icon ( ).

   b. In the Service Mapping Readiness Checklist window, click the link next to the setup category with the error icon ( ).
   You cannot perform discovery and mapping unless there are no errors in the Readiness Checklist window.

   c. Perform the necessary action to complete the setup.
   For example, configure missing credentials.

3. To map all application services suggested by Service Mapping, click **Discover All** or select them individually under the Action on selected rows button then click **Discover Selected**.
4. To map selected application services suggested by Service Mapping:

   a. Click **Additional options** under the **Map** tile.

   b. Select **View Candidates List**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Optional: Display candidates based on traffic-based connections, remove the condition that filters out candidates retrieved from load balancers or imported from a CSV file.

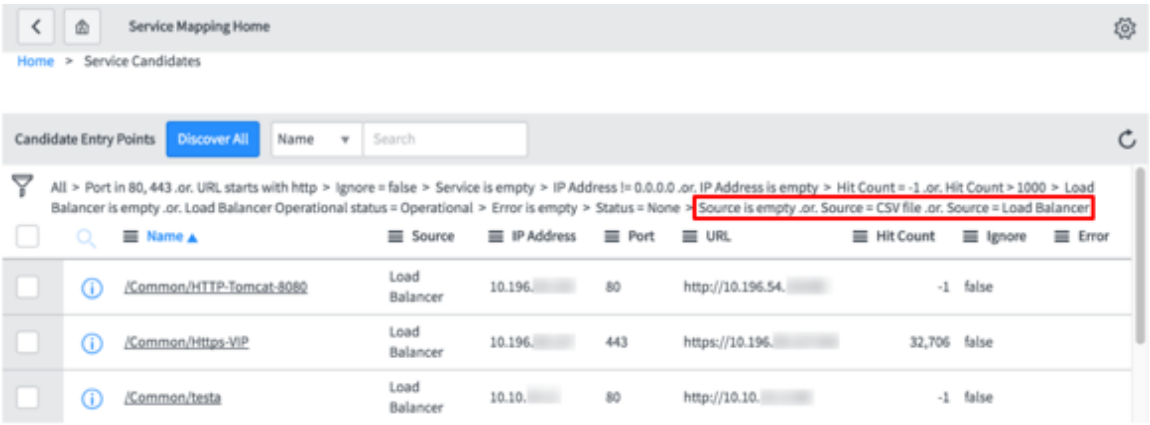

The candidate list shows candidates with the source attribute *netstat*.

d. Review the list of candidates to decide which ones you want to discover as application services.

e. Create application services from the relevant candidates:

If you want to create application services from most of the candidates	**i.** Filter out irrelevant candidates by selecting them and clicking **Actions on selected rows**, and then clicking **Ignore Selected**. The irrelevant candidates are hidden from the list. **Note:** To show candidates you previously filtered out, set the **Ignore** filtering condition to **true**. **ii.** Alternatively, delete irrelevant candidates by selecting them and clicking **Actions on selected rows**, and then selecting **Delete**. The irrelevant candidates are removed from the Candidate Entry Point [sa_cand_entry_point] table containing candidates. **Note:** Service Mapping recreates deleted candidates every time Discovery performs horizontal discovery of the same load balancers or you reimport candidates from the same CSV file. **iii.** Click **Discover All**.
If you want to create application services only from several candidates	**i.** Select the check boxes next to the relevant candidates. **ii.** Click **Actions on selected rows** at the bottom of the screen. **iii.** Click **Discover selected**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping starts discovery of all candidates and creates application services from them. If you have many candidates, the discovery may take some time.

After the discovery process finishes, application services, which Service Mapping discovered without errors and with at least one discovered CI, appear under the **Approve** tile. application services discovered with errors appear under the **Fix** tile.

What to do next

- If the connection suggestions feature is enabled, Service Mapping does not automatically add any traffic-based connections to application services, providing suggestions instead. Manage CI connections in bulk using suggestions.
- Fix application services discovered with errors. For more information, see Fix application service errors in bulk.
- Review the list of created application services. If Service Mapping did not create some application services during bulk mapping as you expected, create them manually as described in Map a single application service using classic Service Mapping.
- Send application services discovered without errors to their respective owners for approval. For more information, see Send application service maps for review.
- Make sure that the service instances aren't too large:
  - Service Mapping doesn't offer to view CI list instead of a map for a service instance.
  - There is no discovery message indicating that the service instance is too large: *`The map does not display the entire service, because it is too large. The number of CI connections exceeded the allowed maximum.`*

  If some application services are too large, resolve the service size issue:
  - If you use customized discovery patterns and mapping results in very large service instances, modify the patterns so that they don't map irrelevant CIs.
  - Split the service into smaller services.

Manage CI connections in bulk using suggestions

If the discovery based on Predictive Intelligence is enabled, the newly mapped application services include only CIs and CI connections added by discovery patterns. Use connection suggestions to decide which configuration items (CIs) to include or exclude globally. Service Mapping then updates all relevant discovered application services to reflect your decisions.

Before you begin

Confirm that fingerprint-based discovery remains enabled, since it is enabled by default. To verify, navigate to **All** > **System Properties** > **All Properties**. Check that the *`sys_property process.clustering.appfingerprint.enabled`* property does not appear on the list.

Confirm that discovery based on Predictive Intelligence is enabled.
Navigate to the System Property [sys_properties] table and verify that the *`sa_ml.connection_suggestions.active`* property is set to **True**.

You can add or remove CIs using connection suggestions only in application services discovered by Service Mapping.

Role required: service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

If Predictive Intelligence is enabled, Service Mapping automatically adds connections to service instances based on connection rules. These suggestions are generated using traffic-related data from the Configuration Management Database (CMDB) and the analysis of application fingerprints, CIs, and processes by Predictive Intelligence. Predictive Intelligence evaluates connections between application fingerprints, CIs, and processes, and ranks their relevancy. Service Mapping uses this information to create connections based on connection rules. It also generates connection suggestions for servers and load balancers for you to decide which connections to add or remove from the service instances. For more information, see Traffic-based discovery in Service Mapping.

If you are not sure which traffic-based connections are relevant, identify a group of application services that your organization uses for the same purpose. Add or remove CI connections for a single application service in this group. Then use the connections you created for this service as an example of traffic-based connections for multiple application services.

Procedure

1. **Optional:** Discover applications based on fingerprints to view information about all application CIs in connection suggestions.

   Important: Not using fingerprint-based discovery may result in mapping some application CIs as generic applications.

2. Navigate to **Service Mapping** > **Administration** > **Connection Suggestions**.
   The Connection Suggestions window opens.
3. Expand the relevant group to see the connection suggestions.
   For example, expand the **python - python** group.
4. Review the suggestion details paying attention to the parameters covered in the Connections Suggestions list.
   See Connection Suggestions list.
5. Select the check boxes next to the relevant suggestions and click the action button at the bottom of the window:

Action	Result
Add	Service Mapping includes both the connections and the CIs that these connections lead to, making them part of the service instance. The *Decision* value for added connections changes to **Added manually**.
Exclude	Service Mapping excludes connections and the CIs they lead to from the service instance. The *Decision* value for removed connections changes to **Excluded manually**.
Reset	Service Mapping reverses earlier decisions, whether made manually or by rule. The *Decision* value for reset connections changes to **Undecided**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Result
	This action is relevant only if you have already used connection suggestions to complete this service instance.

6. Refresh the connection suggestions to see the changes that you made:
   a. Click the List controls icon .
   b. Select **Refresh list**.

Result
Service Mapping applies the changes to traffic-based connections to all relevant application services in the background.

Add CIs to application services using connection rules

Create rules for automatically adding traffic-based connections and the CIs they lead to in discovered service instances. Create rules that add CIs to multiple application services.

Before you begin

- Analyze the infrastructure you are mapping to identify the most common connections. These common connections are good candidates for connection rules.
- Discover applications based on fingerprints.
- Confirm that discovery based on Predictive Intelligence is enabled.
  Navigate to the System Property [sys_properties] table and verify that the *sa_ml.connection_suggestions.active* property is set to **True**.

Role required: service_mapping_admin

About this task
Connection rules enhance the connection suggestions generated using Predictive Intelligence. Each rule contains a condition that describes an aspect of a traffic-based connection. The fields available in conditions come from the Connection Suggestion [sa_ml_connection_suggestion] table, which is only populated when you run top-down discovery. You can base your rules directly on the connection suggestions. For example, create a rule for any connection with confidence level High. Alternatively, if you know which CIs must be part of the service instances, you can choose fields that point at these CIs.

For information about Service Mapping evaluating and applying connection rules, see Traffic-based discovery in Service Mapping.

In addition to connection rules that add CIs to multiple services, you can create rules for adding CIs to specific application services.

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Connection Rules**.
   Alternatively, from a service application map, select **Suggestions** > **Global Rules**.
2. Create a new connection rule by clicking **New**.
   Alternatively, customize an existing rule by clicking the rule name in the list.

   The Connection Rule form opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Active** to start using this connection rule.
   Service Mapping evaluates only active connection rules.
4. Enter a unique and meaningful name in the **Name** field.
5. Enter the value that Service Mapping uses for the sequence in the **Order** field.
6. Configure the filter condition in the **Condition** area.

   a. Select the field, operator, and the value.

   Example
   For example, to add outgoing connections for any Tomcat server, configure the following filtering condition.

   b. To create an additional filtering condition, click **Add Filter Condition**.

   c. To create an additional filtering condition using the OR operator, click **Add "OR" Clause**.

7. Select `Global` from the **Scope** list.
8. **Optional:** Specify services to which this connection rule doesn't apply.

   a. Click **New** in the **Excluded Services** area.
   The Excluded Service form opens.

   b. Select the relevant service from the **Discovered Application Service** list.

   c. Click **Submit**.

9. Click **Update**.

What to do next

Verify that Service Mapping added CIs as expected by performing the following steps:

1. After the scheduled jobs rediscover the application services, open the maps of the application services where you expected Service Mapping to add CIs.
2. Check that CIs are added correctly based on the connection rules.

Related topics

Discovery based on Predictive Intelligence

Map application services using tags with classic Service Mapping

Use tags that help categorize and organize configuration items (CIs) in your organization to map application services. Tag-based mapping doesn't require configuring credentials or providing users with elevated rights.

Before you begin

Analyse, review, and define tags as covered in Prepare for mapping application services based on tags.

Role required: service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

A tag is a label that consists of a key-value pair. Your organization may use tags to categorize its assets, to enhance query and reporting capabilities. Discovery and Cloud Provisioning and Governance can discover tags used by all major cloud providers and container ecosystems. Once the tags are discovered, Service Mapping can create service instances based on these tags. For example, you can use tags to map all application services your organization uses in the production environment in the EMEA region. You can effectively use tags to map multiple service instances.

Note: Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

The result of the tag-based mapping is a hierarchical service map that may be less precise than the application service created using the patterns.

Comparison of pattern-based and tag-based mapping results

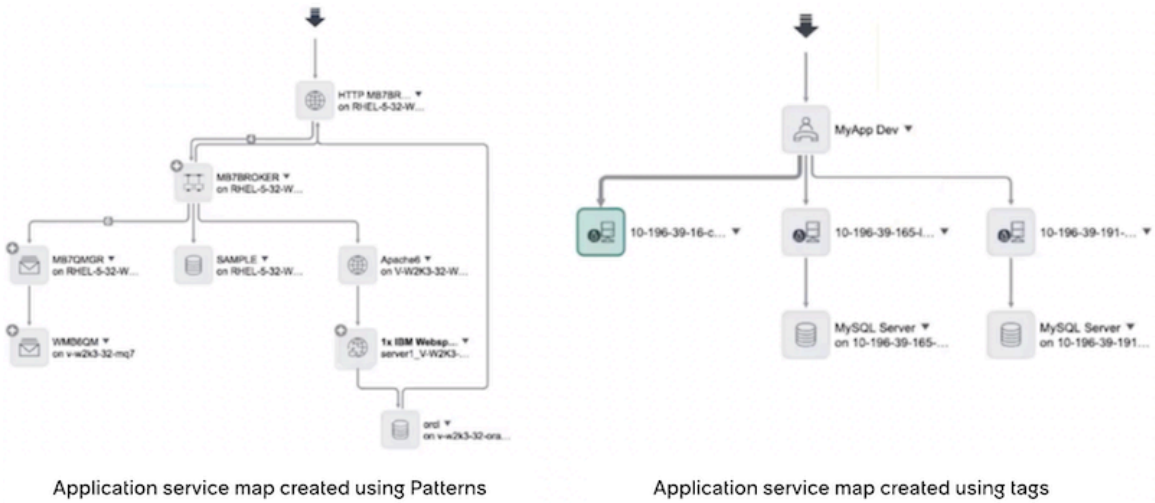

Application service map created using Patterns

Application service map created using tags

Service Mapping creates names for tag-based services using tag values discovered for the tag definitions, to which these tag-based services belong. For example, for Environment and Application tag categories with Production and HR as their respective tag values, the default service name is production::hr. Tag-based service names use low case.

Important: If you add or remove tag categories after mapping services, Service Mapping does not implement these changes for previously mapped services belonging to this service family. For more information, see Modify or update tag definitions for tag-based mapping.

In domain-separated deployments, Service Mapping populates tag-based services based on this service family only with CIs belonging to the leaf domain to which the service family belongs.

To map a single application service using tags without preliminary configuration, use the flow based on Common Service Data Model. For more information, see Populate application services using tags.

Note: Starting with version 1.15.2, there is a limit of 200 application service candidates per family.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Tag-based Service Families**.
2. Click **New**.
3. Define the service family name.
4. **Optional:** Select **Regularly update service candidates** for Service Mapping to regenerate suggestions for service instances that match tag categories you define for the service family.
5. Define the tag categories and tag values to use for mapping service instances belonging to this service family.

   a. Under **Selected tag categories**, select the relevant tag category from the list of tag categories you defined earlier.

   b. **Optional:** To make the mapping process more precise, add the tag value, for example, Production.
   For multiple tag values, use comma as a separator: Staging, Production.

   c. **Optional:** To change the default service naming convention, set the order in which Service Mapping uses the tag values in the candidate name.
   For example, the first category of the service family is Environment and the second category is Application. The default service name that Service Mapping assigns to the service candidate is Production::Finance. To make Service Mapping use the Application category value at the beginning of the name, set the Order value for the Application category to 1.

   d. **Optional:** Select other tag categories as necessary for this service family.
   The maximum number of categories per service family is three.

6. Click **View service candidates.**
7. In the Service candidate window, review the suggested service candidates and select the ones that you want to map.
   If there are no service candidates, in the Key Value [cmdb_key_value] table, ensure that there are tag values for the tag keys you used for mapping.
8. Click **Map selected**.

   Service Mapping creates service instances based on the candidates you selected.

   If there is an error about the service name that exists in the system, fix the candidate name by performing the following steps:

   a. Click **View service candidates**.

   b. Click the name that exists and change it.

   c. Click **OK** to save the change.

   d. View the service candidates and map the relevant ones.

      New candidates contain only the service instance CIs as their entry points. The Populate Calculated Services scheduled job maps new tag-based services once a day.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. **Optional:** To map a new service instance immediately, without waiting for the scheduled job to populate it:
   a. Select the relevant service instance.
   b. In the service instance form, click **Additional Info** on the left.
   c. Under **Related Links**, click **Recalculate Service**.
   d. Click **View Map** at the top of the form and review the mapping result.
10. After the scheduled job maps all new tag-based services, navigate to the relevant tag-based service family and review the newly mapped services listed under **Mapped services**.

Result

All tag-based service instances have their respective service instance CIs as their entry points.

The map for the Finance::Production service instance based on tags

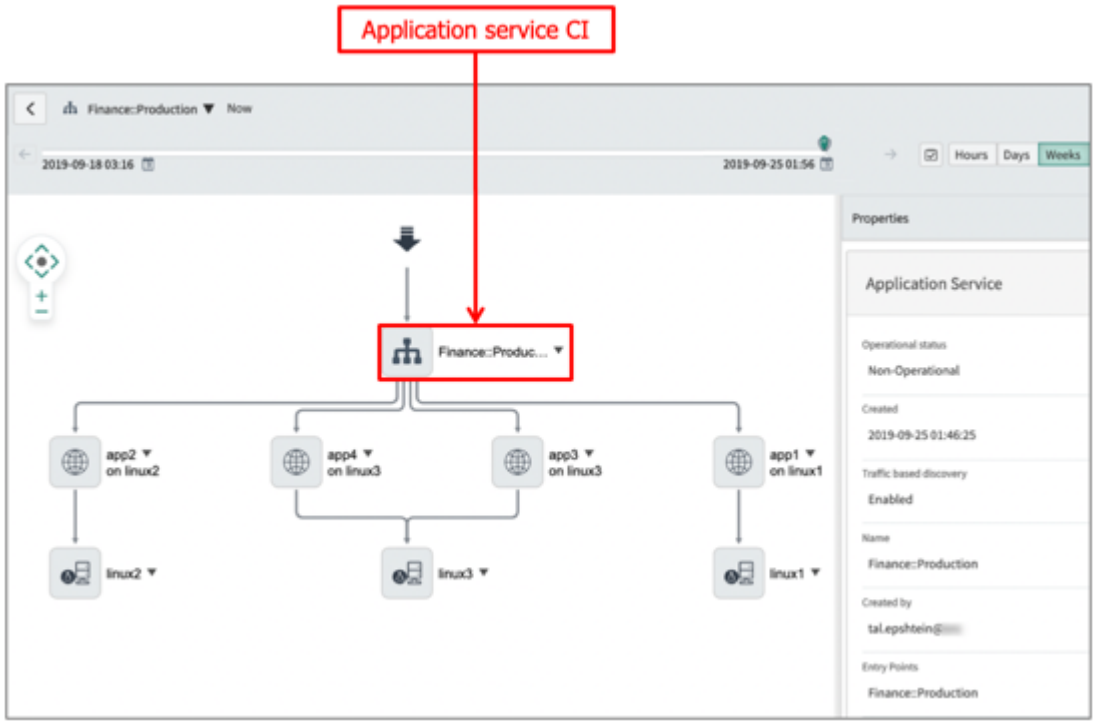

What to do next

1. After the scheduled job maps all new tag-based services, navigate to the relevant tag-based service family and review the newly mapped services listed under **Mapped services**.
2. Make sure that the service instances aren't too large: The following discovery message is displayed in a tag-based service instance: *`This tag-based service instance reached the maximum number of CIs. For more information, refer to the product documentation.`*

   If some application services are too large, resolve the service size issue:

   - Modify tag definitions used for mapping to make sure Service Mapping includes only relevant CIs in the service instance.
   - If necessary, Tag-based discovery configuration to change which CIs Service Mapping includes in tag-based service instances.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. To fine-tune or fix the tag-based services, modify the tag categories that you used for the service family.

> **Important:** The map window displays maps for tag-based and dynamic services in the View mode only. You cannot fine-tune or edit tag-based and dynamic services from the map.

Related topics

Choose the right method for discovering and mapping services

Tag-based discovery configuration

Prepare for mapping application services based on tags

Analyze, review, and define tags in your organization for successful mapping.

Before you begin

Role required: service_mapping_admin

About this task

As a preparation for tag-based mapping, create tag categories that contain tags with similar use.

> **Note:** Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Define a tag-based service family and define tags that you want Service Mapping to use for mapping. Based on the tag definitions for the tag-based service family, Service Mapping creates service candidates - suggested service instances. You review the candidates and decide which ones you want to use to create the actual tag-based service instances. For details on the tag-based mapping process, see Tag-based discovery in Service Mapping.

CIs that have discovered tag values for the tag categories you selected become part of application services. Service Mapping creates a separate service candidate for each tag value combination. If you narrow the criteria down by defining the tag values in addition to tag categories, Service Mapping maps only CIs that have the matching values. CIs that have more than one tag assigned to them, can be part of multiple services. You may want to create a tag category for tags related to different types of environments, if your organization uses "production" and "staging" tag values. Tag-based mapping is not case-sensitive; same key names and key values spelled with upper and lower case are identified as the same.

> **Note:** Service Mapping includes CIs that are part of CI relationships even if these CIs do not have tags assigned to them. For more information, see Tag-based discovery in Service Mapping.

Procedure

1. Analyze the tag usage in your organization and make a list of all tags and their purposes. Use the Key Value [cmdb_key_value] table to see the tags in the CMDB.

   Use Tag Governance to analyze and optimize tag application in your organization. See Tag Governance.

2. If necessary, assign tags to CIs that you want to include in application services.

   Avoid creating very large application services containing more than 5000 CI relations. Such large services may cause performance issues.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Create and define tag categories that you want Service Mapping to use for mapping service instances.

   For example, to map all service instances your organization uses in the production environment in the EMEA region, create the following tag categories: Applications, EMEA region, and Production.

   a. Navigate to **Service Mapping** > **Administration** > **CI Tag Categories**.

   b. Select a **CI tag category** to modify or select **New**.

   c. Define the tag category name, for example, *Environment*.

   d. Under **CI tag keys**, double-click the row and type the tag key, for example *Env*.

      Since tag-based mapping is not case-sensitive, tag key variants using different capitalization, like *Env*, *env* or *ENV*, are also assigned to this same tag category.

      Repeat this step until you add all tag key variants that your organization uses for the same purpose, for example *Env* and *Environment*.

   Note: To see detailed examples of defining tag-based mapping criteria for service families, see Tag-based mapping criteria examples for service families.

4. Run Discovery to horizontally discover CIs. Discovery populates the CMDB with the CI data, including tags.

   If your ServiceNow instance uses domain separation and you have access to the global domain, log in to the relevant domain. The selected domain must be a domain without any child domains.

Map multiple application services from a CSV file using classic Service Mapping

This method suits you if your organization has performed cross-organization mapping and analysis and collected some information about planned service instances. If so, you can organize the collected information in a specific order and save it as a CSV file. Service Mapping extracts information from this file and creates potential service instances referred to as service candidates.

Before you begin

Perform the procedure described in Prepare a CSV file for mapping your candidates.

Role required: service_mapping_admin

About this task

You must have Service Mapping enabled to map or review service instances.

If necessary, you can import service candidates from multiple CSV files.

When reimporting service instance candidates, Service Mapping does not import candidates you chose to ignore earlier.

After Service Mapping discovers configuration items (CIs) belonging to your service instance for the first time, it then rediscovers CIs to find changes and updates. Create or modify discovery schedules to control how often Service Mapping rediscovers services or CIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Service Mapping** > **Home**.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.
2. Under the **Map** tile, select **Additional Options**.
3. Select **Import Service Map List**.

Map tile Additional Options

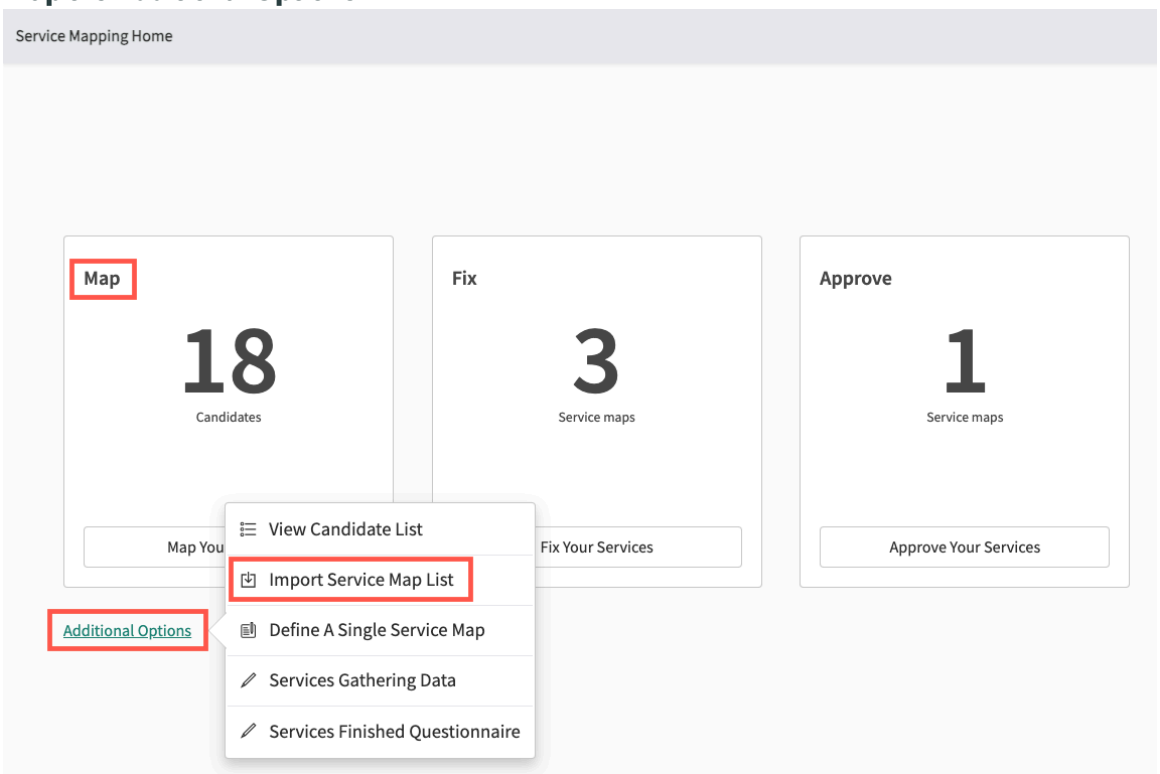

4. In the Import CSV file window, select **Choose file**.
5. Navigate to the CSV file to use for the import and select **Open**.
6. Select **Import**.
   The imported service instance candidates are added to the list of candidates.
7. Check that the overall number of service candidates on the **Map** tile increased by the expected number.
8. On the **Map** tile, select **Map Your Services** to create application services from the candidates you imported from the CSV file.

What to do next

- Review the list of created application services. If Service Mapping did not create some application services during bulk mapping as you expected, create them manually as described in Map a single application service using classic Service Mapping.
- Fix application services discovered with errors. For more information, see Fix application service errors in bulk.
- Send application services discovered without errors to their respective owners for approval. For more information, see Send application service maps for review.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prepare a CSV file for mapping your candidates

Organize information about potential application services (candidates) in your organization and save it in a CSV file.

Before you begin

Role required: service_mapping_admin

About this task

Procedure

1. Create an .xlsx file.
2. Define entry point attributes for service candidates in the .xlsx file.
   Customers use these entry points to access services. For example, they use http://www.google.com:8080 to access the Google page.

Warning: Do not confuse column order in the .xlsx file. Entering attributes in wrong columns causes mapping errors.

Column	Description
A	Enter a unique name of the planned service instance. For example, NY Public Library or Google. **Warning:** Service Mapping imports only application services with unique names. If there is more than one entry with the same name in CSV files you use, the latest imported record overwrites the previous ones.
B	(Optional) If you know the URL that serves as the entry point to your service instance, enter it in this column. For example, enter https://www.nypl.org/ or www.google.com:8080. **Warning:** If you define the URL, you do not have to define the IP address, port, FQDN, and protocol. Service Mapping parses the URL to extract the IP address, port, and FQDN. If there is no URL that serves as the entry point, define other parameters in columns C, and E.
C	(Optional) The IP address of the entry point in the IPv4 format. Enter this parameter only if you did not enter the URL in column B.
D	(Optional) The port of the entry point. For example, 8080 for the Google page. Enter this parameter only if you did not enter URL in column B.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
	If you did not enter either URL or the port, Service Mapping uses port 80 for the http protocol and port 443 for the https protocol.
E	(Optional) The FQDN of the service entry point, for example,www.google.com for Google. Enter this parameter only if you did not enter the URL in column B.

3. Save the file in a drive that you can access during the import with the *.csv* extension.

What to do next

Map multiple application services from a CSV file using classic Service Mapping

Map a single application service using classic Service Mapping

In addition to mapping application services in bulk, you can map individual, single application services by defining attributes for each application service. Use this mapping method if you already know or are planning to find out application service details. You must have Service Mapping enabled to map or review application services.

Before you begin

- Verify that Service Mapping is set up properly.
- If the service instance relies on PaaS supported by Microsoft Azure or Amazon Web Services, discover datacenters and then create horizontal discovery schedules on the cloud service account.

Role required: service_mapping_admin

About this task

You can map single application services in addition to bulk mapping, which is effective at mapping multiple application services, but is not precise. You may need to individually map application services omitted during bulk mapping.

The most important attribute you must know and configure to discover an application service is an entry point.

An entry point is a point where clients access a service instance. Usually, it is either a URL or a combination of the IP address and port. Service Mapping starts the mapping process from this point. For example, to map your electronic mailing service instance, define an IP address or host name of the email server as an entry point.

Entry points vary depending on the nature of the service instance. Service Mapping comes with a wide range of preconfigured entry point types that cover many commonly used applications.

In addition to entry points that Service Mapping can discover, you may add entry points that Service Mapping cannot discover. You may want to add such an entry point as an indication that your service instance contains a device or application.

Also, you must define an owner for your service instance. The service instance owner is a user who is familiar with the infrastructure and applications making up the service. This user is the service instance SME who provides information necessary for a successful creation of a service instance. After a service is mapped, this user reviews the results, and either approves it or suggests changes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After Service Mapping discovers configuration items (CIs) belonging to your service instance for the first time, it then rediscovers CIs to find changes and updates. Create or modify discovery schedules to control how often Service Mapping rediscovers services or CIs.

All service instances, both manually created and discovered by Service Mapping, are stored in the Mapped Service Instance [cmdb_ci_service_discovered] table.

You can also create a single application service manually using the flow based on Common Service Data Model. The Common Service Data Model (CSDM) helps you streamline service types and service offerings. You can add relationships between application services and other service-related objects in the CSDM: Business Application, Technical Service Offerings, or Business Service Offerings. For more information, see Create application services .

Procedure

1. Navigate to **All** > **Service Mapping** > **Home**.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.
2. Click **Additional options** under the **Map** tile.
3. Click **Define A Single Service Map**.
4. Define the basic service instance attributes:

Field	Description
Name	Enter a unique application service name which is not in use by any other type of application service. Use self-explanatory names such as`mailing service` or `printing service`.
Owner	Select an owner who is responsible for this service instance in the organization. If the owner name does not appear in the list, create a user for this person and assign the sm_app_owner role to this user. Alternatively, choose a user with the service_mapping_admin role.

5. If you do not know the entry points, request them from the service instance owner, as described in Request entry point information for service mapping.
6. To add an entry point that Service Mapping can discover, perform these steps.

   a. Click **Discoverable by Service Mapping**.
   The page offers tiles for applications most commonly used for entry points.

   b. Click the tile matching the application to which the entry point belongs.

   For example, **Web Application** if you create an service instance with a web interface. If the tiles do not contain the necessary application, click **See More Types** to choose from a list of all preconfigured entry point types.

   > **Important:** Do not use central IT components like Microsoft MQ or Microsoft Active Directory as entry points. Using such entry points creates large services, which cause performance issues during mapping and monitoring.

   c. **Optional:** If the required entry point type does not appear in the list, you may need to create it as described in Create entry point types for Service Mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Configure attributes for the entry point as described in Entry point attributes.
Entry point attributes depend on the type you select.

e. Click **Add**.
The entry point is added under **Entry Points**.

f. To add more entry points, click **Add Entry Point** and repeat step 6.

7. **Optional:** To add a static entry point that Service Mapping cannot discover, perform these steps.

a. Click **Manually created**.

b. Define the CI to serve as the entry point.

Field	Description
CI Type	Select the CI type (CI class) for the CI you are adding. Every CI belongs to a CI type which contains a set of attributes configured for this kind of CI, for example, *cmdb_ci_appl* for applications.
CI Name	Select the CI from the list of CIs of the selected CI type. **Note:** To eliminate the possibility of delayed results when searching for a specific CI, make your search as specific as possible. A search with *<name> might take a long time and return a large data set.

The CI type list includes only allowed CI types. For example, you cannot add an application cluster.

c. Click **Submit**.

8. To save attributes you have defined up until this moment and carry on configuring attributes for this service instance, click **Save** or **Update**.
Service Mapping saves the new service instance in the Mapped Application Service [cmdb_ci_service_discovered] table. It also starts the discovery process in the background.
9. To save definitions, discover this service instance, and navigate to the map for this service instance, click **View Map**.
Service Mapping starts the discovery process. The map for this service instance page opens and displays the mapping in progress. The service instance status changes to In Progress.

What to do next

- If Service Mapping discovers the service instance with errors, fix errors in individual application services using discovery messages.
- Make sure that the service instances aren't too large:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Service Mapping doesn't offer to view CI list instead of a map for a service instance.
- There is no discovery message indicating that the service instance is too large: *The map does not display the entire service, because it is too large. The number of CI connections exceeded the allowed maximum.*

If some application services are too large, resolve the service size issue:

- If you use customized discovery patterns and mapping results in very large service instances, modify the patterns so that they don't map irrelevant CIs.
- Split the service into smaller services.

- Send application services discovered without errors to their respective owners for approval. For more information, see Send application service maps for review.

Request entry point information for service mapping

The most important attribute you must know and configure to discover an service instance is an entry point. If you do not know the entry points for the service instance, request this information from the service instance owner.

Before you begin

Requesting information about entry points is an optional part of mapping a single application service.

Roles required: service_mapping_admin

About this task

An entry point is a point where clients access a service instance. Usually, it is either a URL or a combination of the IP address and port. Service Mapping starts the mapping process from this point. For example, to map your electronic mailing service instance, define an IP address or host name of the email server as an entry point.

Entry points vary depending on the nature of the service instance. Service Mapping comes with a wide range of preconfigured entry point types that cover many commonly used applications.

Procedure

1. Click the **Request data** link.
   The Request for data window opens.
2. If the **Assigned to** list is empty, select the name of the service instance owner.
3. Enter your questions about this service instance in the **Description** field.
   For example, you can ask what the entry points are.
4. Click **Submit**.
   The system opens ServiceNow task for the owner and automatically sends an email notification.
5. Review information added by the owner on the **Other Entry Points** and **Components** tabs.
   Once the owner provides missing information, you receive an email notification.
6. Click the link in the email to open the Questionnaire screen for the relevant service instance.
7. Continue mapping a single application service.

Provide entry points for mapping an application service

As an application service owner you may receive a request for information about entry points in an email notification. Provide information about entry points to enable administrators to start discovery of an application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: sm_app_owner or service_mapping_admin

About this task

If information about service instance entry points is missing, an administrator cannot start mapping the service instance. In this case, the administrator sends a request for missing data from the service instance form, which creates a service process task assigned to you. You receive an email notification with the link to the Questionnaire page where you must enter the missing data. The most important information is about the entry points, however, any additional data about application service components, their connections, or usage is of help for the administrator. When you finish entering data in the Questionnaire and submit it, the system closes the service process task.

Mapping an application service flow

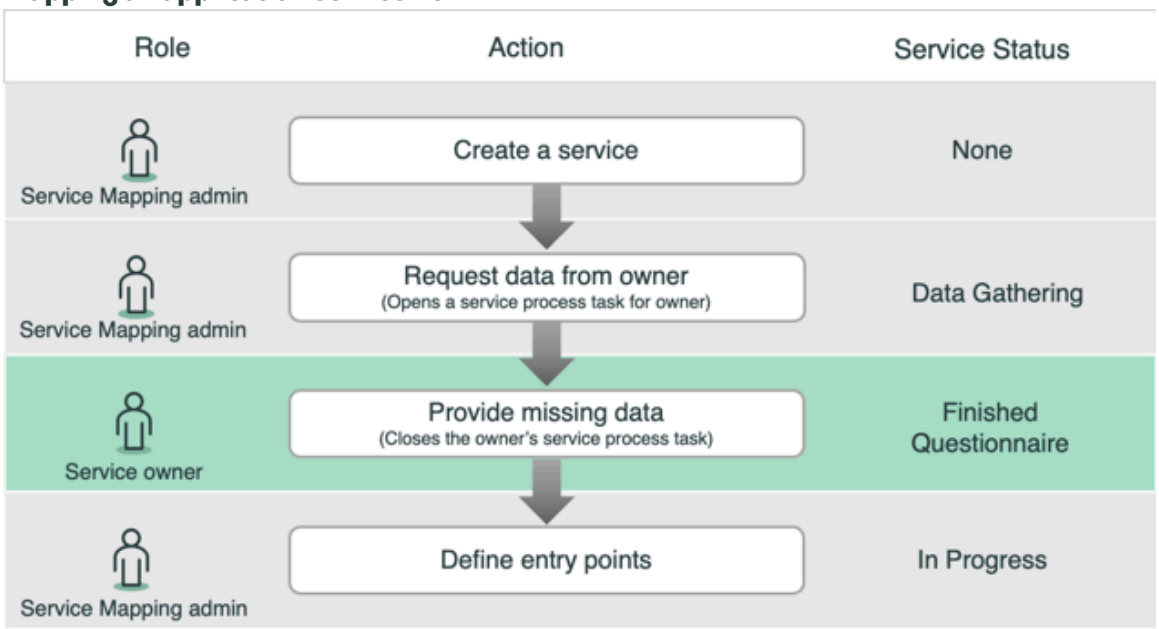

Procedure

1. Click the link in the notification email to access the **Questionnaire** page.
2. Alternatively, open to the **Questionnaire** page from the list of tasks assigned to you.
   a. Navigate to **Service Mapping** > **Administration** > **My Tasks**.
   b. **Optional:** Sort the list of service process tasks as required.
   c. Click the required task.
   d. Click the service instance form link.
   e. Click **Questionnaire** in the left pane.
3. On the **Planned Custom Entry Points** tab, enter information about service instance entry points.

   Note:

   If you cannot specify the exact attributes of the entry points, such as the URL or the IP address, add general information which can guide the administrator. For example, the type of the application, like Tomcat server.

4. To add information about any other components comprising this service instance, click the **Components** tab, and enter the relevant data there.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. When you finished entering the information, click **Actions**, and then select **Submit Questionnaire**.
   The data request task assigned to you closes. The service instance state changes to In Progress.
6. If there is any other useful information concerning this service instance, enter it in the **Notes** field under **Worknotes** and press Enter to post your comment.

Manage CI connections for a single service using suggestions

If discovery based on Predictive Intelligence is enabled, newly mapped application services include only configuration items (CIs) and CI connections added using discovery patterns. Use the suggestions for CI connections, generated by Service Mapping, to complete the application service.

Before you begin

- Confirm that fingerprint-based discovery remains enabled, since it is enabled by default. To verify, navigate to **All** > **System Properties** > **All Properties**. Check that the *sys_property* `process.clustering.appfingerprint.enabled` property does not appear on the list.
- Confirm that discovery based on Predictive Intelligence is enabled.
  Navigate to the System Property [sys_properties] table and verify that the `sa_ml.connection_suggestions.active` property is set to **True**.

Role required: service_mapping_admin

About this task

If discovery based on Predictive Intelligence is enabled, Service Mapping automatically adds connections to service instances based on connection rules. These suggestions are generated from traffic-related data in the Configuration Management Database (CMDB) and an analysis of application fingerprints, CIs, and processes. Predictive Intelligence evaluates and ranks the relevancy of these connections, and Service Mapping then uses this information to create connections and provide suggestions for servers and load balancers. Then you can decide which connections to add or remove from service instances. For more information, see Traffic-based discovery in Service Mapping.

Important: You can add or remove CIs using connection suggestions only in application services discovered by Service Mapping.

Procedure

1. **Optional:** View information about all application CIs in connection suggestions.
   See Discover applications based on fingerprints for more information.

   Important: If you don't use fingerprint-based discovery, some application CIs might be mapped as generic applications, lacking specific identification or detailed information.

2. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
3. Select **Suggestions** > **Connection Suggestions** at the top of the map to see suggestions for the entire application service.
   Alternatively, right-click a CI for which you want to see connection suggestions and select **Show connection suggestions**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Review the suggestion details paying attention to the parameters covered in the Connections Suggestions list.
   See Connection Suggestions list.
5. **Optional:** If you haven't discovered application CIs, discover a standalone CI of the application class based on fingerprints:

   a. Select the **Preview** icon next to the relevant suggestion based on the application fingerprints (AFP).
   The Connection Suggestion list opens showing details of this suggestion.

   b. Select the **Preview this record** icon next to **Target AFP**, and then select **Open Record**.

   c. Select **Discover Application**.

   d. Select **Refresh** in the Review suggested connections window to confirm that the displayed information reflects the discovered CI classes.

6. Select the check boxes next to the relevant suggestions and click the action button at the bottom of the window:

Action	Result
Add	Service Mapping includes both the connections and the CIs that these connections lead to, making them part of the service instance. The *Decision* value for added connections changes to **Added manually**.
Exclude	Service Mapping excludes connections and the CIs they lead to from the service instance. The *Decision* value for removed connections changes to **Excluded manually**.
Reset	Service Mapping reverses earlier decisions, whether made manually or by rule. The *Decision* value for reset connections changes to **Undecided**. This action is relevant only if you have already used connection suggestions to complete this service instance.

7. Select **Close** to leave the Review suggested connections window.

Result

The application service map reflects the changes that you made. The connections you added, along with the CIs they lead to, appear on the map.

Related topics

Traffic-based discovery in Service Mapping

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add CIs to an application service using connection rules

Create rules for automatically adding traffic-based connections and the CIs they lead to in discovered service instances. Add CIs to specific application services using connection rules.

Before you begin

- Analyze the infrastructure you are mapping to identify the most common connections. These common connections are good candidates for connection rules.
- Confirm that discovery based on Predictive Intelligence is enabled.
  Navigate to the System Property [sys_properties] table and verify that the *sa_ml.connection_suggestions.active* property is set to **True**.
- Discover applications based on fingerprints.

Role required: service_mapping_admin

About this task

Connection rules enhance the connection suggestions generated using Predictive Intelligence. Each rule contains a condition that describes an aspect of a traffic-based connection. The fields available in conditions come from the Connection Suggestion [sa_ml_connection_suggestion] table, which is only populated when you run top-down discovery. You can base your rules directly on the connection suggestions. For example, create a rule for any connection with confidence level High. Alternatively, if you know which CIs must be part of the service instances, you can choose fields that point at these CIs.

You can create connection rules that apply to specific application services or to all application services. If the rule is for a specific service, but the rule condition points at a CI that is part of multiple services, Service Mapping adds this CI to all relevant services.

For information about Service Mapping evaluating and applying connection rules, see Traffic-based discovery in Service Mapping.

In addition to connection rules that add CIs to specific services, you can create rules for adding CIs to multiple application services.

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Connection Rules**.
   Alternatively, from a service application map, select **Suggestions** > **Local Rules**.
2. Create a new connection rule by clicking **New**.
   Alternatively, customize an existing rule by clicking the rule name in the list.
3. Select **Active** to start using this connection rule.
   Service Mapping evaluates only active connection rules.
4. Enter a unique and meaningful name in the **Name** field.
5. Enter the value that Service Mapping uses for the sequence in the **Order** field.
6. Configure the filter condition in the **Condition** area.

   a. Select the field, operator, and the value.

      Example
      For example, to add outgoing connections for any Tomcat server, configure the following filtering condition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. To create an additional filtering condition, click **Add Filter Condition**.

c. To create an additional filtering condition using the OR operator, click **Add "OR" Clause**.

7. Select `Local` from the **Scope** list.
8. From the **Discovered Application Service** list, select an existing service to which this rule applies.
9. Click **Update**.

What to do next

Verify that Service Mapping added CIs as expected by performing the following steps:

1. Open the map for the application service you selected for this connection rule.
2. Click **Discover** at the top of the map window.
3. After the application service is rediscovered, check that CIs are added correctly based on the connection rule.

Enable traffic-based discovery for an application service

Make Service Mapping use traffic-based connections to discover CIs making up an application service.

Before you begin

Role required: service_mapping_admin

About this task

Service Mapping can discover and map configuration items (CIs) following their traffic-based connections. This method is referred to as traffic-based mapping and complements pattern-based mapping. For more information about how Service Mapping performs traffic-based discovery, see Traffic-based discovery in Service Mapping.

Procedure

1. In the classic environment, perform the following steps:

   a. Navigate to **All** > **Service Mapping** > **Home**.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.

   b. Click **View Your Services** on the **Completed** tile.

   c. Select the relevant application service.

   d. Click **Additional Info**.

   e. Click the **Traffic based discovery** check box.

   f. Click **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Alternatively, perform the following steps in the Agent Workspace UI:

   a. Navigate to **Agent Workspace** > **Agent Workspace Home**.

   b. Click **Lists** in the left pane.

   c. Navigate to **Services** > **All Services**.
   The list of services displays services in Operational status.

   d. Select the relevant application service.

   e. Click the **Open Preview** icon.

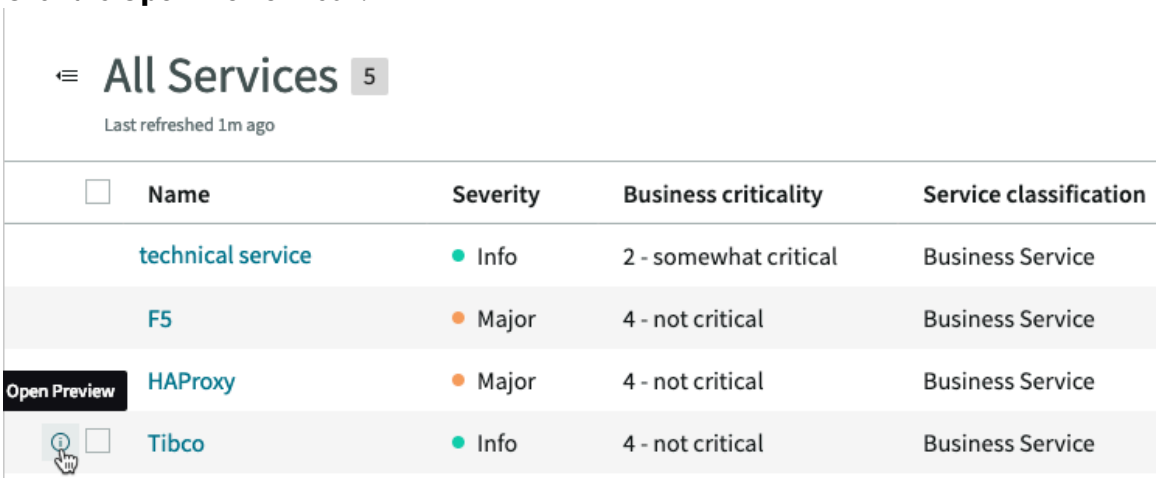

   The Details pane opens on the right.

   f. Click the **Traffic based discovery** check box in the Details pane.

   g. Click **Update**.

Mapping Citrix XenApp and Presentation Server services

Service Mapping can discover application services containing Citrix XenApp, Presentation Server, and Citrix Usage Collector.

Discovered components

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Citrix XenApp Server	Windows	6.x>6.5, 7.x<7.6	WMI WinRM	Horizontal and top-down	Citrix XenApp or Presentation Server
Citrix Presentation Server	Windows	4.x, 5.x, 6.x<6.5	WMI WinRM	Horizontal and top-down	Citrix XenApp or Presentation Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Citrix Usage Collector	Windows	7.x>7.5, 8.x	WMI WinRM	Top-down	Citrix Collector
Generic application based on the Icon Path	Windows	-	WMI WinRM	Top-down	Citrix Application Icon Path

Map Citrix XenApp and Presentation Server services

Map an application service containing Citrix XenApps, Presentation Servers, and Citrix Usage Collectors and verify that Service Mapping discovered all Citrix components correctly.

Before you begin

- Provide a user with permissions to read and query the Citrix repository:
  - For versions 4.5 and later, a PowerShell script from the XenApp server.
  - For versions earlier than 4.5, a Microsoft Visual Basic Script (VBScript) from the Citrix Presentation Server
- On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials.

  Note: To use this credential later for XenAPP, set the **Credential alias** attribute to cmdb_ci_appl_citrix_xenapp. To use this credential later for Citrix, set the **Credential alias** attribute to cmdb_ci_appl_citrix_collector.

- (Optional) To discover using the WinRM protocol, see Configure WinRM trusted hosts on MID Server
- Configure a user with permissions to run PowerShell commands on the Citrix server.
- (For discovery of Citrix Presentation Server only) Find out the icon path for the entry point:
  1. In the Citrix Access Management Console, navigate to **Presentation Server** > **Applications**.
  2. Locate the Icon Path folder in the navigation tree.
  3. Right-click the Icon Path folder and click **Copy path**.
- (For discovery of Citrix XenApp Server only) Find out the icon path for the entry point by checking the Distinguish Name attribute for the Icon Path.

Role required: service_mapping_admin

Procedure

1. Create a service instance.

   a. Navigate to **All** > **Service Mapping** > **Service Instances**.

   b. Select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Discoverable by Service Mapping** option is selected by default.

c. In the **Name** field, enter a descriptive name for the service instance.

2. Define the entry point for this service instance:

a. Select the **Citrix XenApp** tile.

b. Define entry point attributes.

Attribute	Description
Host	The URL of the Citrix XenApp or Presentation Server. The URL must contain a VIP IP or a physical host.
Port	The port that the Citrix XenApp or Presentation Server uses on the host.
Icon Path	The path to the logical application location as it appears in the Citrix Access Management Console. The path must start with the Applications folder. For example, /Application/Jane/Windows Media Player.
URL	The URL must contain a VIP IP or a physical host. For example, http://<host>/Citrix/StoreWeb.

c. Click **Add**.

d. Click **Save**.

3. Click **Additional Info** on the left and clear the **Traffic based discovery** check box.
4. Click **Update**.
5. Verify that the mapping result is satisfactory:

a. Click **View map** and wait until the mapping is complete.
It may take a few minutes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Review the application service map and make sure that it resembles in the following example.

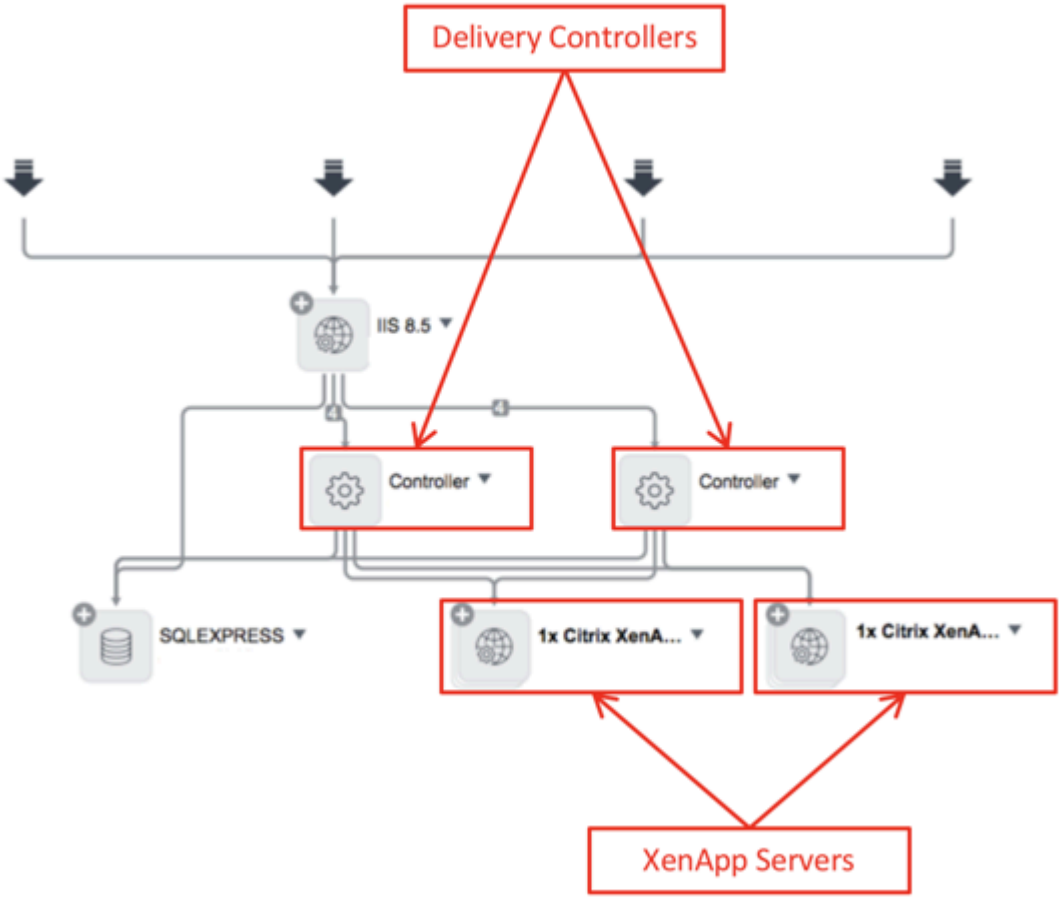

Trouble?

If the mapping result is not as you expected, perform basic troubleshooting.

Symptom	Cause	Solution
Discovery failed. The discovery log displays the message that starts with the following text: You do not have sufficient permissions.	The indication sections of the pattern fail, because the user does not have permissions for running PowerShell commands.	Make sure that you provided a user with permissions to run PowerShell commands on servers that host Citrix XenApp, Presentation Server, or Citrix Usage Collector.
Discovery failed. The discovery log displays the error message that includes the following text: CategoryInfo : NotSpecified: (:) [], PSSecurityException + FullyQualifiedErrorId : RuntimeException	The MID Server cannot run PowerShell commands remotely, because the execution policy is set to RemoteSigned on the target servers.	Use the `set-executionpolicy remotesigned` command on the target servers to enable the MID Server to successfully run PowerShell commands.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
Discovery failed. The discovery log displays the error message that includes the following text: Import-XAFarm : The Citrix Common Commands PowerShell Snapin is not installed	The target host does not have this PowerShell library: Citrix Common Commands PowerShell Snapin.	Download and install PowerShell SDK for Citrix XenApp on all servers hosting Citrix XenApp, Presentation Server, or Citrix Usage Collector.
The application service map shows the Citrix Delivery Controller, but does not include the Citrix XenApp, Citrix Presentation Server or Citrix Usage Collector.	The icon path you defined is incorrect.	Verify the icon path and reenter it in the entry point attributes for the application service.

Application services with Microsoft Exchange Server

Service Mapping can discover service instances containing a Microsoft Exchange Server and the Exchange Server components.

Discovered components

If e-mail services in your organization deploy Microsoft Exchange Server, the discovered application service includes the following components:

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Microsoft Active Directory Domain Services	Windows	13, 30, 31, 44, 47, 52, 56, 69	WMI WinRM	Top-down	Active Directory Service pattern
Microsoft server (as a Microsoft Exchange component)	Windows	2003, 2007, 2010, 2013, 2016	WMI WinRM	Horizontal and top-down	ExchangeFrontEndServer On Windows; PatternExchangeBackEndServ On Windows pattern
HUB Transport Server (as a Microsoft Exchange component)	Windows	2003, 2007, 2010, 2013, 2016	WMI WinRM	Horizontal and top-down	Hub Transport Server On Windows pattern
Client Access Server (CAS) (as a Microsoft Exchange component)	Windows	2003, 2007, 2010, 2013, 2016	WMI WinRM	Horizontal and top-down	CAS On Windows pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Typically, you do not want to discover Outlook clients connected to the Exchange Server to avoid noise. If this is the case, disable the traffic-based discovery for this application service.

For information on the discovery flow, refer to pattern-based discovery in Service Mapping. Failure to provide users with required permissions may cause Service Mapping to perform traffic-based discovery.

Limitations

Discovery and Service Mapping cannot discover the Exchange Server if its configuration does not support public folders, edge transports, and public messaging. Create a pattern or customize the existing pattern to discover the Exchange Server with such configuration.

Map application services containing Microsoft Exchange Servers

Map an application service containing an Microsoft Exchange Server and verify that Service Mapping discovered all Microsoft Exchange Server components correctly.

Before you begin

- On the ServiceNow AI Platform, configure a domain user for accessing the Windows OS as described in Windows credentials. Use the Microsoft Exchange Server administrator's user name and password as Windows credentials. For example, user name: local\administrator and password: Are1300.

  Note: If you configure Windows credentials for a user different from the Microsoft Exchange Server administrator, Service Mapping uses traffic-based discovery that results in mapping all Microsoft Exchange clients.

- (For Exchange CAS 2010 and 2013, 2016) On the ServiceNow AI Platform, configure applicative credentials, setting the CI type to Exchange Mailbox.

  Warning: Do not confuse the Exchange Mailbox CI type with the Exchange Mailbox Server CI type.

  Service Mapping uses these applicative credentials to access the of/Servers/Servers.slab folder on the Key Exchange Admin Center (EAC). Service Mapping also uses the applicative credentials to run the following PowerShell commands against the MID Servers:

  - Get-ExchangeServer
  - Get-ExchangeServer
  - Get-MailboxServer
  - Get-ClusteredMailBoxServerStatus
  - Get-ExchangeServer
  - Get-StorageGroup
  - Get-MailboxDatabase

- (For Microsoft Exchange CAS 2010 and 2016) On the Windows Server hosting Microsoft Exchange CAS, install the Microsoft.Exchange.Management.PowerShell.E2010 module.
- (For Microsoft Exchange CAS 2007) On the Windows Server hosting Microsoft Exchange CAS, install the Microsoft.Exchange.Management.PowerShell.Admin module.
- (For Microsoft Exchange CAS 2007, 2010) On the ServiceNow AI Platform, upload the rctrlx tool for running PowerShell commands as described in Upload the rctrlx.exe file to MID Servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- (For Microsoft Exchange CAS 2007 and 2010) Provide an Exchange Mailbox OS user with the rights to run Exchange services on Windows and has permission to query Exchange repository:

  On the ServiceNow AI Platform, configure Windows credentials setting the *Credential alias* attribute to [cmdb_ci_exchange_mailbox]. Configuring this attribute allows you later to use the credential for this CI in patterns.

- (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
- Verify that you know the entry point necessary for discovering the Microsoft Exchange Server. Use the URL for the Outlook web access in the following format:

  `http(s)://<hostname>:<port>/owa`

  For example, `http(s)://myhost:<443>/owa`

Role required: service_mapping_admin

Procedure

1. Create a service instance.

   a. Navigate to **All** > **Service Mapping** > **Service Instances**.

   b. Select **New**.
   The **Discoverable by Service Mapping** option is selected by default.

   c. In the **Name** field, enter a descriptive name for the service instance.

2. Define the entry point for this service instance:

   a. Select the **Web Application** tile.
   The **Discoverable by Service Mapping** option is selected by default.

   b. In the **URL** field, enter the URL of the Outlook web access.
   For example, http(s)://myhost:<443>/owa.

   c. Click **Add**.

   d. Click **Save**.

3. Click **Additional Info** on the left and clear the **Traffic based discovery** check box.
4. Click **Update**.
5. Verify that the mapping result is satisfactory:

   a. Click **View Map** and wait until the mapping is complete.
   It may take a few minutes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Review the application service map and make sure that the service layout is as in the examples below.

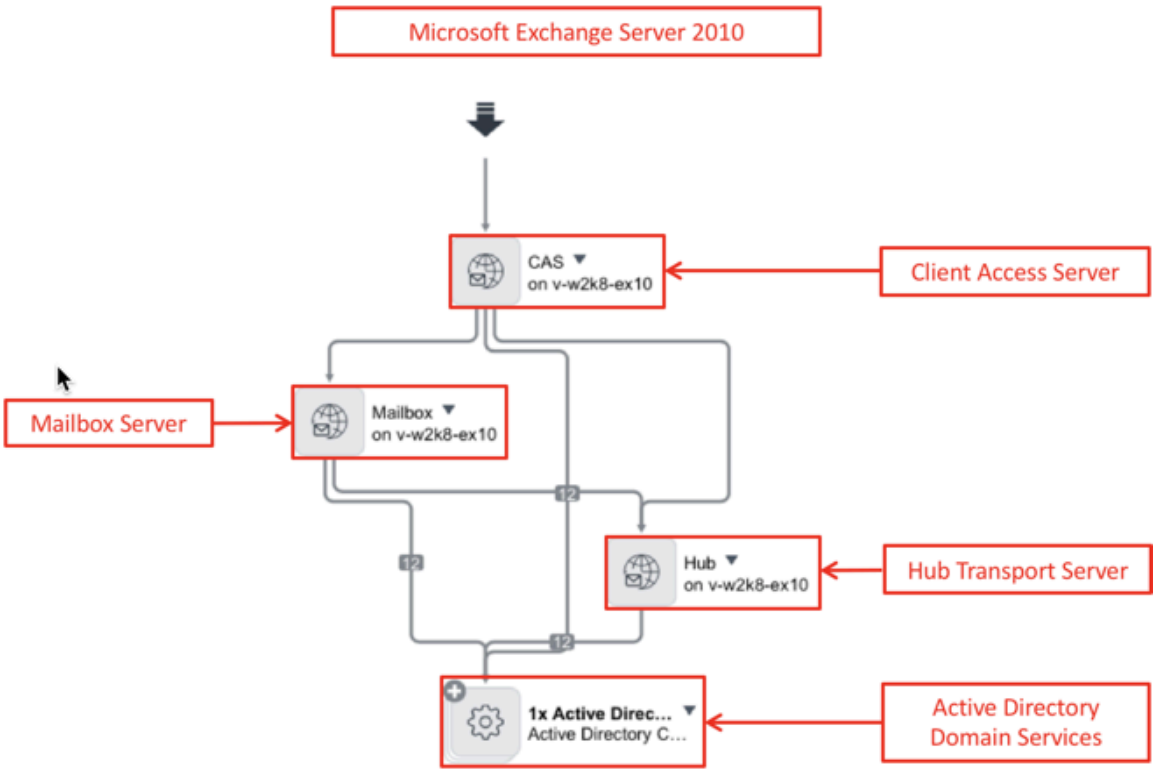

Trouble?

If the mapping result is not as you expected, perform basic troubleshooting.

Symptom	Cause	Solution
Discovery failed. The discovery log displays the message that starts with the following text: You must have the Exchange Common Commands PowerShell Snapin installed to use this script.	The indication sections of the pattern fail, because the user does not have permissions for running PowerShell commands.	Make sure that you provided a user with permissions to run PowerShell commands on servers that host Microsoft Exchange Server components, for example, CAS.
Discovery failed. The discovery log displays the error message that starts with the following text: [servername.destination.com] Connecting to remote server failed with the following error message: The WinRM client can not process the request. The authentication mechanism requested by the client is not supported by the server or unencrypted traffic is disabled in the service configuration. Verify the unencrypted	The indication sections of the pattern fail, because the MID Server is not configured as a trusted host on the component you are discovering.	Add the MID Server to the list of trusted hosts on all hosts that you want to discover.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
traffic setting in the service configuration or specify one of the authentication mechanisms supported by the server.		
The application service map is very large, because it includes connections to all Outlook clients.	(In deployments using Microsoft Exchange Server 2007, 2010 or 2013), Service Mapping created connections using traffic-based discovery, which traces all open connections, instead of using PowerShell commands.	Make sure that you provided a user with permissions to run PowerShell commands on servers that host Microsoft Exchange Server components, for example, CAS.

Tibco BusinessWorks and EMS discovery

Discovery can find Tibco ActiveMatrix BusinessWorks and Enterprise Message Service (EMS). Service Mapping can discover application services containing Tibco BusinessWorks, Enterprise Message Service (EMS), and their components.

Data collected by Discovery for Tibco ActiveMatrix BusinessWorks

The following data is collected in the ActiveMatrix Business Works `[cmdb_ci_appl_tibco_matrix]` table.

Table and field	Description
ActiveMatrix Business Works [cmdb_ci_appl_tibco_matrix]	
Name [name]	The attributes of ActiveMatrix BusinessWorks configuration item.
Version [version]	
TCP port(s) [tcp_port]	
Class [sys_class_name]	
Fully qualified domain name [fqdn]	
IP Address [ip_address]	
Configuration directory [config_directory]	
Configuration file [config_file]	
Installation directory [install_directory]	
Project [project]	

Data collected by Discovery for Tibco Enterprise Message Service

By default, Discovery uses the *`Enterprise Message Service`* pattern to discover Tibco Enterprise Message Service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Table and field	Field Name
Tibco Enterprise Message Services [cmdb_ci_appl_tibco_message]	
Name [name]	Attributes of the Tibco Enterprise Message Services configuration item.
Version [version]	
Class [sys_class_name]	
Fully qualified domain name [fqdn]	
IP Address [ip_address]	
Installation directory [install_directory]	
TCP port(s) [tcp_port]	
Configuration directory [config_directory]	
Configuration file [config_file]	
Installation directory [install_directory]	

Data collected by Service Mapping

Service Mapping can discover application services that contain Tibco ActiveMatrix BusinessWorks (including BusinessWorks Process), Tibco Enterprise Message Service (EMS), Tibco Enterprise Message Service (EMS) Queue.

Map application services containing Tibco BusinessWorks and EMS

Map an application service containing Tibco BusinessWorks and Enterprise Message Service (EMS) and verify that Service Mapping discovered all Tibco components correctly.

Before you begin

- (If hosted on Unix) Configure applicative credentials for the Tibco EMS.
- (If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
- (If hosted on a Unix server) On the ServiceNow AI Platform, configure SSH credentials.
- (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

Role required: service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create a service instance.
   a. Navigate to **All** > **Service Mapping** > **Service Instances**.
   b. Select **New**.
   The **Discoverable by Service Mapping** option is selected by default.
   c. In the **Name** field, enter a descriptive name for the service instance.
2. Define the entry point for this application service:
   a. Select the **Other application** tile.
   The **Discoverable by Service Mapping** option is selected by default.
   b. From the Select Entry Point Type list, select `HTTP(s) Endpoint, TCP Endpoint, Tibco BW Endpoint`.
   c. Define the entry point as follows:
      - The URL to the process that correlates to the Tibco BusinessWorks Project, for example, `http://10.xx.xx.xx:6710/GetMemberActivitiy/Processes/ServiceGetMemberActivitiy.service` .
      - (Optional) The hostname on which the BusinessWorks Project runs as a service.
   d. Click **Add**.
   e. Click **Save**.
3. Click **Additional Info** on the left and clear the **Traffic based discovery** check box.
4. Click **Update**.
5. Verify that the mapping result is satisfactory:
   a. Click **View Map** and wait until the mapping is complete.
   It may take a few minutes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Review the application service map and make sure that the service layout is as in the following examples.

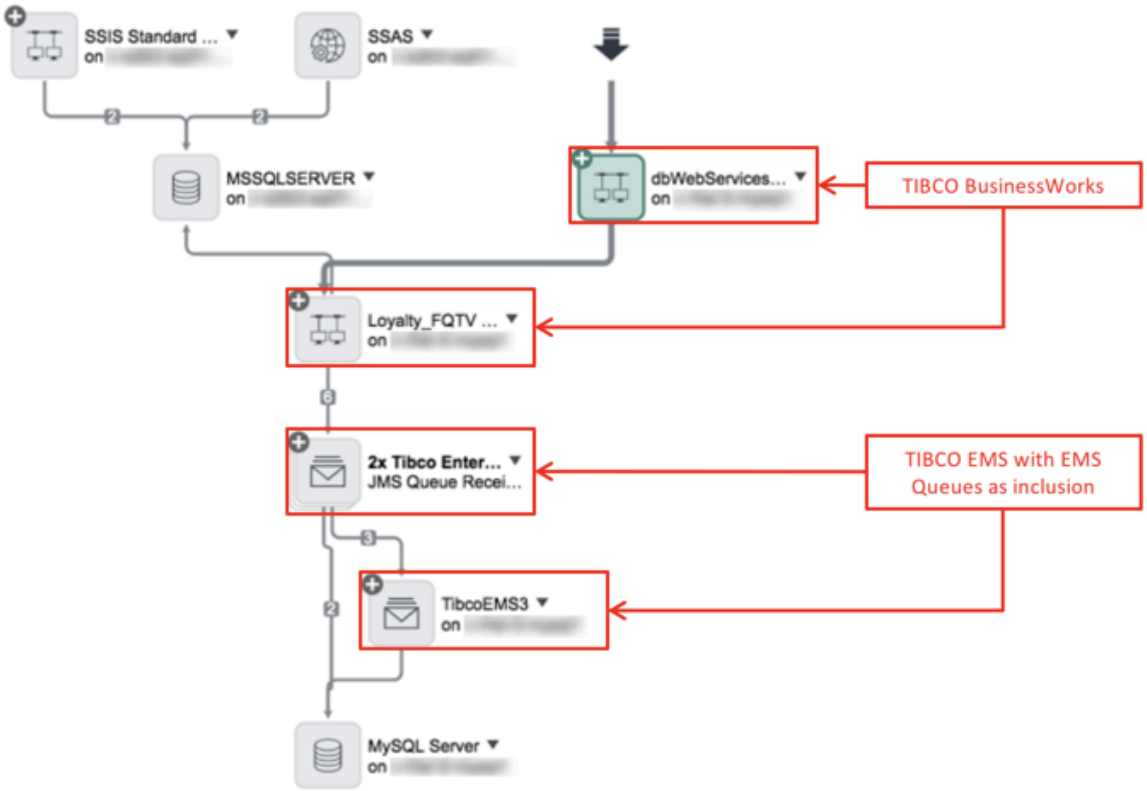

In addition, verify that the BusinessWorks CIs are inclusions each of which contain a BusinesWorks process. Also, verify that the EMS CIs are inclusions containing EMS Queues.

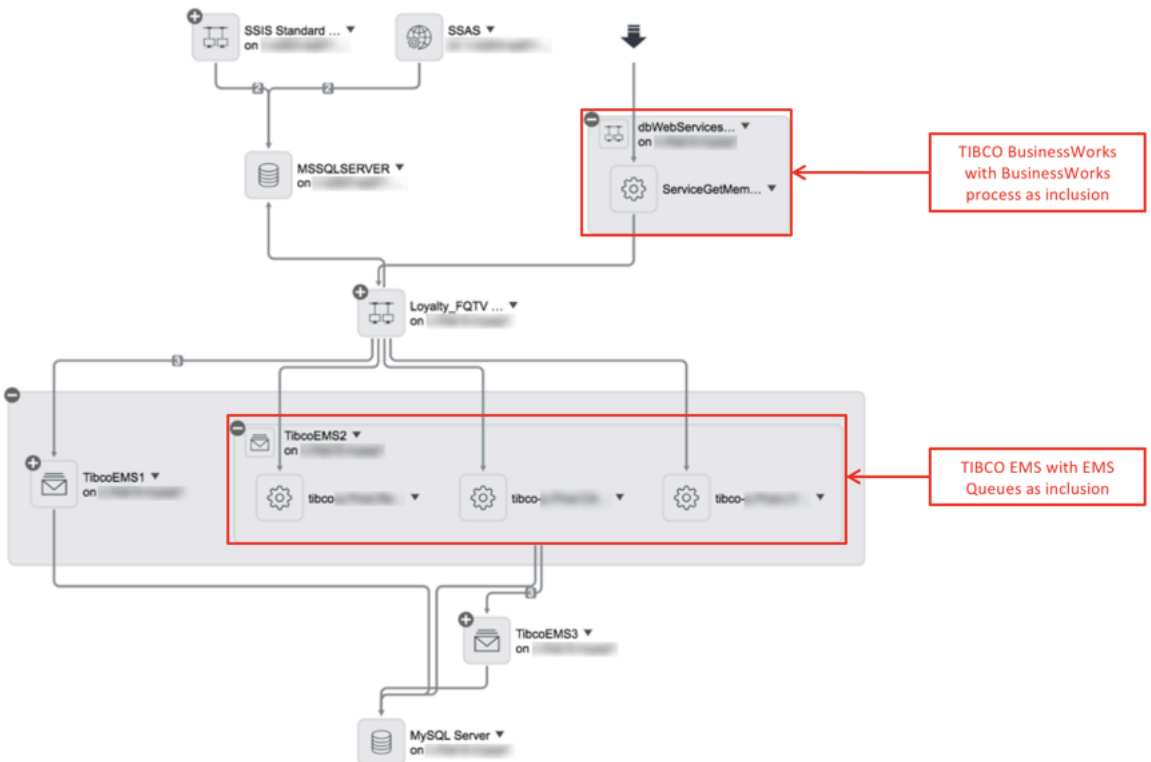

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Trouble?

If the mapping result is not as you expected, perform basic troubleshooting.

Symptom	Cause	Solution
(Only for deployments hosting BusinessWorks on a Unix server) The service instance map shows the BusinessWorks CI without an inclusion of the BusinessWorks process. The discovery log shows the error message stating that the user does not have sufficient credentials.	Service Mapping failed to run the `TibcoFilesParser.ksh` command.	Make sure that you provided a user with elevated rights as described in Service Mapping commands not requiring a privileged user.
Service Mapping failed to discover the outgoing connections of EMS to the EMS process.	Service Mapping failed to run the `tibemsadmin` command.	Make sure that you provided a user with elevated rights as described in Service Mapping commands not requiring a privileged user.

Application services with Oracle GlassFish Servers

Service Mapping can discover application services containing Oracle GlassFish Servers.

Discovered components

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Oracle GlassFish Server	Windows UNIX	3.1, 4.0, 4.1	WMI WinRM SSH	Horizontal and top-down	GlassFish Server
Oracle GlassFish WAR file	Windows UNIX	3.1, 4.0, 4.1	WMI WinRM SSH	Horizontal and top-down	GlassFish WAR

Map application services containing Oracle GlassFish Servers

Map an service instance containing an Oracle GlassFish Server correctly.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- (If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials.
- (If hosted on a Linux or UNIX server) (Optional) on the ServiceNow AI Platform, configure SSH credentials.
- (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

Role required: service_mapping_admin

Procedure

1. Create a service instance.

   a. Navigate to **All** > **Service Mapping** > **Service Instances**.

   b. Select **New**.
   The **Discoverable by Service Mapping** option is selected by default.

   c. In the **Name** field, enter a descriptive name for the service instance.

2. Define the entry point for this application service:

   a. Select the **Other Application** tile.
   The **Discoverable by Service Mapping** option is selected by default.

   b. Define the entry point by filling in the fields.

Field	Description
Select Entry Point Type	Entry point type: ▪ HTTP(S) Endpoint ▪ TCP Endpoint ▪ Tibco BW Endpoint
URL	The URL of the process that correlates with the GlassFish Server, for example, `http://10.196.xx.xx:28080/hello/`. This field appears when **HTTP(S) Endpoint** is selected in **Select Entry Point Type**.
Host	The server that hosts the GlassFish Server. This field appears when **Tibco BW Endpoint** or **TCP Endpoint** is selected in **Select Entry Point Type**.
Host Name	The host name of the server that hosts the GlassFish Server. This field appears when **HTTP(S) Endpoint** is selected in **Select Entry Point Type**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Port	The port, which the process correlating with the GlassFish Server uses. This field appears when **Tibco BW Endpoint**o r **TCP Endpoint** is selected in **Select Entry Point Type**.
Process Name	The name of the process that correlates with the GlassFish Server. This field appears when **Tibco BW Endpoint** is selected in **Select Entry Point Type**.
Comments	Free text comment to provide information on the endpoint. This field appears when **HTTP(S) Endpoint** or **TCP Endpoint** is selected in **Select Entry Point Type**.

c. Click **Add**.

d. Click **Save**.

3. Click **Additional Info** on the left and clear the **Traffic based discovery** check box.
4. Click **Update**.
5. Verify that the mapping result is satisfactory:

a. Click **View Map** and wait until the mapping is complete.
It may take a few minutes.

b. Review the application service map and make sure that the service layout is as in the following example.

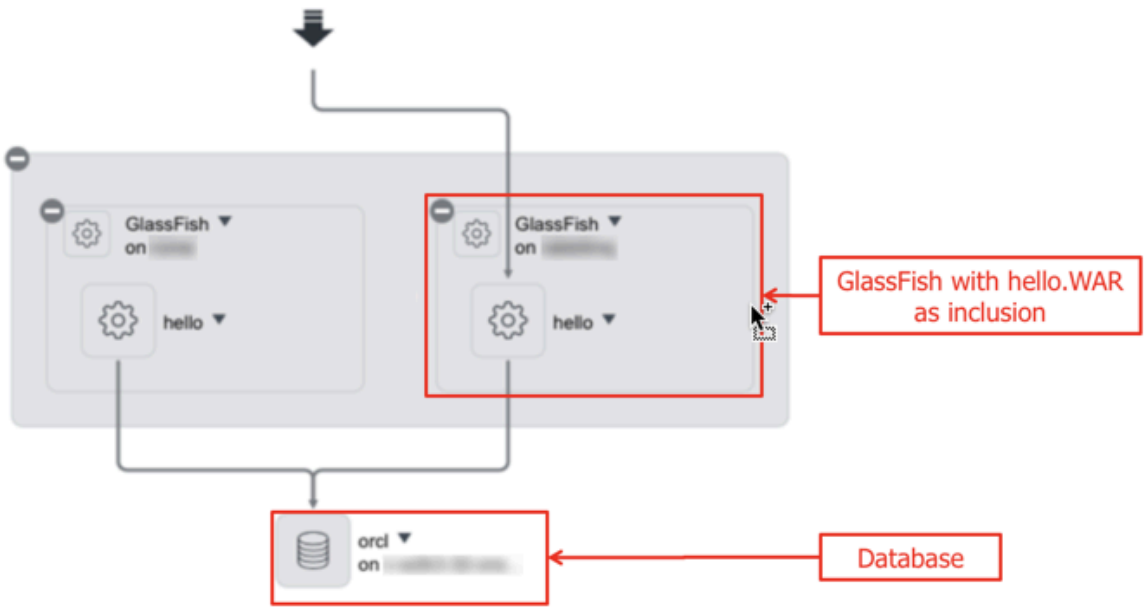

In addition, verify that the GlassFish Server CIs are inclusions, each of which contains a WAR file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Trouble?

If the mapping result is not as you expected, perform basic troubleshooting.

Symptom	Cause	Solution
Service Mapping fails to discover the GlassFish Server CI. The discovery log displays the following message: Process is not running or port is not correctFailed to identify CI due to the host 10.196.xx.xx lacking an application running on port 28080. Make sure that the application is running and run quick discovery on the host.	There is no process running on the port that is part of the entry point.	1. Make sure that the entry point is correct. 2. If the entry point is correct, check that here is a process running on the port by running a quick discovery on the host.

Mapping application services containing Oracle Tuxedo platform

Service Mapping can discover application services containing Oracle Tuxedo platform and its components.

Discovered components

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Oracle Tuxedo	Windows UNIX	10.x, 11.x, 12.x	WMI WinRM SSH	Horizontal and top-down	Tuxedo
Oracle Tuxedo Portal	Windows UNIX	10.x, 11.x, 12.x	WMI WinRM SSH	Horizontal and top-down	Tuxedo Portal pattern

Map application services containing Oracle Tuxedo platform

Map an application service containing an Oracle Tuxedo platform and verify that Service Mapping discovered all Tuxedo components correctly.

Before you begin

- On the ServiceNow AI Platform, configure SSH credentials.
- (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
- (If hosted on Unix) Provide a user with the elevated rights for running commands described in Service Mapping commands requiring a privileged user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: service_mapping_admin

Procedure

1. Create a service instance.
   a. Navigate to **All** > **Service Mapping** > **Service Instances**.
   b. Select **New**.
   The **Discoverable by Service Mapping** option is selected by default.
   c. In the **Name** field, enter a descriptive name for the service instance.
2. Define the entry point for this application service:
   a. Select the **Other Application** tile.
   b. Select `TCP Endpoint` from the list.
   c. Define the entry point as follows:
      - A host of the server that hosts the Tuxedo process.
      - A port on this server that the Tuxedo process opens.
   d. Click **Add**.
   e. Click **Save**.
3. Click **Additional Info** on the left and clear the **Traffic based discovery** check box.
4. Click **Update**.
5. Verify that the mapping result is satisfactory:
   a. Click **View Map** and wait until the mapping is complete.
   It may take a few minutes.
   b. Review the application service map.

Trouble?

If the mapping result is not as you expected, perform basic troubleshooting.

Symptom	Cause	Solution
Service Mapping fails to discover the Tuxedo CI. The discovery log displays the message that states that the credentials are insufficient for running the command including "permissions.tmadadmin -v".	The provided user does not have elevated permissions for running this command.	Make sure that you provided a user with permissions to run commands requiring elevated rights.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
Service Mapping fails to discover the Tuxedo CI. The discovery log displays the message that states that the user cannot run the tux.env script.	The provided user does not have the permissions to run the tux.env script.	Make sure that you provided a user with permissions to run the tux.env script.

Modify or update tag definitions for tag-based mapping

You can change tag definitions that Service Mapping uses for tag-based mapping. Service Mapping uses updated tag definitions to create new tag-based services without applying tag-related changes to services you mapped earlier.

Before you begin

If you need to create a new tag category or modify an existing tag category, follow the prerequisites in Map application services using tags with classic Service Mapping.

Review recommendations in Prepare for mapping application services based on tags.

Role required: service_mapping_admin

About this task

Note: Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

You may decide to add or remove tag categories from a tag-based service family in the following cases:

- If you are not satisfied with the results of the tag-based mapping, for example, if the service is too large, because it includes irrelevant CIs.
- If there are changes in the way your organization uses tags, for example there are new tags.

When you modify tag categories defined for a tag-based service family or delete a tag-based family itself, the system detaches all mapped services from this tag-based service family. These services still exist and Service Mapping continues to update them using the original tag definition. To avoid having tag-based services based on outdated tag definitions, delete existing tag-based services.

Service Mapping generates new tag-based service candidates that you use to map services based on modified tag definition.

Important: New services based on modified tag definitions do not overwrite previously mapped services.

Procedure

1. Delete existing tag-based services:
   a. Navigate to **Service Mapping** > **Administration** > **Tag-based Service Families**.
   b. Select the tag-based service family for which you want to modify the tag definition.
   c. Under **Mapped services**, click **Delete mapped services** to delete all services.
   d. Alternatively, click the **Mark for deletion** icon next to the relevant service in the list, and then click **Delete**.
2. Navigate to **Service Mapping** > **Administration** > **Tag-based Service Families**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select the tag-based service for which you want to modify the tag definition.
4. Under **Selected tag categories**, perform one of the following actions:
   - Add a new tag category by double-clicking the empty row and selecting the relevant tag category from the list of tag categories you defined earlier.
   - Remove a tag category by clicking the **Mark for deletion** icon next to it.
   - Modify the use of a tag category already associated with the service family by changing the tag values or the service naming order.

     For multiple tag values, use comma as a separator. The service naming order defines how Service Mapping uses the tag category name in the names of application services generated for the service family.

   - Modify the tag category itself by clicking the tag category name and then defining the CI tag keys.
5. Click **Update**.

What to do next
Remap tag-based application services to reflect tag changes

Remap tag-based application services to reflect tag changes

Recreate tag-based services if the way your organization uses tags have changed or after modifying tag definitions for tag-based service families.

Before you begin
If necessary, modify the tag definitions for the relevant tag-based service family.

Role required: service_mapping_admin

About this task
Remap existing tag-based services after you have manually modified tag definitions. If your organization uses new tag keys or tag values, you may also want to remap existing services. In both cases, Service Mapping does not automatically update existing services to reflect tag changes.

Procedure
1. Navigate to **Service Mapping** > **Administration** > **Tag-based Service Families**.
2. Select the tag-based service family for which tag definitions have changed.
3. Under **Related Links**, click **Manually update candidates** to generate new service candidates using the modified tag definitions.
4. In the Service candidate window, review the suggested service candidates and select the ones that you want to map.
   If there are no service candidates, ensure that you selected the tag values that exist in the Key Value [cmdb_key_value] table.
5. Click **Map selected**.
   Service Mapping creates application services based on the candidates you selected.
6. Under **Mapped services**, review the newly created services to ensure that the result is as you expected.
   If Service Mapping offers to view the CI list instead of a map for an application service, this service is too large. Fine-tune it as described in modify the tag categories.

Result
You have mapped the application service based on the latest tag definitions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remapped tab-based application service

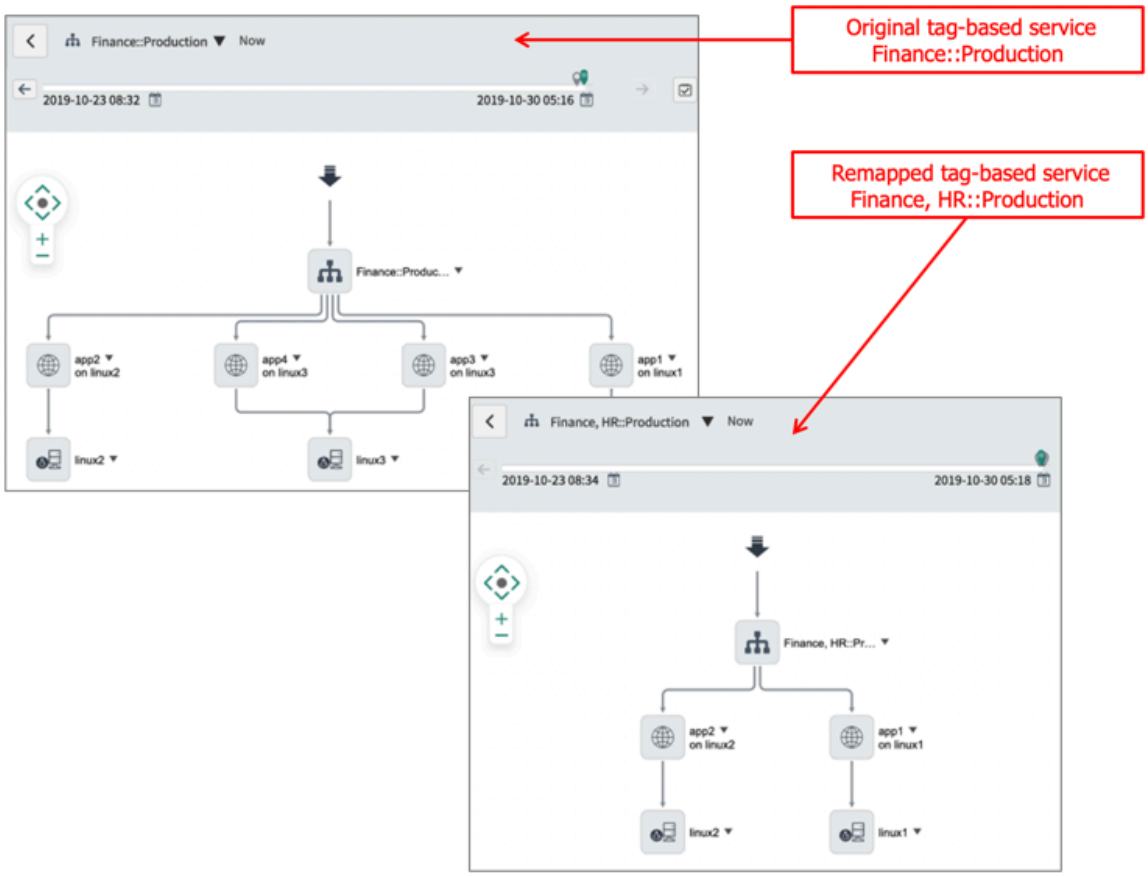

What to do next

Make sure that the service instances aren't too large:

- Service Mapping doesn't offer to view CI list instead of a map for a service instance.
- There is no discovery message indicating that the service instance is too large: *`The map does not display the entire service, because it is too large. The number of CI connections exceeded the allowed maximum.`*

If the service is too large, perform the following actions:

- Modify tag definitions used for mapping to make sure Service Mapping includes only relevant CIs in the service instance.
- If necessary, Tag-based discovery configuration to change which CIs Service Mapping includes in tag-based service instances.

Fix application service errors in bulk

Service Mapping classifies errors by their root cause, for example, missing credentials or task timeout. For a fast and efficient process, fix errors belonging to the same category in bulk.

Before you begin

Role required: service_mapping_admin or admin

About this task

You can view errors by category, for example, **Configuration**, **Network**, and **Credentials**. Each category tile also displays the error counter.

As part of the Service Mapping workflow, after you map application services in bulk, you can fix multiple application service errors in bulk (as a group).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fixing errors in bulk workflow

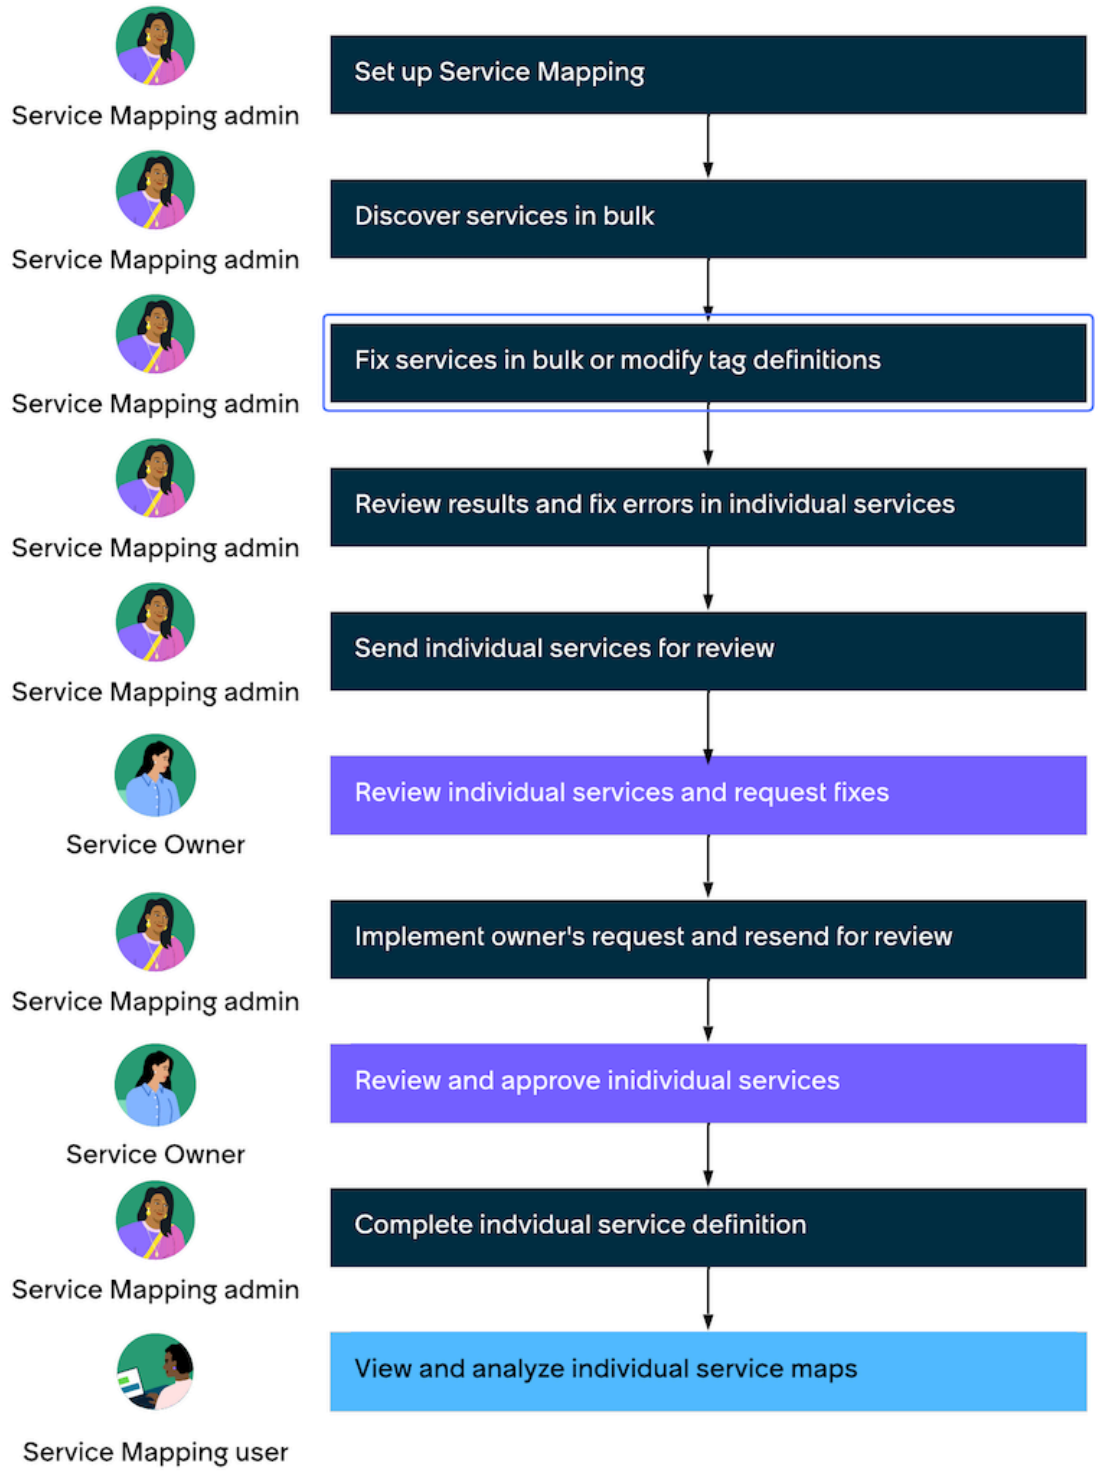

After performing initial mapping, Service Mapping rediscovers configuration items (CIs) regularly, constantly updating the list of errors. Discovery schedules define how often Service Mapping runs the discovery process for CIs. Service Mapping discovers only CIs belonging to application services with Operational status.

There may be discovery errors not assigned to any category:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- After upgrading Service Mapping from Jakarta release or earlier.
- If there are configuration items (CIs) that have errors associated with them, and that were first removed from application services and then readded to application services.

Procedure

1. Navigate to **All** > **Service Mapping** > **Home** and then click **Fix Your Services** in the **Fix** box.

   The number in the **Fix** box indicates the number of affected service maps. Click **Update** ( ) to update the display.

   Note: To fix the errors for a particular service instead, click **Manage Errors by Service**.

2. If there are uncategorized errors, you can categorize them by service instance:

   a. Click **View Uncategorized Errors by Service**.

   b. Click the relevant service name.

   c. On the service instance map, click **Run Discovery**.
   Service Mapping discovers all CIs discovered with errors and belonging to this service instance.

3. If there are uncategorized errors, categorize them by clicking **Rediscover Relevant CIs**, and clicking **Continue** in the confirmation message.
   Service Mapping discovers all CIs discovered with errors and belonging to application services with Operational status.
4. Click an error category box.
   In the example below, configuration errors are selected.
   - The Errors by Category page organizes errors by high, medium, and low impact on the services.
   - For each impact category, errors are grouped into boxes by error code (possibly the root cause).
   - If error tasks have been assigned, the task numbers appear in the box. Click a task to view details. Point to a task to view comments that the assigner might have added. A check mark indicates that a task is completed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

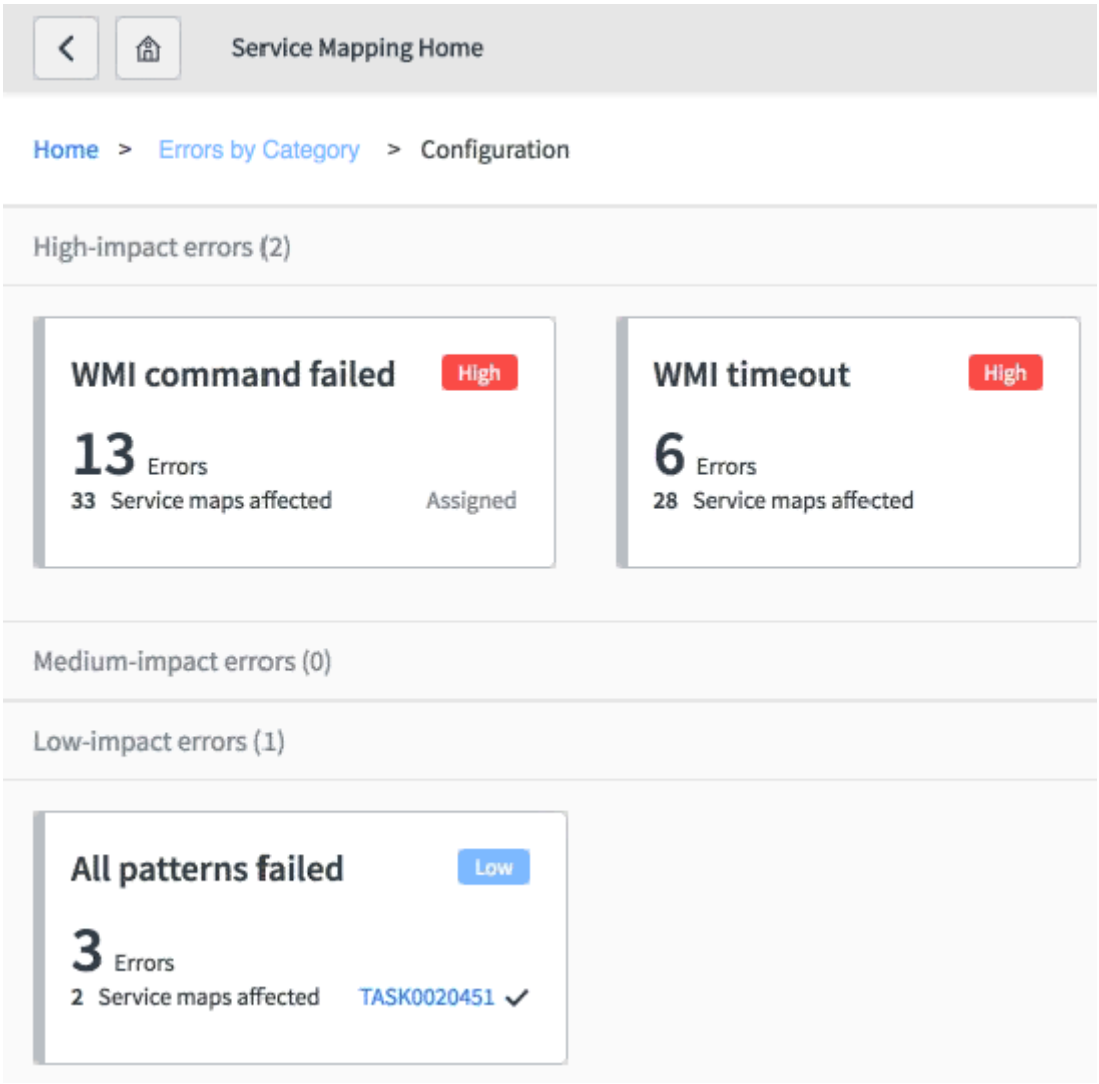

5. To view the list of errors that share an error code, click the box.
   You can perform one of the **Recommended Actions** on items that you select or on all items in the list.

   In the example, **All patterns failed** is selected.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

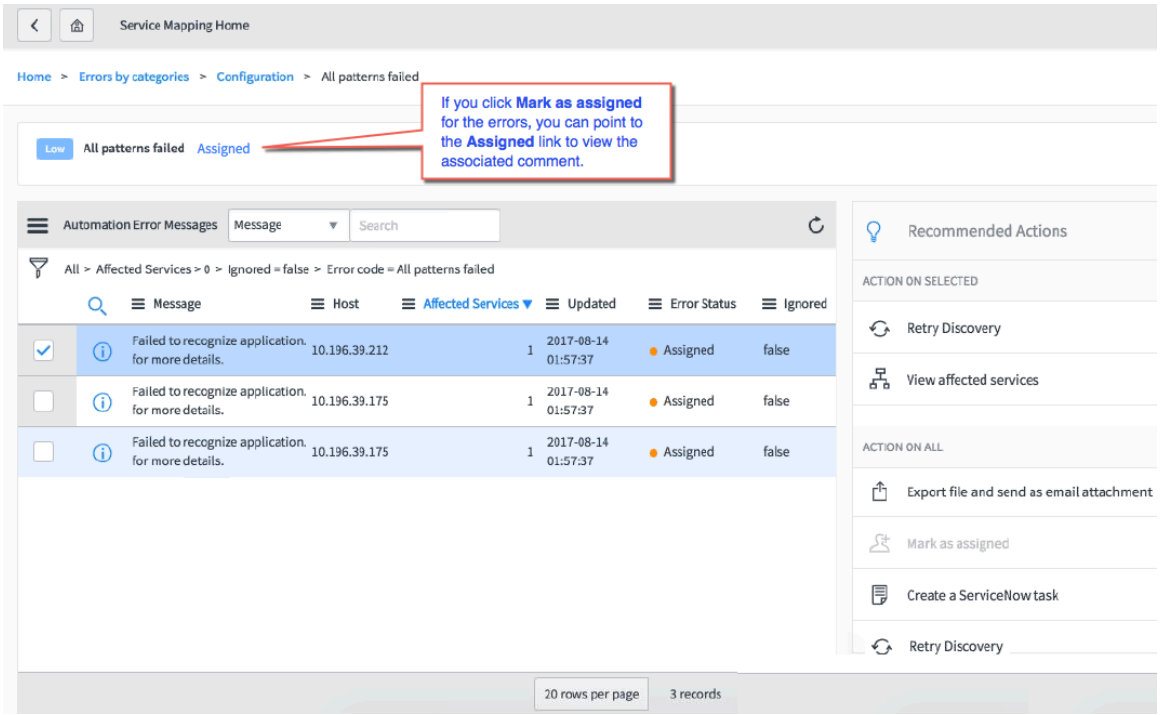

6. Click an action in the **Action on All** section, or select errors and click an action in the **Action on Selected** section.

Retry Discovery

Performs the full Discovery process. You typically perform Discovery after following the fix instructions that appear when you click **View instructions**.

View affected services

Displays a popup list of all services that are affected by the selected error. In the list, click a service to view service details.

Mark as assigned

Apply the **Mark as assigned** action when you plan to perform an out-of-band action to notify a person to fix the error (typically a person who does not use the ServiceNow instance). You can enter a **Comment** in the **Mark as Assigned** popup. Click **Update** to set the **Error Status** field in the error record as **Assigned** and associate the **Comment** with the error. To view the comment, point to the **Assigned** link.

Note: The **Create ServiceNow task** action also sets the **Error Status** field in the error record as **Assigned**.

Create a ServiceNow task

Opens the ServiceNow Error task form in a pop-up window so you can create and assign a task for the selected errors. Sets the **Error Status** field in the error record as **Assigned**. For information on creating error tasks, see Create a table and Create a task from an incident, problem, or change request.

Note: The **Mark as assigned** action also sets the **Error Status** field in the error record as **Assigned**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

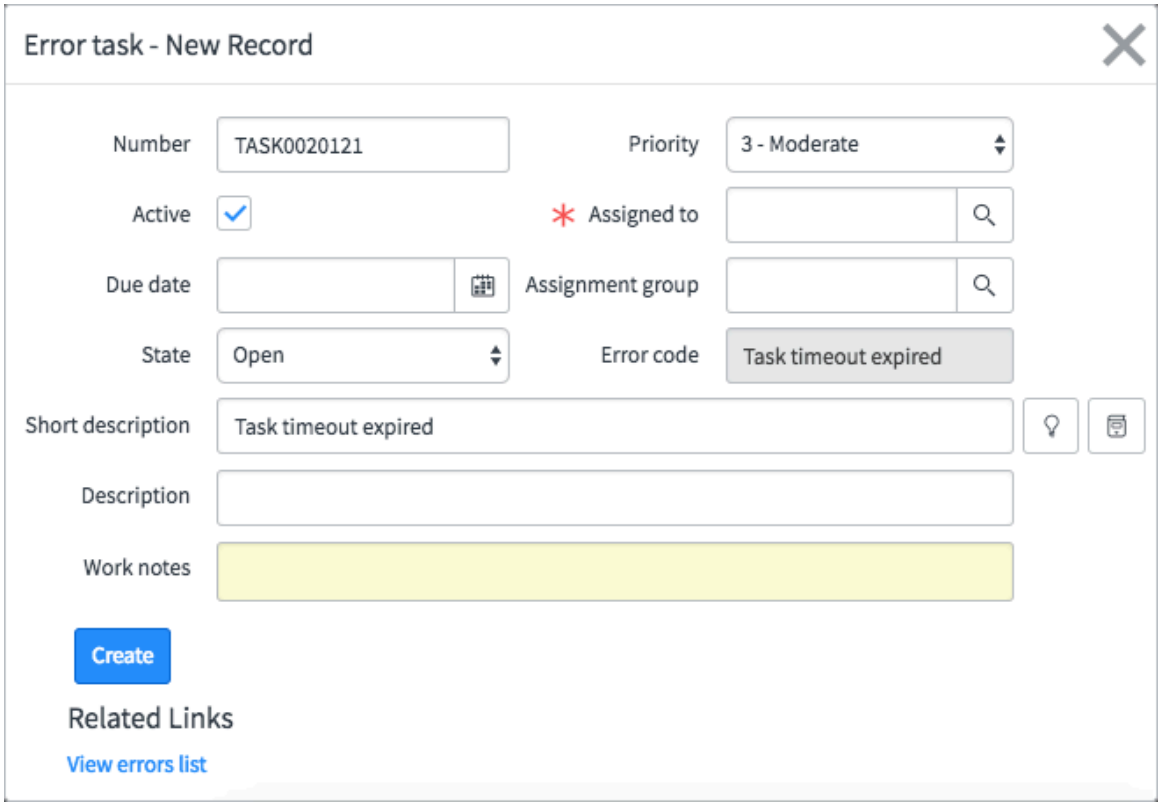

Retry Discovery

Performs the full Discovery process. You typically perform Discovery after following the fix instructions that appear when you click **View instructions**.

View instructions

Displays a troubleshooting procedure in a pop-up window. Follow the steps to try to resolve the error.

Ignore Error

Right-click an error and then select **Ignore Error** to set the **Ignored** value to **true**.

The system performs the action on the items, displays **In progress**, and then displays the result (success or failure) of the action. If the error is resolved, the **Error Status** field in the error record is set to **Resolved**.

Note: If an action (for example, Discovery) is in progress, but taking too long, you can click **Stop** to stop the action.

7. Click **Done** to close the status message.
8. If you performed an action on a selection and the error was resolved, you can click **Resolve All** to apply the same action to all items in the list.

Fix errors in individual application service maps

You can address discovery issues in each application service map individually.

Before you begin
Role required: service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

You can fix errors in individual application services at any time. When you perform initial mapping of application services in your organization, you typically fix errors in individual application services after resolving errors in multiple services, in bulk. You may also need to fix errors in an approved and completed application service at a later stage.

Service Mapping workflow

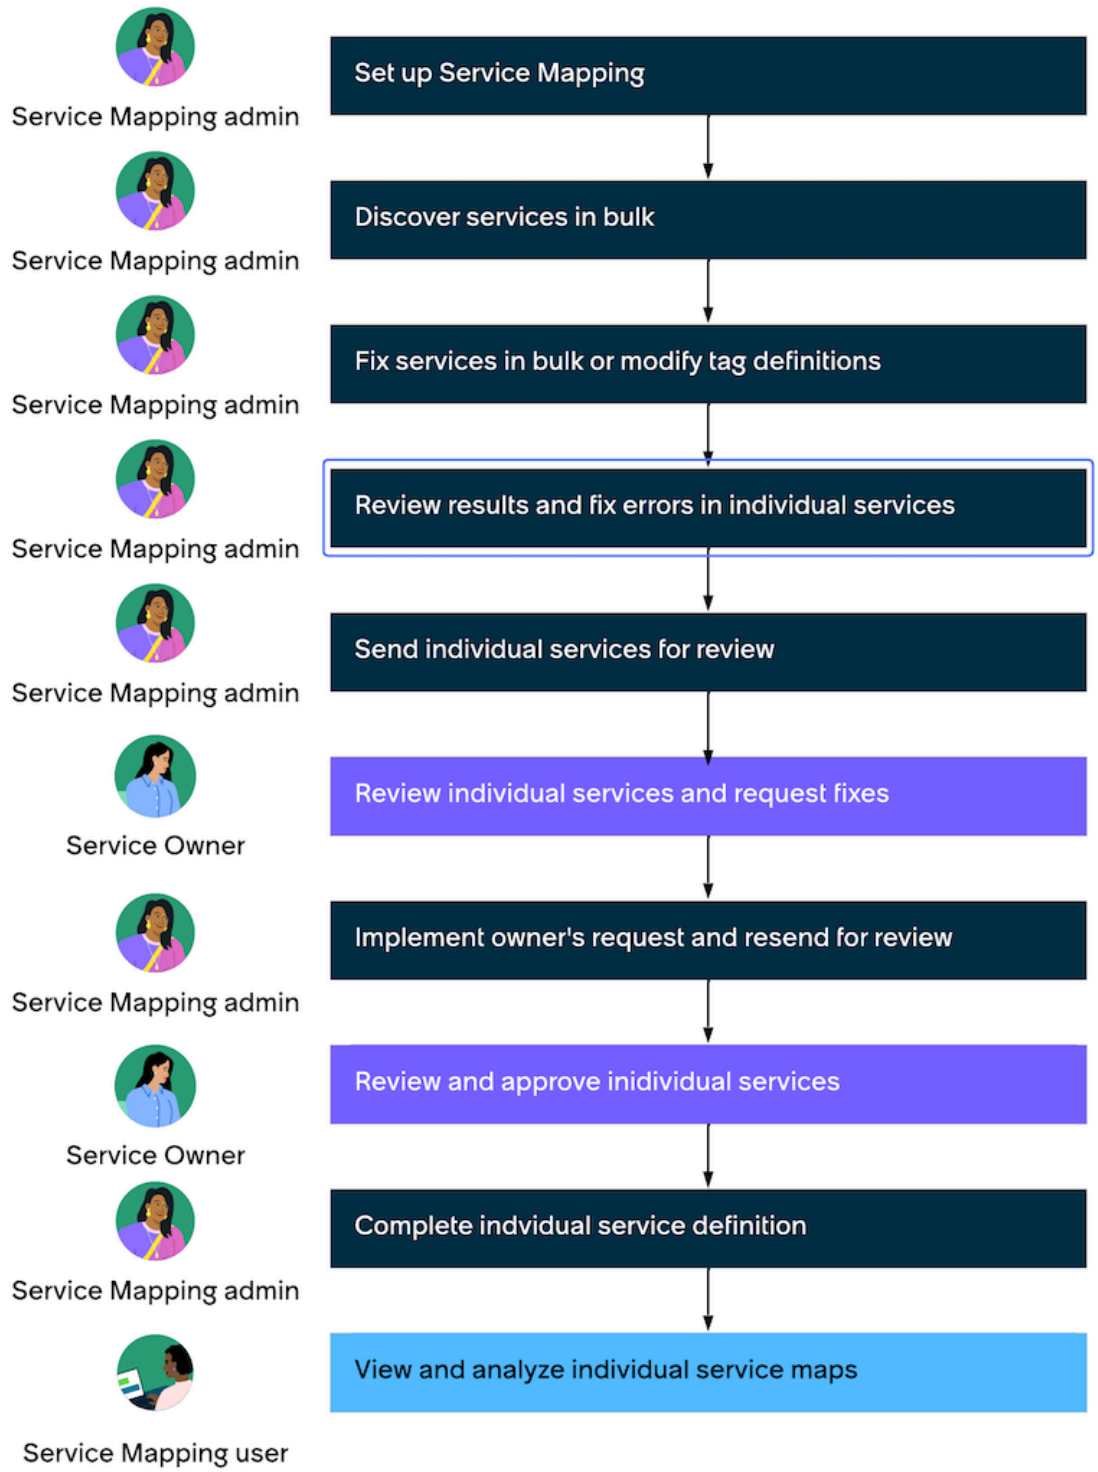

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Ensure that the map opens in Edit mode.

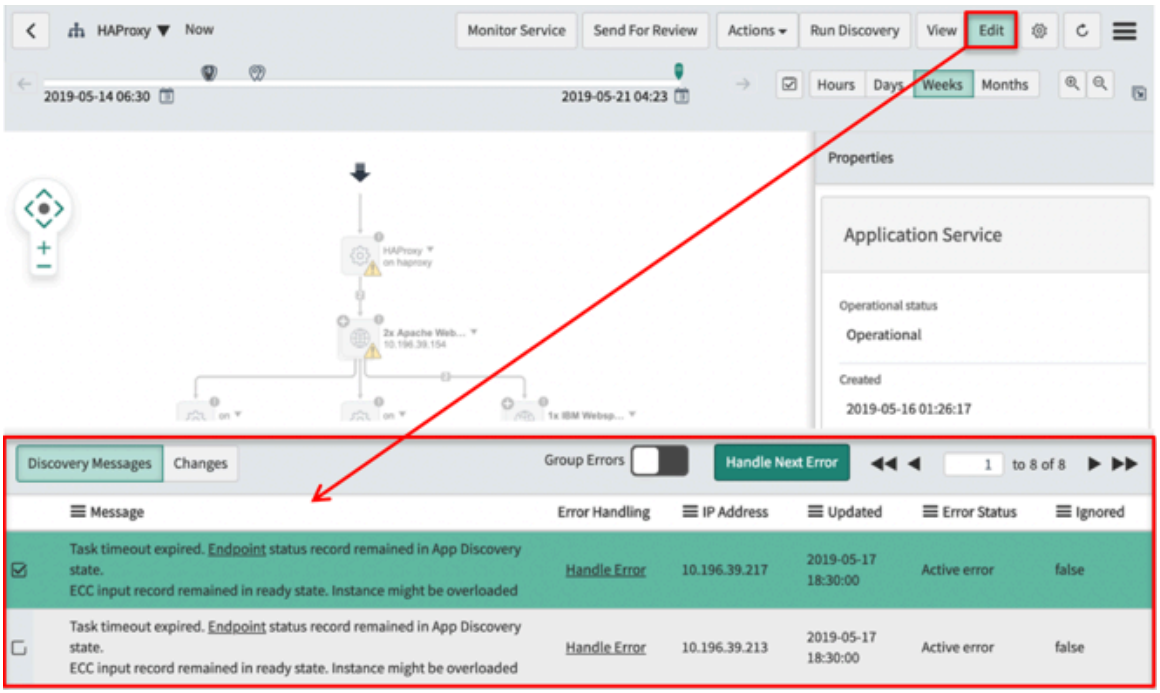

   CIs discovered with errors appear with the warning icon (⚠) on the map. Discovery messages on the **Discovery Messages** tab give short error descriptions.

3. **Optional:** Group errors by types by clicking **Group Errors**.
   The discovery error messages on the **Discovery Messages** tab appear grouped by predefined types, error codes.
4. To review and fix an error from the list, click **Handle Error** next to the error description.
   Or
5. To review the first error in the largest error group, click **Handle Next Error** at the left top corner of the map.
6. To learn how to resolve the error, click **Open instructions** under **Action on Selected** in the right pane.

   Note: Instructions are available only for the most common errors.

7. Click the relevant semi-automated resolution option to resolve the error under **Action on Selected** in the right pane.
   Error resolution options vary, because they depend on the specific discovery error associated with the selected CI.

Use this option	To do this
Add Management IP	Enter the management IP to trigger discovery on the host resources and connections. Typ

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use this option	To do this
	ically, you need this solution for application services that have undiscovered devices like load balancers as their entry points. The man agement IP address of a network device is the address that you use to reach the device using the SNMP or any other management protocol. **Note:** The system does not save the management IP you configure.
Create new WMI credentials	Open the Windows Credentials form and con figure missing Windows credentials for the se lected CI. See
Create new SSH credentials	Open the SSH Credentials form and config ure missing SSH credentials for the selected CI. See
Create new SSH private key	Open the SSH private keys form and config ure missing credentials for the selected CI. See
Create new SNMP credentials	Open the SNMP credentials form and config ure missing SNMP credentials for the select ed CI. See
Show Discovery Log	See the log to understand at what stage an er ror occurred and what caused it.
Skip and resume	Manually add the CI whose discovery failed and configure attributes of the entry point from which Service Mapping continues dis covery. For more information, see Skip errors to continue discovering an application ser vice.
Retry discovery	Run the discovery process for the selected CI after trying to resolve a discovery error associ ated with it.
View affected services	Display a list of all services affected by the se lected error. In the list, click a service to view service details.

8. If there are no instructions or semi-automated resolution options available, fix errors based on symptoms and discovery messages.
9. To ignore an irrelevant error, right-click its error message on the **Discovery Messages** tab and select `Toggle Ignore State`.
   The ignored error disappears from the map and from the list of discovery error messages.

   Note: If necessary, you can still display ignored errors by switching on the **Include Ignored Errors** option under **More Options** menu. For more info, see Modify view for an application service map.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Click **Handle Next Error** to review the next error in the largest error group.
11. When finished reviewing and handling errors, click anywhere in the map to close the **Resolve Error** pane.

What to do next

If you resolved most errors, send the service instance for review.

Related topics

Prerequisites for performing top-down discovery using Service Mapping

Fix service mapping errors using discovery messages

Service Mapping does not offer semi-automated resolution options for errors that require advanced resolution. Fix such errors using symptoms and discovery messages.

Before you begin

Try fixing errors using semi-automated resolution options as described in Fix errors in individual application service maps.

Role required: service_mapping_admin

About this task

You can fix errors in individual application services at any time.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Select **View Map** next to the relevant service instance.
3. Ensure that the map opens in Edit mode.

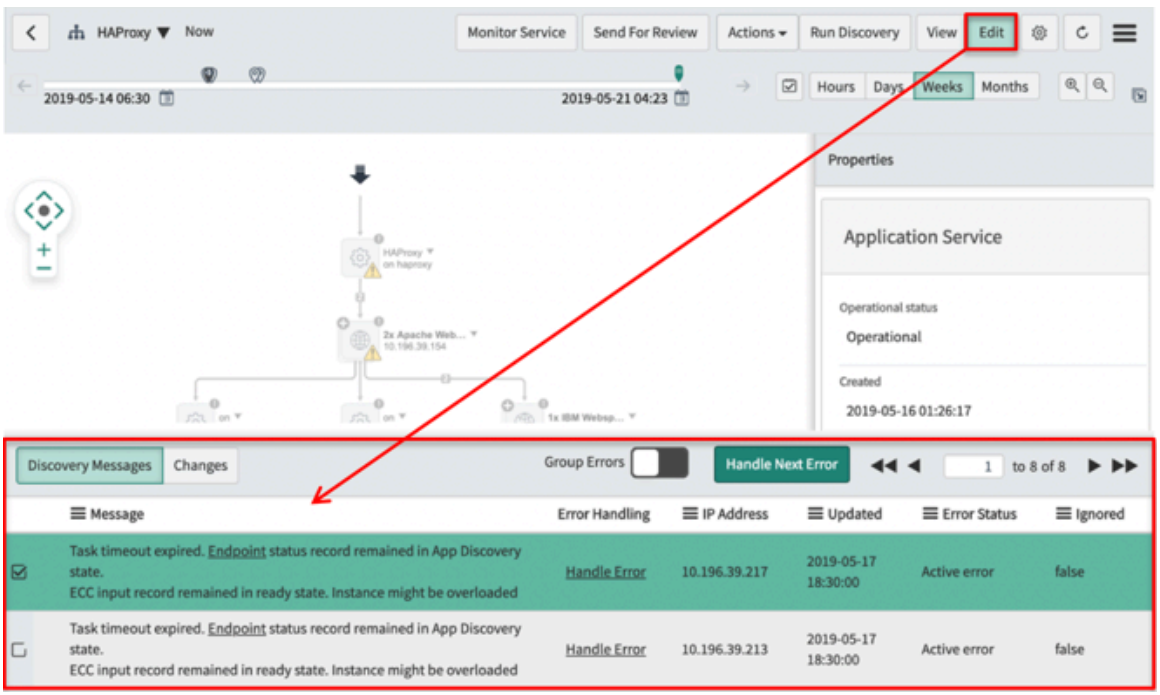

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CIs discovered with errors appear with the warning icon (⚠) on the map. Discovery messages on the **Discovery Messages** tab give short error descriptions.

4. Fix errors based on the symptoms and discovery messages.

Symptom	Resolution
The following discovery message is displayed: `The map does not display the entire service, because it is too large. The number of CI connections exceeded the allowed maximum.`	Reduce the application service size. ○ For an application service mapped using pattern-based and traffic-based discovery, perform the following actions: ▪ If you use customized discovery patterns and mapping results in very large service instances, modify the patterns so that they don't map irrelevant CIs. ▪ Split the service into smaller services. ○ For an application service converted from an application service, perform the following actions: ▪ Review the converted service instance to identify CI relations irrelevant or redundant for this service. Remove such CI relations in the CMDB. ▪ Decide how many levels of related CIs you must include into this service instance. If necessary, change the number of levels used in conversion to reduce the service size.
The following discovery message is displayed in a tag-based service instance: *`This tag-based service instance reached the maximum number of CIs. For more information, refer to the product documentation.`*	Perform the following actions: ○ Modify tag definitions used for mapping to make sure Service Mapping includes only relevant CIs in the service instance. ○ If necessary, Tag-based discovery configuration to change which CIs Service Mapping includes in tag-based service instances.
○ The service instance map displays the warning icon (⚠) on top or instead of the configuration item. ○ The following discovery message displays for the configuration item: `Failed to execute command using sudo on host <host IP address>`.	KB0621669: Resolving failure to execute commands using sudo during Discovery in Service Mapping
The following error message displays for the configuration item: `Access is denied`.	KB0564283: Resolving an issue of denied access to a Windows Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Resolution
○ The service instance map displays the warning icon (⚠) on top or instead of the Windows Server. ○ The following discovery message displays for the Windows Server: `RPC Server unavailable`.	KB0564282: Resolving issue of RPC Server unavailable during Windows Server discovery in Service Mapping
○ The service instance map displays the warning icon ⚠ instead of the load balancer CI. ○ The following discovery message is displayed: `Service Mapping triggered the horizontal discovery to find the host x.x.x.x, because this host was not in the CMDB. The horizontal discovery failed. See discovery status for more info.`	KB0610414: Resolving failure to discover operating system for a load balancer
Instead of the map, the following discovery message displays: `Cannot display the map. Topology too large`.	KB0596671: Resolving a large topology issue in service instance maps
A Virtual IP (VIP) for a load balancer service CI is sometimes updated with the IP addresses of another load balancer.	KB0610412: Resolving load balancers merging into one CI in Service Mapping
The map displays incorrect virtual IP names for load balancers.	KB0610413: Resolving incorrect virtual IP names of load balancers in Service Mapping
The service instance map consists of different CIs every time you run the mapping process on the same entry points to discover the same service instance.	KB0595227: Troubleshooting inconsistent mapping results
There is no load balancer CI on the service instance map. There are connections going directly from the entry point of the service instance to the next tier CIs. The rest of the CIs also disappear from the map.	KB0621529: Resolving the issue of a load balancer disappearing from the application service map
○ The service instance map displays the warning icon (⚠) on top or instead of the configuration item. ○ The following discovery message displays for the configuration item: `SSH command timed out on host`.	KB0621670: Resolving the issue of SSH command time-out during discovery in Service Mapping
A cluster CI appears detached from all CIs in map tiers above or below it.	KB0621531: Resolving disconnected cluster CIs on the service instance map

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Resolution
○ The map displays either the load balancer configuration item (CI) with a warning icon or just the warning icon ⚠ ○ The following discovery message appears for the CI that is expected to be the load balancer in the service instance: `Failed to recognize application. See the discovery log for more details.`	KB0621576: Resolving failure to discover VIP for a load balancer
The service instance map shows a different load balancer from the one you expected to see in this application service.	KB0621616: Business service map displays a wrong load balancer in Service Mapping
○ The service instance map displays the warning icon (⚠) on top or instead of the configuration item. ○ The following discovery message displays for the configuration item: `SSH command timed out on host`.	KB0621670: Resolving the issue of SSH command time out during discovery in Service Mapping
○ The service instance map displays the warning icon (⚠) on top or instead of the configuration item. ○ The following discovery message displays for the configuration item: `No active MID Server found for IP X.X.X.X.`	KB0622808: Resolving a failure to find a MID Server for an IP address
○ The service instance map displays the warning icon (⚠) on top or instead of the configuration item. ○ The following discovery message displays for the configuration item: `'64\rctrlx' is not recognized as an internal or external command, operable program or batch file.`	Perform the procedure described in Upload the rctrlx.exe file to MID Servers.
A CI or a segment of the map is grayed out after Service Mapping runs the top-down discovery on a previously mapped service instance. In this application service Service Mapping failed to rediscover the CI or connections leading to this CI.	Perform the following steps: **a.** (Optional for grayed out segments) Identify the faulty CI that causes the segment to be grayed out. **b.** Use the discovery message associated with this faulty CI to resolve the problem.

5. Click **Run Discovery**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Skip errors to continue discovering an application service

If you know what configuration items (CIs) and connections make up your service instance, you can enable Service Mapping to continue discovery of the service instance even if there are some errors. You can skip errors to troubleshoot later so you can complete mapping most of the service instance, even if some CIs are missing.

Before you begin

Familiarize yourself with the structure of the application service to know which CIs make up the next tier after the problematic CI. If necessary, consult the application service owner to collect necessary information.

Role required: service_mapping_admin or admin

About this task

Normally, discovery of an application service stops when Service Mapping fails to identify a CI. In this case, the map shows a warning icon instead of the undiscovered CI. There are no CIs or connections going out of the undiscovered CI.

Undiscovered CI on the map

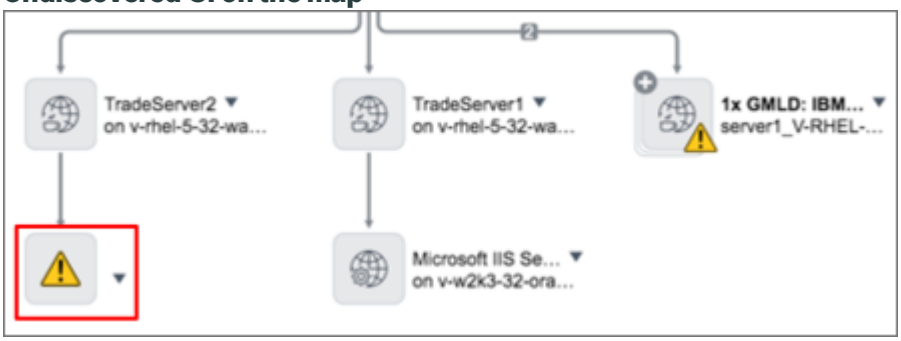

You can overcome the obstacle of the undiscovered CI and restart discovery from the CIs belonging to the next tier after the problematic CI. To do so, add a CI manually if you know what the undiscovered CI is and what the incoming connection of the next tier CI is. Service Mapping uses incoming connections to access a CI for discovery. These connections are referred to as endpoints.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Ensure that the map opens in Edit mode.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

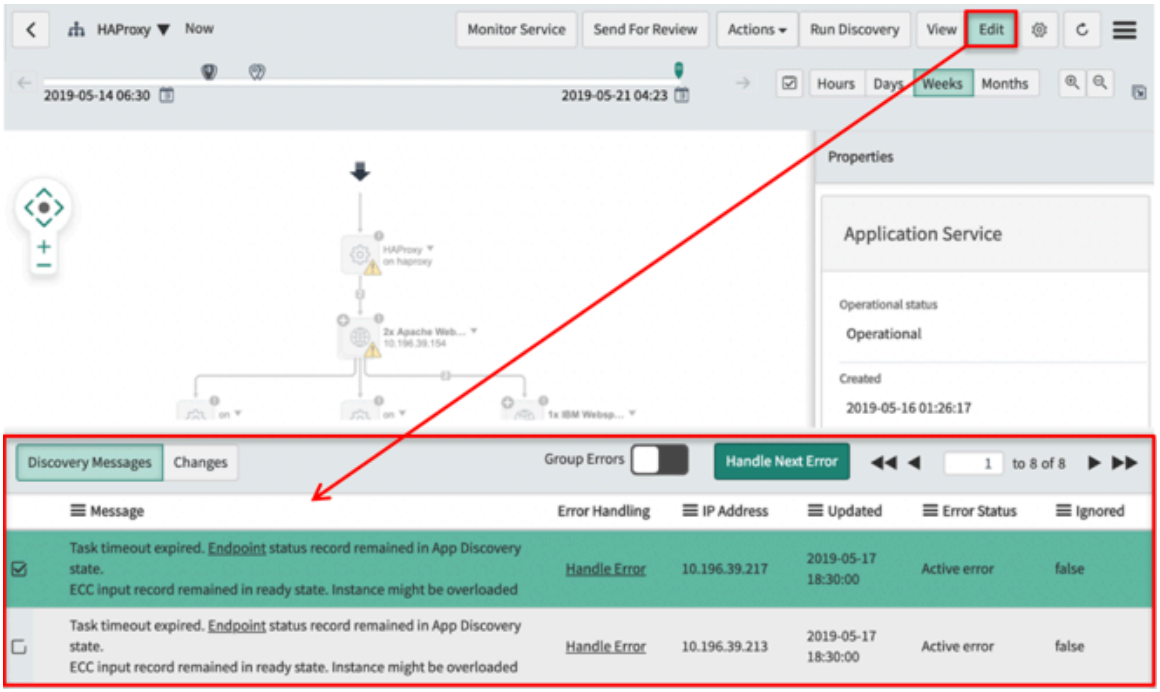

3. Ensure that the map is not in the Host View.
4. Click the arrow next to the undiscovered CI and select **Skip and Resume**.
5. Define attributes of the CI you want to manually place instead of the undiscovered CI.

Field	Description
CI Type	Select the CI type (CI class) for the CI you are adding. Every CI belongs to a CI type which contains a set of attributes configured for this kind of CI, for example, *cmdb_ci_appl* for applications.
CI Name	Select the CI from the list of CIs of the selected CI type. **Note:** To eliminate the possibility of delayed results when searching for a specific CI, make your search as specific as possible. A search with *<name> might take a long time and return a large data set.

Note:

If the CI type does not appear in the list, Service Mapping does not allow adding CIs of this type manually. For example, you cannot add an application cluster or an endpoint manually.

6. Configure attributes of the entry point from which Service Mapping continues discovery by selecting the entry point type and entering relevant fields.
   If the required entry point type does not appear in the Select Entry Point Type list, create it as described in Create entry point types for Service Mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Click **Submit**.
   The map shows the manually added CI instead of the undiscovered CI. Service Mapping starts discovery for CIs of the next tiers. When the discovery process finishes, the map shows all discovered CIs below the manually added CI.

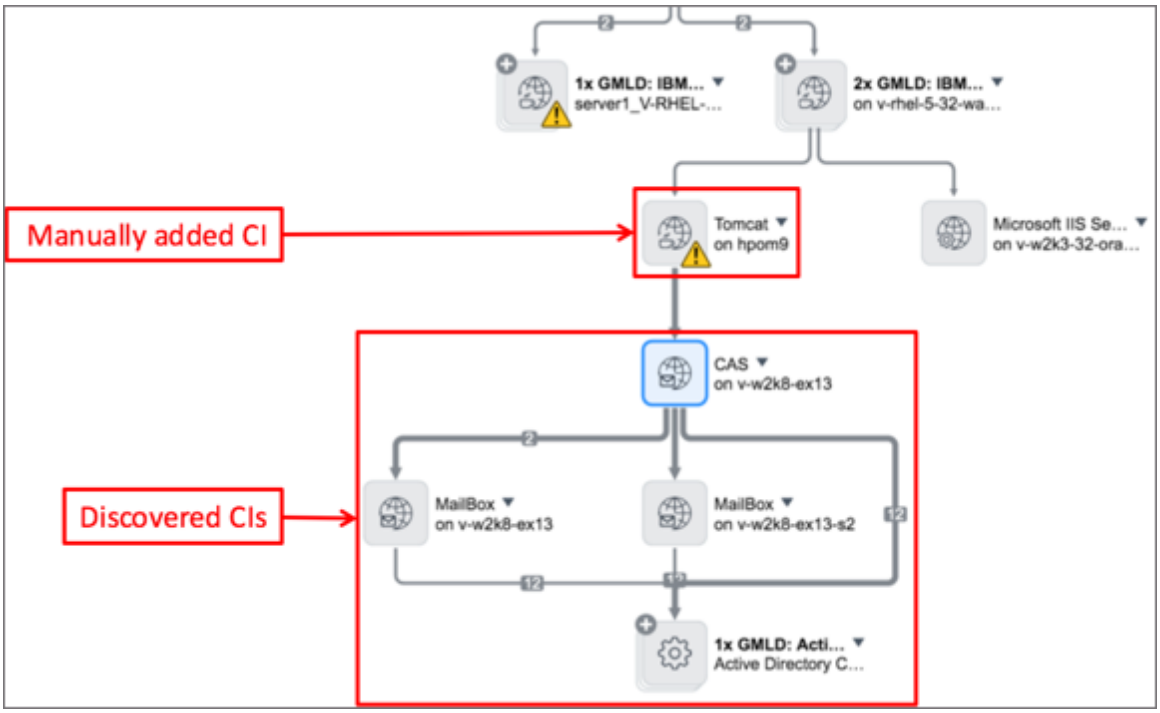

What to do next

To discover the problematic CI automatically:

1. Perform necessary troubleshooting to enable Service Mapping to discover the problematic CI. For example, you may need to customize the relevant pattern, provide missing credentials or fine-tune traffic-based discovery.
2. Navigate to the relevant service instance map.
3. Click **Run discovery**.
4. After the discovery process finishes, verify that Service Mapping discovered the CI. If Service Mapping discovered the CI, the map shows the actual CI without the Warning icon (⚠).

Related topics

- Discovery patterns used by ITOM Visibility
- Enable traffic-based discovery for CI types or specific CIs
- Prerequisites for performing top-down discovery using Service Mapping

Resolve pattern-related mapping errors

You can troubleshoot mapping errors caused by patterns.

Before you begin

Verify that the mapping error is caused by an inaccurate pattern by checking the discovery message. If the message says "Failed to recognize application", the error is pattern-related.

You need to be familiar with the Pattern Designer module of Service Mapping.

Role required: admin and service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

If there are configuration items (CIs) that Service Mapping could not map correctly, they appear on your service instance map with warning icons ⚠. Typically, mapping errors happen when Service Mapping fails to connect to a CI or fails to recognize it due to an inaccurate pattern.

You can identify problematic steps in your pattern and fix them without reviewing all steps and operations contained in the pattern. It allows you to troubleshoot pattern-related errors quickly and effectively.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Ensure that the map opens in Edit mode displaying discovery messages and errors.

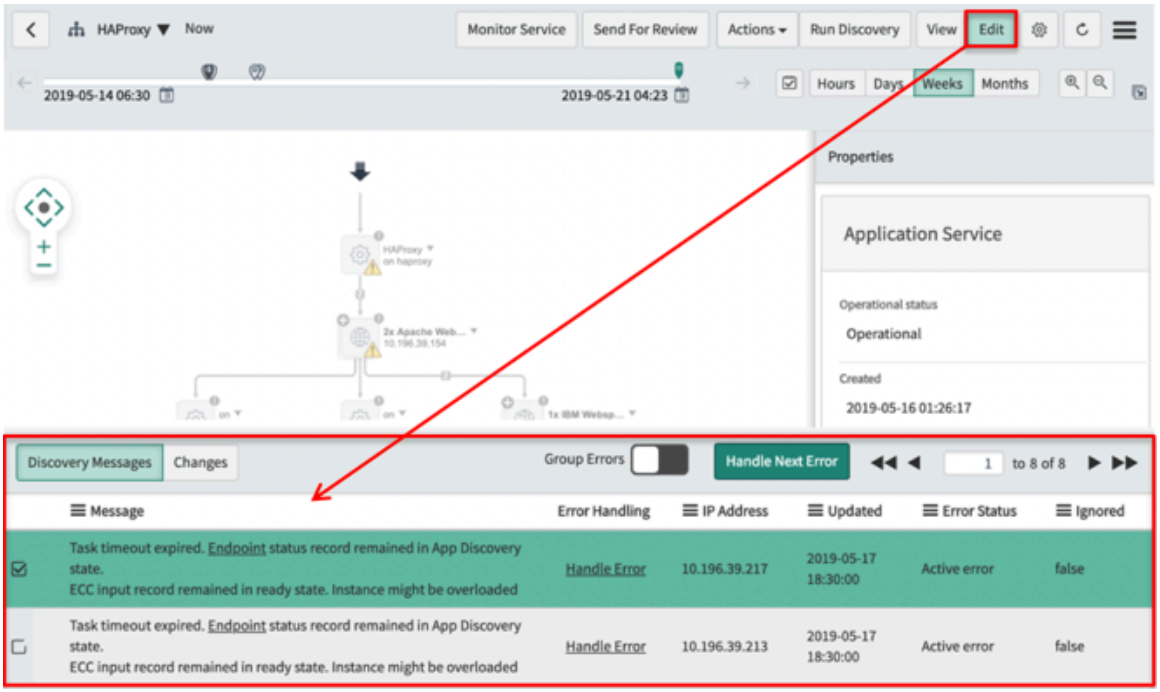

3. Right-click the CI with the warning icon (⚠) and select **Show discovery log**.
4. Expand the failed pattern which appears above the result marked red.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

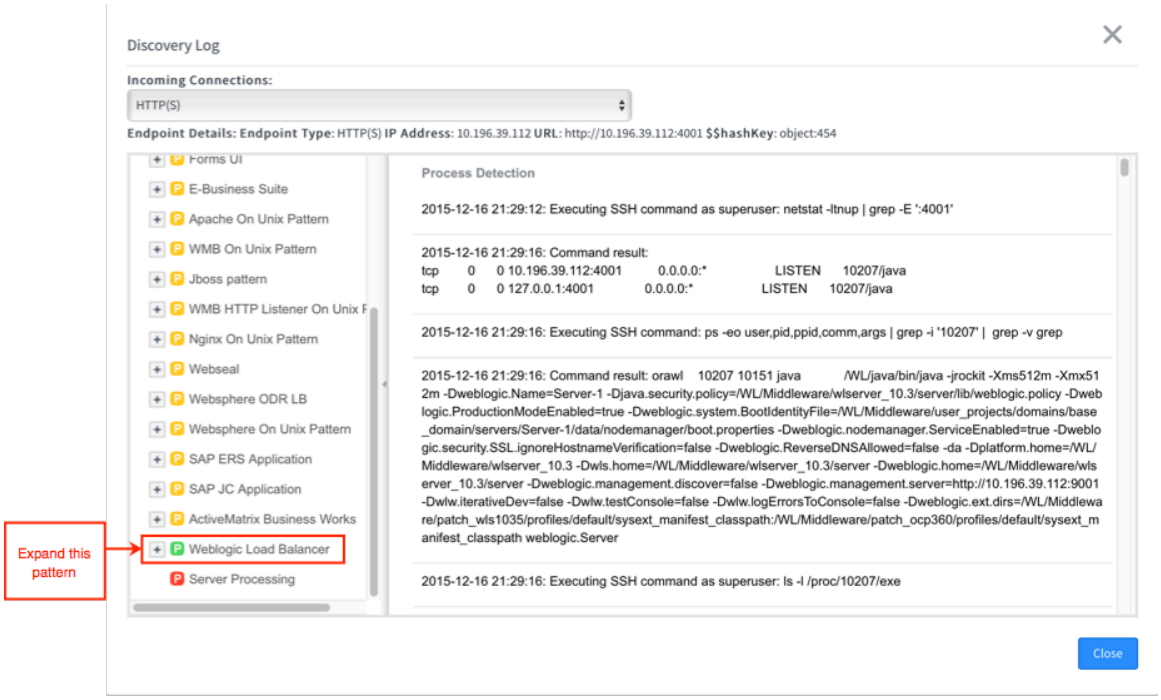

5. Expand the failed section and click the failed step.

6. Click **Debug**.
7. Click the first step and wait for Pattern Designer to run it.

   Repeat this action for other steps following the order until you reach the faulty step that causes the error.

   In the example below the port number is 8081 instead of 8080 which causes the problem.

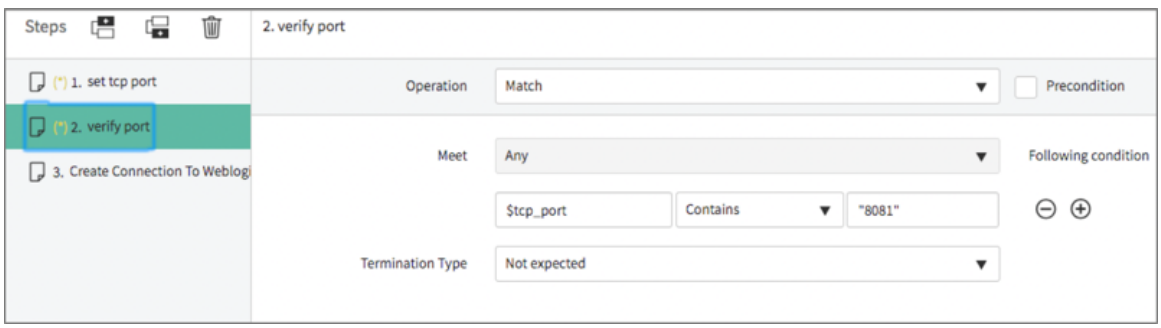

8. Click **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. Fix the pattern attribute that causes the problem.
10. Click **Test**.
11. Verify that Pattern Designer returns a success message.
12. Click **Save**.
13. Click **Activate**.
14. Return to the service instance map which contained the mapping error.
15. Right-click the CI for which the pattern was fixed and select **Resume discovery**.
16. Verify that the CI is discovered and mapped correctly.

Review and approval of application service maps

After the Service Mapping administrator maps IT services and fixes errors in them, the administrator and the owner collaborate to review and approve the service maps. The review and approval process is available only for discovered and manually created service instances.

You must have Service Mapping enabled to map or review service instances.

Users who perform review and approval have the roles of a Service Mapping administrator and an Service Mapping administrator is responsible for mapping, fixing, and maintaining application services. The service instance owner checks that the service instance maps are complete and all major components comprising it are correctly represented. If necessary, the owner leaves comments, referred to as reject messages, on service instance maps for the Service Mapping administrator to implement.

The review and approval process is performed in the following stages:

1. After the initial mapping of service instances, the administrator fixes errors and sends individual service instances for review. The system creates a service process task assigned to the service instance owner and sends an email notification about it.
2. The service instance owner checks that the service instance maps are complete and all major components comprising it are correctly represented. If necessary, the owner leaves comments, referred to as reject messages, on service instance maps for the Service Mapping administrator to implement. The service process task assigned to the owner closes. The system sends an email notification to the administrator that the owner posted comments.
3. The administrator receives the email notification, views owner requests and implements them.
4. The administrator resends the service instance maps for review. The system again creates a service process task assigned to the service instance owner and sends an email notification about it.
5. If the revised service instance maps are satisfactory, the service instance owner approves them. If not, the owner requests further fixes, which the administrator must address.

To see documentation for a review phase, click the relevant box in the diagram.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

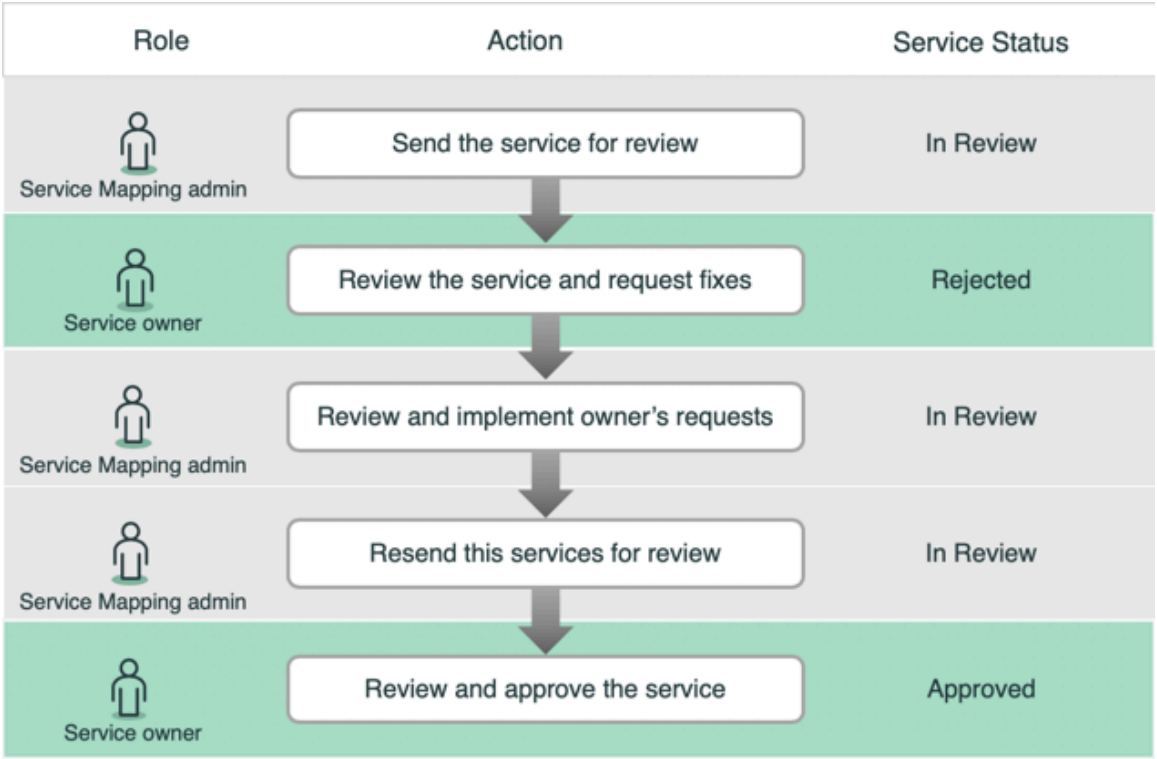

1. Send the service instance for review
2. Review the business service and request fixes
3. Review and implement owner's requests
4. Resend the application services for review
5. Review and approve the application services

The process of application service review may take some time as it requires making changes and repeatedly running the mapping process on the application service. Typically, it takes several iterations to arrive at the desired result.

Once the owner approves an application service, the application service status changes to Approved and it appears in the list of completed application services on the **Home** page.

Send application service maps for review

After you map an application service, send it to the application service owner for review to make sure that the map is accurate.

Before you begin

Perform initial error fixing as described in Fix application service errors in bulk and Fix errors in individual application service maps.

Role required: service_mapping_admin

About this task

While you can fix errors in bulk, you always send application services for review individually, one by one.

Sending application service maps for review is part of the review and approval process. Typically, you send each application service map for review twice: The first time for the initial owner review and the second time after you implemented owner feedback.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After the initial mapping of application services, you fix errors and send individual application services for review. The system creates a service process task assigned to the service instance owner and sends an email notification about it.

To see documentation for another review phase, click the relevant box in the diagram.

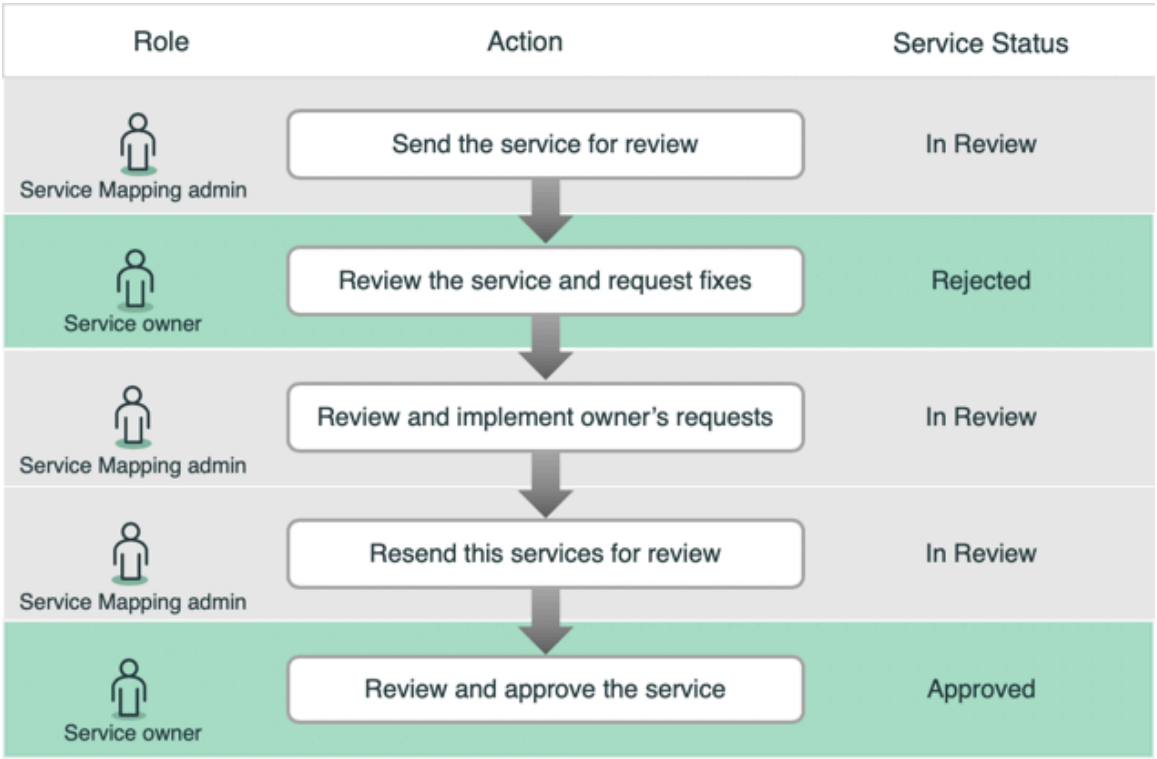

1. Review the service instance and request fixes
2. Review and implement owner's requests
3. Resend the service instance for review
4. Review and approve the service instance

Procedure

1. To send an service instance for review from the service instance form:
   a. Open the service instance map.
   b. Sort the list of planned application services by status and scroll to see services in **In progress** status, or use the filter to narrow the list.
   c. Click the required service instance.
   d. Click **Actions** > **Send For Review**.
2. Alternatively, send an service instance for review from its map:

   a. Navigate to **Service Mapping** > **Home**.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.

   b. Click the **Approve** tile.

   c. Click **Send For Review**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Click the required service instance.
The map for this service instance opens.

e. Click **Send For Review** in the upper corner of the window.

What to do next
When you receive an email notification that the owner sent comments for this service instance, review and implement the owner's requests.

Review application service maps

As the application service owner for the application service map, you receive an email notification that the application service map is assigned to you for review. Review mapping results for correctness and either provide your feedback or approve the application service map. The review and approval process is available only for discovered and manually created service instances.

Before you begin
Role required: sm_app_owner

About this task
Ideally, only the application service owner reviews the planned application service, however, users with the sm_app_owner or service_mapping_admin role can also perform this task.

Reviewing application service maps is part of the review and approval process. The process of application service review may take some time as it requires making changes and repeatedly running the mapping process on the application service. Typically, it takes several iterations to arrive at the desired result. Once you are satisfied with the discovery result, you approve the application service.

After you request fixes or approve the application service, the service process task assigned to you closes.

To see documentation for another review phase, click the relevant box in the diagram.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

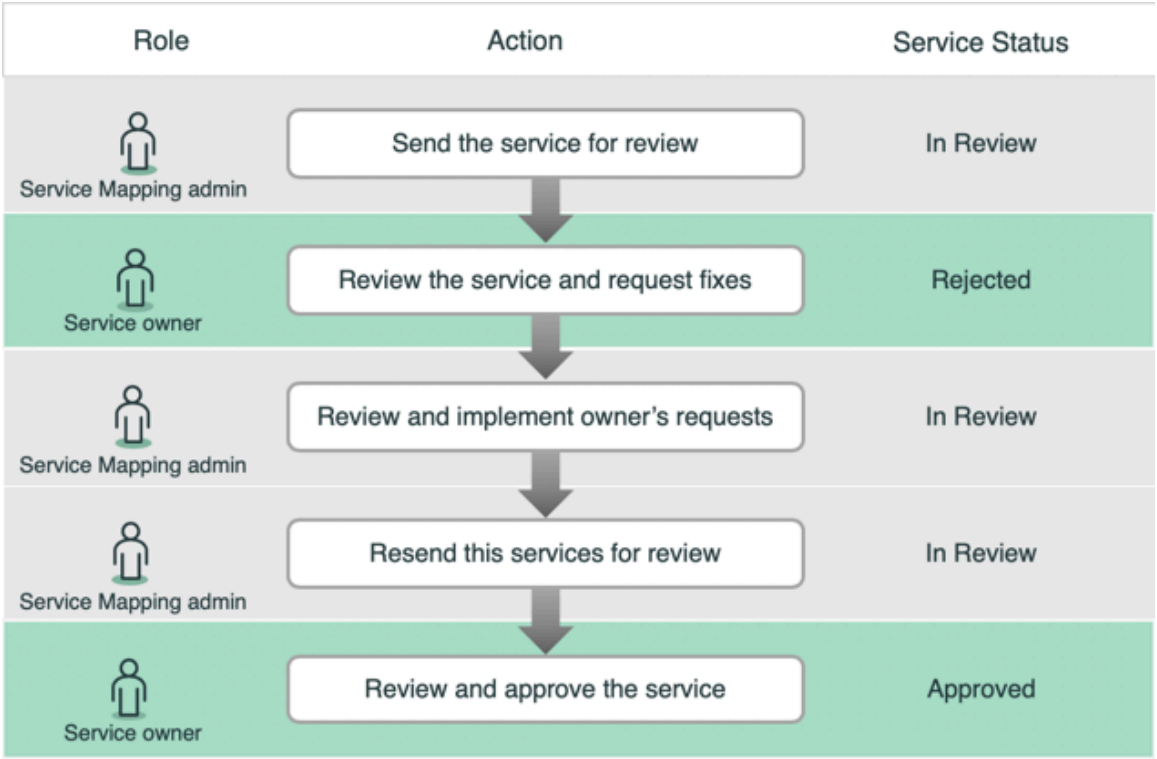

1. Send the application service for review
2. Review and implement owner's requests
3. Resend the application services for review

Procedure

1. Click the link to the application service map in the email notification.
   The map of the application service opens in the ServiceNow instance.
2. Alternatively, open to the application service to review from the list of tasks assigned to you.
   a. Navigate to **Service Mapping** > **Administration** > **My Tasks**.
   b. **Optional:** Sort the list of service process tasks as required.
   c. Click the required task.
   d. Click the application service form link.
   e. Click **View map**.
3. On the application service map, check that all essential CIs comprising the application service are discovered and mapped correctly.

Purpose	Action
Verify that there are no missing CI connections.	Pay attention to CIs with no connectors from it to other CIs. For example, in the following figure, the Web Server on win2K12 CI is missing connections.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Purpose	Action
	 If connections and CIs to which they lead are missing, the map does not reflect the re al state of the service instance and its opera tion. Inaccurate data can also be transferred to Event Management, causing imprecise monitoring.
Verify that there are no CIs that do not be long in the service.	Check all CIs comprising the application ser vice map to identify CIs not belonging to this application service. Typically, it is applica tions supporting internal services. For exam ple, Microsoft Internet Information Services (IIS) connecting to Active Directory, may be not part of the application service.
Check that the connections between CIs are correct.	View CI connection attributes in Service Map ping.
Check that clusters are reflected correctly.	Click the plus (+) icon next to a CI. There are two types of clusters: **Application cluster** In an application cluster, two or more applications or devices are configured to work together and serve the same purpose. Typical ly, the clusters provide high avail ability or load balancing. For ex ample, web servers working be hind a load balancer. This type of cluster appears as a stack of CIs with a label showing the number of CIs in this cluster with the multiplication sign. 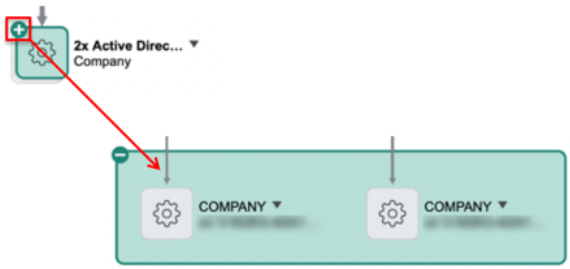

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Purpose	Action
	Application clusters can have a cluster within a cluster, but never more than two levels: 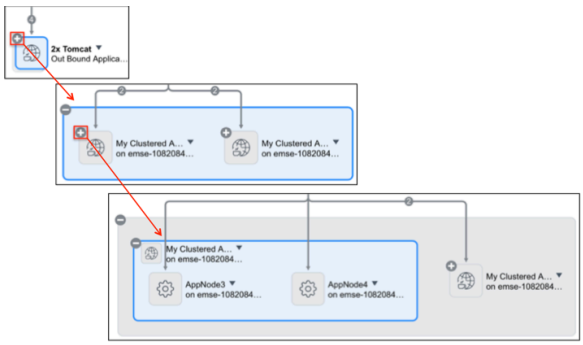 **OS cluster** In an OS cluster, a server hosts two identical operating systems. Typically, the cluster provides failover. This type of cluster appears as a CI with a plus sign and the num ber of CIs in this cluster. 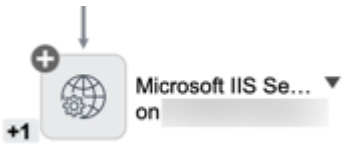
Check that inclusions are reflected correct ly.	Click the plus (+) icon next to a CI. In an inclusion, a server hosts applications that are treated as independent objects. For example, IIS Virtual Directory can run on a Windows Service as its host. 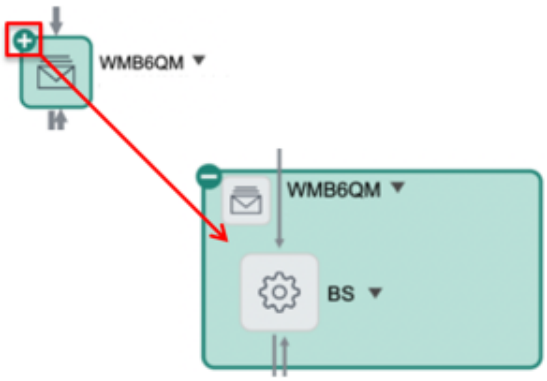

4. To notify the service mapping admin of a necessary change or fix, perform the following steps:

a. Click **Actions** above the map.

b. Select **Reject Service Map**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. In the Add Reject Message window, enter your comment, and click **Add**.
The new comment appears on the **Discovery Messages** tab under the map. The application service status changes to Rejected.

d. If necessary, add more comments.

5. If you are satisfied with the results, approve this application service by clicking **Actions**, and then **Approve service**.
The application service status changes to Approved and it appears in the list of completed application services on the **Home** page. The operational status changes to Operational. From this point, other applications can use data collected and organized by Service Mapping.

Related topics

Fix errors in individual application service maps

Fine-tune application services to implement owner requests

As an administrator, you tweak and fine-tune maps based on comments from the application service owner.

Before you begin

Role required: service_mapping_admin

About this task

You can edit discovered and manually created service instances.

Important: You cannot fine-tune or edit tag-based and dynamic services from the map.

Reviewing and implementing owner requests is part of the review and approval process.

To see documentation for another review phase, click the relevant box in the diagram.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

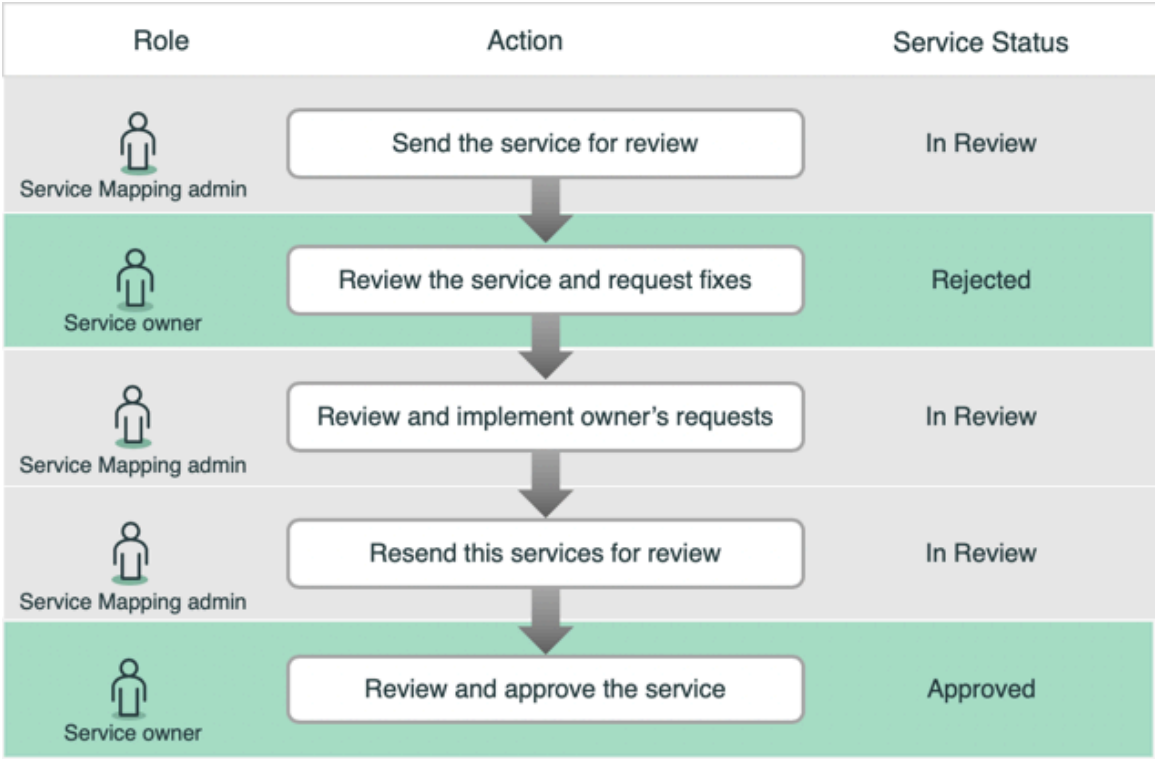

1. Send the service instance for review
2. Review the service instance and request fixes
3. Resend the application services for review
4. Review and approve the service instance

Procedure

1. Navigate to the relevant application service:

   a. Click **Approve** on the **Home** page.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.

   b. Click **Fix Rejected**.

   c. Right-click the required application service and select **Open Service Form**.
   The form for this application service opens.

2. Click **Reject Messages** on the left.
   The owner comments are displayed under the **Service Discovery Messages tab**.
3. Review the reject messages to see comments that the application service owner left.
4. Click **View Map**.
5. Implement changes requested by the application service owner.
6. Click the discovery message describing the service defect that you have fixed and click **Mark as Resolved** under **Action On Selected** on the right.
7. If necessary, resend the application service for review.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Manually add CIs to an application service

Add configuration items to manually created application services or to services discovered by Service Mapping.

Before you begin

- Verify that the CI type for the configuration item (CI) that you are planning to add, exists. If necessary, create the CI type as described in Create CI types for Service Mapping and Discovery.
- Add CIs to the CMDB for the device or application that you want to add, if necessary. See Populate the CMDB for more information.

Role required: app_service_admin or service_mapping_admin

About this task

Adding a CI to an application service requires creating a relationship between the new CI and a CI in the application service. You can add CIs to an application service that was created manually, by either:

- Adding a method to populate the application service.

  Navigate to **CSDM** > **Manage Technology Management Services** > **Application Service**. Select an application service and then use the **Populate the Application Service** tab to choose a method to populate the application service. For more details, see Create an application service.

- Using the application service service map as described in the steps below.

  The default relationship type of the added connection in this case is *Depends on::Used by*. You can modify this default relationship type by changing the value of the *sa.it_service.manual_ci_rel_type* property. See Components installed with application services for more information.

Important: You cannot fine-tune or edit tag-based and dynamic services from the map.

Manually adding a new CI to an existing CI in a service instance prompts the CMDB to update the information about both CIs, including their relationship type. If other application services use the same applicative flow, the CMDB recognizes the new CI and automatically adds it to those services as well.

For example, you manually add an IBM WebSphere Message Broker to an IBM WebSphere HTTP Listener in the Bank Customer Portal service. The system also adds this IBM WebSphere Message Broker to the same HTTP Listener in the Bank Internal Portal service, because it uses this HTTP Listener. Similarly, when you remove a CI you added manually, the system removes it from all application services. This includes services where you manually added it and services where it was automatically added by analogy.

You can manually connect a CI only to actual CIs existing in the CMDB, not to a visualization of other items on the map such as clusters or boundaries. Also, you cannot add CIs of these CI types to an application service:

- NAT [cmdb_ci_translation_rule]
- Endpoint [cmdb_ci_endpoint]
- Qualifier [cmdb_ci_qualifier]
- Application cluster [cmdb_ci_application_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

There are system records in the Manual CI Exclusions/Inclusions [svc_manual_ci_exclusions_inclusions] table. See Components installed with application services for more information.

Note: The Manual CI Exclusions/Inclusions [svc_manual_ci_exclusions_inclusions] table doesn't include CIs added using traversal rules.

In environments with domain separation, only CIs belonging to the same domain as the service instance are added to the service instance. If there is a domain hierarchy, CIs must belong to the same child domain as the service instance.

If working with an service instance discovered by Service Mapping, manually add a CI:

- To indicate that an service instance contains a device or application, which Service Mapping cannot discover. For example, add an A/C unit to the Production Floor service.
- To add a temporary placeholder for a CI, which Service Mapping did not discover. In this case you are planning to perform necessary troubleshooting to ensure that Service Mapping discovers this CI in the future. For example, add an IBM WebSphere Message Broker to the Bank Customer Portal service.
- To create an service instance that combines entry points and CIs automatically discovered by Service Mapping with entry points and CIs from the CMDB. After you manually add an entry point, you can update the service instance with CIs from the CMDB based on the relationships defined there.

For additional information related to Service Mapping, see Pattern customization and Enable traffic-based discovery for CI types or specific CIs.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. If needed, click **Edit** to ensure that the map is in Edit mode.

   If Service Mapping is deployed, then in Edit mode, the Discovery Messages section appears below the map.
3. To connect a CI to another CI on the map, right-click the CI to which you want to connect the new CI, and then select **Add a CI**.
4. In the Add A CI dialog box, specify the CI to add:

Field	Description
CI Type	Select the CI type (CI class) for the CI you are adding. Every CI belongs to a CI type which contains a set of attributes configured for this kind of CI, for example, *cmdb_ci_appl* for applications.
CI Name	Select the CI from the list of CIs of the selected CI type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: To eliminate the possibility of delayed results when searching for a specific CI, make your search as specific as possible. A search with *<name> might take a long time and return a large data set.

The CI type list includes only allowed CI types. For example, you cannot add an application cluster.

5. Click **Submit**.
   The manually added CI appears on the map.

   Note: When you manually add a CI, which is an application, as a child to a service that already includes its parent application CI, the newly added child application CI is hidden inside the inclusion. Click the plus (+) symbol next to the parent application CI to see the child application CI.

6. **Optional:** If Service Mapping is activated, add a discoverable outgoing connection for the manually added CI:

   a. Right-click the manually added CI.

   b. Select **Manually add a connection**.

      Note: If you do not see the **Manually add a connection** option in the right-click menu, check that you are logged in with the user that belongs to the same domain as the application service.

   c. Configure attributes for the entry point as described in Entry points attributes.

   d. Click **Submit**.
      Discovery and Service Mapping attempt to discover this CI. If successful, the CI appears on the map. Otherwise, a warning icon (⚠) appears.

7. **Optional:** If Service Mapping is activated and you want Service Mapping to automatically discover a CI, which you previously added manually:

   a. Customize the relevant pattern or fine-tune traffic-based discovery to enable Service Mapping to discover the CI.

   b. Navigate to the relevant service instance map.

   c. Click **Run discovery**.

   d. After the discovery process finishes, verify that Service Mapping discovered the CI by checking the connector leading to the CI.
      If Service Mapping discovered the CI, then two connectors, a manual and automatically discovered, appear for the CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. Right-click the CI you added manually.
In the example, it is IBM WebSphere Message Broker.

f. Select **Remove manually added CI**.
The map shows the CI with only one connector leading to it. If this CI had any manually added connections, they are removed together with the manually added CI.

Related topics

Link application services

addCI()

Remove CIs not belonging to application services

Remove CIs erroneously mapped as part of an application service by Service Mapping. Unnecessary CIs included in the map can generate irrelevant alerts in Event Management. For example, when creating an application service for a web portal, Service Mapping might automatically discover a connection to unaffiliated external services, such as PayPal.

Before you begin

You remove CIs only from application services discovered by Service Mapping.

Important: You cannot fine-tune or edit tag-based and dynamic services from the map.

Role required: service_mapping_admin

About this task

Removing CIs from the map also removes all their connections. For example, if you remove a CI of the second tier on the map, all CIs of the third tier connected to it through outgoing connections, get removed too.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Depending on why Service Mapping mapped an unwanted CI, you can remove it by the following methods:

- Add mapping boundaries to remove the unnecessary CIs. A boundary makes Service Mapping stop discovery from this point and not follow outgoing connections. Use this method to remove CIs that Service Mapping discovered using patterns or using traffic-based discovery.
- Tweak traffic-based discovery. Service Mapping can discover and map configuration items (CIs) following their traffic-based connections. This method is referred to as traffic-based mapping and complements pattern-based mapping. Service Mapping generates connection suggestions based on traffic-based connections and on Predictive Intelligence analysis of application fingerprints. If the connection suggestions feature is enabled, Service Mapping doesn't automatically add any traffic-based connections to application services, providing suggestions. Review the suggestions and choose which CIs to add or to remove.

  If the connection suggestions feature is disabled, Service Mapping automatically adds traffic-based connections and CIs to application services. Remove connections that lead to irrelevant CIs to declutter application services.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Remove CIs that do not belong to this service instance by completing the following actions.

Action	Steps
Add mapping boundaries	**a.** Right-click the connector leading to the CI you want to remove. **b.** Select **Mark boundary**. The CI is marked as boundary on the map ( ). **Note:** Ensure that you are looking at the map in Edit view.
(If the connection suggestions feature is disabled) Remove CIs discovered using traffic-based connections	**a.** Right-click the CI on the map. **Note:** You can see traffic-based connections for a CI, even if you disabled the traffic-based discovery for an service instance the CI is part of. **b.** Select **Show traffic based connections**. The Traffic Based Connections List opens. See Traffic-based connections list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Steps
	c. In the **Traffic Based Connections List**, identify the connection used to discover the unwanted CI. **d.** Select `Remove` from the **User decision** field for this connection. **e.** Click **Close**.
(If the connection suggestions feature is enabled) Remove CIs discovered using traffic-based connections	Perform the procedure described in Manage CI connections for a single service using suggestions.

Related topics

Enable traffic-based discovery for CI types or specific CIs

Transfer a map segment into another application service

You can remove a branch of a service and place it into another application service, either new or existing. Transfer map segments to split large services or when you want to organize services differently from the initial mapping result.

Before you begin

You can edit discovered and manually created service instances.

Important: You cannot fine-tune or edit tag-based and dynamic services from the map.

Role required: service_mapping_admin

About this task

When you transfer a map segment into another application service, you create a link between the original application service and the application service into which you planted the map segment. The connection to the top CI in the segment becomes an entry point when you transfer it into another, existing or new, application service.

The original service map shows the link to the other application service instead of the map segment. The service that contains the reference, becomes a dependent service. The service that appears as a reference is a contained service.

When you transfer a map segment into another application service, the information about the connection to the top CI in the transferred segment, is updated in the CMDB. If other application services use the same ingoing connection (endpoint), the CMDB recognizes it and replaces map segments coming out of the same connection with connected service CIs by analogy.

You can't insert links to multiple application services from the same connection.

If you use Event Management for service monitoring in your organization, dependent application services include alerts generated for their contained services.

Procedure

1. Open the application service map containing the segment that you want to transfer:
   a. Navigate to **Service-Mapping** > **Application Services**.
   b. Click **View map** next to the relevant service instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. **Optional:** Click **Edit** to be in Edit mode.

d. If the Host view is on, switch it off by clicking **More Options** and clicking the **Display in Host View** toggle.

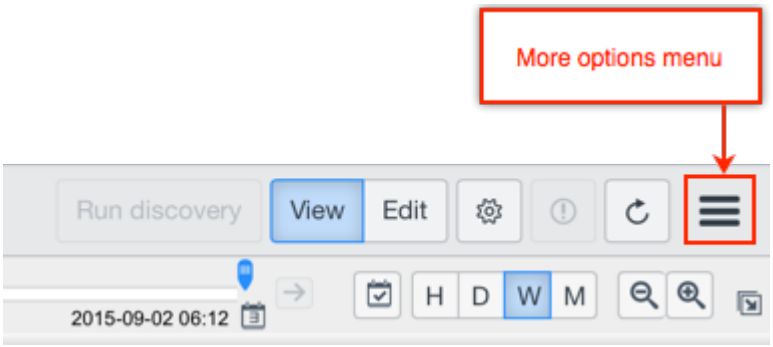

2. Identify the segment of the map that you want to transfer.
3. Right-click the connection leading to the CI at the top of this segment.

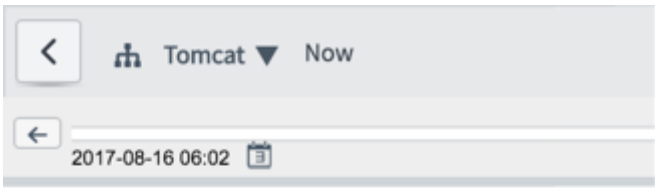

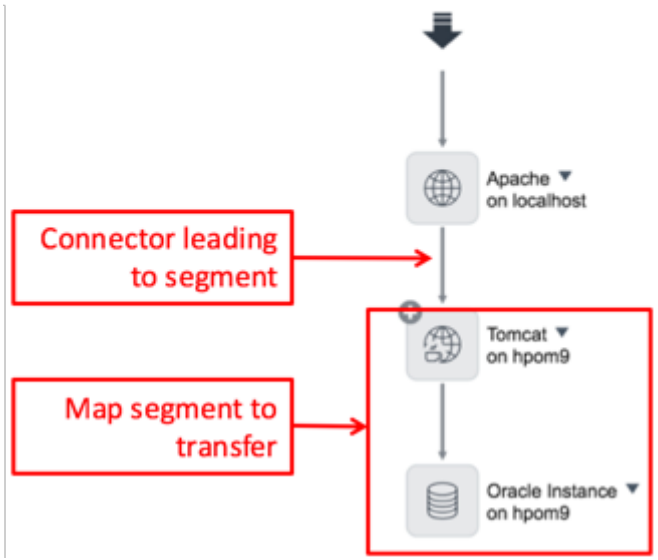

4. To transfer the segment into an existing application service map and connect to it:

a. Select **Connect to existing service**.

b. In the Add To Existing Discovered Service window, select the application service to which you want to transfer the segment.

c. Click **Choose**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The map displays the contained application service icon instead of the transferred segment. The application service map to which you transferred the segment, shows it under the newly added entry point. For example, the Tomcat application service map includes the Trade application service as its contained service. The Tomcat application service becomes dependent on the Trade service. The Trade application service map shows the segment transferred from the Tomcat application service.

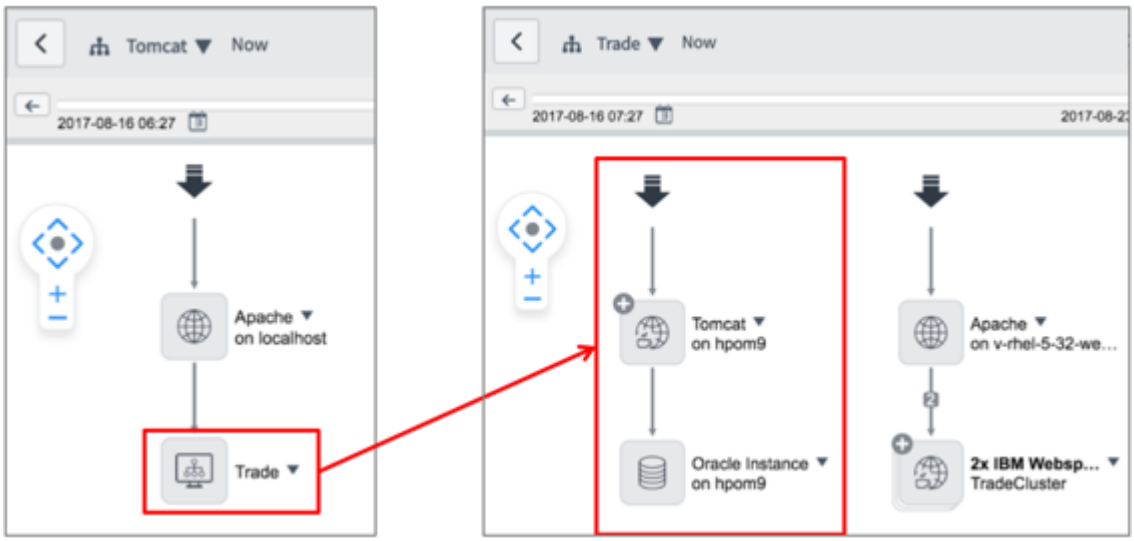

5. To transfer the segment into a new service instance map and connect to it:

a. Select **Create new service from here**.

b. In the Create New Discovered Service window, define attributes for the service instance that you want to create for this segment:

Field	Description
Name	Enter the service instance name. This name must be unique. Use self-explanatory names such as `mailing service` or `printing service`.
Group	(Optional) Restrict access to a service instance by adding it to an service instance group. Users must then have the service group role to access the service instance.
Criticality	(Optional) Select the option that reflects how important this service instance is to your organization operation. For more information about service instance criticality, see Define criticality for application services.

c. Click **Create**.
The map displays the contained application service icon instead of the transferred segment. The new application service appears in the list of application services and contains the transferred segment of the map. For example, the Tomcat application service includes the TP rooms service as its contained service. The Tomcat application service becomes a service

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

dependent on the TP rooms service. The TP rooms application service map contains the segment transferred from the Tomcat application service.

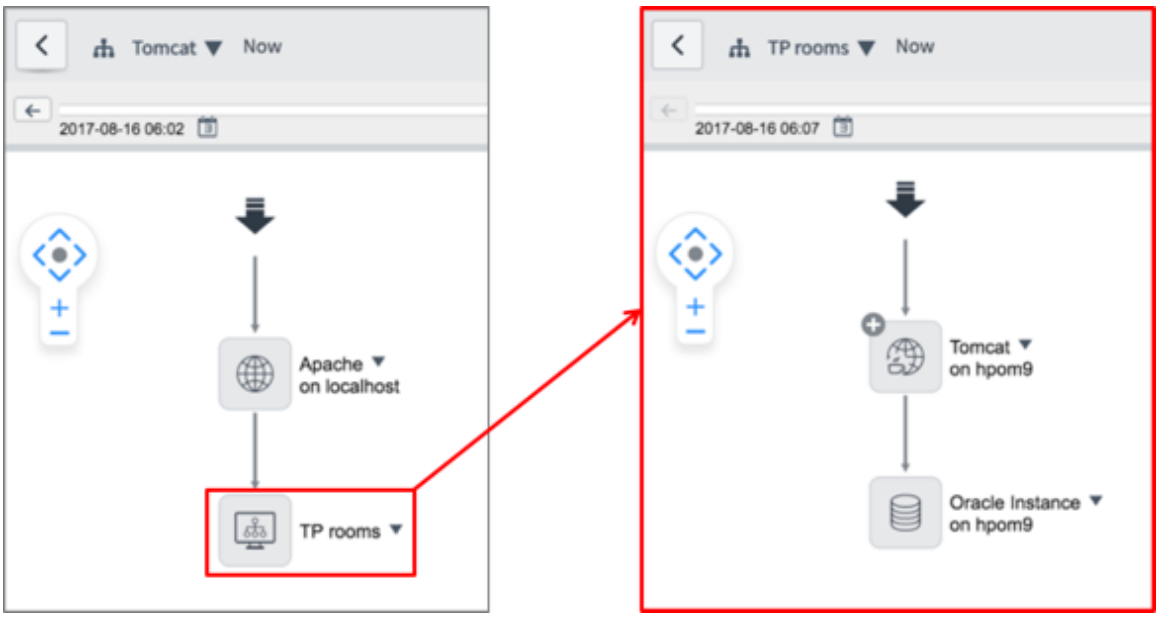

What to do next

To revert this operation:

1. On the service instance map containing the connected service CI, click the connection leading to it.
2. In the Properties pane, note the URL attribute of the connection.

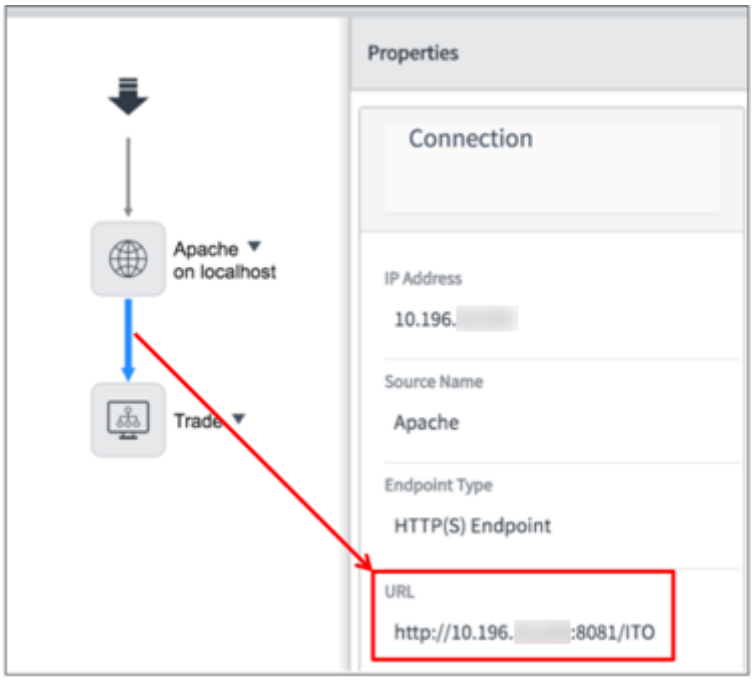

3. Double-click the connected service CI.

   The service instance map for the CI containing the transferred segment opens.

4. Click **Service Map Form**.
5. Under Entry Points on the left, locate the entry point with the same URL you took note of.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Click **Delete** on the entry point tab, and click **Remove** to confirm.

   This entry point disappears from this service instance form.

   The transferred segment disappears from the map of this service instance.

7. Navigate to the map that contained the connected service CI.

   This map displays a boundary ( ) instead of the connected service CI.

8. Click the boundary CI and click **Unmark boundary**.

   The map is refreshed and displays the segment of the service instance.

Related topics

View dependent application services in classic Service Mapping

View contained application services in classic Service Mapping

Control user access to application services

Link application services

You can manually link two application services by adding a reference to one application service into another application service. The service that contains the reference, becomes a dependent service. The service that you include as a reference is a contained service. You can link application services to create dependencies for impact monitoring in Event Management.

Before you begin

You can edit discovered and manually created service instances.

Important: You cannot fine-tune or edit tag-based and dynamic services from the map.

Ensure that you know the name and the service type of the application service, to which you want to add a reference.

Role required: app_service_admin or service_mapping_admin

About this task

To create a link, add a reference to the relevant application service as an outgoing connection of the relevant CI inside another application service. For example, you can add the UK Portal application service as a link to the Online Store application service. In this case, the Online Store service becomes dependent on the UK Portal service that it contains. The Online Store service reflects discovery errors for its contained service in the Edit map mode, as well as alerts in Event Management.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example of linked application services

When you link an application service to another application service, the information about the CI, to which you linked the service, is updated in the CMDB. The CMDB recognizes other application services that use the same applicative flow, and adds the contained application service to these application services by analogy. The same logic applies when you remove a contained application service: The system removes it from all application services where you either manually linked this service or the system linked this service by analogy.

When using Service Mapping, you may want to link application services to create:

- A dependency between two application services.
- A placeholder for a map branch that Service Mapping failed to discover. If you create or customize a pattern to discover the configuration item (CI) serving as an entry point for the contained service instance, Service Mapping can discover this contained service.
- An indication that an service instance contains a branch, which Service Mapping cannot discover.

You can add an service instance as a contained service to as many service instance as necessary.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Service Mapping** > **Application Services**.
2. Click **View map** next to the relevant application service.
3. If needed, click **Edit** to ensure that the map is in Edit mode.
4. Right-click the CI to which you want to link an application service as a reference.
5. Select **Add A CI**.
6. In the Add a CI dialog box, select the application service you want to add as a contained service:

Field	Description
CI Type	Select the relevant service type from this list: ○ `Tag-Based Application Service` ○ `Mapped Application Service` for discovered or manually created application services ○ `Calculated Application Service` for dynamic services ○ `Dynamic CI Group`
CI Name	Select the name of the application service that you want to link as a contained service.

7. Click **Submit**.
   The icon for the contained service appears on the map.

Related topics

Application services

View dependent application services in classic Service Mapping

View contained application services in classic Service Mapping

Discover CIs identified as generic applications

Service Mapping identifies application configuration items (CIs) that it failed to properly discover as generic applications. Correctly identify generic applications by creating a simplified discovery pattern from an application service map, rather than creating a fully functional pattern from scratch using the Pattern Designer.

Before you begin

You can edit discovered and manually created service instances.

Important: You cannot fine-tune or edit tag-based and dynamic services from the map.

Make sure that you understand what discovery process classification is.

Request information about the CI discovered as generic application from the application service owner.

Role required: pd_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

ServiceNow applications refer to devices and applications that comprise a service instance as configuration items (CIs).

As part of creating a simplified pattern for complete and proper application CI discovery, you create these necessary articles:

CI type

All applications and hosts in your organization must have a corresponding configuration item (CI) type, which is necessary for discovering and processing applications and hosts correctly. A CI type (or class) contains several important definitions that apply to all CIs belonging to it. Preconfigured CI types form a hierarchy where child CI types derive attributes from their parents.

Define the CI type name and label for Service Mapping to create a child CI type for the application [cmdb_ci_appl] CI type. For more information about CI types, see Create CI types for Service Mapping and Discovery.

Process classification

Service Mapping relies on process information collected and classified by Discovery to run relevant patterns to discover an application CI. During the horizontal discovery process, Discovery checks if the attributes of the processes running by the application match classification criteria you define. When a match is found, Discovery uses the existing Application Rule identifier to check if the application CI is new or if it already exists in the CMDB. Then Discovery either updates an existing record or creates a new record for this application CI in the CMDB.

Prior to starting the top-down discovery of an application CI, Service Mapping checks its process classification in the CMDB. It then runs only those patterns that are relevant for this application CI.

Create process classification by defining the classification criteria based on preconfigured process attributes. The new process classification is added to the Process Classification [discovery_classy_proc] table. To learn about process classification, refer to Discovery classifiers.

Pattern

A pattern is a sequence of commands designed to detect attributes of a CI and its outbound connections. When you create a simplified pattern for a generic application, Service Mapping assigns the new simplified pattern to the CI type you define as part of this process. The default name for the pattern starts with "ga_".

Association of an icon with this CI type

Each CI type is associated with an icon. When a CI is discovered and classified as belonging to a CI type, the map displays the associated icon to visualize this CI.

You cannot create simplified patterns from generic applications that are based on multiple processes. In this case, you have to either create a fully functional pattern using the Pattern Designer or create a new CI type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. **Optional:** Click **Edit** to be in Edit mode.

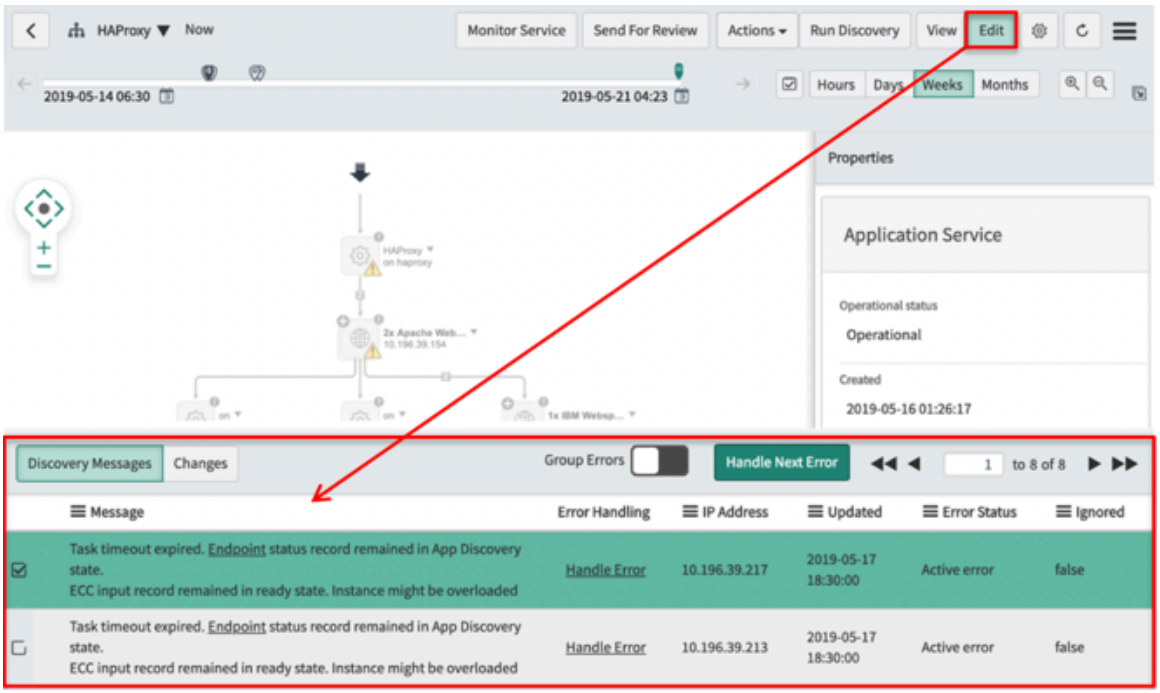

3. On the map, right-click the relevant generic application CI.

   The *CI Type* attribute in the **Application** pane is Generic Application.

   In the base system, the map displays generic application CIs with the gear icon.

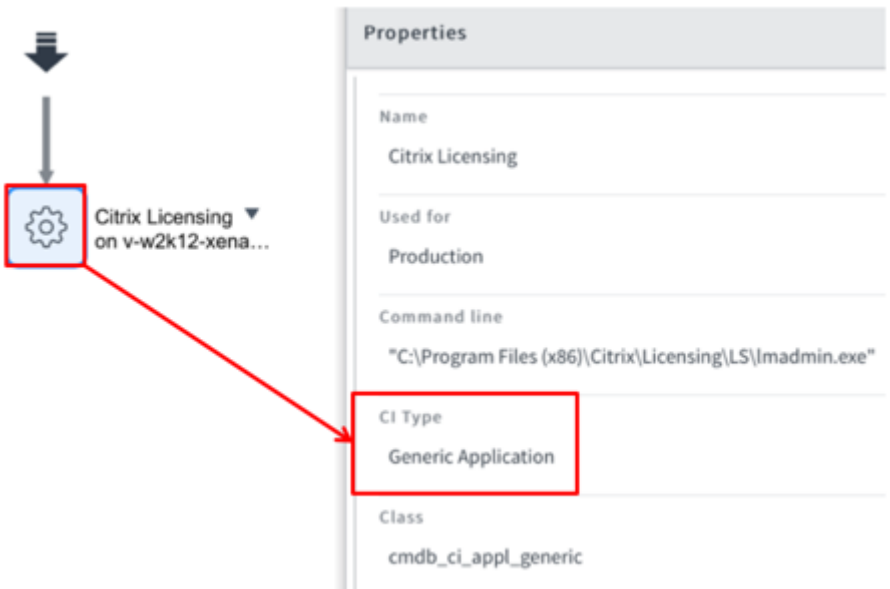

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Select **Create pattern for this generic application**.
5. Define the CI type for the application.

Field	Description
Label	CI type name. For example, Citrix Licensing Server.
Name	The table name containing data for this CI type. For example, cmdb_ci_app_license_server. Use a name that is consistent with existing CI type names. Typically, application CI type names start from "cmdb_ci_app".
Icon	Select the name of the existing icon to represent this application on maps.

6. Create the process classification for this CI by defining classification criteria for identifying the process that the generic application is running.
   In the base system, there are two conditions based on the running process name and the running process command. The values are filled in by the latest process discovery.

   a. Select the CI attribute from the list.
   The attributes in the list are from the Running Process [cmdb_running_process] table.

   b. Select an operator from the list.
   If you select `Is Empty`, the second field is rendered irrelevant and disappears.

   c. Specify the value to match.
   The value field displays a dynamically generated suggestion from the latest process discovery.

   d. To add a dependent condition, click **OR** or **AND**.

   Note: If the value in the first condition is not unique like Java.exe, use additional conditions to narrow down the results.

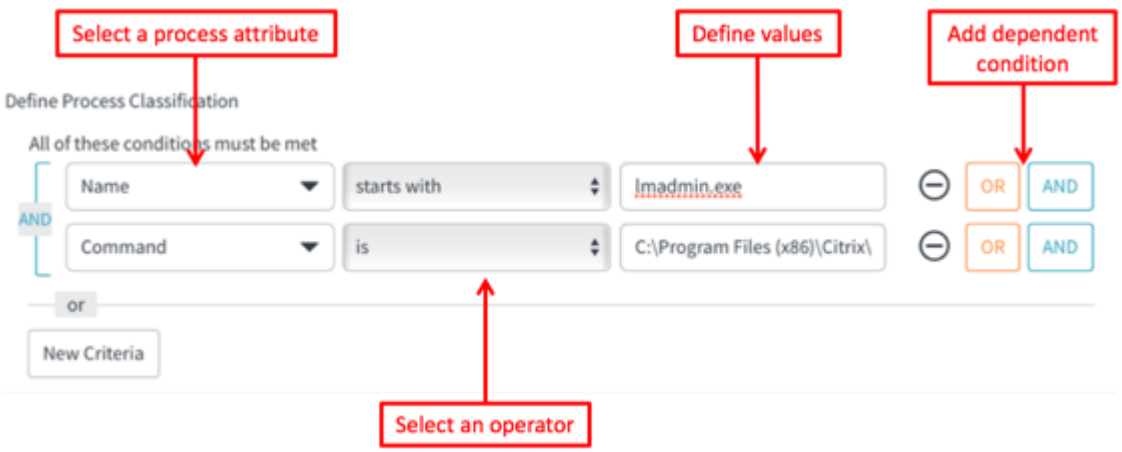

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, to discover a Citrix Licensing server, use the following classification criteria: the process name `Imadmin.exe` with the command `C:\Program Files (x86)\Citrix\Licensing\LS\lmadmin.exe`.

e. To add a top-level condition, click **New Criteria**.

7. Click **Create**.
   Once Service Mapping finished creating the CI type, the pattern, and the process classification, it resumes discovery of the CI initially identified as a generic application.
8. Verify that Service Mapping discovered the CI as you expected.

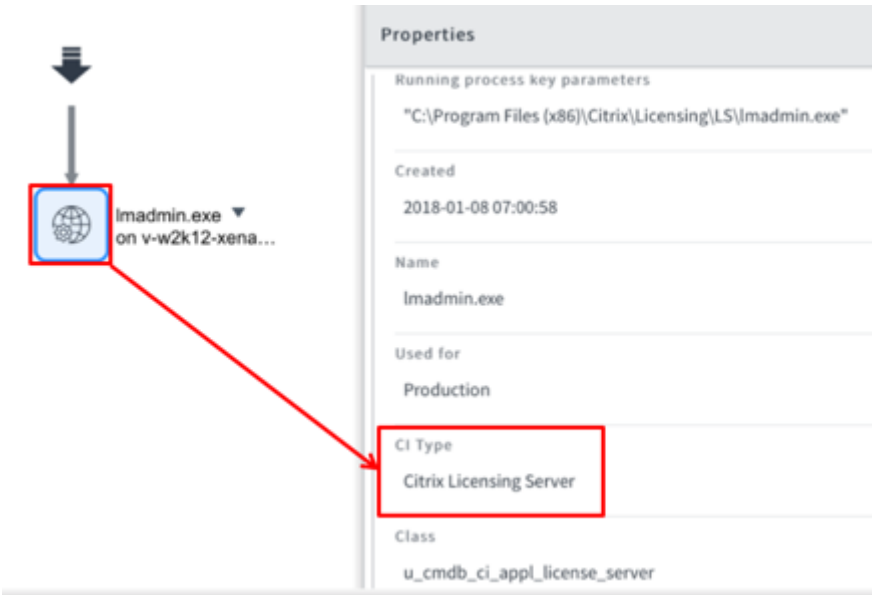

Notice that Service Mapping uses the name of the process for the CI label on the map.

What to do next

If Service Mapping still has not discovered the CI as you expected, fine-tune the match and classification criteria for the process classifier and configure advanced classification parameters.

If you are familiar with modifying patterns in the Pattern Designer, you can also enrich the simplified pattern by defining operations that discover more attributes for the CI.

You may configure Discovery to use a different identifier. For more information, see Discovery identifiers.

Related topics

Pattern-based discovery in Service Mapping

Application service completion

After an application service is reviewed and approved, you can define attributes that enhance its discovery, reflect its importance, and control access to it.

There are several configurations you must perform to complete definition of an application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Mapping workflow

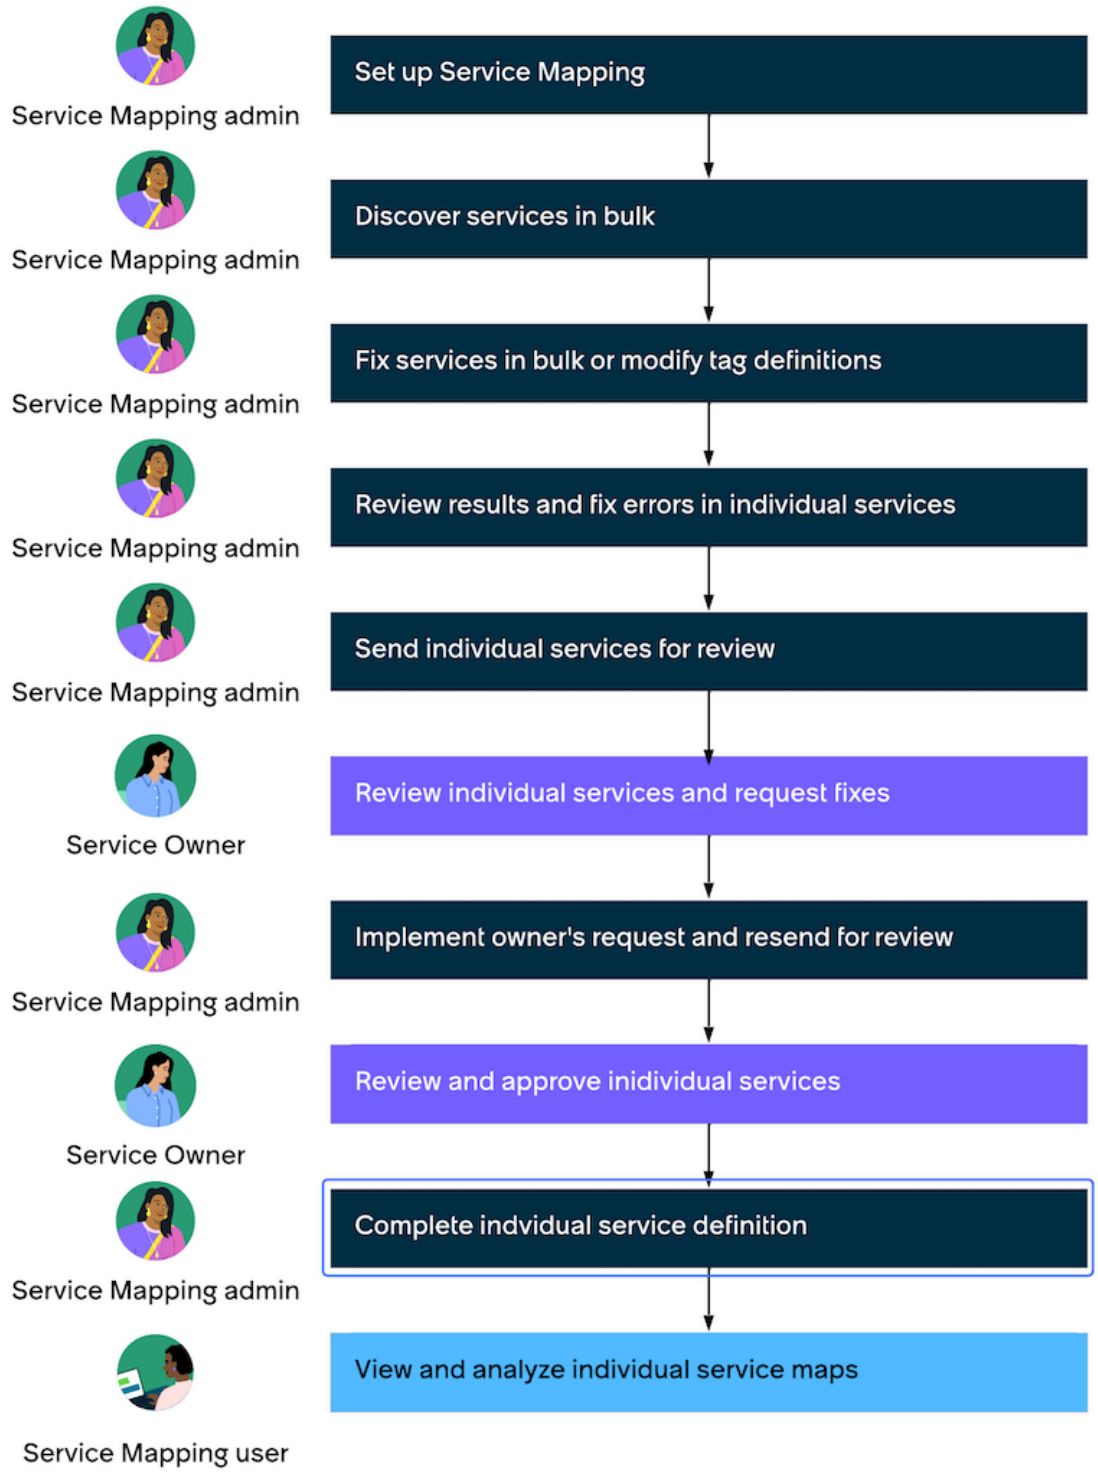

Define criticality for application services

For each service instance, select a criticality level that reflects how important it is to your organization operations.

Before you begin

Determine how important the application service is to your organization.

Role required: service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Typically, you might assign high criticality to service instances, which impact your customer base or sales revenue, while giving low criticality to internal service instances causing minor inconvenience to organization employees. For example, you could set high criticality for a service instance that supports sales functionality using the organization web site. A service instance that provides internal printing for the organization employees could receive low criticality.

You can also use a criticality value to define recovery strategies.

Criticality levels displayed in Operator Workspace page

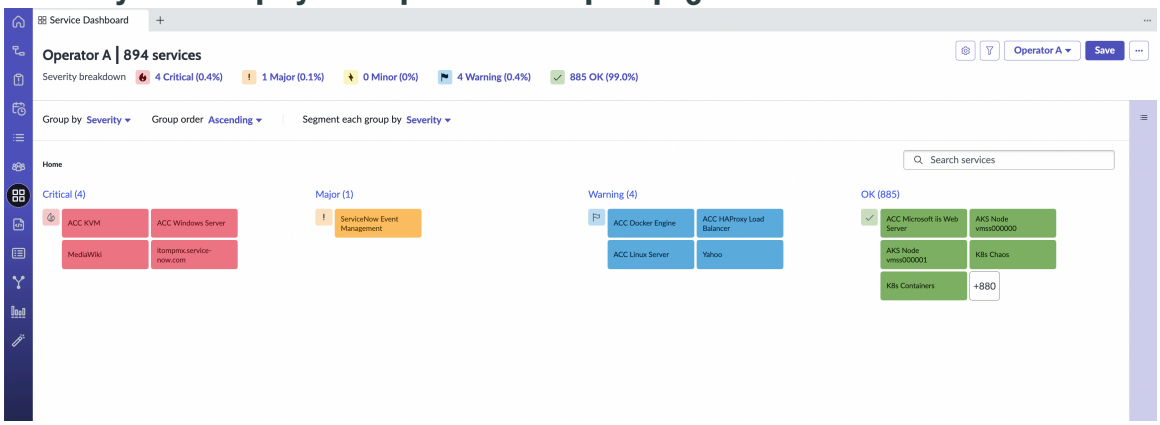

Procedure

1. In the classic environment, perform the following steps:

   a. Navigate to **All** > **Service Mapping** > **Home**.
   The Home page displays only information on service instances that Service Mapping can discover or already discovered. The Home page does not display information on service instances that are created manually or using the API.

   b. Click **View Your Services** on the **Completed** tile.

   c. Select the relevant application service.

   d. Click **Additional Info**.

   e. Select the appropriate option from the **Business criticality** list.

   f. Click **Update**.

2. Alternatively, perform the following steps in the Agent Workspace UI:

   a. Navigate to **Agent Workspace** > **Agent Workspace Home**.

   b. Click **Lists** in the left pane.

   c. Navigate to **Services** > **All Services**.
   The list of services displays services in Operational status.

   d. Select the relevant application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. Click the **Open Preview** icon.

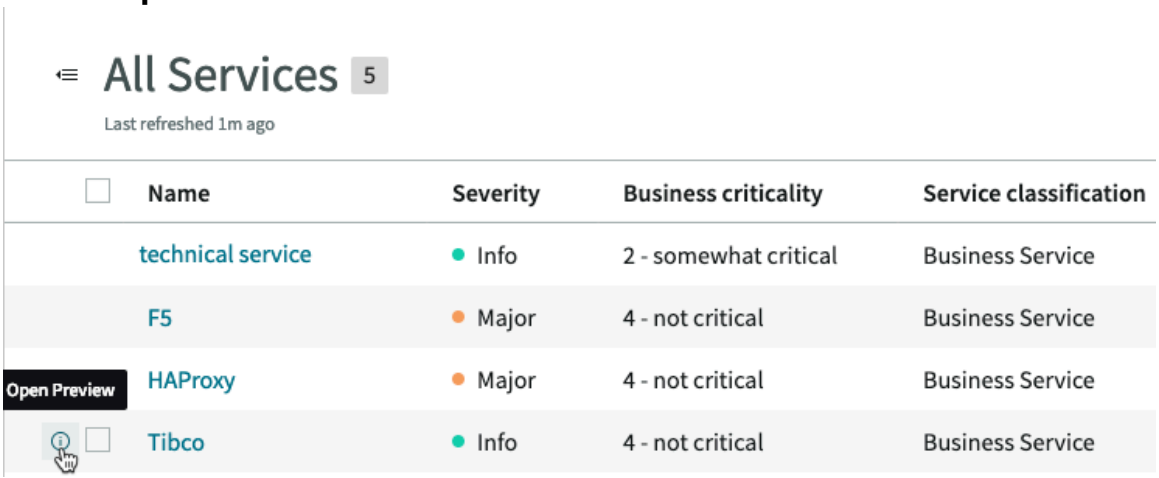

The Details pane opens on the right.

f. Select the appropriate option from the **Business criticality** list.

g. Click **Update**.

Group application services

Organize application services by groups to perform actions simultaneously on multiple services, and to control user access to services. You can use Event Management to track service health by service groups.

Before you begin

Role required: service_mapping_admin or app_service_admin

About this task

Typically, enterprises have hundreds of services which makes it impractical to manage them individually. Service groups can make service lists much shorter and easier to manage, especially in large organizations or service providers.

How you group service instances depends on the user and on service provisioning policies in your enterprise. The relation between service instances in groups is purely logical and the same service instance can belong to multiple groups. For example, the Mobile service can be part of the following service groups: Sales, Beijing, and Telephony.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example of a service instance belonging to different groups

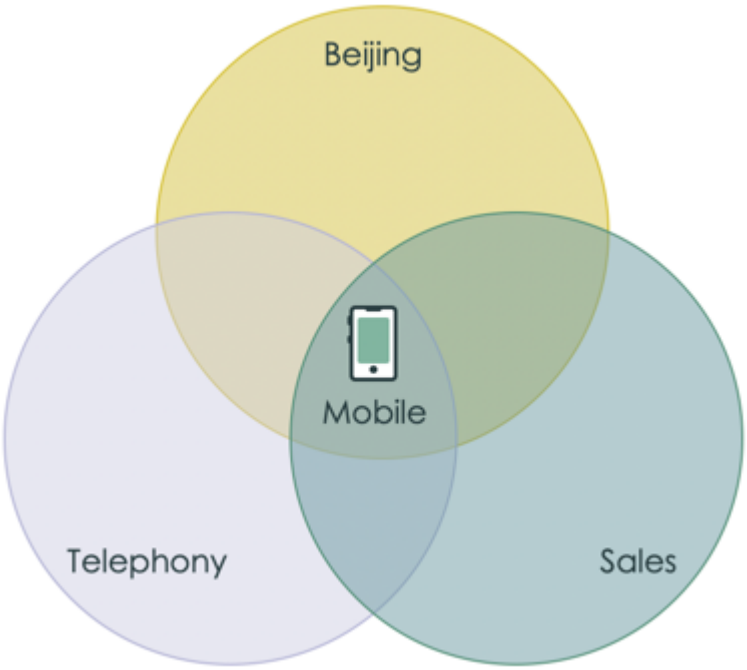

You can embed a service group within another service group to create a hierarchy of service groups. If users have access to a parent service group, they automatically have access to all its child groups. By default, all new services are assigned to the **All** service group that lets all users view and manage service instances. When you assign a role to a service group, the users with this role can access service instances in this service group and in the **All** service group. To enable users with this role to access other services, assign this role to the respective service group. Do not assign user roles directly to the **All** service group.

If Service Mapping is activated, service groups can contain a mixture of manually created application services and application services discovered by Service Mapping.

You can use ServiceNow AI Platform Notifications to alert users if the service group severity changes to critical. The overall severity of the group is determined by the highest alert severity within the group.

Procedure

1. Navigate to **All** > **Configuration** > **Application Services** > **Service Groups**.
2. Click **New**.
3. Enter the name of the new service instance group in the **Name** field.
4. To embed this group in another group, enter the name of the other group in the **Parent Group** field.
5. Right-click the form header and click **Save**.
6. Add a service instance to the newly created service group.

   a. In the Service Group Members section, click **New**.

   b. In the **Name** field, enter the name of the service instance.
   If you are using Event Management, you can also enter an alert group name.

   c. Click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Alternatively, add a service instance to a group from the service instance form.
   a. Navigate to **All** > **Configuration** > **Service Instances** > **Service Instances**.
   b. Select the service instance you want to add to a service group.
   c. In the **Service Group Members** section, double-click **Insert a new row**.
   d. Enter the name of the service group to which you want to add the selected service instance.
   e. Click the **OK** icon ( ).
   f. Click **Update**.

Control user access to application services

Assign user roles to service groups to grant users access to application services in your organization. Your organization may restrict access to some services for security or secrecy reasons.

Before you begin

Make sure that you have performed the user provisioning tasks for the users you want to grant access:

1. Add users to user groups .
2. Create new roles .
3. Assign roles to users or user groups .

Also, make sure that you have created service groups as described in Group application services.

Role required: app_service_admin or service_mapping_admin

About this task

In the base system, the following roles provide access to application services:

app_service_admin

Creates and modifies service instances, creates service groups, views, and edits service instance maps.

app_service_user

Views maps for operational service instances and retrieves service content using the getContent - GET REST API. The itil role that serves as the basic helpdesk technician role contains the app_service_user role.

Service Mapping provides these preconfigured roles:

service_mapping_admin

Sets up the Service Mapping application. Maps, fixes, and maintains service instances. Also performs advanced configuration and customization of the product. Assign this role to application administrators.

service_mapping_user

Views maps for operational service instances to plan change or migration, as well as analyze the continuity and availability of services. Assign this role to application users.

sm_app_owner

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provides information necessary for successful mapping of a service instance. Once a service is mapped, this user reviews the results and either approves it or suggests changes. Assign the sm_app_owner role to users who own service instances and are familiar with the infrastructure and applications that make up the services.

Note: Users with the sn_cmdb_editor or itil role only can view all service instances.

Event Management provides these preconfigured roles:

evt_mgmt_admin

Has read and write access to all Event Management features to configure Event Management.

evt_mgmt_operator

In addition to the evt_mgmt_user permissions, can also activate operations on alerts such as acknowledge, close, open incident, and run remediations.

evt_mgmt_user

Has read access to all Event Management features. Has write access to alerts to manage the alert life. Has the itil role to be able to manage incidents that are created from alerts.

evt_mgmt_integration

Has create access to the Event [em_event] and Registered Nodes [em_registered_nodes] tables to integrate with external event sources.

Typically, enterprises have hundreds of services which makes it impractical to manage them individually. Service groups can make service lists much shorter and easier to manage, especially in large organizations or service providers. In a hierarchy of service groups, access to a parent service group automatically grants access to all the child service groups.

Users inherit permissions from roles that are assigned to them. You can assign some roles directly to service groups to allow all users with this role to access all application services belonging to this group. However, most enterprises choose to organize their roles as a hierarchy. It helps to manage roles across multiple ServiceNow applications. For example, the Service Mapping administrator [service_mapping_admin] can be part of a broader administrator role like administrator [admin]. You can add users to user groups and then assign roles to the user groups to give permissions of this role simultaneously to all the group users.

Assigning a role to an application service group

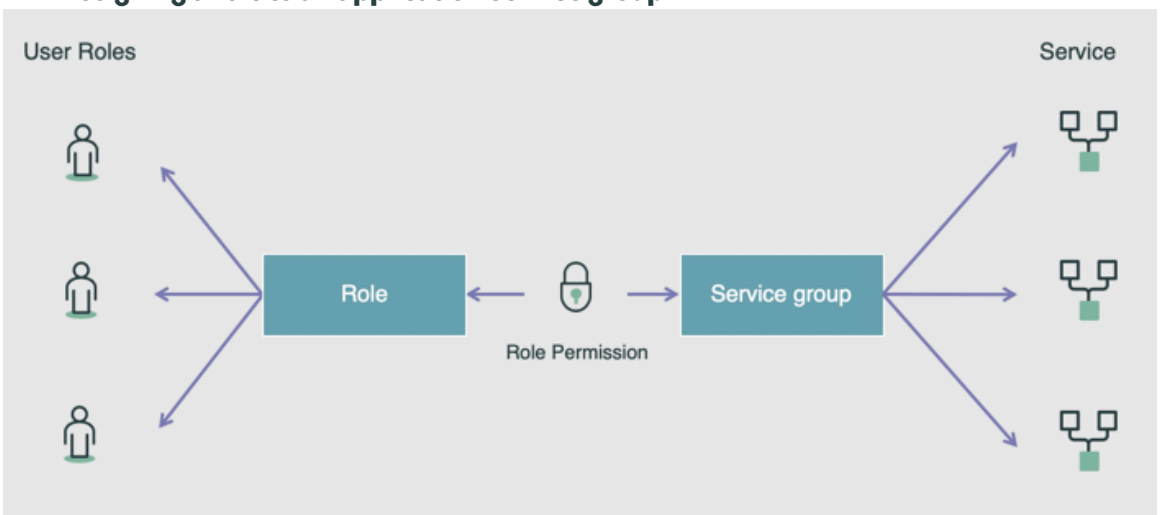

By default, all new services are assigned to the **All** service group that lets all users view and manage service instances. When you assign a role to a service group, the users with this role can

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

access service instances in this service group and in the **All** service group. To enable users with this role to access other services, assign this role to the respective service group. Do not assign user roles directly to the **All** service group.

Procedure

1. Navigate to either of the following:
   - **Configuration** > **Application Services** > **Service Group Responsibilities**.
   - If Service Mapping is activated: **Service Mapping** > **Services** > **Service Group Responsibilities**.
   - If Event Management is activated: **Event Management** > **Services** > **Service Group Responsibilities**.
2. Click **New** and fill out the Application Service Group Responsibilities form.

Field	Description
Application Service Group	Service group to which you want to assign a role.
Role	Role you want to assign to the selected service group. For example, financial_services_admin.

3. Click **Submit**.

Example:

To manage access to services that contain sensitive financial information in your organization:

1. Organize the services into the Financial Services group.
2. Create a new user role, financial services administrator [financial_services_admin] role, that contains the [app_service_admin] role.
3. Assign the Financial Services administrator role to the Financial Services group.

As a result, only users with the Financial Services administrator role can access application services belonging to the Financial Services group.

Schedule a top-down discovery by Service Mapping

After Service Mapping discovers configuration items (CIs) belonging to your service instance for the first time, it then rediscovers CIs to find changes and updates. Create or modify discovery schedules to control how often Service Mapping rediscovers services or CIs. For example, you may create custom discovery schedules to avoid redundant stress on the infrastructure.

Before you begin

Define your application services prior to defining discovery schedules for discovering CIs which are part of these services.

When creating a schedule for CI types or specific CIs, make sure that at least one application service containing the relevant CI is in the *`operational`* state.

Learn about different types of top-down discovery schedules.

Role required: service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Procedure

1. Navigate to **Service Mapping** > **Administration** > **Discovery Schedules**.
2. To create a custom discovery schedule, click **New**.
   Or

   To customize the generic schedule for all application CIs, click **All Applications** or **Load Balancer Service**.

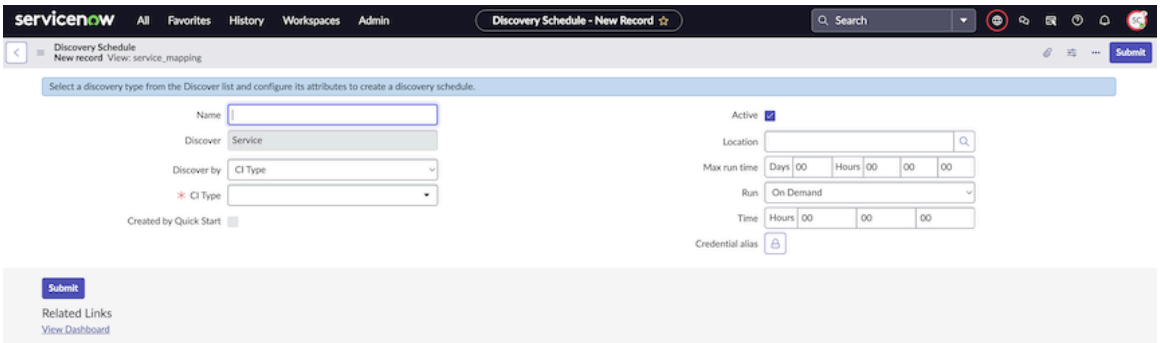

3. In fresh install or reset instances, fill in the fields in the Discovery Schedule form.

Field	Description
Name	For a new specific schedule, enter a unique and descriptive name.
Discover	By default, the value for this field is `Service`, indicating that you are creating a top-down discovery schedule. You cannot change this value.
Discover by	The schedule types for specific top-down discovery: ○ CI Type ○ Specific CI ○ Service Attributes ○ Service Group Attributes
CI type	This field is available when CI type is selected from Select Discovery. For a new CI type-based schedule, select the relevant application CI type from the list.
CI	This field is available when `Specific CI` is selected from **Discover by**. For a new schedule for a specific CI, select the relevant CI from the list.
Filter Services to Discover	This area contains the condition builder and is available when `Service Attributes` is selected from **Discover by**. Use the filter to create the filtering criteria for services. Service Mapping discovers only services that answer the defined criteria.
Filter Groups to Discover	This field contains the condition builder and is available when `Service Group Attributes` is selected from **Discover by**. Use the filter to create the filtering criteria for service groups. Service Mapping discovers only services belonging to the groups that answer the defined criteria.

4. **Optional:** For a discovery schedule based on service attributes, click **View Selected Services** to preview the list of services that answer the filtering criteria you entered, and then select **Close**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. **Optional:** For a discovery schedule based on service group attributes, click **View Selected Groups** to preview the list of groups that answer the filtering criteria you entered, and then select **Close**.
6. Define the rediscovery frequency by selecting the relevant option from the **Run** list. See Run options for discovery schedules.
7. Click **Submit**.

Related topics

Pattern-based discovery in Service Mapping

Top-down discovery schedules

Learn about schedules that trigger top-down discovery of application services.

When you define a new application service, Service Mapping performs discovery of all CIs that participate in this application service and creates its map. After the initial mapping is complete, Service Mapping regenerates application service maps regularly by rediscovering the CIs making up an application service.

By default, Service Mapping is preconfigured with these generic schedules:

All Applications

This schedule triggers the top-down discovery of all CIs of the application type [cmdb_ci_appl].

Load Balancer Service

This schedule starts the top-down discovery of all CIs of the load balancer service type [cmdb_ci_lb_service].

The generic schedules trigger discovery of all applications and load balancer services in your organization once a day. Typically, these preconfigured schedules are enough to update information for application services. However, if it is necessary to discover specific CIs or services more frequently, you can create the following custom discovery schedules:

Schedule types for top-down discovery

Schedule type	Description	Example
Services	Service Mapping discovers all CIs belonging to certain services that answer filtering criteria.	In your organization, some application services are more critical than others and it is important for you to rediscover such services with high criticality more frequently. Create a custom discovery schedule to discover all services with the business criticality value set to *`most critical`*. Alternatively, certain services in your organization are undergoing changes more often than once a day. You may want to create custom schedules to rediscover such services more frequently. You can also create a schedule of this type to discover application

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Schedule types for top-down discovery (continued)

Schedule type	Description	Example
		services in a sub-production instance, where services are not in the *operational* state yet.
Services belonging to a group	Service Mapping discovers all CIs belonging to services assigned to groups that answer filtering criteria. Service Mapping discovers only services located directly inside the group that answers defined filtering criteria. Service Mapping does not discover services inside embedded groups. Notice that you can define schedules based on service groups, even if your role does not have access to these service groups.	In a distributed organization with offices in several geographic locations, you may want to create a custom discovery schedule to rediscover services in one of these locations. For example, one schedule can trigger rediscovery of all services relevant for the EMEA site at 7am GMT. Another schedule can start rediscovery of services for the US headquarters at 7am PT.
For CI types	Service Mapping discovers all application CIs belonging to this application type. Some application CI types are prone to more frequent changes and updates than others, so you can manage the load by adjusting the discovery schedule to match the nature of each CI type. When you define discovery schedules based on application CI types, several schedules may apply to the same CI. To avoid discovering the same CIs more than once, the most specific schedule takes precedence. If you create separate schedules for a parent CI type and its child CI type, CIs belonging to the child type use the child schedule. If no child schedule exists, the parent schedule applies. This schedule discovers only CIs belonging to application services in the *operational* state.	If you modify a certain application more often than the rest of your applications, you may want to discover it more frequently than other CIs. In that case, create a custom discovery schedule for such an application CI type.
Specific CIs	Service Mapping discovers only one CI that you specified for this schedule. If you define a discovery schedule for a specific CI and a schedule for its CI type, Service Mapping uses the CI-specific schedule. The	For rare cases of rediscovering a specific CI that causes discovery errors frequently.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Schedule types for top-down discovery (continued)

Schedule type	Description	Example
	generic CI type schedule does not apply. This schedule discovers only CIs belonging to application services in the *operational* state.	

Part of discovering an application CI is identifying its host. Service Mapping checks if the device hosting this application CI exists in the CMDB. For a load balancer service, Service Mapping checks if the load balancing device hosting this load balancer service exists in the CMDB. If not, Service Mapping triggers Discovery to perform the horizontal discovery. As a result, Discovery performs host detection and updates the information on hosts in the CMDB. If necessary, you can manage the schedules that trigger horizontal discovery as described in:

- Schedule a horizontal discovery — for standard horizontal discovery.
- Create a schedule for standard serverless Discovery — for discovering CIs using a protocol such as REST, without using the IP address of any devices or hosts.
- Create a schedule for host-based serverless Discovery — for discovering CIs that reside on an already discovered host using this host as a Proxy server.

Related topics

Schedule a top-down discovery by Service Mapping

Service definition transfer from one instance to another

Save time and effort by exporting definitions of service instances from the source instance and importing these definitions into the target instance. You can use this method to copy definitions of service instances only of the discovered type.

Typically, you maintain two ServiceNow instances in your organization: for production and for development. Perform the initial mapping of application services in your organization on the development instance, which you use for testing. The results of initial mapping often require some basic troubleshooting or fine-tuning. Once you finish mapping application services on your source instance, transfer service definitions into the target (production) instance.

Notice that you cannot export the actual application services, only service definitions. Mapping is a dynamic process, and the mapping result reflects the actual infrastructure and its configuration. Therefore, the services you exported from one instance and imported into another may look different.

By default, you export the following service definitions:

- Service name
- Traffic-based discovery usage
- Entry point
- Business criticality
- Group to which the service belongs

  If this service group is part of a service group hierarchy, all parent service groups in this hierarchy are exported as well.

- Comments

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The system uses the Sys ID [sys_id] attribute to identify services on the source and target instances and implements service definition import according to these rules:

- If a service Sys ID does not exist on the target instance, the system creates a new service on the target instance.
- If a service Sys ID and name already exist on the target instance, the system updates this service by overriding the existing service definitions with the imported definitions.
- If a service Sys ID exists on the target instance, but is associated with a different service name, the system updates the service with this Sys ID. The system overrides service definitions, including the service name.
- If the same name is used for different services on the source and target instances, the system does not import service definitions.
- You cannot export manual services. Likewise, you cannot export manually added entry points. When importing definitions of a service containing manual entry points, the system does not make any changes to the entry points.

Exporting service definitions flow

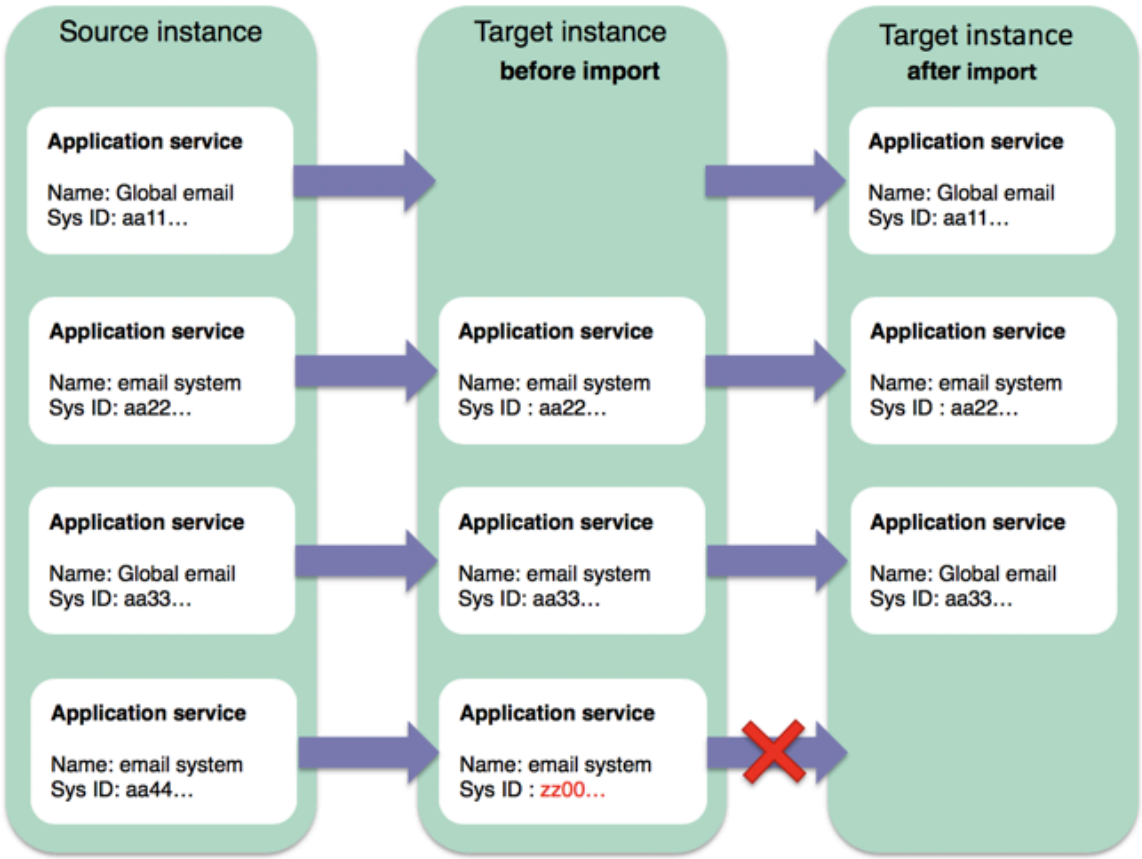

When transferring service definitions between instances, imported service instances and service groups are assigned to the user's domain.

By default, the system assigns *Non-operational* status to application services whose definitions you imported.

A service definition transfer deploys the standard import set feature available on the ServiceNow AI Platform. When you export service definitions on the source instance, the system creates a JSON file containing exported data. When you copy service definitions into the target instance, you import the data from this JSON file, which becomes an import set.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Export service definitions

Create an import set containing definitions of services. You use this import set to import service definitions into another instance.

Before you begin

- Decide which service definitions you want to export.

  Important: You can use this method to copy definitions of application serviced of the discovered type only.

- Verify that the source and the target instances are on the same ServiceNow release.
- Open the source instance from which you want to copy service definitions.
- If necessary, fine-tune the import set creation by modifying relevant properties under **Service Mapping** > **Administration** > **Properties**. For more information, see Components installed with Service Mapping.

Role required: service_mapping_admin

Procedure

1. If the instance uses domain separation and you have access to the global domain, select the domain to which the service instance belongs from the domain picker (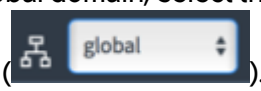

). The selected domain must be a domain without any child domains.
2. Navigate to **Service Mapping** > **Services** > **Application Services**.
3. Select the application services that you want to export.
4. Click **Actions on selected rows** at the bottom of the page and select **Export Services**. The import set appears in the `Downloads` folder. The file name follows this convention: export_service_<datetime>.json.
5. If necessary, copy the import set to a different location on the computer.

Import service definitions from one instance to another

Import definitions of application services from the source instance.

Before you begin

- Verify that the source and the target instances are on the same ServiceNow release.
- Open the target instance into which you want to import service definitions.

Role required: import_admin

Procedure

1. If the instance uses domain separation and you have access to the global domain, select the domain to which the service instance belongs from the domain picker (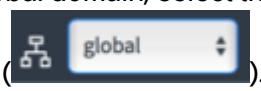

). The selected domain must be a domain without any child domains.
2. Navigate to **Service Mapping** > **Services** > **Import Service Definitions**.
3. In the Import Service Definitions window, click **Choose File**, navigate to the JSON file that contains service definitions, and click **OK**.
4. **Optional:** Click **Show Preview** and check the definitions in the chosen JSON file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Click **Import**.
6. Click **Yes** to confirm that you are ready to import service definitions.

> **Important:** During import, the system irreversibly overrides the current service definitions.

The Progress window displays the success message.

Service Mapping automatically starts the discovery of services containing imported entry points.

7. Review imported or updated services after Service Mapping finishes the mapping process. If the transfer result is not as you expected, perform basic troubleshooting as described in the following table.

Trouble?

Symptom	Cause	Solution
Changes you have recently made to an service instance disappear.	Definitions of existing application services are always overridden during transfer.	1. Make necessary changes to the definitions of this service instance again. 2. Run the top-down discovery of this service instance by clicking **Run Discovery** in the service map.
An imported service group does not appear.	The import of an application service assigned to this group failed and the system did not create the group because it appears to be empty.	1. Resolve the issue that prevented the system from importing the service instance belonging to this group. 2. Import the import set again. 3. Check that the service group appears.
An imported service group appears to have wrong application services assigned to it.	This is a different service group with the same name, not the service group you imported. The system failed to import the service group. The source instance already has a service group with the same name or with the same name and Sys ID.	1. Check the import log. 2. If an error states that a record creation failed while running the Check Name Uniqueness business rule, perform the following steps: a. Rename the service group on the source instance. b. Re-export the application services assigned to this group. c. Import the latest import set on the target instance. 3. If an error states that a record creation failed because the Sys ID already exists, perform the following steps:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
		a. Recreate the service group on the source instance, ensuring you provide a different group name and Sys ID to it. b. Re-export the application services assigned to this group. c. Import the latest import set on the target instance. 4. Verify that the service group appears correctly on the source instance.

Application service analysis and maintenance using classic Service Mapping

Service Mapping creates maps to help you see the architecture and organization of application services. These maps are useful for planning change or migration, as well as analyzing the continuity and availability of services.

The maps are useful in the following cases:

Planning the migration of an entire service instance or its segments

Use the map to see which configuration items (CIs) need relocating or which CIs become redundant.

Planning the upgrade or replacement of a CI

Check the map to see what other CIs are affected when the CI is non-operational during maintenance. You can assess and plan the downtime or make provisions to avoid the downtime.

Because the same CI may be used for multiple application services, check the CI using the Dependency Views application. With this application, you can see all application services that a CI belongs to.

Analyzing a service instance for high availability and continuity

Use the map to identify CIs crucial for the service performance. Decide if you want to fortify these CIs by creating clusters.

Troubleshooting application services

Use the map to understand the impact of an issue and determine which CI is causing the problem.

Application service maps in classic Service Mapping

Maps offer you a visualization of data on configuration items (CIs) comprising application services, and relations and connections between these CIs.

A new map is created every time you manually create an application service or map an application service using Service Mapping. The system updates the map every time you add CIs to the application service or when Service Mapping runs discovery for CIs belonging to this application service.

Note:

Information and actions available on a map depend on your role:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- As a Service user [app_service_user] or Service Mapping user [service_mapping_user], you can view the map in the View mode.
- As a Service administrator [app_service_admin] or Service Mapping administrator [service_mapping_admin], you can view application services in View mode and modify discovered and manually created services in Edit mode.

Important: You cannot fine-tune or edit tag-based and dynamic services from the map.

- As a Service Mapping administrator [service_mapping_admin], you can view discovery messages and errors generated by Service Mapping and modify services in the Edit mode. You can also reload the map to ensure the map reflects the latest information.
- As a Service owner [sm_app_owner] role, you can review, provide feedback, and approve services discovered using Service Mapping in the View mode.

Note: The review and approval process is available only for discovered and manually created service instances.

Also, actions and information related to automatic discovery using Service Mapping are not available for CIs that Service Mapping cannot discover. For example, you do not see the tab for Discovery Messages for a manually created service instance that Service Mapping does not discover.

Map window

Aside from the map itself, the map window also displays the Properties pane and tabs with additional information.

Map window of Service Mapping

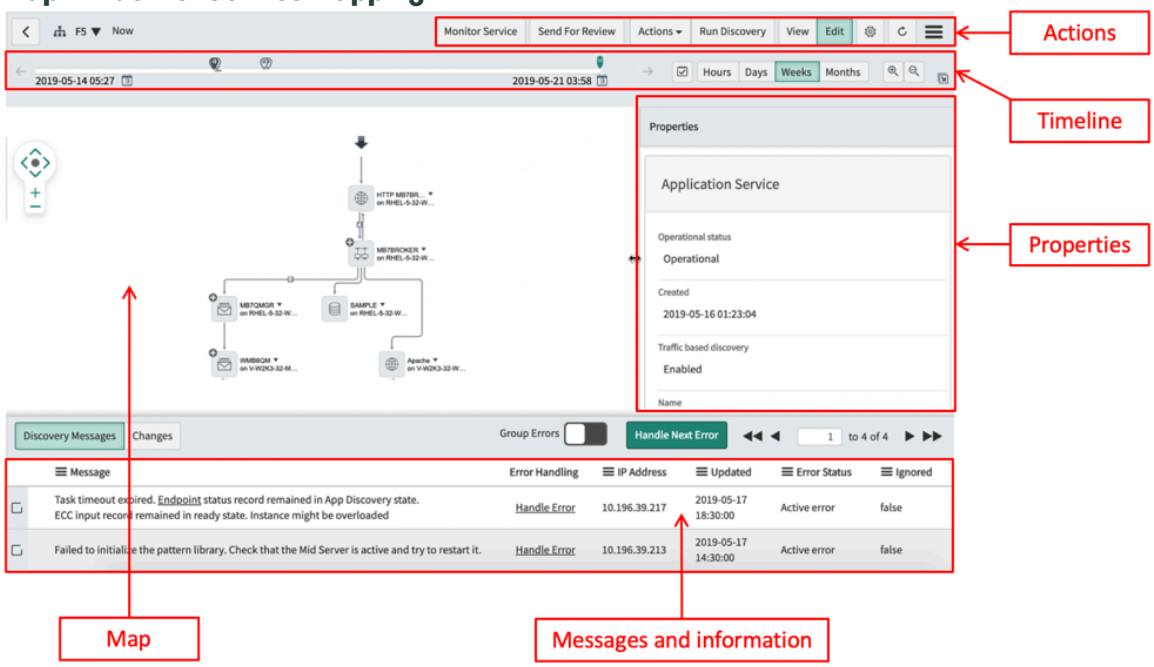

You can navigate to a different application service map directly from this window by selecting it from the list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service list

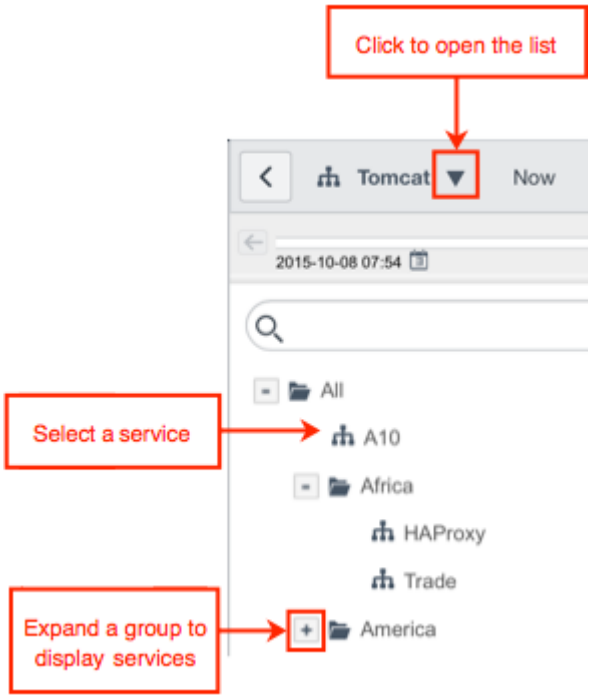

By default, the entire map is shown in the center of the visible map area of the window. You can zoom in and out as well as position the map depending on which segment of the map you must see using these controls:

Map navigation

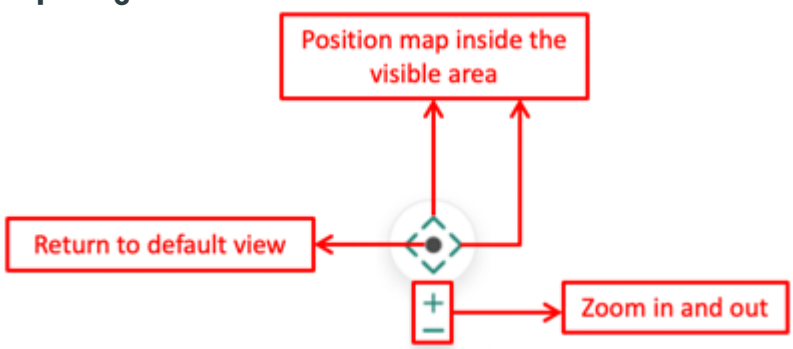

You can also click anywhere in the map area and drag the required segment of the map into the visible area.

You can view changes made to an service instance as a whole and to individual CIs belonging to a service by choosing a time range. For more information, see View the change history of application services in classic Service Mapping.

Map elements and their appearance

Every map consists of icons representing CIs and arrows that represent the connections between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connected CIs

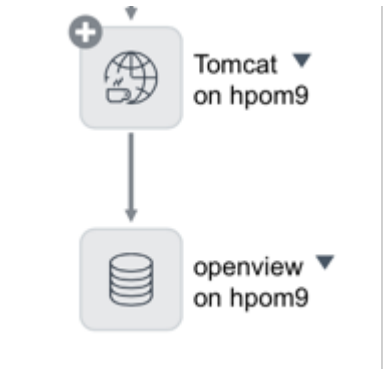

Some map elements can contain other elements such as in the case of application clusters. On the map, application clusters appear as a stack of CIs with the plus (+) symbol next to it. The cluster label shows the number of CIs in this cluster with the multiplication sign.

CI cluster

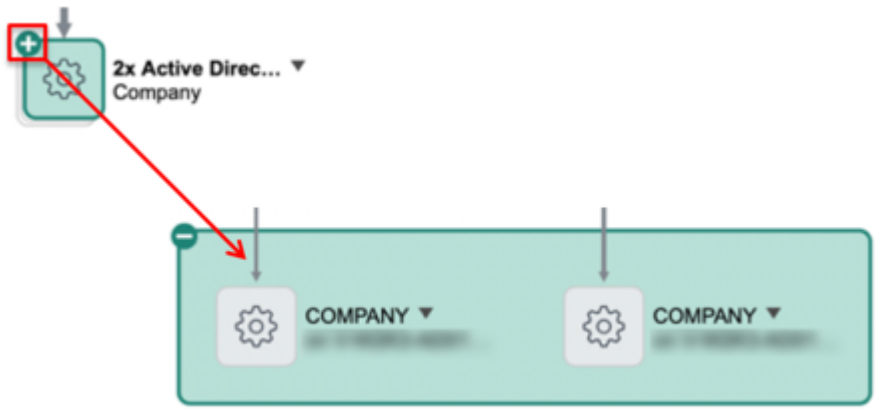

When the CI has inclusions, the CI icon has the plus (+) symbol next to it and shows the contained application when you expand it.

Inclusion

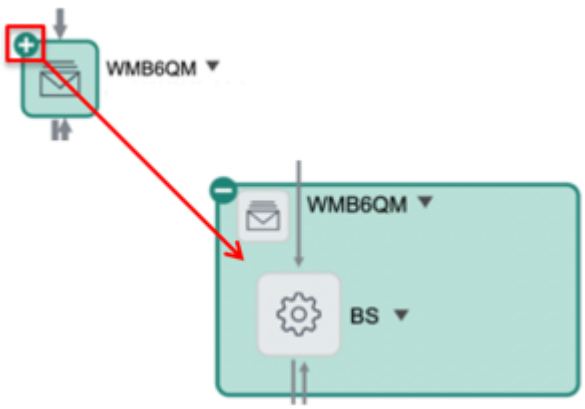

An OS cluster appears as a CI with the plus (+) symbol and the number of CIs in this cluster.

OS cluster CI

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI attributes and CI-related messages

When a map is loaded and no elements are selected, the Properties pane shows the details of this service instance.

A selected device, application, or connector appears in blue and is highlighted. Information about the selected map element is displayed in the Properties pane on the right of the map. The Properties pane contains the links to the Detailed Properties page.

Connection information in the Properties pane

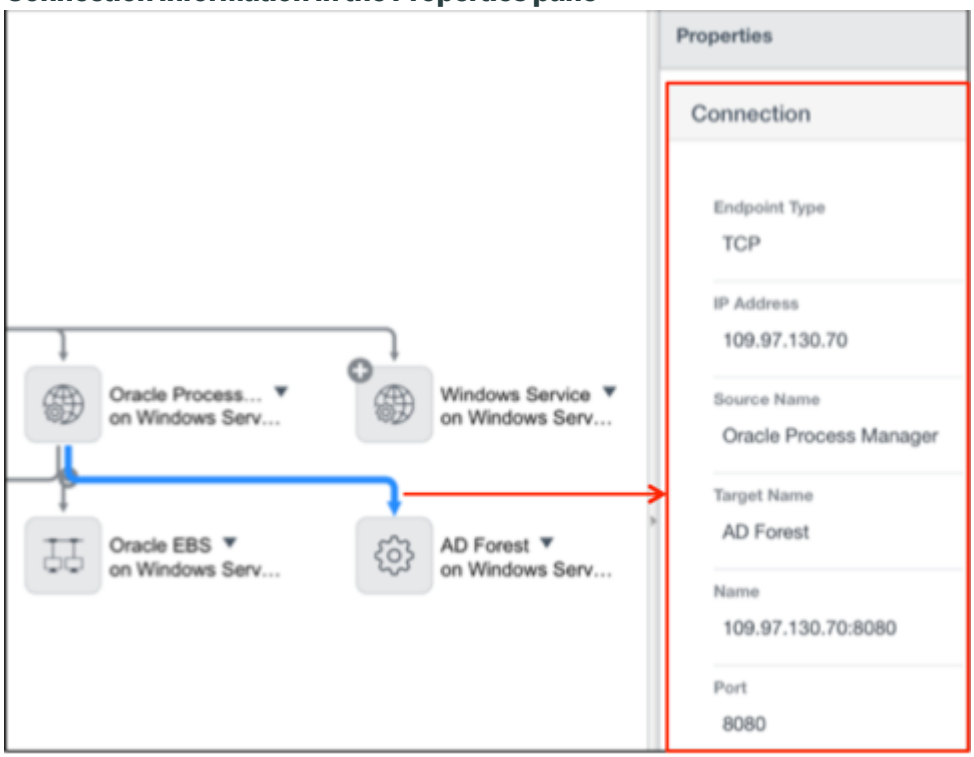

The attributes for the server hosting applications and applications themselves are shown separately inside the Properties pane.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Server information in the Properties pane

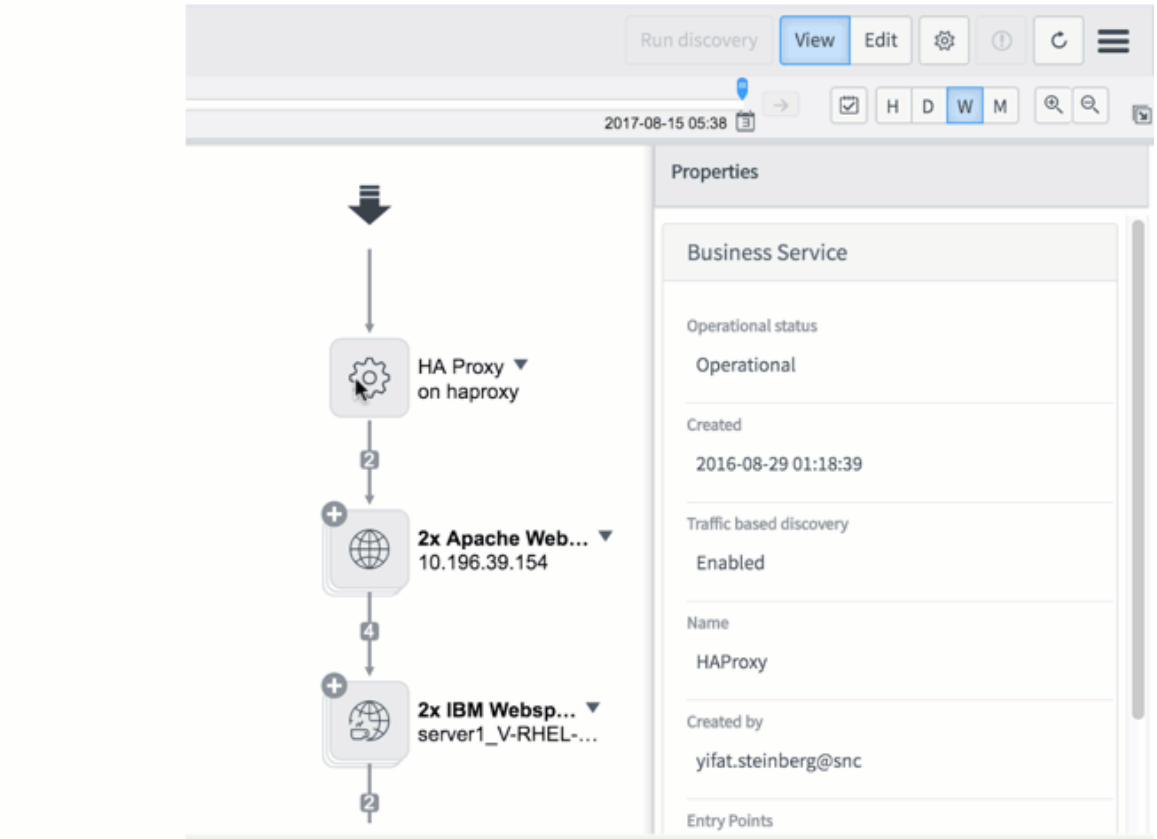

If there is any information related to the selected CI or connection, it is highlighted on the tab at the bottom of the window.

The CI and the related discovery message

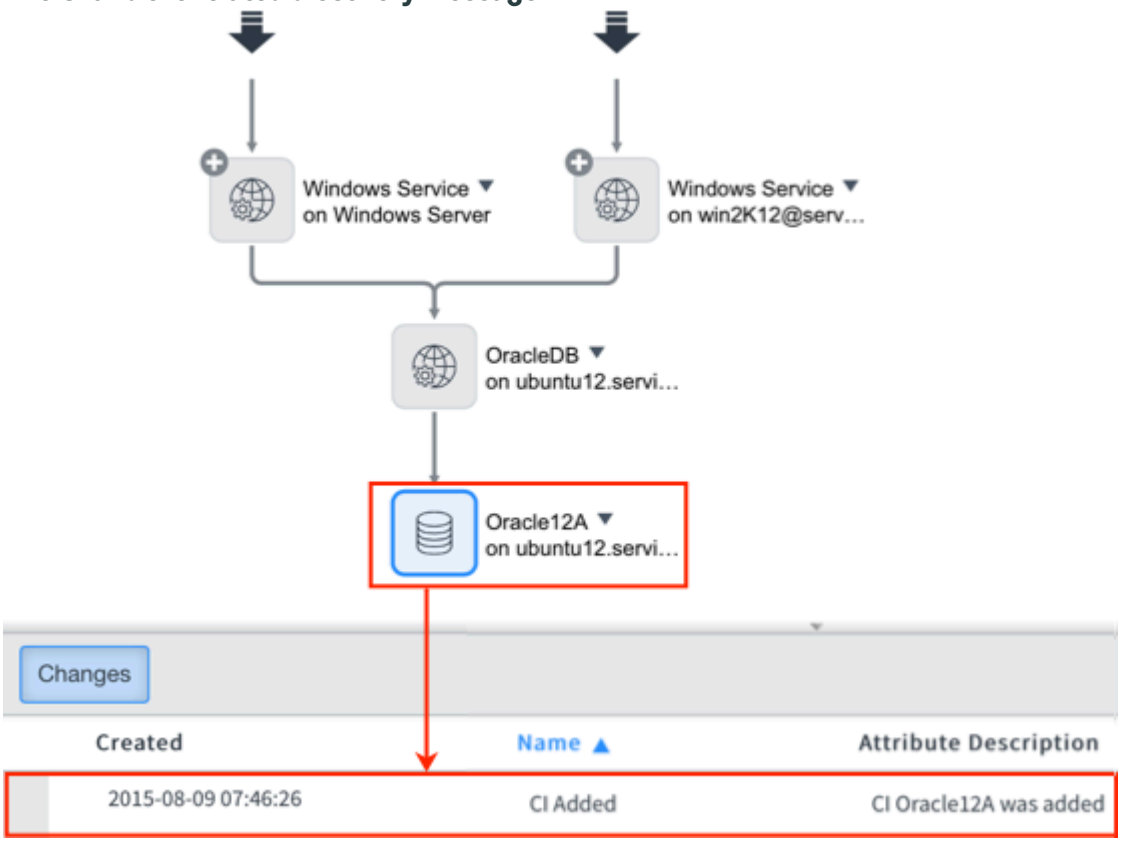

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When you select information on a tab at the bottom of the map window, the related CI displays in yellow.

The message and the CI it is generated for

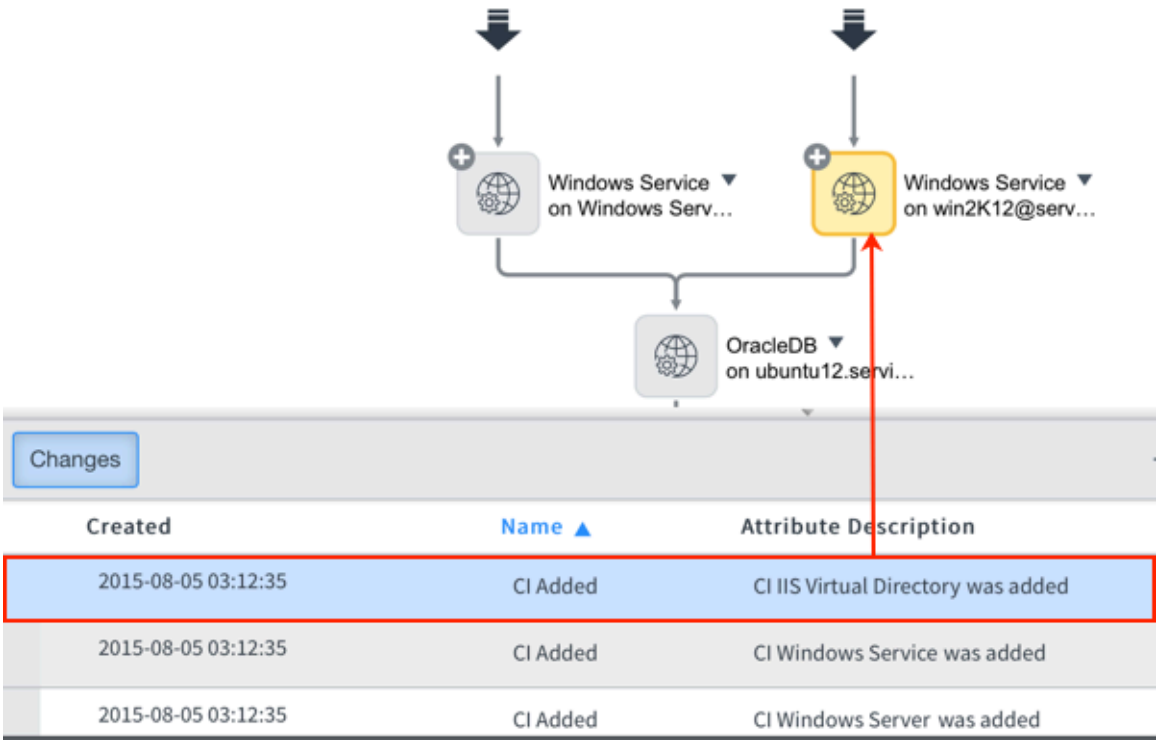

Map indicators

You can show additional information for a CI or service instance itself by displaying its related records such as alerts, outages, incidents, and problems. If you enable this feature, record indicators appear next to relevant CIs and next to the service instance name. Records themselves are displayed on the tabs at the bottom of the map window.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application service with indicators

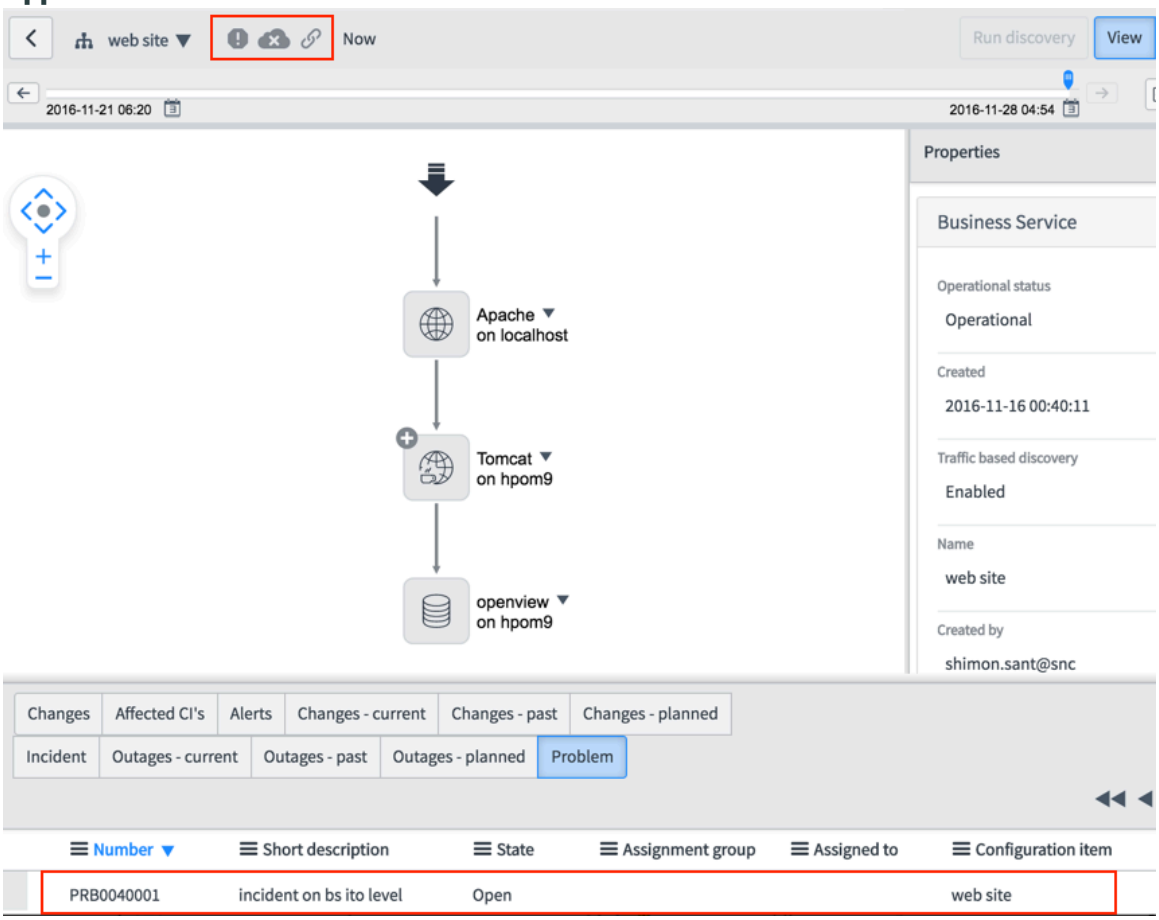

Map views

Customize a map view to exclude irrelevant information and increase clarity, as described in Modify view for an application service map.

To open this service instance in Event Management and view alerts on CIs that are part of the service instance, click **Monitor Service** at the top of the window.

Monitor Service button

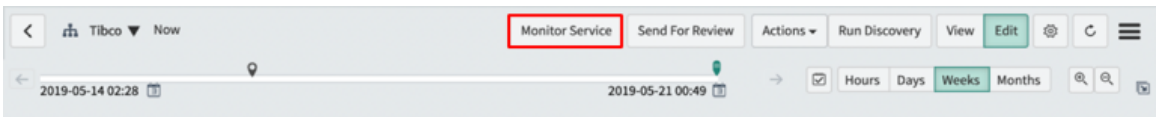

Related topics

- View CI attributes in an application service map in classic Service Mapping
- Modify display for CI attributes
- Modify view for an application service map
- View an alert impact on CIs in a service map

View CI attributes in an application service map in classic Service Mapping

An application service map displays attributes for each configuration item (CI) that is part of the application service, as well as for the application service itself. The attributes come from the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: service_mapping_admin or service_mapping_user

About this task
You can view the following information for each CI:

Name label

The CI name. This attribute is either preconfigured on the CI or configured during CI installation.

Basic attributes

A summary of the most important CI attributes.

Detailed attributes

A complete list of all attributes collected for the CI.

Each CI type (CI class) has different attributes. For example, the Linux Server type has different attributes than the SQL Instance type.

If Service Mapping is deployed, the way CIs appear on the map depends on the view you select for the map. Attributes available for viewing also depend on the Service Mapping setup. For more information, see description of components installed with Service Mapping.

Procedure
1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. If needed, click **Edit** to ensure that the map is in Edit mode.

   If Service Mapping is deployed, then in Edit mode, the Discovery Messages section appears below the map.
3. To see the full name of a CI whose name has been shortened on the map, point to the CI.

   A tooltip displays the full CI name.
4. Click a CI to see its details in the **Properties** pane.

   The attributes of applications and the servers that host them appear separately.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

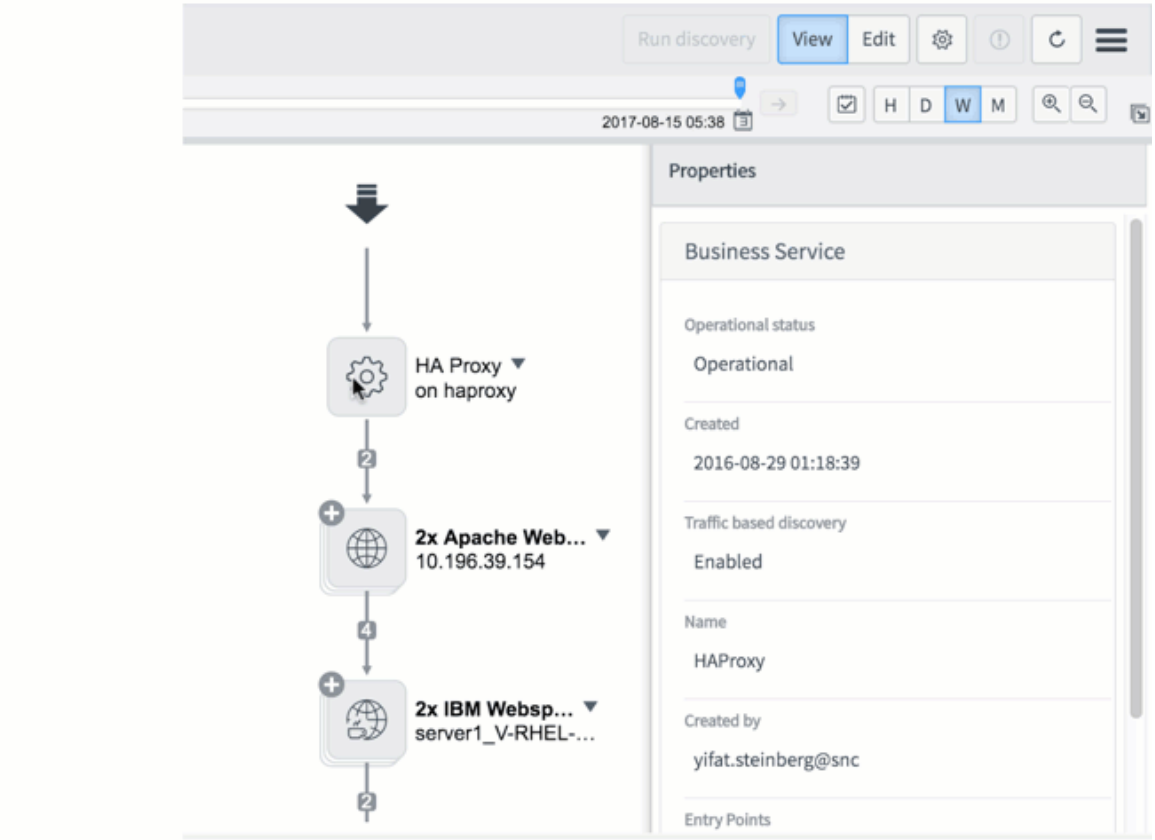

5. To view more detailed attributes for the CI, click **Detailed properties** at the bottom of the Properties pane.
6. **Optional:** To view configuration files associated with a CI in environments where Service Mapping or Discovery are enabled and tracking changes to CI configuration files is enabled):
   - Review the list of files under **Tracked Files** in the **Properties** pane. Click the file name to open the actual file.
   - Click the **Affected CIs** tab and view the list of configuration files. Click the file name to open the actual file.

View CI connection attributes in Service Mapping

You can view attributes of connections between configuration items (CIs) in application service maps in classic Service Mapping.

Before you begin

Role required: service_mapping_admin or service_mapping_user

About this task

You can view the following information for each connection:

- The source and target CIs of this connection displayed in the Context (right-click) menu.
- A complete list of all attributes collected about the connection in the Properties pane.

By default, Service Mapping merges connection lines for the same CI to declutter an application service map. It helps to make the map more readable.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Merged connection lines

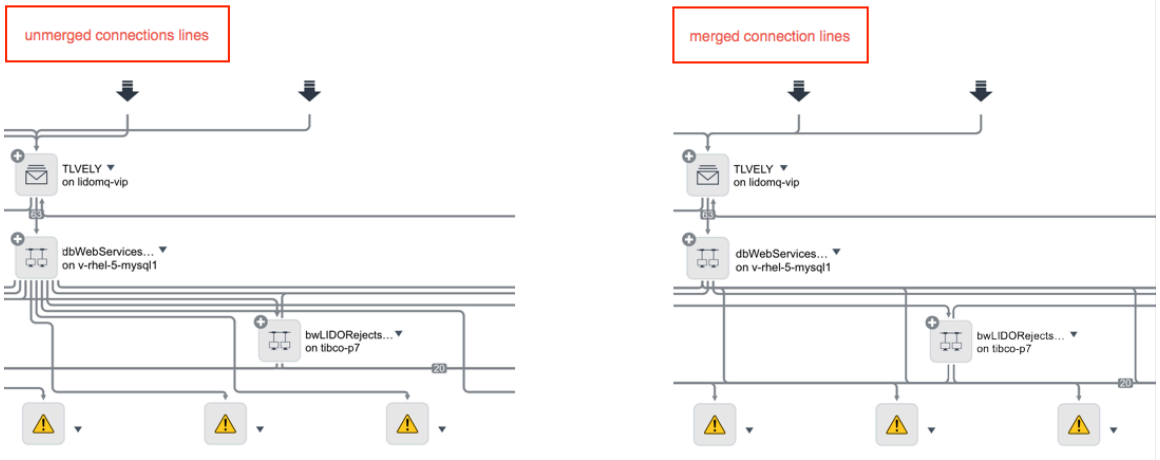

Each connection type has different attributes. For a merged connection line, all underlying connections are listed.

The way an service instance map displays connection attributes depends on the following factors:

- Views you configure for CIs define which connection attributes appear in the Properties pane of the map.
- Connection lines on the map depend on the view you select for the map.
- Attributes available for viewing also depend on the Service Mapping setup, for more information see description of components installed with Service Mapping.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. View connection attributes as follows:

To view	Do this
The source and target CIs of a regular connection	Right-click the connection. 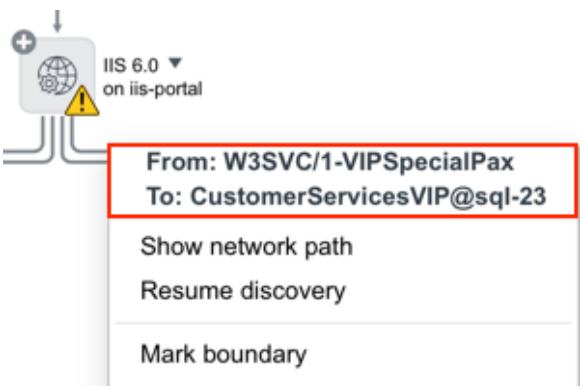

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To view	Do this
(Available for application services discovered with Service Mapping) The source and target CIs for a merged connection	○ Right-click the merged connection line with a number and select the relevant connection. Or ○ If a segment of a connection line is shared by more than one merged connections, there is no indication that it is a merged line. Right-click the connection line coming out or going into a CI, and select the relevant connection. 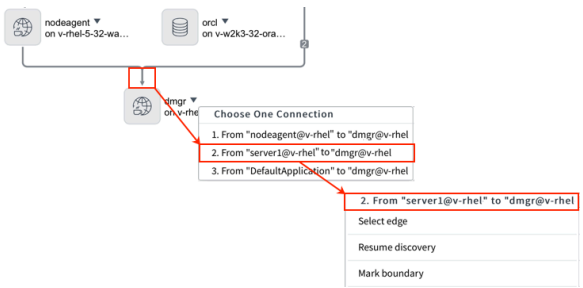
The source and target CIs for a connection if the spanning tree view is enabled. The spanning tree view simplifies an service instance map concealing most of the connection lines.	**a.** Click the CI whose connections you want to view. All concealed connections for this CI appear on the map. **b.** Right-click the relevant connection line.
Detailed attributes of a regular connection in the **Properties** pane	Click the connection line.
(Available for application services discovered with	Click the merged connection line indicated by a number.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To view	Do this
Service Mapping) Detailed attributes of a merged connection in the **Properties** pane.	All underlying connections are displayed separately in the **Properties** pane. 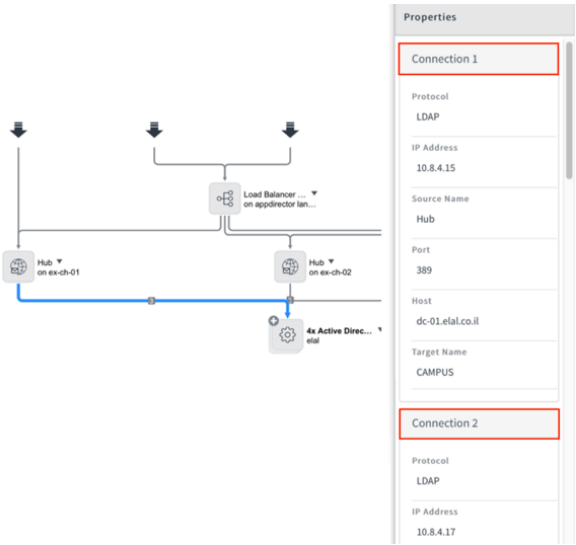
Detailed attributes of a connection segment shared by more than one merged connection in the **Properties** pane.	Right-click the connection line coming out or going into a CI, select the relevant connection, and then choose **Select edge**. The detailed attributes are displayed in the **Properties** pane. **Note:** Shared segments do not have an indication that it is a merged line.
(Available for application services discovered with Service Mapping) Detailed properties of a connection in the **Properties** pane, when the spanning tree view is enabled. The spanning tree view simplifies an service instance map concealing most of the connection lines.	**a.** Click the CI whose connections you want to see. All concealed connections for this CI appear on the map. **b.** Click the required connection line. 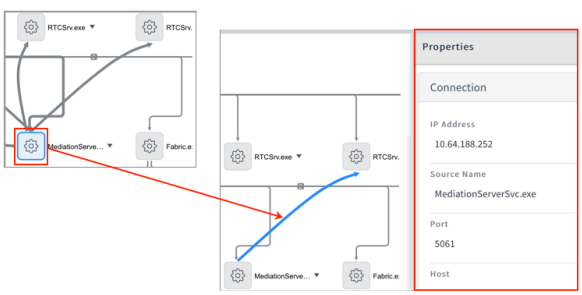

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To view	Do this
(Available for application services discovered with Service Mapping) Traffic-based connections for a CI. **Note:** You can see traffic-based connections for a CI, even if you disabled the traffic-based discovery for an service instance the CI is part of.	**a.** Right-click the CI whose traffic-based connections you want to check. **b.** Select **Show traffic based connections**. The Traffic Based Connections List opens displaying the information in Traffic Based Connections List table. **c.** Click **Close** when finished viewing this list.
Manually added connection. You can either add a connection or a CI manually. In both cases, the connection you manually add an entry point to the next tier of the application flow.	Click the manually added connection. The **Properties** pane shows Manual Endpoint under **Endpoint Type**. 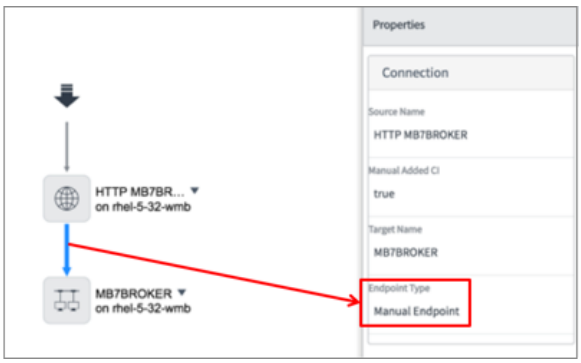

Traffic Based Connections List (when discovery based on Predictive Intelligence is disabled)

Field	Description
IP	The IP address of the application connected to the selected CI.
Port	The port on the selected CI that is used to communicate to the other application.
Process	The ID of the process in the selected CI.
Already on map	○ **Yes** — if this connection shows on the map. ○ **No** — if this connection is not part of the service instance and not on the map.
System decision	The setting defines if Service Mapping keeps the discovered traffic-based connection

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	or removes it. The value comes from the algorithm that Service Mapping uses.
User decision	(Optional) The setting overrides the System decision setting, which defines if Service Mapping keeps the discovered traffic-based connection or removes it. For example, if the System decision setting for a connection is Keep, and you want to remove this connection, select `Remove`.

View the change history of application services in classic Service Mapping

You can view the changes made to an application service as a whole and to the individual configuration items (CIs) comprising the service. Change history is useful for maintenance, planning, or troubleshooting procedures.

Before you begin

Role required: sn_cmdb_admin, admin, service_mapping_admin, service_mapping_user, app_service_admin, or app_service_user

About this task

Details about changes to a service instance and to its CIs are stored in the CMDB. Typically, these changes reflect adding or removing CIs from a service instance, upgrading or updating CIs, or modifying CI configuration files. The system gathers this data by querying CMDB tables and then creating the change history view. In deployments where Service Mapping is activated, the type of change information Service Mapping queries depends on discovery patterns that Service Mapping uses to discover CIs.

Changes to configuration files are associated with CIs to which these files belong. Maps show configuration file changes as changes to related CIs.

While you can see change records for a specific CI in the context of application services, you can also see detailed history of a specific CI separate from its service instance as described in History Timeline.

If the ServiceNow AI Platform is configured to validate changes, all changes are evaluated and rendered as valid or not. If a change is valid, its change record on the service instance map is marked as approved. For more information about configuring the platform for change validation, see Managing proposed changes.

Changes to the service instance appear on the history timeline.

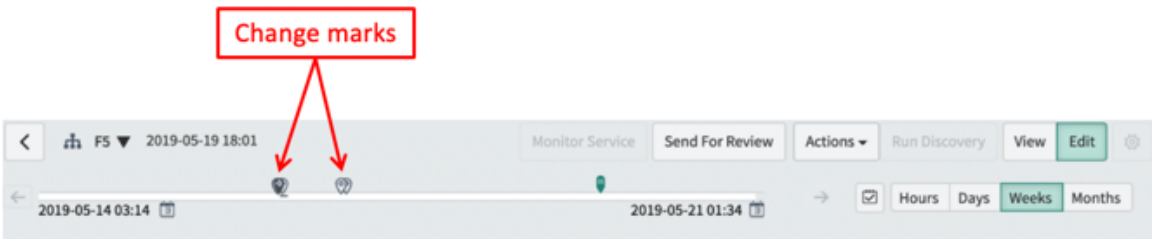

The type of change mark depends on the nature of changes that it represents:

Light gray balloon ( )

Unapproved change that does not influence the service instance behavior. For example, a change in a network path or adding a node to a cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dark gray balloon ( )

Unapproved change that changes the service instance behavior.

Green balloon ( )

An approved change in deployments where the ServiceNow AI Platform is configured to validate changes.

Double balloon ( )

Multiple separate changes that happened a short time from each other.

You can mark times on the history scale by creating baselines to quickly return to the marked view.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Review change records created for this service instance on the **Changes** tab at the bottom of the page.

   If Service Mapping is deployed, then in Edit mode, the Discovery Messages section appears below the map.
3. On the history timeline, set the time range of changes that you want to view.

Option	Action
To set the time range of the history timeline	Click the hour, day, week, or month icons. 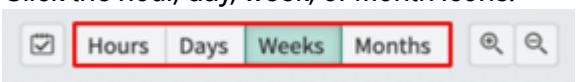
To increase or decrease the time range	Click the zoom in and zoom out icons. 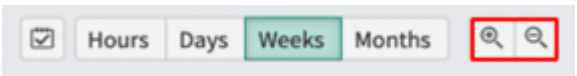
To change the upper limit on your history range	Click the history scale. 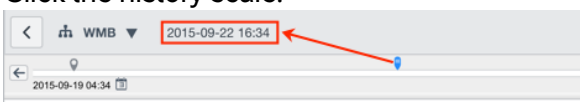 The time that serves as the upper limit ap pears above the history timeline.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
	Note: You cannot set the lower limit on your history range to a time before this ser vice instance was created. This time is marked with the **IT Service Created** event on the history timeline. 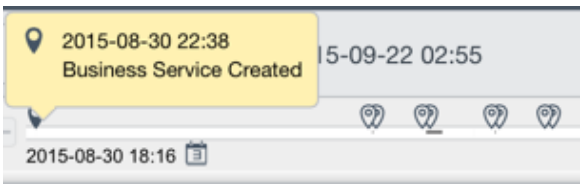

The map shows the history view of the service instance for the time you selected.

Note: The **Change** tab shows all change records, even the ones which are filtered out of the history view.

4. To mark a time on the time scale, set a baseline:

   a. Click the **Compare** icon.

   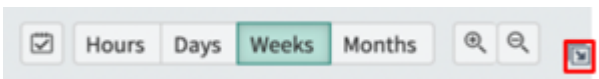

   b. Navigate to the time you want to mark as a baseline on the history scale.

   c. Click **Set baseline**.

   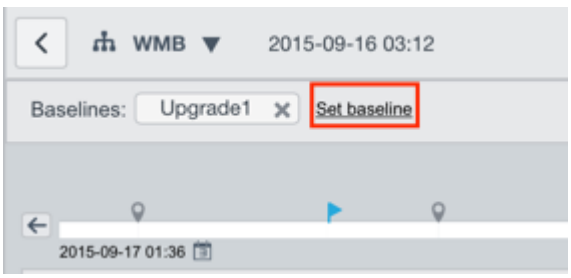

   d. Enter the name of the baseline and click **OK**.
   The new baseline appears as a button above the history scale and as a blue flag on the history scale.

   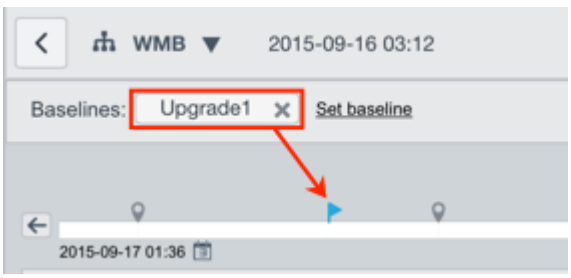

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. View the change history.

Option	Action
To see the CI responsible for a change record	Select a change record on the **Changes** tab. The related CI is marked yellow in the map. 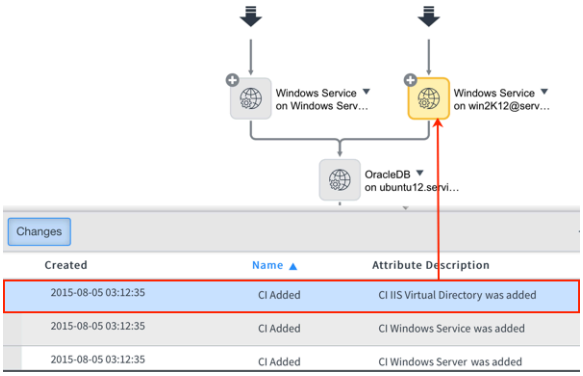
To see only change records related to a CI	Select the required CI or the connection on the map. The **Changes** tab displays only change records related to the selected CI or connec tion. 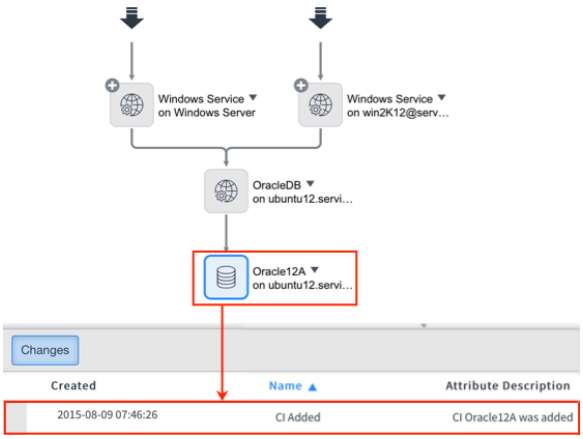
To see the configuration file at the selected moment in the past	**a.** Set the time on the history scale. **b.** In the **Properties** pane, scroll to **Tracked Configuration Files**, and click the file name. The new tab opens displaying the content of the tracked configuration file at the se lected time.
To see the network at the selected moment in the past	**a.** Set the time on the history scale. **b.** Right-click the connection and select **Show network path**. The new tab opens displaying the network or storage path map for the time you select ed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
	Note: You cannot view the network path for connections marked as bound aries to this service instance.

6. To exit the history view and see the current status of the service instance, click the current icon.

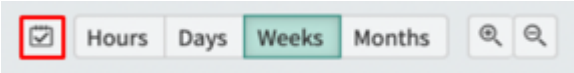

Related topics

View an application service map in base system

Compare two versions of an application service in classic Service Mapping

Modify tracking changes in configuration files

Fine-tune tracking changes for the change history

Compare two versions of an application service in classic Service Mapping

You can see a summary of application service changes at a glance by comparing two versions of an application service. This feature is useful for checking the application service status before and after a certain change or problem.

Before you begin

Role required: sn_cmdb_admin, admin, app_service_admin, app_service_user, service_mapping_admin, or service_mapping_user

About this task

Specify two points in time for which to compare the two versions of an application service. You can use the change indicators on the timeline to specify one point in time that is before and another that is after a change for which to see the details. For example, if you know that the application service started to fail at a certain time, you can compare two versions of the application service, one before and one after the problem started. This comparison lets you see the summary of changes that possibly led to the problems.

Service Mapping, if deployed, tracks and shows all changes to a CI including configuration files associated with a CI. When you compare two versions of an application service, you can see changes made to configuration files as changes to CIs. You can also compare two versions of a configuration file to see the actual changes in the files, during the time range specified for the comparison.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. If needed, click **Edit** to ensure that the map is in Edit mode.

   If Service Mapping is deployed, then in Edit mode, the Discovery Messages section appears below the map.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. On the history timeline, set the time range of changes that you want to view.

Option	Action
To set the time range of the history timeline	Click the hour, day, week, or month icons. 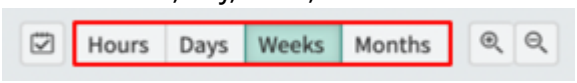
To increase or decrease the time range	Click the zoom in and zoom out icons. 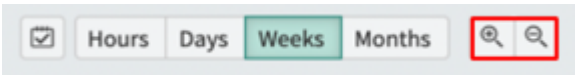
To change the upper limit on your history range	Click the history scale. 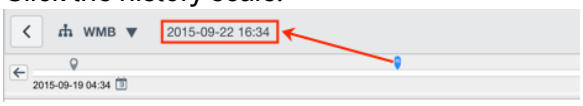 The time that serves as the upper limit ap pears above the history timeline. **Note:** You cannot set the lower limit on your history range to a time before this ser vice instance was created. This time is marked with the **IT Service Created** event on the history timeline. 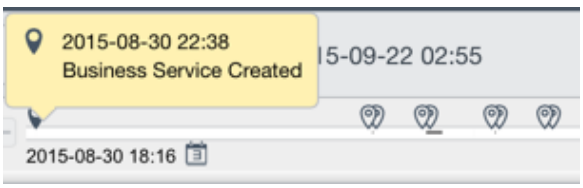

The map shows the history view of the service instance for the time you selected.

Note: The **Change** tab shows all change records, even the ones which are filtered out of the history view.

4. Click the **Compare** icon.

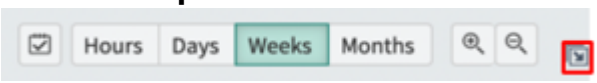

5. Set **Compare point 1** and **Compare point 2** as the two points in time for the comparison.

You can drag the pointers on the history scale to set corresponding time points.

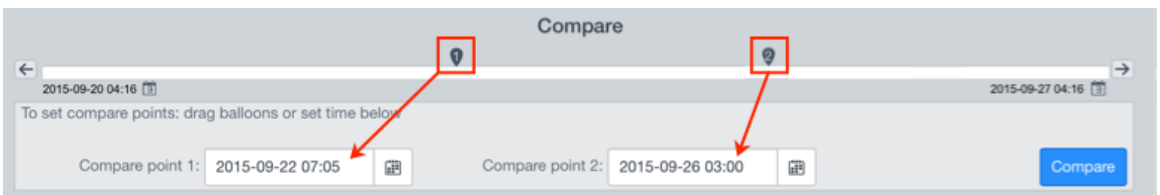

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the history scale does not include the time set for comparison, then its corresponding pointer appears next to the compare point in yellow:

Note: If there are no changes to the service during the time interval specified by **Compare point 1** and **Compare point 2**, then no change details are displayed.

6. Click **Compare**.
   The comparison view opens in a separate tab.
7. Select a marked CI to see the relevant change record on the **Changes** tab.

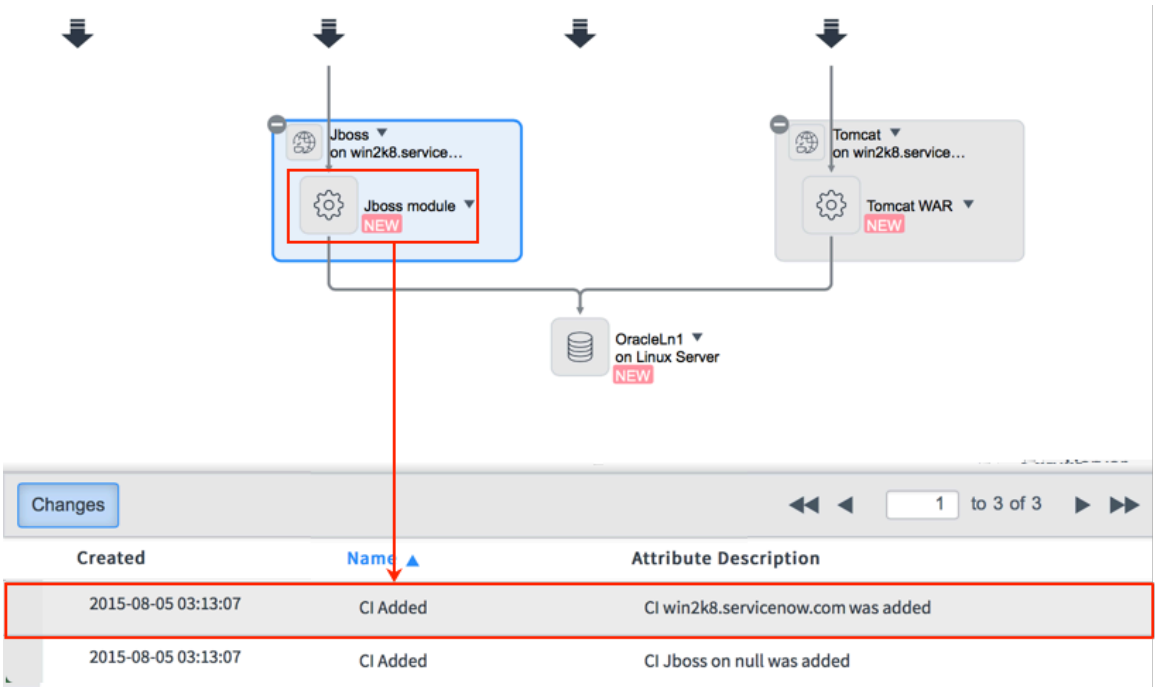

8. **Optional:** If Service Mapping is deployed, you can compare two versions of a configuration file that appears on the map as Updated:

   a. Select the CI that is associated with the updated configuration file.

   b. In the **Properties** pane, click the link to the updated file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

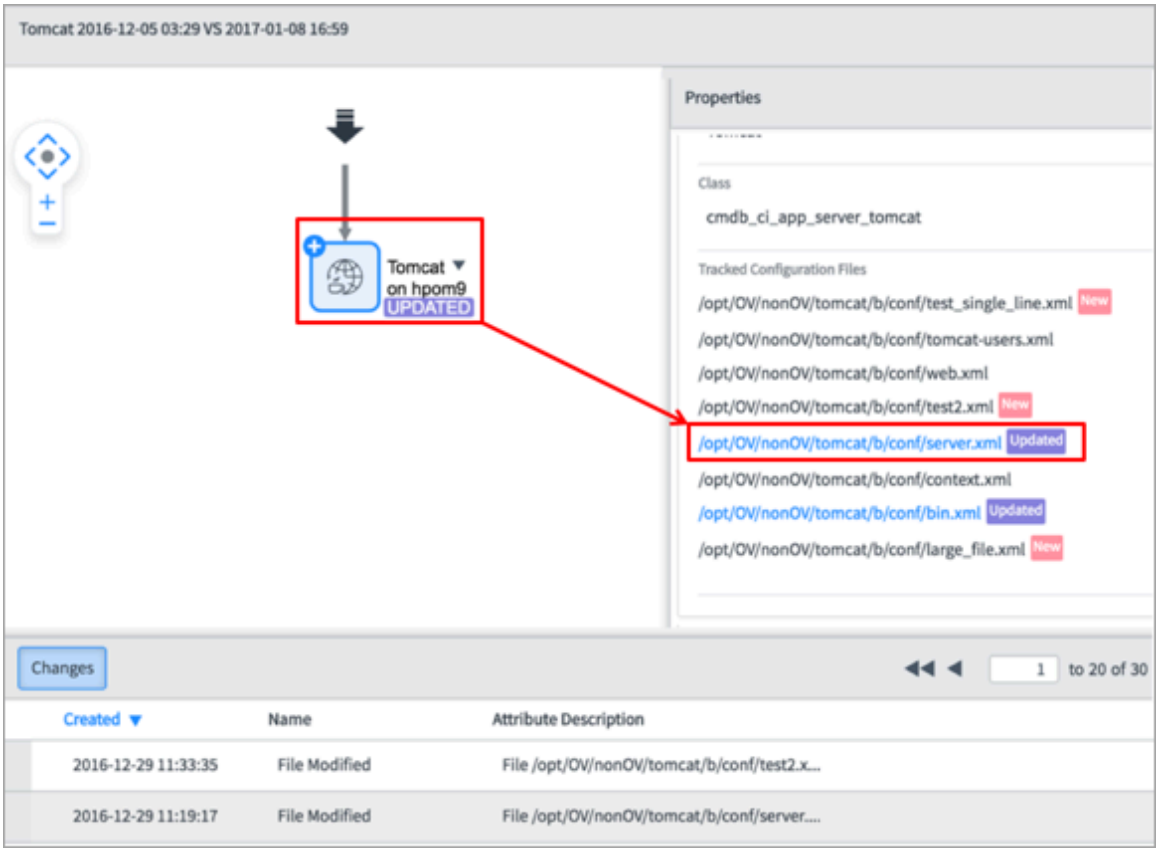

The **Tracked Configuration Files Version Compare** tab opens showing two versions of the configuration file side by side.

c. Review actual changes.

(Optional) Highlight colors indicate the type of change:

- Purple — Updated line
- Pink — New line
- Gray — Deleted line

d. Navigate between the changes using the Next difference icon and the Previous difference icon .

e. Close the **Tracked Configuration Files Version Compare** tab when finished.

9. Close the comparison view when finished.

Related topics

View the change history of application services in classic Service Mapping

Compare versions of CI configuration files

View dependent application services in classic Service Mapping

Check which application services depend on an application service and open maps for dependent application services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Roles required: service_mapping_user

About this task
The service that contains a reference to another service instance, becomes a dependent service. The service that you include as a reference is a contained service.

Dependent service map with a reference to a contained application service

Procedure
1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Right-click the entry point and select **Show dependent services**.
   The pop-up window displays the links to all dependent services for this application service.
3. Select the link for the dependent service whose map you want to view.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

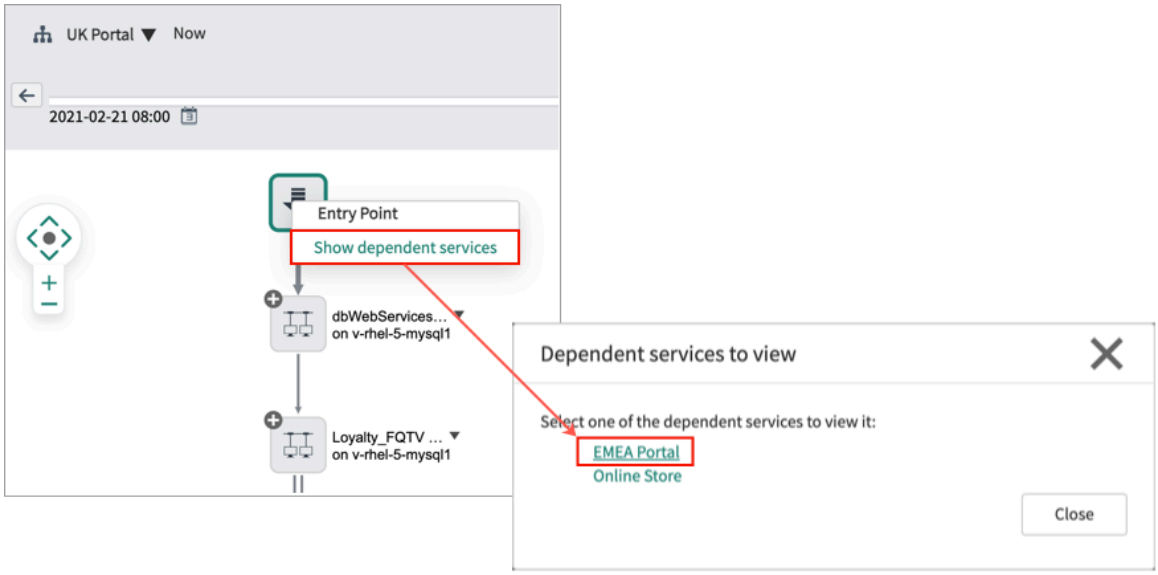

The map window shows the map for the dependent service you selected.

Related topics

Link application services

View contained application services in classic Service Mapping

A contained application service is a reference included into another, dependent application service. You can open the contained service map directly from within the dependent service map.

Before you begin

Roles required: service_mapping_user

About this task

The service that contains a reference to another service instance, becomes a dependent service. The service that you include as a reference is a contained service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dependent service map with a reference to a contained application service

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Right-click the CI that is the contained service, and select **Show contained service**.
   The map shows an automatic redirection to the web address for the contained service. No additional tabs need to open.

Check CI dependencies

You can see if a particular configuration item (CI) is part of other application services and check if it depends on other CIs.

Before you begin

Role required: admin, app_service_admin, app_service_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

While Service Mapping shows position of a CI in a particular application service, your organization may also use the same CI in other application services. If you are planning to perform changes to this CI, make sure that the changes do not affect other application services.

Procedure

1. In the map, right-click the relevant CI and select **Open dependency views**.
   The Dependency Views page opens in the new tab showing the selected CI in the center.
2. Click **Details**.

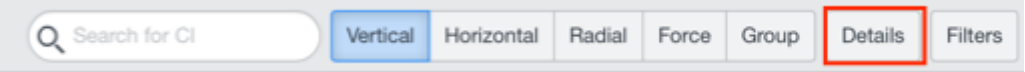

3. Click the **Related Services** tab.
   The list of links to application services containing this CI appear on the **Related Services** tab.
4. To view the map of another application service containing this CI, click the relevant link.

Related topics

Application service maps in classic Service Mapping

Dependency Views

View the network or storage path in classic Service Mapping

You can drill down to see objects behind a connection on a map for troubleshooting service performance or maintaining your network. This feature is not available for manually created services or for services on instances using Edge Encryption.

Before you begin

Role required: admin, service_mapping_admin or service_mapping_user

About this task

Displaying a network or storage path allows you to see routers, switches, and host ports which are part of the connection. The following connection types exist:

Network path

The infrastructure route between two CIs of the host type.

Storage path

The infrastructure route between a CI of the host type and a CI of the storage type.

Service Mapping discovers and displays network devices located on your organization local network, not on WAN or the public cloud. When Service Mapping encounters a connection between two applications, it uses traceroute on the source host to figure out the path to the target host. Then Service Mapping use SNMP credentials to discover the devices it found on the path.

Service Mapping can show devices for storage paths using the following protocols:

- Fiber Channel: Storage Area Network (SAN)
- IP Ethernet: Internet Small Computer Systems Interface (iSCSI) or Network Attached Storage (NAS)

Note: For storage paths using Fiber Channel, Service Mapping does not display Fiber Channel switches.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Between a host CI and a storage CI there may be two separate storage paths using different protocols: Fiber Channel and IP Ethernet. If so, the service instance map displays one connection line for both protocols.

You can see the change history for network and storage paths for application services. Checking path history for manual or technical services is not supported.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. On the map, right-click the connection for the network or storage path you want to see.
3. Select **Show network path** or **Show storage path**.
4. If necessary, select the type of protocol and click **Choose** in the Choose Protocol Type dialog box.
   The view showing the network or storage path opens in a new tab.

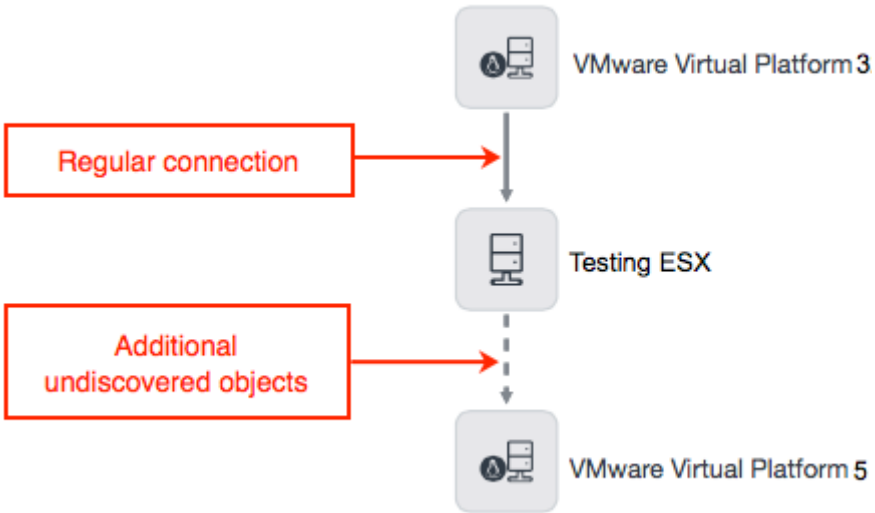

   If Service Mapping cannot discover all objects in a connection, the connector line on the network or storage path appears dotted as shown above.

5. For storage paths, use the Storage Mapping tab at the bottom of the screen to view correlation between file systems on the host and the storage volume or the shared folder on the storage device.
6. To view the change history for network paths, perform steps in View the change history of application services in classic Service Mapping.
7. If you access the network or storage path map from Event Management, you can see the list of alerts related to this path appear at the bottom of the screen.

Related topics

Application service maps in classic Service Mapping

SNMP-based queries

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View the list of CIs belonging to an application service

View a complete list of CIs that make up an application service, including CIs not shown on the application service map: tracked configuration files, endpoints, processes, and network devices.

Before you begin

Role required: service_mapping_admin or service_mapping_user

About this task

By default, application service maps show application CIs and connections between them. The rest of the CIs behind applications and connections between them are hidden. The CI list displays the following additional CIs:

- Servers that host applications.
- Network devices like routers or switches.
- Processes running on hosting servers. Processes are used for identification and classification of applications.
- Tracked configuration files for applications that are part of the service.
- Endpoints belonging to the applications. Endpoints contain information necessary for connecting to applications, like hosts and ports.
- The application service CI for this application service.

The CI lists for empty application services, that do not contain any CIs contain only one CI: For the application service itself.

Procedure

1. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
2. Navigate to **Application Services** > **View:csdm_view** > **View** > **Default view**.

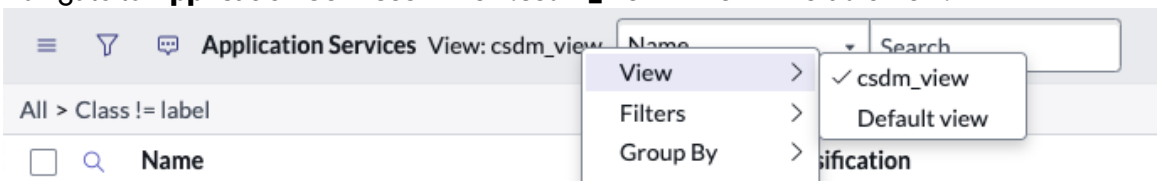

3. Click the name of the application service for which you want to view the list of CIs.
   The discovery state of the selected service must be **Done**.
4. Click **List CIs** under Related Links.

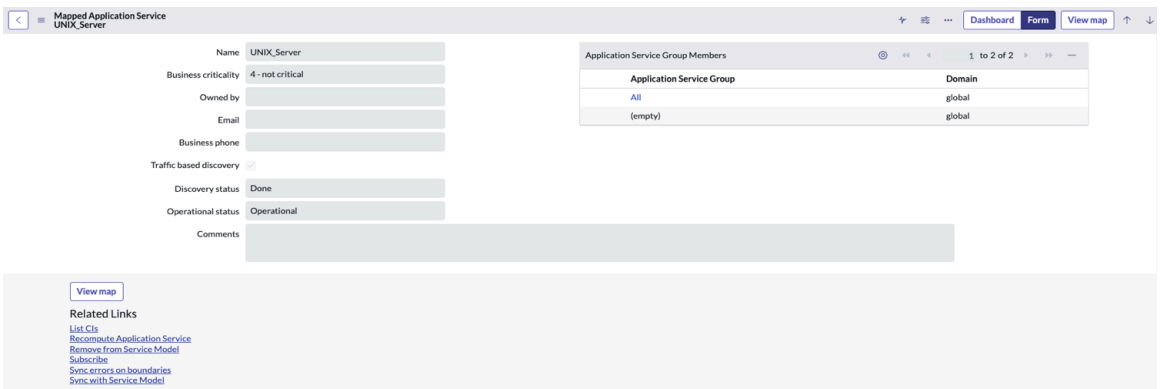

5. View the complete list of CIs that are part of the selected application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Click the back button on your browser to return to the Mapped Application Service page when you have finished reviewing the list of CIs.
7. Select **View map** to navigate to the application service map.

Convert application service maps into PDFs

You may want to create a PDF for a map to share information about the service content with other people. PDFs are especially useful during the review and approve process for application services discovered by Service Mapping.

Before you begin

Role required: app_service_admin, app_service_user

About this task

The PDF of an application service map shows the name of the application service and the exact time when this PDF was created.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Select the **Additional actions** menu.

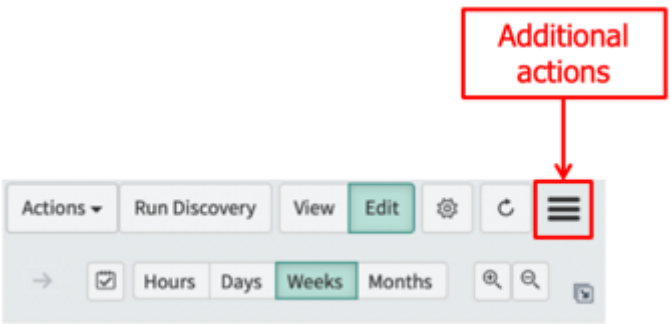

3. Click the button next to the **Export to PDF**option.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

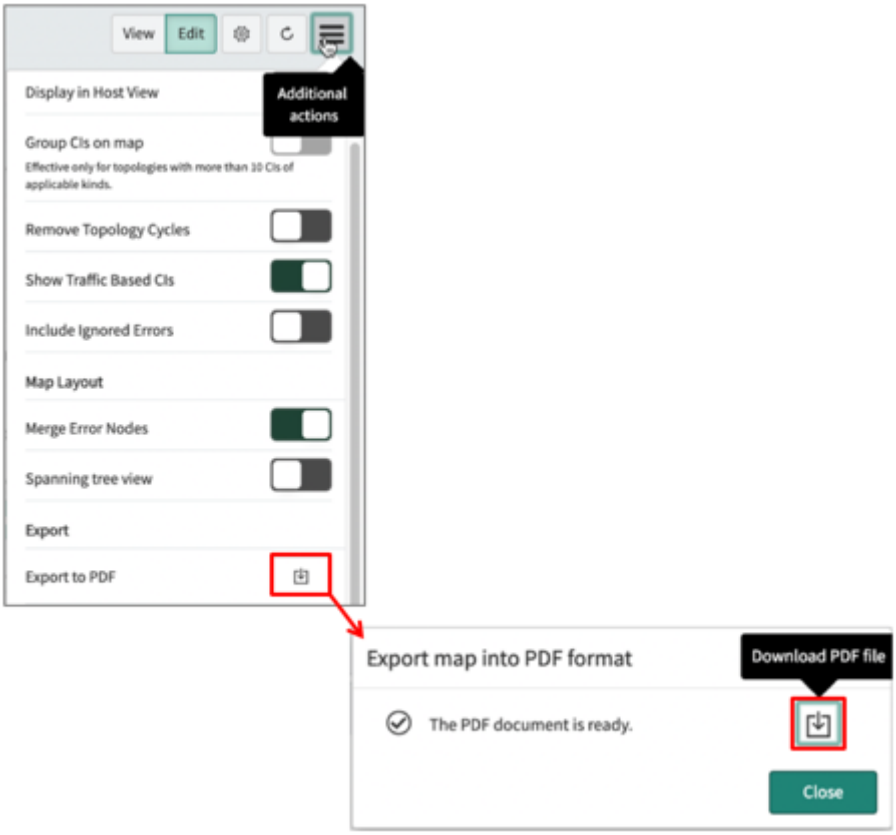

4. When the PDF is ready, click the **Download PDF** button in the Export Map into PDF Format dialog box.
   The PDF appears in the downloads folder on your computer.

Modify view for an application service map

Customize a map view to display relevant information and increase clarity.

Before you begin

Role required: app_service_admin, app_service_user, service_mapping_admin, or service_mapping_user

About this task

By default, service maps are concise and contain every detail about all CIs that you added manually or that Service Mapping mapped. While it provides a complete picture, using such a map might be difficult. Customizing maps allows you to:

- make it easier to navigate the map.
- hide irrelevant information.
- show essential information.

Attributes you show or hide when you customize maps are not removed permanently and can be displayed or hidden again as you need. The system saves the changes to the map view and displays the last map view when you reopen a service instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open the service instance map.
   a. Navigate to **All** > **CSDM** > **Manage Technology Management Services** > **Service Instance**.
   b. Select the needed service instance.
   c. On the service instance page, select **View Map**.
2. Select the **Additional actions** menu.

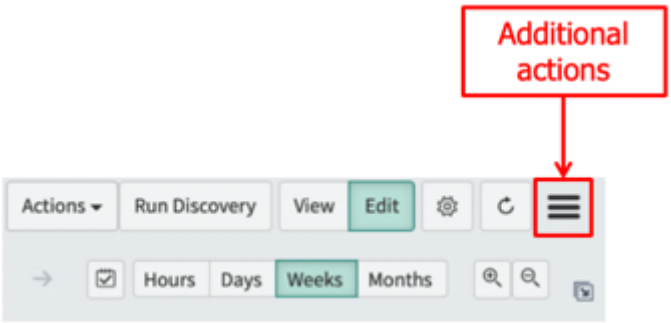

3. Select the relevant option:

Option	When Enabled	When Disabled
Display in Host View This option is available only when the Service Mapping plugin is activated. Switch between the CI view that hides hosts and the host view that hides CIs. The host view can help determine that Service Mapping discovered and mapped all hosts correctly as well as help locate connections to each host. The host view also simplifies the map especially when many applications are running on a few hosts. The Host View is useful only for service instances discovered using Service Mapping. This view does not offer additional information for manually created service instances.	The service map shows host CIs. 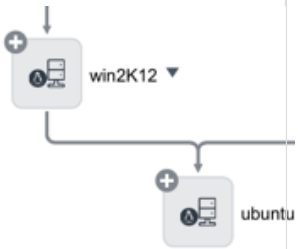 Click the plus icons to see application CIs running on the hosts. 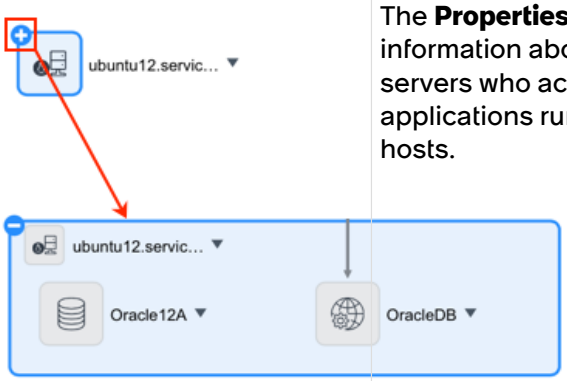	The service map shows application CIs. 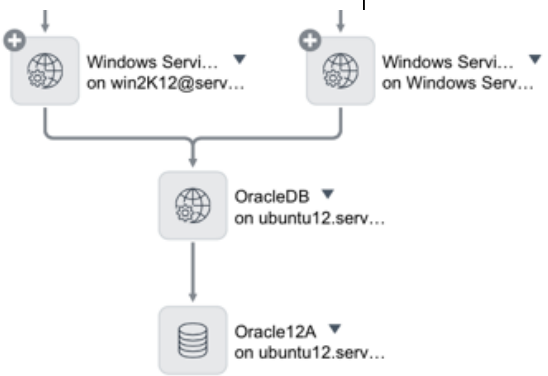 The **Properties** pane shows information about both servers who act as hosts and applications running on the hosts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	When Enabled	When Disabled
	The **Properties** pane shows only information about servers who act as hosts.	
Group CIs on map Simplify maps by grouping CIs belonging to the same type and hosted on servers sharing a prefix and domain name. The map still reflects the application flow correctly, because grouped CIs belong to the same tier. This option is relevant for service instances consisting of more than 10 CIs.	The service instance map shows groups of CIs. The map is more compact.	The service instance map shows all CIs separately. The map is larger.
Remove Topology Cycles This option is available only for users with these roles: Service Mapping administrator [service_mapping_admin] and Service Mapping user [service_mapping_user]. Hide connections returning to the same configuration item (CI) that originated them to reduce map cluttering. Usually these returning connections are not important for analysis or troubleshooting of service instances.	The service instance map hides topology cycles.	The service instance map shows topology cycles. 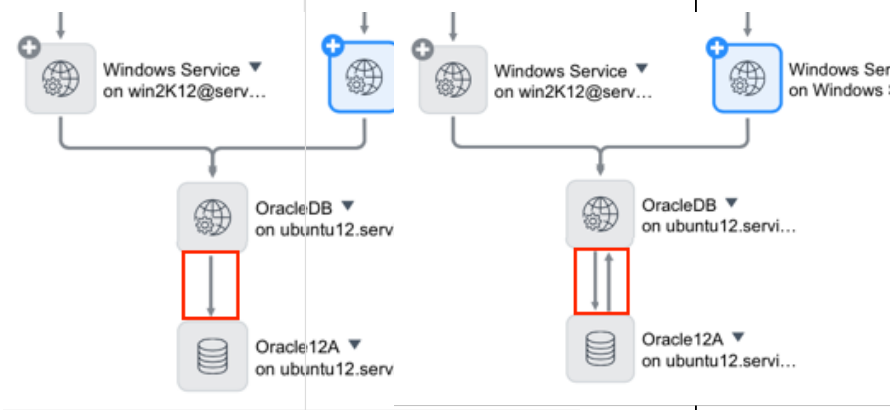
Show Traffic Based CIs This option is available only for users with these roles: Service Mapping administrator [service_mapping_admin] and Service Mapping user [service_mapping_user]. Display configuration items (CIs) found using traffic-based mapping.	The service instance map displays CIs and hosts that Service Mapping discovered using both patterns and traffic-based discovery. While detecting CIs inbound and outbound traffic creates an inclusive map, it may also result in mapping many redundant CIs that do not influence the service instance operation.	The service instance map shows only CIs and hosts discovered using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	When Enabled	When Disabled
Merge error nodes This option is available only for users with the Service Mapping administrator [service_mapping_admin] role. Merge CIs with errors caused by the same source CI to declutter a service instance map. A discovery or operational error in a CI is often propagated to all its connections down the hierarchy crowding your view. You can remove identical errors caused by the same CI by consolidating them in one object on the map. For example, if one CI causes an error in six CIs that are connected to it, you can show only one CI with this error on the map. The connector line shows the number of actual connections.	The service instance map merges all CIs with identical errors into one. The number on the connector line indicates the number of CIs with this error. 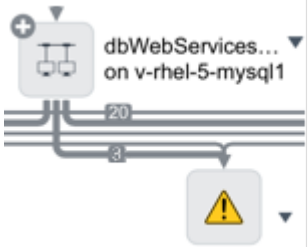	The service instance map shows all CIs with identical errors separately. 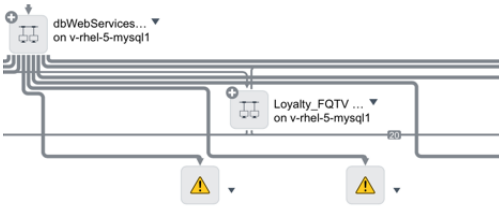
Map indicators This option is available only for users with these roles: Service user [app_service_user] and Service Mapping user [service_mapping_user]. Display additional information for a CI or service instance itself by displaying related records such as alerts, outages, incidents, and problems. The base system provides map indicators for the following record types: ○ Affected CIs ○ Open alerts ○ Current, planned, or recent change requests ○ Open incidents	For each record type that is set to display, the corresponding indicator and the tab display. For example, the **Incidents** tab appears underneath the map and the **Incidents** indicator shows next to CIs that have incidents associated with them. 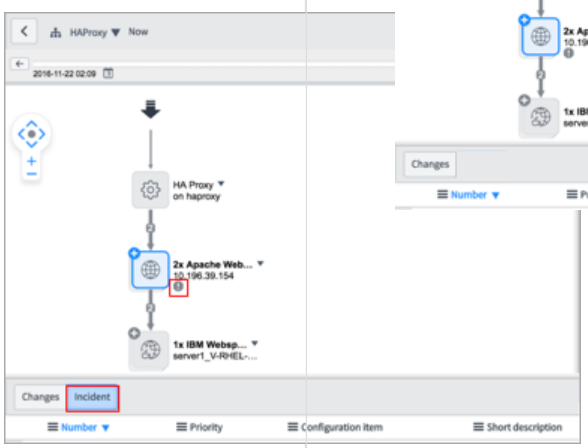	The service instance map shows CIs without related indicators. 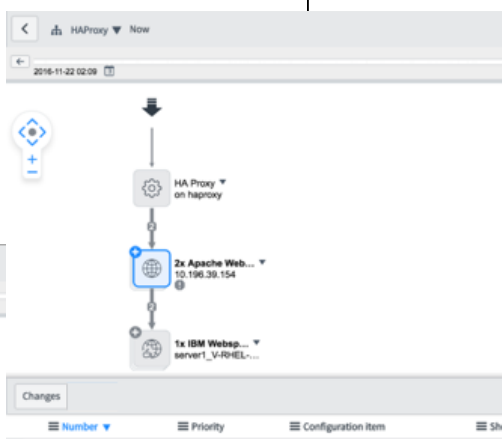

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	When Enabled	When Disabled
○ Current, planned, or recent outages ○ Open problems For more information, see Event Management Map Indicators (Video).	If a record is associated with the service instance itself, the indicator shows next to the service instance name. 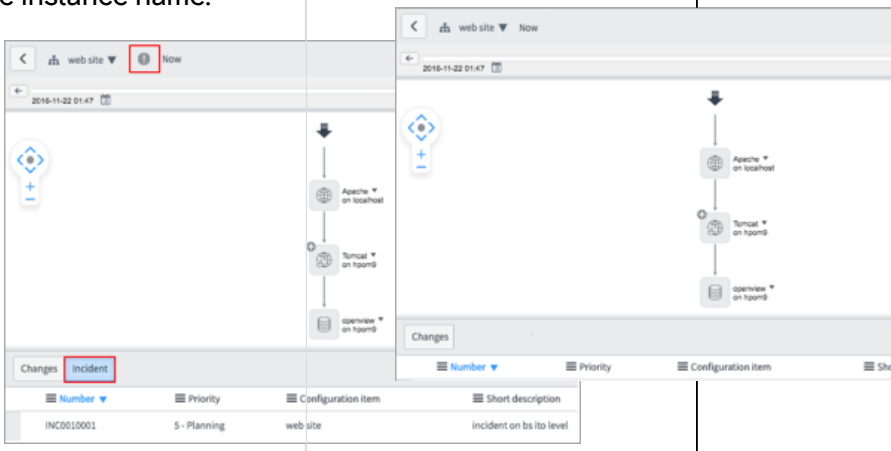	No map indicators appear next to the service instance name.
Spanning tree view Simplify the map by organizing CIs into a tree structure and concealing some connection lines. This option is especially useful for very large maps.	The map displays a service instance as a tree. 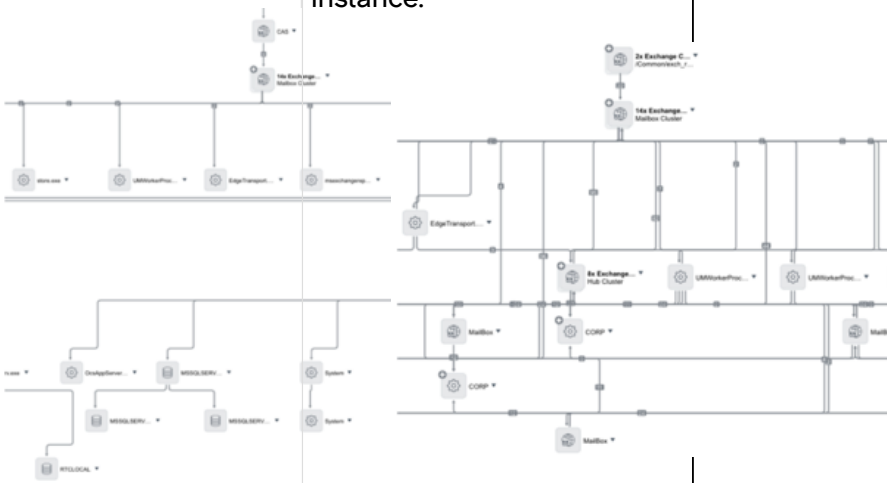 To view actual connections, point to a CI. For more information on viewing connections, see View CI connection attributes in Service Mapping.	The map reflects the actual structure of a service instance.
Include ignored errors This option is available only for users with the Service Mapping administrator	The map shows all errors: both active and ignored.	The map does not show ignored errors.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	When Enabled	When Disabled
[service_mapping_admin] role. You can display errors which were previously hidden from the map.		

Application service analysis and maintenance using legacy Agent Workspace

Service Mapping creates maps to help you see the architecture and organization of application services. These maps are useful for planning change or migration, as well as analyzing the continuity and availability of services.

Application service maps in legacy Agent Workspace

Maps offer you a visualization of data on configuration items (CIs) comprising application services, and relations and connections between these CIs.

A new map is created every time you manually create an service instance or map an service instance using Service Mapping. The system updates the map every time you add CIs to the service instance or when Service Mapping runs discovery for CIs belonging to this service instance.

Map window

Aside from the map itself, the Service Map may also display the basic service information like the severity, status, and application owner of the service. You can also view additional service information on the Service properties pane. Click the Related items icon to display all alerts or changes associated with CIs comprising this application service or the application service itself.

The service instance map in the Event Management view reflects the severity status of CIs belonging to the service instance.

Service Mapping

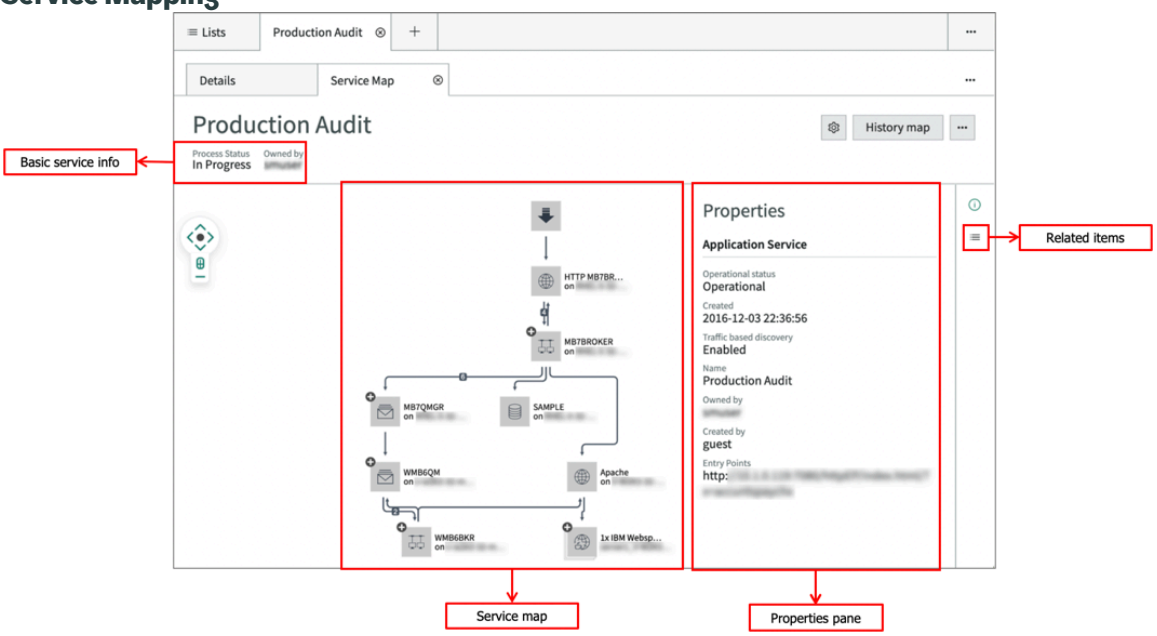

By default, the entire map is shown in the center of the visible map area of the window. Position the map and zoom in and out using these controls.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Control zoom and positioning of your map

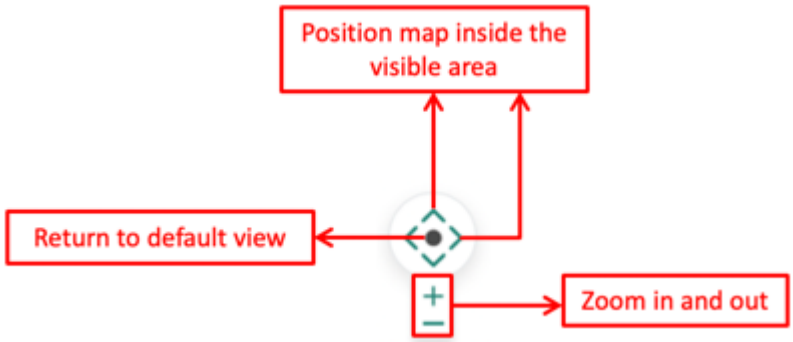

You can also click anywhere in the map area and drag the required segment of the map into the visible area.

CI attributes and CI-related messages

When a map is loaded and no elements are selected, the Service properties pane shows the details of this service instance.

A selected device, application, or connector appears in green and is highlighted. Information about the selected map element is displayed in the Properties pane on the right of the map.

The Properties pane shows attributes for connections and CIs

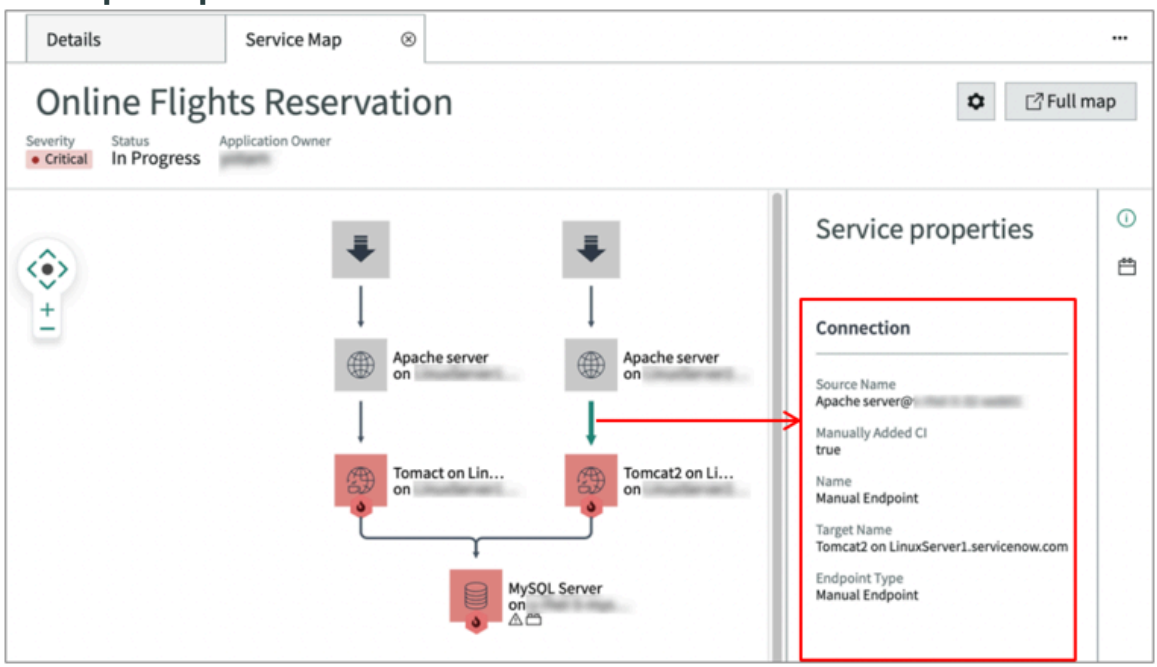

The attributes for the server hosting applications and applications themselves are shown separately inside the Properties pane.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Properties pane shows attributes for hosts and applications separately

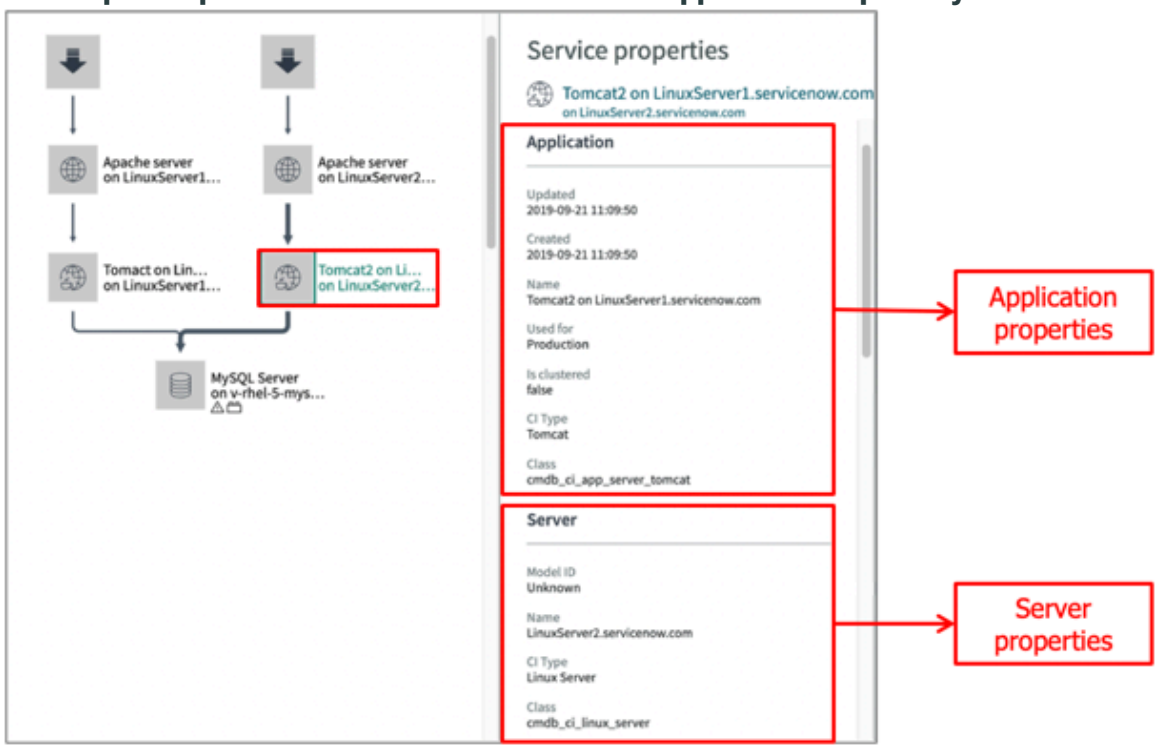

Related Items

Click the **Related Items** icon to display alerts or changes associated with this application service or its CIs.

If there is any information related to the selected CI or connection, it is highlighted in the Related Items pane on the right of the window.

Select a CI to display information related to it

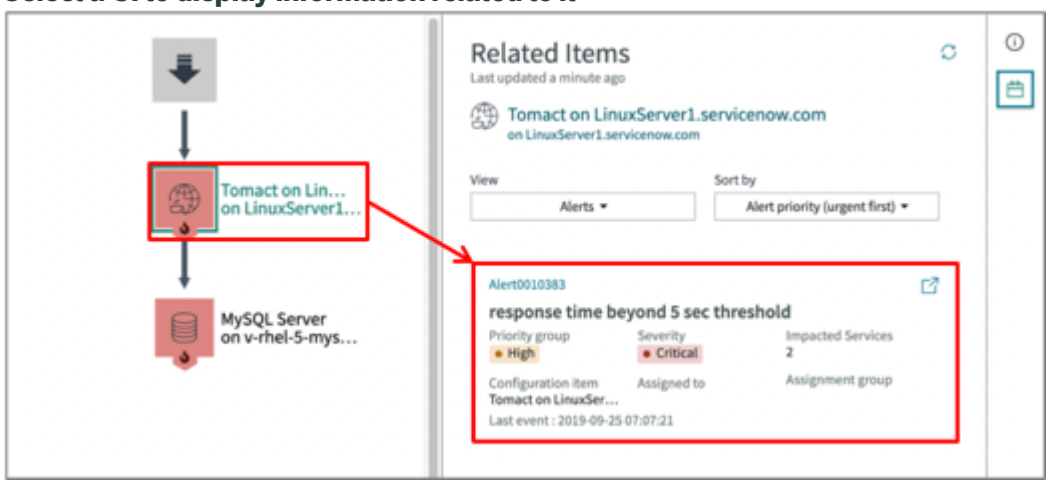

When you select information in the Related Items pane on the right of the map window, the related CI displays in yellow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Selecting message marks the related CI yellow

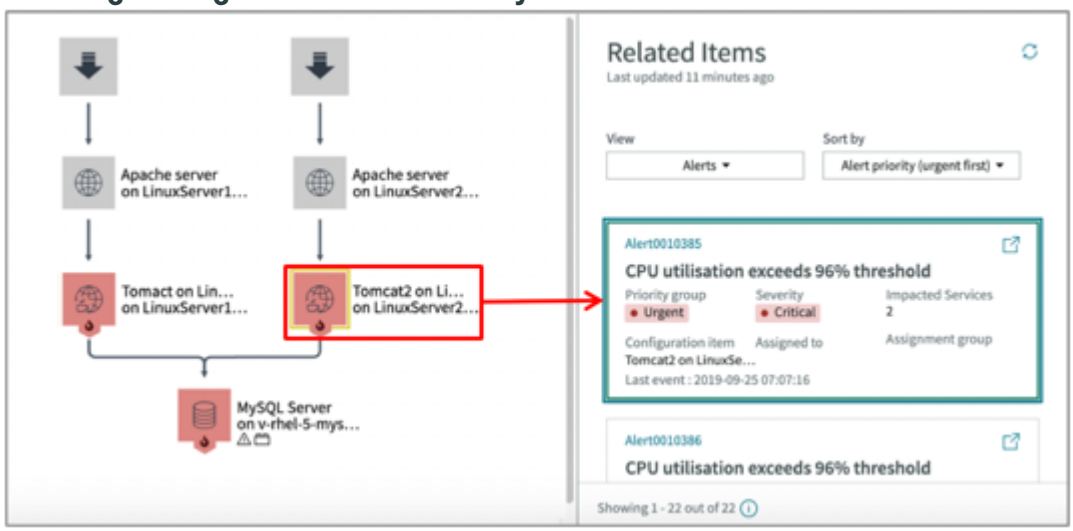

You can display other related items, such as planned changes or cases, associated with the application service. For more information, see Modify view for an application service map in legacy Agent Workspace.

Related topics

View an alert impact on CIs in a service map

View maps for application services in legacy Agent Workspace

Use maps to see the details of application services. Maps show configuration items (CIs) making up application services, connections between these CIs and other service-related details.

Before you begin

Role required: sm_app_owner, service_mapping_user, evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

The service instance map in the Event Management view reflects the severity status of CIs belonging to the service instance. For more information about viewing service details on maps, see Application service maps in legacy Agent Workspace.

Procedure

1. In Agent Workspace, click the **Lists** icon on the left pane.
2. Navigate to **All** > **Services** > **All Services**.
   The list of services displays services in Operational status.
3. Click the name of the relevant service instance.
   The service tab opens showing service details on the Details tab.
4. Click **Service Map**.

   The Service Map tab opens displaying the map of this application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

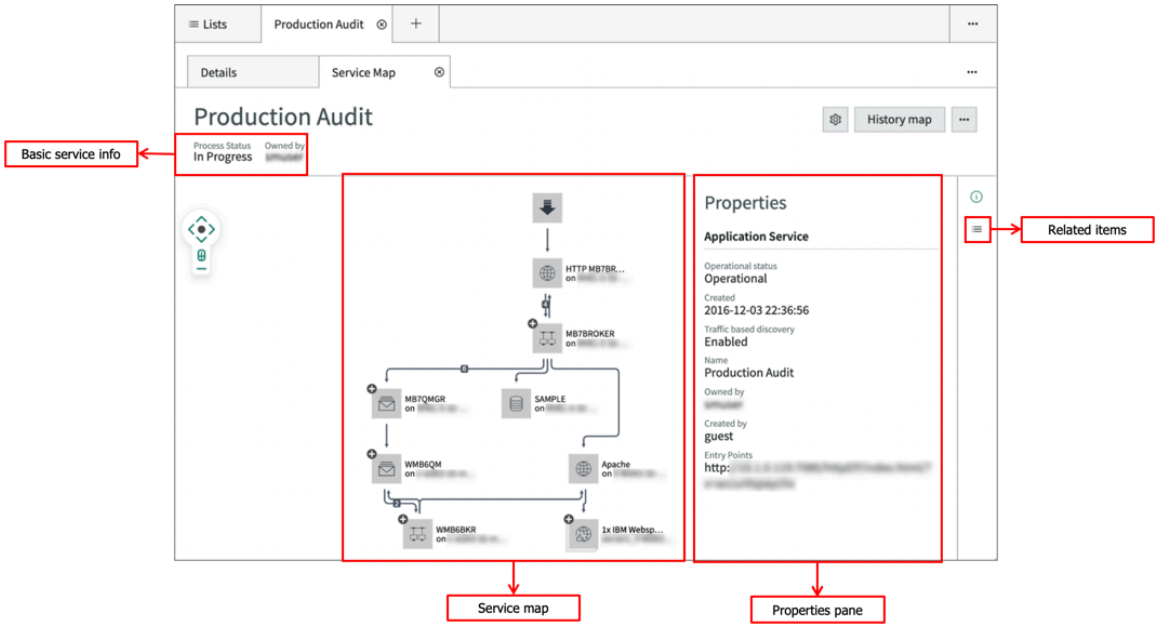

5. Click on CIs or their connections to display their details in the Service properties pane on the right.
6. Display additional information for a CI or application service itself by displaying related records such as alerts, outages, incidents, and problems.
   For more information, see Modify view for an application service map in legacy Agent Workspace.
7. To access the service map in the standard Event Management UI that provides full map-related functionality, click **Full map**.

Related topics

Service Mapping

View the change history of application services in legacy Agent Workspace

View changes made to a service instance and to the individual configuration items (CIs) comprising the service in Agent Workspace. You can also compare the service status before and after a certain change or problem. The change history is useful for maintenance, planning, or troubleshooting procedures.

Before you begin

Role required: admin, service_mapping_admin, service_mapping_user, app_service_admin, or app_service_user

About this task

Details about changes to a service instance and to its CIs are stored in the CMDB. Typically, these changes reflect adding or removing CIs from a service instance, upgrading or updating CIs, or modifying CI configuration files. The system gathers this data by querying CMDB tables and then creating the change history view. In deployments where Service Mapping is activated, the type of change information Service Mapping queries depends on discovery patterns that Service Mapping uses to discover CIs.

Changes to the application service appear on the history timeline.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The history scale shows changes to the application service

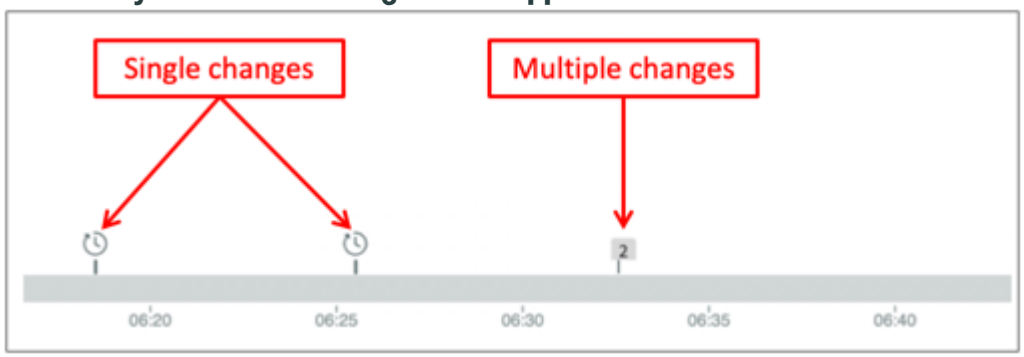

Specify two points in time for which to compare the two versions of an application service. You can use the change marks on the timeline to specify points in time that are before and after a change for which you want to see details. For example, if you know that the application service started to fail at a certain time, you can compare two versions of the application service, one before and one after the problem started. This comparison lets you see the summary of changes that possibly led to the problems.

The default timeline range is the recent week. The upper limit of the time range is the moment when you opened the map in the Change History view. If the application service was created less than a week ago, the time range starts from the time of service creation.

The time range cannot contain more than 1 thousand changes. If there are more changes within the selected range, the history scale automatically adjusts to display 1 thousand changes.

To view validation status of changes and changes to related items, like tracked configuration files, use the classic Service Mapping UI. For more information, see View the change history of application services in classic Service Mapping and Compare two versions of an application service in classic Service Mapping.

Procedure

1. In Agent Workspace, click the **Lists** icon on the left pane.
2. Navigate to **All** > **Services** > **All Services**.
   The list of services displays services in Operational status.
3. Click the name of the relevant service instance.
   The service tab opens showing service details on the Details tab.
4. Click **Service Map**.

   The Service Map tab opens displaying the map of this application service.
5. Click **Change History Map**.
   The History Map tab displays the map in the history view. By default, the time range is a week.
6. Compare the application service before and after a change, by performing the following steps:

   a. Change the time range on the timeline by clicking the **Zoom in** or **Zoom out** buttons next to the timeline.
   Alternatively, point at the timeline and scroll.

   The map shows changes to CIs. CIs are marked as added, updated, or deleted.

   The **Properties History** pane shows the comparison summary of changes to this application service. Click **Change history** to view change record cards for this application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

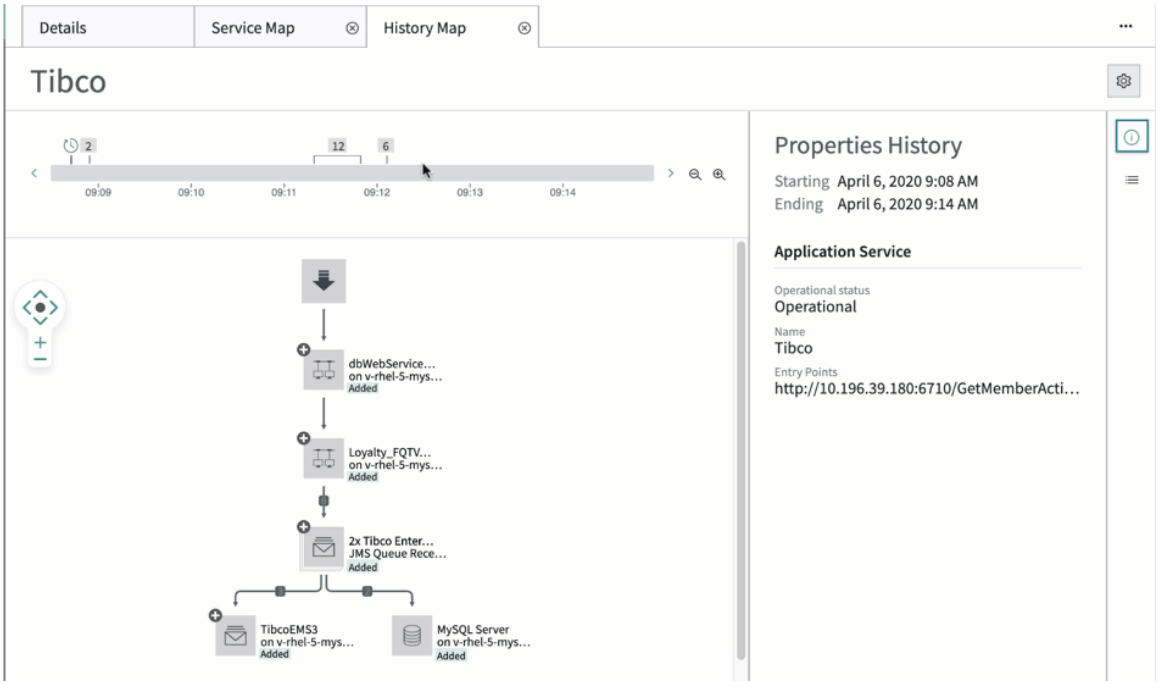

b. Focus on changes on the map using one of the following methods:

- To highlight the CI responsible for the change record and to show the change record card, click the change mark on the timeline. You can highlight multiple CIs by clicking the mark for multiple changes.

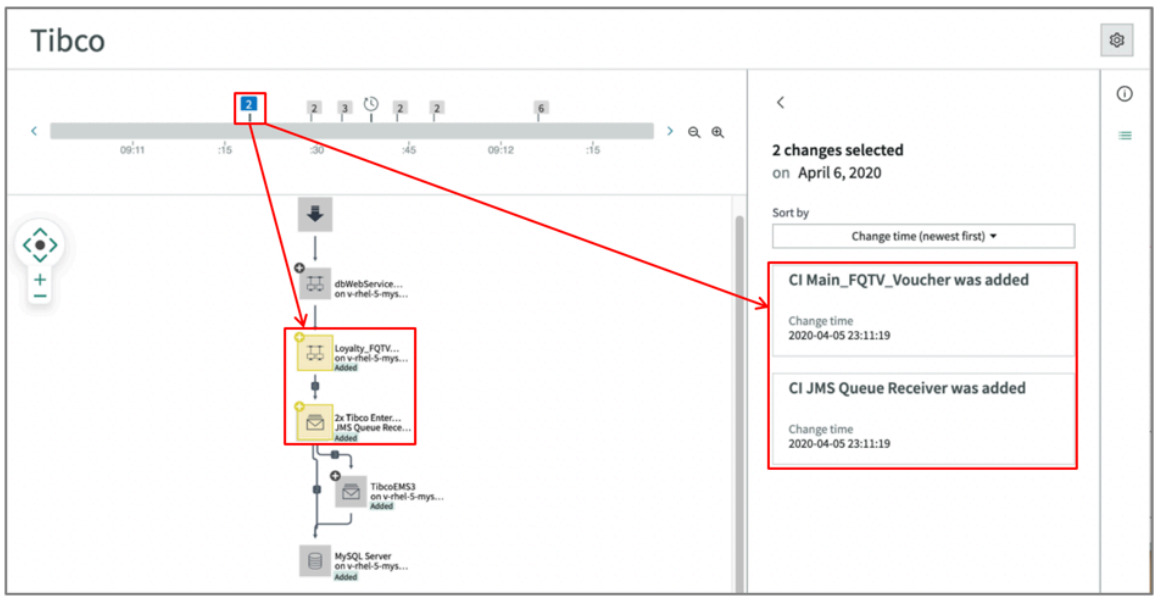

- To highlight the CI responsible for the change in the map, click the change record card in the Change history pane.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

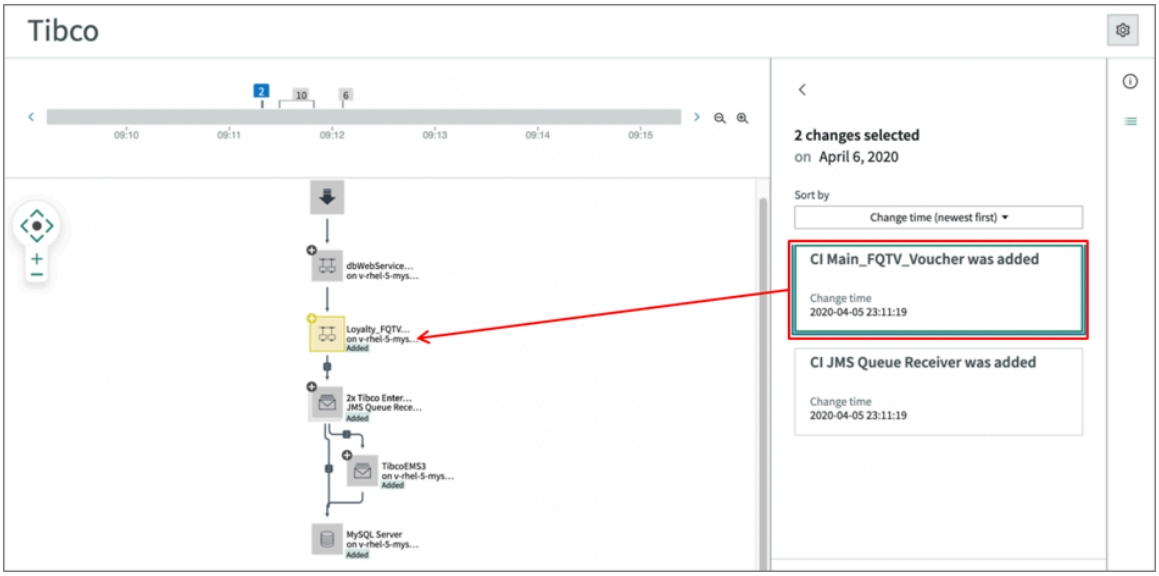

- To display the related change record card in the Change history pane, click the CI on the map.

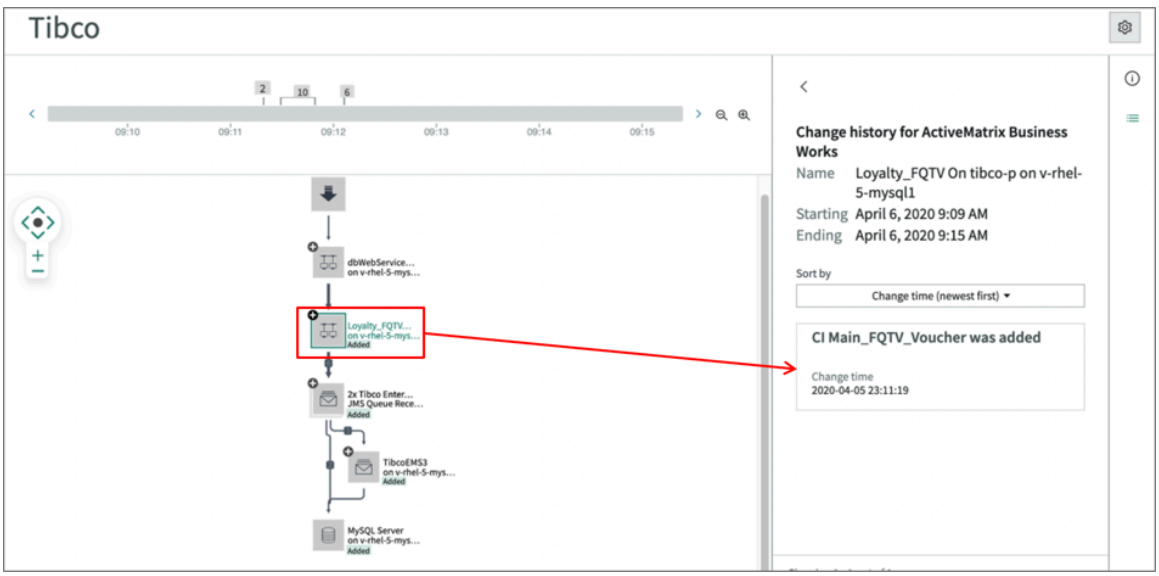

Modify view for an application service map in legacy Agent Workspace

Customize a map view to display relevant information and increase clarity.

Before you begin

Role required: service_mapping_user, sm_app_owner, evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

By default, service maps are concise and contain every detail about all CIs that you added manually or that Service Mapping mapped. While it provides a complete picture, using such a map might be difficult. Customizing maps allows you to:

- make it easier to navigate the map.
- hide irrelevant information.
- show essential information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attributes you show or hide when you customize maps are not removed permanently and can be displayed or hidden again as you need. The system saves the changes to the map view and displays the last map view when you reopen a service instance.

The service instance map in the Event Management view reflects the severity status of CIs belonging to the service instance.

Procedure

1. In Agent Workspace, click the **Lists** icon on the left pane.
2. Navigate to **Services** > **All Services**.
   The list of services displays services in Operational status.
3. Click the name of the relevant service instance.
   The service tab opens showing service details on the Details tab.
4. Click **Service Map**.
5. Click the **Open Map Display Settings** button.

6. Select the relevant option:

Option	When Enabled	When Disabled
Display in Host View The host view also simplifies the map especially when many applications are running on a few hosts. The Host View is useful only for service instances discovered using Service Mapping. This view does not offer additional information for manually created service instances.	The service map shows host CIs.	The service map shows application CIs.
Group CIs on map Simplify maps by grouping CIs belonging to the same type and hosted on servers sharing a prefix and domain name. The map still reflects the application flow correctly,	The service instance map shows groups of CIs. The map is more compact.	The service instance map shows all CIs separately. The map is larger.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	When Enabled	When Disabled
because grouped CIs belong to the same tier. This option is relevant for service instances consisting of more than 10 CIs.		
Remove Topology Cycles Hide connections returning to the same configuration item (CI) that originated them to reduce map cluttering. Usually these returning connections are not important for analysis or troubleshooting of service instances.	The service instance map hides topology cycles.	The service instance map shows topology cycles.
Spanning tree view Simplify the map by organizing CIs into a tree structure and concealing some connection lines. This option is especially useful for very large maps.	The map displays a service instance as a tree. 	The map reflects the actual structure of a service instance. 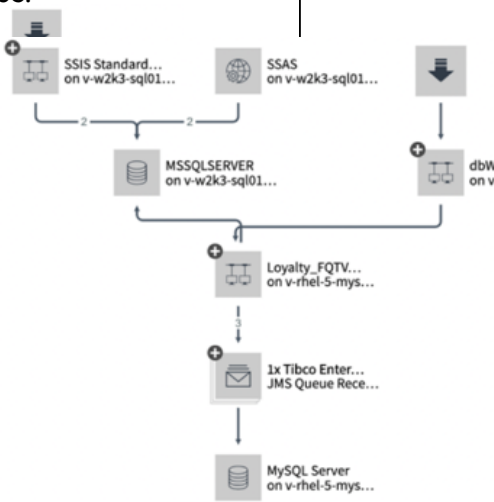
Show Traffic Based CIs This option is available only for users with these roles: Service Mapping administrator [service_mapping_admin] and Service Mapping user [service_mapping_user]. Display configuration items (CIs) found using traffic-based mapping.	The service instance map displays CIs and hosts that Service Mapping discovered using both patterns and traffic-based discovery. While detecting CIs inbound and outbound traffic creates an inclusive map, it may also result in mapping many redundant CIs that do not influence the service instance operation.	The service instance map shows only CIs and hosts discovered using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	When Enabled	When Disabled
Map indicators Display additional information for a CI or service instance itself by displaying related records such as alerts, outages, incidents, and problems. The base system provides map indicators for the following record types: ○ Affected CIs ○ Changes: Current, past, or planned ○ Outages: Current, past, or planned ○ Problem For more information, see Event Management Map Indicators (Video).	For each record type that is set to display, the corresponding indicator and the tab display. For example, the **Incidents** tab appears underneath the map and the **Incidents** indicator shows next to CIs that have incidents associated with them.	The service instance map shows CIs without related indicators.
	If a record is associated with the application service itself, the indicator shows next to the application service name.	No map indicators appear next to the application service name.

Service Mapping reference

Reference topics provide additional information about mapping and fine-tuning application services using Service Mapping lists and forms.

Service Mapping Workspace

The Service Mapping workspace provides a central location to streamline the process of mapping your application services. Use the visualizations and reports to analyze, monitor, and update your resources, create application services, and manage your service mapping tasks efficiently.

Access

Navigate to **Workspaces** > **Service Mapping**.

Required roles

Roles required for the Service Mapping Workspace

Role	Description
service_mapping_admin	Required to access the Service Mapping Workspace.

Key features and capabilities

Machine Learning (ML) readiness

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Confirm that your Machine Learning environment is prepared before you start to map application services.

Detailed service maps

Drill down for detailed information on service maps across various categories.

Automated Service Suggestions

Leverage Automated Service Suggestions to evaluate suggested application service candidates, using them to create application services or enhance existing ones.

Improved resource utilization

Use unmapped servers linked to an application service candidate to create application services.

Tag-based service mapping

Use the Tag-based mapping dashboard to map your organization's configuration items (CIs) into application services.

Create a service instance or a unified service

- Create a service instance to leverage your resources in unmapped services.
- Create a unified service from various sources. For more information, see: Multi-source service mapping.

Important:

Access to tag-based service mapping in the Service Mapping workspace requires the installation of Service Mapping Plus version 1.16.3. For more information, see Install Service Mapping Plus.

Key metrics

Mapped application services

The number of application services previously mapped. Select to view specific information about these services such as operational status and business criticality.

ML-powered candidates

The number of application service candidates identified using the Automated Service Suggestions feature.

Note: Starting with Service Mapping Plus version 1.16.3, the filtering process has been updated to include only the most relevant candidates, excluding any ignored or irrelevant ones. This enhancement requires the installation of ITOM Content Service.

Mapped servers

The number of servers currently linked to an application service. Select the tile to view information such as operational status and owner. The **Mapped servers** tile is replaced with ServiceNow Otto metrics, once ServiceNow Otto is installed.

Maps created by ServiceNow Otto

The number of service maps created by ServiceNow Otto. Select the tile to view the Service Mapping AI Agent activity. This tile is available only when ServiceNow Otto is installed.

Business app linked to service instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The number of business applications with a CSDM relationship linked to a service instance. Select to view the Business App Map Agent activityThis tile is available only when ServiceNow Otto is installed.

Service Mapping AI Agents

When ServiceNow Otto is installed, the **Service Mapping home** page displays the AI Agent activation panel. Use this panel to activate or deactivate AI agents that automate service mapping tasks. For more information, see AI Agents for Service Mapping.

Reports

Select the relevant section of each report to view more detailed information.

Title	Type	Source table	Description
Service maps by operation type	Bar report	Mapped Application Service [cmdb_ci_service_discovered]	A bar report that presents the mapped application services grouped by operation type.
Service maps by creation method	Pie report	Mapped Application Service [cmdb_ci_service_discovered]	A pie report that presents the mapped application services grouped by creation method.
Service maps by criticality	Bar report	Mapped Application Service [cmdb_ci_service_discovered]	A bar report that presents the mapped application services grouped by the level of criticality.
Unmapped servers **Note:** This view is applicable to users with Service Mapping Plus versions earlier than 1.16.3.	Bar report	Server [cmdb_ci_server] Service Configuration Item Association [svc_ci_assoc] Application service candidate [sa_ml_application_service_candidate]	A bar report that presents unmapped servers that have been matched to an application service candidate as well as those servers with no candidate. Select **Unmapped servers with candidate** to view servers you can map to an application service. For more information, see Create an application

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
			service for unmapped servers.
Unmapped servers **Note:** This view is available starting with Service Mapping Plus version 1.16.3.	Bar report 	Server [cmdb_ci_server] Service Configuration Item Association [svc_ci_assoc] Application service candidate [sa_ml_application_service_candidate]	A bar report that presents the number of unmapped servers, divided into two categories: • **Unmapped servers with candidate** - servers that have been matched to an application service candidate. • **Unmapped servers without relevant candidate** - servers that have no candidate. **Note:** Starting with Service Mapping Plus version 1.16.3, the filtering process has been updated to include only the most relevant candidates, excluding any ignored or irrelevant ones. This enhancement requires the installation of ITOM Content Service.

Additional components

Refresh icon

You can select this button to refresh the page and update the displayed information to reflect the most recent data.

Search for server

You can use this field to locate a specific server and review it's detailed information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag-based service mapping

Select the Tag-based service mapping icon to access the Tag-based service mapping dashboard. Here you can categorize and map your organizaton's on-prem and cloud resources into application services. For more information, see Tag-based mapping in the Service Mapping Workspace.

Application service readiness dashboard

Select the Application service readiness icon to access the Application service readiness dashboard.

Use this dashboard to ensure you have enabled Predictive Intelligence, verify Machine Learning readiness, and confirm that your data has been trained.

For more information about application readiness, see Application service readiness dashboard in configurable workspace.

Related topics

- Automated Service Suggestions
- Map application services based on Automated Service Suggestions
- Name suggestions for application service candidates
- Map your application services using tags in the Service Mapping Workspace

Service Mapping AI Agent activity list tables

Use this reference to interpret the tables in the Service Mapping AI Agent activity list.

Service Mapping AI Agent activity [sn_sm_gen_ai_svc_create_run]

Column	Description
Name	The name of the service map created or evaluated by the AI Agent.
Mapping results	The outcome of the mapping attempt: • **Map created** – The AI Agent successfully created a new service map. • **Already exists** – A service topology already exists for this candidate; the agent skipped creation. • **Missing service context** – The agent could not identify a valid service from the candidate data. • **Missing data** – Insufficient candidate data was available to create a map. • **Failed** – The mapping attempt failed.
Created	The date and time the activity record was created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Business App Map Agent activity [sn_sm_gen_ai_ba_to_as_conn_run]

Column	Description
Name	The name of the business application.
Application service mapped	The application service that the Agent identified as a match.
Mapping results	The outcome of the mapping attempt: • **Linked** – The AI Agent found a matching application service for the business application and automatically created a "Uses::Used by" relationship between them. • **Already connected** – The relationship between the business application and the application service was created manually. • **Needs review** - The AI Agent identified a possible match between the business application and an application service. The confidence level was not high enough to create the relationship automatically. • **No match found** - The AI Agent did not find an application service that could be linked to the business application with sufficient confidence. No relationship was created. • **Failed** – The AI Agent has failed to link between the business application and the service instance.
Confidence	The AI confidence rating for the mapping: Very low, Low, Medium, High, or Very high.
Created	The date and time the activity record was created.

Predefined traversal rules for application services

Service Mapping uses traversal rules to bring in related configuration items and expand the application service map. Several predefined rules are provided by default to help find and correctly map these relationships.

About traversal rules

Predefined traversal rules in Service Mapping are key to identifying the relationships between configuration items (CIs) and building accurate application service maps. To maintain integrity and avoid issues, such as incomplete maps or service disruptions, these rules are fixed. You can't change, customize, or remove them. To create custom traversal rules to meet your organization's needs, see Traversal Rules for Application Services form completion.

Predefined traversal rules

Parent Type	Relationship	Child Type	Priority Level	Is Reverse
cmdb_ci	Runs on::Runs	Hardware [cmdb_ci_hardware]	100	false
cmdb_ci	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]	200	false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parent Type	Relationship	Child Type	Priority Level	Is Reverse
cmdb_ci_lb_service	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	300	false
cmdb_ci_cluster_node	Hosted on::Hosts	Server [cmdb_ci_server]	400	true
cmdb_ci	Contains::Contained by	Tracked Configuration File [cmdb_ci_config_file_tracked]	500	false
cmdb_ci_cluster_node	Cluster of::Cluster	Cluster [cmdb_ci_cluster]	600	false
cmdb_ci_cluster_node	Cluster of::Cluster	Cluster [cmdb_ci_cluster]	700	true
cmdb_ci_cluster_node	Hosted on::Hosts	Server [cmdb_ci_server]	800	false

For example, the first row of the table illustrates a traversal rule that identifies and adds the hardware (cmdb_ci_hardware) supporting a given software application (cmdb_ci). When **Is Reverse** is set to false, the discovery process starts with the application and identifies the hardware it runs on-traversing from the parent CI to the child CI. The relationships are then added to the application service map.

To learn more about traversal rule components and related terminology, see Traversal rules in Service Mapping.

Traversal Rules for Application Services form completion

To edit an existing CI relationship or add a new CI relationship during the tag-based discovery process, complete the Traversal Rules for Application Services form.

Field	Description
Rule Definition	To create a new CI relationship: **1.** Select **New**. **2.** Select the Lookup using list icon and to view the available CI Relationship Type Rule Definitions. Continue as follows, depending on your situation: ○ If you see a CI relationship that you want to use: **a.** Select that relationship. **b.** Select **Submit**. The new CI relationship is added to the CI Relationship Type Rule Definitions [cmdb_rel_type_rule_definitions_list] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ If you don't see the desired CI relationship: **a.** Select **New** in the CI Relationship Type Rule Definitions window. **b.** Select the CI Class from the **Parent Type** list. **c.** Begin entering the relationship in the **Relationship Type** field and select the appropriate relationship type from the list. **d.** Select the CI Class from the **Child Type** list. **e.** Select **Submit**. **f.** Select **Submit** on the Traversal Rules for Application Services - New Record form. The new CI relationship is added to the CI Relationship Type Rule Definitions [cmdb_rel_type_rule_definitions_list] table. To edit a preconfigured CI relationship: **1.** Select the preconfigured CI relationship in the **Parent Type** column. **2.** Select the Preview this record icon, then select Open Record. **3.** Modify the **Parent Type**, **Relationship Type**, or **Child Type** as needed. **4.** Select **Update**. For example, to discover Linux servers hosting storage devices, select the following values: • Parent Type: Linux Server [cmdb_ci_linux_server] • Relationship Type: Contains::Contained by • Child Type: Storage Device [cmdb_ci_storage_device]
Is reverse	Select this check box to traverse a CI relationship starting from the child CI instead of the parent CI. Using the previous example, select this check box to begin with the storage devices and find the Linux servers that host these specific storage devices.
Used By	Select only **Service by Tags**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: Don't select **Service Recomputation**. This selection helps prevent editing any existing traversal rules or adding new ones and produces an error message.
Active	Select this check box to use this traversal rule in the tag-based discovery. Alternatively, clear this check box to exclude this traversal rule from the tag-based discovery. **Note:** By default, you can have only five traversal rules participating in tag-based discovery. To change the maximum number of traversal rules, edit the value for the svc_by_tags.max.traversal.rules.active property in the System Property [sys_properties] table.
Order	Enter a number for the order. The system uses traversal rules for discovery in sequence from low order to high. **Note:** When configuring traversal rules with dependencies, verify that prerequisite rules are executed before dependent ones by assigning them a lower-order number.

Name suggestions for application service candidates

Automated Service Suggestions provides meaningful name suggestions for service identification, which you can access on the **Application Service Suggestions** tab in the Service Mapping workspace.

Application service candidate name suggestions

Name suggestion	Description
Candidate name suggestion	The primary candidate name suggestion prioritizes the Service Fingerprints name, then defers to other name suggestions as needed.
Service Fingerprint	Name suggestions taken from a maximum of three Service Fingerprints related to the candidate.
Load balancer	Name suggestions taken from the names of the top three load balancer services associated with the candidate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application service candidate name suggestions (continued)

Name suggestion	Description
Process	Name suggestions taken from the process entry point command lines for the candidate. An algorithm identifies the most relevant words and provides the top three name suggestions.
Entry point	Name suggestions taken from all the entry point IP addresses associated with the candidate. Entry point name suggestions is deprecated beginning with version 1.12.3. This column is still displayed but is no longer populated with data. It will eventually be removed.
AFP (Application fingerprint)	Name suggestions taken from the top three application fingerprints associated with the candidate.

Related topics

Automated Service Suggestions

Map application services based on Automated Service Suggestions

Application service readiness dashboard in configurable workspace

Review the information on the dashboard to confirm that you're ready to discover and map application services based on machine learning (ML). Service Mapping uses data processed by Predictive Intelligence to generate suggestions for traffic-based connections.

The Application service readiness dashboard is part of Service Mapping Plus, available on the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Predictive Intelligence evaluates connections between application fingerprints, CIs, and processes, and ranks their relevancy. Service Mapping uses this information to create connections based on connection rules. It also generates connection suggestions for servers and load balancers for you to decide which connections to add or remove from the service instances.

Widgets on the ML Dashboard page show the information about prerequisites and issues related to service discovery based on Predictive Intelligence. Select links inside the widgets and reports to navigate to the related list or form.

Required ServiceNow AI Platform roles

service_mapping_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Access the Application service readiness dashboard

To open the dashboard, navigate to **Workspaces** > **Service Mapping**. Then select the Application service readiness icon .

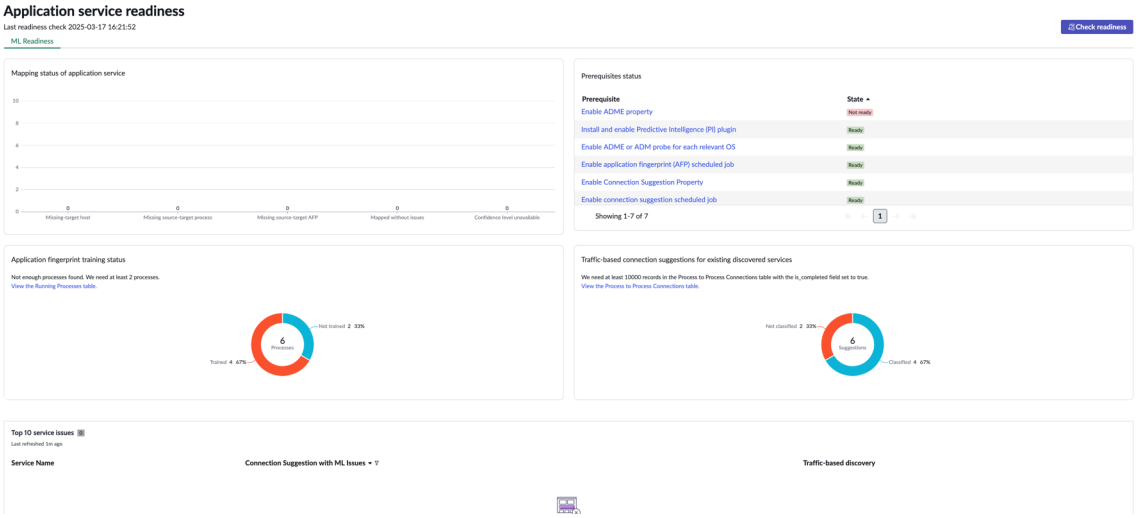

Reports

The dashboard includes the following reports.

Title	Type	Source table	Description
Mapping status of application service	A bar report	ML-Related Service Status [ml_related_service_status]	A bar report that provides the summary of ML-related issues in mapped application services. For detailed information, see Mapping status of application services.
Application fingerprints training status	A donut report	AFP Training Status [afp_training_status]	A donut report that shows the status of application fingerprint training. Predictive Intelligence trains predictive models and machine-learning solutions. View the training status for application fingerprints to understand if your deployment is ready for mapping using Predictive Intelligence.
Traffic-based connection suggestions for	A donut report	Connection Suggestions [sa_ml_connection_suggestion]	A donut report that reflects the ratio of classified connections to valid connections

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
existing discovered services			in the Connections Suggestions table. This table is only populated during top-down discovery.

Mapping status of application services

Review the summary of ML-related issues in mapped application services.

Category	Description
Mapped without issues	The number of service instances discovered without ML-related issues.
Missing source-target AFP	The number of service instances missing some application fingerprints for a source or target process. To solve issues, calibrate the fingerprint-based discovery.
Missing source-target process	The number of service instances missing some source-target process information. To solve these issues, rediscover the target hosts and ensure that the relevant processes are discovered.
Confidence level unavailable	The confidence level indicates the likelihood of this connection being part of the service instances. If the confidence level appears as **N/A**, wait until the application fingerprints training is complete.
Missing target host	The number of service instances with some CI connections not fully discovered, because the horizontal discovery didn't discover host CIs. To solve these issues, rediscover the target hosts and ensure that the relevant processes are discovered.

Prerequisites status

Service instance mapping requires the integration of several modules and applications: credentials, Predictive Intelligence, enhanced Application Dependency Mapping (ADME) discovery, and scheduled jobs. Review the list of prerequisites and ensure that the state of all prerequisites is **Ready**.

Prerequisite	Description
Install and enable Predictive Intelligence (PI) plugin	Service Mapping uses Predictive Intelligence to generate connection suggestions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisite	Description
	1. Ensure that the Predictive Intelligence plugin appears **Installed**. 2. Click the Predictive Intelligence tile and verify that the **Status** is **Active**. If the status is **Inactive**, select the **Activate/Repair** link under **Related Links**.
Enable ADME property	Service Mapping uses ADME probes. Ensure that the *glide.discovery.enable_adme* property is set to **True**.
Enable ADME or ADM probe for each relevant OS	There are multiple ADME probes for different types of operating systems. Enable all probes necessary to discover configuration items (CIs) in your environment.
Allow IP address expansion	This setting helps manage the expansion of global addresses (such as '0.0.0.0' or '*') into individual IP addresses. The limit for this expansion must be greater than 0. To verify and set this limit: 1. Navigate to the System Property [sys_property] table. 2. Locate the *sn.adm.ip_expansion_limit* property. 3. Ensure that this property is set to a positive number.
Enable application fingerprint (AFP) scheduled job	Ensure that the *Applications suggestion - ITOM Autodisco* scheduled job that controls the fingerprint-based discovery is set to **Active**.
Enable Connection Suggestion property	Confirm that discovery based on Predictive Intelligence is enabled. Navigate to the System Property [sys_properties] table and verify that the *sa_ml.connection_suggestions.active* property is set to **True**.
Enable connection suggestion scheduled job	Ensure that the **Status** of the *Service Mapping - Traffic Process to Process* scheduled job is **Active**. This schedule job triggers generation of connection suggestions.

Service issues

Review the list of service instances most affected by ML-related issues. The list of most affected services is available if the connection suggestions feature is enabled in your deployment. The

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

list shows service names and the number of ML-related issues for each of them. It also indicates if the traffic-based feature is enabled for the services.

Related topics

- Learn about ADME probes
- Enable and configure discovery using ADME probes
- Calibrate fingerprint-based discovery
- Verify fingerprint-based discovery generates suggestions
- Troubleshooting guide for Service Mapping ML Connection Suggestions [KB0963421]

Components installed with Service Mapping Plus

Several types of components are installed with activation of the Service Mapping Plus plugin, including tables and scheduled jobs.

Note: The Service Mapping Plus application is available in the ServiceNow Store.

Scheduled jobs installed

Scheduled job	Description
Update statistics for application service readiness dashboard	Refreshes the information displayed on the Application service readiness dashboard. Runs hourly. In addition, runs when the **Check Readiness** button is selected on the **Application service readiness dashboard**. The system doesn't trigger this scheduled job sooner than five minutes since the last run.

Tables installed

Table	Description
ML-Related Service Status [ml_related_service_status]	Contains the information about the ML-related issues in mapped application services.
Top ML Error Services [top_ml_err_services_list]	Contains the list of up to 10 application services containing the largest number of ML-related issues.
ML Dashboard Prerequisites [ml_dashboard_prerequisites]	Contains the links to forms and tables that require setting up for mapping based on ML and the status information for these configurations.
AFP Training Status [afp_training_status]	Contains the status of application fingerprint training.
Connection Suggestions Training Status [connection_suggestion_training_status]	Contains the status of connection suggestions.
ML P2P ETA Calculation Constants [sa_ml_p2p_eta_values_list]	Stores the information on how the system calculates the time remaining for connection suggestion training.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties Installed

Property	Description
sa.asc.logger_level	Adds messages to the system log for the application service candidate. • **Type:** string • **Default value:** error • **Other possible values:** debug **Note:** Setting the property to debug displays both debug and error messages.
sa.topdown.enable_cmdb_based_mapping	This property enables CMDB-based mapping using the CMDB instead of the MID Server. • **Type:** boolean • **Default value:** false • **Other possible values:** true
sn_sm_scoped_app_app.sa.unified_map.enable	Enables access to the Unified Map through the Service Mapping workspace. • **Type:** boolean • **Default value:** false • **Other possible values:** true
sn_cmdb_ws_config_property.unifiedmap.map_search.max_nodes	To ensure high performance, this property limits the number of nodes that can appear on a map. • **Type:** integer • **Default value:** 250 Higher values might decrease performance because Unified Map loads all elements and related details when opened. **Note:** The *unifiedmap.map_search.max_nodes* property in the workspace-specific [sn_cmdb_ws_config_property] table replaces the deprecated *sn_cmdb_ws.node.map.max.edge.count* property.
sn_sm_scoped_app.svc_by_tags.max_candidates_per_family	Limits the maximum number of tag-based service candidates created for a single family. If the family definition includes more candidates, the most recently updated CIs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	take precedence. This value applies to all tag-based service families. • **Type:** integer • **Default value:** 200
sn_sm_scoped_app.max_valid_discovery_period_in_days	Sets the maximum number of days a host is considered valid when generating application service candidates. Setting this value to any non-positive integer deactivates this feature. ⚠ **Warning:** If this value is set lower than or too close to the discovery cycle, it might result in the removal of all hosts. Consequently, no new candidates are created and existing candidates are deleted. • **Type:** integer • **Default value:** 0

Components installed with Service Mapping

Several types of components are installed with activation of the Service Mapping plugin, including tables, user roles, properties, and scheduled jobs.

Demo data is available for this feature.

Roles installed

Role title [name]	Description	Contains roles
[app_service_admin]	Creates and modifies service instances, creates service groups, views, and edits service instance maps.	[itil]
[app_service_user]	Views maps for operational service instances and retrieves service content using the getContent - GET REST API. The itil role that serves as the basic helpdesk technician role contains the app_service_user role.	[cmdb_read]
[itom_admin]	Configure IT Operations Management applications using the ITOM Guided Setup.	[sn_bm_client.benchmark_data_viewer] [sm_itom_license.reader]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description	Contains roles
	Note: The ITOM Guided Setup does not cover Service Mapping.	
[sam_core_admin]	An internal role.	[sam_core_user]
[sam_core_user]	An internal role.	[itil]
[service_mapping_admin]	Sets up the Service Mapping application. Maps, fixes, and maintains service instances. Also performs advanced configuration and customization of the product. Assign this role to application administrators.	• [app_service_admin] • [app_service_owner] • [discovery_admin] • [itil_admin] • [itom_admin] • [sam_core_admin] • [personalize_dictionary] • [pd_admin] • [pd_user] • [service_mapping_user]
[sm_app_owner]	Provides information necessary for successful mapping of a service instance. Once a service is mapped, this user reviews the results and either approves it or suggests changes. Assign the sm_app_owner role to users who own service instances and are familiar with the infrastructure and applications that make up the services.	[service_mapping_user]
[sm_itom_license.reader]	Views subscription information and statistics for IT Operations Management applications.	[cmdb_read]
[service_mapping_user]	Views maps for operational service instances to plan change or migration, as well as analyze the continuity and availability of services. Assign this role to application users.	• [sam_core_user.service_mapping_user] • [itil.service_mapping_user]
[sn_svc_err.admin]	Fixes discovery errors in bulk using the error framework feature in Service Mapping and Discovery.	[itil]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description	Contains roles
[pd_admin]	Creates new and modifies existing patterns, which Service Mapping and Discovery use for performing horizontal and top-down discovery.	[pd_user]

Note: Users with the sn_cmdb_editor or itil role only can view all service instances.

Scheduled jobs installed

Scheduled job	Description
Deferred Discovery Cleanup [sysauto_script_ba4bee737f0022008f1c3b19befd9108]	Runs every minute. Detects and cancels discovery schedules on TCP connections that are not in the active state anymore (5 minutes of no activity).
Update Query Based Services [sysauto_script_08745913f31131002e6bae4716612bff1]	Runs every 10 minutes. Updates views for all defined technical services.
Update App Service Status [sysauto_script_5641702ac32102003e76741e81d3ae9e]	Runs every minute. Updates the discovery status attribute for application services. The discovery status is calculated as the overall status over all endpoints that an service instance contains.
Flow Discovery Scheduler [sysauto_script_74c676f0dbb0220060ff742eaf9f19f2]	Runs every 10 minutes. Triggers the network flow discovery and writes the data into the [sa_flow_connector] table. This scheduled job is disabled by default. Enable this scheduled job to discover CIs using Netflow as described in Data collection and discovery using Netflow.
Generate entry point candidates from netflow and netstat data [sysauto_script_b0b95067c3d0030039a3553a81d3aeb5]	Runs daily. Scans the TCP Connection [cmdb_tcp] table and uses this data to generate entry points and store them in the Candidate Entry Points [sa_cand_entry_point] table. This scheduled job supports generating entry points necessary for mapping multiple application services (in bulk). This scheduled job is enabled by default.
Save thread dumps [sysauto_script_a49cd030c302020039a3553a81d3aefd]	Not active by default. When enabled, collects thread dumps and stores the data in the [sys_thread_dump] table.
Slow steps stats persist [sysauto_script_dc0de690c302020039a3553a81d3aeb0]	Runs every minute. Collects data on slow execution of pattern steps and writes this data into the [sys_step_pattern] table.
Calculate number of servers [sys_trigger_db438f407f510200f128134f8dfa91c6]	Runs every Sunday. Updates the licensing metric with the overall number of discovered servers on instance. This information is used for billing. On domain separated instances, the number is updated for each domain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scheduled job	Description
Run DNS Lookup For All Entry Points and Manual Connection Endpoints [f178ab457fb213001952baf8befa91d3]	Runs the DNS lookup probe on all entry points and manual connection endpoints and updates information on their IP addresses in the CMDB. This scheduled job is disabled by default.
[sm dedup tracker for specific discovery sources]	In some cases, Service Mapping and third-party discovery or monitoring tools identify the same configuration item (CI) differently. It results in the CMDB application table containing two records for the same CI. This scheduled job scans the CMDB application tables that contain information about CIs identified as generic applications. If the scheduled job identifies duplicate records based on the CI type and the discovery source, it creates a deduplication task.
[sys_auto_flush_5211e7e1c320220039a3553a81d3aae6]	Cleans the [sa_flow_connection] table from records not updated for the last seven days.
[sys_auto_flush_956d498b938503006a7031f63b7ffbc2]	Cleans the [sa_performance_statistics] table from records containing statistics of overall event processing time and created earlier than 30 days ago.
[sys_auto_flush_ce1a978193001200e0e931f635...	Cleans the [sa_performance_statistics] table from records created earlier than 2 days ago.
[sys_auto_flush_b24d9c30c302020039a3553a81d3aa58]	Cleans the [sys_thread_dump] table from records created earlier than 7 days ago.
[sys_auto_flush_cf129ebe7f223200995cbaf8befa91a6]	Cleans the [sa_ml_feature_data] table from records not updated in the last 7 days, and have a special state (empty label field).
[sys_auto_flush_d2516be1c320220039a3553a81d3aae2]	Cleans the [sa_flow_service] table from records not updated in the last 7 days.
[sys_auto_flush_db416be1c320220039a3553a81d3aae5]	Cleans the [sa_flow_server_comm] table from records not updated in the last 7 days.
[sys_auto_flush_41d9b007b71310105984eb25de...	Cleans the Connection Suggestion [sa_ml_connection_suggestion] table from the old records that match the following conditions: • They weren't updated in the last 7 days. • The suggestion record has an active discovered TCP connection (the TCP connection record is not absent). • The record doesn't have a user decision “Add” or “Remove” associated with it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables installed

Table	Description
Pre Service Mapping Task Scripts [sa_pre_task_script]	Contains scripts that Service Mapping runs on the instance and collect information that helps identify the relevant pattern for a CI.
Traffic Based Connection Qualifiers [cmdb_ci_qualifier_traffic_based_connection]	Legacy content. Not used in the current version.
Dynamic CI Group [cmdb_ci_query_based_service]	Contains technical services.
Components [sa_bs_components]	Contains information on components that are part of the service instance. Customers enter this information manually in the Questionnaire tab of the service instance form.
Candidate Entry Points [sa_cand_entry_point]	Contains information on potential entry points that Service Mapping retrieved from load balancers or traffic-based connections.
Dashboard View [sa_dashboard_view]	Contains configuration that defines which services to display in the Service tree.
Service Discovery Messages [sa_discovery_message]	Contains discovery messages and errors that Service Mapping generates during the discovery process.
ITOM Error task [sa_error_handler_task]	Contains information that Service Mapping and Discovery use for error-handling functionality.
Error EP AS Manual EP [sa_errors_as_manual_ep]	Contains data on errors that customers skip during the following procedure: Skip errors to continue discovering an application service.
Manual connections [sa_manual_connections]	Contains information on endpoints that customers use while adding manual connections.
ML Feature Data [sa_ml_feature_data]	Internal table.
ML Model [sa_ml_model]	Internal table
ML Traffic Validation [sa_ml_traffic_validation]	Internal table
Performance Statistics	Internal table for storing performance data on impact calculation in Event Management.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sa_performance_statistics]	
Service Maps [sa_service_counters]	Stores data about number of CIs and discovery errors per each discovered service.
Simulation entry points [sa_simulation_entry_points]	Internal table.
Service Process Tasks [service_process_task]	Stores data requests and review requests for application services.
Slow Steps [sys_step_pattern]	Stores information about slow execution of pattern steps. This data is used for monitoring and troubleshooting.
Thread dump [sys_thread_dump]	Internal table for storing performance data used for troubleshooting in Event Management.
Automation Error Category [automation_error_category]	Stores all supported categories for errors.
Automation Error Categories Statistics [automation_error_category_stats]	Stores the latest error statistics on all error categories.
Automation Error Code [automation_error_code]	Stores all supported error codes for discovery errors.
Automation Error Code Statistics [automation_error_code_stats]	Stores the latest error statistics on all error codes.
Automation Error Discovery Categories Statistics [automation_error_discovery_category_stats]	Deprecated.
Automation Error Discovery Code Statistics [automation_error_discovery_code_stats]	Deprecated.
Automation Error Discovery Message [automation_error_discovery_msg]	Stores the latest error messages.
Automation error instance category statistics [automation_error_instance_category_stats]	Stores the latest error statistics on all categories per instance (discovery status).
Automation Error Categories Statistics	Stores the latest error statistics on all error codes per instance (discovery status).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[automation_error_instance_code_stats]	
Automation error instance message [automation_error_instance_msg]	Stores all latest errors associated with the instance (discovery status).
Automation error instance message [automation_error_stats_update_time]	Stores the last statistics update time for every application.
Automation Error Suggestion [automation_error_suggestion]	Contains all the configured suggestions that presented in the UI per different error codes.
Automation Source Application [automation_source_application]	Contains the source applications that are using the Discovery and Service Mapping error framework.
Mapping between suggestion and error code [suggestion_error_code_m2m]	Contains information on mapping between error code and suggestion.
ITOM Error task [sa_error_handler_task]	Contains the tasks configured per different groups of errors.
Applicative Credentials [sa_applicative_credentials]	Contains credentials used when running application-specific commands on servers.
Service Mapping CI attributes [sa_ci_attr]	Contains CI attributes relevant only for discovery process.
Related CI Types [sa_ci_to_pattern]	Assigns CI types to patterns (relevant only for the horizontal discovery performed by Discovery).
Pattern Debugger Session [sa_debug_session]	This internal table is used for the Pattern Designer components.
Pattern Debugger Session Status for UI [sa_debug_session_status]	This internal table is used for the Pattern Designer components.
Sa Deferred Service Discovery [sa_deferred_service_discovery]	This internal table is used for Service Mapping components.
Service Discovery Log [sa_discovery_log]	Contains discovery logs collected by Service Mapping.
Service Discovery Messages [sa_discovery_message]	Contains messages created by Service Mapping during service discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Sa Endpoint Status [sa_endpoint_status]	Contains discovery status for endpoints in Service Mapping.
Find Host Scripts [sa_find_host_scripts]	Contains the names of script includes that find hosts based on the endpoint glide records.
nfdump file [sa_flcon_file_nfdump]	This internal table contains information used for Netflow-based discovery.
nfdump install [sa_flcon_local_nfdump]	This internal table contains information used for Netflow-based discovery.
AWS VPC flow logs [sa_flcon_vpc_flow_log]	This internal table contains information used for Netflow-based discovery.
SA Flow Connection [sa_flow_connection]	This internal table contains information used for Netflow-based discovery.
Flow Connector [sa_flow_connector]	This internal table contains information used for Netflow-based discovery.
Flow Server Communication [sa_flow_server_comm]	This internal table contains information used for Netflow-based discovery.
Flow Services IP/Port and Statistics [sa_flow_service]	This internal table contains information used for Netflow-based discovery.
Host Class to Capability [sa_host_class_to_capability]	Maps server type to capabilities required by Service Mapping.
Infra Path To Elements [sa_infra_path_assoc]	Contains data on components which are part of a network path.
Mapping External Discovery Commands [sa_mapping_ext_commands]	Contains data on customized commands created by user.
Network Paths [sa_network_paths]	Contains data on network paths discovered by Service Mapping.
Horizontal Discovered Paged Payloads [sa_paged_payload]	Service Mapping uses data in this table for splitting payload sent from the MID Server if the discovery results in a very large payload.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Discovery Patterns [sa_pattern]	Contains preconfigured and customized patterns that Service Mapping uses for discovery.
Service Mapping Draft Pattern [sa_pattern_draft]	Contains non-activated discovery patterns.
Service Mapping Relation attributes [sa_rel_attr]	Contains relation attributes used during service discovery process.
Rel to Infra Path [sa_rel_to_infra_path]	Maps relations to network or storage paths.
Service Mapping Server Usage [sa_server_usage]	Contains count of servers used by Service Mapping.
File System To Storage Path [sa_storage_paths]	Contains data on storage paths discovered by Service Mapping.
Storage Path Connectivity Map [sa_storage_path_connectivity_map]	Maps file system to storage volume.
Traffic Based Discovery Rules [sa_traffic_based_rules]	Contains data on configuration of traffic-based discovery.
Uploaded File [sa_uploaded_file]	Contains data on files uploaded to the MID Server and used by Service Mapping.
Configuration Item [cmdb_ci]	Stores data on all configuration items (CIs) in the CMDB.
Application [cmdb_ci_appl]	Stores data on all configuration items (CIs) of application type in the CMDB.
IP Service [cmdb_ip_service]	Stores data for credentialless discovery support for the original table.
Application Category [discovery_category_appl]	Stores data on application categories used by Pattern Designer for managing pattern libraries.
Device Classification Category [discovery_category_device_class]	Stores information on relationship between a device classifier and protocol used during horizontal discovery performed by Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Device Info Category [discovery_category_device_info]	Stores information on device category used during horizontal discovery performed by Discovery.
Protocol Category [discovery_category_protocol]	Stores information on protocol category used during horizontal discovery performed by Discovery.
HTTP Classification [discovery_classy_http]	Stores the HTTP classification definition.
HTTP Classification Match [discovery_classy_http_match]	Stores the HTTP classification conditions and the information about the pattern to launch if classification condition is matched.
Discovery Pattern Launcher [discovery_pattern_launcher]	Stores information about the pattern launcher mechanism.
Cloud Execution Pattern [discovery_ptrn_cld_lchr]	Stores information about the pattern launcher mechanism for discovering cloud applications.
Serverless Execution Pattern [discovery_ptrn_hostless_lchr]	Stores information about the pattern launcher mechanism for serverless discovery.
Discovery Pattern Parameters [discovery_ptrn_lnch_param]	Stores parameters for serverless/cloud application execution.
Pattern Input Parameters [discovery_ptrn_lnch_param_def]	Stores the key parameters collected by patterns.
Applicative Credentials [sa_applicative_credentials]	Stores the applicative credentials.
Related CI Types [sa_ci_to_pattern]	Stores related CIs within patterns.
Custom operation [sa_custom_operation]	Stores information on discovery pattern operations defined by users.
Custom operation parameters [sa_custom_operation_param]	Stores custom operation parameters for discovery patterns.
Custom parsing strategy [sa_custom_parsing_strategy]	Stores information on pattern parsing strategies that are defined by users.
Debug session	Stores information collected from debug sessions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sa_debug_session]	
Pattern Debugger Session Status for UI [sa_debug_session_status]	Stores information on status of the debug session for the UI purposes.
Sa Deferred Service Discovery [sa_deferred_service_discovery]	Contains information used for starting the top-down debug session when the relevant host is still not discovered.
Discovery log [sa_discovery_log]	The log of several discoveries.
Find Host Scripts [sa_find_host_scripts]	Contains scripts for finding hosts in the CMDB based on the entry point information.
Host Class to Capability [sa_host_class_to_capability]	Contains information that maps the type of host to MID Server capabilities. For example, the Linux OS is mapped to the SSH WMI capability.
Mapping Discovery Commands [sa_mapping_ext_commands]	Contains customizable operations for the Pattern Designer.
Horizontal Discovered Paged Payloads [sa_paged_payload]	Contains the IE output payload if the split mechanism is deployed.
Discovery Patterns [sa_pattern]	Stores published patterns.
Service Mapping Draft Pattern [sa_pattern_draft]	Stores modified version of patterns.
Pattern Extensions [sa_pattern_extension]	Stores pattern extension sections.
Pattern Pre/Post Script [sa_pattern_prepost_script]	Stores pre/post/pre execution scripts for patterns.
Horizontal Discovered Payloads Snapshot [sa_payload_snapshot]	Stores the IE output payload for the delete per main CI.
Tracked File Definition [sa_tracked_file_definition]	Stores the definitions of the file tracking feature.
Uploaded File	Stores scripts and files that are used by the pattern engine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sa_uploaded_file]	

Properties installed

Note: To open the System Property [sys_properties] table, enter `sys_properties.list` in the navigation filter.

Property	Description
sa.max_relations_in_service	The maximum number of CI relations for a service. Exceeding the value automatically prohibits some services from displaying in the map view. Very large application services may cause performance issues. Instead of maps, users can view CI lists for large application services. • **Type**: integer • **Default value**: 5000 • **Limit**: 10000 • **Location**: System Property [sys_properties] table
best_practice.check_connectivity.expand_listening_pid	This property controls which PIDs Service Mapping checks during the hop operation in the connectivity check. This property is true, then Service Mapping only checks connectivity on PIDs whose ports are in the `listening` state, not `connected`. • **Type**: boolean • **Default value**: true • **Other possible values**: false • **Location**: System Property [sys_properties] table
com.snc.sw.mapping.reaper.checkpoints_to_keep	The maximum number of change marks maintained in the history timeline of the service. Setting the value to zero or any negative integer overrides the default setting and maintains all change marks for a three-month maximum. Using the property 'com.snc.sw.mapping.reaper.ignored_services_from_optimizatio allows the exclusion of specific services and returns to the removal of change marks based upon intervals.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Type**: integer • **Default value**: 1000
com.snc.sw.mapping.reaper.ignored_services_from_optimization	This property contains the service IDs separated by commas. Instead of the number specified by the property 'com.snc.sw.mapping.reaper.checkpoints_to_keep' these services retain all change marks for a three-month maximum. **Type**: integer
com.snc.sw.mapping.reaper.max_reap_size	The maximum number of change marks deleted each time the system removes old and unused records from the history of changes to the service. Setting the value to zero or any negative integer overrides the limit and allows the removal of all unnecessary change marks. • **Type**: integer • **Default value**: 100000
glide.service_mapping.computation_depth	This property controls the maximal depth of the service instance map. • **Type**: integer • **Default value**: 25 • **Other possible values**: any number higher than 1 • **Location**: System Property [sys_properties] table
sa_ml.connection_suggestions.active	Enable generation of the connection suggestions feature that is based on traffic-based discovery. Service Mapping generates connection suggestions using data analyzed by Predictive Intelligence. When this feature is enabled, you can use connection suggestions to manage CI connections in application services. • **Type**: boolean • **Default value**: true **Note:** If traffic-based discovery was enabled in your deployment prior to Quebec and you used it to discover at least one application service, the default value is *false*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Other possible values**: false • **Location**: System Property [sys_properties] table
sa.debugger.max_timeout	Maximum timeout (in seconds) since the last server activity during a Pattern Debugger run. • **Type**: integer • **Default value**: 120 • **Other possible values**: any number higher than 60 • **Location**: **Service Mapping** > **Administration** > **Properties**. For more information, see KB0854603: Pattern Debugger Session Timed Out.
sa.discovery_task_timeout_min	Maximum time for a Service Mapping task in minutes (including waiting for execution in internal queues and ECC queue). • **Type**: string • **Default value**: 120 • **Other possible values**: any number higher than 30 • **Location**: System Property [sys_properties] table
sa.enable_gray_out_on_resume_discovery	Enable sub-tree to be grayed out after Service Mapping resumes discovery on an edge, or rerun discovery on the entire topology. • **Type**: boolean • **Default value**: true • **Other possible values**: false • **Location**: System Property [sys_properties] table
sa.excluded_discovery_sources	Add this property to exclude discovery sources during host detection. • **Default value**: EventManagement • **Other possible values**: additional discovery sources, added after EventManagement, separated by a comma
sa.map.enable_loops_on_service_split	Enable you to connect a service instance to another, already existing, service instance even if the CI connection (endpoint) that you are connecting is part of this other service

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	instance. This property helps to avoid loops or cycles in service maps. • **Type**: boolean • **Default value**: false • **Other possible values**: true • **Location**: **Service Mapping** > **Administration** > **Properties**.
sa.mapping.legacy_mid_selection	Controls the MID Server selection algorithm used during top-down discovery. Keep this legacy property set to **False** to use the current MID Server selection algorithm. When set to true, Service Mapping uses the legacy MID Server selection algorithm and the MID affinity features have no effect, regardless of their configuration. This can result in incorrect MID Server identification in networks with overlapping IP ranges or reused IP addresses. • Type: boolean • Default value: false • Location: **Service Mapping** > **Administration** > **Properties**.
sa_ml.connection_suggestions.lb.routing.enable	State if a target host is a load balancer member or not in the Connection Suggestions list. • **Type**: true/false • **Default value**: true • **Other possible values**: false • **Location**: System Property [sys_properties] table.
sa.network_path_calculation.arp_enableadded per DOC 1074513 in Vancouver	Controls queries to the ARP table. Data is only calculated when the value is set to true. **Note:** For more information about querying the ARP table, see the ARP property documentation for Service Mapping [KB1311037] article in the Now Support Knowledge Base. • **Type**: true/false • **Default value**: false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Other possible values**: true • **Location**: System Property [sys_properties] table
sa.network_path_calculation.active	Enable network path calculation. • **Type**: true/false • **Default value**: true • **Other possible values**: false • **Location**: **Service Mapping** > **Administration** > **Properties**.
sa.rediscovery.batch_size	Number of discovery tasks executed in a single batch. • **Type**: integer • **Default value**: 100 • **Other possible values**: any number higher than 10 • **Location**: **Service Mapping** > **Administration** > **Properties**.
sa.service_history.max_allowed_bubbles	Number of changes that the history scale and the Changes tab display. • **Type**: integer • **Default value**: 3000 • **Other possible values**: any number higher than 3000 **Important:** Increasing the value for this property may cause performance issues with loading service instance maps. • **Location**: **Service Mapping** > **Administration** > **Properties**.
sa.service.max_ci_service_population	The maximum number of CI connections added to application services during the following operations: Converting manual services created in Event Management into application services with changes from the CMDB. • **Type**: integer • **Default value**: 1000 • **Limit**: 2000

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Location**: Add to System Properties [sys_properties] table. Increasing the default value may cause performance issues.
sa.storage_path_calculation.active	Enable storage path calculation. • **Type**: true/false • **Default value**: true • **Other possible values**: false • **Location**: **Service Mapping** > **Administration** > **Properties**.
sa.topdown.mapping_with_acc_mode	This property enables top-down discovery and service mapping using Agent Client Collector by defining the MID Server selection mode for top-down discovery. • **Type:** integer • **Default value:** 0 • **Other possible values:** 1,2 0: Top-down discovery runs using the MID Server exclusively. 1: Top-down discovery initially attempts to run using Agent Client Collector. If no agent is available, then it defaults to using the MID Server. 2: Top-down discovery runs using Agent Client Collector exclusively. If unsuccessful, the process terminates.
sa.traffic_based_discovery.active	Traffic-based discovery • **Type**: true/false • **Default value**: false • **Other possible values**: true • **Location**: **Service Mapping** > **Administration** > **Properties**. **Note:** This property applies only to legacy traffic-based connections. If machine-learning (ML) connection suggestions are enabled, ensure that the property value remains set to `false`.
sa.traffic_based_discovery.conn_aging_time	Time period in hours for a Traffic Based Connection to remain active since last discovered.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Type**: integer • **Default value**: 72 • **Other possible values**: any number higher than 24 • **Location**: **Service Mapping** > **Administration** > **Properties**.
sa.traffic_based_discovery.ignored_ports	Ports to ignore when found by traffic-based discovery. This property is available in the System Property [sys_properties] table. Change this property to define ports that Service Mapping ignores while performing traffic-based discovery. It makes discovery more efficient since resources are not wasted on discovering irrelevant connections. • **Type**: string • **Default value**: 445, 139, 111, 2049, 860, 3260, 135, 53 • **Other possible values**: any relevant port numbers • **Location**: System Property [sys_properties] table
sa.traffic_based_discovery.max_connections	Maximum number of traffic-based connections from a single CI. This property is available in the System Property [sys_properties] table. This property helps to keep the map size reasonable by limiting the number of possible CI connections. • **Type**: integer • **Default value**: 30 • **Other possible values**: any number higher than 1 • **Location**: System Property [sys_properties] table
sa.update_paths_in_service_model.active	Update network and storage paths in Service Model. • **Type**: true/false • **Default value**: true

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Other possible values**: false • **Location**: **Service Mapping** > **Administration** > **Properties**.
service_watch.implied_match.endpoint.black_list	A list of endpoint types for which Service Mapping does not apply process matching logic utilized during the top-down pattern execution. • **Type**: string • **Default value**: cmdb_ci_endpoint_ssis_mssql, cmdb_ci_endpoint_ssas_mssql • **Other possible values**: comma-separated list of endpoints • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.blacklist	Excludes these fields from the record in the service definition export file. • **Type**: string • **Default value**: [empty] • **Other possible values**: comma-separated list of fields • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.blacklist.cmdb_ci_endpoint	Excludes these fields from the entry point record in the service definition export file. • **Type**: string • **Default value**: [empty] • **Other possible values**: comma-separated list of fields • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.blacklist.cmdb_ci_service_discovered	Excludes these fields from the application service record in the service definition export file. • **Type**: string • **Default value**: discovery_status,operational_status,process_status,work_not

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Other possible values**: comma-separated list of fields • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.blacklist.cmdb_ci_service_group	Excludes these fields from the service group record in the service definition export file. Use comma to separate fields. • **Type**: string • **Default value**: [empty] • **Other possible values**: comma-separated list of fields • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.max_services_limit	Defines the maximum allowed number of services to export. • **Type**: integer • **Default value**: 100 • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.skip_empty_fields	Excludes fields with an empty value from the record in the service definitions export file. • **Type**: boolean • **Default value**: true • **Other possible values**: false • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.skip_reference_fields	Excludes reference fields from the record in the service definitions export file. • **Type**: boolean • **Default value**: true • **Other possible values**: false • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.export.skip_sys_fields	Excludes fields beginning with “sys_” from the record in the service definitions export file. • **Type**: boolean • **Default value**: true

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Other possible values**: false • **Location**: **Service Mapping** > **Administration** > **Properties**.
sm.service.import.auto_start_discovery	Enable automatic discovery on all successfully imported entry points. • **Type**: boolean • **Default value**: false • **Other possible values**: true • **Location**: **Service Mapping** > **Administration** > **Properties**.
sn_cloud_migration.categories.inclusion_list	A comma-separated allow list of categories for cloud migration. Only resources assigned to categories that are listed here can participate in the migration process. • **Type**: string • **Default value**: Physical,VMware,HyperV,IBM HMC,OpenStack,Solaris,Nutanix,Amazon AWS,Microsoft Azure,Google Cloud,IBM Cloud • **Location**: **Cloud Migration** > **Cloud Migration Workspace** > **Administration** > **Properties**
sn_cloud_migration.server_status.inclusion_list	A comma-separated allow list of server status for cloud migration. Only servers with the listed here install status can participate in the migration process. • **Type**: ? • **Default value**: 1 • **Other possible value**: Installed=1, On Order=2, In Maintenance=3, Pending Install=4, Pending Repair=5, In Stock=6, Retired=7, Stolen=8, Absent=100 • **Location**: **Cloud Migration** > **Cloud Migration Workspace** > **Administration** > **Properties**
sn_cloud_migration.os.exclusion_list	A comma-separated deny list of operating systems for cloud migration. The servers whose operating system is listed here are excluded from the migration process.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Type**: String • **Default value**: ESX, Windows 2000,Windows XP. • **Location**: **Cloud Migration** > **Cloud Migration Workspace** > **Administration** > **Properties**
sn_cloud_migration.vm_state.inclusion_list	A comma-separated allow list of VM instance states for cloud migration. Only VMs with states that are listed here can participate in the migration process. • **Type**: ?? • **Default value**: on • **Other possible values**: on, off, paused, scheduled, starting, stopping, retired, terminated, error, cancelled, pausing, terminating • **Location**: **Cloud Migration** > **Cloud Migration Workspace** > **Administration** > **Properties**
	The number of CI relations, exceeding which application service is not displayed automatically in the map view. Very large application services may cause performance issues. Instead of maps, users can view CI lists for large application services. Setting a negative value results in not limiting the number of added CI relations. Setting an invalid value, like a string, is equivalent to the default value, 5000. • **Type**: Integer • **Default value**: 5000 • **Location**: System Property [sys_properties] table
svc_by_tags.classes.whitelist	This property contains the list of allowed CI classes separated by commas. Tag-based mapping includes the CIs belonging to these classes. An empty 'value' field indicates the inclusion of all CI classes.
svc_by_tags.install_status.blacklist	This property contains the list of exclusions based upon the install status separated by commas. Tag-based mapping excludes any tagged CIs with the following install statuses: **Default value**: 7,100

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	Other possible values: Installed=1, On order=2, In maintenance=3, Pending install=4, Pending repair=5, In stock=6, Retired=7, Stolen=8
svc_by_tags.max.categories.per.family	Set the maximum number of categories allowed per service family. • **Type**: integer • **Default value**: 3
svc_by_tags.max.traversal.rules.active	Set the maximum allowed number of CI relationships that can be in the active state. Service Mapping uses only active CI relationships for tag-based discovery. • **Type**: integer • **Default value**: 5 • **Location**: System Property [sys_properties] table
svc_by_tags.max_ci_limit	Set the maximum allowed number of CIs that a tag-based application service can contain. • **Type**: integer • **Default value**: 1000 • **Limit**: 2000 • **Location**: System Property [sys_properties] table
sm.topdown.reuse_mid_from_hd_schedule	Use the MID Server configured during horizontal discovery for top-down discovery. • **Type:** boolean • **Default value:** false • **Other possible values:** true • **Location: Service Mapping** > **Administration** > **Properties**
sm.network_location.detect_host_by_location	Select the subnetwork location to allow for the identification of the appropriate host and MID Server. **Important:** For best results, you must enable this property and sm.topdown.reuse_mid_from_hd_schedule concurrently.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Type:** boolean • **Default value:** false • **Other possible values:** true • **Location: Service Mapping** > **Administration** > **Properties**

Domain separation and Service Mapping

Domain separation is supported in Service Mapping. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

How domain separation works in Service Mapping

In domain-separated environments, switch into relevant leaf domains to access actions related to mapping performed by Service Mapping. A leaf domain is a child domain of the lowest level, meaning it does not have any child domains itself.

Domain separation is handled in a slightly different way than simply configuring separate domains. Notice that you must select or specify the relevant leaf domain to perform the following mapping actions in domain-separated environments:

- When creating an service instance, the service instance is assigned to the user's domain.
- When manually adding a CI to an service instance, you can choose only CIs that belong to the service domain.
- When manually adding connections to CIs to an service instance, users must belong to the same domain as the application service.
- When transferring service definitions between instances, imported service instances and service groups are assigned to the user's domain.

Domain separation and Automated Service Suggestions

Automated Service Suggestions supports domain separation. This means that:

- The source and target process must reside in the same domain to create a connection.
- The user can view only the data in the current domain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: After installing Service Mapping Plus with domain separation support, verify that all content is in the appropriate domain. For a detailed explanation of the required steps, see the Service mapping plus for domain separated instance [KB1303398] article in the Now Support Knowledge Base.

Automated Service Suggestions supports domain separation starting with version 1.9.0. No support is currently offered for on-premise instances. Reach out to product management for further information.

For more context, refer to Discovery patterns used by ITOM Visibility topic and MID Server configuration for Service Mapping. The following topics provide operational information on how to use Service Mapping in deployments with domain separation:

- Enable traffic-based discovery for CI types or specific CIs
- Map application services using tags with classic Service Mapping
- Create or customize patterns
- Create entry point types for Service Mapping
- Enhance patterns without modifying identification sections

Related topics

Domain separation

Domain separation recommended practices for service providers

Service Mapping Recomputation

Keep application services up to date and track changes to services through the recomputation capability of the Application Service feature of Service Mapping. Recomputation jobs recalculate mapped application services when changes occur to an associated configuration item (CI) within the Configuration Management Database (CMDB).

The most common changes that can cause recomputation are:

- Topology changes: The relationship of the associated CI to other CIs changes. These changes might include the addition of a new relationship or the removal of a previous one.
- CI changes: A field value of the CI changes.

These changes might be caused by horizontal discovery, top-down discovery, service graph connectors, or manual operations. If the CI change doesn't affect the topology or service associations, the recomptuation process stops with no changes to the service.

Two types of recomputation jobs are available in the system scheduler: Service Mapping recomputation and Service Mapping fast recomputation. Fast recomputation selects services with only CI changes as part of their most recent batch and doesn’t include topology changes.

Recomputation jobs are configured to run every five seconds. After the jobs are initiated, they begin with the service that has waited the longest. When the job is complete, the effect of topology changes might be reflected in the topology of the map or a change in the service associations. You can view CI changes by selecting the map node of the CI. Successful recomputations appear as a single checkpoint on the application service timeline regardless of whether it represents one change or a batch of changes.

You can configure the recomputation type, frequency, and other factors as needed. For more information about Service Mapping, see the Service Mapping Recomputation Jobs: Configuration and Best Practices [KB1001151] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection Suggestions list

Review details of connection suggestions to decide which connections are relevant for application services.

Connection Suggestions list columns

Column	Description
From	The CI class or fingerprint process of the CI that is the connection source. This field appears only when you access connection suggestions by navigating to **Service Mapping** > **Administration** > **Connection Suggestions**.
To	The CI class, or fingerprint, or process of the CI that is the connection target.
Target Host	The host of the CI that is the connection target. If the **VIP** column shows **true**, the target host is a load balancer member. When you select such a target host, the system adds the load balancer and all its members as CIs in the application service.
IP	The IP of the CI that is the connection target.
Port	The port of the CI that is the connection target.
Confidence	The confidence level indicates the likelihood of this connection being part of the service instances. • **High** - Commonly a connection to an internal component that is only used by this specific component or application. • **Medium** - Commonly a connection to a middleware component used by multiple services, such as IBM MQ or Oracle WebLogic Server. • **Low** - Commonly a connection to an application deployed on numerous servers in the organization, typically for monitoring purposes, such as Tivoli Enterprise Monitoring Agents or MID Server. • **Very low** - Commonly a connection to a central application used by the entire organization, for example, Microsoft Active Directory or Okta. If the confidence level appears as **N/A**, wait until the application fingerprints training is complete. See the Application fingerprint training status report.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection Suggestions list columns (continued)

Column	Description
	Confidence level is available only if fingerprint based discovery is enabled. See Discover applications based on fingerprints for more information.
Decision	This value indicates the outcome of the previous action taken on this connection suggestion: • **Added manually** - You added a connection based on this suggestion. • **Excluded manually** - You removed a connection based on this suggestion. • **Undecided** - You reset the previous decision to add or remove a connection based on this suggestion. • **Added by rule** - Service Mapping added a connection based on a connection rule.
Rule	The rule that Service Mapping applied to add a connection.
VIP	The target host is a load balancer member. The system shows this information only for load balancers that you previously discovered using horizontal discovery. This field appears only when you access connection suggestions by navigating to **Service Mapping** > **Administration** > **Connection Suggestions**.
Decision changed by	The user who changed the decision last time. If the decision was changed by a rule, this field shows the MID Server user. For more information about setting up MID Server user role, see Create the MID Server user and grant the role . This field appears only when you access connection suggestions by navigating to **Service Mapping** > **Administration** > **Connection Suggestions**.
Updated	The timestamp of when this connection last changed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Discovery based on Predictive Intelligence
- Manage CI connections for a single service using suggestions
- Manage CI connections in bulk using suggestions

Entry point attributes

Service Mapping comes with a wide range of preconfigured entry point types that cover many commonly used applications. Check attribute definitions to correctly add entry points to your application services.

An entry point is a point where clients access a service instance. Entry points are stored as records in the Endpoint [cmdb_ci_endpoint] tables.

There are some general attributes shared by almost all entry point types as described in the following table:

General attributes

Attribute	Description
Host/Hostname/Host Name	The value of the target server on which the service is running. This value can contain a real host name, alias, IP, or VIP. **Important:** Do not use a localhost or loopback address (127.0.0.1) for discovery and mapping.
Port	The port number of the service that you want to discover.
Name	The name of the service that you want to discover.
URL	The URL of the service that you want to discover in the following format: http(s)://(address):(port). If you do not provide the port, the system uses the default port 443.
Network location	The location of the applicable subnetwork. **Note:** This field is optional and visible only if Network location is enabled.

To configure entry points for subnetworks, see Fine-tune Service Mapping with MID affinity and IP reuse.

There are some attributes which you configure differently depending on what entry point they relate to, as described in the following table:

Entry point attributes specific for some entry points

Entry point type	Attribute	Description
Active Directory Domain to Domain	Domain name	The AD domain name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Entry point attributes specific for some entry points (continued)

Entry point type	Attribute	Description
Controllers Endpoint	Forest name	The AD forest name.
	Domain controller distinguished name	The distinguished name of the AD domain controller. Usually it is a compilation of the domain controller name and the hostname.
	Domain Controller Name	The AD Domain Controller name.
Active Directory Forest Endpoint	Forest name	The AD Forest name.
AWS Resource Name Endpoint	ARN	The Amazon Resource name.
Application to Storage Endpoint	File in use by application	The path to the file used by an application. For example, mssql uses c:\applications\mssql\users.db. Service Mapping populates this attribute using patterns.
	File system name	The file system containing the file used by an application. For example, if a mssql uses c:\applications\mssql \users.db, the value is "c:".
	Connected array ids	N/A
Advanced Queue	Scheme	The name of the Oracle DB table where the Queue is defined.
	Instance	The name of the Oracle instance.
	Queue	The name of the queue.
BizTalk Connection	Adapter address	The connect string to the Biztalk adapter. It can be a path, URL, host, or port.
	Transport Type	The type of the Biztalk adapter. For example a URL, file SOAP, or FTP.
CRM Component	Role	The role of the CRM. Leave empty if unknown.
DB2 UDB	Instance	The name of the DB2 instance.
DCTM Connection to broker	Repository	The name of documentation repository.
DCTM Index Agent Connection	Instance Name	The name of the instance of Documentation Index Agent.
	Repository	The name of the documentation repository.
DCTM Job Processor	Repository	The name of documentation repository.
Citrix Delivery Controller	Farm Type	The type of the Delivery Controller.
	Controller Name	The name of the Citrix Controller.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Entry point attributes specific for some entry points (continued)

Entry point type	Attribute	Description
DocBase Connection	Repository	The name of the documentation repository.
EJB	JNDI Name	The jndi name of target EJB.
EMS	Queue	The name of Tibco EMS queue.
EMS JNDI	JNDI Name	The jndi name of Tibco EMS connection.
F5 Mirror	Origin IP	The IP address of the second pair F5 device.
	Origin host	The hostname of the second pair F5 device.
Fast Search	Repository	The name of repository of fast search.
ITAM Asset Center	Folder	The path to configuration folder.
JMS Flow	JMS project	The name of Tibco project.
JMS Flow	Queue Name	The name of Tibco EMS queue.
JMS Server	Queue Name	The name of the queue defined on JMS server.
HTTP Listener to WMB Dependency	Directory	The path to the WMB installation directory.
LDAP	Protocol	The protocol name: `ldap`.
	Type	The LDAP query type. This attribute is not supported. Leave empty.
MQ	Queue Name	The name of the IBM MQ Queue.
MQ Flow	Source IP	The IP of the source IP create connection to MQ.
	Queue Name	The name of IBM MQ Queue.
MQ Cross-memory	Command line	The command line of MQ.
	Queue Name	The name of IBM MQ Queue.
MSMQ Flow	Queue Name	The name of MS MQ queue.
MS SQL Server	Instance	The name of MSSQL instance.
MySQL Cluster Data Node	Node ID	The ID of MySQL instance.
	Node group	The group of MySQL instance
MySQL Cluster MGM	Node ID	The ID of MySQL instance.
MySQL Secondary Server	Config IP	The IP of the MySQL Config server.
	Config Port	The port of the MySQL Config server.
Oracle DB	Schema	The name of table on the Oracle instance.
	Instance	The name of oracle instance.
Oracle ESB Connection	Adapter data	The content of the adapter.
	Adapter Type	The type of the adapter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Entry point attributes specific for some entry points (continued)

Entry point type	Attribute	Description
Oracle RAC DB	Service Name	The name of the Oracle RAC service.
PostgreSQL DB	Instance	The name of the PostgreSQL DB.
Remote TCP Cross-memory	Source_ IP	The IP of the source host process.
	Source Port	The port of source process.
SAP APP EP	SID	The SAP instance SID name.
	Instance	The instance name.
SAP BO Servers	Server name	The name of the BO server.
Sharepoint connection	Origin URL	The URL of the IIS.
SSAS	Project	The name of the SSAS project (can be empty).
SSAS for MSSQL	Instance	The name of the MSSQL instance.
	Client Type	The type of the client connected.
SSIS	Job	The name of SSIS job can be empty.
SSIS file	Package file name	The name of the package file.
SSIS for MSSQL	Instance	The name of the MSSQL instance.
	Client Type	The type of the client connected.
StoreParameters Components	Front	The parameters of the Citrix application.
	Farm	The farm host the Citrix application.
	Application name	The Citrix application name.
XenApp or Presentation Server	Type	The type of the Citrix application.
	Farm	The farm host the Citrix application.
	Icon path	The logic path of the Citrix application.
	Argument	The parameters of Citrix application.
	Path	The logic path of the Citrix application.
	Exec	The EXE name of Citrix application.
	IP Address	The IP of host is member of Citrix farm.
Sybase	Instance	The name of Sybase instance.
Tibco Conf File	Configuration File	The path to Tibco BW project.
	Process Name	The name of Tibco BW main process.
	Queue Port	The port of the queue used by the Tibco process.
	Config IP	The IP of the source of the queue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Entry point attributes specific for some entry points (continued)

Entry point type	Attribute	Description
Tibco BW	Process name	The name of Tibco BW main process.
Tibco File Listener	File	The path of the file is used by Tibco BW for integration.
Tibco Hawk	Configuration directory	The path to Tibco Hawk Configuration folder.
WMB Dependency	Command line	The command line of the WMB.
XenApp or Presentation Server Components	Farm	The farm host the Citrix application.
	Icon path	The logic path of the Citrix application.

Related topics

Map a single application service using classic Service Mapping

Create entry point types for Service Mapping

MID Server properties used by Service Mapping

Service Mapping uses the MID Server parameters during discovery and mapping. Do not change default values for these parameters unless you are troubleshooting the discovery process.

General properties

Property	Description
mid.file_tracking.file_path_forbidden_strings	Comma-separated list of strings to search for in tracked file paths. Once any string in this list is found in a tracked file path, the file is removed from further processing. • **Default value**: job finished with status ERROR • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.

SSH-related properties

Property	Description
mid.sa.ssh.path	A location at the target system where the MID Server checks for commands to execute.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Default value**: /usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.truncate_last_line_regex	A regular expression of line that the MID Server removes from the end of the SSH command response. • **Default value**: Command.*performed by.* • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.remove_first_line_regex	A regular expression of line that the MID Server removes from the start of every SSH command response. • **Default value**: [empty string] • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.privileged_command_message	A prompt the MID Server sends with commands with sudo-level privileges. For example, a sudo-level command is sent with "Password:". • **Default value**: [empty string] • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.restricted_shell_regex	A regular expression text included on opening message of a restricted shell. Using a restricted shell, the SSH provider does not run some commands during the session initialization. • **Default value**: restricted shell • **Type**:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.terminal_cols	The number of columns in the terminal windows opened with the SSH session. • **Default value**: 300 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.terminal_rows	The number of rows in the terminal windows opened with the SSH session. • **Default value**: 800 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.time_between_change_shell_and_prompt	The time to wait between changing a shell during initialization and the time the shell prompt appears (in milliseconds). • **Default value**: 1000 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.connect_timeout	The SSH initial connection timeout (in milliseconds). • **Default value**: 2000 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.max_chars_single_response	The maximum number of characters allowed in a single SSH command response.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Default value**: 10000000 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.break_session_chars	The ASCII value of character used to break a session. • **Default value**: 7 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.no_change_shell_regex	The regular expression of text within command response after sending 'su' in order to change to superuser. If the regular expression matches, the MID Server does not change the shell. • **Default value**: pfsh • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.permission_denied_regex	The regular expression of command responses which the MID Server recognizes as an access problem and generates an appropriate message. • **Default value**: permission denied • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.more_line_command_regex	When the SSH response exceeds the terminal size, the system sends 'more' line which matches the regular expression. When the MID Server recognizes the 'more' line, it sends the space character. This is relevant for proprietary systems sending a non-standard 'more' prompt.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Default value**: [empty] • **Type**: • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.become_superuser_command	An SSH command for turning into a superuser. • **Default value**: su • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.prompt_changing_command_regex	SSH commands that change the shell prompt. When the command the MID Server sends matches the regular expression, the MID Server knows to expect a change in the server prompt. • **Default value**: changeto\\s+[^\\s]+ • **Type**: • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.ignore_errors_on_commands_regex	A regular expression for the command on which the MID Server ignores permission errors. • **Default value**: stat.*|istat.*|perl.* • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.prompt_last_char_regex	A regular expression of the last character in the default prompt of the server. • **Default value**: [$#>\\]\\:] • **Type**: string

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.timeout	The SSH command timeout (in milliseconds). • **Default value**: 60000 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.port	The SSH port number. To add more than one port, separate port numbers with a comma. For example, 22,33,44. • **Default value**: 22 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.connect.retry	The SSH connection retry count. • **Default value**: 1 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.additional_setup	The command the MID Server sends after it creates the SSH session. • **Default value**: [empty] • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.expect_after_su	A regular expression of the last character in the prompt after sendind 'su'.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Default value**: [:$#>\\]] • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.privileged_command_last_char	A character indicating the last expected character in a command result String. Once this character is found in the result, the MID Server does not wait for additional output. • **Default value**: [empty] • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.prompt	A string representing the non-standard SSH prompt. The MID Server uses this string when parsing results of SSH CLI commands. • **Default value**: SERVNOW> • **Type**: • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.proprietary_system_trial_count	Number of additional waiting periods on proprietary non-UNIX/Linux systems. The MID Server waits for a response after the initial waiting period of 5 seconds. After the initial 5 seconds, the MID Server waits for the response the number of times as defined by this parameter. The additional waiting time starts from 1 second and increases by 1 sec for every iteration. • **Default value**: 10 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.restart_after_unstopped_threads	The number of allowed unstopped threads before the MID Server is restarted, when it attempts to close the SSH session.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Default value**: 50 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.stty_command	The stty command for configuring appropriate settings. The MID Server runs this command when initializing the SSH terminal. • **Default value**: [empty] • **Type**: • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.inocrrect_password_regex	The expected string which the SSH terminal returns when the MID Server uses an incorrect password during the login process. • **Default value**: Sorry, try again|incorrect password attempt • **Type**: • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.max_unstopped_threads	Same as the mid.sa.ssh.restart_after_unstopped_threads parameter. • **Default value**: 50 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.ssh.idle_after_sudo	The period of time (in milliseconds) that the MID Server waits after sending a sudo command and expects to have no activity on the line.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Default value**: 3000 • **Type**:integer • **Location:**This property is not available in the base system. You can add the property as described in Create a MID Server property.

WMI-related properties

Important: WMI Collector is deprecated in New York patch 10 and later. See the KB0791835 WMI Collector deprecation details for more information.

Properties	Description
mid.servicewatch.wmi.host	The Host/IP of the server running the ServiceNow WMI collector. • **Default value**: localhost • **Type**: string • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.servicewatch.wmi.port	The port on which the ServiceNow WMI collector service listens to requests from the MID server. • **Default value**: 8585 • **Type**: integer • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.sa.wmi.timeout	The timeout on command execution in milliseconds. This property controls also the Powershell commands. • **Default value**: 60000 • **Type**: integer • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.sa.wmi.max_bytes_in_response	The maximum number of characters allowed in a response. • **Default value**: 20971520 • **Type**: integer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties	Description
	• **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.servicewatch.wmi.get_file_on_shared_drive_from_collector	The MID Server attempts to access the requested file on a shared drive directly from the MID Server. • **Default value**: true • **Type**: true/false • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.sa.wmi.access_denied_regex	The regular expression of error messages returned by server that the MID Server considers as an access problem and triggers trying other credentials. • **Default value**: Access is denied|Access denied|service installation failed|User credentials cannot be used for local connections|System error 1312|generic failure|Insufficient privilege|user rights| HRESULT: 0x800706BA • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.wmi.send_retry_count	The number of attempts to commiunicate with the ServiceNow WMI Collector service on each requests. • **Default value**: 1 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.servicewatch.wmi.mui	The MID Server supports MUI (Multilingual UI) configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties	Description
	• **Default value**: false • **Type**: true/false • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.powershell_api.winrm.remote_port	Set the port used by WinRM to connect to remote hosts. • **Default value**: 5985 • **Type**: integer • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.windows.management_protocol	The management protocol the MID Server uses to connect to Windows servers. Allowed values are WMI and WinRM. • **Default value**: WMI • **Type**: string • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.servicewatch.wmi.protocol	Communication protocol between the MID Server and the WMI collector. Allowed values: http/https. • **Default value**: http • **Type**: string • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.use_powershell	Allows using PowerShell for executing commands on Windows servers. • **Default value**: false • **Type**: true/false • **Location:** This property is available under specific MID Server parameters as described in Add a MID Server parameter.
mid.sa.prefer_powershell	Makes PowerShell the preferred method of executing commands on Windows servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties	Description
	Note: If you set this parameter to true, you must also set the mid.sa.use_powershell parameter to true. • **Default value**: false • **Type**: true/false • **Location: MID Server** > **Properties**
mid.win.baseboard.or.bios.insertion	Makes the *Windows OS - Servers* pattern discover the BIOS Serial Number instead of the Baseboard Serial Number. • **Possible values**: false - to populate the BaseBoard number; true - to populate the BIOS serial number. • **Type**: true/false • **Default value**: false • **Location: MID Server** > **Properties**

SNMP-related properties

Property	Description
mid.snmp.request.timeout	The time (in milliseconds) the MID Server waits for a response for each OID request. • **Default value**: 1500 • **Type**: Integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.port	The SNMP port. • **Default value**: 161 • **Type**: Integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.retry	The SNMP retry count. • **Default value**: 1 • **Type**: integer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.table_size_limit	Maximum table size to bring in the SNMP table requests. • **Default value**: 20000 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.table_walk_limit	Maximum requests in the SNMP walk session. • **Default value**: 10000 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.get_table_timeout_sec	Maximum time (in seconds) allowed for load of a whole SNMP table. • **Default value**: 600 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.walk_timeout_sec	Maximum time (in seconds) allowed for the SNMP walk session. • **Default value**: 600 • **Type**: integer • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.discover_engine	Discovery authoritative engine ID on SNMP V3 sessions. • **Default value**: true • **Type**: true/false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.snmp_v2_disabled	Disable usage of SNMP version 2 • **Default value**: false • **Type**: true/false • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.
mid.sa.snmp.failed_status_oid	OID indicating the Failed status. If an SNMP result starts with this OID, the MID Server considers it a failure. • **Default value**: 1.3.6.1.6.3.15.1.1.3 • **Type**: string • **Location:** This property is not available in the base system. You can add the property as described in Create a MID Server property.

Preconfigured CI relationships in tag-based discovery

Several preconfigured CI relationships participate in tag-based discovery, by default. The Traversal Rules for Application Services [svc_traversal_rules] table stores these relationships.

Important: You cannot delete or modify preconfigured CI relationships used for tag-based discovery from the Traversal Rules for Application Services [svc_traversal_rules] table.

Preconfigured CI relationships participating in tag-based discovery by default

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Virtualizes::Virtualized by	Hardware [cmdb_ci_hardware]
Hardware [cmdb_ci_hardware]	Runs::Runs on	Application [cmdb_ci_appl]
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Contains::Contained by	Operating-system-level Virtualization Container [cmdb_ci_oslv_container]

Additional preconfigured CI relationships for tag-based discovery

CI	Relationship	CI
Configuration Item [cmdb_ci]	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Additional preconfigured CI relationships for tag-based discovery (continued)

CI	Relationship	CI
Logical Datacenter [cmdb_ci_logical_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Kubernetes Service [cmdb_ci_kubernetes_service]	Contained By::Contains	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
OpenShift Project [cmdb_ci_openshift_project]	Contained by::Contains	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Contained by::Contains	Kubernetes Cluster [cmdb_ci_kubernetes_cluster]
Kubernetes Pod [cmdb_ci_kubernetes_pod]	Cluster::Cluster of	Kubernetes Service [cmdb_ci_kubernetes_service]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Deployed Config [cmdb_ci_openshift_dep_conf]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Build Config [cmdb_ci_openshift_build_conf]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Route [cmdb_ci_openshift_route]
Kubernetes Namespace [cmdb_ci_kubernetes_namespace]	Contains::Contained By	OpenShift Image Stream [cmdb_ci_openshift_images_stream]

Properties related to tag-based discovery

The Service Mapping plugin includes several properties specifically related to tag-based discovery.

Note: To open the System Property [sys_properties] table, enter `sys_properties.list` in the navigation filter.

Property	Description
sn_sm_scoped_app.svc_by_tags.max_candidates_per_family	Limits the maximum number of tag-based service candidates created for a single family. If the family definition includes more candidates, the most recently updated CIs take precedence. This value applies to all tag-based service families. • **Type:** integer • **Default value:** 200
svc_by_tags.classes.whitelist	This property contains the list of allowed CI classes separated by commas. Tag-based mapping includes the CIs belonging to these classes. An empty 'value' field indicates the inclusion of all CI classes.
svc_by_tags.install_status.blacklist	This property contains the list of exclusions based upon the install status separated by

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	commas. Tag-based mapping excludes any tagged CIs with the following install statuses: **Note:** In this context, CIs brought by traversal rules are not considered tagged by default. This property applies only to original tagged CIs unless svc_by_tags.install_status.filter_traversed_cis is also set to true. **Default value**: 7,100 **Other possible values**: Installed=1, On order=2, In maintenance=3, Pending install=4, Pending repair=5, In stock=6, Retired=7, Stolen=8
svc_by_tags.install_status.filter_traversed_cis	Extends the svc_by_tags.install_status.blacklist exclusion to CIs that traversal rules bring into a tag-based service. When set to true, the next service recalculation removes traversed CIs with an excluded install status. • **Type:** boolean • **Default value:** false
svc_by_tags.max.categories.per.family	Set the maximum number of categories allowed per service family. • **Type**: integer • **Default value**: 3
svc_by_tags.max.traversal.rules.active	Set the maximum allowed number of CI relationships that can be in the active state. Service Mapping uses only active CI relationships for tag-based discovery. • **Type**: integer • **Default value**: 5 • **Location**: System Property [sys_properties] table
svc_by_tags.max_ci_limit	Set the maximum allowed number of CIs that a tag-based application service can contain. • **Type**: integer • **Default value**: 1000 • **Location**: System Property [sys_properties] table

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Keyboard shortcuts for operating Service Mapping in accessibility mode

You can access all service instance maps and service forms using a keyboard in the accessibility mode.

When the whole map is selected, the information message at the top of the window offers guidelines on how to navigate inside the map.

Action	Keyboard shortcut
Access right-click menus.	Space
Expand clusters.	Enter
Move into the map pane from the **Properties** pane.	Escape
In the More Options menu, toggle options on and off.	Space
Select CIs and connections inside the map.	Arrow
In the timeline, you can edit the time inside the current time selection.	Number (1-9)
Move through UI elements in the service instance map window.	Tab

Tag-based mapping criteria examples for service families

When preparing to map application services based on tags, create tag categories that contain tags with similar use. Define a tag-based service family and the tags you want to use for mapping. In addition to tag categories, you can also define tag values to narrow the criteria for the service candidates created by Service Mapping. Use these examples for guidance.

For detailed information about defining service families in tag-based mapping, see Prepare for mapping application services based on tags.

Note: Starting with Service Mapping Plus version 1.16.3, take advantage of the Tag-based Service Mapping workspace to efficiently map you application services. For more information, see Tag-based mapping in the Service Mapping Workspace and Tag-based discovery for the Service Mapping Workspace.

Examples of defining tag-based mapping criteria for service families

Definitions for service family	Category	Key	Value (optional)	Result
Example 1: one tag category, no values	Environment	Env	Not defined	Service Mapping creates a service candidate for each tag value for the Env key. Each service candidate contains only CIs with the same value for the Env key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Examples of defining tag-based mapping criteria for service families (continued)

Definitions for service family	Category	Key	Value (optional)	Result
Example 2: one tag category, one value	Environment	Env	Production	Service Mapping creates the service candidate containing only CIs that have the Env tag key with the Production tag value.
Example 3: one tag category, two values	Environment	Env	Staging,Production	Service Mapping creates two service candidates: One containing CIs that have the Production tag value; another containing CIs that have the Staging tag value.
Example 4: two tag categories, no values	Environment Application	Env App	Not defined	Service Mapping creates a service candidate for each combination of tag values for all defined tag keys.
Example 5: two tag categories, two values for each category	Environment Application	Env App	Production,Staging HR,Finance	Service Mapping creates four service candidates: production::hr, production::finance, staging::hr, staging::finance.

Traffic-based connections list

Refer to this chart for information about traffic-based connections when you remove CIs from an application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Traffic Based Connections List (when discovery based on Predictive Intelligence is disabled)

Field	Description
IP	The IP address of the application connected to the selected CI.
Port	The port on the selected CI that is used to communicate to the other application.
Process	The ID of the process in the selected CI.
Already on map	• **Yes** — if this connection shows on the map. • **No** — if this connection is not part of the service instance and not on the map.
System decision	The setting defines if Service Mapping keeps the discovered traffic-based connection or removes it. The value comes from the algorithm that Service Mapping uses.
User decision	(Optional) The setting overrides the System decision setting, which defines if Service Mapping keeps the discovered traffic-based connection or removes it. For example, if the System decision setting for a connection is Keep, and you want to remove this connection, select `Remove`.

Map application service form

Use this reference to define the name, description, owner, and group for a new application service created through service mapping.

Field	Description
Name	Pre-populated, unique name suggested for the new application service. Use this name. enter a unique name, or select **Other name suggestions** for additional options. For more information, see Name suggestions for application service candidates.
Number	Pre-populated unique ID for the application service.
Description	Description of the new application service.
Service owner	The user who approves of or rejects any change to the mapped application service.
Service group	A specific application services group to which you want to add this service. The default

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	All setting means the application service is available to all groups.

Discovery patterns used by ITOM Visibility

Service Mapping and Discovery use patterns in their discovery process that cover most industry standard network devices and applications. You can customize these patterns and create new ones.

ServiceNow applications refer to devices and applications that comprise a service instance as configuration items (CIs).

What discovery patterns are

A pattern is a sequence of commands designed to detect attributes of a CI and its outbound connections. Service Mapping and Discovery share a set of preconfigured patterns that cover most of the commonly used devices and applications. Patterns can be of the infrastructure or application type. Infrastructure patterns are used only by Discovery for creating lists of devices. Application patterns serve both Service Mapping and Discovery, which use the same application patterns for their purposes. For example, Discovery runs the horizontal discovery with the Apache Web Server pattern to find and list all Apache Web Servers in your organization. Service Mapping runs the top-down discovery using the same pattern to discover a specific Apache Web Server and place it on a service instance map.

Note: Currently, pattern operations do not support multi-languages. If values returned from pattern operations are not in English, the returned data cannot be parsed properly and the pattern discovery will fail.

Pattern usage by Service Mapping and Discovery

Product	Pattern type	Result
Discovery	Infrastructure pattern	Inventory list of devices
	Application pattern	Inventory list of applications
Service Mapping	Application pattern	Service instance map

For discovering devices that act as hosts for applications, Service Mapping relies on Discovery. As part of the top-down discovery process, Service Mapping triggers Discovery to perform its horizontal discovery behind the scenes. Service Mapping then uses information on hosts provided by the horizontal discovery to create its service instance maps.

Patterns of all types are stored in the Discovery Patterns [sa_pattern] table.

Discovery uses a combination of probes and patterns. For more information, see Horizontal discovery process flow with probes and sensors.

Correlation between pattern and CI type

Patterns are assigned to the CI types that they serve to discover. If necessary, you can assign more than one CI type per pattern. In that case, you define one main CI type and multiple related CI types. For example, a pattern for discovering BIG-IP Global Traffic Manager (GTM) F5 has F5 BigIP GTM as its main CI type. It also has related CI types for the Domain Name System (DNS) name, network adapter, and other components.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For the top-down discovery performed by Service Mapping, each application pattern serves to discover only the main CI type.

Main CI type discovery during top-down discovery

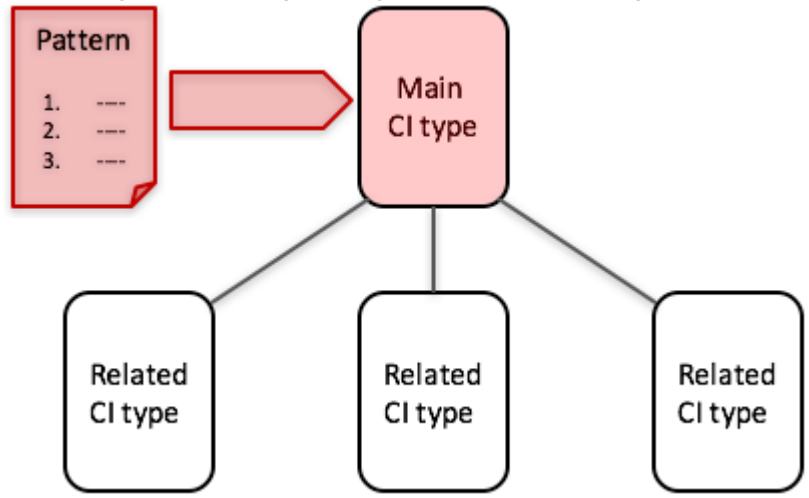

However, Service Mapping usually uses more than one pattern to discover the same CI type, because a CI type can use different protocols, operating systems, entry points, and so on.

Unlike top-down discovery, the process of horizontal discovery uses each pattern to discover a main CI type with all related CI types.

Discovery of main and related CI types during horizontal discovery

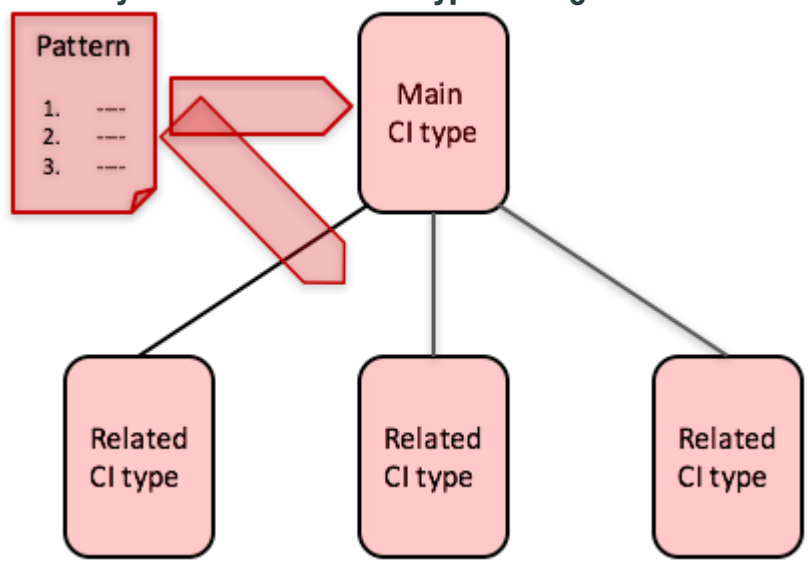

Why install patterns from the ServiceNow Store

ServiceNow releases new discovery patterns on the ServiceNow Store on a monthly basis to ensure that your organization can discover the latest industry-standard devices and applications. Major ServiceNow versions incorporate patterns previously released on the ServiceNow Store.

ServiceNow releases all discovery patterns using the following applications:

Discovery and Service Mapping Patterns (sn_itom_pattern)

This application provides the latest versions of discovery patterns the original version of which ServiceNow released on ServiceNow Store.

Visibility Content (sn_pattern_design)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This application supplies the updated version of the patterns that were part of the family releases up until Tokyo.

You can install the latest available versions of the pattern applications from the ServiceNow Store. Alternatively, you can install these applications as plugins on your ServiceNow instance.

Why customize patterns

You can customize patterns in the following cases:

- If your organization uses proprietary devices and applications, create patterns for these items to enable Discovery and Service Mapping to discover them.
- If you modify key attributes of CI types that had corresponding patterns, modify the relevant patterns to reflect the change.

Pattern versions

When you customize a pattern, you actually create a copy of the original preconfigured pattern. While Service Mapping or Discovery use the customized version, the original version isn't deleted. When you download an update of the pattern from the ServiceNow Store, the original pattern is updated, not the customized copy of it.

Exclusion of a customized pattern from an upgrade

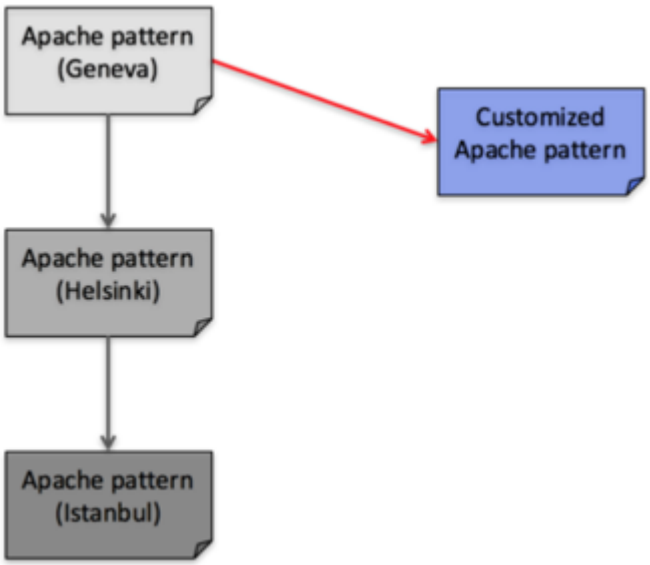

If at some point you want to abandon the customized pattern and start using the updated original pattern, you can revert to the original pattern as described in Choose the pattern version.

Patterns users

The following user roles have access to patterns or pattern-related modules and can perform various actions. Note that customizing patterns requires basic knowledge of programming.

Patterns users and access

User	Description
Discovery admin	Can view, create, edit, and publish patterns. The role enables users to run discovery, migrate probes or CAPI to patterns, and access discovery logs and dashboards.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Patterns users and access (continued)

User	Description
PD user	Has read-only access to **Discovery Pattern Log**.
PD admin	Can view, create, edit, and publish patterns.
PDE viewer	Starting with Pattern Designer Enhancements version 3.9.0, users can view **Command Validation Tasks**, **Command Validation Tasks Results**, and **Command List**. The pde_viewer can view the **Command Validation Tool** modules and related tables, but doesn't have permissions to modify or edit them. The pde_viewer role can view the following tables only: • Command List [pd_command_list] • Command Validation Task [pd_command_validation] • Command Validation Task Results [pd_command_validation_results] • Pattern Shared Library Mapping [pd_pattern_to_shared_library_mapping] • Temporary Variable Mappings [pd_temp_variable_value_mapping]
PD MID	Not assigned to a user directly but to the MID Server record or the user under which the MID Server runs. The role enables the MID Server to interpret and run pattern-based probes.
Discovery operator	Can run or cancel existing discovery schedules and view discovery-related data, including logs, ECC queue payloads, device history, and MID Server logs. Can't create, edit, or delete records. Not supported for Discovery Admin Workspace.
MID Server	Can grant the MID Server access to the instance.

Patterns for instances using domain separation

In instances that use domain separation, patterns might be domain-specific, covering only the domains that you created them for, or global, applying to all domains.

Patterns belong to domains. By default, all preconfigured patterns are assigned to the global domain and apply to all domains of all levels.

You can create patterns for specific domains. In that case, the new pattern is used only for this domain and doesn’t exist in any other domains. If you customize an existing pattern in the global domain and assign it to a specific domain, you create a copy of the global pattern. The original global pattern is still used for all other domains except for the domain that has the customized version of this pattern. Likewise, if you customize the pattern belonging to the global domain, the change affects all domains except for the one that uses a customized copy of this pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Global or domain-specific patterns

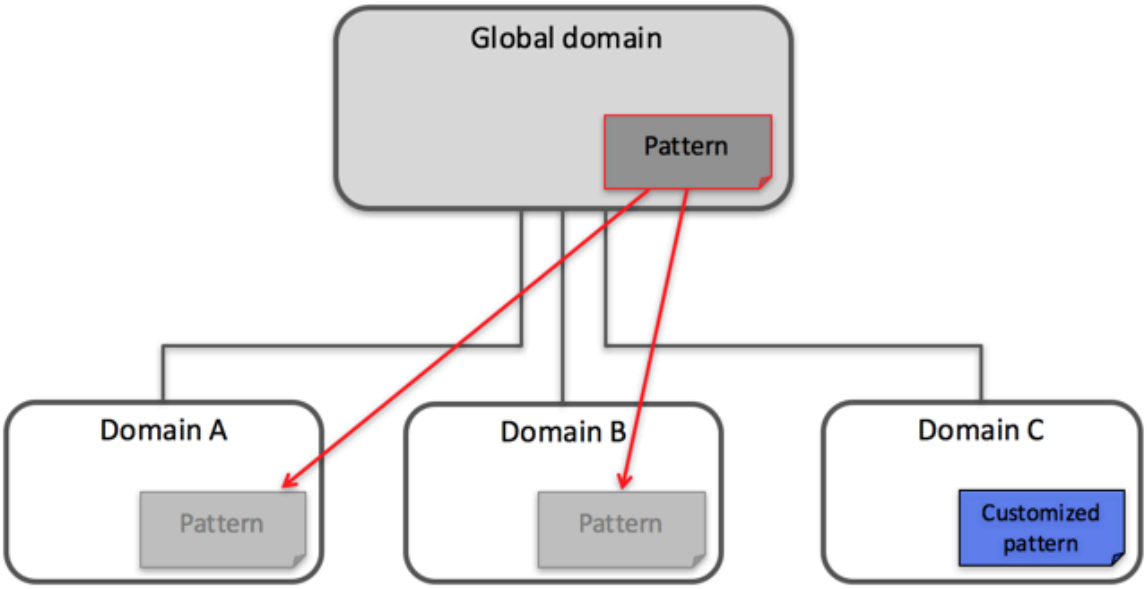

Pattern creation or modification flow

Typically, you maintain two ServiceNow instances in your organization: for production and for development. Create or modify patterns, test them, and verify results in the development instance. When you are satisfied with the discovery results, export relevant patterns from the development instance to create an update set. Then you retrieve and commit the update set in your production instance.

If you're creating a pattern for applications and devices that aren't supported in the ITOM Visibility global content application, start from creating CI types for them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern creation flow

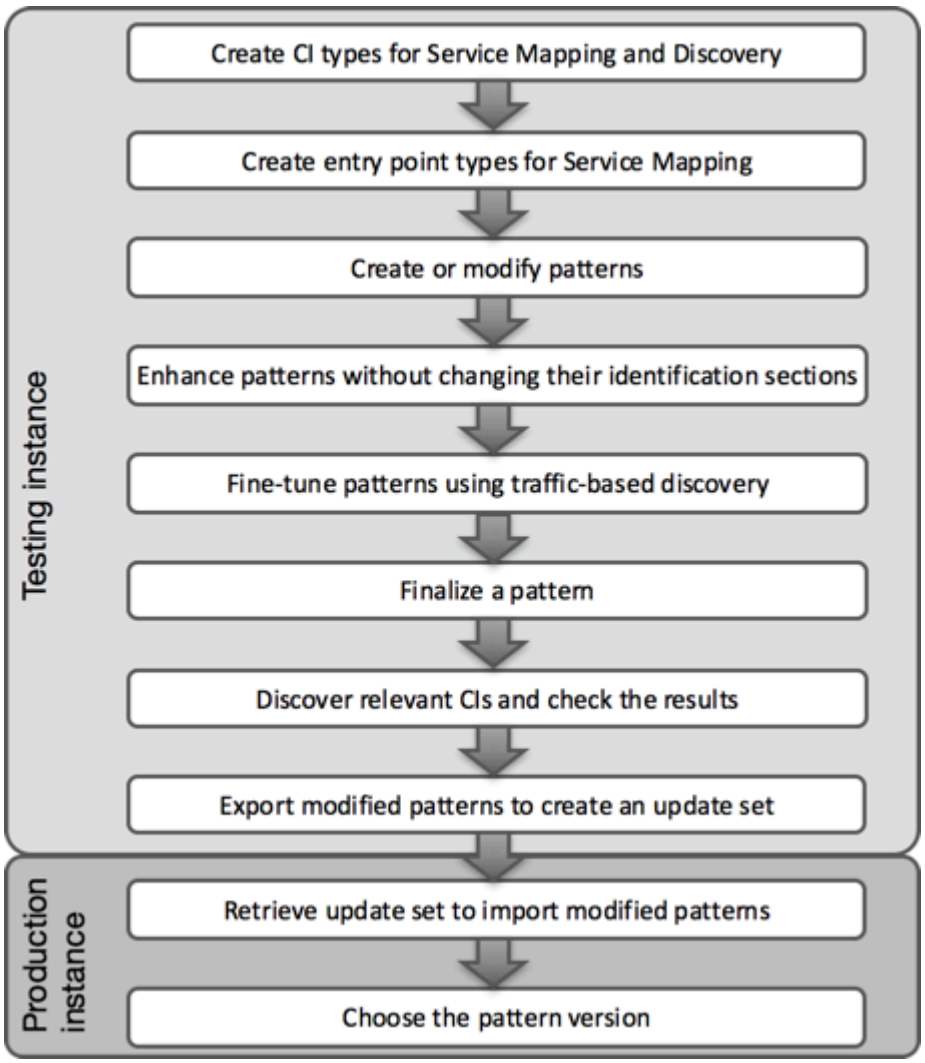

Related topics

- Detailed information on products discovered by ITOM Visibility
- Class diagrams and class attributes for key CMDB classes
- Description of key CMDB tables in the base system
- CMDB CI Class Models: ServiceNow Store app that adds class models that extend the base CMDB class hierarchy
- Information about the various options for populating the CMDB
- ServiceNow Store apps that provide pre-defined Service Graph connectors for importing and integrating common third-party data into CMDB classes
- KB0747679: Pre/Post Processing Scripts for patterns
- Application patterns for the Agent Client Collector
- Certificate Inventory and Management patterns

Available on-premise discovery patterns

ITOM Visibility comes with an extensive library of patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow releases new discovery patterns on the ServiceNow Store on a monthly basis to ensure that your organization can discover the latest industry-standard devices and applications. Major ServiceNow versions incorporate patterns previously released on the ServiceNow Store.

For a complete list of devices and applications that ITOM Visibility can discover, see Detailed information on products discovered by ITOM Visibility.

Apache Cassandra database discovery

ServiceNow Discovery finds the Apache and the DataStax Cassandra database on UNIX using the Cassandra Distributed DB pattern. Discovering some of these resources requires installing the Discovery and Service Mapping Patterns from the ServiceNow® Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Provide credentials for discovering the host

Identify the servers hosting a Cassandra database and provide credentials necessary for horizontal discovery of these servers. For more information, see Credentials required for host discovery.

Verify CMDB CI Class Models store app is up to date

Ensure that your instance supports the Cassandra CI class models by deploying the 1.14.0 version of the CMDB CI Class Models store app from the ServiceNow Store.

Verify permissions for Cassandra commands for UNIX

For the list of Cassandra commands, see the Cassandra (for UNIX) section in Service Mapping commands requiring a privileged user.

Verify read access to cassandra.yaml

Possible locations of the cassandra.yaml file:

- `$config_directory + "/cassandra.yaml"`
- `$process.environmentVariables.CASSANDRA_CONF.value + "/cassandra.yaml"`
- `$install_directory + "/conf/cassandra.yaml"`

Data collected by Discovery for Cassandra database

Discovery populates the data in the CMDB when running the Cassandra Distributed DB pattern.

Cassandra Cluster [cmdb_ci_cassandra_cluster]

Field	Description
Name [name]	Name of the cluster as well serves as a display name

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cassandra Cluster Node [cmdb_ci_cassandra_cluster_node]

Field	Description
Cluster [cluster]	Reference of the Cassandra cluster CI in which this node is part of.
Cluster Name [cluster_name]	Name of the cluster.
IP Address [ip_address]	IP Address of the cluster node (host machine).
Name [name]	Display name in the following format: cluster_name@hostname.
Operational status [operational_status]	Operational status of the cluster node.
Server [server]	Reference of server (host machine) CI

Cassandra Instance [cmdb_ci_cassandra_instance]

Field	Description
Cluster Name [cluster_name]	Name of the cluster.
Configuration directory [config_directory]	Directory path where configuration files are available.
Configuration file [config_file]	File path of base init configuration file where other configuration-related paths can be configured.
Data Center [data_center]	Data center in which the cluster node is part of a cluster.
Installation directory [install_directory]	The installation path.
IP Address [ip_address]	IP address of the host machine where the process is running.
Name [name]	Display name which is the same as the <uid>.
PID [pid]	Process ID running on the host machine.
Rack [rack]	The rack information in which the cluster node is part of.
Description [short_description]	The short description of the instance.
TCP port(s) [tcp_port]	List of TCP ports configured as in “cassandra.yaml” configuration file (Format: “:port1:port2:port3:”)
UID [uid]	The unique ID to identify the instance (Format: <cluster_name>__<data_center>__<rack>@hostname).
Vendor [vendor]	Vendor name for reference from “core_company” table if already available. If it is not available, a new entry is logically created with the name “Apache” or “DataStax”.
Workload Running [workload_running]	The name of the workload running on the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cassandra Keyspace [cmdb_ci_cassandra_keyspace]

Field	Description
Cluster [cluster]	The reference of Cassandra cluster CI in which the keyspace is part of.
Cluster Name [cluster_name]	Name of the cluster in which the keyspace is part of.
Keyspace Name [keyspace_name]	Name of the keyspace.
Name [name]	The display name in the following format: <keyspace_name>@<cluster_name>.

Dependency Views showing discovered Apache Cassandra deployment

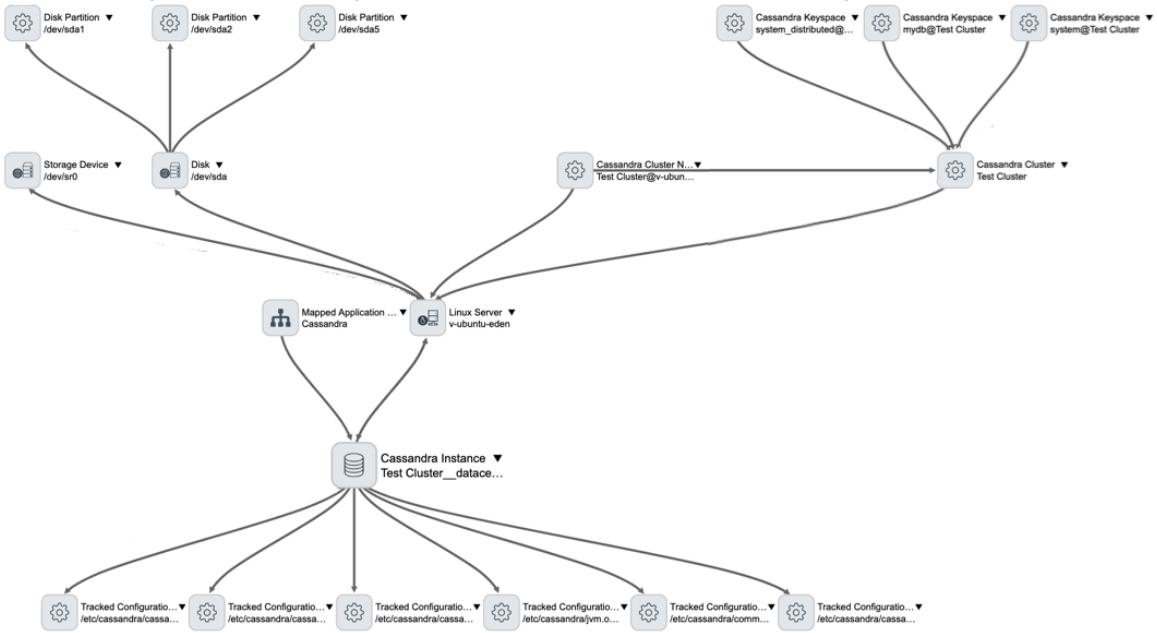

CI relationships

Discovery creates these relationships to support the Apache Cassandra database discovery.

CI	Relationship	CI
Cassandra Instance [cmdb_ci_cassandra_instance]	Runs on::Runs	Server [cmdb_ci_server]
Cassandra Cluster Node [cmdb_ci_cassandra_cluster_node]	Hosted on::Hosts	Server [cmdb_ci_server]
	Cluster of::Cluster	Cassandra Cluster [cmdb_ci_cassandra_cluster]
Cassandra Keyspace [cmdb_ci_cassandra_keyspace]	Hosted on::Hosts	Cassandra Cluster [cmdb_ci_cassandra_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data discovered by Service Mapping during top-down discovery

Service Mapping performs the top-down discovery of Cassandra, which typically runs as a distributed database on a cluster of server nodes. During top-down discovery, Service Mapping discovers instances running on nodes and displays them on the application service map.

In addition to the tables and fields discovered during horizontal discovery, the top-down discovery also discovers the fields listed in the following table.

Cassandra Instance [cmdb_ci_cassandra_instance]

Field	Description
Keyspaces [keyspaces]	Keyspace(s) the instance can access (populated only in top-down discovery).

For incoming connections, the instance is discovered, as shown in the following image.

Cassandra discovery top-down service map

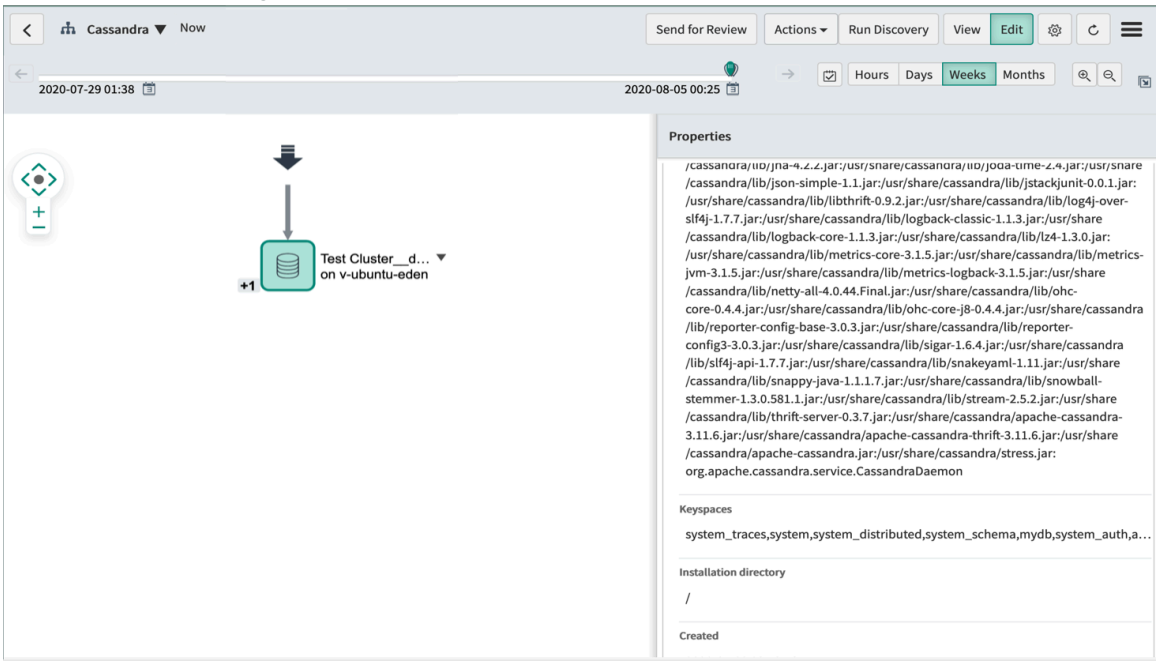

Related topics

Database Administrator (DBA) report discovery

Apache Kafka and Apache Zookeeper discovery

ServiceNow Discovery uses the Kafka and Zookeeper discovery pattern to find Kafka data built on the Zookeeper synchronization service.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

CMDB CI Class Models store app

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify the CMDB CI Class Models store app (version 1.20.0) is installed.

Discovery and Service Mapping Patterns store app

Verify the Discovery and Service Mapping Patterns store app (version 1.0.72) is installed.

Certificate Authority process

Ensure that the CA process is up and running on the host machine.

Discovery credentials

Ensure you have the credentials for running Discovery on the host machine running the CA process.

Kafka and Zookeeper process should be running on the host.

Note: Kafka Broker Discovery should run first followed by Zookeeper. In the Zookeeper Configuration file, the order of nodes forming the cluster should be consistent across all of the nodes in the cluster.

Data collected by Discovery during horizontal discovery

Resources discovered on the [cmdb_ci_appl_kafka_connect] table

Fields	Description
name	Name of the CI
node	Node where the zookeeper is installed
tier	Tier of the zookeeper
install_directory	Path of the install directory
config_file	Path of the config file
config_directory	Path of the config directory

Resources discovered on the [cmdb_ci_appl_kafka_topic] table

Fields	Description
name	Name of the CI
isr	In Sync Replica
replicas	Number of replicas of the topic
partition_count	Partition count of the topic
replication_factor	Replication factor of the topic

Resources discovered on the [cmdb_ci_appl_zoo_keeper] table

Fields	Description
name	Name of the CI
node	Node where the zookeeper is installed
tier	Tier of the zookeeper
config_file	Path of the config file

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resources discovered on the [cmdb_ci_appl_zoo_keeper] table (continued)

Fields	Description
install_directory	Path of the install

Resources discovered on the [cmdb_ci_kafka_cluster] table

Fields	Description
Cluster_id	Unique id specific to cluster

Resources discovered on the [cmdb_ci_zoo_keeper_cluster] table

Fields	Description
Cluster_id	Unique id specific to cluster

Resources discovered on the [cmdb_ci_kafka_cluster_node] table

Fields	Description
Name	Name of the Node
Clustername	Name of the cluster the node belongs to

Resources discovered on the [cmdb_ci_zoo_keeper_cluster_node] table

Fields	Description
Name	Name of the Node
clustername	Name of the cluster the node belongs to

CI Relationships

The Kafka and Zookeeper pattern does not create any CI relationships.

Apigee Edge Enterprise edition discovery

The ServiceNow Discovery application uses the APIGee pattern to find Apigee Edge Enterprise edition versions 4.x.x. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The APIGee pattern uses a built-in cache mechanism that collects data from each API service into a local cache file. Service Mapping performs top-down discovery on these cache files to find outgoing Apigee Edge connections.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Linux user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide the Linux operating system (OS) user with elevated rights for running the cat, ls, netstat, and stat commands. For more information, see Service Mapping commands requiring a privileged user.

In addition, provide this Linux user with permissions to run the following commands on the OS user home folder:

Command	Description
"date +%s"	Gets the current date.
"ls -d $HOME/APIGee_Cache/cache_folder*"	Gets the last cache folder.
"mkdir-p $HOME/APIGee_Cache/cache_folder_" + $today + ";ls -d $HOME/APIGee_Cache/cache_folder_" + $today	Creates the cache file if the cache file is empty.
"ls -d $HOME/APIGee_Cache/cache_folder*"	Verifies that the cache file has been created.
"echo `expr " + $today + " -" + $current_date+ "`"	Calculates the gap between the folder creation time and the current date.
"if [" + $diff + " -ge" + $week_seconds+ "]; then echo true; fi"	Checks if it is necessary to recreate the folder.
"ls" + $cache_folder+ "/* | sort | xargsgrep -l " + $entry_point.source_url	Checks if existing cache files contain the source URL that you enter as the entry point for the service.
"rm-rf" + $cache_folder+ ";mkdir-p $HOME/APIGee_Cache/cache_folder_" + $today + ";ls -d $HOME/APIGee_Cache/cache_folder_" + $today	Removes old cache files.
"curl -k -X GETurl-H '<ClientID>:<Token>'"	Gets the API response for the given token.
"ls" + $cache_folder+ "/* | sort | xargsgrep -l " + $entry_point.source_url	Gets the URL definition from the cache file.

Applicative credential

Configure the applicative credential as follows:

Applicative Credentials form

Field	Description
Name	The credential name, for example Apigee token.
Active	Check box for enabling this credential for discovery.
User name	The name of the user of this applicative credential. You can use any user for the credential for this pattern, since the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Applicative Credentials form (continued)

Field	Description
	information is extracted from a local cache.
Password	The actual user password of this applicative credential.
CI type	The CI type for which this credential is used: APIGee Service [cmdb_ci_appl_apigee_srv]. **Note:** ServiceNow applications refer to devices and applications that comprise a service instance as configuration items (CIs).
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
Order	Enter the order (sequence) in which the platform tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

curl command

Make sure that the curl command is installed on the Linux server hosting Apigee Edge.

Configure system properties

- To create top-down outgoing Apigee Edge connections, configure the Apigee URLs (apigee.urls) system property using the following syntax:

  ```
  01#DVL#https://<url1>.com/dashboard/v1/targetServers/dvl### 02#QA#https://<url2>.com/dashboard/v2/targetServers/qa###
  ```

  Where

  01 is the URL number.

  DVL is the environment type, like development.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

is the delimiter between the attributes. Place it exactly where it is in the example above.

is the delimiter between URLs.

- Configure the Apigee API Client ID (apigee.api.client.id) system property by entering the Apigee API Client ID in the Value field. The pattern uses this property as header for the curl command with the token:

```
curl -k -X GET " + url + " -H ' apigee.api.client.id: " + token
```

The pattern populates these two system properties using the *Add Apigee URLs and Apigee api client id* sa_pre_task_script.

Data collected by Discovery for Apigee Edge Enterprise edition

Discovery populates the data in the CMDB when running the APIGee pattern.

APIGee Service [cmdb_ci_appl_apigee_srv]

Field	Description
Name [name]	The name of the server hosting the Apigee Edge.
Version [version]	The version of the Apigee Edge.
Installation directory [install_directory]	The folder containing all the Apigee Edge libraries and executable files.
Configuration directory [config_directory]	The folder containing all the Apigee Edge configuration files.

CI relationships

The APIGee pattern does not create any CI relationships.

Data collected by Service Mapping during top-down discovery

The APIGee pattern identifies connections from the Apigee Edge to other CIs based on the URL you configure in the connection section of the pattern.

Avi Vantage load balancer discovery

The ServiceNow Discovery and Service Mapping applications use the Avi load balancer discovery patterns to find Avi Vantage load balancer components. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses the Avi Load Balancer Controller pattern and the Avi Load Balancer Controller - Session-based pattern to run horizontal discovery. Both patterns use set of REST APIs to find the load balancers. They differ only in the type of authentication used when running the API queries.

For top-down discovery, Service Mapping uses the Avi Load Balancer VIP pattern.

The patterns support Avi infrastructure residing on AWS and VMware. They relate the discovered CIs to the infrastructure only when run on VMware . The patterns also support Global Server Load Balancing (GSLB).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can use these patterns on the ServiceNow AI Platform® using the release Madrid patch 5, New York or later.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Ensure that your instance supports the Avi Load Balancer CI class models by deploying the latest version of the CMDB CI Class Models store app. Install the application from the ServiceNow Store.
- For the Avi Load Balancer Controller pattern:
  - In the Avi Controller, ensure that Basic Authentication is allowed.
  - Configure the following credentials:
    1. Navigate to **Discovery** > **Credentials**
    2. Click **New**.
    3. Click **Basic Auth Credentials**, fill out the form, and then click **Submit**.

Credentials form

Field	Description
Name	Credential name.
Active	Option for enabling this credential for discovery. Select this check box to enable discovery.
Applies to	Select **All MID servers**.
Order	Order in which the platform tries this credential as it attempts to log on to devices. A smaller number indicates that the credential appears higher in the list. Establish the credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number, or none, the instance tries the credentials in a random order. The default value is 100.
User name	Name of the user for performing Avi Load Balancer discovery.
Password	Password for this user.

    4. Run a Quick Discovery on the IP address of the Avi Controller Cluster.
- For the Avi Load Balancer Controller - Session-based pattern:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Create a Basic Authentication credential with an alias:
  1. Navigate to **Connections & Credentials** > **Connection & Credential Aliases**.
  2. Click **New**.
  3. Enter a unique name for the alias and select **Credential** for the alias type.
  4. Click **Submit**.
  5. Open the new alias.
  6. On the Credentials tab, click **New**.
  7. Select **Basic Auth Credentials** for the credential type and fill out the form.

     Credentials form

Field	Description
Name	Credential name.
User name	Name of the user for performing Avi Load Balancer discovery.
Password	Password for this user.

  8. Click **Update**.
- Create a serverless discovery schedule.

  Then do the following:

  1. In the scheduler, click the **Serverless Execution Patterns** tab.
  2. Click **New**.
  3. Fill out the form and then click **Submit**.

Field	Description
Name	Provide a descriptive name.
Pattern	Select the Avi Load Balancer Controller - Session-based pattern.

  4. Open the record created in the previous steps, and configure these parameters:

Parameter	Description
IP address	The Avi IP address.
Credential alias	The name of the credential alias created above.

- For the Avi Load Balancer Controller pattern and the Avi Load Balancer Controller - Session-based pattern, provide a user with permissions to send the following API queries.
  - /api/cluster
  - /api/serviceengine
  - /api/pool

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- /api/vimgrvmruntime?controller_vm=true
- /api/cloud
- /api/virtualservice
- /api/serviceenginegroup
- /api/gslbservice
- /api/tenant

- For the Avi Load Balancer VIP pattern:
  - Create an service instance .
  - Before using the Avi Load Balancer VIP pattern, use one of the horizontal Avi load balancer discovery patterns to find the Avi infrastructure.

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

Table and field	Description
Avi Controller [cmdb_ci_avi_controller]	
Name [name]	Name of the Avi Controller.
Fully qualified domain name [fqdn]	Hostname of the Virtual Machine (VM).
IP address [ip_address]	IP address of the FQDN.
Avi Service Engine [cmdb_ci_avi_service_engine]	
Name [name]	Name of the Avi Service Engine.
IP address [ip_address]	IP address of the Avi Service Engine.
Version [version]	Version of the Avi Service Engine.
Object ID [object_id]	Unique identifier of the Avi Service Engine.
Avi Controller Cluster [cmdb_ci_avi_controller_cluster]	
Cluster ID [cluster_id]	Unique identifier of the Avi Controller Cluster.
Name [name]	Name of the Avi Controller Cluster.
IP address [ip_address]	IP address of the Avi Controller Cluster host.
Load Balancer Pool [cmdb_ci_lb_pool]	
Name [name]	Name of the Load Balancer Pool.
Server ID [server_id]	Unique identifier of the Load Balancer Pool.
Load Balancing Method [load_balancing_method]	Method that Avi Vantage uses to make its load balancing decision. The pattern translates the algorithms as follows: • "LB_ALGORITHM_LEAST_CONNECTIONS": "least-connections" • "LB_ALGORITHM_ROUND_ROBIN": "round-robin" • "LB_ALGORITHM_FASTEST_RESPONSE": "fastest-response"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
	• "LB_ALGORITHM_CONSISTENT_HASH": "consistent-hash" • "LB_ALGORITHM_LEAST_LOAD": "least-load" • "LB_ALGORITHM_FEWEST_SERVERS": "fewest-servers" • "LB_ALGORITHM_RANDOM": "random" • "LB_ALGORITHM_FEWEST_TASKS": "fewest-tasks" • "LB_ALGORITHM_NEAREST_SERVER": "nearest-server"
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	
Name [name]	Name of the Pool Member @ name of the Pool.
IP address [ip_address]	IP address of the Pool Member host.
Load Balancer Service [cmdb_ci_lb_service]	
Name [name]	Name of the Virtual IP (VIP).
IP address [ip_address]	IP address of the VIP.
Object ID [object_id]	Unique identifier of the VIP.
Port [port]	Port of the host VM.
Avi Tenant [cmdb_ci_avi_tenant]	
name	Name of the Tenant.
object_id	The object_id of the Tenant.
tenant_url	URL to the Tenant.

The Dependency Views map shows discovered load balancer CIs and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

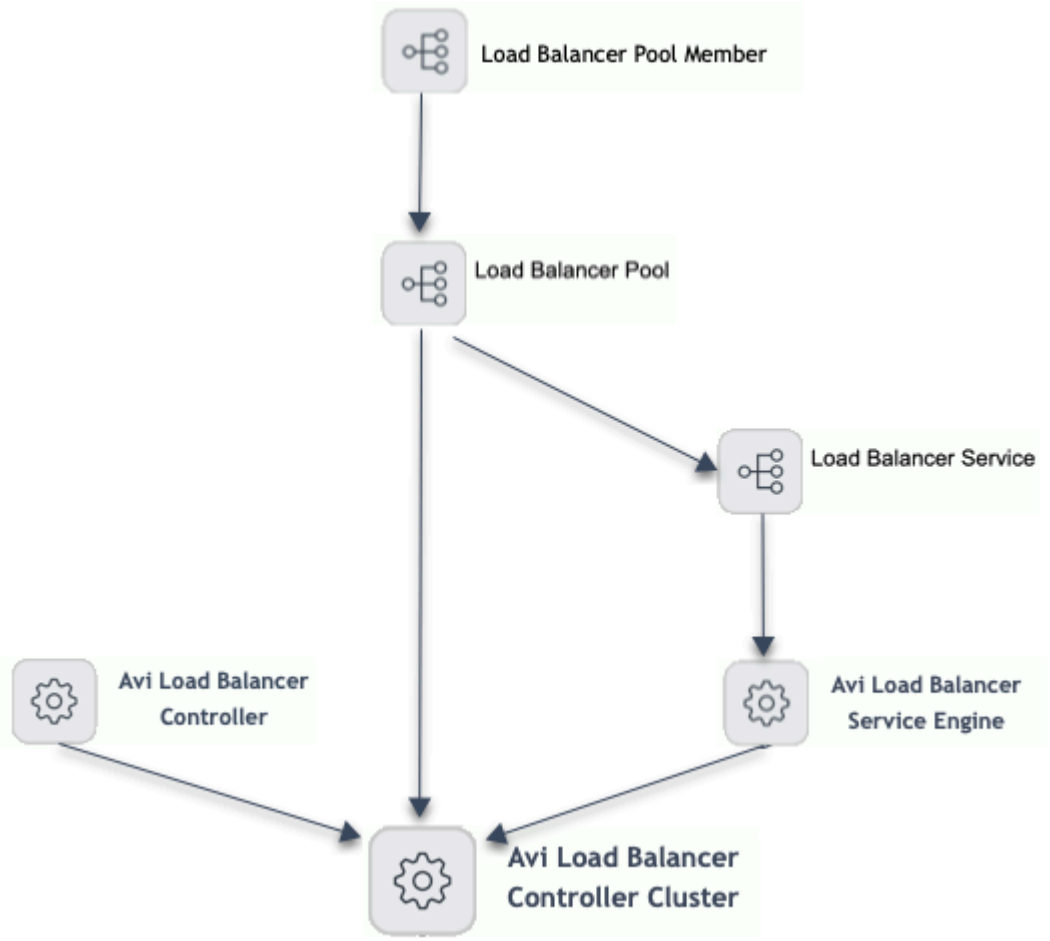

The following image shows the results of top-down discovery using the Avi Load Balancer VIP pattern.

Top-down discovery using the Avi Load Balancer VIP pattern

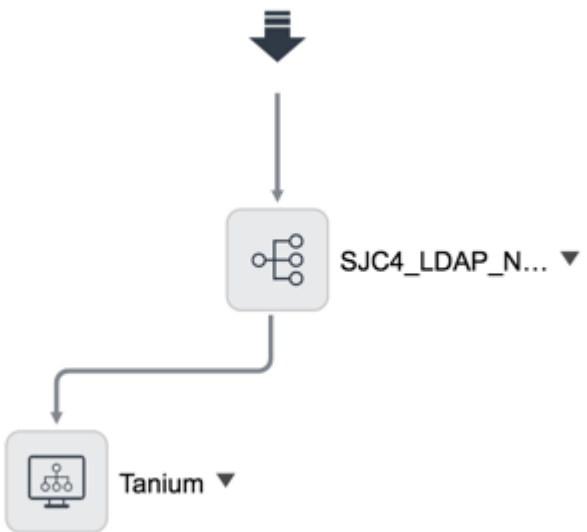

CI relationships

These relationships are created to support Avi Load Balancer discovery:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Avi Controller Cluster [cmdb_ci_avi_controller_cluster]	Cluster of::Cluster	Avi Controller [cmdb_ci_avi_controller]
Avi Tenant [cmdb_ci_avi_tenant]	Hosted on::Hosts	Avi Controller Cluster [cmdb_ci_avi_controller_cluster]
	Contains::Contained by	Load Balancer Service [cmdb_ci_lb_service]
Load Balancer Pool [cmdb_ci_lb_pool]	Used by::Uses	Load Balancer Service [cmdb_ci_lb_service]
	Hosted on::Hosts	Avi Controller Cluster [cmdb_ci_avi_controller_cluster]
	Owns::Owned by	Load Balancer Pool Member [cmdb_ci_lb_pool_member]
Avi Service Engine [cmdb_ci_avi_service_engine]	Hosted on::Hosts	Avi Controller Cluster [cmdb_ci_avi_controller_cluster]
Avi Controller [cmdb_ci_avi_controller]	Contains::Contained by	Avi Service Engine [cmdb_ci_avi_service_engine]
Load Balancer Service [cmdb_ci_lb_service]	Used by::Uses	Avi Service Engine [cmdb_ci_avi_service_engine]
	Runs on::Runs	Avi Controller Cluster [cmdb_ci_avi_controller_cluster]
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	Reference only	Avi Controller Cluster [cmdb_ci_avi_controller_cluster]

Troubleshooting

If Avi load balancer top-down discovery fails, see KB article KB0824019 for a workaround.

If the mapping process does not proceed as you expected, follow the following suggestions.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again. **Note:** If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

Next-Generation Cisco Firewall discovery

The ServiceNow Discovery application uses the Next Generation Cisco Firewall pattern to find Cisco firewalls. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The discovery pattern uses a set of SNMP calls to find the Cisco firewalls. Discovery uses the pattern to run horizontal discovery.

You can use this pattern on the ServiceNow AI Platform® using the Madrid release or later.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Ensure that your network firewall device has SNMP access.
- On the ServiceNow instance, configure SNMP credentials. For more information, see SNMP credentials.
- Add the SNMP system OID record for the Cisco device to the ServiceNow instance. Update the following:
  - Classifier: Cisco Firewall
  - Class: Cisco Firewall Device
- Deploy the pattern as follows:
  1. Download and install Firewall extension classes from the ServiceNow Store. The app adds the new CMDB classes required for network firewall discovery.
  2. Download and install the Discovery and Service Mapping Patterns application from the ServiceNow Store.
  3. Sync the pattern with the appropriate MID Server.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Next Generation Cisco Firewall pattern.

Cisco Firewall Device [cmdb_ci_firewall_device_cisco]

Field	Description
Serial number [serial_number]	Serial number of the device.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cisco Firewall Device [cmdb_ci_firewall_device_cisco] (continued)

Field	Description
Name [name]	Administratively assigned name for this managed node. By convention, this is the node's fully qualified domain name (FQDN).
IP Address [ip_address]	IP address of the device.
Firmware version [firmware_version]	Firmware version.
Fully qualified domain name [fqdn]	FQDN of the device.
Manufacturer [manufacturer]	Device manufacturer.
Model ID [model_id]	Device model name.
Harware OS [hardware_os]	Operating system (OS) running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.
Description [short_description]	Short description of the Cisco device.
Operational status [operational_status]	Indicates if the device is in active state.
Hardware Status [hardware_status]	Detailed description of the current status of the resource.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address of the Cisco firewall.
Netmask [netmask]	Netmask of the Cisco firewall.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
IP Address [ip_address]	IP address of the network adapter.
Alias [alias]	User-assigned name for the network adapter.
Netmask [netmask]	Netmask of the network adapter.
MAC Address [mac_address]	MAC address of the network adapter.
Name [name]	Name of the network adapter.
Configuration Item [cmdb_ci]	References the Cisco Firewall Device [cmdb_ci_firewall_device_cisco] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Name of the Domain Name System (DNS).
IP Address [ip_address]	IP address of the DNS.

This Dependency Views map on the Cisco Firewall Device CI shows the Cisco Firewall Device.

CI relationships

These relationships are created to support Cisco firewall discovery.

CI	Relationship	CI
Cisco Firewall Device [cmdb_ci_firewall_device_cisco]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Cisco Firewall Device [cmdb_ci_firewall_device_cisco]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
Cisco Firewall Device [cmdb_ci_firewall_device_cisco]	Uses::Used by	Router Interface [dscy_router_interface]
IP Address [cmdb_ci_ip_address]	References	Network Adapter [cmdb_ci_network_adapter]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Network Adapter [cmdb_ci_network_adapter]	References	Cisco Firewall Device [cmdb_ci_firewall_device_cisco]
Router Interface [dscy_router_interface]	References	Cisco Firewall Device [cmdb_ci_firewall_device_cisco]
Serial Number [cmdb_serial_number]	References	Cisco Firewall Device [cmdb_ci_firewall_device_cisco]

Cisco Switch Wireless Access Point (WAP) discovery

The ServiceNow Discovery application uses the Cisco WAPs extension to find Wireless Access Points (WAPs) that are controlled by a Cisco network switch. Discovering some of these

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

This extension extends the Network Switch discovery pattern. For more information, see Network switch and router discovery.

Discovery uses the Cisco WAPs extension to run horizontal discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

For prerequisites related to network switch discovery, see Network switch and router discovery.

In addition, these prerequisites must be in place for the Cisco WAPs extension:

- Permission to query the OID: 1.3.6.1.4.1.14179.2.2.1.1.
- The client must be able to run discovery patterns.
- The client must use a Cisco network switch.
- The client must have SNMP access to the Cisco network switch that controls the WAP devices.

Data collected by Discovery during horizontal discovery

The data discovered by the Cisco WAP extension includes the following tables and fields.

Tables and fields	Description
Main CI Class [cmdb_ci_wap_network]	
Name [name]	Name assigned to the WAP. Collected from bsnAPName.
MAC Address [mac_address]	The MAC address of the WAP. Collected from bsnAPDot3MacAddress.
IP address [ip_address]	IP address of the WAP. Collected from bsnApIpAddress.
Firmware version [firmware_version]	IOS version of the IOS Cisco WAP. Collected from bsnAPIOSVersion.
Serial Number [serial_number]	WAP serial number. Collected from bsnAPSerialNumber.
Model number [model_number]	WAP model. Collected from bsnAPModel.
snmp_sys_location	User-specified location of the WAP. Collected from bsnAPLocation.
Manufacturer [manufacturer]	Manufacturer of the WAP (Cisco).

The Dependency Views map shows all discovered Cisco WAP CIs and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

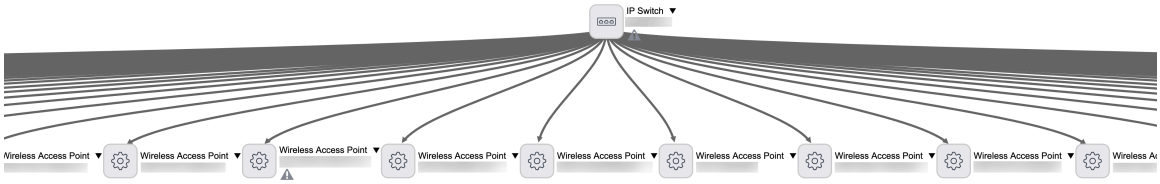

CI relationships

These relationships are created to support Cisco WAPs discovery:

CI	Relationship	CI
cmdb_ci_ip_switch	Controller for :: Controlled by	cmdb_ci_wap_network

Citrix NetScaler SDX discovery

The ServiceNow Discovery application can discover Citrix NetScaler SDX devices using the Citrix NetScaler SDX pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The classification mechanism introduced in Quebec triggers the pattern. The classification mechanism prevents the whole probe from failing if a sub probe in a multi-probe fails.

The classification mechanism includes a trigger_probe_m2m table, which defines which classification probe to trigger in relation to the relevant protocol. For this pattern, it triggers the Interactive Probe Shell over the *ssh* protocol.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the following applications are up to date

- CMDB CI Class Models
- Discovery and Service Mapping Patterns
- Visibility Content

Create SSH credentials

Create SSH credentials for the Citrix NetScaler SDX device. For more information, see SSH credentials.

Verify permissions for the following commands

Verify that the Discovery user has admin permissions and can run the following commands. The *show hardware* command requires shell access as well.

Command	Description
show vmdevice	Get virtual machine (VM) instances.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Description
show hardware	Provides a quick overview of all configurable settings for the Citrix NetScaler SDX devices.

If the Discovery user doesn't have shell access or admin permissions, verify that the user can run the following read-only commands in the NetScaler CLI instead:

Command	Description
show systemstatus	Retrieve the current system status of the SDX appliance, including serial number, build number, and product name.
show hostname	Retrieve the hostname of the SDX appliance.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Citrix NetScaler SDX pattern.

Citrix NetScaler SDX [cmdb_ci_citrix_netscaler_sdx]

Field	Description
Name [name]	The name of the Citrix NetScaler SDX device.
Serial number [serial_number]	The serial number of the Citrix NetScaler SDX device.
IP Address [ip_address]	The IP address of the Citrix NetScaler SDX device.
Object ID [object_id]	The CI identifier of the Citrix NetScaler SDX device, which includes the serial_number and ip_address.
Description [short_description]	The description of the Citrix NetScaler SDX device.
OS Version [os_version]	The version of the Citrix NetScaler SDX device.

Citrix Netscaler [cmdb_ci_lb_netscaler]

Field	Description
Name [name]	The name of the Citrix NetScaler load balancer.
Serial number [serial_number]	The serial number of the Citrix NetScaler load balancer.
IP Address [ip_address]	The IP address of the Citrix NetScaler load balancer.

On the Dependency Views map you can see all discovered Citrix NetScaler SDX resources in your organization, and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

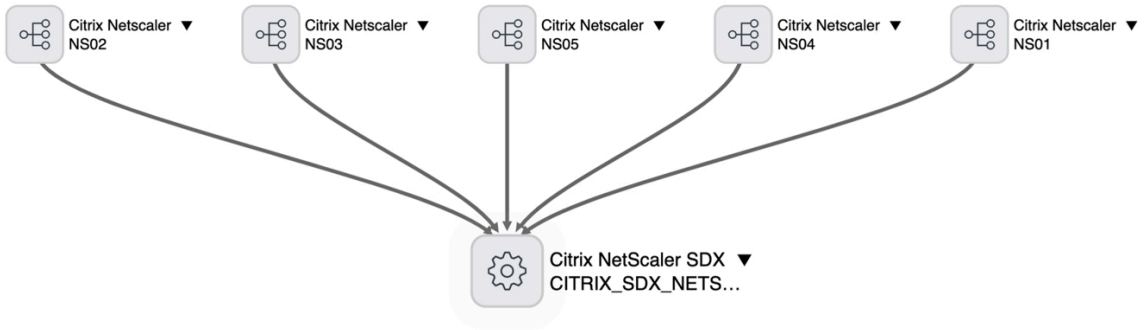

CI relationships

These relationships are created to support Citrix NetScaler SDX discovery:

CI	Relationship	CI
Citrix Netscaler [cmdb_ci_lb_netscaler]	Registered on::Has registered	Citrix NetScaler SDX [cmdb_ci_citrix_netscaler_sdx]

Citrix Xen Hyper-V Discovery

ServiceNow Discovery uses the Citrix Xen Hyper-V pattern to find and map instances. Discover these resources by installing the pattern applications from the ServiceNow Store and add Citrix Xen Hyper-V to the Linux Server Pattern's Extension Section.

Use the Xen (XE) and (XL) commands to discover Citrix Xen Hyper-V, its components, and respective attributes. Discover information regarding the Hyper-V Server, Associated Pools, VMs, Network and Storage and their relationships with new and existing CIs.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the following applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models

Verify pool admin RBAC role

This role is required for SSH access to the Linux server hosting the hypervisor and for executing all the necessary commands to extract information from the hypervisor environment. For more information on RBAC role permissions or adding Citrix users to a domain and assigning roles, see the Citrix documentation.

Verify MID Server reachability

Verify MID Server has reachability to the target device.

Create SSH credentials with access to the target Linux host

For more information, see SSH credentials

Create an extension section after the Linux Server pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create an extension section after the Linux Server pattern and associate the **Citrix Xen HyperV** pattern to be triggered after the execution of the Linux Server discovery pattern. For more information on creating extension sections, see Enhance patterns without modifying identification sections.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Citrix Xen HyperV pattern.

Hyper-V Resource Pool [cmdb_ci_hyper_v_resource_pool]

Field	Description
object_id	UUID of the Hyper-V resource pool
name	Name of the Hyper-V resource pool
default_sr	Default storage repository is where the VDIs are stored by default

Hyper-V Server [cmdb_ci_hyper_v_server]

Field	Description
object_id	UUID of the Hyper-V Server
name	Name of the Hyper-V Server in format NAME@Hyper-V Server
serial_number	Serial number of the Linux Server where the Citrix Xen Hyper-V server is running, in the format SERIAL-NUMBER_hyper_v_server in order to distinguish from the Linux Server
manufacturer	Vendor of the Hyper-V Server
model_id	Model ID of the Hyper-V Server
short_description	Description of the Hyper-V Server
os	Operating system of the Hyper-V Server
os_version	Version of the operating system of the Hyper-V Server
host_name	Host name of the Hyper-V Server
ip_address	IP address of the Hyper-V Server
virtual	Set to false, as this value is used by calculators to populate the IsVirtual flag on the dashboard (SAM).
free_memory	Available memory of the Hyper-V Server
total_memory	Total memory of the Hyper-V Server
edition	Edition of the Hyper-V Server
cpu_core_count	The number of cores per physical CPU
cpu_count	Count of the total CPUs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hyper-V Server [cmdb_ci_hyper_v_server] (continued)

Field	Description
ram	Total RAM memory in MB
cpu_speed	CPU speed in MHz
cpu_manufacturer	Manufacturer of the CPU
cpu_type	Type of the CPU
cpu_core_thread	The number of threads per CPU core

Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance]

Field	Description
object_id	UUID of the Hyper-V Server
name	Name of the Hyper-V Server
state	State of the Hyper-V instance
ip_address	IP address of the Hyper-V instance
host_id	The UUID of the host where the instance is residing
baseboard_serial	UUID of the Hyper-V instance
chassis_serial	UUID of the Hyper-V instance
bios_serial	UUID of the Hyper-V instance

Hyper-V Virtual Network [cmdb_ci_hyper_v_network]

Field	Description
object_id	UUID of the Hyper-V Virtual Network
pif_uuid	UUID of the physical NIC
name	Name of the Hyper-V Virtual Network
bridge	Type of the network setup: either bridged or not

Hyper-V Virtual Network Component [cmdb_ci_hyper_v_network_comp]

Field	Description
object_id	UUID of the Hyper-V network component
mac_address	MAC address of the Hyper-V network component
vm	UUID of the Hyper-V instance, to which the network component is associated
name	Name of the Hyper-V network component

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hyper-V Storage [cmdb_ci_hyper_v_storage]

Field	Description
object_id	LUN of the Hyper-V storage
name	Name of the Hyper-V storage
description	Description of the Hyper-V storage
type	Type of the storage (SR)
vm_object_id	UUID of the Hyper-V instance, to which the storage is associated
vm_id	UUID of the Hyper-V instance, to which the storage is associated

Hyper-V Storage Component [cmdb_ci_hyper_v_storage_comp]

Field	Description
object_id	UUID of the virtual disk image (VDI)
name	Name of the virtual disk image (VDI)
description	Description of the virtual disk image (VDI)
type	Type of the storage component (VDI)
capacity	Total capacity of the virtual disk image (VDI)
location_id	Location of the virtual disk image (VDI)
is_a_snapshot	Indication if the virtual disk image (VDI) is snapshot or not
snapshot_of	Indication of VDI snapshot of the disks
snapshot_time	Time of the snapshot creation
sr_object_id	UUID of the storage repository to which the virtual disk image (VDI) is associated
vm_object_id	UUID of the virtual machine associated with the virtual disk image (VDI)

CI Relationships

Relationships created by Citrix Xen Hyper-V Discovery

Parent	Relationship	Child
cmdb_ci_hyper_v_resource_pool	Contains :: Contained by	cmdb_ci_hyper_v_server
cmdb_ci_hyper_v_server	Runs on :: Runs	cmdb_ci_linux_server
cmdb_ci_hyper_v_network	Provides :: Provided by	cmdb_ci_hyper_v_server
cmdb_ci_hyper_v_network_comp	Provides :: Provided by	cmdb_ci_hyper_v_server
cmdb_ci_hyper_v_instance	Connected by :: Connects	cmdb_ci_hyper_v_network_comp
cmdb_ci_hyper_v_resource_pool	Reference	cmdb_ci_hyper_v_storage

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parent	Relationship	Child
cmdb_ci_hyper_v_instance	Contains :: Contained by	cmdb_ci_hyper_v_storage
cmdb_ci_hyper_v_storage_comp	Managed by :: Manages	cmdb_ci_hyper_v_storage
cmdb_ci_hyper_v_storage_comp	Reference	cmdb_ci_hyper_v_instance
cmdb_ci_hyper_v_server	Define resources for :: Gets resources from	cmdb_ci_hyper_v_resource_pool
cmdb_ci_hyper_v_instance	Registered on :: Has registered	cmdb_ci_hyper_v_server

Cloudian Storage Discovery

ServiceNow Discovery uses the Cloudian Storage discovery pattern to find Cloudian servers and related disks. Discovering these resources requires installing the Patterns application from the ServiceNow Store.

Discovery uses Cloudian Storage System patterns to find Cloudian Storage system data. Discovering these resources requires installing the Discovery and Service Mapping Patterns and the CMDB CI Class Models applications from the ServiceNow Store, as each application is dependent on the other. You can use the patterns for discovering the Cloudian Storage components on the ServiceNow AI Platform using the Paris release or later.

Cloudian is a software-defined data storage company. Cloudian provides turnkey distributed software-defined S3-compatible object storage solutions.

Cloudian Hyperstore Software : Deploy the CentOS operating system, related packages, and Hyperstore software binary files on the Hyperstore Node.

Cloudian Hyperstore Node: This is an independent hardware server in which Cloudian Hyperstore software is deployed on an industry standard server, on a VM, or it is pre-installed on a Cloudian Hyperstore appliance.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

To discover a device as a Cloudian server, the **/bin/systemctl status cloudian-s3** command must return valid responses and contain the **Cloudian HyperStore Storage System**. If this is not the case, then the device is discovered as a Linux server.

- Make sure the client is using Discovery patterns.
- Set **SSH** credentials for the Cloudian Storage Server.
- Make sure the MID server has access to the target device.
- Ensure that the Discovery user has the correct privileges to run commands.
- Install the following plugins:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- CMDB CI Class Models plugin (Version 1.31.0 or later)
- Discovery and Service Mapping Patterns plugin (Version 1.0.84 or later)
- Discovery and Service Mapping Patterns: (sn_itom_pattern) version 1.0.84 and dependent plugin **CMDB CI Class Models** (sn_cmdb_ci_class) version 1.31.0 (October content release)

Credentials

Make sure the Cloudian Hyperstore Software is deployed on either CentOS or RHEL Servers. The credentials configurations are the same as for a Linux server.

Configure the **SSH** credentials for the Cloudian Storage Server.

Running Discovery

To discover the Cloudian Storage System, create a Discovery schedule that targets the IP address of the Cloudian Storage Server. Make sure that the **SSH** credentials that you configured can access the Cloudian Server.

Data collected by Discovery during horizontal discovery

Storage Server [cmdb_ci_storage_server]

Field	Description
name	
serial_number	
ip_address	
manufacturer	
model_id	
short_description	
os_version	
os	
ram	
cpu_manufacturer	
cpu_name	
cpu_core_count	
cpu_type	

Disk [cmdb_ci_disk]

Field	Description
name	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk [cmdb_ci_disk] (continued)

Field	Description
device_id	
storage_type	
size_bytes	
manufacturer	
model_id	
computer [Reference cmdb_ci_storage_server]	

File System [cmdb_ci_file_system]

Field	Description
name	
size_bytes	
free_space_bytes	
file_system	
mount_point	
computer [Reference cmdb_ci_storage_server]	
provided_by [Reference cmdb_ci_disk]	

Network Adapter [cmdb_ci_network_adapter]

Field	Description
name	
ip_address	
netmask	
mac_address	
cmdb_ci [Reference cmdb_ci_storage_server]	

IP Address [cmdb_ci_ip_address]

Field	Description
ip_address	
netmask	
nic [Reference cmdb_ci_network_adapter]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Cloudian**probe is created in the **Unix-Classify** Multiprobe. The **Cloudian** probe runs the **/bin/systemctl status cloudian-s3** command and checks that the output contains **Cloudian HyperStore Storage System**.

Cloudian Probe

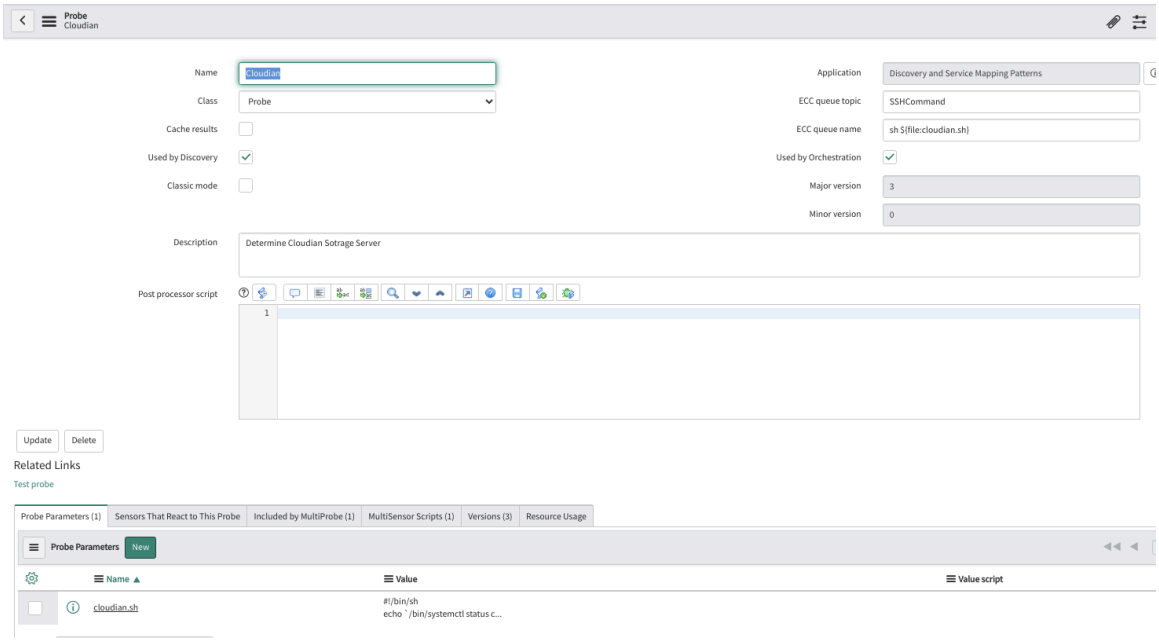

In Unix classification, the **Cloudian Server** classifier runs before the **Linux Classifier**. The Cloudian classifier creates records in the [cmdb_ci_storage_server] table. This classifier invokes the Cloudian Storage System Pattern.

Classification Criteria

Classification Criteria

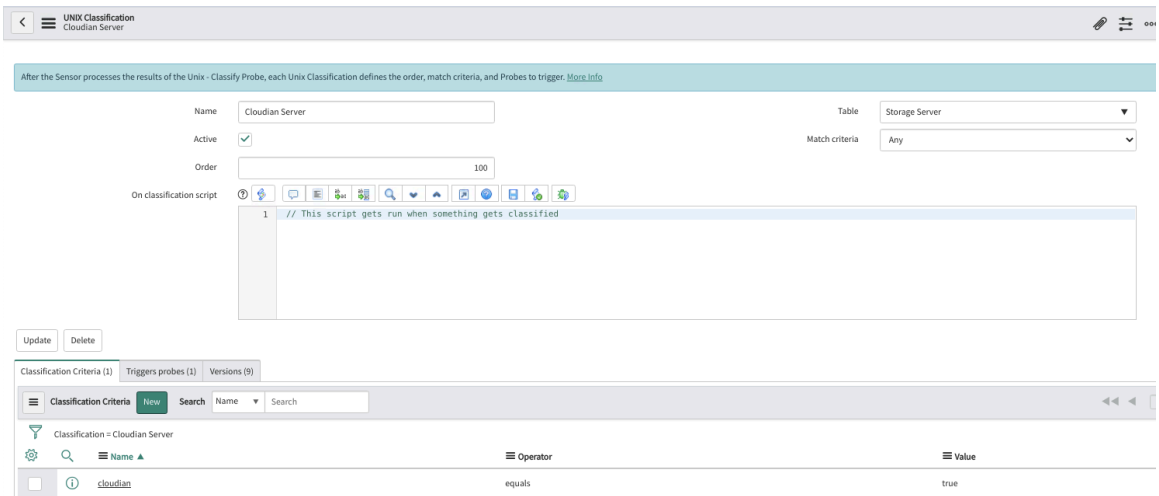

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CMDB Data Model

CMDB Data Model

CMDB Data Model

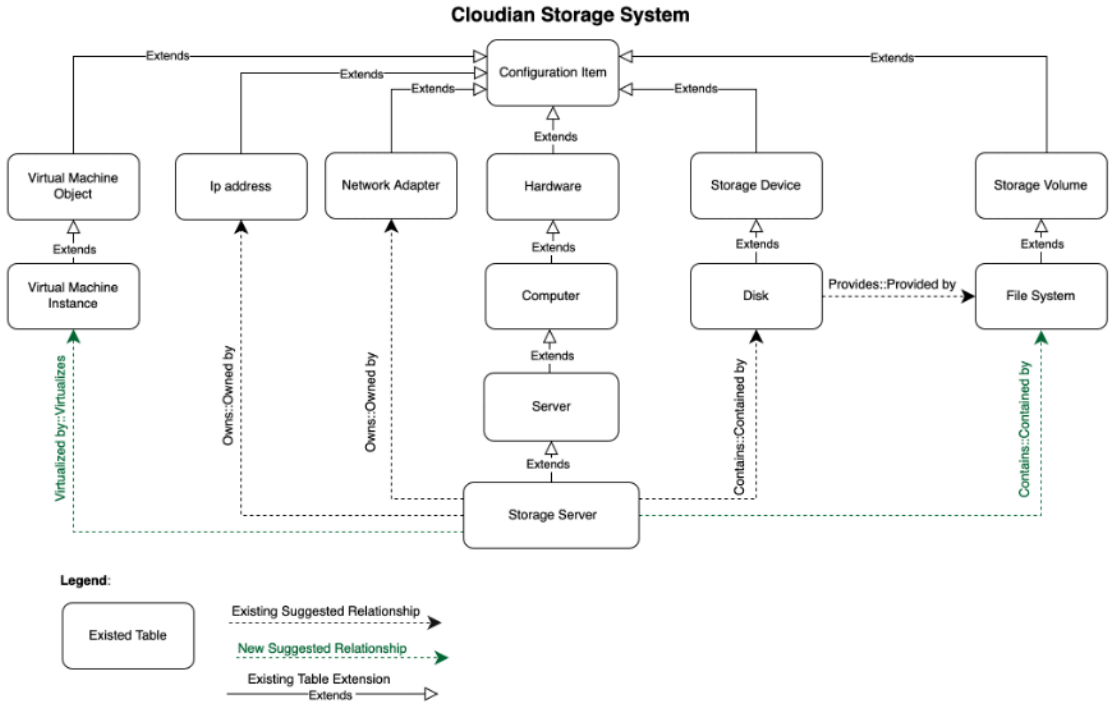

Storage Server to Network Adapter

Storage Server to Network Adapter

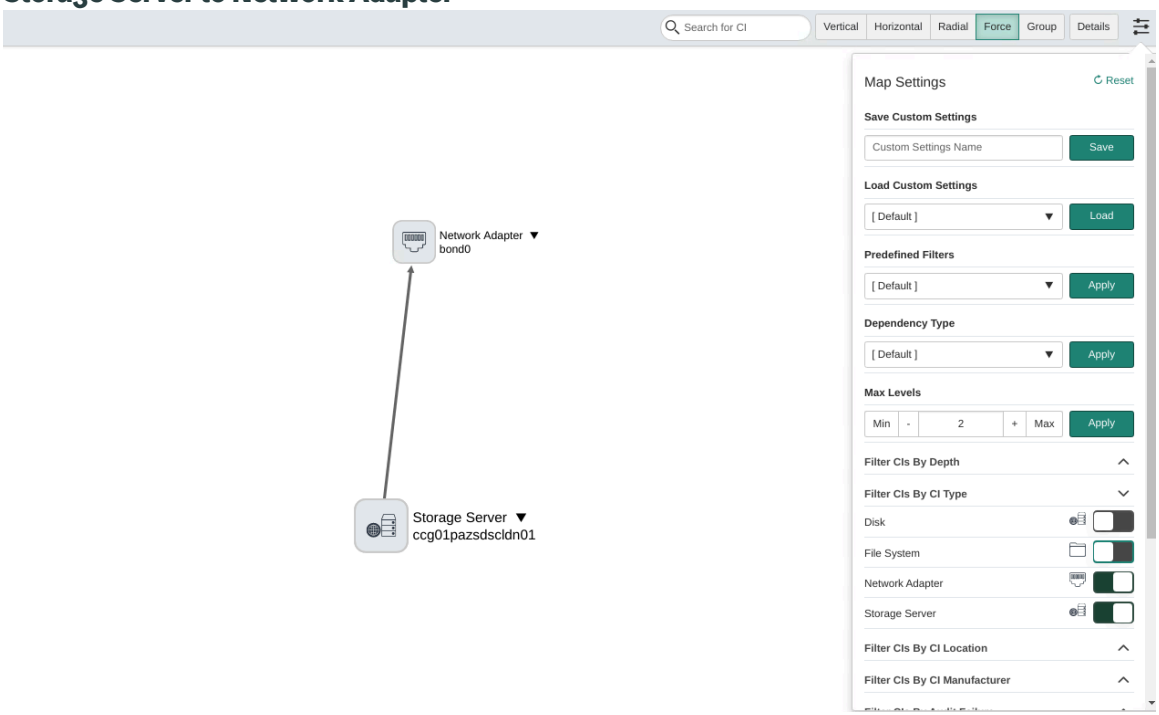

Storage Server to Disks

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server to Disks

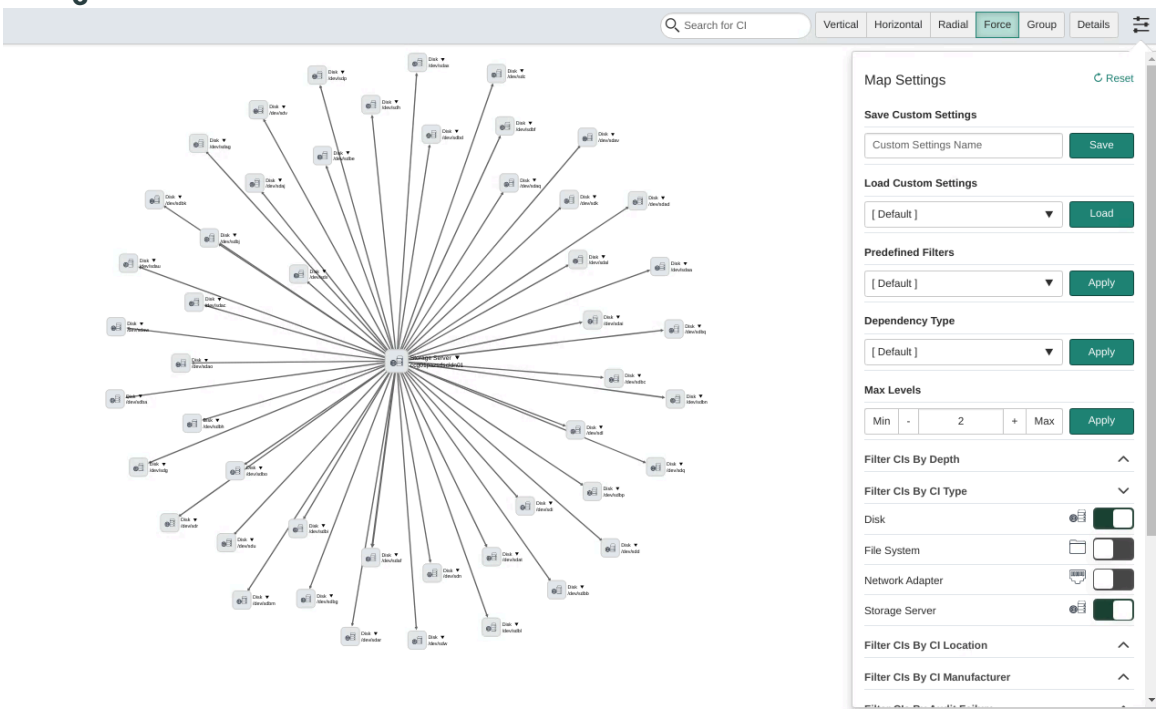

Storage Server to Disks and File Systems

Storage Server to Disks and File Systems

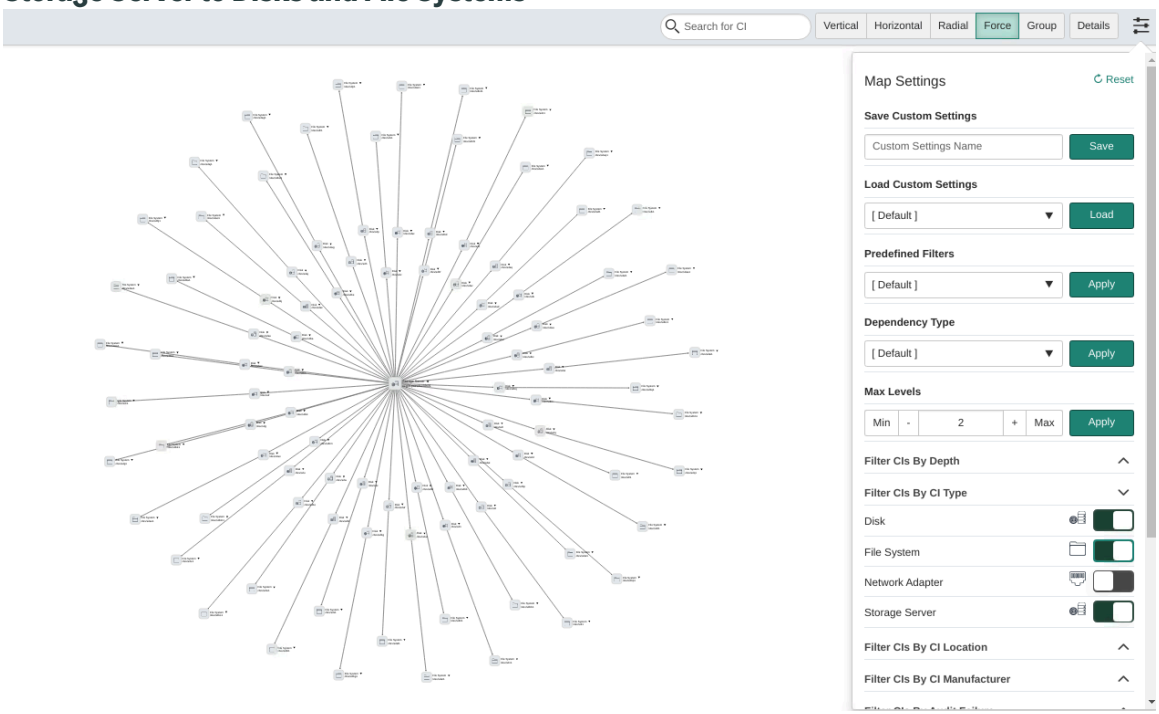

CI Relationships

Relationships created by the Cloudian Storage pattern

Parent	Relationship	Child
cmdb_ci_storage_server	Owns::Owned by	cmdb_ci_network_adapter

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parent	Relationship	Child
cmdb_ci_storage_server	Owns::Owned by	cmdb_ci_ip_address
cmdb_ci_storage_server	Contains::Contained by	cmdb_ci_disk
cmdb_ci_storage_server	Contains::Contained by	cmdb_ci_file_system
cmdb_ci_storage_server	Virtualized by::Virtualizes	cmdb_ci_vm_instance
cmdb_ci_disk	Provides::Provided by	cmdb_ci_file_system

Cohesity storage system discovery

The Discovery and Service Mapping Patterns application uses the Cohesity Storage System pattern to find clusters, nodes, and chassis for the Cohesity DataPlatform. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Cohesity develops software that are used to consolidate and simplify data management, and it includes analytics capabilities. Their software also solves the problem of mass data fragmentation with data that extends across multiple systems or cloud providers.

The company's main product, DataPlatform, is a hyper-converged software that allows businesses to consolidate a variety of workloads onto a single cloud-native platform. These workloads include backups, archives, test, development, and analytics data. It works with physical servers as well as virtual machines.

Cohesity extends backup capabilities to NoSQL workload distributed databases like MongoDB, Apache Cassandra, Couchbase, and Apache HBase. It also supports Apache Hadoop data on Hadoop Distributed File System (HDFS) datastores.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date

- CMDB CI Class Models, from version 1.37.0
- Discovery and Service Mapping Patterns, from version 1.0.91

Verify MID Server configuration

Verify the MID Server is set to Active and is reachable to the Cohesity storage cluster.

Verify API permissions

Verify you have a basic authentication account configured with proper permissions for the following API calls:

- `https://" + ${cluster_ip} + "/irisservices/api/v1/public/cluster?fetchStats=true`
- `https://" + ${cluster_ip} + "/irisservices/api/v1/nexus/cluster/status`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `https://" + ${cluster_ip} + "/irisservices/api/v1/public/nodes?fetchStats=true`
- `https://" + nodeIp + "/irisservices/api/v1/nexus/node/hardware_info`

Create an alias for the basic authentication credential

For more information, see Create an alias for a basic authentication credential for Cohesity storage system discovery.

Create a serverless discovery schedule

For more information, see Create a serverless schedule for Cohesity storage system discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Cohesity Storage System pattern.

Cohesity Cluster [cmdb_ci_cohesity_cluster]

Label	Description
IP Address [ip_address]	IP address of the cluster.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) of the cluster.

Cohesity Node [cmdb_ci_cohesity_node]

Label	Field Name
Manufacturer [manufacturer]	Manufacturer of the cluster node, which could be Cohesity or third-party.
Fully qualified domain name [fqdn]	FQDN of the node within the cluster.
Operational status [operational_status]	Operational status. Default value: Installed.
Description [short_description]	Description of the cluster node: Cohesity Backup Server.

Chassis [cmdb_ci_chassis]

Label	Field Name
Manufacturer [manufacturer]	Manufacturer of the chassis.
Description [short_description]	Description of the chassis: Cohesity Chassis.
Serial number [serial_number]	Serial number of the chassis.

IP Address [cmdb_ci_ip_address]

Label	Field Name
Name [name]	Name of the Cohesity node.
IP Address [ip_address]	IP address of the Cohesity node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Address [cmdb_ci_ip_address] (continued)

Label	Field Name
IP version [ip_version]	IP version of the Cohesity node.

CI relationships

Discovery creates these relationships to support the Cohesity storage system discovery.

CI	Relationship Type	CI
Cohesity Node [cmdb_ci_cohesity_node]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Cohesity Node [cmdb_ci_cohesity_node]	Cluster of::Cluster	Cohesity Cluster [cmdb_ci_cohesity_cluster]
Cohesity Node [cmdb_ci_cohesity_node]	Hosted on::Hosts	Chassis [cmdb_ci_chassis]

Create an alias for a basic authentication credential for Cohesity storage system discovery

Create an alias and add it to a basic authentication credential to discover Cohesity storage systems.

Before you begin

Verify you have a basic authentication account configured with proper permissions. For more information, see the Verify API permissions section in Cohesity storage system discovery.

Role required: discovery_admin

Procedure

1. Create an alias.
   a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   b. Select **New**.
   c. Enter a unique name for the alias and select **Credential** for the alias type in the **Type** field.
   d. Select **Submit**.
2. Configure a basic authentication credential for the new alias.
   a. Navigate to **Discovery** > **Credentials**.
   b. Select **New**.
   c. Select **Basic Auth Credentials**.
   d. Unlock the **Credential alias** field and select the alias you created.
   e. Configure the rest of the **Basic Auth Credentials** form fields.
   f. Select **Submit**.
   For more information about Basic Auth Credentials form fields, see Basic authentication credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Create a serverless discovery schedule. For more information, see Create a serverless schedule for Cohesity storage system discovery.

Related topics

Cohesity storage system discovery

Create a serverless schedule for Cohesity storage system discovery

Create a serverless discovery schedule to discover Cohesity storage systems.

Before you begin

- Verify the MID Server is set to Active and is reachable to Cohesity storage cluster.
- Create an alias for the basic authentication credential. For more information, see Create an alias for a basic authentication credential for Cohesity storage system discovery.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule New record form

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type, which should be Serverless.
MID server	Name of the MID Server to use for this schedule.
Active	Enables this schedule for discovery.

   c. Select **Submit**.
3. Create the execution pattern.
   a. In the Discovery Schedules page, select the record you created.
   b. In the **Serverless Execution Patterns** tab, select **New**.
   c. On the form, fill in the fields.

Serverless Execution Pattern New record form

Field	Description
Name	Descriptive name for this record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Pattern	Pattern to be used for this schedule, which should be the Cohesity Storage System pattern.
Proxy Host	Fully qualified domain name of the machine on which a proxy server is installed, when applicable. This field should be left empty.
Active	Enables this schedule for discovery.

d. Select **Submit**.

4. Set the pattern launcher parameters.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Parameter	Value
cluster_ip	The Cohesity cluster URL or IP address.
credentialAlias	Name of the credential alias you created.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

Cohesity storage system discovery

ColdFusion discovery

The ServiceNow Discovery application finds Adobe ColdFusion servers and the instances of ColdFusion applications running on them. Only the 2016 version of ColdFusion is supported. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Components

Discovery and Service Mapping use the *ColdFusion Application* and *ColdFusion Application Server* patterns to run horizontal and top-down discovery. The *ColdFusion Application Server* pattern is triggered from the ColdFusion Application Server process classifier for horizontal discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Credentials

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure the following credentials, depending on what type of host machine the ColdFusion server is installed on:

- Windows credentials
- SSH credentials (for Linux machines)

File access

The pattern must be able to access these ColdFusion files:

File	Description
`neo-runtime.xml`	This file provides the mapping between the URL and the working directory path of the application.
`application.cfc` and `application.cfm`	These files provide application names.
`neo-datasource.xml`	This file provides the name of the datasource that is used in the configuration of the ColdFusion application. The datasource reference is configured in the `application.cfc` and `application.cfm` files. **Note:** The datasource is used for top-down discovery only.

Entry point

For top-down discovery, use the full URL or IP address of the web application as the HTTP(S) endpoint. For example: `https://my-website-on-coldfusion/path` or `https:10.120.255.2555:8500/path`

Data collected during horizontal discovery

Table and Field	Description
Cold Fusion Server [cmdb_ci_coldfusion_server]	
Installation directory [installation_directory]	The name and version of the ColdFusion server, and the directory where it is installed.
Name [name]	
Version [version]	
Installation directory [installation_directory]	The name of the application and the directory where it is installed.
Name [name]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Cold Fusion applications [cmdb_ci_cf_application] are custom applications that this pattern does not discover.

CI relationships

CI	Relationship	CI
cmdb_ci_cf_application	Contains::Contained by	cmdb_ci_coldfusion_server
cmdb_ci_coldfusion_server	Runs on::Runs	cmdb_ci_linux_server cmdb_ci_win_server cmdb_ci_osx_server

Connections discovered by Service Mapping during the top-down discovery

Service Mapping performs the top-down discovery of the Adobe ColdFusion in the context of application services. Service Mapping discovers the outgoing datasource connections from ColdFusion servers to instances of database.

Example

In this example, the application service map shows the results of a top-down discovery of the ColdFusion Web application service. The application service is comprised of three CIs:

- The ColdFusion application server is named **cfusion**.
- The ColdFusion application **hdStreetOracle**.
- A database named **XE**, which the ColdFusion application connects to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ColdFusion service map

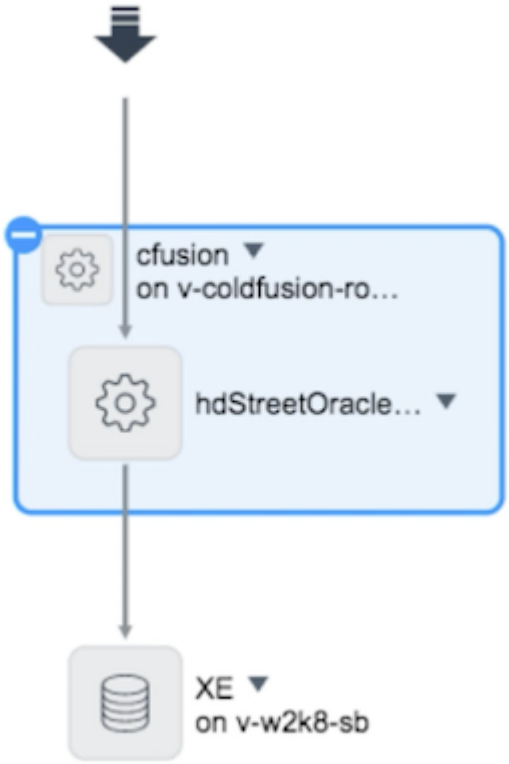

Couchbase Database discovery

The ServiceNow Discovery application uses the Couchbase Instance discovery pattern to find and map Couchbase Database instances and their clusters. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Couchbase Database is an open source, distributed database that manages JSON documents. One or more instances of a Couchbase Server constitute a cluster. Data and services are shared across the cluster.

Discovery uses the Couchbase Instance discovery pattern to run horizontal and top-down discovery. You can use this pattern on the ServiceNow AI Platform® using the Quebec release or later.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- The Couchbase Database must be installed on a Linux Server: Red Hat 7.9 or Ubuntu 20.
- Ensure that CLI tools are installed on the servers.
- On the ServiceNow AI Platform, configure SSH credentials. Configure the credentials with the following permissions:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Discovery permissions
- Permissions to list files, such as VERSION.txt, in the Couchbase installation folders on the Linux Server.
- Permissions to execute the command, `systemctl status couchbase-server`, on the Linux Server.

On the ServiceNow AI Platform, configure applicative credentials for the table cmdb_ci_couchbase_instance. The credentials must have sufficient permissions to execute the following commands on the Couchbase Server:

- `<couchbase install dir>/bin/curl -u <username>:'<password>' http://"<couchbase server>:<access_port>/pools/`
- `<couchbase install dir>/bin/curl -u <username>:'<password>' http://"<couchbase server>:<access_port>/pools/<default Pool name>`
- `<couchbase install dir>/bin/curl -u <username>:'<password>' http://"<couchbase server>:<access_port>/pools/<default Pool name>/buckets`
- `<couchbase install dir>/bin/curl -u <username>:'<password>' http://"<couchbase server>:<access_port>/pools/<default Pool name>/remoteClusters`

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

Table and field	Description
Main CI (Couchbase Instance) [cmdb_ci_couchbase_instance]	
name	The name of the Couchbase instance.
install_directory	The installation folder or directory of the Couchbase instance.
config_file	The main configuration file of the Couchbase instance.
version	The version of the Couchbase instance.
tcp_port	The port used for command line and Rest API access.
edition	Either the Basic or Enterprise edition of the Couchbase instance.
start_date	The date from which the Couchbase instance has been active.
Couchbase Cluster [cmdb_ci_couchbase_cluster]	
name	The name of the Couchbase cluster.
cluster_id	The ID of the Couchbase cluster.
install_status	The installed status, set to "1"/"installed".
cluster_type	The operational status, set to "1"/"operational".

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Couchbase Cluster Node [cmdb_ci_couchbase_cluster_node]	
name	The name of the Cluster node.
install_status	The installed status, set to "1"/"installed".
operational_status	The operational status, set to "1"/"operational".
node_status	The Cluster node health status.
node_encryption	Field that indicates whether the node is encrypted.
configured_host_name	The configured host name of the Cluster node.
host_name	The actual host name of the Cluster node.
total_memory	The total memory.
free_memory	The free memory.
mcd_reserved_memory	The reserved MCD memory.
mcd_allocated_memory	The allocated MCD memory.
couch_api_base	The Couchbase base API URL.
couch_api_base_https	The Couchbase HTTPS API URL.
recovery_type	The recovery type.
otpnode	The OTP node.
node_uuid	The UUID of the Cluster node.
cluster_compatibility	The cluster compatibility of the Cluster node.
version	The version of the Cluster node.
os	The OS of the Cluster node.
cpu_count	The CPU count of the Cluster node.
edition	The Cluster node edition. The edition is the same as that of the Couchbase instance.
Couchbase Bucket [cmdb_ci_couchbase_bucket]	
name	The name of Couchbase bucket.
cluster	Reference to the Couchbase cluster.
install_status	The installed status, set to "1"/"installed".
Couchbase Cluster Resource [cmdb_ci_couchbase_cluster_resource]	
name	The name of the Couchbase service.
install_status	The installed status, set to "1"/"installed".
operational_status	The operational status, set to "1"/"operational".
cluster	Reference to the Couchbase cluster.

CI relationships

These relationships are created to support Couchbase DB discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Couchbase Instance [cmdb_ci_couchbase_instance]	Runs on::Runs Reference	Linux Server [cmdb_ci_linux_server]
Couchbase Cluster Node [cmdb_ci_couchbase_cluster_node]	Cluster of::Cluster Reference	Couchbase Cluster [cmdb_ci_couchbase_cluster]
Couchbase Cluster Node [cmdb_ci_couchbase_cluster_node]	Hosted on::Hosts Reference	Linux Server [cmdb_ci_linux_server]
Couchbase Bucket [cmdb_ci_couchbase_bucket]	Uses::Used by Reference	Couchbase Instance [cmdb_ci_couchbase_instance]
Couchbase Bucket [cmdb_ci_couchbase_bucket]	Hosted on::Hosts Reference	Couchbase Cluster [cmdb_ci_couchbase_cluster]
Couchbase Cluster [cmdb_ci_couchbase_cluster]	Replicates to::Replicated by	Couchbase Cluster [cmdb_ci_couchbase_cluster]

CMDB data model

Couchbase DB discovery CMDB data model

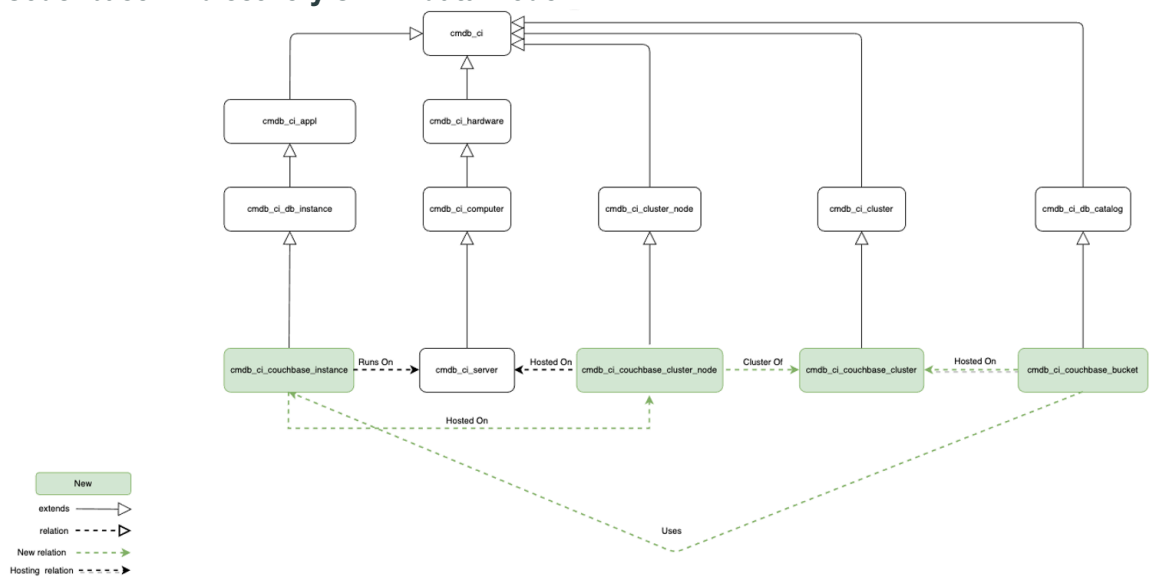

Dell EMC XtremIO storage array discovery

ServiceNow Discovery uses the Dell EMC XtremIO storage array pattern to find Dell EMC XtremIO storage array components. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Discovery and Service Mapping Patterns store app

Verify the Discovery and Service Mapping Patterns store app, version 1.0.96, is installed.

Verify the CMDB CI Class Models, version 1.42.0, is installed.

Note: The Dell EMC XtremIO storage array pattern 1.0.96 store version was tested only for XtremIO cluster type X2.

Credentials

Grant Basic Auth permissions by configuring Basic Auth credentials for RESTful API. Save the admin credentials as Basic Auth credentials to ensure a proper access.

Run Discovery

Run Quick Discovery by Target IP. Provide the IP of the XMS (that manage the storage).

Permissions for REST API queries

During the discovery, the pattern uses the following API queries. The user that is added to the credentials in the instance must have permissions to run these queries.

- Get details of XMS by index_id https://<IP>/api/json/v2/types/xms/<index_id>
- Get details of Storage Server (Cluster) by index_id https://<IP>/api/json/v2/types/clusters/<index_id>
- Get details of Storage Node Element (Bricks) by index_id https://<IP>/api/json/v2/types/bricks/<index_id>
- Get details of Storage Disks (SSDs) by index_id https://<IP>/api/json/v2/types/ssds/<index_id>
- Get details of Storage Enclosure (DAEs) by index_id https://<IP>/api/json/v2/types/daes/<index_id>
- Get details of Storage Switches (Infiniband Switches) by index_id https://<IP>/api/json/v2/types/infiniband-switches/<index_id>
- Get details of Battery Backup Units (BBUs) by index_id https://<IP>/api/json/v2/types/bbus/<index_id>
- Get details of Storage Controllers by index_id https://<IP>/api/json/v2/types/storage-controllers/<index_id>
- Get details of Storage Ports (targets) by index_id https://<IP>/api/json/v2/types/targets/<index_id>
- Get details of Disks (local disks) by index_id https://<IP>/api/json/v2/types/local-disks/<index_id>
- Get details of Initiator Groups by index_id https://<IP>/api/json/v2/types/initiator-groups/<index_id>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Get details of Fibre Channel Ports (initiators) by index_id https://<IP>/api/json/v2/types/initiators/<index_id>
- Get details of Fibre Channel Disks (Volumes) by index_id https://<IP>/api/json/v2/types/volumes/<index_id>

The new Configuration Item (CI) that is classified:

Initiator Group.

Data collected by Discovery during horizontal discovery

The discovered Dell EMC XtremIO storage array Configuration Management Database (CMDB) includes the following tables and fields.

Table and field	Description
Main CI [cmdb_ci_linux_server]	
name	The name of the Linux/XMS server
ip_address	The IP address of the Linux/XMS server
host_name	The host name of the Linux/XMS server
object_id	The identifier attribute of the Linux/XMS server
ram	The total available RAM of the Linux/XMS server
Storage Server [cmdb_ci_storage_server]	
name	The name of the Storage Server (Cluster)
object_id	The IP address of the Storage Server (Cluster)
serial_number	The serial number of the Storage Server (Cluster)
Storage Node Element [cmdb_ci_storage_node_element]	
name	The name of the Storage Node Element (X-Brick)
object_id	The identifier attribute of the Storage Node Element (X-Brick)
serial_number	The Storage Node Element (X-Brick) serial number
Storage Controller [cmdb_ci_storage_controller]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
name	The name of the Storage Controller
device_id	The Storage Controller ID
serial_number	The serial number of the Storage Controller
model_number	The model number of the Storage Controller
model_id	The model ID of the of the Storage Controller
Storage Port [cmdb_ci_storage_port]	
wwpn	The WWPN
name	The name of the Storage Port (target)
correlation_id	The Storage Port (target) correlation ID
speed	The speed of the Storage Port (target)
port_type	The Storage Port (target) type
Storage Switch [cmdb_ci_storage_switch]	
object_id	The identifier attribute of the Storage Switch (Infiniband Switch)
name	The name of the Storage Switch (Infiniband Switch)
serial_number	The serial number of the Storage Switch (Infiniband Switch)
model_id	The model ID of the Storage Switch (Infiniband Switch)
Storage Enclosure [cmdb_ci_storage_enclosure]	
name	The name of the Storage Enclosure (Disk Array Enclosure)
object_id	The identifier attribute of the Storage Enclosure (Disk Array Enclosure)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
serial_number	The serial number of the Storage Enclosure (Disk Array Enclosure)
model_number	The model number of the Storage Enclosure (Disk Array Enclosure)
UPS [cmdb_ci_ups]	
name	The name of the Battery Backup Unit (BBU)
serial_number	The serial number of the Battery Backup Unit (BBU)
model_number	The model number of the Battery Backup Unit (BBU)
model_id	The model ID of the Battery Backup Unit (BBU)
estimated_charge_remaining	The estimated charge remained for the Battery Backup Unit (BBU)
battery_voltage	The battery voltage of the Battery Backup Unit (BBU)
Initiator Group [cmdb_ci_initiator_group]	
object_id	The identifier attribute of the Initiator Group
name	The name of the Initiator Group
number_of_initiators	The number of initiators in a the group
number_of_volumes	The number of volumes
Fibre Channel Disk [cmdb_ci_fc_disk]	
name	The name of the Fibre Channel Disk (Volume)
device_id	The ID of the Fibre Channel Disk (Volume)
size	The available storage size of the Fibre Channel Disk (Volume)
Fibre Channel Port [cmdb_ci_fc_port]	
wwpn	The WWPN
wwnn	The WWNN.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
name	The name of the Fibre Channel Port (initiator)
port_type	The port type of the Fibre Channel Port (initiator)

CI relationships

CI	Relationship Type	CI
cmdb_ci_linux_server	Manage::Managed by	cmdb_ci_storage_server
cmdb_ci_storage_switch	Connects to::Connected by	cmdb_ci_storage_server
cmdb_ci_storage_server	Allocated to::Allocated from	cmdb_ci_storage_node_element
cmdb_ci_storage_enclosure	Contains::Contained By	cmdb_ci_storage_node_element
cmdb_ci_disk	Contains::Contained By	cmdb_ci_storage_node_element
cmdb_ci_storage_controller	Contains::Contained By	cmdb_ci_storage_node_element
cmdb_ci_storage_controller	Contains::Contained By	cmdb_ci_storage_port
cmdb_ci_ups	Connects to::Connected by	cmdb_ci_storage_node_element
cmdb_ci_storage_server	Contains::Contained By	cmdb_ci_Initiator _group
cmdb_ci_Initiator _group	Member::Member of	cmdb_ci_fc_port
cmdb_ci_Initiator _group	Connects to::Connected by	cmdb_ci_fc_disk

Related topics

Detailed information on products discovered by ITOM Visibility

Dell EMC Isilon discovery

The ServiceNow Discovery application uses the Dell EMC Isilon pattern to find components of Dell EMC Isilon. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

For a complete list of devices and applications that ITOM Visibility can discover, see Detailed information on products discovered by ITOM Visibility.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify Dell EMC Isilon user permissions

Once the application is classified as an Isilon cluster, the discovery process uses REST queries to find information about it. Provide the Dell EMC Isilon user with the following REST API endpoints:

- "https://" + $host + ":8080/platform/3/cluster/config”
- "https://" + $host + ":8080/platform/3/network/interfaces”
- "https://" + $host_ip + ":8080/platform/3/cluster/nodes/”
- “https://” + $host + ":8080/platform/3/zones”
- "https://" + $host + ":8080/platform/3/network/pools”
- "https://" + $host + ":8080/platform/3/storagepool/nodepools”
- "https://" + $host + ":8080/platform/3/storagepool/storagepools”
- "https://" + $host + ":8080/platform/3/protocols/nfs/exports”
- "https://" + $host + ":8080/platform/3/protocols/smb/shares"

Verify Dell EMC Isilon user roles

The Dell EMC Isilon user must have the following roles:

- ISI_PRIV_LOGIN_PAPI
- ISI_PRIV_AUTH
- ISI_PRIV_DEVICES
- ISI_PRIV_NETWORK
- ISI_PRIV_NFS
- ISI_PRIV_SMARTPOOLS
- ISI_PRIV_SMB

Create basic authentication credentials

In the Credentials module of the ServiceNow platform, create a basic authentication credential for the same Dell EMC Isilon user. For more information, see Basic authentication credentials.

Data collected by Discovery for Dell EMC Isilon

Discovery populates the data in the CMDB when running the EMC Isilon pattern.

Table and Field	Description
Storage Server [cmdb_ci_storage_server]	The server hosting the Isilon cluster.
short_description	The description configured during the Isilon server installation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and Field	Description
firmware_version	The version of the firmware installed on the Isilon server
ip_address	The IP address of the Isilon server.
manufacturer	The company that produced the device: EMC.
model_id	Preconfigured identification string of the device.
host_name	The host name of the Isilon server.
Storage Cluster [cmdb_ci_storage_cluster]	The name of the IP address of the cluster that Isilon servers form.
name	
ip_address	
Storage Cluster Node [cmdb_ci_storage_cluster_node]	The name and other attributes of the node that is part of the Isilon storage cluster.
name	
cluster	
operational_status	
Storage Node Element [cmdb_ci_storage_node_element]	The name and other attributes of the physical node that are hosted by the storage cluster.
name	
manufacturer	
cpu_type	
model_id	
Network Adapter [cmdb_ci_network_adapter]	The attribute showing if the network adapter installed on the cluster node is active.
active	
name	
ip_address	
netmask	
IP Address [cmdb_ci_ip_address]	The attributes of the IP address of the cluster node.
name	
netmask	
Disk [cmdb_ci_disk]	The storage disk installed on the Isilon Server.
device_id	
name	
model_id	
interface	
disk_interface	
Storage File Share [cmdb_ci_storage_fileshare]	The fileshare server attributes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and Field	Description
path	
name	
fileshare_id	
correlation_id	
Storage Volume [cmdb_ci_storage_volume]	The attributes of the storage volume belonging to the storage cluster.
media_type	
object_id	
size_bytes	
free_space_bytes	
Storage Pool [cmdb_ci_storage_pool]	
pool_id	The attributes of the storage pool to which the storage cluster belongs.
free_space_bytes	
size_bytes	
name	
object_id	
Serial Number [cmdb_serial_number]	The attributes of the serial number of the storage node.
serial_number	
serial_number_type	
valid	

The graphic illustrates CIs that are part of Isilon discovery.

Note: This Dependency Views map was simplified for clarity. Your Isilon deployments can contain many more CIs.

Dependency Views showing Isilon components

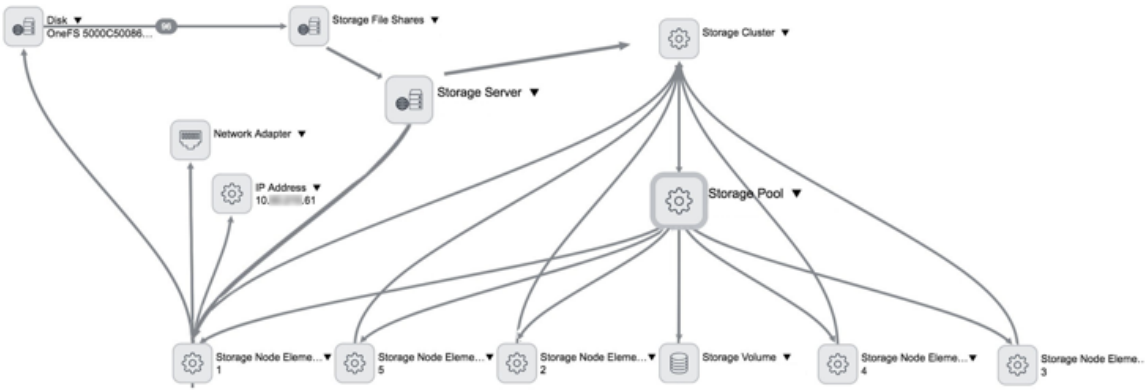

CI Relationships

CI	Relationship/Reference	CI
cmdb_ci_storage_server	Runs on::Runs	cmdb_ci_storage_cluster

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship/Reference	CI
cmdb_ci_storage_server	Allocated to::Allocated from	cmdb_ci_storage_node_element
cmdb_ci_storage_cluster_node	Cluster of::Cluster	cmdb_ci_storage_cluster (reference *cluster*)
cmdb_ci_storage_cluster_node	Hosted on::Hosts	cmdb_ci_storage_node_element (reference *server*)
cmdb_ci_storage_node_element	Managed by::Manages	cmdb_ci_storage_cluster
cmdb_ci_network_adapter	Owns::Owned by	cmdb_ci_storage_node_element (reference *cmdb_ci*)
cmdb_ci_network_adapter	Used by::Uses	cmdb_ci_storage_server (reference *cmdb_ci*)
cmdb_ci_ip_address	Owns::Owned by	cmdb_ci_storage_node_element
cmdb_ci_ip_address	Used by::Uses	cmdb_ci_storage_server
cmdb_ci_disk	Contains::Contained by	cmdb_ci_storage_node_element (reference *computer*)
cmdb_ci_storage_fileshare	Provided by::Provides	cmdb_ci_storage_server
cmdb_ci_storage_fileshare	Allocated from::Allocated to	cmdb_ci_storage_volume
cmdb_ci_storage_volume	Uses::Used by	cmdb_ci_storage_server
cmdb_ci_storage_volume	Owns::Owned by	cmdb_ci_storage_cluster
cmdb_ci_storage_pool	Owns::Owned by	cmdb_ci_storage_cluster (reference *hosted_by*)

Dell PowerMax storage discovery with Patterns

ServiceNow Discovery uses the **EMC PMAX phase1** and **EMC PMAX phase2** patterns to find Dell PowerMax storage components. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about PowerMax storage components and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

Prerequisites

Verify that the applications are up to date

Discovery and Service Mapping Patterns

CMDB CI Class Models

Verify the creation of the credential alias

1. Set the scope to: **Global**.
2. Navigate to **Connections & Credentials** > **Connection & Credential Aliases**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **New**.
4. Fill in the fields in the form.

Field	Description
Name	Enter a unique name for the alias. For example: `powermax_discovery_ceredential_alias`. Make note of the alias name to use it for configuring the **Basic Auth** credentials.
Type	From the list, select **Credential**.
ID	The **ID** field is populated automatically when submitting the form. Make note of the alias ID to use it for configuring the parameters for the **EMC PMAX phase1** pattern. For example: `powermax_credential_alias_id`.

5. Select **Submit**.

Verify the configuration of the Basic Auth credentials

Configure the **Basic Auth** credentials to enable the connection to the PowerMax REST API. For the procedure, see Create Basic Auth credentials.

Fill in the form.

Field	Description
Name	Fill in a unique name for the credentials. For example: `powermax_discovery_basic_auth_ceredentials`.
User name	Fill in the user name that can access the PowerMax. The user should have minimum read-only permissions.
Password	Fill in the password that can access the PowerMax.

Verify the configuration of a Serverless Discovery schedule

Make sure to select `Serverless` in the **Discover** field. For information about the procedure, see: Schedule a horizontal discovery

Associate the EMC PMAX Phase1 and EMC PMAX Phase2 patterns with the Serverless Discovery schedule

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Select the schedule that you configured.
3. Select the **Serverless Execution Patterns** tab.
4. Verify that the **Run Child Patterns** field is set to **true**.
5. Select **New**.
6. Fill in the fields in the form as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
IP Address [emc_ip_address]	The IP Address of the Navisphere, which serves as the EMC Management Console capable of managing multiple storages.
Alias [credential_alias]	The Alias ID that is configured for the EMC PowerMax credentials. For example: `powermax_credential_alias_id`.

7. Select **Submit**.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the **EMC PowerMAX** Patterns.

Storage Server [cmdb_ci_storage_server]

Field	Description
Name [name]	The user-assigned identifier or label for a specific storage array or component within the system.
Model ID [model_id]	The particular variant or model of the storage array.
Serial Number [serial_number]	The serial number that uniquely identifies each individual storage array unit.
OS [os]	The operating system environment hosted on the storage array. For example, embedded management software or operating systems.
OS Version [os_version]	The software version that's installed on the storage array, which includes the core operating system and management software.
Firmware Version [firmware _version]	The software embedded in the hardware components of the storage array, including controllers, drives, and other subsystems.
Manufacturer [manufacturer]	The company responsible for designing, producing, and supporting the PowerMax storage array which is Dell EMC.
Physical Capacity (GB) [physical_capacity]	The maximum physical capacity that can be utilized for storing data or other content, measured in gigabytes (GB).
Used Physical Capacity (GB) [used_physical_capacity]	Amount of physical storage space currently allocated and utilized within the PowerMax system, measured in GB. It indicates how much of the available physical storage resources are actively storing data or other content.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server [cmdb_ci_storage_server] (continued)

Field	Description
Free Physical Capacity (GB) [free_physical_capacity]	Available physical storage space within the PowerMax system, measured in GB.
Virtual Capacity (GB) [virtual_capacity]	Total virtual storage capacity provisioned within the PowerMax system, measured in GB.
Free Virtual Capacity (GB) [free_virtual_capacity]	Available virtual storage capacity within the PowerMax system, measured in GB.
Used Virtual Capacity (GB) [used_virtual_capcity]	Used virtual storage capacity within the PowerMax system, measured in GB.

Storage Cluster Node [cmdb_ci_storage_cluster_node]

Field	Description
Name [name]	The cluster node name.
Cluster [cluster]	The host group associated with this cluster node in the PowerMax environment.
Manufacturer [manufacturer]	The manufacturer name, which is Dell EMC.

Storage Cluster [cmdb_ci_storage_cluster]

Field	Description
Name [name]	The storage cluster name.
Correlation ID [correlation_id]	Combines the cluster name and the storage's serial number to create a unique identifier.

Storage Controller [cmdb_ci_storage_controller]

Field	Description
Name [name]	The combination of asset and asset tag.
Assigned [assigned]	The availability status.
Asset [asset]	The director slot number.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Controller [cmdb_ci_storage_controller] (continued)

Field	Description
Asset Tag [asset_tag]	The director number.
Device ID [device_id]	A unique label for the controller.
Short Description [short_description]	The director slot and number.
Manufacturer [manufacturer]	The manufacturer name, which is Dell EMC.

Storage Port [cmdb_ci_storage_port]

Field	Description
WWPN [wwpn]	The World Wide Port Name (WWPN) of the storage port.
Name [name]	The port identifier.
PO Number [po_number]	The port number.
State [state]	The port status.
Speed [speed]	The port transmission speed.
Manufacturer [manufacturer]	The manufacturer name, which is Dell EMC.
Short Description [short_description]	A list of the port IP addresses.
IP Address [ip_address]	The connected IP address.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Object ID [object_id]	The World Wide Name (WWN) unique identifier.
State [state]	The current operational status of the volume. For example: **Write Disabled** or **Ready**.
Size [size]	The capacity in Gigabytes (GB).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk [cmdb_ci_disk]

Field	Description
Name [name]	The Spindle ID.
Manufacturer [manufacturer]	The disk manufacturer.
Model ID [model_id]	The disk model.
Size [size]	The disk capacity in Gigabytes (GB).
Device ID [device_id]	The combination of the storage system serial number with the disk name that uniquely identifies each disk within the PowerMax environment.

FC Export [cmdb_ci_fc_export]

Field	Description
Export ID [export_id]	The Fibre Channel identifier.
Initiator WWPN [initiator_wwpn]	A unique identifier combining the Director ID, Port Number, and Initiator ID.
PO Number [po_number]	The port number.
Assigned to [assigned_to]	The specific host assigned to the export and manages the Fibre Channel connection.
Attributes [attributes]	The director ID that identifies the specific director responsible for managing the Fibre Channel exports.
Name [name]	The initiator ID.
Manufacturer [manufacturer]	The manufacturer name, which is Dell EMC.
Lun [lun]	The Logical Unit Number (LUN) that is connected within PowerMax systems, indicating the specific storage unit allocated for the Fibre Channel export.
Asset Tag [asset_tag]	A combination of the Director ID with the Port Number that uniquely identifies the asset within the Fibre Channel exports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

iSCSI Export [cmdb_ci_isci_export]

Field	Description
Export ID [export_id]	The unique identifier for an iSCSI export within PowerMax systems, facilitating the management and tracking of iSCSI connections.
Initiator WWPN [initiator_wwpn]	A combination of the Director ID, Port Number, and Initiator ID that forms a unique identifier for the initiator within the iSCSI exports.
Name [name]	The Initiator ID in PowerMax systems, serving as a unique identifier for the initiator in the context of iSCSI exports.
PO Number [po_number]	The port number identifying the specific port used for the iSCSI export.
Attributes [attributes]	The Director ID, which identifies the specific director responsible for managing the iSCSI export.
Assigned to [assigned_to]	The host assigned to the iSCSI export within the PowerMax environment and specifies the particular host system that is connected to the export.

Storage Pool [cmdb_ci_storage_pool]

Field	Description
Name [name]	The name of the pool.

CI relationships

Discovery creates these relationships to support the PowerMax discovery with the **EMC PMAX phase2** pattern.

CI	Relationship Type	CI
cmdb_ci_disk	Contains::Contained by	cmdb_ci_storage_server
cmdb_ci_fc_export	Allocated from::Allocated to	cmdb_ci_storage_volume
cmdb_ci_fc_export	Exports to::Imports from	cmdb_ci_storage_cluster_node
cmdb_ci_isci_export	Hosted::Hosts	cmdb_ci_storage_server
cmdb_ci_isci_export	Allocated from::Allocated to	cmdb_ci_storage_volume
cmdb_ci_isci_export	Exports to::Imports from	cmdb_ci_storage_cluster_node

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship Type	CI
cmdb_ci_storage_cluster_node	Cluster of::Cluster	cmdb_ci_storage_cluster
cmdb_ci_storage_controller	Contains::Contained By	cmdb_ci_storage_server
cmdb_ci_storage_controller	Contains::Contained By	cmdb_ci_storage_port
cmdb_ci_storage_port	Owns::Owned by	cmdb_ci_fc_export
cmdb_ci_storage_port	Owns::Owned by	cmdb_ci_isci_export
cmdb_ci_storage_port	Owns::Owned by	cmdb_ci_isci_export
cmdb_ci_storage_port	Owns::Owned by	cmdb_ci_fc_export
cmdb_ci_storage_server	Owns::Owned by	cmdb_ci_storage_port
cmdb_ci_storage_server	Runs on::Runs	cmdb_ci_storage_cluster
cmdb_ci_storage_server	Allocated from::Allocated to	cmdb_ci_storage_cluster_node
cmdb_ci_storage_server	Contains::Contained By	cmdb_ci_storage_pool

Dell Data Domain storage discovery using Patterns

Discovery and Service Mapping use the **Dell EMC Data Domain** serverless pattern to find the Data Domain storage systems registered in the Data Domain Management Center (DDMC). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about Data Domain components and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility

Prerequisites

Verify that the applications are up to date

Discovery and Service Mapping Patterns

CMDB CI Class Models

Verify the creation of the credential alias

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

role: DDMC admin.

1. Set the scope to: **Global**.
2. Navigate to **Connections & Credentials** > **Connection & Credential Aliases**.
3. Select **New**.
4. Fill in the fields in the form.

Field	Description
Name	Enter a unique name for the alias. For example: `ddmc_discovery_ceredential_alias`. Make note of the alias name to use it for configuring the **Basic Auth** credentials.
Type	From the list, select **Credential**.
ID	The **ID** field is populated automatically when submitting the form. Make note of the alias ID to use it for configuring the parameters for the **Dell EMC Data Domain** pattern. For example: `ddmc_credential_alias_id`.

5. Select **Submit**.

Verify the configuration of the Basic Auth credentials

role: DDMC admin.

Note: To run discovery with Patterns, you can create only one credential per one alias.

1. Navigate to **Connections & Credentials** > **Connection & Credential Aliases**.
2. In the list, search for the alias unique name and open the record.
3. From the list, select **Basic Auth** Credentials.
4. Fill in the fields in the form.

Field	Description
Name	Fill in a unique name for the credentials. For example: `ddmc_discovery_basic_auth_ceredentials`.
User name	Fill in the user name that can access the DDMC. This user is created in the DDMC. The user should have minimum read-only permissions. Alternatively, fill in the REST API interface for the DDMC.
Password	Fill in the password that can access the DDMC. This password is created in the DDMC.
Active	Enabled by default.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Credential alias	If you entered the form through the alias, you don't need to configure this field.

5. Select **Submit**.

Verify the configuration of a Serverless Discovery schedule

role: Discovery_admin.

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Select **New**.
3. Fill in the fields in the form.

Field	Description
Name	Fill in the name of your choice.
Discover	In the list, select `Serverless`.
MID Server	Select the search button to open the list of MID Servers and select the MID Server that has network access to the DDMC.
Active	Enabled by default.
Credential alias	No need to configure this field.

For more information about setting the **Max run time**, **Run**, and **Time** fields, see: Schedule a horizontal discovery

4. Select **Submit**.

Associate the Dell EMC Data Domain pattern with the Serverless Discovery schedule

role: Discovery_admin.

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Select the schedule that you configured.
3. Select the **Serverless Execution Patterns** tab.
4. Select **New**.
5. Fill in the fields in the form.

Field	Description
Name	Fill in the name of your choice. For example: `dell_emc_data_domain_serverless_discovery`
Pattern	Select the search button to open the list of patterns, and select the **Dell EMC Data Domain** pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Proxy Host	No need to configure this field.
Extended Parameters	No need to configure this field.

6. Select **Submit**.

Configure the parameters for the Dell EMC Data Domain pattern

role: Discovery_admin.

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Select the schedule that you configured.
3. Select the **Serverless Execution Pattern** tab.
4. Select the record created for the **Dell EMC Data Domain** pattern.
5. In the **Discovery Pattern Launcher parameters** tab, set the value for the fields as follows:

Field	Description
Credential Alias	Fill in the alias ID configured for DDMC. You can find the alias in the **Connection & Credential Aliases** list. For example: `ddmc_credential_alias_id`.
DDMC Port	(Optional) by default the pattern uses port 3009.
DDMC IP Address	Fill in the host name or the IP Address of the DDMC. These details can be provided by the DDMC Admin. For example: `101.10.10.101` or `ddmc.demo.local`.

6. To run the discovery, you can select **Discover now** from the **Related Links** or wait for the run time.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the **Dell EMC Data Domain** Patterns.

Storage Server [cmdb_ci_storage_server]

Field	Description
Name [name]	The unique identifier of the storage server. The name is typically the Domain Name System (DNS) hostname or the computer name, which corresponds to the **host_name** attribute.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server [cmdb_ci_storage_server] (continued)

Field	Description
Host Name [host_name]	The DNS hostname of the Data Domain system. The host that is used to register the system with the Data Domain Management Center (DDMC).
FQDN [fqdn]	The Fully Qualified Domain Name of the Data Domain system. The complete domain name used to register the system with the DDMC.
IP Address [ip_address]	The IP address of the Data Domain storage server. The name is the specific IP address used to register the Data Domain system with the DDMC.
OS Version [os_version]	The version of the operating system running on the Data Domain system. The versions supported are from DD OS 7.3.
Model [model]	The specific model of the Data Domain (DD) system. For example, the name for the DD Virtual Edition (VE) is **DELL EMC DD VE**.
Manufacturer [manufacturer]	The manufacturer is Dell EMC.
Serial Number [serial_number]	The unique serial number assigned to the DD system.
Operating System [os]	The Data Domain OS is: **DD OS**.

Storage Cluster [cmdb_ci_storage_cluster]

Field	Description
Name [name]	The name of the DDMC associated with the storage cluster.
Cluster ID [cluster_id]	The universally unique identifier (UUID) of the DDMC, which uniquely identifies the cluster.
Cluster Version [cluster_version]	The software version of the DDMC that is currently installed on the cluster.
Serial Number [serial_number]	The unique serial number assigned to the Virtual Data Domain Management Center (VDDMC) system.
Model ID [model_id]	The specific model of the Data Domain system in use. For the DDMC, the model ID is 'DELL EMC DD Management Center'.
Manufacturer [manufacturer]	The manufacturer is Dell EMC.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk [cmdb_ci_disk]

Field	Description
Name [name]	The name assigned to the disk device within the storage system.
Device ID [device_id]	The unique identifier for the disk device, which is the same as the device name.
Model ID [model_id]	The model of the disk device. For example: **VMware Virtual_disk**.
Size [size]	The capacity of the disk device, measured in gigabytes (GB).

The Disk Shelf Chassis database is populated only for physical DD systems that support disk expansion shelves.

Disk Shelf Chassis [cmdb_ci_diskshelf_chassis_cluster]

Field	Description
Name [name]	The name of the disk expansion shelf.
Serial Number [serial_number]	The unique serial number of the disk expansion shelf.

Serial Number [cmdb_ci_serial_number]

Field	Description
Serial Number [serial_number]	The unique serial number of the Dell EMC Data Domain storage system, as registered in the DDMC.
Serial Number Type [serial_number_type]	The format of the serial number for the Dell EMC Data Domain storage system, indicating whether it's a system serial number or a Universally Unique Identifier (UUID).
Valid [valid]	The validity of the data, with a value of **true** confirming that the serial number has been successfully verified.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	The identifier or label of the network adapter, such as **eth1**.

The DDMC REST API doesn't provide the MAC address of network adapters. As a result, the MAC address field for network adapters remains empty.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	The IP address that is assigned to the device.
Netmask [netmask]	The netmask associated with the IP address, defining the subnet to which the IP address belongs.

CI relationships

Discovery and Service Mapping Patterns create these relationships to support the Data Domain discovery.

CI	Relationship Type	CI
cmdb_ci_storage_server	Managed by::Manages	cmdb_ci_storage_cluster
cmdb_ci_storage_server	Contains::Contained By	cmdb_ci_disk
cmdb_ci_disk_shelf_chassis	Hosted on:: hosts	cmdb_ci_storage_server
cmdb_ci_serial_number	reference	cmdb_ci_storage_server
cmdb_ci_storage_server	Owns::Owned by	cmdb_ci_network_adapter
cmdb_ci_storage_server	Owns::Owned by	cmdb_ci_ip _address

F5 certificate discovery

The ServiceNow® Discovery application uses **The F5-SSH-SSL Certification** pattern extension to find all associated certificates on F5 load balancers that use IPv4 addresses, IPv6 addresses, or both.

The **The F5-SSH-SSL Certification** pattern extension is part of the **F5 load balancer** and **F5 load balancer SSH** patterns that discover F5 BIG-IP load balancer. The **The F5-SSH-SSL Certification** pattern extension discovers F5 load balancer instances running on the hosts via SNMP/SSH, and then discovers certifications used by the running node of the load balancer.

Note: For load balancers running IPv6, only SSH is supported.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

For details on system requirements and family compatibility, view the application listing on the ServiceNow Store website.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IPv6 support limitations

- F5 SNMP is not supported.
- *MID Server* cannot run REST APIs with IPv6.

  To run discovery with F5 REST, see KB0864769.

Prerequisites

Verify the applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models

Verify the configuration of F5 load balancer

- Ensure that F5 load balancer instances are up and running.
- Make sure the host (that has the F5 instance running on it) can discover successfully with the Discovery user credentials.

Verify the configuration of Discovery

Ensure that the Discovery user added in the ServiceNow AI Platform instance can run the following commands:

Command	Description
modify	Modifies the TMSH components. You can modify one or more property settings in multiple components. The *modify* command uses the following option: *display-threshold pager*
display_threshold	Allows you to re-enable a *display-threshold* in your script.
list	Displays components that you have permission to view or are passed as arguments. The *list* command looks for the following arguments: • *ssl-cert* • *certificate-key-size* • *checksum* • *create-time* • *expiration-string* • *issuer* • *subject* • *version* • *fingerprint* • *serial-number*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Description
	• *subject-alternative-name* • *size*

Verify the configuration of the patterns

Make sure that the **F5-SSH-SSL Certification** shared library is added to the extension section of the **F5 Load Balancer** and **F5 Load Balancer SSH** patterns, in order to collect the certification attributes.

F5 certificate discovery class model

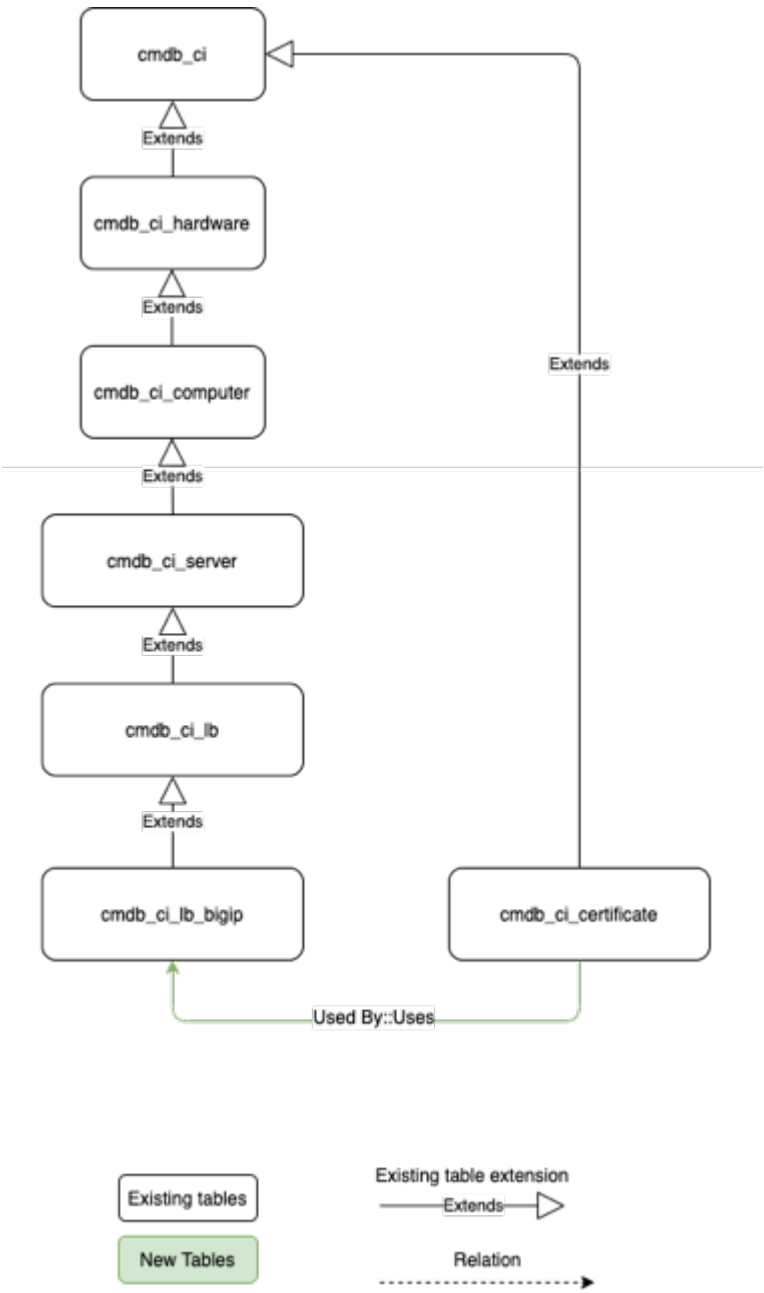

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

Table and field	Description
Base Configuration Item Cluster [cmdb]	
Serial number	Serial Number associated with the CI.
Configuration Item [cmdb_ci]	
operational_status	Operational status of the cluster node.
Unique Certificate [cmdb_ci_unique_certificate]	
fingerprint	Hash value of the certificate.
fingerprint algorithm	Algorithm used to hash the certificate.
subject common name	Identifies the hostname/domain associated with the certificate.
subject distinguished name	Identifying information of the subject.
issuer distinguished name	Distinguished name of the issuer.
comments	
renewal tracking	Indicates whether to create any priority 1 or priority 3 tasks for the expiring certificates.
issuer common name	Common name of the issuer.
valid From	Validity start period of the certificate.
serial Number	Serial Number associated with the CI.
subject country	The subject's two letter country code.
subject organization	Subject’s organization.
Version	X.509 version of the certificate.
Issuer	Entity that signed and issued the certificate.
subject organizational unit	Subject's organizational unit.
subject alternative name	List of fully qualified domain names secured by the certificate.
valid to	Validity end period of the certificate.
name	
State	Lifecycle states of the certificate.
root issuer	Root entity that signed and issued the immediate certificate.
key size	Size of the key used by the signing algorithm.
subject locality	Subject's locality.
subject state	Subject's state.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The **The F5-SSH-SSL Certification** pattern extension does not create any CI relationships.

Fortinet firewall and FortiGate VDOM REST-based discovery

The Discovery and Service Mapping Patterns application uses the Next Generation Fortinet Network Firewall - REST pattern to find Fortinet firewalls through REST API calls. Additionally, the pattern extension VDOM Discovery finds FortiGate Virtual Domains (VDOMs). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The Next Generation Fortinet Network Firewall - REST pattern uses a set of REST API calls to find the Fortinet firewalls. For FortiGate VDOM discovery, only multi-VDOM mode is supported.

Note: Only the REST-based Fortinet firewall discovery method finds FortiGate VDOMs. The SNMP-based Fortinet firewall discovery method doesn't discover them. For information about the default SNMP-based Fortinet firewall discovery, see Next-Generation Fortinet Network Firewall SNMP-based discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about Fortinet Firewalls and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

Fortinet firewall and FortiGate VDOM data model

The Next Generation Fortinet Network Firewall - REST pattern and VDOM Discovery extension introduce the following CI classes that extend existing CMDB classes.

CI classes introduced by these patterns

CI class	Extends from
Fortinet Firewall Cluster [cmdb_ci_firewall_cluster_fortinet]	Firewall Cluster [cmdb_ci_firewall_cluster]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Firewall Device [cmdb_ci_firewall_device]
Fortinet Firewall Interface [cmdb_ci_fortinet_firewall_interface]	Network Interface [cmdb_ci_ni_interface]
Fortinet Firewall Policy [cmdb_ci_fortinet_firewall_policy]	Firewall Security Policy [cmdb_ci_firewall_sec_policy]
Fortinet Virtual Domain [cmdb_ci_fortinet_vdom]	IP Firewall [cmdb_ci_ip_firewall]

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create API Token

Create an API Token in Fortinet. For instructions, go to the Fortinet Document Library. Search for the "Connect FortiGate device via API Token - Online Help" article under the FortiConverter Tool product family.

Verify API access and permissions

- Verify that the MID Server can access the Fortinet APIs.
- Verify that the API Token has sufficient permissions to retrieve the required information from the Fortinet devices.

Required Fortinet APIs:

- `v2/cmdb/system/global`
- `api/v2/monitor/system/status`
- `api/v2/cmdb/system/ha`
- `api/v2/cmdb/router/static`
- `/api/v2/cmdb/firewall/policy`
- `api/v2/cmdb/system/vdom-property`
- `api/v2/cmdb/system/interface`
- `/api/v2/monitor/system/interface/select`

Disable SNMP-based Fortinet firewall discovery

For more information, see Disable SNMP-based Fortinet firewall discovery.

Create an alias for the API Key Credentials

For more information, see Create an alias for the API key credential for Fortinet firewall REST-based discovery.

Create a serverless discovery schedule

For more information, see Create a serverless schedule for Fortinet firewall REST-based discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Next Generation Fortinet Network Firewall - REST pattern.

Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]

Field	Description
Name [name]	**Hostname** field of the Fortinet device.
Serial number [serial_number]	Serial number of the Fortinet device.
Fully qualified domain name [fqdn]	Fully qualified domain name of the Fortinet device.
Operational status [operational_status]	Indicates whether the Fortinet device is in active state.
IP Address [ip_address]	IP address of the Fortinet device.
Manufacturer [manufacturer]	Fortinet device manufacturer.
Description [short_description]	Short description of the Fortinet device.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet] (continued)

Field	Description
Model Number [model_number]	Fortinet device model number.
Firmware version [firmware_version]	Fortinet device firmware version.
Hardware OS [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Fortinet Firewall Cluster [cmdb_ci_firewall_cluster_fortinet]

Field	Description
Name [name]	**Hostname** field of the Fortinet firewall cluster.
Fully qualified domain name [fqdn]	Fully qualified domain name of the firewall cluster.
IP address [ip_address]	IP address of the firewall cluster.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description of the firewall cluster.
Model Number [model_number]	Device model number.
Hardware OS [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
IP Address [ip_address]	IP address of the network adapter.
Netmask [netmask]	Netmask of the network adapter.
Alias [alias]	User-assigned name for the network adapter.
MAC Address [mac_address]	MAC address of the network adapter.
Name [name]	Name of the network adapter.
Configuration Item [cmdb_ci]	References the Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet] table.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address of the Fortinet firewall.
Netmask [netmask]	Netmask of the Fortinet firewall.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery populates the data in the CMDB when running the Next Generation Fortinet Network Firewall - REST pattern extension VDOM Discovery.

Fortinet Virtual Domain [cmdb_ci_fortinet_vdom]

Field	Description
Vdom Index [vdom_index]	Index of the VDOM in the list.
Name [name]	Name of the VDOM.
Description [short_description]	Description of the VDOM property that provides additional context or information about the purpose of the property.
IP Address [ip_address]	IP address of the Fortinet device associated with this VDOM.

Fortinet Firewall Interface [cmdb_ci_fortinet_firewall_interface]

Field	Description
ID [id]	Unique identifier for the interface. For example: port1.
Name [name]	Name of the interface. For example: LAN.
IP Address [ip_address]	IP address assigned to the interface.
Description [short_description]	Description of the interface, often used for documentation or identification purposes. For example: Main LAN interface.
MAC Address [mac_address]	MAC address of the interface.
Access Type [access_type]	Type of interface. For example: physical, VLAN, or aggregate.

Network Topology [cmdb_ci_network_topology]

Field	Description
Name [name]	Virtual LAN (VLAN) ID associated with the interface, if applicable. The format is: `VLAN-{Vlan ID}`. For example: VLAN-310.

Fortinet Firewall Policy [cmdb_ci_fortinet_firewall_policy]

Field	Description
Policy ID [policy_id]	Unique ID assigned to the policy in VDOM or device level.
UUID [uuid]	Global unique identifier (GUID) for the firewall policy.
Name [name]	Name of the policy.
Source interface [source_interface]	Network interface from which the traffic originates.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fortinet Firewall Policy [cmdb_ci_fortinet_firewall_policy] (continued)

Field	Description
Destination interface [destination_interface]	Network interface to which the traffic is directed.
Source address [source_address]	Source address or address group from which traffic originates.
Destination address [destination_address]	Destination address or address group to which traffic is directed.
Internet Service [internet_service]	Service or application being managed by the policy, often represented by a service group or name.

CI relationships

The Next Generation Fortinet Network Firewall - REST pattern creates the following relationships and references to support Fortinet firewall discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Fortinet Firewall Cluster [cmdb_ci_firewall_cluster_fortinet]	Hosted on::Hosts	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Uses::Used by	Router Interface [dscy_router_interface]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	IP Address [cmdb_ci_ip_address]

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration item [configuration_item]	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
Router Interface [dscy_router_interface]	Configuration Item [cmdb_ci]	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]

The VDOM Discovery extension creates the following relationships to support FortiGate VDOM discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Fortinet Virtual Domain [cmdb_ci_fortinet_vdom]	Contains::Contained by	Fortinet Firewall Interface [cmdb_ci_fortinet_firewall_interface]
Fortinet Virtual Domain [cmdb_ci_fortinet_vdom]	Contains::Contained by	Fortinet Firewall Policy [cmdb_ci_fortinet_firewall_policy]
Fortinet Virtual Domain [cmdb_ci_fortinet_vdom]	Hosted on::Hosts	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
Fortinet Firewall Interface [cmdb_ci_fortinet_firewall_interface]	Members::Member of	Network Topology [cmdb_ci_network_topology]

Disable SNMP-based Fortinet firewall discovery

Disable the default SNMP-based Fortinet firewall discovery and use REST-based discovery instead to discover FortiGate Virtual Domains (VDOMs).

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **SNMP**.
2. In the **Name** field, search for `Fortinet Firewall`.
3. Select the Fortinet Firewall record.
4. Clear the **Active** check box.
5. Select **Update**.

What to do next

Create an alias and add it to an API key credential. For more information, see Create an alias for the API key credential for Fortinet firewall REST-based discovery.

Related topics

Fortinet firewall and FortiGate VDOM REST-based discovery

Create an alias for the API key credential for Fortinet firewall REST-based discovery

Create an alias and add it to an API key credential to discover Fortinet firewalls and FortiGate Virtual Domains (VDOMs) through REST-based discovery.

Before you begin

Create an API Token in Fortinet. For more information, see the Create API Token prerequisite in Fortinet firewall and FortiGate VDOM REST-based discovery.

Role required: discovery_admin

Procedure

1. Create an alias.
   a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   b. Select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Enter a unique name for the alias and select **Credential** for the alias type in the **Type** field.

d. Select **Submit**.

2. Configure an API key credential for the new alias.

a. Navigate to **Discovery** > **Credentials**.

b. Select **New**.

c. Select **API Key Credentials**.

d. Unlock the **Credential alias** field and select the alias you created.

e. Configure the rest of the **API Key Credentials** form fields.

f. Select **Submit**.
For more information about API Key Credentials form fields, see API key credentials.

What to do next

Create a serverless schedule for Fortinet Firewall REST-based discovery. For more information, see Create a serverless schedule for Fortinet firewall REST-based discovery.

Related topics

Fortinet firewall and FortiGate VDOM REST-based discovery

Create a serverless schedule for Fortinet firewall REST-based discovery

Create a serverless discovery schedule to discover Fortinet firewalls and FortiGate Virtual Domains (VDOMs) through REST-based discovery.

Before you begin

- Verify you have the IP address for the Fortinet firewall device.
- Create an alias for the API key credential. For more information, see Create an alias for the API key credential for Fortinet firewall REST-based discovery.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.

a. Select **New**.

b. On the form, fill in the fields.

Discovery Schedule New record form

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type, which should be Serverless.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
MID server	Name of the MID Server to use for this schedule.
Active	Enables this schedule for discovery.

c. Select **Submit**.

3. Create the execution pattern.

a. In the Discovery Schedules page, select the record you created.

b. In the **Serverless Execution Patterns** tab, select **New**.

c. On the form, fill in the fields.

Serverless Execution Pattern New record form

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern to be used for this schedule, which should be the Next Generation Fortinet Network Firewall - REST pattern.
Proxy Host	Fully qualified domain name of the machine on which you're installing the proxy server, which should be Global.
Active	Enables this schedule for discovery.

d. Select **Submit**.

4. Set the pattern launcher parameters.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Parameter	Value
ip_address	IP address of the firewall device.
credentialsAlias	ID of the credential alias you created.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

Fortinet firewall and FortiGate VDOM REST-based discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Next-Generation Fortinet Network Firewall SNMP-based discovery

The Discovery and Service Mapping Patterns application uses the Next Generation Fortinet Network Firewall pattern to find Fortinet firewalls through a series of SNMP calls. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The Next Generation Fortinet Network Firewall pattern uses a set of SNMP calls to find the Fortinet firewalls. Discovery uses the pattern to run horizontal discovery.

Note: Only the REST-based Fortinet firewall discovery method finds FortiGate VDOMs. The SNMP-based Fortinet firewall discovery method doesn't discover them. For information on REST-based Fortinet firewall and FortiGate Virtual Domains (VDOMs) discovery, see Fortinet firewall and FortiGate VDOM REST-based discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about Fortinet firewalls and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

Fortinet Firewall data model

The Next Generation Fortinet Network Firewall pattern introduces the following CI classes that extend an existing CMDB class.

CI classes introduced by this pattern

CI class	Extends from
Fortinet Firewall Cluster [cmdb_ci_firewall_cluster_fortinet]	Firewall Cluster [cmdb_ci_firewall_cluster]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Firewall Device [cmdb_ci_firewall_device]

Prerequisites

Verify the applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models

Ensure SNMP access

Ensure that your Fortinet firewall device has SNMP access.

Configure SNMP credentials

On the ServiceNow instance, configure SNMP credentials. For more information, see SNMP credentials.

Add SNMP system OID record to ServiceNow instance

Add the SNMP system OID record for the Fortinet device to the ServiceNow instance. Update the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Classifier: Fortinet Firewall
- Class: Fortinet Firewall Device

Run a horizontal discovery

For more information, see Running discoveries in your network.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Next Generation Fortinet Network Firewall Pattern.

Fortinet Firewall Cluster [cmdb_ci_firewall_cluster_fortinet]

Field	Description
Name [name]	Hostname.
Fully qualified domain name [fqdn]	Fully qualified domain name.
IP address [ip_address]	IP address.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description of the Fortinet firewall cluster.
Model Number [model_number]	Device model number.
Hardware Operating System [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]

Field	Description
Name [name]	Hostname.
Serial Number [serial_number]	Serial number of the device.
Fully qualified domain name [fqdn]	Fully qualified domain name.
Operational Status [operational_status]	Indicates if the device is in active state.
IP address [ip_address]	IP address.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description of the device.
Model Number [model_number]	Device model number.
Firmware [firmware_version]	Firmware version.
Hardware Operating System [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address of the Fortinet firewall.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Address [cmdb_ci_ip_address] (continued)

Field	Description
Netmask [netmask]	Netmask of the Fortinet firewall.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
IP Address [ip_address]	IP address of the network adapter.
Netmask [netmask]	Netmask of the network adapter.
Alias [alias]	User-assigned name for the network adapter.
MAC Address [mac_address]	MAC address of the network adapter.
Name [name]	Name of the network adapter.
Configuration Item [cmdb_ci]	References the Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet] table.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Name of the Domain Name System (DNS).
IP Address [ip_address]	IP address of the DNS.

This Dependency Views map on the Fortinet Firewall Device CI shows the Fortinet Firewall Cluster to which it belongs.

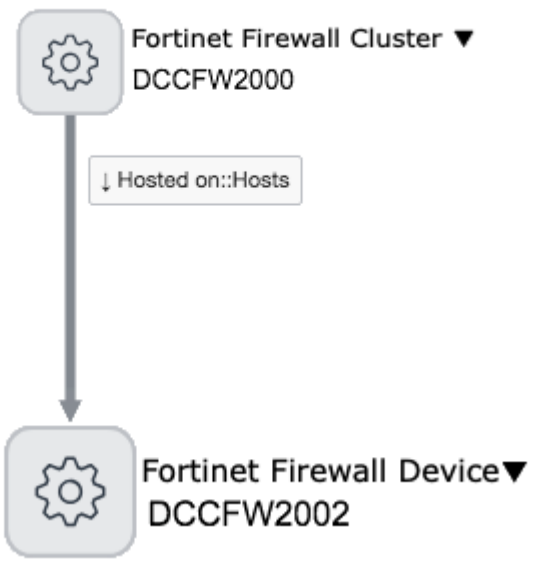

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Next Generation Fortinet Network Firewall pattern creates the following relationships and references to support Fortinet firewall discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Fortinet Firewall Cluster [cmdb_ci_firewall_cluster_fortinet]	Hosted on::Hosts	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]	Uses::Used by	Router Interface [dscy_router_interface]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	IP Address [cmdb_ci_ip_address]

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration item [configuration_item]	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
Router Interface [dscy_router_interface]	Configuration Item [cmdb_ci]	Fortinet Firewall Device [cmdb_ci_firewall_device_fortinet]
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]

Gunicorn pattern-based discovery

Discovery and Service Mapping Patterns finds Gunicorn WSGI HTTP server instances running on Linux servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify Gunicorn is installed and running on the Linux server

Verify the Gunicorn software is installed and running on a Linux server. The Gunicorn process must be active at the time of discovery.

Verify `sudo` access to run network commands

Verify the discovery credentials have permission to run the following command: *`sudo netstat -tulnp`*.

Create SSH credentials

For more information, see SSH credentials.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the GUNICORN pattern.

Gunicorn [cmdb_ci_gunicorn]

Field	Description
Name [name]	Name of the Gunicorn instance, in the format `gunicorn@<hostname>`.
Version [version]	Version of the Gunicorn server.
Installation directory [install_directory]	Installation directory of the Python environment containing Gunicorn.
TCP port(s) [tcp_port]	TCP port or ports on which the Gunicorn server is listening.

HAProxy load balancer discovery with patterns

The Discovery and Service Mapping Patterns application uses the HA Proxy pattern to find HAProxy Community load balancers running on a Linux server. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

HAProxy Community edition is an open-source load balancer that can run in either TCP or HTTP mode. Note that although Discovery supports HAProxy for HTTP load balancing, TCP load balancing isn't supported.

Starting with Discovery and Service Mapping Patterns 1.15.0 version, pattern-based discovery is the default method for discovering HAProxy load balancers. For information on probe-based HAProxy discovery, see HAProxy load balancer discovery using probes.

Note: HAProxy discovery supports one-to-one relationship mapping between a front-end port and a back-end port. If multiple front-ends are connected to the same back-end, the discovery process will only map the relationship from the lowest front-end port number in a defined range to the back-end port.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To learn about HAProxy load balancers and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

Verify HAProxy is installed on the Linux server

Verify the HAProxy software is installed and running on a Linux server.

Verify permission to run HAProxy command

Verify you have permission to run the following command: *<HAproxy Installation Path>/haproxy -v*.

Verify read access to HAProxy configuration file

Verify you have read access to the `haproxy.cfg` file.

Create SSH credentials

For more information, see SSH credentials.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the HA Proxy Pattern.

HAProxy Load Balancer [cmdb_ci_lb_haproxy]

Field	Description
Name [name]	Name of the HAProxy load balancer.
Version [version]	Version of the HAProxy load balancer.
Description [short_description]	Description of the HAProxy process, which is obtained by running the command: `haproxy -version`.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Name [name]	Name of the load balancer in the following format: `<LB name>:<LB port>`.
Port [port]	Port on which the load balancer is listening.
IP Address [ip_address]	IP address of the load balancer service.
Last port [last_port]	Highest port number used by the load balancer service.
Load balancer [load_balancer]	References the HAProxy Load Balancer [cmdb_ci_lb_haproxy] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Name [name]	Name of the load balancer pool.
Load balancer [load_balancer]	References the HAProxy Load Balancer [cmdb_ci_lb_haproxy] table.
Load balancing method [load_balancing_method]	Method used to load balance the pool, for example: Round Robin.
Most recent discovery [last_discovered]	Timestamp of the most recent discovery.
Install Status [install_status]	Install status of the load balancer pool, for example: Installed.

Load Balancer Pool Member [cmdb_ci_lb_pool_member]

Field	Description
Name [name]	Name of the load balancer pool member.
Service port [service_port]	Port used by the pool member.
IP Address [ip_address]	IP address of the pool member host.
Pool [pool]	References the Load Balancer Pool [cmdb_ci_lb_pool] table.
Load balancer [load_balancer]	References the HAProxy Load Balancer [cmdb_ci_lb_haproxy] table.

CI relationships

Discovery creates these relationships to support the HAProxy load balancer discovery.

CI	Relationship	CI
Load Balancer Pool [cmdb_ci_lb_pool]	Allocated to::Allocated from	Load Balancer Service [cmdb_ci_lb_service]
Load Balancer Pool [cmdb_ci_lb_pool]	Owns::Owned by	Load Balancer Pool Member [cmdb_ci_lb_pool_member]
Load Balancer Pool [cmdb_ci_lb_pool]	References	HAProxy Load Balancer [cmdb_ci_lb_haproxy]
Load Balancer Pool [cmdb_ci_lb_pool]	References	Load Balancer Pool Member [cmdb_ci_lb_pool_member]
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	References	HAProxy Load Balancer [cmdb_ci_lb_haproxy]
Load Balancer Service [cmdb_ci_lb_service]	References	HAProxy Load Balancer [cmdb_ci_lb_haproxy]
Load Balancer Service [cmdb_ci_lb_service]	Runs on::Runs	Hardware [cmdb_ci_hardware]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HPE BladeSystem Enclosure Discovery

ServiceNow Discovery uses the HPE BladeSystem Enclosure discovery pattern to discover BladeSystem, which is a line of Hewlett Packard Enterprise blade server machines. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the pattern version

Verify you're using the latest version of Discovery and Service Mapping Patterns.

Set SNMP credentials

Set SNMP credentials for the enclosure (to be used by the Classifier). For more information, see SNMP credentials.

Set SSH credentials

Set the SSH credentials for the enclosure's Onboard Administrator (to be used by the Discovery pattern). The account used for accessing the Onboard Administrator must have permissions to view all of the enclosure's bays, so it can gather the blade servers' information. For more information, see SSH credentials.

Verify the Onboard Administrator IP address

Verify the Onboard Administrator's IP address, which Discovery runs against.

Verify host machine reachability

Verify that the host machine is reachable from the MID Server.

Populate only the serial number in the Serial Number field

Starting with version 1.35.0 of Discovery and Service Mapping Patterns, you can populate the **Serial Number** field with only the serial number value, without the UUID suffix, by setting the *sn_itom_pattern.clear_hpe_serial_numbers* MID Server property to true. For more information, see Populate only the serial number for HPE BladeSystem.

Highlights and supported versions

The pattern has been tested against:

- HPE BladeSystem c7000 Enclosure G2 and G3
- HPE BladeSystem blades: HPE ProLiant BL460c
  - Generations 8, 9 and 10
  - ROM Versions: I31, I36, I41 (2014-2021)

Blades managed by a Virtual Connect Manager/Virtual Connect Enterprise Manager (VCM/VCEM):

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Have a logical serial number instead of a physical one. To relate the blade to the OS server, the system matches both physical and logical serial numbers against OS Server records.
- When discovering blade servers with the logical serial numbers, their logical serial numbers are inserted to **cmdb_serial_number** with the **logical** type. This is a new serial number type that was created to support these new serial numbers introduced by HP.

Note: The WMI provides the logical serial number as the bios serial number. ESXi servers are discovered by probes. The ESXiserver CI gets the logical serial number as **chassis**.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the HPE BladeSystem Enclosure Discovery pattern.

HPE BladeSystem Enclosure [cmdb_ci_hpe_bladesystem_enclosure]

Field	Description
Name [name]	The name of the enclosure.
UUID [uuid]	The unique identifier of the enclosure (HPE's proprietary identifier).
IP Address [ip_address]	The management IP of the enclosure (Onboard Administrator address).
Model ID [model_id]	The enclosure's type (model).
Manufacturer [manufacturer]	The enclosure's manufacturer name ("HP" or "HPE", depending upon the model).
Serial number [serial_number]	The serial number of the enclosure. The default format is **<enclosure's original serial>::<enclosure's UUID>**.*
Asset tag [asset_tag]	The asset tag as extracted from the enclosure's data.
Model number [model_number]	The enclosure's part number (HPE proprietary).
Midplane Spare Part Number [midplane_spare_part_number]	The spare part number (HPE proprietary) of the mid-plane.
PDU Type [pdu_type]	The type of Power Distribution Unit.
PDU Spare Part Number [pdu_spare_part_number]	The Power Distribution Unit spare part number (HPE proprietary).
Onboard Administrator Tray Type [onboard_admin_tray_type]	The type of the Onboard Administrator tray.
Onboard Administrator Tray Spare Part Number [onboard_admin_tray_spare_part_number]	The Onboard Administrator tray's spare part number (HPE proprietary).
Onboard Administrator Tray Serial Number [onboard_admin_tray_serial_number]	The Onboard administrator tray's serial number.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

* To populate the **Serial number** field with only the serial number value, see Populate only the serial number for HPE BladeSystem.

HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade]

Field	Description
Name [name]	The generated value, according to this convention: **<enclosure's name>/bay-<bay number>**
UUID [uuid]	The unique identifier of the blade server (HPE's proprietary identifier).
Logical UUID [uuid_logical]	The logical unique identifier of the HPE server blade (HPE's proprietary identifier).
Model ID [model_id]	The blade server's product name (model).
Manufacturer [manufacturer]	The blade server's manufacturer name ("HP" or "HPE" depending on the model).
Serial number [serial_number]	The serial number of the blade server. The default format is **<blade's original serial>::<blade's UUID>**.*
Logical serial number [serial_number_logical]	The logical serial number of the HPE server blade (HPE's proprietary identifier).
Asset tag [asset_tag]	The asset tag as extracted from the blade's data.
Model number [model_number]	The blade's part number (HPE proprietary).
RAM (MB) [ram]	The blade's RAM size.
ROM Version [rom_version]	The version of the blade's ROM.
System Board Spare Part Number [system_board_spare_part_number]	The blade's system board spare part number (HPE proprietary).
Boot Mode [boot_mode]	The blade's boot mode.
Secure Boot Status [secure_boot_status]	The blade's secure board status.
Management Processor Type [management_processor_type]	The blade's management processor type (iLOversion).
Management Processor Name [management_processor_name]	The blade's management processor hostname.
Management Processor Firmware Version [management_processor_firmware_version]	The blade's management processor firmware version.
Management Processor IP Address [amanagement_processor_ip_address]	The blade's management processor IP address.
Management Processor MAC Address [management_processor_mac_address]	The blade's management processor MAC address.
Power Management Controller Version [power_management_controller_version]	The blade's power management controller's version.
iLO Federation Capable [ilo_federation_capable]	The iLofederation capability enabled/disabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade] (continued)

Field	Description
Bay Number [bay_number]	The enclosure's bay number in which the blade is installed.

* To populate the **Serial number** field with only the serial number value, see Populate only the serial number for HPE BladeSystem.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
MAC Address [mac_address]	The blade's management processor's ethernet address (= iLOnetwork interface MAC address).
Name [name]	The blade's management processor's hostname.
IP Address [ip_address]	The blade's management processor's IP address.
Configuration Item [cmdb_ci]	References the HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade] table.

IP Address [cmdb_ci_ip_address]

Field	Description
Name [name]	The blade's management processor (iLO) IPv4 IP address.
IP Address [ip_address]	The blade's management processor's IP address.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.

The implementation of HPE BladeSystem Enclosure CMDB model includes three new CI Classes:

CMDB Model: (New CI Classes)

Field	Description
Enclosure [cmdb_ci_enclosure]	A general class to represent any enclosure entity. Extends Hardware CI class.
HPE BladeSystem Enclosure [cmdb_ci_hpe_bladesystem_enclosure]	A CI class for representing the enclosure entity.
HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade]	A CI class for representing the blade server entity.

CI relationships

Discovery creates these relationships to support HPE BladeSystem Enclosure discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
HPE BladeSystem Enclosure [cmdb_ci_hpe_bladesystem_enclosure]	Contains::Contained by	HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade]
HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
IP Address [cmdb_ci_ip_address]	references	Network Adapter [cmdb_ci_network_adapter]
Network Adapter [cmdb_ci_network_adapter]	references	HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade]
Serial Number [cmdb_serial_number]	references	HPE BladeSystem Enclosure [cmdb_ci_hpe_bladesystem_enclosure]
Serial Number [cmdb_serial_number]	references	HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade]

Related topics

SNMP-based queries

Populate only the serial number for HPE BladeSystem

Clear UUID serial numbers from existing HPE BladeSystem CIs

Populate only the serial number for HPE BladeSystem

You can exclude the universally unique identifier (UUID) from the **Serial number** field for discovered HPE BladeSystem Enclosure and HPE BladeSystem Blade configuration items (CIs).

Before you begin

Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.

Role required: discovery_admin

About this task

By default, Discovery populates the **Serial number** field for HPE BladeSystem Enclosure [cmdb_ci_hpe_bladesystem_enclosure] and HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade] CIs in the format **serial_number::UUID**. The *sn_itom_pattern.clear_hpe_serial_numbers* MID Server property controls this behavior. When set to true, Discovery populates the field with only the serial number value. The default value is false.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** column, search for `sn_itom_pattern.clear_hpe_serial_numbers`.
3. Select the *sn_itom_pattern.clear_hpe_serial_numbers* MID Server property.
4. In the **Value** field, enter `true`.
   To revert to the default behavior, set the value to `false`.
5. Select **Update**.

What to do next

Run Discovery again to apply the change.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To delete existing CI records with the **serial_number::UUID** serial number format, run the **Clear HPE Blade Serial Numbers** scheduled job. For more information, see Clear UUID serial numbers from existing HPE BladeSystem CIs.

Related topics

HPE BladeSystem Enclosure Discovery

Clear UUID serial numbers from existing HPE BladeSystem CIs

Clear UUID serial numbers from existing HPE BladeSystem CIs

Delete records with UUID-formatted serial numbers from HPE BladeSystem discovery tables.

Before you begin

Verify that you have at least version 1.94.3 of CMDB CI Class Models.

Verify that you have set the *sn_itom_pattern.clear_hpe_serial_numbers* MID Server property to true and run HPE BladeSystem discovery. For more information, see Populate only the serial number for HPE BladeSystem.

Role required: admin

About this task

After setting the *sn_itom_pattern.clear_hpe_serial_numbers* MID Server property to true and running HPE BladeSystem discovery, new CI records are created with only the serial number value. You can run the **Clear HPE Blade Serial Numbers** scheduled script to delete existing records with UUID-formatted serial numbers. The script affects the HPE BladeSystem Enclosure [cmdb_ci_hpe_bladesystem_enclosure], HPE BladeSystem Blade [cmdb_ci_hpe_bladesystem_blade], and Serial Number [cmdb_serial_number] tables.

Procedure

1. In the navigation filter, enter `sysauto_script.list`.
2. In the **Name** column, search for `Clear HPE Blade Serial Numbers`.
3. Select **Clear HPE Blade Serial Numbers**.
4. Select **Execute Now**.

Related topics

HPE BladeSystem Enclosure Discovery

Populate only the serial number for HPE BladeSystem

IBM Db2 on Linux database discovery

ServiceNow Discovery uses the Db2 on Linux discovery pattern to find IBM Db2 instances on Linux servers.

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Discovery uses this pattern to run horizontal discovery. The pattern is triggered automatically by classification using these Db2 identifiers: tcp_port, name, installed_dir.

Prerequisites

Configure SSH credentials for the Linux server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In addition, ensure that the user has the permissions to run the following commands:

Identification for Db2 UDB entry point types

```
". /" + $process.environmentVariables.HOME.value +
"/sqllib/db2profile; " +
$process.environmentVariables.HOME.value + "/sqllib/bin/
db2level"

"source " + $installed_dir + "db2profile"

"db2 list database directory | grep \"Database name\" |
grep -i \"" + $entry_point.instance + "\" | head -1"

"db2level | grep \"DB2 v\""
```

Identification for Db2 UDB entry point types - NONE

```
"cat " + $installed_dir + "/db2dump/db2diag.log | grep
\"INSTANCE:\" | grep " + $entry_point.instance + " | head
-1"

"db2level | grep \"DB2 v\""
```

Identification for Db2 on z/OS

```
"netstat -a | grep " + $entry_point.port

"ls -d /*/* |grep " + $id
```

Data collected by Discovery during horizontal discovery

The **Identification for Db2 UDB entry point types** and **Identification for Db2 UDB entry point types – NONE** sections collect the following tables and fields:

DB2 Subsystem [cmdb_ci_db_db2_instance]

Field	Description
Name [name]	The Db2 instance name.
TCP port(s) [tcp_port]	TCP port.
Install Directory [installed_dir]	Installation directory.
Version [version]	Db2 version.

The Identification for Db2 on Z/OS section collects the following tables and fields:

DB2 Subsystem [cmdb_ci_db_db2_instance]

Table and field	Description
Name [name]	The Db2 instance name.
TCP port(s) [tcp_port]	TCP port.
Install Directory [installed_dir]	Installation directory.

The Dependency Views map shows all discovered IBM Db2 instances in your organization and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

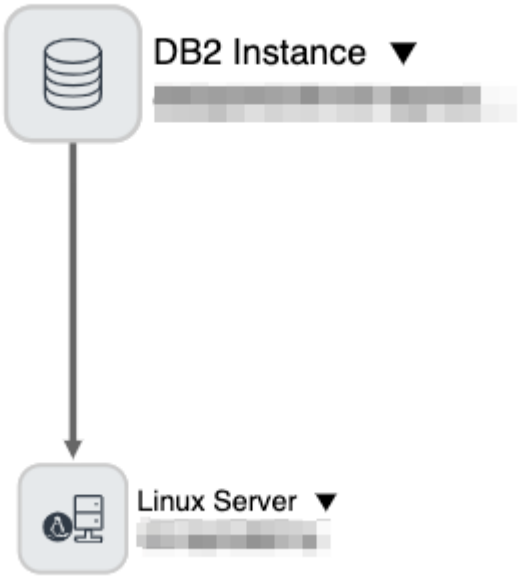

CI relationships

The IBM Db2 on Linux pattern does not create any CI relationships.

IBM Db2 on Windows database discovery

ServiceNow Discovery uses the Db2 on Windows discovery pattern to find IBM Db2 instances on Windows servers.

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Discovery uses this pattern to run horizontal discovery. The pattern is triggered automatically by classification using these Db2 identifiers: tcp_port, name, installed_dir

Prerequisites

Ensure that the user has the permissions to run the following commands:

- `"|"" + $installed_dir + "|bin|db2cmd|" /c /w /i db2 list database directory"`
- `"|"" + $installed_dir + "|bin|db2cmd|" /c /w /i db2 list database directory | findstr |"Database name|" | findstr /i " + $entry_point.instance`
- `"|"" + $installed_dir + "|bin|db2level|""`

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

Table and field	Description
name	Name of the Db2 instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
tcp_port	TCP port.
installed_dir	The installation directory.
version	Db2 version.
db_name	Name of the database.
pid	Process ID of the Db2 instance.

The Dependency Views map shows all discovered IBM Db2 instances in your organization and the relationships between them.

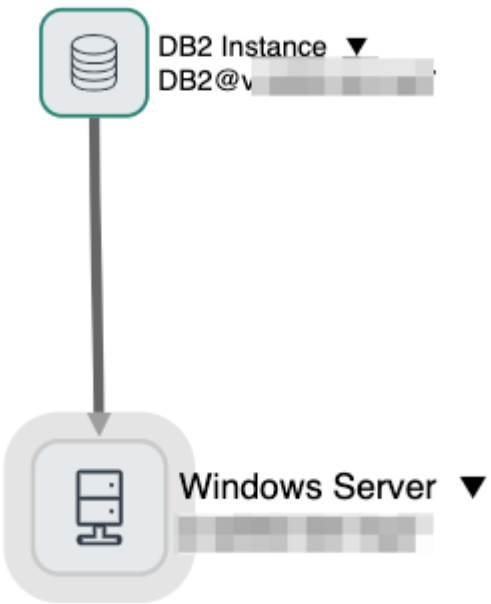

CI relationships

The IBM Db2 on Windows pattern does not create any CI relationships.

IBM Flash System pattern-based discovery

Discovery and Service Mapping Patterns finds IBM Flash System storage servers in your environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

IBM Flash System data model

The following diagram illustrates the tables and relationships that the Discovery and Service Mapping Patterns application creates when discovering IBM resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables and relationships populated during IBM Flash System discovery

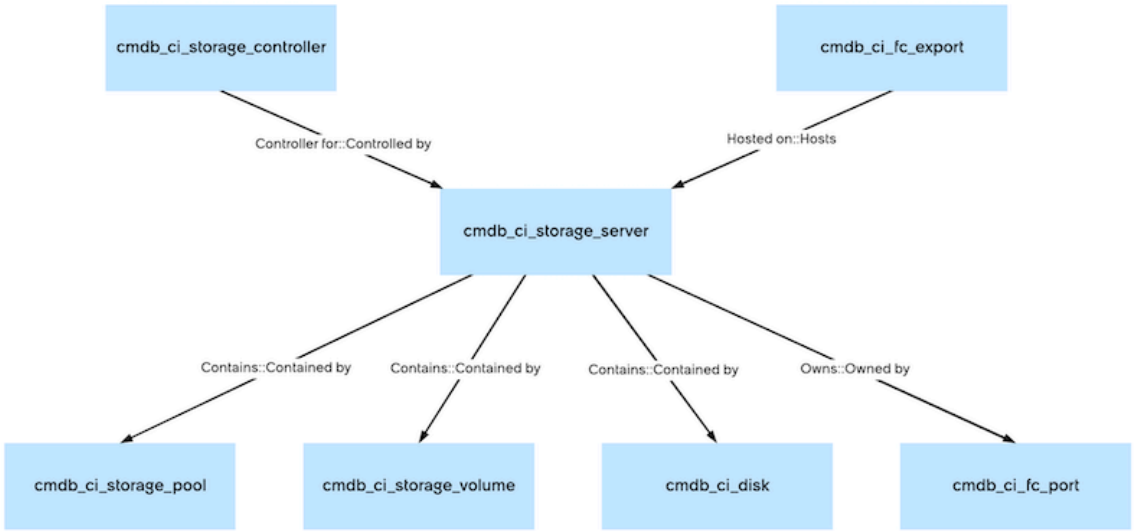

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models

Verify REST API access

Verify that the MID Server is active and can reach the following REST API endpoints on the IBM Flash System array:

- `POST /rest/v1/auth`
- `POST /rest/v1/lssystem`
- `POST /rest/v1/lsmdiskgrp`
- `POST /rest/v1/lsvdisk`
- `POST /rest/v1/lsportfc`
- `POST /rest/v1/lsenclosurecanister`
- `POST /rest/v1/lsdrive`
- `POST /rest/v1/lsfabric`

Create an alias for a basic authentication credential

For more information, see Create an alias for a basic authentication credential for IBM Flash System discovery.

Schedule a serverless discovery schedule

For more information, see Create a serverless discovery schedule for IBM Flash System discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the IBM Flash System pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server [cmdb_ci_storage_server]

Field	Description
Name [name]	Name of the storage server.
Model ID [model_id]	Hardware model of the storage server.
Manufacturer [manufacturer]	Name of the manufacturer. The value is set to **IBM**.
IP Address [ip_address]	Console IP address of the storage server.
Serial number [serial_number]	Serial number of the storage server.
OS Version [os_version]	Software version of the storage server.
Firmware version [firmware_version]	Software version of the storage server.
Physical Capacity (GB) [physical_capacity]	Total physical storage capacity, in gigabytes (GB).
Used Physical Capacity (GB) [used_physical_capacity]	Amount of physical storage capacity in use, in GB.
Free Physical Capacity (GB) [free_physical_capacity]	Amount of physical storage capacity available, in GB.
Virtual Capacity (GB) [virtual_capacity]	Total virtual storage capacity, in GB.
Free Virtual Capacity (GB) [free_virtual_capacity]	Amount of virtual storage capacity available, in GB.
Used Virtual Capacity (GB) [used_virtual_capacity]	Amount of virtual storage capacity in use, in GB.
Operational status [operational_status]	Operational status of the storage server. Default value is Operational.

Storage Pool [cmdb_ci_storage_pool]

Field	Description
Name [name]	Name of the storage pool.
Pool ID [pool_id]	Unique identifier of the storage pool.
Size [size]	Total capacity of the storage pool.
Size bytes [size_bytes]	Total capacity of the storage pool, in bytes.
Physical Capacity (GB) [physical_capacity]	Total physical capacity of the storage pool, in GB.
Free space [free_space]	Amount of free space available in the storage pool.
Free space bytes [free_space_bytes]	Amount of free space available in the storage pool, in bytes.
Operational status [operational_status]	Operational status of the storage pool.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	Name of the storage volume.
Serial number [serial_number]	Serial number of the storage volume.
Object ID [object_id]	Unique identifier of the storage volume.
State [state]	Current state of the storage volume. For example: Available or Offline.
Size bytes [size_bytes]	Total capacity of the storage volume, in bytes.
Operational status [operational_status]	Operational status of the storage volume.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.
Provided by [provided_by]	References the Storage Pool [cmdb_ci_storage_pool] table.

Fibre Channel Port [cmdb_ci_fc_port]

Field	Description
Name [name]	Name of the Fibre Channel port, derived from the port ID.
WWPN [wwpn]	World Wide Port Name of the Fibre Channel port.
Speed [speed]	Port speed of the Fibre Channel port.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.
Controller [controller]	References the Storage Controller [cmdb_ci_storage_controller] table.

Storage Controller [cmdb_ci_storage_controller]

Field	Description
Name [name]	Name of the storage controller node.
Device ID [device_id]	Canister ID of the storage controller.
Correlation ID [correlation_id]	Node ID used to correlate the storage controller with associated Fibre Channel ports.
Operational status [operational_status]	Operational status of the storage controller.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

Disk [cmdb_ci_disk]

Field	Description
Name [name]	Name of the disk.
Device ID [device_id]	Identifier of the disk drive.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk [cmdb_ci_disk] (continued)

Field	Description
Size bytes [size_bytes]	Capacity of the disk, in bytes.
Correlation ID [correlation_id]	Identifier used to correlate the disk with its managed disk group.
Device bus ID [device_bus_id]	Member ID of the disk within its enclosure.
Device interface [device_interface]	Slot ID of the disk within its enclosure.
Device target ID [device_target_id]	Drive class ID of the disk.
Operational status [operational_status]	Operational status of the disk.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

Fibre Channel Export [cmdb_ci_fc_export]

Field	Description
Name [name]	Name of the Fibre Channel export.
Initiator WWPN [initiator_wwpn]	WWPN of the remote initiator.
Export ID [export_id]	Local WWPN that identifies this Fibre Channel export.
Operational status [operational_status]	Operational status of the Fibre Channel export.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] table.

CI relationships and references

The IBM Flash System pattern creates the following relationships and references to support IBM Flash System storage server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Storage Server [cmdb_ci_storage_server]	Contains::Contained by	Storage Pool [cmdb_ci_storage_pool]
Storage Server [cmdb_ci_storage_server]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Storage Server [cmdb_ci_storage_server]	Contains::Contained by	Disk [cmdb_ci_disk]
Storage Server [cmdb_ci_storage_server]	Owns::Owned by	Fibre Channel Port [cmdb_ci_fc_port]
Storage Controller [cmdb_ci_storage_controller]	Controller for::Controlled by	Storage Server [cmdb_ci_storage_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Fibre Channel Export [cmdb_ci_fc_export]	Hosted on::Hosts	Storage Server [cmdb_ci_storage_server]

CI references

CI	Field	Referenced CI
Storage Pool [cmdb_ci_storage_pool]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server]
Storage Volume [cmdb_ci_storage_volume]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Storage Volume [cmdb_ci_storage_volume]	Provided by [provided_by]	Storage Pool [cmdb_ci_storage_pool]
Fibre Channel Port [cmdb_ci_fc_port]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Storage Controller [cmdb_ci_storage_controller]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Fibre Channel Port [cmdb_ci_fc_port]	Controller [controller]	Storage Controller [cmdb_ci_storage_controller]
Disk [cmdb_ci_disk]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Fibre Channel Export [cmdb_ci_fc_export]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server]

Create an alias for a basic authentication credential for IBM Flash System discovery

Create an alias for a basic authentication credential to run IBM Flash System discovery.

Before you begin

- Verify you have a basic authentication account with read-only access configured on the IBM Flash System array.
- Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.
- Select the Application scope icon and verify that the scope is set to **Discovery and Service Mapping Patterns**. The pattern launcher only resolves credentials created in the **Discovery and Service Mapping Patterns** scope.

Role required: discovery_admin

About this task

The IBM Flash System pattern uses basic authentication credentials to authenticate with the storage system. The pattern exchanges the username and password for a session token, which it uses for all subsequent REST calls to retrieve storage data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create an alias.
   a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   b. Select **New**.
   c. Enter a unique name for the alias and select **Credential** for the alias type in the **Type** field.
   d. Select **Submit**.
2. Configure a basic authentication credential for the new alias.
   a. Navigate to **Discovery** > **Credentials**.
   b. Select **New**.
   c. Select **Basic Auth Credentials**.
   d. Unlock the **Credential alias** field and select the alias you created.
   e. Configure the rest of the **Basic Auth Credentials** form fields.
   For more information about Basic Auth Credentials form fields, see Basic authentication credentials.
   f. Select **Submit**.

What to do next

Create a serverless discovery schedule. For more information, see Create a serverless discovery schedule for IBM Flash System discovery.

Related topics

IBM Flash System pattern-based discovery

Create a serverless discovery schedule for IBM Flash System discovery

Create a serverless discovery schedule to run IBM Flash System storage discovery.

Before you begin

- Verify that the MID Server is active and can reach the following REST API endpoints on the IBM Flash System array:
  - `POST /rest/v1/auth`
  - `POST /rest/v1/lssystem`
  - `POST /rest/v1/lsmdiskgrp`
  - `POST /rest/v1/lsvdisk`
  - `POST /rest/v1/lsportfc`
  - `POST /rest/v1/lsenclosurecanister`
  - `POST /rest/v1/lsdrive`
  - `POST /rest/v1/lsfabric`
- Create an alias for a basic authentication credential. For more information, see Create an alias for a basic authentication credential for IBM Flash System discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.
- Retrieve the following values from your IBM Flash System array:
  - Management IP address
  - REST API port (typically 7443)

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule form

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type, which should be `Serverless`.
MID Server selection method	Method that Discovery uses to select a MID Server: ▪ `Specific MID Cluster`: Use a preconfigured cluster of MID Servers. The MID Server can't be part of multiple clusters ▪ `Specific MID Server`: Use only one MID Server. If that MID Server is part of a cluster, only that MID Server is used. The cluster isn't used.
MID server	Name of the MID Server to use for this schedule. This field is available if **MID Server selection method** is set to `Specific MID Server`.
MID Server cluster	Name of the MID Server cluster to use for this schedule. This field is available if **MID Server selection method** is set to `Specific MID Cluster`.
Active	Option to enable this schedule for discovery.

   c. Select **Submit**.
3. Create the execution pattern.
   a. In the Discovery Schedules page, select the record you created.
   b. In the **Serverless Execution Patterns** tab, select **New**.
   c. On the form, fill in the fields.

Serverless Execution Pattern form

Field	Description
Name	Descriptive name for this record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Pattern	Pattern to use for this schedule, which should be `IBM Flash System Storage`.
Proxy Host	Fully qualified domain name of the machine on which a proxy server is installed, when applicable. Leave this field empty.
Active	Option to enable this schedule for discovery.

d. Select **Submit**.

4. Set the pattern launcher parameters.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Parameter	Value
credential_alias	ID of the credential alias you created.
ipaddress	Management IP of the IBM Flash System array.
port	REST API port of the IBM Flash System array.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

IBM Flash System pattern-based discovery

IBM Informix Dynamic Server discovery

The ServiceNow Discovery and Service Mapping applications can find and map the Informix Dynamic Server. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

User

Provide a user with permissions to run the following commands without elevated rights:

- `$install_directory/bin/onstat -g dis`

  This command gets the status and configuration of Informix Dynamic server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `$install_directory/bin/dbaccess sysmaster`

  This command runs an sql query against the master schema.

The same user must have read permission for the following directories:

- install_directory + "/etc/sqlhosts

  Access to this directory is necessary to get information about the port, host, and service definition of the Informix Dynamic server.

- /etc/services

  Access to this directory is required to get information about the host service definition.

Credentials

Configure SSH credentials in the Credentials module of the ServiceNow platform for accessing the server hosting the Informix Dynamic server.

Data collected by Discovery during horizontal discovery

Table and Field	Description
Informix Instance [cmdb_ci_db_informix_instance]	The Informix instance attributes.
tcp_port	
running_process	
version	
install_directory	
config_file	
name	
Informix Catalog [cmdb_ci_db_informix_catalog]	The database included in the Informix Instance.
database	

CI Relationships

CI	Relationship/Reference	CI
Informix Instance [cmdb_ci_db_informix_instance]	Contains::Contained by	Informix Catalog [cmdb_ci_db_informix_catalog]

Top-down

During top-down discovery, Service Mapping discovers connections between the Informix Instance and the Informix Catalogs that the instance contains. The connections display as an inclusion on the application service map.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM PowerHA SystemMirror for AIX discovery

The ServiceNow Discovery application uses the IBM PowerHA Cluster (HACMP) pattern to find IBM PowerHA SystemMirror for AIX (formerly HACMP) high-availability clusters on AIX UNIX and Linux systems running on IBM platforms. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses this pattern to run horizontal discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the following application is up to date

Discovery and Service Mapping Patterns

Verify the IBM PowerHA Cluster (HACMP) pattern is in Triggers probes tab

1. Navigate to **CI Classification** > **UNIX**.
2. Select **AIX**.
3. In The **Triggers probes** tab, verify that the IBM PowerHA Cluster (HACMP) pattern is listed as an exploration that runs after the AIX Server pattern.

For more information on AIX Server discovery, see AIX server discovery.

Verify server IP address or hostname

Verify that the server IP address or hostname are used during discovery. For more information, see Schedule a horizontal discovery.

Verify MID Server access

Verify that the MID Server has access to the target host and can execute all the commands that are implemented in the pattern.

Configure a UNIX user with the following permissions

Configure a UNIX user with permissions to run all commands defined in the official IBM documentation, and the following commands on AIX.

Command	Description
`lssrc -s clstrmgrES`	Verify that the clstrmgrES process is running
`lslpp -l cluster.license`	Get the license and version data of the cluster
`lslpp -l cluster.license | awk '{$1=$2=$3=\"\"; print $0}' |sed -e 's/^[[:space:]]*//'`	Get the software label version data of the cluster
`/usr/es/sbin/cluster/utilities/clRGinfo -a | awk`	Get resource groups name and type

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Description
`'{print $1,$2}' | grep -E -v 'denied|error'`	
`/usr/es/sbin/cluster/utilities/clRGinfo -s | grep -E -v 'denied|error'`	Get more information about resource groups
`/usr/es/sbin/cluster/utilities/clfindres -s | grep -E -v 'denied|error'`	Alternative command to get more information about resource groups
`/usr/es/sbin/cluster/utilities/clshowres | grep -E 'Resource Group Name|Volume' | grep -E -v 'Concurrent| Automatically|error|denied' | sed 's/Volume Groups/Volume Groups Names/' | awk '{print substr($0, index($0,$4))}' | sed 'N;s/\n/ /'`	Get volume groups associated to resource groups
`/usr/es/sbin/cluster/clstat -o | grep 'Resource Group' | grep -E -v 'error|denied' | awk '{print $3, $5, $6}'`	Get the state of the resource group(s)
`/usr/es/sbin/cluster/utilities/cltopinfo -c`	Get the cluster name
`/usr/es/sbin/cluster/utilities/cllscf`	Alternative command to get the cluster name
`lssrc -ls clstrmgrES`	Get cluster current state
`/usr/es/sbin/cluster/clstat -o |awk '{$1=$2=\"\"; print $0}' |sed -e 's/^[[:space:]]*//'`	Get the cluster ID
`/usr/es/sbin/cluster/utilities/cllsclstr | grep -E -v 'error|denied' | awk '{print $1}'`	Alternative command to get the cluster ID
`/usr/bin/odmget HACMPcluster| grep -w id | grep -E -v 'error|denied'`	Alternative command to get the cluster ID
`/usr/es/sbin/cluster/utilities/cltopinfo -n | grep -E -v 'error|denied'`	Get the cluster node(s) name(s)
`/usr/es/sbin/cluster/utilities/cllsnode | grep -E -v 'error|denied'`	Alternative command to get cluster node(s) name(s)
`/usr/es/sbin/cluster/clstat -o | grep -i 'Node:' | awk '{print $2, $4}'`	Get cluster node(s) state(s)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Description
`lscluster -m | grep -E 'Node name|State of node' | grep -E -v 'ERROR| lscluster' | awk '{print $2, $3, $4}' | sed 'N;s/\n/ /'`	Alternative command get cluster node(s) state(s)
`/usr/es/sbin/cluster/utilities/cllsif | grep -E -v 'error|denied' | awk '{print $1, $2, $6, $7}'`	Get network interface information
`lspv | grep -E -v 'error| denied'`	Get physical volumes information
`lsvg | grep -E -v 'error| denied'`	Get volume groups information
`lsvg -o | grep -E -v 'error|denied'`	Get information about the active volume groups
`/usr/es/sbin/cluster/utilities/clshowres | grep 'Concurrent Volume Group' | awk '{$1=$2=$3=\"\"; print $0}' |sed -e 's/^[[:space:]]*//'`	Get information about the concurrent volume groups
`hostname`	Get current node host name
`nslookup hostname`	Convert hostname to FQDN

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the IBM PowerHA Cluster (HACMP) pattern.

Cluster Resource Group [cmdb_ci_cluster_resource_group]

Field	Description
Name [name]	The name of the resource group configured into the cluster context.
Server [server]	The name of the PowerHA/HACMP cluster.
Resource Group Status [resource_group_status]	The status of the resource group into the cluster context. Example: Online.
Resource Group Type [resource_group_type]	The type of the resource group. Examples: Concurrent, non-concurrent.
Cluster [cluster]	The name of the PowerHA/HACMP cluster to which the resource group belongs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

UNIX Cluster Resource [cmdb_ci_unix_cluster_resource]

Field	Description
Name [name]	The name of the volume group into the cluster context.
Resource Type [resource_type]	The resource type (volume group).
Resource Status [resource_status]	The status of the volume group. Examples: Online, Offline.
Cluster [cluster]	The name of the PowerHA/HACMP cluster to which the volume group belongs.

UNIX Cluster Node [cmdb_ci_unix_cluster_node]

Field	Description
Name [name]	The name of the node that participates in the cluster.
Node State [node_state]	The state of the cluster node in the context of the CI state. For example: UP, DOWN, UNKNOWN.
Node Status [node_status]	The status of the node in the context of the resource group to which it belongs.
IP Address [ip_address]	The IP addresses assigned to the node.
Cluster [cluster]	The name of the cluster to which the cluster node belongs.

AIX Server [cmdb_ci_aix_server]

Field	Description
Name [name]	Name of the AIX server, representing the cluster node into the context of the PowerHA/HACMP cluster.

UNIX Cluster [cmdb_ci_unix_cluster]

Field	Description
Name [name]	The name of the PowerHA/HACMP cluster.
Cluster ID [cluster_id]	The ID of the PowerHA/HACMP cluster. Example: 0123456789.
Cluster Status [cluster_status]	The current state of the PowerHA/HACMP cluster. Examples: ST_STABLE, NOT_CONFIGURED.
Cluster Version [cluster_version]	The version of the PowerHA/HACMP cluster. Examples: 5.5; 6.1; 7.2.
Cluster Type [cluster_type]	The type of the PowerHA/HACMP cluster. Relevant only for versions 7.2 and above. Example: Stretched.
Description [short_description]	Description of the PowerHA/HACMP cluster (version, ID).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cluster Virtual IP [cmdb_ci_cluster_vip]

Field	Description
Name [name]	The cluster node interface name.
IP Address [ip_address]	Hostname or IP address where the Hostname represents the cluster node network interface name and the IP address the IPv4 address associated to the network interface name. (As the IBM PowerHA/HACMP cluster does not have dedicated VIP address).
Cluster [cluster]	The name of the PowerHA/HACMP cluster.

On the Dependency Views map you can see all discovered cluster CIs in your organization and the relationships between them.

In this example, UNIX Cluster Node represents a cluster of UNIX Clusters. UNIX Cluster Node gets resources from UNIX Cluster Resource and uses Cluster Virtual IP. UNIX Cluster Node contains three Cluster Resource Groups.

Dependency Views sample map

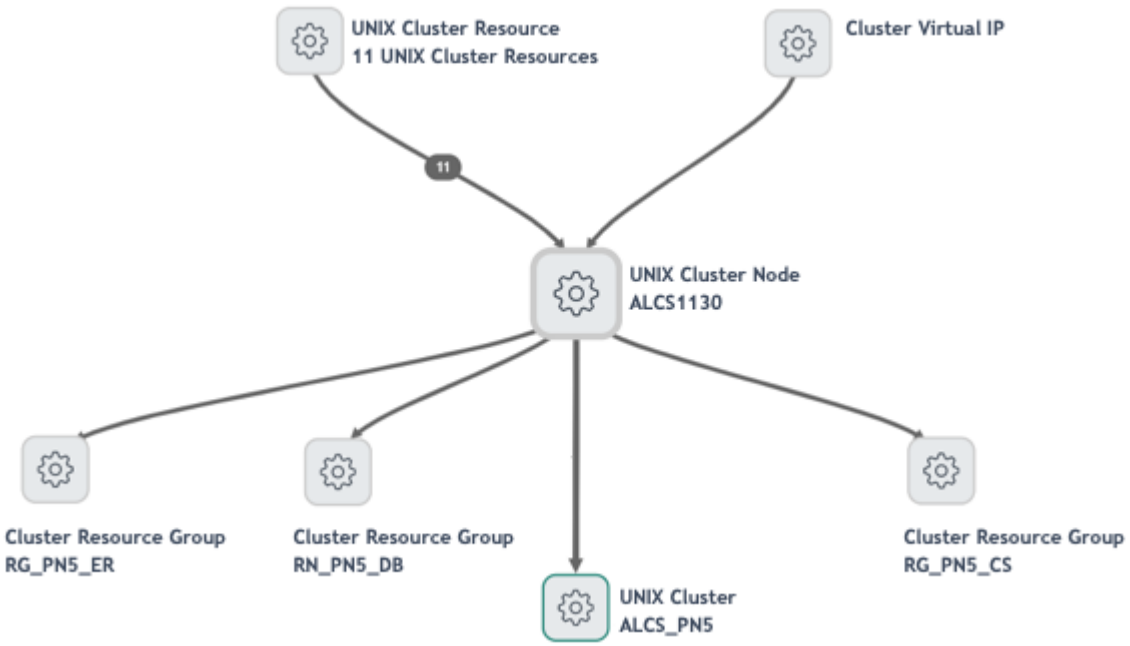

CI relationships

These relationships are created to support IBM HACMP/PowerHA cluster discovery:

CI	Relationship	CI
UNIX Cluster Node [cmdb_ci_unix_cluster_node]	Cluster of::Cluster	UNIX Cluster [cmdb_ci_unix_cluster]
UNIX Cluster Node [cmdb_ci_unix_cluster_node]	Contains::Contained by	Cluster Resource Group [cmdb_ci_cluster_resource_group]
Cluster Resource Group [cmdb_ci_cluster_resource_group]	Contains::Contained by	UNIX Cluster Resource [cmdb_ci_unix_cluster_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
UNIX Cluster Resource [cmdb_ci_unix_cluster_resource]	Defines resources for::Gets resources from	UNIX Cluster Node [cmdb_ci_unix_cluster_node]
UNIX Cluster Resource [cmdb_ci_unix_cluster_resource]	Defines resources for::Gets resources from	UNIX Cluster [cmdb_ci_unix_cluster]
Cluster Virtual IP [cmdb_ci_cluster_vip]	Used by::Uses	UNIX Cluster Node [cmdb_ci_unix_cluster_node]
UNIX Cluster [cmdb_ci_unix_cluster]	Virtualized by::Virtualizes	Cluster Virtual IP [cmdb_ci_cluster_vip]
UNIX Cluster [cmdb_ci_unix_cluster]	Contains::Contained by	Cluster Resource Group [cmdb_ci_cluster_resource_group]
UNIX Cluster Node [cmdb_ci_unix_cluster_node]	Hosted on::Hosts	AIX Server [cmdb_ci_aix_server]

IBM Virtualization and Hardware Management Console discovery

The Discovery and Service Mapping Patterns application uses the IBM HMC Server pattern to find information about Hardware Management Console (HMC), frame, and logical partition (LPAR) servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery supports the following configuration items (CIs):

- HMC server
- IBM frame
- LPAR instance

  An LPAR is the virtual machine equivalent of IBM.

- LPAR resource

Discovery is used on the HMC server and the individual AIX machines (LPARs).

The IBM HMC Server pattern runs using SSH. It's supported to run on IBM AIX OS for LPAR and on Linux for HMC servers only.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models

Verify permissions for HMC and AIX servers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Verify you have permissions for the following HMC commands:
  - `lssyscfg`
  - `lshmc`
  - `lshwres`
  - `lssysconn -r all -F type_model_serial_num:ipaddr`
- Verify you have permissions for the following AIX (LPAR) command: `lparstat -i`.

> **Note:** The `lparstat -i` command is in addition to the commands required for AIX server discovery. For more information on AIX server discovery, see AIX server discovery. For a list of the AIX commands, see Service Mapping commands requiring a privileged user.

Configure SSH credentials

Configure SSH credentials on the ServiceNow AI Platform. For more information, see SSH credentials.

Run a horizontal discovery for HMC Server and LPARs

1. Run a discovery on the HMC IP address.
2. Only when the HMC Server discovery has successfully completed, run a discovery on the LPARs using the AIX machine IP address.

For more information on running a discovery, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the IBM HMC Server pattern.

IBM HMC Server [cmdb_ci_hmc_server]

Field	Description
Name [name]	Name of the HMC.
Frame count [frame_count]	Number of frames managed for the HMC. A frame is the physical machine on which the LPARs reside. A frame is the equivalent of an ESX machine.
Serial number [serial number]	Serial number of the HMC server.
IP Address [ip_address]	IP address of the HMC.
OS Version [os_version]	Operating system version.
CPU type [cpu_type]	CPU type.
CPU speed (MHz) [cpu_speed]	CPU speed.
CPU count [cpu_count]	CPU count.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM HMC Server [cmdb_ci_hmc_server] (continued)

Field	Description
CPU core count [cpu_core_count]	CPU core count.
RAM (MB) [ram]	Memory RAM.

IBM Frame [cmdb_ci_ibm_frame]

Field	Description
Name [name]	Name of the frame.
Serial number [serial number]	Serial number of the Frame server.
Current available processor units [curr_avail_sys_proc_units]	Current available processor units.
Configurable processor units [configurable_sys_proc_units]	Configurable processor units.
Installed processor units [installed_sys_proc_units]	Installed processor units.
Current available memory units [curr_avail_sys_mem]	Current available memory units.
Configurable memory units [configurable_sys_mem]	Configurable memory units.
Installed memory units [installed_sys_mem]	Installed memory units.
IP Address [ip_address]	Management IP address of the IBM frame.

IBM LPAR Instance [cmdb_ci_lpar_instance]

Field	Description
Name [name]	Name of the LPAR instance.
State [state]	Operational state. Possible values are On or Off. • On: running • Off: not activated or not available
Serial number [serial_number]	Serial number of the LPAR.
Object ID [object_id]	Unique object identifier. This identifier is concatenated from the frame serial number and the LPAR identifier.
Is VIO [is_vio]	Identifies the LPAR as a VIO server. A VIO server is the LPAR that manages the physical components for other LPARs, such as network interface cards (NICs) and optical drives.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM LPAR Instance [cmdb_ci_lpar_instance] (continued)

Field	Description
VIO Servers [vio_servers]	• Client LPARs: Identifies the VIO servers that manage the client LPARs. • VIO servers: Identifies the Client LPARs the VIO servers manage.
Current Min Proc Units [curr_min_proc_units]	Minimum number of processing units that can be assigned to the partition dynamically.
Running Proc Units [run_proc_units]	Number of processing units that are varied on for the partition.
Current Memory [curr_mem]	The amount of memory in megabyte (MB) assigned to the partition.
Running Memory [run_mem]	Current amount of memory in MB that the partition has varied on.
IP Address [ip_address]	IP address assigned to the VIOS LPAR instance.
Model number [model_number]	OS version of the VIOS LPAR instance.
Frame Name [frame_parent]	References the IBM Frame [cmdb_ci_ibm_frame] table.

LPAR Resource [cmdb_ci_lpar_resource]

Field	Description
High Watermark VCPU [high_watermark_vcpu]	Highest recorded value of the **Online Virtual CPUs** field, as collected by the `lparstat -i` command. Because the value is dynamic, the results of a discovery could affect the value: • If the value of the **Online Virtual CPUs** field is lower than the value collected by the `lparstat -i` command, the value doesn't change. • If the value of the **Online Virtual CPUs** field is higher than the value collected by the `lparstat -i` command, the higher value is used.
Node name [node_name]	Name of the instance, along with the field that refers to the LPAR instance CI.

Serial Number [cmdb_serial_number]

Field	Description
Serial number [serial_number]	Serial number of the CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI Relationships

Discovery creates these relationships to support the IBM Virtualization and HMC discovery.

CI	Relationship	CI
IBM HMC Server [cmdb_ci_hmc_server]	Defines resources for::Gets resources from	IBM HMC Processor Pool [cmdb_ci_processor_pool]
IBM LPAR Instance [cmdb_ci_lpar_instance]	Virtualized by::Virtualizes	IBM Frame [cmdb_ci_ibm_frame]
IBM Frame [cmdb_ci_ibm_frame]	Managed by::Manages	IBM HMC Server [cmdb_ci_hmc_server]
IBM HMC Processor Pool [cmdb_ci_processor_pool]	Members::Member of	IBM LPAR Instance [cmdb_ci_lpar_instance]
IBM LPAR Instance [cmdb_ci_lpar_instance]	References	IBM Frame [cmdb_ci_ibm_frame]

IBM WebSEAL discovery

The ® Discovery application uses the IBM WebSEAL patterns to find WebSEAL applications, web application servers, and junctions on your infrastructure. Discovering some of these resources requires installing the Discovery and Service Mapping Patterns application from the ® Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Detailed information on IBM WebSEAL

For information, see Detailed information on products discovered by ITOM Visibility

Authentication

The WebSEAL authentication process includes HTTP header authentication with basic credentials Basic authentication credentials.

If you're running discovery for the first time, configure Credential affinity for Discovery and Orchestration Credential affinity for Discovery and Orchestration.

Note: Credential affinity isn't available In debug mode.

Configured MID Server

Ensure the MID Server has access and permissions to send HTTPS requests to the WebSEAL application.

Configured user permissions to execute API calls

Ensure that the user has the permission to run the following API:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `/net/general`
- `/net/dns`
- `/wga/reverseproxy`
- `/wga/reverseproxy/<reverse_proxy_id>/configuration/stanza/junction/entry_name/match-vhj-first`
- `/wga/reverseproxy/<reverse_proxy_id>/junctions`
- `/wga/reverseproxy/"<reverse_proxy_id>/junctions?junctions_id=<junction_id>`
- `/wga/widgets/health.json`

HTTP Classification enabled

The IBM WebSEAL Pattern is triggered when you run the discovery on the server that the WebSEAL is on. Ensure that the HTTP Classification is present in the instance. For more information, see Create an HTTP classification and Run discovery through an HTTP or HTTPS REST call

Data collected during horizontal discovery

Field	Description
ISAM Server [cmdb_ci_isam_server]	
name	The name of the server as returned by the `/net/general` API call. If the field is empty, `nslookup` on the IP address would be attempted to populate this field.
IP Address [ip_address]	The IP address according to the discovery schedule/configuration.
fqdn	The `search_domain` returned by the `/net/dns` API call.
WebSEAL [cmdb_ci_app_server_webseal]	
name	The name of the WebSEAL server as returned by the API.
IP Address [ip_address]	The IP address according to the discovery schedule/configuration.
fqdn	The `search_domain` returned by the `/net/dns` API call.
install_status	The status is set to **installed** by default.
operation_status	The status is set to **operational** by default.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
install_directory	According to the identification requirement, this field is populated with the FQDN of the server.
WebSEAL Reverse Proxies [cmdb_ci_webseal_reverse_proxy]	
name	The name of the resource according to the API response.
object_id	The ID of the resource, of the API response.
install_status	Indicates if the installation is enabled.
operational_status	Indicates if the operation has started.
service_type – hard-coded	The status is set to `Webseal Reverse Proxy` by default.
webseal_health_status	The WebSEAL reported health status regarding the `/wga/widgets/health.json` API call
WebSEAL Junctions [cmdb_ci_webseal_junction]	
name	The name of the resource according to the API response.
object_id	The id of the resource, created as a unique hash number corresponding to the case-sensitive name of the resource. **Note:** The case insensitivity of the queries may create duplicate fields when the **name** value is identical to the **Object ID** value. One of the identification attributes must be changed to have a unique value.
install_status	Hard-coded to installed status
operational_status	Hard-coded to operational status
stateful_junction	Boolean value indicating whether it's a stateful Junction
transparent_path_junction	Boolean value indicating whether it's a transparent path Junction
junction_type	List for the type of the Junction
WebSEAL Backend Servers [cmdb_ci_webseal_backend_server]	
name	the name of the resource, as per API response
object_id	the id of the resource, as per API response
install_status	the status of the resource as per the "enabled" value in the API response
operational_status	the status of the resource as per the "operation state" value

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
server_id	The UUID used to identify the junction web server
http_port	HTTP port of the back-end third-party server. Applicable when the junction type is tcp.
service_port	TCP port of the back-end third-party server. Default is 80 for TCP junctions and 443 for SSL junctions
Priority	The priority of the server (1-9). Default is 9.

Pattern Relations

CI	Relationship Type	CI
Webseal [cmdb_ci_app_server_webseal]	Runs on::Runs	ISAM Server [cmdb_ci_isam_server]
Webseal Reverse Proxy [cmdb_ci_webseal_reverse_proxy]	Hosted on::Hosts Reference[load_balancer]	Webseal [cmdb_ci_app_server_webseal]
Webseal Reverse Proxy [cmdb_ci_webseal_reverse_proxy]	Runs on::Runs Reference[isam_server]	ISAM Server [cmdb_ci_isam_server]
Webseal Junction [cmdb_ci_webseal_junction]	Allocated to::Allocates Reference [service]	Webseal Reverse Proxy [cmdb_ci_webseal_reverse_proxy]
Webseal Junction [cmdb_ci_webseal_junction]	Owns::Owned by Reference [pool]	Load Balancer Pool Member [cmdb_ci_webseal_backend_server]
Webseal Junction [cmdb_ci_webseal_junction]	Reference only [load_balancer]	Webseal [cmdb_ci_app_server_webseal]

Note: Currently, the relation between the back-end servers and actual server CIs in the CMDB isn't possible for HD.

IBM WebSphere Application Server discovery using patterns

The Discovery and Service Mapping Patterns application uses the Websphere On Windows and Websphere On Unix patterns to find IBM WebSphere Application Servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The IBM WebSphere Application Server is a software framework with middleware that hosts Java-based web applications. Discovery identifies and classifies information about the IBM WebSphere Application Server on both Windows and Linux computers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Starting with Visibility Content version 6.1.0, pattern-based discovery is the default method for discovering IBM WebSphere Application Servers. The WebSphere probes are disabled by default. For more information, see IBM WebSphere server discovery using probes.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites for Windows discovery

Verify that the following applications are up to date:

- Discovery and Service Mapping Patterns
- Visibility Content

Enable PowerShell on the MID Server

For more information, see Set up MID Servers to use PowerShell.

Verify access to files and directories

Verify access to the following files and directories:

- Installation directory of the product
- Configuration directory of the product
- `versionInfo.bat`
- `WAS.product`
- `BASE.product`
- `profile.version`
- `Express.product`
- `application.xml`
- `server.xml`
- `cell.xml`

Create Windows credentials

For more information, see Windows credentials.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Prerequisites for Linux discovery

Verify that the following applications are up to date:

- Discovery and Service Mapping Patterns
- Visibility Content

Enable SSH on WebSphere Application Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see the "Setting up SSH on Linux or UNIX servers" article in the "WebSphere Automation" section of the IBM documentation (https://www.ibm.com/docs/en/products).

Verify access to files and directories

Verify access to the following files and directories:

- Installation directory of the product
- Configuration directory of the product
- `oeconsol`
- `WAS_HOME`
- `CONFIG_ROOT`
- `versionInfo.sh`
- `WAS.product`
- `BASE.product`
- `profile.version`
- `Express.product`
- `application.xml`
- `server.xml`
- `cell.xml`

Create SSH credentials

For more information, see SSH credentials .

Verify elevated rights for IBM WebSphere Application Server commands

For more information, see IBM WebSphere Application Server (on UNIX) table in Service Mapping commands requiring a privileged user.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Websphere On Windows and Websphere On Unix patterns.

IBM Websphere [cmdb_ci_app_server_websphere]

Field	Description
Name [name]	Name of the WebSphere Application Server.
Version [version]	Current version of the WebSphere Application Server.
Node [node]	Name of the Node grouping the WebSphere Application Server or servers.
Cell [cell]	Name of the cell grouping one or more nodes of the WebSphere Application Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Websphere EAR [cmdb_ci_app_server_ws_ear]

Field	Description
Name [name]	Name of the enterprise application archive (EAR) file.
Installation directory [install_directory]	Installation directory of the EAR file.

Web Service [cmdb_ci_web_service]

Field	Description
Name [name]	Name of the web service.
Service ID [service_id]	ID of the web service.
Attributes [attributes]	Additional attributes of the web service.

Websphere Cell [cmdb_ci_websphere_cell]

Field	Description
Name [name]	Name of the WebSphere cell.
Cell ID [cell_id]	ID of the WebSphere cell.
Cell type [cell_type]	Type of the WebSphere cell.
Cell discovery protocol [cell_discovery_protocol]	Cell discovery protocol.
Description [short_description]	Description of the cell, if one has been provided. This field is populated in Linux discovery only.

CI relationships

Discovery creates these relationships to support the WebSphere Application Server discovery for both Windows and Linux systems.

CI	Relationship	CI
IBM Websphere [cmdb_ci_app_server_websphere]	Contains::Contained by	Websphere Cell [cmdb_ci_websphere_cell]
IBM Websphere [cmdb_ci_app_server_websphere]	Contains::Contained by	Websphere EAR [cmdb_ci_app_server_ws_ear]
Web Service [cmdb_ci_web_service]	Runs on::Runs	IBM Websphere [cmdb_ci_app_server_websphere]

InfiniBox Discovery

ServiceNow Discovery uses the InfiniBox pattern to find and map instances. Discovering these resources requires installing the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses multiple patterns to find InfiniBox storage system data. Discovering these resources requires installing the Discovery and Service Mapping Patterns and the CMDB CI Class Models applications from the ServiceNow Store, as each application is dependent on the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

other. You can use the patterns for discovering the InfiniBox components on the ServiceNow AI Platform using the Paris release or later.

The InfiniBox enterprise storage array delivers faster-than-all-flash performance, high availability, and multi-petabyte scale for mixed application workloads. The InfiniBox hardware redundancy design enables rapid recovery from any component failures.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Note: Discovery of 'Hosts' and 'Clusters' associated to the InfiniBox system are not supported from version 1.0.83 of the Discovery and Service Mapping Patterns (sn_itom_pattern) plugin.

- Install the following plugins:
  - CMDB CI Class Models plugin (Version 1.28.0 or later)
  - Discovery and Service Mapping Patterns plugin (Version 1.0.81 or later)
- Set Basic Auth credentials for the InfiniBox storage system
- Set a credential alias and associate it with the InfiniBox system credential record.
- Set discovery schedule type to Serverless.
- Make sure that the MID server has access to the target device.

Authentication

Configure the credentials of the Basic Auth type for authenticating the REST API calls. They are used for discovery of the CI types associated with the InfiniBox Storage system.

1. Navigate to **Connections & Credential Aliases**.
2. Create a new **Credential** record.
3. Create a new **Basic Auth** credential record from the Credentials related list form. Configure the **Basic Auth** parameters:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

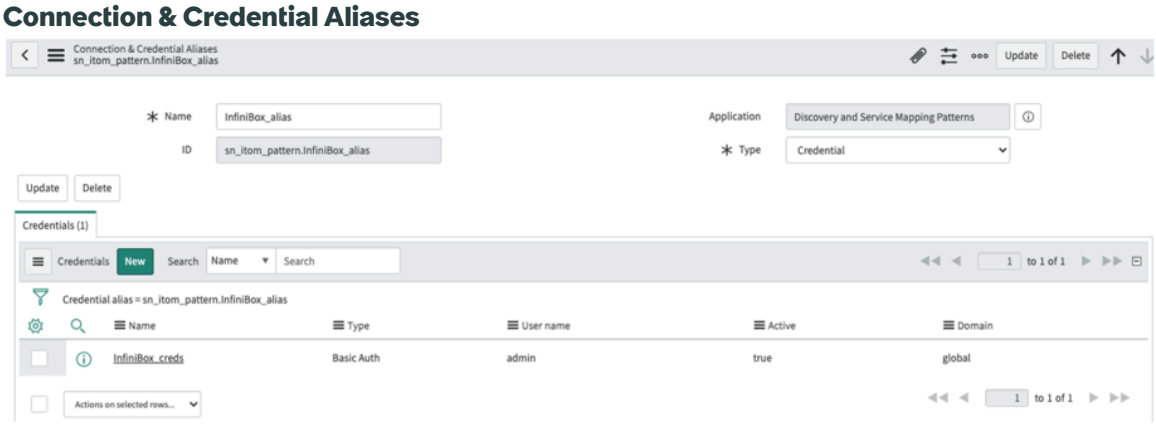

Configure the credentials of the **Basic Auth** type and the credential Alias associated with the credential record.

You can make the same configuration from the credential record form where you previously created the Credential alias. Associate the Connections and Credential Alias records respectively.

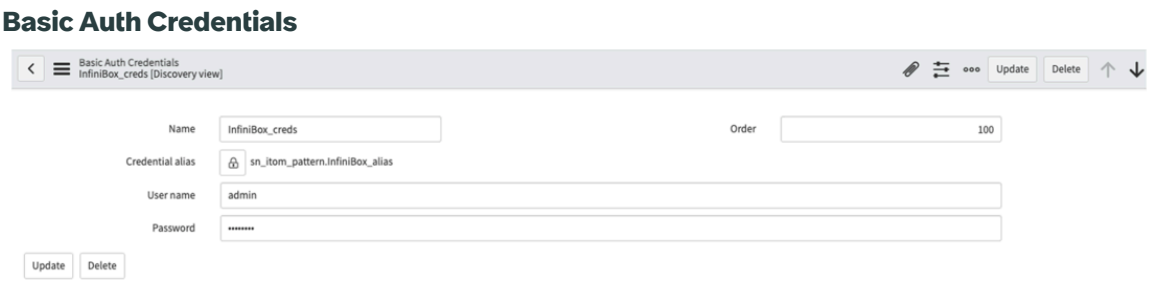

Discovery schedule

1. Create the Discovery schedule with the Serverless Discovery type.

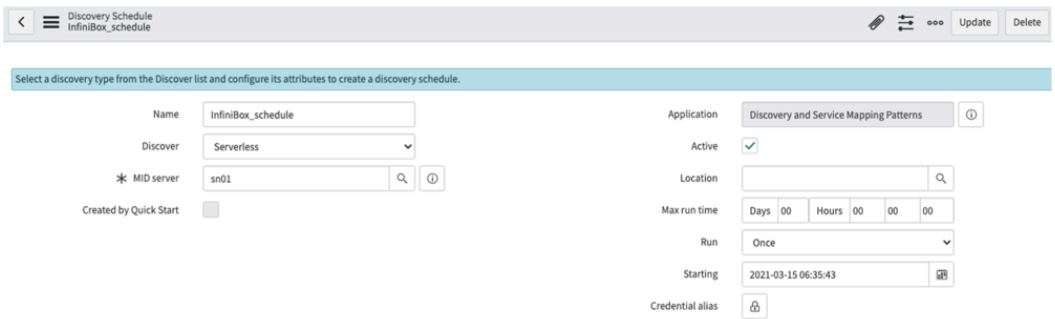

2. For the Discovery schedule with the Discover type of Serverless, configure the InfiniBox Storage System LP pattern within the Serverless Execution Patterns tab.
3. Set the Run Child Patterns to True.

   This enables the system to trigger all the patterns associated Child Patterns with the InfiniBox Storage System (LP). This is necessary due to the Large Payload (LP) mechanism. This is used for the discovery of the InfiniBox storage system and all the corresponding CI types.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

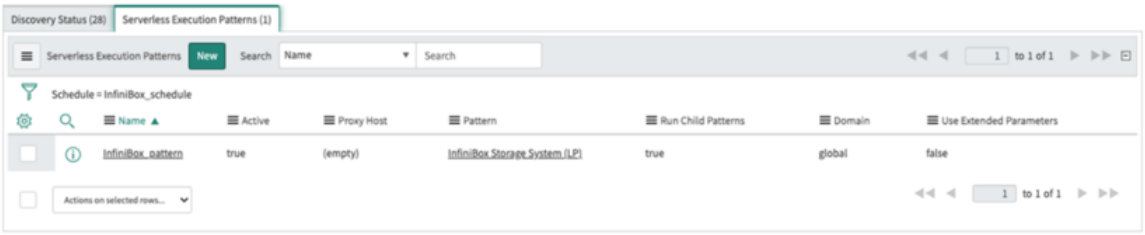

Apply the values for **ip_address** and **credential_alias** from the Serverless Execution Pattern related section. These serve as input parameters for triggering the discovery schedule and the associated InfiniBox discovery pattern.

Apply the IP Address of the InfiniBox system and the ID of the previously created Connections & Credential Aliases record. These are related to the InfiniBox credential **Basic Auth** type record.

Serverless Execution Pattern

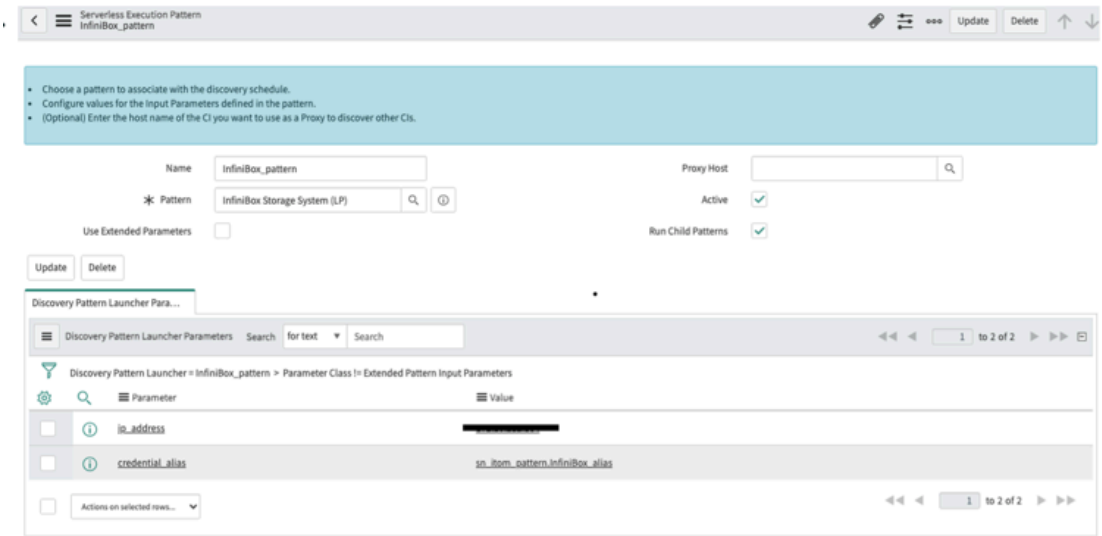

Data collected by Discovery during horizontal discovery

Storage Server [cmdb_ci_storage_server]

Field	Description
ip_address	The IP address of the InfiniBox system.
name	The name of the InfiniBox system.
os_version	The current version of the InfiniBox system.
model_id	The model ID of the InfiniBox system.
manufacturer	The product name of the InfiniBox system.
serial_number	The serial number of the InfiniBox system.
operational_status	The operational state of the InfiniBox system.
physical_capacity	The total physical capacity of the InfiniBox system in gigabytes.
virtual_capacity	The total virtual capacity of the InfiniBox system in gigabytes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server View Map

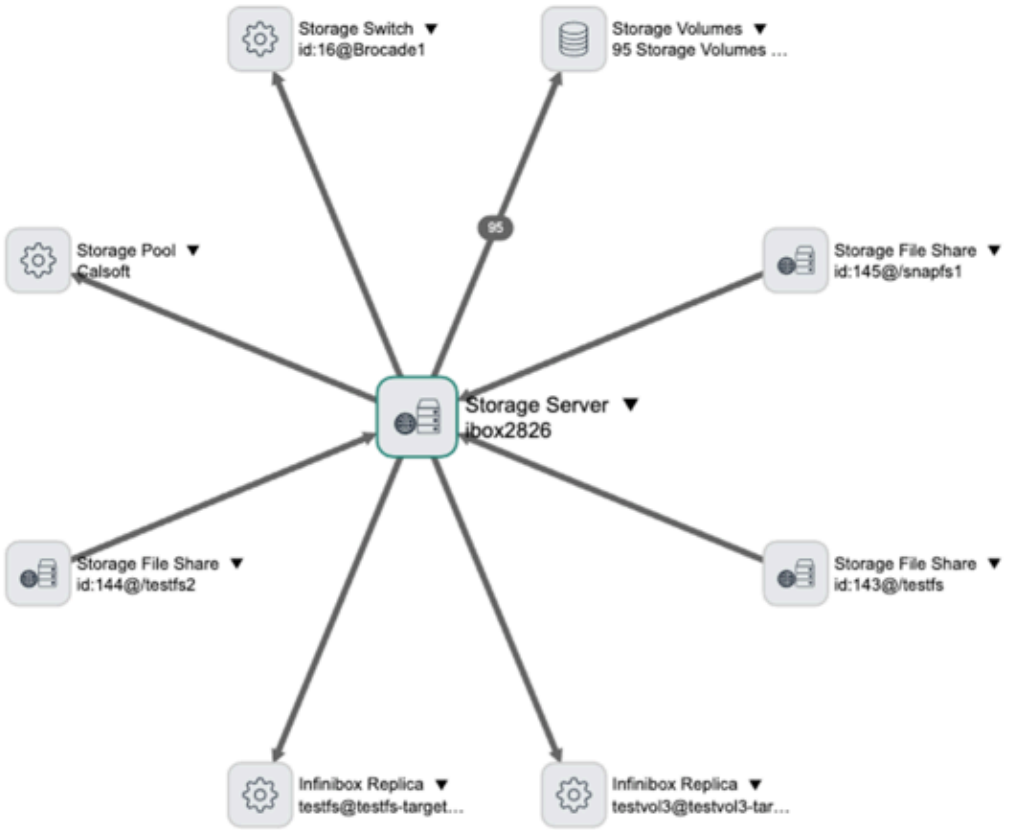

Disk [cmdb_ci_disk]

Field	Description
device id	The ID of the physical disk device, combination of NODE_ID@DISK_ID.
name	The name of physical disk device, combination of INFINIBOX_NAME:NODE_ID@DISK_ID
drive_type	The drive type of the physical disk device.
model_id	The model ID of the physical disk device.
manufacturer	The vendor of the physical disk device.
operational_status	The operational state of the physical disk device.
serial_number	The serial number of the physical disk device.
computer	Reference to any of the physical nodes part of the InfiniBox system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fibre Channel Port [cmdb_ci_fc_port]

Field	Description
wwpn	The world wide port name of fibre channel port.
name	The name of fibre channel port (wwpn).
speed	The speed value of the fibre channel port in gigabits per second, if applied.
wwnn	The world wide node name of fibre channel port.
port_role	The role of fibre channel port.
state	The state of fibre channel port.
link_state	The state of the connectivity of fibre channel port.
manufacturer	The vendor of the fibre channel port.
model_id	The model ID of the of fibre channel port.
computer	The reference to any of the physical nodes part of the InfiniBox system.

Fibre channel view map

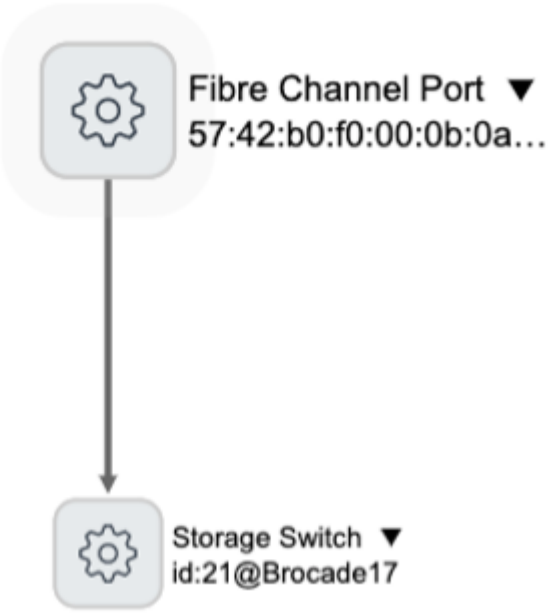

Storage Switch [cmdb_ci_storage_switch]

Field	Description
object_id	The ID of the storage switch.
name	The name of storage switch, combination of id:SWITCH_ID@SWITCH_NAME.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Switch [cmdb_ci_storage_switch] (continued)

Field	Description
wwnn	The world wide node name of the storage switch.
manufacturer	The vendor of the storage switch.

Storage Switch View Map

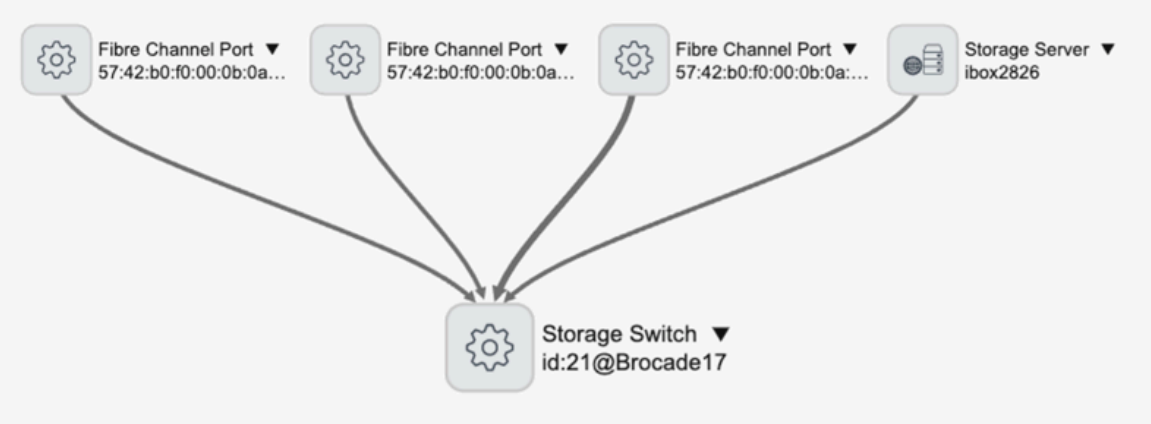

Network Adapter [cmdb_ci_network_adapter]

Field	Description
mac_address	The ethernet address of the network adapter.
name	The ethernet name of the network adapter.
operational_status	The ethernet link state of the network adapter.
manufacturer	The ethernet vendor of the network adapter.
model_id	The ethernet model ID of the network adapter.
ip_address	The IPv4 IP address associated with the ethernet network adapter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network Adapter View Map

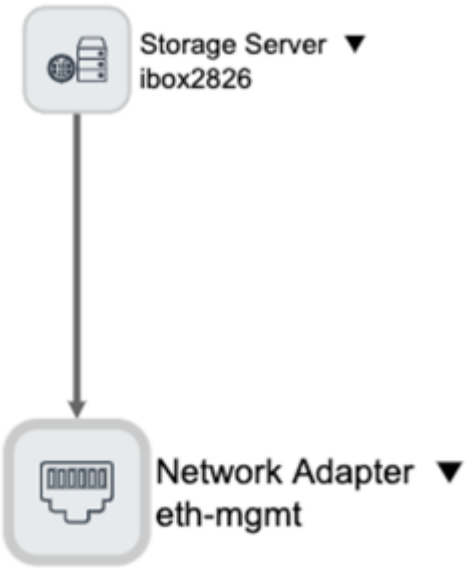

IP Address [cmdb_ci_ip_address]

Field	Description
netmask	The netmask associated with the InfiniBox storage system.
name	The IPv4 IP address of any of the InfiniBox nodes or the InfiniBox storage system.
ip_address	The IPv4 IP address of any of the InfiniBox nodes.

Storage Node Element [cmdb_ci_storage_node_element]

Field	Description
object_id	The name of the InfiniBox Storage Node, combination of id:NODE_ID@NODE_NAME.
name	The name of the InfiniBox Storage Node, combination of INFINIBOX_NAME@NODE_NAME.
node_state	The operational state of the InfiniBox Storage Node.
manufacturer	The vendor of the InfiniBox Storage Node.
model_id	The model of the InfiniBox Storage Node.
ip_address	The IPv4 IP address assigned to the InfiniBox Storage Node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Node Element View Map

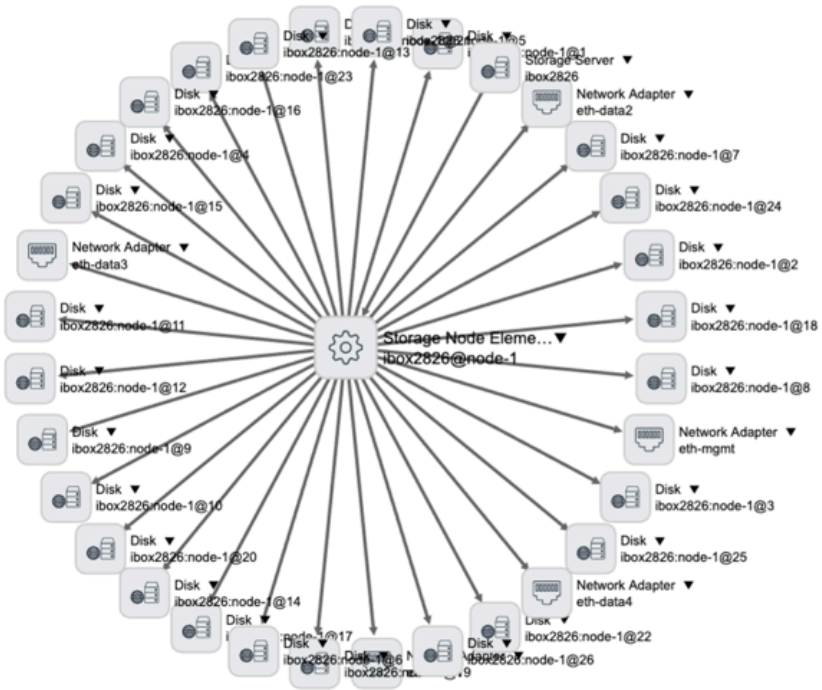

Serial Number [cmdb_serial_number]

Field	Description
serial_number	The serial number associated to the InfiniBox storage system.
serial_number_type	The serial number type of the InfiniBox storage system.
valid	The serial number valid type of the InfiniBox storage system.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
object_id	The ID of the storage volume, combination of id:VOLUME_ID@VOLUME_NAME.
volume_id	The ID of the storage volume, combination of id:VOLUME_ID@VOLUME_NAME.
name	The name of the storage volume.
storage_type	The role of the storage volume.
size_bytes	The total size of the storage volume in bytes.
serial_number	The serial number of the storage volume.
pool_id	The ID of the pool to which the storage volume is associated.
state	The lock state of the storage volume.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume [cmdb_ci_storage_volume] (continued)

Field	Description
provisioning_type	The type of provisioning of the storage volume.
size	The total size of the storage volume in gigabytes
provided_by	The reference to the storage pool, to which the storage volume is associated.

Storage Volume View Map

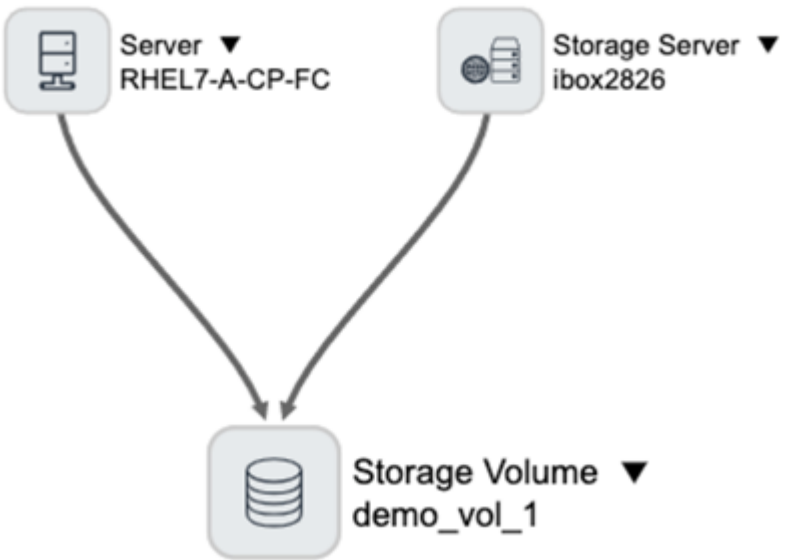

Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]

Field	Description
object_id	The ID of the storage volume snapshot, combination of id:VOLUME_ID@VOLUME_NAME.
storage_server	The name of InfiniBox system, to which the volume snapshot is associated.
name	The name of the storage volume snapshot.
size_bytes	The total size of the storage volume snapshot in bytes.
size	The total size of the storage volume snapshot in gigabytes.
serial_number	The serial number of the storage volume snapshot.
pool_id	The ID of the pool, to which the storage volume snapshot is associated.
parent_id	The ID of the storage volume, parent of the storage volume snapshot.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot] (continued)

Field	Description
state	The lock state of the storage volume snapshot.
provisioning_type	The type of provisioning of the storage volume snapshot.
provided_by	The reference to the storage pool, to which the storage volume is associated.

Storage Volume Snapshot

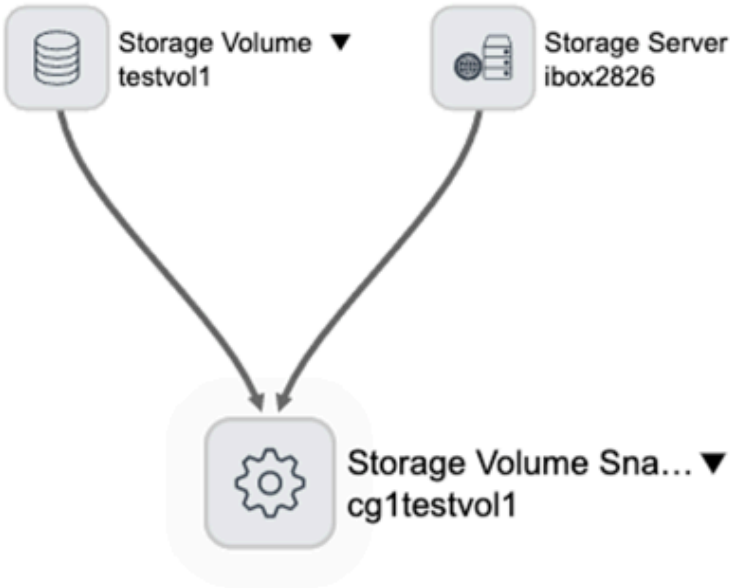

Storage Pool [cmdb_ci_storage_pool]

Field	Description
object_id	The ID of the storage volume, combination of id:POOL_ID@POOL_NAME.
name	The name of the storage pool.
physical_capacity	The total physical capacity of the storage pool in gigabytes.
virtual_capacity	The total virtual capacity of the storage pool in gigabytes.
free_phys_capacity	The available physical capacity of the storage pool in gigabytes.
free_virt_capacity	The available virtual capacity of the storage pool in gigabytes.
num_vols	The number of the volumes participating into the storage pool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Pool [cmdb_ci_storage_pool] (continued)

Field	Description
num_file_sys	The number of the file systems participating into the storage pool.
num_file_sys_snap	The number of the file system snapshots participating into the storage pool.
num_vol_snap	The number of the volume snapshots participating into the storage pool.
pool_id	The ID of the storage pool.
pool_state	The state of the storage pool.

Storage Pool View Map

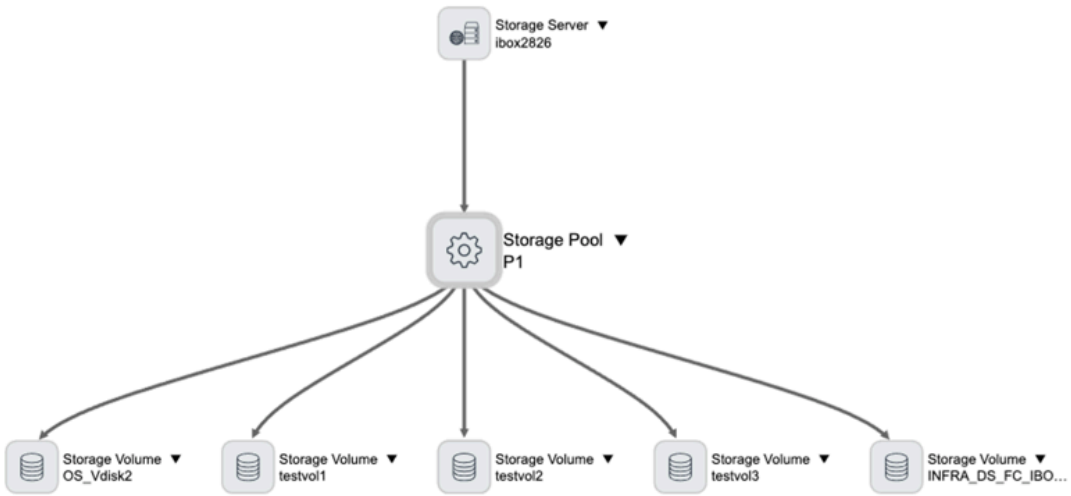

File System [cmdb_ci_file_system]

Field	Description
object_id	The ID of the primary file system, combination of id:FILE_SYSTEM_ID.
server_name	The name of InfiniBox system, to which the primary file system is associated.
name	The name of the primary file system, combination of id:FILE_SYSTEM_ID@FILE_SYSTEM_NAME.
size_bytes	The total size of the primary file system in bytes.
size	The total size of the primary file system in gigabytes.
free_space_bytes	The free space of the primary file system in bytes.
serial_number	The serial number of the primary file system.
storage_type	The type of the primary file system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File System [cmdb_ci_file_system] (continued)

Field	Description
pool_id	The ID of the pool, to which the primary file system is associated.
provisioning_type	The type of provisioning of the primary file system.
provided_by	The reference to the storage pool, to which the primary file system is associated.

File System View Map

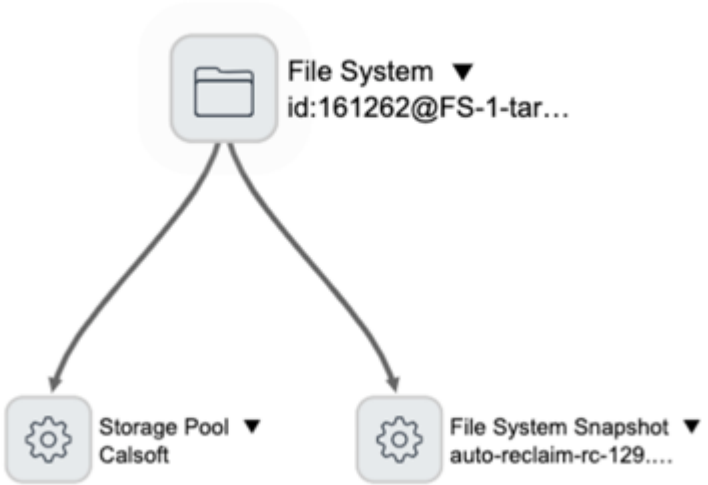

File System Snapshot [cmdb_ci_file_system_snapshot]

Field	Description
object_id	The ID of the snapshot file system, combination of id:FILE_SYSTEM_ID@FILE_SYSTEM_NAME.
volume_id	The ID of the snapshot file system, combination of id:FILE_SYSTEM_ID@FILE_SYSTEM_NAME.
server_name	The name of InfiniBox system, to which the snapshot file system is associated.
name	The name of the snapshot file system.
size_bytes	The total size of the snapshot file system in bytes.
free_space_bytes	The free space of the snapshot file system in bytes.
serial_number	The serial number of the snapshot file system.
storage_type	The type of the snapshot file system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File System Snapshot [cmdb_ci_file_system_snapshot] (continued)

Field	Description
pool_id	The ID of the pool, to which the snapshot file system is associated.
parent_id	The ID of the primary file system, to which the snapshot file system is associated.
provisioning_type	The type of provisioning of the snapshot file system.

File System Snapshot View Map

Storage File Share [cmdb_ci_storage_fileshare]

Field	Description
path	The export path of the InfiniBox export entity.
name	The name of the InfiniBox export entity, combination of id:EXPORT_ID@EXPORT_PATH.
fileshare_id	The ID of the export entity in the context of the InfiniBox system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Export View Map

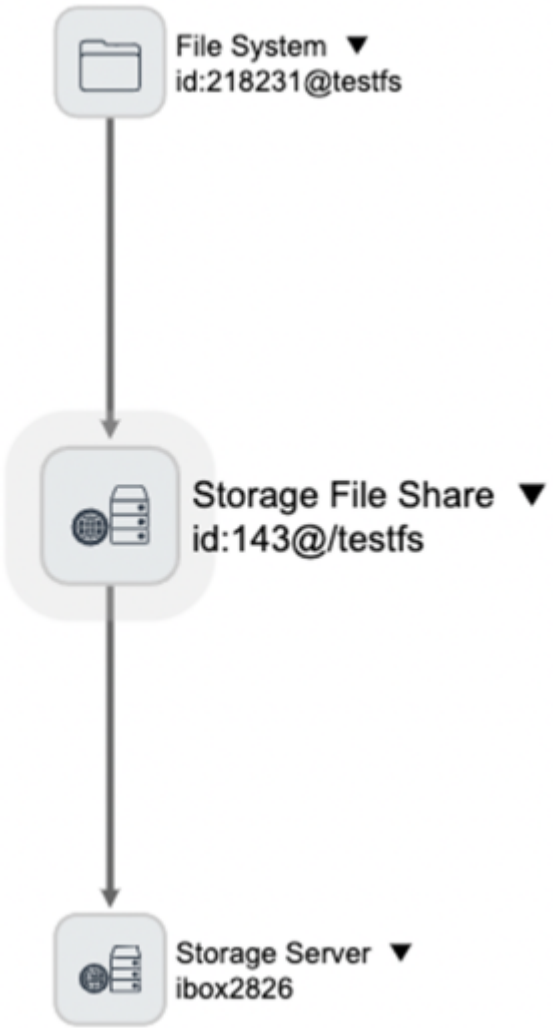

InfiniBox Primary Consistency Group [cmdb_ci_infinibox_cg]

Field	Description
object_id	TheID of the primary InfiniBox consistency group, combination of id:CONSISTENCY_GROUP_ID@ CONSISTENCY_GROUP_NAME.
created_time	The creation time of the primary InfiniBox consistency group in format YY-MM-DD HH:mm::SS.
updated_time	The updated time of the primary InfiniBox consistency group in the format YY-MM-DD HH:mm::SS.
name	The name of the primary InfiniBox consistency group.
type	The type of the primary InfiniBox consistency group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

InfiniBox Primary Consistency Group [cmdb_ci_infinibox_cg] (continued)

Field	Description
members_count	The count of the entities participating in the primary InfiniBox consistency group.
pool_id	The ID of the pool, to which the primary InfiniBox consistency group is associated.
lock_state	The lock state of the primary InfiniBox consistency group.

Primary Consistency Group View Map

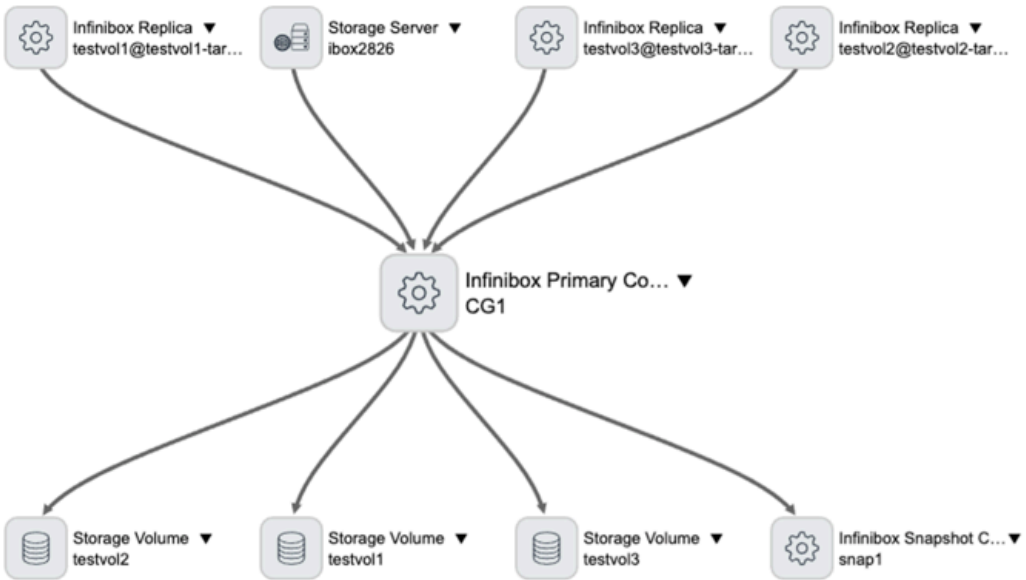

InfiniBox Snapshot Consistency Group[cmdb_ci_infinibox_snapshot_cg]

Field	Description
object_id	The ID of the snapshot InfiniBox consistency group, combination of id:CONSISTENCY_GROUP_ID@ CONSISTENCY_GROUP_NAME.
created_time	The creation time of the snapshot InfiniBox consistency group in format YY-MM-DD HH:mm::SS.
updated_time	The updated time of the snapshot InfiniBox consistency group in format YY-MM-DD HH:mm::SS.
name	The name of the snapshot InfiniBox consistency group.
type	The type of the snapshot InfiniBox consistency group.
members_count	The number of entities participating in the snapshot InfiniBox consistency group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

InfiniBox Snapshot Consistency Group[cmdb_ci_infinibox_snapshot_cg] (continued)

Field	Description
pool_id	The ID of the pool, to which the snapshot InfiniBox consistency group is associated.
lock_state	The lock state of the snapshot InfiniBox consistency group.

Snapshot Consistency Group View Map

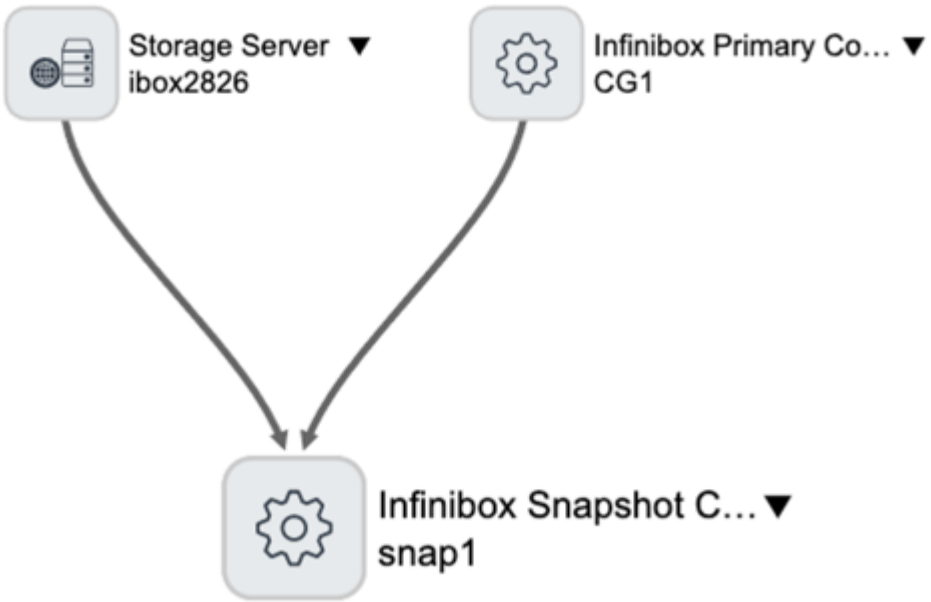

InfiniBox Link [cmdb_ci_infinibox_link]

Field	Description
object_id	The ID of the InfiniBox link.
name	The name of the InfiniBox link.
remote_version	The version of the remote InfiniBox storage system.
remote_host	The host name or IP address of the remote InfiniBox storage system.
remote_ip	The IP address of the remote InfiniBox storage system.
remote_link_id	The link ID of the remote InfiniBox storage system.
remote_serial_number	The serial number of the remote InfiniBox storage system.
remote_sys_name	The system name of the remote InfiniBox storage system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

InfiniBox Link [cmdb_ci_infinibox_link] (continued)

Field	Description
link_state	The state of the link between the InfiniBox storage systems.
link_config_id	The link configuration ID of the remote InfiniBox storage system in GUID format.
link_mode	The mode of the link between the InfiniBox storage systems.
resiliency_mode	The resiliency mode between the InfiniBox storage systems.
local_witness_state	The local witness state of the InfiniBox storage system.
witness_address	The local witness address of the InfiniBox storage system, if assigned.
last_connection_time	The time stamp of the established connection between the Infinibox storage systems in format YY-MM-DD HH:MM::SS.
local_replica_net_space_id	The local replica network space ID of the Infinibox storage system.
link_replication_type	The replication type of the link between the Infinibox storage systems.

Link View Map

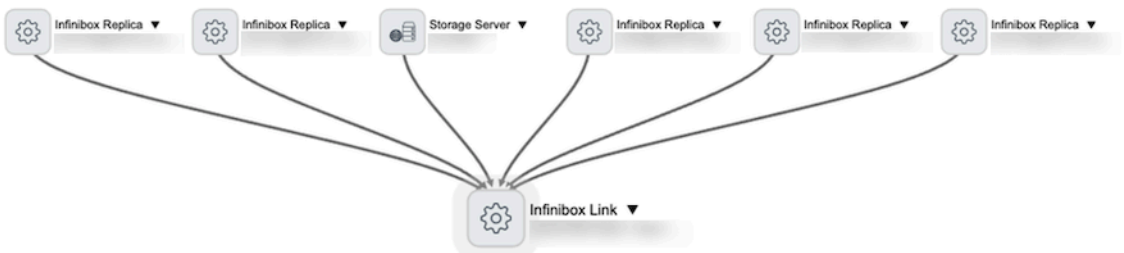

InfiniBox Replica [cmdb_ci_infinibox_replica]

Field	Description
object_id	The ID of the InfiniBox replica, combination of LOCAL_ENTITY_NAME@REMOTE_ENTITY_NAME@replica:REPLICA_ID.
name	The ID of the InfiniBox replica, combination of LOCAL_ENTITY_NAME@REMOTE_ENTITY_NAME@replica:REPLICA_ID.
replica_role	The role of the InfiniBox storage system replica.
entity_type	The type of the entity replicated in the context of the InfiniBox storage system.
replica_state	The state of the replica in the context of the InfiniBox storage system

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

InfiniBox Replica [cmdb_ci_infinibox_replica] (continued)

Field	Description
replication_type	The type of the replication of the context of the InfiniBox storage system, if applied.
sync_state	The synchronization state of the InfiniBox storage system replica.
link_id	The ID of the link associated to the replica in the context of the InfiniBox storage system.
entity_id	The ID of the replicated entity in the context of the InfiniBox storage system.
remote_entity_id	The ID of the remote replicated entity of the InfiniBox storage system.
entity_name	The name of the local replicated entity in the context of the InfiniBox storage system.
remote_entity_name	The name of the remote replicated entity in the context of the InfiniBox storage system.
dataset_type	The local dataset type in the context of the InfiniBox storage system.

Replica View Map

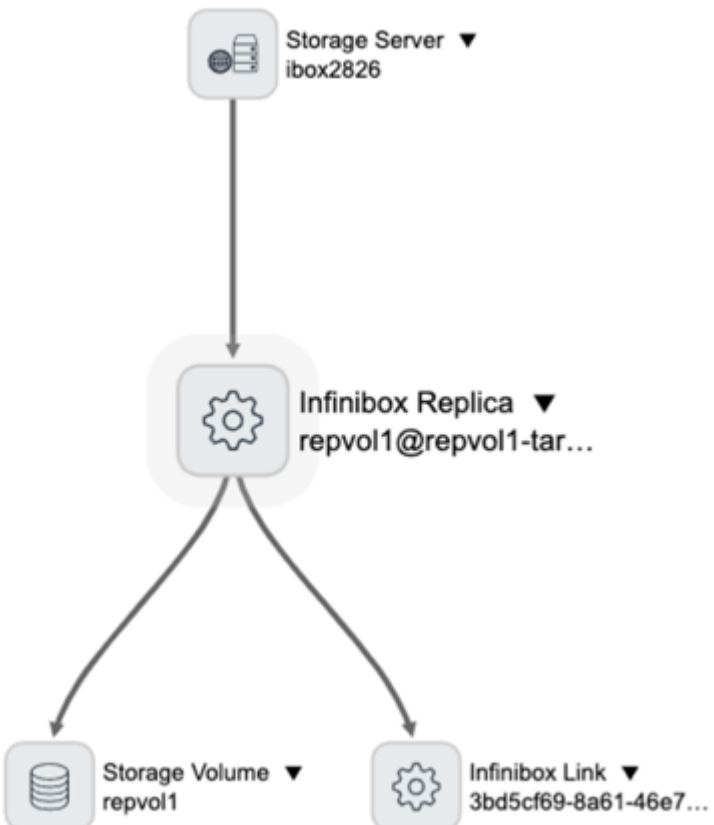

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI Relationships

Relationships created by the InfiniBox Storage System (LP) pattern

CI	Relationship	CI
cmdb_ci_storage_server	"Contains :: Contained by"	cmdb_ci_storage_switch
cmdb_ci_fc_port	“Connects to :: Connected by”	cmdb_ci_storage_switch
cmdb_ci_storage_server	“Uses :: Used by”	cmdb_ci_network_adapter
cmdb_ci_storage_node_element	“Owns :: Owned by”	cmdb_ci_network_adapter
cmdb_ci_ip_address	“reference by column name nic”	cmdb_ci_network_adapter
cmdb_ci_storage_node_element	“Owns :: Owned by”	cmdb_ci_fc_port
cmdb_ci_storage_server	“Owns :: Owned by”	cmdb_ci_ip_address
cmdb_ci_storage_node_element	“Owns :: Owned by”	cmdb_ci_ip_address
cmdb_ci_storage_node_element	“Contains :: Contained by”	cmdb_ci_disk
cmdb_serial_number	“reference”	cmdb_ci_storage_server
cmdb_ci_storage_server	“Allocated to :: Allocated from”	cmdb_ci_storage_node_element

Relationships created by the InfiniBox Storage Volume (LP) pattern

CI	Relations	CI
cmdb_ci_storage_volume	“Provisioned From :: Provisioned”	cmdb_ci_storage_vol_snapshot
cmdb_ci_storage_pool	“reference by column name provided by”	cmdb_ci_storage_volume
cmdb_ci_storage_server	"Contains :: Contained by and reference by column name computer”	cmdb_ci_storage_volume
cmdb_ci_storage_server	“Contains :: Contained by”	cmdb_ci_storage_vol_snapshot
cmdb_ci_infinibox_cg	“Contains :: Contained by”	cmdb_ci_storage_volume
cmdb_ci_infinibox_snapshot_cg	“Contains :: Contained by”	cmdb_ci_storage_vol_snapshot

Relationship created by the InfiniBox Storage Pool (LP) pattern

CI	Relations	CI
cmdb_ci_storage_server	“Contains :: Contained by”	cmdb_ci_storage_pool

Relationship created by the InfiniBox Storage Export (LP) pattern

CI	Relations	CI
cmdb_ci_storage_fileshare	“Provided By :: Provides”	cmdb_ci_storage_server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships created by the InfiniBox Storage File System (LP) pattern

CI	Relations	CI
cmdb_ci_storage_server	"Contains :: Contained by"	cmdb_ci_file_system
cmdb_ci_storage_server	"Contains :: Contained by"	cmdb_ci_file_system_snapshot
cmdb_ci_file_system	"Provided By :: Provides"	cmdb_ci_storage_pool
cmdb_ci_file_system_snapshot	"Provided By :: Provides"	cmdb_ci_storage_pool
cmdb_ci_file_system	"Provisioned From :: Provisioned"	cmdb_ci_file_system_snapshot
cmdb_ci_file_system	"Provided By :: Provides"	cmdb_ci_storage_fileshare
cmdb_ci_file_system_snapshot	"Provided By :: Provides"	cmdb_ci_storage_fileshare

Relationships created by the InfiniBox Storage Consistency Group (LP) pattern

CI	Relations	CI
cmdb_ci_storage_server	"Contains :: Contained by"	cmdb_ci_infinibox_cg
cmdb_ci_infinibox_cg	"Provisioned From :: Provisioned"	cmdb_ci_infinibox_snapshot_cg
cmdb_ci_storage_server	"Contains :: Contained by"	cmdb_ci_infinibox_snapshot_cg
cmdb_ci_infinibox_cg	"Contains :: Contained by"	cmdb_ci_storage_pool
cmdb_ci_infinibox_snapshot_cg	"Contains :: Contained by"	cmdb_ci_storage_pool

Relationships created by the InfiniBox Storage Replica (LP) pattern

CI	Relations	CI
cmdb_ci_storage_server	"Contains :: Contained by"	cmdb_ci_infinibox_replica
cmdb_ci_infinibox_replica	"Replicates to :: Replicated by"	cmdb_ci_infinibox_cg
cmdb_ci_infinibox_replica	"Uses :: Used by"	cmdb_ci_infinibox_link
cmdb_ci_infinibox_replica	"Replicates to :: Replicated by"	cmdb_ci_storage_volume
cmdb_ci_infinibox_replica	"Replicates to :: Replicated by"	cmdb_ci_file_system

Relationship created by the InfiniBox Storage Link (LP) pattern

CI	Relations	CI
cmdb_ci_storage_server	"Contains :: Contained by"	cmdb_ci_infinibox_link

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Infinibox Class Model

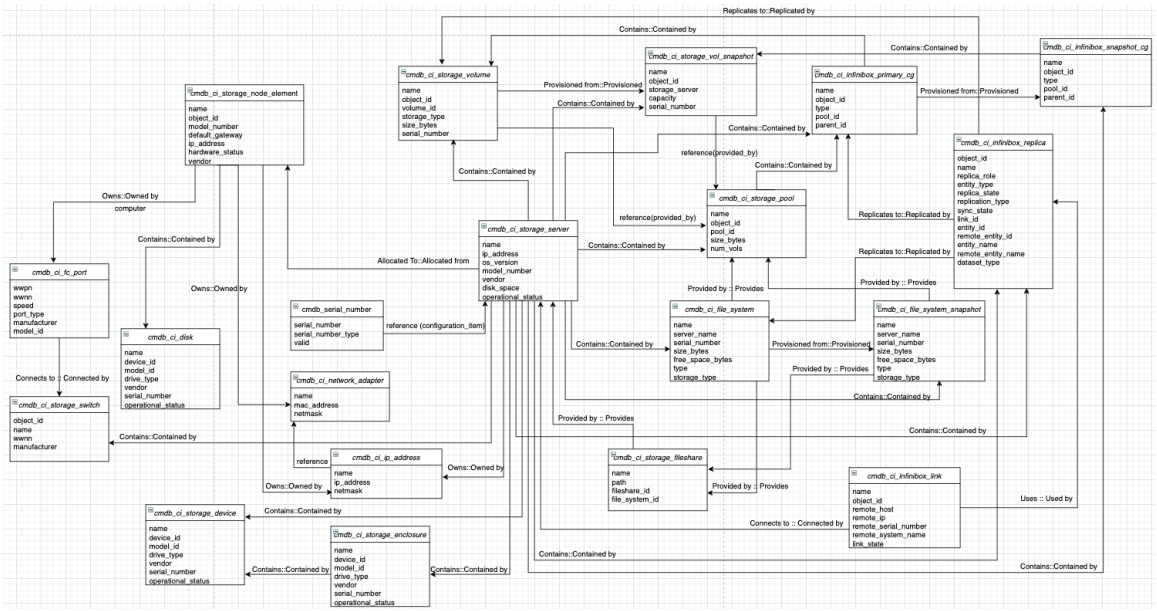

Next-Generation Juniper Network Firewall discovery

The ServiceNow Discovery application uses the Next-Generation Juniper Network Firewall discovery pattern to find Juniper network firewalls. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The discovery pattern uses a set of SNMP calls to find the Juniper network firewalls. Discovery uses the pattern to run horizontal discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Juniper Network Firewall data model

The Next-Generation Juniper Network Firewall pattern introduces the following CI classes that are part of the Juniper firewall data model.

CI class hierarchy

CI class	Extends from
Firewall Device [cmdb_ci_firewall_device]	IP Firewall [cmdb_ci_ip_firewall]
Juniper Firewall Device [cmdb_ci_firewall_device_juniper]	Firewall Device [cmdb_ci_firewall_device]
Firewall Cluster [cmdb_ci_firewall_cluster]	Cluster [cmdb_ci_cluster]
Juniper Firewall Cluster [cmdb_ci_firewall_cluster_juniper]	Firewall Cluster [cmdb_ci_firewall_cluster]
Firewall Device Group [cmdb_ci_firewall_device_group]	Configuration Item [cmdb_ci]
Firewall Manager [cmdb_ci_firewall_manager]	Configuration Item [cmdb_ci]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI class hierarchy (continued)

CI class	Extends from
Firewall Security Policy [cmdb_ci_firewall_security_policy]	Configuration Item [cmdb_ci]

Prerequisites

- Ensure that your network firewall device has SNMP access.
- On the ServiceNow instance, configure SNMP credentials.
- Add the SNMP system OID record for the Juniper device to the ServiceNow instance. Update the following:
  - Classifier: Juniper Firewall
  - Class: Juniper Firewall Device

Deploy the pattern as follows:

1. Download and install the CMDB CI Class Models: Release 1.10.0 from the ServiceNow Store. The app adds the new CMDB classes required for network firewall discovery. For more information, see Firewall extension classes.
2. Download and install the discovery pattern from the ServiceNow Store.
3. Sync the pattern with the appropriate MID Server.

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

IP Firewall [cmdb_ci_ip_firewall]

Field	Description
Name [name]	Hostname.
Serial Number [serial_number]	Device serial number.
Operational Status [operational_status]	Indicates if the device is in active state.
IP address [ip_address]	IP address.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description.
Model Number [model_number]	Device model number.
Firmware [firmware_version]	Firmware version.
Hardware Operating System [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Firewall Device [cmdb_ci_firewall_device]

Field	Description
Name [name]	Hostname.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Firewall Device [cmdb_ci_firewall_device] (continued)

Field	Description
Serial Number [serial_number]	Serial number of the device.
Fully qualified domain name [fqdn]	Fully qualified domain name.
Operational Status [operational_status]	Indicates if the device is in active state.
IP address [ip_address]	IP address.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description.
Model Number [model_number]	Device model number.
Firmware [firmware_version]	Firmware version.
Hardware Operating System [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Juniper Firewall Device [cmdb_ci_firewall_device_juniper]

Field	Description
Name [name]	Hostname.
Serial Number [serial_number]	Serial number of the device.
Fully qualified domain name [fqdn]	Fully qualified domain name.
Operational Status [operational_status]	Indicates if the device is in active state.
IP address [ip_address]	IP address.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description.
Model Number [model_number]	Device model number.
Firmware [firmware_version]	Firmware version.
Hardware Operating System [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Firewall Cluster [cmdb_ci_firewall_cluster]

Field	Description
Name [name]	Hostname.
Fully qualified domain name [fqdn]	Fully qualified domain name.
IP address [ip_address]	IP address.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description.
Model Number [model_number]	Device model number.
Hardware Operating System [hardware_os]	OS running on the hardware.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Firewall Cluster [cmdb_ci_firewall_cluster] (continued)

Field	Description
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Juniper Firewall Cluster [cmdb_ci_firewall_cluster_juniper]

Field	Description
Name [name]	Hostname.
Fully qualified domain name [fqdn]	Fully qualified domain name.
IP address [ip_address]	IP address.
Manufacturer [manufacturer]	Device manufacturer.
Description [short_description]	Short description.
Model Number [model_number]	Device model number.
Hardware Operating System [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	Name of the network adapter.
IP Address [ip_address]	IP address of the network adapter.
MAC Address [mac_address]	MAC address of the network adapter.
Configuration Item [cmdb_ci]	References the Juniper Firewall Device [cmdb_ci_firewall_device_juniper] associated with this network adapter.

The Dependency Views map on the Juniper Firewall Device CI shows the Juniper Firewall Cluster to which it belongs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

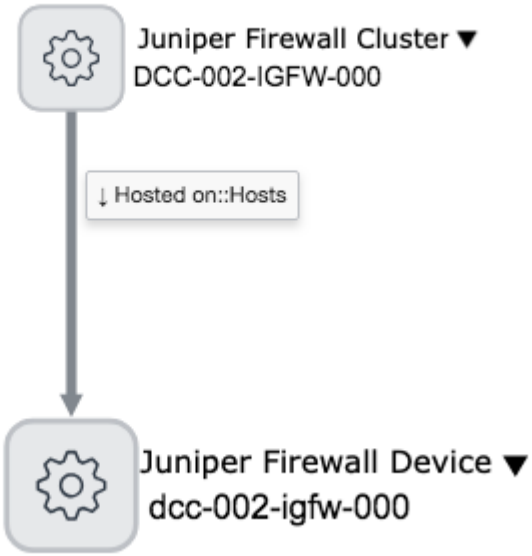

CI relationships

The Next-Generation Juniper Network Firewall pattern creates the following relationships and references to support Juniper network firewall discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Juniper Firewall Cluster [cmdb_ci_firewall_cluster_juniper]	Hosted on::Hosts	Juniper Firewall Device [cmdb_ci_firewall_device_juniper]
Juniper Firewall Device [cmdb_ci_firewall_device_juniper]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
Juniper Firewall Device [cmdb_ci_firewall_device_juniper]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Juniper Firewall Device [cmdb_ci_firewall_device_juniper]	Uses::Used by	Router Interface [dscy_router_interface]
Firewall Device [cmdb_ci_firewall_device]	Contains::Contained by	Firewall Security Policy [cmdb_ci_firewall_security_policy]
Firewall Device Group [cmdb_ci_firewall_device_group]	Members::Member Of	Firewall Device [cmdb_ci_firewall_device]
Firewall Device Group [cmdb_ci_firewall_device_group]	Contains::Contained by	Firewall Security Policy [cmdb_ci_firewall_security_policy]
Firewall Manager [cmdb_ci_firewall_manager]	Manages::Managed by	Firewall Device [cmdb_ci_firewall_device]
Firewall Manager [cmdb_ci_firewall_manager]	Contains::Contained by	Firewall Security Policy [cmdb_ci_firewall_security_policy]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration item [configuration_item]	Juniper Firewall Device [cmdb_ci_firewall_device_juniper]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	Juniper Firewall Device [cmdb_ci_firewall_device_juniper]
Router Interface [dscy_router_interface]	Configuration Item [cmdb_ci]	Juniper Firewall Device [cmdb_ci_firewall_device_juniper]

Linux Pacemaker Cluster discovery

The ServiceNow Discovery application uses the Linux Pacemaker Cluster discovery pattern to find high-availability cluster data and populate the CMDB with the discovered information. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

This pattern is an extension section of the Linux Server pattern. Discovery uses the pattern to run horizontal discovery.

You can use this pattern on the ServiceNow AI Platform® using the Madrid release or later.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Configure SSH credentials for the Linux server.
- Provide a user with permissions to read the following files:
  - `/etc/hosts`
  - `/etc/corosync/corosync.conf`
  - `/var/lib/heartbeat/crm/cib.xml`, or `/var/lib/pacemaker/cib/cib.xml` from Linux SP3
- Provide a user with permissions to run the following commands:

```
ps -ef | grep corosync | grep -v grep
```

```
ps -ef | grep pacemaker | grep -v grep
```

```
ifconfig | grep 'inet addr:' |awk '{print $1, $2 }'
```

```
hostname -s
```

```
hostname -f
```

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Main CI - UNIX Cluster [cmdb_ci_unix_cluster]	
name	Cluster name.
cluster_id	Cluster ID. Contains the nodes ID.
cluster_type	Cluster type.
cluster_status	Cluster status: Online/offline.
cluster_version	Cluster version.
UNIX Cluster Resource [cmdb_ci_unix_cluster_resource]	
name	Resource name.
resource_type	Resource type.
resource_status	Resource status. For example, Started.
ip_address	IP address of the machine where the resource is running.
fqdn	FQDN of the machine where the resource is running.
Cluster Virtual IP [cmdb_ci_cluster_vip]	
name	Cluster VIP name.
ip_address	IP address of the VIP.
Cluster Resource Group [cmdb_ci_cluster_resource_group]	
name	Resource group name.
node	Resource group node.
resource_group_status	Resource group status.
server	Resource group server.
fqdn	FQDN of the machine where the resource group is running.
UNIX Cluster Node [cmdb_ci_unix_cluster_node]	
name	Node name and node ID.
node_status	Node status.
node_state	Node state.
fqdn	FQDN of the machine where the node is running.

The Dependency Views map shows discovered load balancer CIs and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

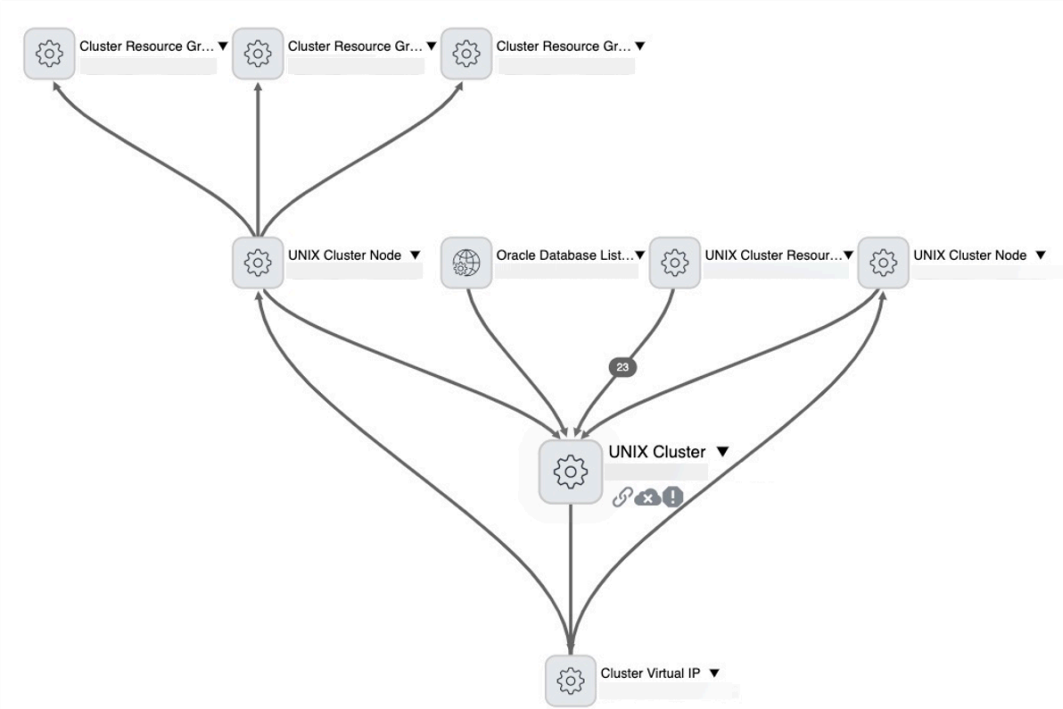

CI relationships

These relationships are created to support Linux Pacemaker Cluster discovery:

CI	Relationship	CI
cmdb_ci_unix_cluster	Hosts::Hosted on	cmdb_ci_linux_server
cmdb_ci_unix_cluster_node	Hosted on::Hosts	cmdb_ci_linux_server
	Cluster of::Cluster	cmdb_ci_unix_cluster
cmdb_ci_cluster_vip	Uses::Used by	cmdb_ci_unix_cluster_node
	Virtualized by::Virtualized	cmdb_ci_unix_cluster
cmdb_ci_unix_cluster_resource	Defines resources for::Gets resources from	cmdb_ci_unix_cluster
	Defines resources for::Gets resources from	cmdb_ci_unix_cluster_node
cmdb_ci_cluster_resource_group	Contains::Contained by	cmdb_ci_unix_cluster
	Contains::Contained by	cmdb_ci_unix_cluster_node
	Contains::Contained by	cmdb_ci_unix_cluster_resource

Microsoft Certificate Authority (CA) certificates discovery

The ServiceNow Discovery application can discover Microsoft Certificate Authority (CA) certificates using the Microsoft CA - Certificate Management pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Starting with version 1.27.0, Discovery and Service Mapping Patterns supports discovering the subject alternative name (SAN) for both the all request IDs and the template ID discovery methods.

Prerequisites

Verify the following applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models

Verify the certificate Authority process

Ensure that the CA process is up and running on the host machine.

Verify previous Windows discovery

Verify a Windows host was discovered during a previous horizontal discovery. For more information, see Windows discovery.

Create a serverless discovery schedule

Create a discovery schedule to perform targeted discovery of Microsoft CA certificates using the host on which the CA process is running. For more information, see Create a serverless schedule for Microsoft CA discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Microsoft CA - Certificate Management pattern.

Unique Certificate [cmdb_ci_certificate]

Field	Description
Fingerprint [fingerprint]	Hash value of the certificate.
Fingerprint algorithm [fingerprint_algorithm]	Algorithm used to hash the certificate.
Subject common name [subject_common_name]	Identifies the host name/domain associated with the certificate, listed in the certificate's **Subject** field.
Subject distinguished name [subject_distinguished_name]	Identifying information of the subject listed in the certificate's **Subject** field.
Issuer distinguished name [issuer_distinguished_name]	Distinguished name of the issuer.
Comments [comments]	Includes a readable description for users.
Renewal tracking [renewal_tracking]	Indicates whether to create any priority 1 or priority 3 tasks for the expiring certificates.
Certificate Template [cert_template]	Certificate template that defines the policies and rules that a CA uses when a request for a certificate is received.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Unique Certificate [cmdb_ci_certificate] (continued)

Field	Description
Certificate ThumbPrint [thumb_print]	Unique identifier for certificates in applications when making trust decisions, in configuration files, and when displayed in interfaces.
Issuer common name [issuer_common_name]	Common name of the issuer.
Valid from [valid_from]	Validity start period of the certificate.
Serial Number [serial_number]	Serial number of the certificate.
Subject country [subject_country]	Two-letter country code in the certificate's **Subject** field.
Subject organization [subject_organization]	Organization listed in the certificate's **Subject** field.
Issuer [issuer]	Entity that signed and issued the certificate.
Subject organizational unit [subject_organizational_unit]	Organizational unit listed in the certificate's **Subject** field..
Subject alternative name [subject_alternative_name]	List of fully qualified domain names secured by the certificate, listed in the certificate's **Subject** field.
Valid to [valid_to]	Validity end period of the certificate.
Name [name]	Name of the CI.
State [state]	Life-cycle state of the certificate. For example: Installed, Issued, or Revoked.
Root issuer [root_issuer]	Root entity that signed and issued the immediate certificate.
Subject locality [subject_locality]	Locality listed in the certificate's **Subject** field.
Subject state [subject_state]	State listed in the certificate's **Subject** field.
Operational status [operational_status]	Indicates whether the certificate is valid or was revoked. Possible values are: • **Operational**: Valid certificate • **Non-Operational**: Revoked certificate
Signature algorithm [signature_algorithm]	Signature algorithm of the certificate. For example: SHA-256, sha256RSA, or SHA1withRSA.
Request Type [request_type]	Format used for requesting the certificate.
Request Submission Date [request_submission_date]	Date that the certificate request was submitted.
Request Resolution Date [request_resolution_date]	Date that the certificate request was resolved by the CA.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Unique Certificate [cmdb_ci_certificate] (continued)

Field	Description
Request Revocation Date [request_revocation_date]	Date that the certificate was revoked. This field is populated only for revoked certificates.
Effective Revocation Date [effective_revocation_date]	Date that the certificate was effectively revoked by being added to the Certificate Revocation List (CRL). This field is populated only for revoked certificates.
Revocation Reason [revocation_reason]	Reason for the certificate's revocation. This field is populated only for revoked certificates.
Requester Name [requester_name]	Name of the person who requested the certificate.
Template Enrollment Flags [template_enrollment_flags]	Information about the certificate that needs to be acted on by the CA or the certificate's owner. For more information, search for the error code in the official Microsoft documentation site.

CI relationships

The Microsoft CA - Certificate Management pattern doesn't create any CI relationships.

Create a serverless schedule for Microsoft CA discovery

Create a serverless discovery schedule to discover Microsoft Certificate Authority (CA) certificates.

Before you begin

- Verify that the CA process is up and running on the host machine.
- Verify that a Windows host was discovered during a previous horizontal discovery. For more information, see Windows discovery.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule New record form

Field	Description
Name	Unique name for this discovery schedule. For example: Discover MS CA.
Discover	Scan type, which should be Certificates.
Certificate Discovery Type	Certificate type, which should be CA Trust Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
MID server	Name of the MID Server to use for this schedule.
Active	Enables this schedule for discovery.

c. Select **Submit**.

3. Create the execution pattern.

a. In the Discovery Schedules page, select the record you created.

b. In the **Serverless Execution Patterns** tab, select **New**.

Discovery uses each execution pattern to discover up to 20,000 certificates. For deployments using more than 20,000 certificates, create several execution patterns.

c. On the form, fill in the fields.

Serverless Execution Pattern New record form

Field	Description
Name	Descriptive name for this record. For example: Discover MS CA.
Pattern	Pattern to be used for this schedule, which should be the Microsoft CA - Certificate Management pattern.
Proxy Host	CI name of the Windows host running the Microsoft CA Service, discovered in a previous discovery and populated in the CMDB.
Active	Option to enable this schedule for discovery.

d. Select **Submit**.

4. Set the pattern launcher parameters.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Parameter	Description
template_list	Determines whether to search on all the request IDs or a template ID to discover certificates. Possible values are **all** or the template ID number. ▪ all: Restricts discovery to all certificate request IDs in ascending order. The request IDs are limited either by the default execution pattern limitation of 20,000 or values set in *`start_offset`* or *`limit`*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
	▪ Template ID: Searches for certificates only in the specific template ID. Can be a single template ID or multiple IDs, separated by commas. For example: 1.2.3.4,5.6.7.8,9.10.11.12
start_offset	This optional parameter specifies the number of the certificate from which to start discovery and is relevant only when *template_list* is set to all and for deployments with over 20,000 certificates. If you create multiple serverless execution patterns for deployments with over 20,000 certificates, use 1 for the first execution, 20001 for the second, and so on. If this parameter remains empty, *start_offset* defaults to 1.
limit	This optional parameter limits the certificate discovery to the specified value and is relevant only when *template_list* is set to all. If the parameter remains empty, the limit defaults to 20,000 certificates.
ip	IP of the server on which the CA process (certsrv) is running.
discover_SAN_for_template	When set to true, this optional parameter enables the discovery of the subject alternative name (SAN) of the certificates when a template ID is provided in the *template_list* parameter. This parameter executes separate commands for each certificate within the template ID, which could impact performance. When searching on all request IDs, SAN is automatically discovered, so this parameter should remain empty.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

Microsoft Certificate Authority (CA) certificates discovery

NetApp SolidFire storage system discovery

ServiceNow Discovery uses the NetApp SolidFire storage system discovery pattern to find clusters and nodes on the SolidFire storage system. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the following apps are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models

Create credential alias for basic authentication credential

Configure a credential alias for basic authentication credentials for a SolidFire Cluster Admin user. For more information, see Create a basic authentication credential alias for NetApp SolidFire discovery.

Verify that the MID Server has access

Verify the MID Server can access the target device.

Create a serverless discovery schedule

Create a serverless discovery schedule to perform targeted discovery of SolidFire clusters and nodes. For more information, see Create a serverless discovery schedule for NetApp SolidFire discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the NetApp SolidFire Storage System pattern.

Storage Server [cmdb_ci_storage_server]

Field	Description
Name [name]	The name of the CI.
Serial number [serial_number]	The serial number associated with the CI.
Firmware version [firmware_version]	The firmware version of the CI.
IP Address [ip_address]	The IP address of the CI.
Description [short_description]	A short textual description of the CI.
Manufacturer [manufacturer]	The device manufacturer.
Model ID [model_id]	The device model ID.
OS Version [os_version]	The version of the OS running on the CI.
Operating System [os]	The OS running on the CI.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Cluster [cmdb_ci_storage_cluster]

Field	Description
Name [name]	The name of the CI.
IP Address [ip_address]	The IP address of the storage cluster.
Description [short_description]	A short textual description of the CI.
Manufacturer [manufacturer]	The device manufacturer.
Cluster Version [cluster_version]	Element OS API version running on the cluster.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Storage Cluster Node [cmdb_ci_storage_cluster_node]

Field	Description
Name [name]	The name of the CI.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Cluster [cluster]	References the Storage Cluster [cmdb_ci_storage_cluster] table.

Storage Node Element [cmdb_ci_storage_node_element]

Field	Description
Name [name]	The name of the CI.
Description [short_description]	A short textual description of the CI.
Manufacturer [manufacturer]	The device manufacturer.
Serial number [serial_number]	The serial number associated with the CI.
Model ID [model_id]	The device model ID.
Chassis type [chassis_type]	The type of chassis that contains the server.
RAM (MB) [ram]	The memory size of the CI.
CPU type [cpu_type]	The CPU type.
CPU manufacturer [cpu_manufacturer]	The CPU manufacturer.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Disk [cmdb_ci_disk]

Field	Description
Name [name]	The name of the CI.
Device ID [device_id]	The ID of the device.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk [cmdb_ci_disk] (continued)

Field	Description
Serial number [serial_number]	The serial number associated with the CI.
Storage type [storage_type]	The storage type.
Size bytes [size_bytes]	The available storage size in bytes.
Manufacturer [manufacturer]	The device manufacturer.
Model ID [model_id]	The device model ID.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	The name of the CI.
IP Address [ip_address]	The IP address of the CI.
MAC Address [mac_address]	The MAC address of the network adapter.
Netmask [netmask]	The netmask of the server hosting the network adapter.
Configuration Item [cmdb_ci]	References the Storage Server [cmdb_ci_storage_server] table.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	The IP address of the CI.
Netmask [netmask]	The netmask of the server hosting the CI.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.

The Dependency Views map shows all discovered SolidFire storage system clusters and nodes in your organization and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

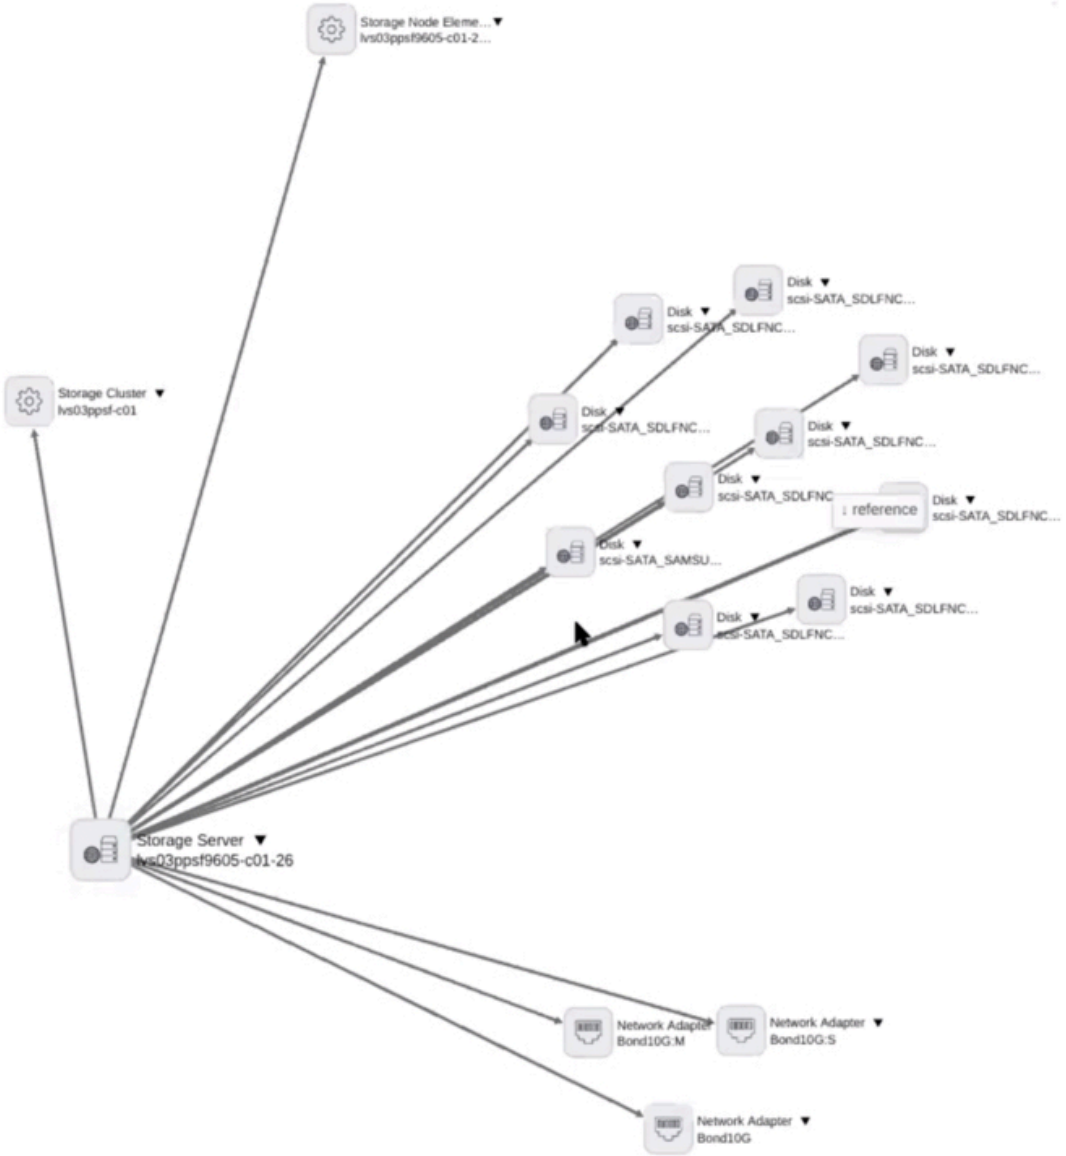

CI relationships

The NetApp SolidFire Storage System pattern creates the following relationships and references to support storage system discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Storage Cluster [cmdb_ci_storage_cluster]	Cluster of::Cluster	Storage Cluster Node [cmdb_ci_storage_cluster_node]
Storage Cluster Node [cmdb_ci_storage_cluster_node]	Hosted on::Hosts	Storage Node Element [cmdb_ci_storage_node_element]
Storage Node Element [cmdb_ci_storage_node_element]	Managed by::Manages	Storage Cluster [cmdb_ci_storage_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Storage Server [cmdb_ci_storage_server]	Allocated to::Allocated from	Storage Node Element [cmdb_ci_storage_node_element]
Storage Server [cmdb_ci_storage_server]	Runs on::Runs	Storage Cluster [cmdb_ci_storage_cluster]
Storage Server [cmdb_ci_storage_server]	Contains::Contained by	Disk [cmdb_ci_disk]
Storage Server [cmdb_ci_storage_server]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
Storage Server [cmdb_ci_storage_server]	Owns::Owned by	IP Address [cmdb_ci_ip_address]

CI references

CI	Field	Referenced CI
Storage Cluster Node [cmdb_ci_storage_cluster_node]	Cluster [cluster]	Storage Cluster [cmdb_ci_storage_cluster]
Disk [cmdb_ci_disk]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	Storage Server [cmdb_ci_storage_server]
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	Storage Server [cmdb_ci_storage_server]
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	Storage Node Element [cmdb_ci_storage_node_element]

Create a basic authentication credential alias for NetApp SolidFire discovery

Create an alias and add it to a basic authentication credential to discover NetApp SolidFire clusters.

Before you begin

Verify that a basic authentication account is configured with SolidFire Cluster Admin credentials and includes the following permissions:

- read
- clusterAdmins
- reporting
- drives

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create an alias.
   a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   b. Select **New**.
   c. Enter a unique name for the alias and select **Credential** for the alias type in the **Type** field.
   d. Select **Submit**.
2. Configure a basic authentication credential for the new alias.
   a. Navigate to **Discovery** > **Credentials**.
   b. Select **New**.
   c. Select **Basic Auth Credentials**.
   d. Unlock the **Credential alias** field and select the alias you created.
   e. Configure the rest of the **Basic Auth Credentials** form fields.
   For more information, see Basic authentication credentials.
   f. Select **Submit**.

What to do next

Create a serverless discovery schedule. For more information, see Create a serverless discovery schedule for NetApp SolidFire discovery.

Related topics

NetApp SolidFire storage system discovery

Create a serverless discovery schedule for NetApp SolidFire discovery

Set up a dedicated serverless discovery schedule for NetApp SolidFire cluster and node discovery.

Before you begin

- Verify the MID Server is set to Active and can reach the target NetApp SolidFire cluster.
- Create an alias for the basic authentication credential. For more information, see Create a basic authentication credential alias for NetApp SolidFire discovery.
- Obtain the management IP address of the target SolidFire cluster.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule form

Field	Description
Name	Unique name for this discovery schedule. For example, `Discover SolidFire`.
Discover	Scan type. Select **Serverless**.
MID server	Name of the MID Server to use for this schedule.
Active	When selected, activates this schedule for discovery runs.

c. Select **Submit**.

3. Create the execution pattern.

a. In the Discovery Schedules page, select the record you created.

b. In the **Serverless Execution Patterns** tab, select **New**.

c. On the form, fill in the fields.

Serverless Execution Pattern form

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern to use for this schedule. Select the NetApp SolidFire Storage System pattern.
Active	When selected, activates this schedule for discovery runs.

d. Select **Submit**.

4. Set the pattern launcher parameters.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Pattern launcher parameters

Parameter	Value
`cluster_ip`	Management IP address of the SolidFire cluster.
`credentialsAlias`	ID of the credential alias you created for the SolidFire basic authentication credential record.

c. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

NetApp SolidFire storage system discovery

Network router discovery using patterns

The Discovery and Service Mapping Patterns application uses the Network Router pattern to find network routers in your environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about network routers and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

For information on probe-based network switch and router discovery, see Network switch and router discovery.

Cisco Nexus Virtual Routing and Forwarding (VRF) data model

The pattern introduces the following CI class starting with Discovery and Service Mapping Patterns version 1.35.0.

CI class hierarchy

CI class	Extends from
Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]	Configuration Item [cmdb_ci]

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

Verify IP address information

The router must contain IP address information for successful discovery.

Verify that exit interface routing rules are configured and reachable

The router that provides the IP address must have exit interface routing rules configured and be reachable for successful discovery.

Create SNMP credentials

Create SNMP Community Credentials or SNMPv3 credentials. For more information, see SNMP credentials.

Configure properties for router discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Navigate to `sys_properties.list` and verify that the *`glide.discovery.L3_mapping`* system property is set to **true**.

You can configure these additional properties for network router discovery.

Optional network router discovery properties

System property	Description
glide.discovery.max_concurrent_invocations_per_schedule	Sets a maximum number of scheduled invocations for the same Discovery schedule prevents a backlog of scheduled runs if Disc does not finish before the next one starts. The value is an integer that represents the m number of automated invocations that can ru concurrently for the same schedule. If the lim reached, subsequent scheduled invocations canceled. A value of 0 or any negative numb disables this property. • **Type:** integer • **Default value:** 3 **Note:** This property does not apply to schedules that have a 'Run after' configuration set to 'Even if canceled'.
glide.discovery.bgp_router_disable	Disables running the SNMP – Routing patter when discovering a router running the BGP protocol. If you must populate the Device Neighbors [discovery_device_neighbors] table during horizontal layer 2 discovery of the BGP-enabled devices, set the *`glide.discovery.bgp_router_dis`* property to **false**. Notice that enabling this property can cause performance issues inclu out-of-memory issues on the MID Server. • **Type:** Boolean • **Default value:** true
glide.discovery.disable_next_hop_data	Skips populating the Next Hop Routing Rule table for BGP-enabled routers during horizon discovery. The device, its physical interfaces its neighbors are still discovered. This property takes effect only when the*`glide.discovery.bgp_router_d`* property is set to false. • **Type:** Boolean • **Default value:** false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Optional network router discovery properties (continued)

System property	Description
	For more information, see Turn off next-hop r data collection for BGP routers.

For Cisco Nexus VRF discovery

Verify the following requirements:

- Discovery and Service Mapping Patterns starting from version 1.35.0.
- Read permissions for the following commands on the Cisco Nexus device:
  - `sh vrf`
  - `sh hsrp br`
  - `sh ip interface vrf all`
  - `sh ipv6 interface vrf all`
- Create SSH credentials for the Cisco Nexus user. For more information, see SSH credentials.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Network Router pattern.

IP Router [cmdb_ci_ip_router]

Field	Description
Name [name]	Name of the router.
Model ID [model_id]	Hardware model ID of the router.
Manufacturer [manufacturer]	Router manufacturer.
Serial number [serial_number]	Unique hardware serial number of the router.
IP Address [ip_address]	IP addresses assigned to interfaces.
Description [short_description]	Device description.
Can route IP [can_route]	Indicates whether the IP pockets can be routed between networks. Possible values are true or false.
Can partition VLANs [can_partitionvlans]	Indicates whether the local area network (LAN) can be partitioned. Possible values are true or false.
Can switch IP [can_switch]	Indicates whether a network device has Layer 2 switching capabilities. Possible values are true or false.
Physical interface count [physical_interface_count]	Number of physical interfaces.
Ports [ports]	Individual network interface (physical or logical) used to connect devices such as computers, servers, access points, or uplinks to switches or routers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Router [cmdb_ci_ip_router] (continued)

Field	Description
Discovery Protocol ID [discovery_proto_id]	Unique identifier of the device participating in the Cisco Discovery Protocol (CDP) or Link Layer Discovery Protocol (LLDP) protocol.
Discovery Protocol Type [discovery_proto_type]	Protocol type. Possible values are CDP or LLDP.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address of the router.
Netmask [netmask]	Netmask of the router.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.
Owned By Configuration Item [owned_by_cmdb_ci]*	References the Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding] table

* Populated for Cisco Nexus VRF discovery only.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	Name of the network adapter.
IP Address [ip_address]	IP address of the network adapter.
Netmask [netmask]	Netmask of the network adapter.
MAC Address [mac_address]	MAC address of the network adapter.
Alias [alias]	User-assigned name for the network adapter.
Configuration Item [cmdb_ci]	References the IP Router [cmdb_ci_ip_router] table.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Name of the Domain Name System (DNS).
IP Address [ip_address]	IP address of the DNS.

Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]*

Field	Description
Name [name]	VRF name as configured on the network device.
VRF ID [vrf_id]	VRF identifier assigned by the network device operating system (OS).
Display Name [display_name]	Display name composed of the VRF name and the associated network device name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]* (continued)

Field	Description
Configuration Item [configuration_item]	References the IP Router [cmdb_ci_ip_router] table.

* Populated for Cisco Nexus VRF discovery only.

CI relationships

Discovery creates these relationships and references to support the network router discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
IP Router [cmdb_ci_ip_router]	Contains::Contained by	Switch Forwarding Rule [dscy_swtch_fwd_rule]
IP Router [cmdb_ci_ip_router]	Contains::Contained by	Switch Partition [dscy_swtch_partition]
IP Router [cmdb_ci_ip_router]	Contains::Contained by	Switchport [dscy_switchport]
IP Router [cmdb_ci_ip_router]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
IP Router [cmdb_ci_ip_router]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
IP Router [cmdb_ci_ip_router]	Uses::Used by	Exit Interface Routing Rule [dscy_route_interface]
IP Router [cmdb_ci_ip_router]	Uses::Used by	Next Hop Routing Rule [dscy_route_next_hop]
IP Router [cmdb_ci_ip_router]	Uses::Used by	Router Interface [dscy_router_interface]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
IP Router [cmdb_ci_ip_router]	Owns::Owned by	Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]*
Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]	Contains::Contained by	IP Address [cmdb_ci_ip_address]*

* Created for Cisco Nexus VRF discovery only.

CI references

CI/Table	Field	Referenced CI
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI/Table	Field	Referenced CI
Network ARP Table [discovery_net_arp_table]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Next Hop Routing Rule [dscy_route_next_hop]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Router Interface [dscy_router_interface]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Switch Bridge Port Table [discovery_switch_bridge_port_table]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Switch Forwarding Rule [dscy_swtch_fwd_rule]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Switch Forwarding Table [discovery_switch_fwd_table]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Switch Partition [sdcy_swtch_partition]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Switch Spanning Tree Table [discovery_switch_spanning_tree_table]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Switchport [dscy_switchport]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
Exit Interface Routing Rule [dscy_route_interface]	Configuration Item [cmdb_ci]	IP Router [cmdb_ci_ip_router]
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]
IP Address [cmdb_ci_ip_address]	Owned By Configuration Item [owned_by_cmdb_ci]	Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]*
Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]	Configuration Item [configuration_item]	IP Router [cmdb_ci_ip_router]*

* Populated for Cisco Nexus VRF discovery only.

Turn off next-hop route data collection for BGP routers

You can turn off Next Hop Routing Rule [dscy_route_next_hop] table collection for BGP-enabled routers to reduce MID Server memory load during router discovery.

Before you begin

- Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.
- Navigate to `sys_properties.list` and verify that the *`glide.discovery.bgp_router_disable`* property is set to false.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

By default, the *glide.discovery.bgp_router_disable* property is set to true, blocking the entire BGP routing branch including interfaces and neighbors. Setting the property to false enables interface and neighbor collection but also triggers a full route-table walk. On routers with large route tables, this can cause MID Server memory and performance issues.

The *glide.discovery.disable_next_hop_data* property controls route-table collection. When this property is set to true, Discovery collects the device, physical interfaces, and neighbors but skips populating the Next Hop Routing Rule [dscy_route_next_hop] table. The default value is false.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `glide.discovery.disable_next_hop_data`.
3. Select the *glide.discovery.disable_next_hop_data* system property.
4. In the **Value** field, enter `true`.
   To also collect next-hop route table data, set the value to `false`.
5. Select **Update**.

What to do next

Run Discovery again against a BGP-enabled router to apply the changes.

Related topics

Network router discovery using patterns

Network switch discovery using patterns

The Discovery and Service Mapping Patterns application uses the Network Switch pattern to find network switches in your environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about network switches and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

For information on probe-based network switch and router discovery, see Network switch and router discovery.

Cisco Nexus Virtual Routing and Forwarding (VRF) data model

The pattern introduces the following CI class starting with Discovery and Service Mapping Patterns version 1.35.0.

CI class hierarchy

CI class	Extends from
Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]	Configuration Item [cmdb_ci]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

Verify IP address information

The switch must contain IP address information for successful discovery.

Verify that exit interface routing rules are configured and reachable

The switch that provides the IP address must have exit interface routing rules configured and be reachable for successful discovery.

Create SNMP credentials

Create SNMP Community Credentials or SNMPv3 credentials. For more information, see SNMP credentials.

Configure properties for switch discovery

Navigate to `sys_properties.list` and verify that the *`glide.discovery.L3_mapping`* system property is set to **true**.

You can configure these additional properties for network switch discovery.

Optional network switch discovery properties

System property	Description
glide.discovery.max_concurrent_invocations_per_schedule	Sets a maximum number of scheduled invocations for the same Discovery schedule prevents a backlog of scheduled runs if Disc does not finish before the next one starts. The value is an integer that represents the m number of automated invocations that can ru concurrently for the same schedule. If the lim reached, subsequent scheduled invocations canceled. A value of 0 or any negative numb disables this property. • **Type:** integer • **Default value:** 3 **Note:** This property does not apply to schedules that have a 'Run after' configuration set to 'Even if canceled'.
glide.discovery.bgp_router_disable	Disables running the SNMP – Routing patter when discovering a router running the BGP protocol. If you must populate the Device Neighbors [discovery_device_neighbors] table during horizontal layer 2 discovery of the BGP-enabled devices, set the *`glide.discovery.bgp_router_dis`* property to **false**. Notice that enabling this

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Optional network switch discovery properties (continued)

System property	Description
	property can cause performance issues incl out-of-memory issues on the MID Server. • **Type:** Boolean • **Default value:** true

For Cisco Nexus VRF discovery

Verify the following requirements:

- Discovery and Service Mapping Patterns starting from version 1.35.0.
- Read permissions for the following commands on the Cisco Nexus device:
  - `sh vrf`
  - `sh hsrp br`
  - `sh ip interface vrf all`
  - `sh ipv6 interface vrf all`
- Create SSH credentials for the Cisco Nexus user. For more information, see SSH credentials.

Discovering stacked switches

Network switch discovery supports Cisco stacked switches. Multiple switches in the stack are identified by the primary IP address. Discovery identifies each member of the stack by elements such as the primary IP address, a name associated with the switch, and the stack mode. The relationship is created between a primary switch and the secondary switches.

The primary switch has all the fields listed in the IP Switch [cmdb_ci_ip_switch] table. The secondary switch has name, serial_number, model_id, manufacturer, stack, and stack_mode.

Note: If all ports on a switch are open for port 161, the switch has an IP per port, but the same serial number is used across all ports. The configuration item (CI) identification is based on the serial number, therefore, discovery updates the existing record every time instead of creating a new one.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Network Switch pattern.

IP Switch [cmdb_ci_ip_switch]

Field	Description
Name [name]	Name of the switch.
Model ID [model_id]	Hardware model ID of the switch.
Manufacturer [manufacturer]	Switch manufacturer.
Install Status [install_status]	Install status of the device. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Switch [cmdb_ci_ip_switch] (continued)

Field	Description
Serial number [serial_number]	Unique hardware serial number of the router.
IP Address [ip_address]	IP addresses assigned to interfaces.
Description [short_description]	Device description.
Can route IP [can_route]	Indicates whether the IP pockets can be routed between networks. Possible values are true or false.
Can partition VLANs [can_partitionvlans]	Indicates whether the local area network (LAN) can be partitioned. Possible values are true or false.
Can switch IP [can_switch]	Indicates whether a network device has Layer 2 switching capabilities. Possible values are true or false.
Physical interface count [physical_interface_count]	Number of physical interfaces.
Ports [ports]	Individual network interface (physical or logical) used to connect devices such as computers, servers, access points, or uplinks to switches or routers.
Discovery Protocol ID [discovery_proto_id]	Unique identifier of the device participating in the Cisco Discovery Protocol (CDP) or Link Layer Discovery Protocol (LLDP) protocol.
Discovery Protocol Type [discovery_proto_type]	Protocol type. Possible values are CDP or LLDP.
Stack [stack]	Indicates whether a network device is part of a stacked group of switches operating as one unit. Possible values are true or false.
Stack Mode [stack_mode]	Stacking architecture or role of a network device within a stack. For example: standalone, stack_master, or stack_member.
Location [location]	Location where the switch is installed.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address of the network switch.
Netmask [netmask]	Netmask of the network switch.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.
Owned By Configuration Item [owned_by_cmdb_ci]*	References the Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding] table

* Populated for Cisco Nexus VRF discovery only.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	Name of the network adapter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network Adapter [cmdb_ci_network_adapter] (continued)

Field	Description
IP Address [ip_address]	IP address of the network adapter.
Netmask [netmask]	Netmask of the network adapter.
MAC Address [mac_address]	MAC address of the network adapter.
Alias [alias]	User-assigned name for the network adapter.
Configuration Item [cmdb_ci]	References the IP Switch [cmdb_ci_ip_switch] table.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Name of the Domain Name System (DNS).
IP Address [ip_address]	IP address of the DNS.

Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]*

Field	Description
Name [name]	VRF name as configured on the network device.
VRF ID [vrf_id]	VRF identifier assigned by the network device operating system (OS).
Display Name [display_name]	Display name composed of the VRF name and the associated network device name.
Configuration Item [configuration_item]	References the IP Switch [cmdb_ci_ip_switch] table.

* Populated for Cisco Nexus VRF discovery only.

CI relationships

Discovery creates these relationships and references to support the network switch discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
IP Switch [cmdb_ci_ip_switch]	Contains::Contained by	Switch Forwarding Rule [dscy_swtch_fwd_rule]
IP Switch [cmdb_ci_ip_switch]	Contains::Contained by	Switch Partition [dscy_swtch_partition]
IP Switch [cmdb_ci_ip_switch]	Contains::Contained by	Switchport [dscy_switchport]
IP Switch [cmdb_ci_ip_switch]	Owns::Owned by	IP Address [cmdb_ci_ip_address]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
IP Switch [cmdb_ci_ip_switch]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
IP Switch [cmdb_ci_ip_switch]	Uses::Used by	Exit Interface Routing Rule [dscy_route_interface]
IP Switch [cmdb_ci_ip_switch]	Uses::Used by	Next Hop Routing Rule [dscy_route_next_hop]
IP Switch [cmdb_ci_ip_switch]	Uses::Used by	Router Interface [dscy_router_interface]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
IP Switch [cmdb_ci_ip_switch]	Owns::Owned by	Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]*
Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]	Contains::Contained by	IP Address [cmdb_ci_ip_address]*

* Created for Cisco Nexus VRF discovery only.

CI references

CI/Table	Field	Referenced CI
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Network ARP Table [discovery_net_arp_table]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Next Hop Routing Rule [dscy_route_next_hop]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Router Interface [dscy_router_interface]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Switch Bridge Port Table [discovery_switch_bridge_port_table]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Switch Forwarding Rule [dscy_swtch_fwd_rule]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Switch Forwarding Table [discovery_switch_fwd_table]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Switch Partition [sdcy_swtch_partition]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Switch Spanning Tree Table [discovery_switch_spanning_tree_table]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
Switchport [dscy_switchport]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI/Table	Field	Referenced CI
Exit Interface Routing Rule [dscy_route_interface]	Configuration Item [cmdb_ci]	IP Switch [cmdb_ci_ip_switch]
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]
IP Address [cmdb_ci_ip_address]	Owned By Configuration Item [owned_by_cmdb_ci]	Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]*
Virtual Routing and Forwarding (VRF) [cmdb_ci_virtual_routing_forwarding]	Configuration Item [configuration_item]	IP Switch [cmdb_ci_ip_switch]*

* Populated for Cisco Nexus VRF discovery only.

Nutanix Acropolis discovery

The ServiceNow Discovery application uses the Nutanix Components pattern to find components of the Nutanix Acropolis solution containing Nutanix Prism Central version 2024.3.1.8 or Nutanix Prism Element 7.0.1.6. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The Nutanix Components pattern uses REST APIs to collect information about the Acropolis deployment components.

Starting with version 1.29.0, Discovery and Service Mapping Patterns supports Nutanix v4 discovery. When you switch to v4, Discovery creates new pool and container records, while existing records stop being updated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Nutanix events

The Nutanix VM Event pattern finds events for virtual machines (VMs) in the Nutanix deployment and triggers rediscovery when a state change is detected. Starting with Discovery and Service Mapping Patterns version 1.31.0, the VM Event pattern also supports API v4, and the Nutanix Host Event pattern detects state changes for Nutanix hosts using API v4. Both patterns run every five minutes by default, but are inactive until you enable the relevant scheduled job.

Patterns used for event discovery

Pattern	CI	API version
Nutanix VM Event	Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]	v1–3 and v4
Nutanix Host Event	Nutanix Host [cmdb_ci_nutanix_host]	v4 only

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models

Verify MID Server requirements

Deploy and connect a dedicated MID Server with network access to Prism Element and Prism Central.

Install Nutanix end-point SSL certificates for the MID Server JRE

1. Collect the certificates by browsing the Nutanix Prism console on Internet Explorer.
2. Add SSL certificates for the MID Server, as described in Add SSL certificates for the MID Server.
3. Install Nutanix end-point SSL certificates for the MID Server JRE.

Verify API permissions for v1-3

Verify that the user has read-only access to the following API requests:

- `/PrismGateway/services/rest/v1/vms`
- `/PrismGateway/services/rest/v1/storage_pools`
- `/PrismGateway/services/rest/v2.0/tasks/list`
- `/PrismGateway/services/rest/v2.0/hosts`
- `/PrismGateway/services/rest/v2.0/clusters`
- `/PrismGateway/services/rest/v2.0/vms`
- `/PrismGateway/services/rest/v2.0/storage_containers`
- `/api/nutanix/v3/clusters/list`
- `/api/nutanix/v3/categories/list`
- `/api/nutanix/v3/categories/{category_keys}/list`
- `/api/nutanix/v3/category/query`

Verify API permissions for v4

Verify that the user has read-only access to the following GET APIs. Note that the API call referencing `/PrismGateway/services/rest/v1/vms` uses a legacy endpoint because it's not available yet in the newer API.

- `/api/clustermgmt/v4/config/clusters`
- `/api/clustermgmt/v4/config/clusters/{clusters_ID}/hosts`
- `/api/clustermgmt/v4/config/storage-containers`
- `/api/vmm/v4/ahv/config/vms`
- `/PrismGateway/services/rest/v1/vms`
- `/api/prism/v4/config/categories?$expand=associations`
- `/api/prism/v4/config/categories/{categories_ID}?$limit=extId&$expand=detailedAssociations`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify Prism Central version for v4 API

Ensure that Prism Central is at least version 7.3 to use the v4 API.

Set MID Server property to discover v4

Enable Nutanix v4 discovery by setting the MID Server property value to the exact v4 version. For more information, see Enable Nutanix v4 discovery.

Configure basic authentication credentials

For more information, see Basic authentication credentials.

Create a serverless discovery schedule

Create a discovery schedule to perform targeted discovery of Nutanix components. For more information, see Create a serverless schedule for Nutanix Acropolis discovery.

Link Nutanix components to ESX clusters

The Nutanix Components pattern can discover Nutanix components hosted on ESX clusters, like storage pools, storage containers, and Controller Virtual Machine (CVM), but not ESX clusters themselves. Run VMware probe discovery before the Nutanix pattern discovery to link the Nutanix components to their ESX clusters. If you don't, the Nutanix components are linked to Nutanix Prism instead. For more information on VMware probe discovery, see Discovery for VMware virtualization.

(Optional) Disable CI tag discovery

The Nutanix Components pattern discovers tags for Nutanix hosts and virtual machines. If you have dozens of tags in your deployment, it might slow down the discovery causing performance issues. You can disable the CI tag discovery by performing the following steps:

1. Navigate to **Pattern Designer** > **Discovery Patterns**.
2. Select the **Nutanix Components** pattern.
3. Under **Extension Section**, select the check box for **Categories (Tags)** and select **Remove**.
4. Select **Save**.

Adjust the batch size for removing redundant relationship data

The *sn_itom_pattern.nutanix_relations_deletion_batch_size* property sets the batch size for deleting relations between Nutanix Virtual Machine Instance and Nutanix Host. The default value is 1000; you can adjust this batch size as needed.

Configure Nutanix events

- Verify that Nutanix Components discovery has run and that the Nutanix Prism Central [cmdb_ci_nutanix_prism_central] table is populated.
- Verify that the MID Servers for Nutanix event discovery have **All** or **Nutanix** capability.
- For v4 event discovery, verify that at least version 1.31.0 of Discovery and Service Mapping Patterns is installed.
- To restrict v4 event discovery to specific Prism Central instances, see Restrict Nutanix v4 event discovery to specific Prism Centrals.
- Enable the event discovery scheduled job. For more information, see Enable Nutanix event discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Nutanix Components pattern.

Note: API v4 calls might retrieve or populate different values compared to v1-3. For more information, search the Nutanix documentation for the "API Reference v4 Introduction" article.

Nutanix Controller VM [cmdb_ci_nutanix_controller_vm]

Field	Description
Name [name]	The name of the Nutanix Controller VM.
VM ID [vm_id]	The universally unique identifier (UUID) of the Nutanix Controller VM.
Object ID [object_id]	The ID of the Nutanix Controller VM.
Memory (MB) [memory]	The memory size of the Nutanix Controller VM in Megabytes (MB).
State [state]	The power state of the Nutanix Controller VM (on/off).
Hypervisor Type [hypervisor_type]	The type of the hypervisor used for the Nutanix Controller VM.
Description [short_description]	The description defined for the virtual machine in the Nutanix console: Prism Central or Prism Element.

Nutanix Cluster [cmdb_ci_nutanix_cluster]

Field	Description
Name [name]	The name of the Nutanix Cluster.
IP Address [ip_address]	The IP address of the Nutanix Cluster.
Cluster ID [cluster_id]	The ID of the Nutanix Cluster.
Object ID [object_id]	The object ID of the Nutanix Cluster.
Version [version]	The Nutanix Cluster version.
Full Version [full_version]	The full version of the Nutanix Cluster.
Internal Subnet [internal_subnet]	The subnet of the internal IP address of the Nutanix Cluster.
External Subnet [external_subnet]	The subnet of the external IP address of the Nutanix Cluster.
Timezone [timezone]	The time zone of the Nutanix Cluster.
Hypervisor Types [hypervisor_types]	The type of the hypervisor used for the Nutanix Cluster.
NTP Servers [ntp_servers]	The Network Time Protocol (NTP) servers that are related to the Nutanix Cluster.
Number of Nodes [num_nodes]	The number of nodes connected to the Nutanix Cluster.
NCC Version [ncc_version]	The Nutanix Cluster Check (NCC) version.
Block Serial Numbers [block_serial_numbers]	The serial numbers of blocks of data connected to the Nutanix Cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Nutanix Host [cmdb_ci_nutanix_host]

Field	Description
Name [name]	The name of the Nutanix host.
Object ID [object_id]	The ID of the Nutanix host.
CPU core count [cpu_core_count]	The number of Nutanix host cores.
CPU core thread [cpu_core_thread]	The number of threads of the Nutanix host.
CPU speed (MHz) [cpu_speed]	The CPU speed of the Nutanix host in MHz.
CPU count [cpu_count]	The number of CPU sockets of the Nutanix host.
RAM (MB) [ram]	The memory of the host in Megabites (MB).
Serial number [serial_number]	Serial number of the Nutanix host. This field is not applicable for v4.

Nutanix Storage Pool [cmdb_ci_nutanix_storage_pool]

Field	Description
Name [name]	• v1-3: The name of the Nutanix storage pool. • v4: Cluster name.
Pool ID [pool_id]	• v1-3: The UUID of the Nutanix storage pool. • v4: The external ID (extId) of the pool retrieved from the Nutanix Storage Container.
Object ID [object_id]	• v1-3: The ID of the Nutanix storage pool. • v4: The external ID (extId) of the pool retrieved from the Nutanix Storage Container.
Size [size]	The size of the Nutanix storage pool. This field is not applicable for v4.

Nutanix Storage Container [cmdb_ci_nutanix_storage_container]

Field	Description
Object ID [object_id]	The ID of the Nutanix container.
Container ID [container_id]	The UUID of the Nutanix container.
Name [name]	The name of the Nutanix container.
Replication Factor [replication_factor]	The number of copies created for the data stored on this Nutanix storage container.
Deduplication [deduplication]	The deduplication is enabled or disabled for the Nutanix container.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Nutanix Storage Container [cmdb_ci_nutanix_storage_container] (continued)

Field	Description
Compression [compression]	The compression is enabled or disabled for the Nutanix container.
Erasure Code [erasure_code]	The erasure of code is enabled or disabled for the Nutanix container.
Size [size]	The Nutanix container size calculated dynamically during discovery.

Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]

Field	Description
Name [name]	The name of the Nutanix Controller VM.
VM Instance ID [vm_inst_id]	The universally unique identifier (UUID) of the Nutanix Controller VM.
Object ID [object_id]	The ID of the Nutanix Controller VM.
CPUs [cpus]	The number of CPUs that the VM has.
Memory (MB) [memory]	The memory size of the Nutanix Controller VM in Megabytes (MB).
State [state]	The power state of the Nutanix Controller VM (on/off).
Description [short_description]	The description defined for the virtual machine in the Nutanix console: Prism Central or Prism Element.

Nutanix Prism Central [cmdb_ci_nutanix_prism_central]

Field	Description

Name [name]	The name of the Nutanix Prism Central multi-cluster manager.
	• **PrismCentral@<IP_ADDRESS>**: When the IP address is associated with the Prism Central instance, in the following scenarios:
	  ○ During quick discovery using HTTP classification.
	  ○ During serverless discovery if IP address was entered in **PRISM_Address** parameter value.
	• **PrismCentral@<DNS_NAME>**: When the DNS name is associated with the Prism Central instance and was entered in **PRISM_Address** parameter value during serverless discovery.

Object ID [object_id]	The ID of the Nutanix Prism Central. This field has the same value as the **Name** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Nutanix Prism Central [cmdb_ci_nutanix_prism_central] (continued)

Field	Description

	• **PrismCentral@<IP_ADDRESS>**: When the IP address is associated with the Prism Central instance, in the following scenarios:
	  ○ During quick discovery using HTTP classification.
	  ○ During serverless discovery if IP address was entered in **PRISM_Address** parameter value.
	• **PrismCentral@<DNS_NAME>**: When the DNS name is associated with the Prism Central instance and was entered in **PRISM_Address** parameter value during serverless discovery.

The graphic illustrates CIs that are part of the Acropolis discovery in deployments that use Nutanix AHV for virtualization.

Nutanix components

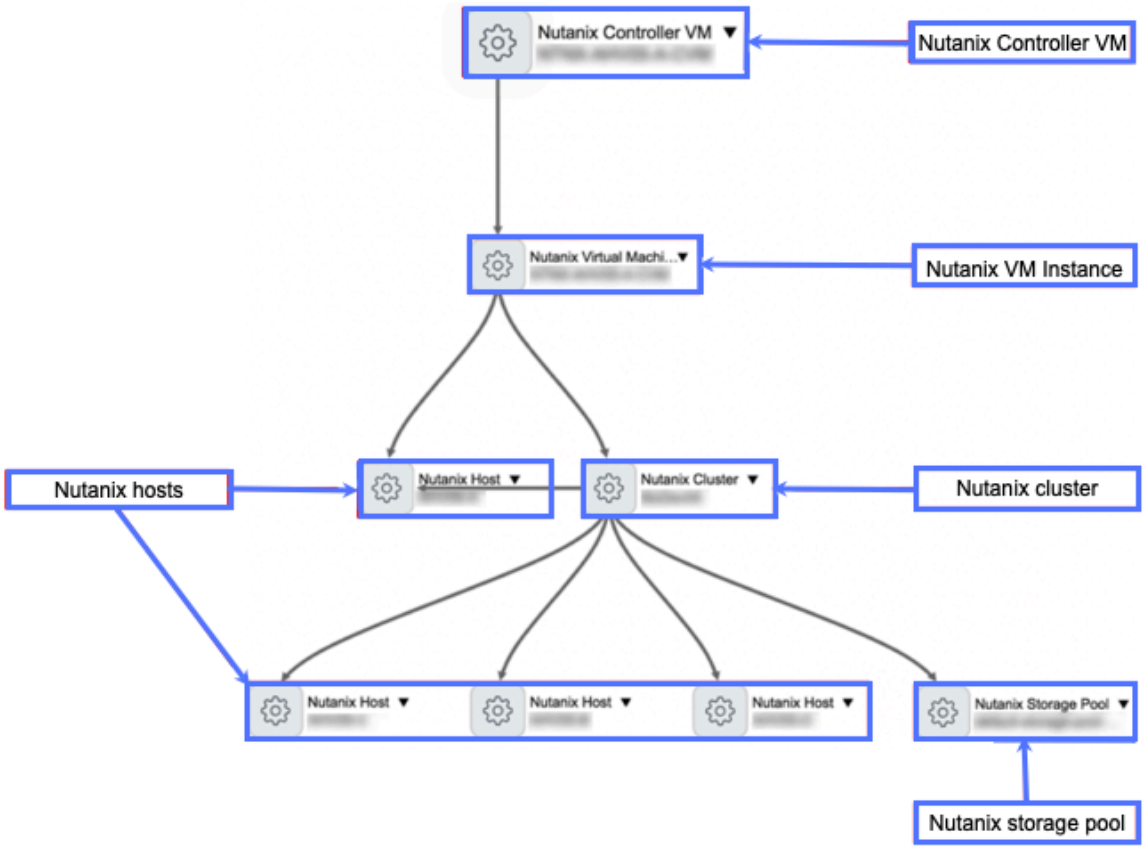

The graphic illustrates CIs that are part of the Acropolis discovery in deployments that use VMware software for virtualization.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Nutanix components in deployments using VMware software

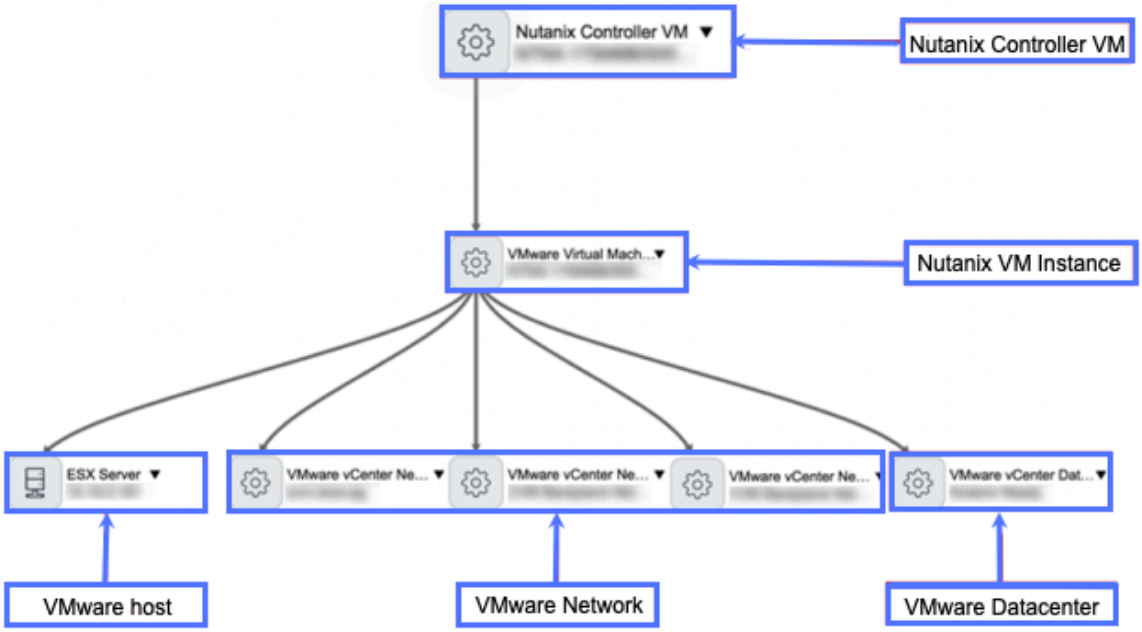

The graphic illustrates the Server and Nutanix VM Instance relationship.

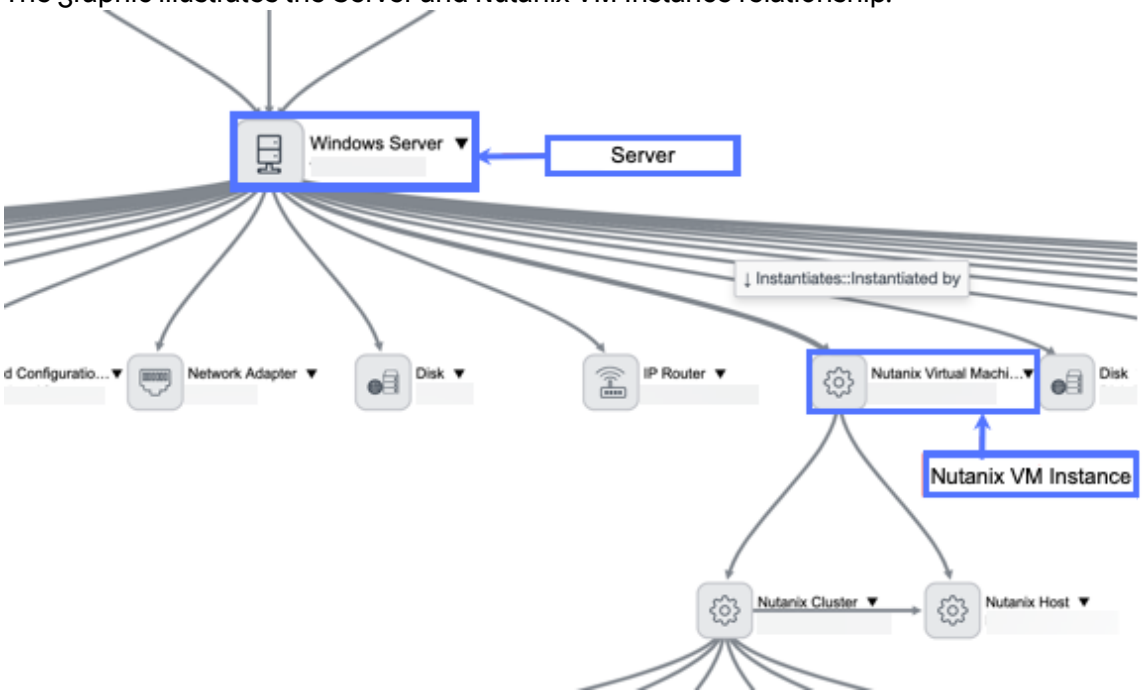

CI relationships

These relationships are created to support all Nutanix discovery, regardless of hypervisor.

CI	Relationship	CI
Nutanix Cluster [cmdb_ci_nutanix_cluster]	Members::Member of	Nutanix Host [cmdb_ci_nutanix_host]
Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]	Registered::Has registered	Nutanix Host [cmdb_ci_nutanix_host]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]	Hosted on::Hosts	Nutanix Cluster [cmdb_ci_nutanix_cluster]
Nutanix Storage Pool [cmdb_ci_nutanix_storage_pool]	Contains::Contained by	Nutanix Storage Container [cmdb_ci_nutanix_storage_container]
Servers [cmdb_ci_server]	Instantiates::Instantiated by	Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]
Nutanix Cluster [cmdb_ci_nutanix_cluster]	Managed by::Manages	Nutanix Prism Central [cmdb_ci_nutanix_prism_central]

These relationships are created to support Nutanix discovery for components hosted on AHV clusters.

CI	Relationship	CI
Nutanix Controller VM [cmdb_ci_nutanix_controller_vm]	Runs on::Runs	Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]
Nutanix Cluster [cmdb_ci_nutanix_cluster]	Defines resources for::Gets resources from	Nutanix Storage Pool [cmdb_ci_nutanix_storage_pool]

These relationships are created to support Nutanix discovery for components hosted on ESX clusters if you ran a VMware discovery.

CI	Relationship	CI
Nutanix Controller VM [cmdb_ci_nutanix_controller_vm]	Runs on::Runs	VMware Virtual Machine Instance [cmdb_ci_vmware_instance]
VMware Cluster [cmdb_ci_vmware_cluster]	Defines resources for::Gets resources from	Nutanix Storage Pool [cmdb_ci_nutanix_storage_pool]

These relationships are created to support Nutanix discovery for components hosted on ESX clusters if you didn't run a VMware discovery.

CI	Relationship	CI
Nutanix Controller VM [cmdb_ci_nutanix_controller_vm]	Hosted on::Hosts	Nutanix Prism Central [cmdb_ci_nutanix_prism_central]
Nutanix Storage Pool [cmdb_ci_nutanix_storage_pool]	Hosted on::Hosts	Nutanix Prism Central [cmdb_ci_nutanix_prism_central]

Tag discovery

The Nutanix Components pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The category key.
Value [value]	The category value.
Configuration item [configuration_item]	References the Nutanix Host [cmdb_ci_nutanix_host] table.

Dashboard

The dashboard contains reports that query the Configuration Management Database (CMDB) and display the results. The information is displayed only after the pattern has discovered the infrastructure.

To display the dashboard, navigate to **Self-Service** > **Dashboards** > **All** > **Nutanix**.

After upgrading to Discovery Admin Workspace version 1.3.1 (August 2024 Store), you can navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** and use the enhanced dashboard.

Create a serverless schedule for Nutanix Acropolis discovery

Create a serverless schedule to discover Nutanix Acropolis using REST API.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule New record form

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type, which should be Serverless.
MID server	Name of the MID Server with access to the Nutanix environment.
Active	Enables this schedule for discovery.
Run	Determines the run schedule of the discovery.

   c. Select **Submit**.
3. Create the execution pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the Discovery Schedules page, select the record you created.

b. In the **Serverless Execution Patterns** tab, select **New**.

c. On the form, fill in the fields.

Serverless Execution Pattern New record form

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern to be used for this schedule, which should be Nutanix Components.
Proxy Host	Fully qualified domain name (FQDN) of the machine where the proxy server is installed, which should be Global. This field is optional.
Active	Enables this schedule for discovery.

d. Select **Submit**.

4. Set the pattern launcher parameters.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Parameter	Value
PRISM_Address	IP Address or the FQDN associated with the Prism Central instance.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

Nutanix Acropolis discovery

Enable Nutanix v4 discovery

Enable Nutanix v4 discovery by setting the MID Server property value to the exact v4 version.

Before you begin

Verify that you have Discovery and Service Mapping Patterns installed (starting with version 1.29.0).

Role required: discovery_admin

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** column, search for the `sn_itom_pattern.nutanix_prism_version` property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select the **sn_itom_pattern.nutanix_prism_version** property.
4. In the **Value** field, enter the exact v4 version to discover.
   For example: `v4.0.b1`.
5. Select **Update**.

Restrict Nutanix v4 event discovery to specific Prism Centrals

Create a MID Server property to target specific Prism Central instances for Nutanix v4 event discovery.

Before you begin

- Verify that you have at least version 1.31.0 of Discovery and Service Mapping Patterns.
- Verify that Nutanix Components discovery has run and that the Nutanix Prism Central [cmdb_ci_nutanix_prism_central] table is populated.
- Verify that the MID Servers for Nutanix event discovery have **All** or **Nutanix** capability.

Role required: discovery_admin.

About this task

By default, Nutanix v4 event discovery runs against all Prism Central instances in the Nutanix Prism Central [cmdb_ci_nutanix_prism_central] table, using any MID Server with the **All** or **Nutanix** capability. Starting with Discovery and Service Mapping Patterns version 1.31.0, you can create a MID Server property to target specific Prism Central instances. You can create multiple properties for different Prism Central and MID Server combinations.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. Select **New**.
3. Fill in the form.

MID Server Property form

Field	Description
Name	Property name. Enter `sn_itom_pattern.nutanix_events_v4.prism_ips`.
Value	Comma-separated list of IP addresses or hostnames of the Prism Central instances to target. The instances must already be present in the Nutanix Prism Central [cmdb_ci_nutanix_prism_central] table.
MID server	MID Server to use for these Prism Central instances. Leave empty to use any MID Server with the **All** or **Nutanix** capability.

4. Select **Submit**.

What to do next

Enable the Nutanix event discovery scheduled job. For more information, see Enable Nutanix event discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Nutanix Acropolis discovery

Enable Nutanix event discovery

Enable Nutanix event discovery

Activate the Nutanix events scheduled job to trigger rediscovery of Nutanix components when a state change is detected.

Before you begin

- Verify that Nutanix Components discovery has run and that the Nutanix Prism Central [cmdb_ci_nutanix_prism_central] table is populated.
- Verify that the MID Servers for Nutanix event discovery have **All** or **Nutanix** capability.
- For Nutanix v4 event discovery, verify that you have at least version 1.31.0 of Discovery and Service Mapping Patterns.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. In the **Name** column, search for the scheduled job.
   - For API v1–3, search for `Nutanix-Events-job`.
   - For API v4, search for `Nutanix-Events-job-V4`.
3. Select the scheduled job.
4. Select the **Active** check box.
   The job runs every five minutes by default. You can change the run frequency using the **Run** field.
5. Select an option to run the job:
   - Select **Update** to enable the job to run on its configured schedule.
   - Select **Execute Now** to run the job immediately.

Related topics

Nutanix Acropolis discovery

NVIDIA GPU discovery using patterns

The Discovery and Service Mapping Patterns application uses the Linux Server pattern extension **Discover Nvidia GPU** to find NVIDIA graphics processing units (GPUs). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The Discover Nvidia GPU pattern is an extension of the Linux Server pattern. For more information, see Linux discovery.

Starting with Visibility Content version 6.32.0, NVIDIA GPU discovery requires the AI Agent Topology Mapping application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models
- Visibility Content
- AI Agent Topology Mapping

Verify permissions for the following SSH command

Verify that you have **read access** permission for the SSH command: `nvidia-smi -q`.

Create SSH credentials

For more information, see SSH credentials.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Discover Nvidia GPU Pattern.

Graphics Processing Unit [cmdb_ci_gpu]

Field	Description
Name [name]	Model name of the GPU.
Manufacturer [manufacturer]	Vendor of the GPU.
Model ID [model_id]	Product name of the GPU.
Serial number [serial_number]	Serial number of the GPU.

Additionally, Discovery populates the data in a non-CMDB table when running the Discover Nvidia GPU Pattern.

GPU Details [gpu_details]

Field	Description
Board ID [board_id]	Identifier for the GPU board, used for diagnostic purposes.
Board Part Number [board_part_number]	The part number of the GPU board.
Bus ID [bus_id]	Identifier for the peripheral component interconnect (PCI) bus on which the GPU is located.
CUDA Version [cuda_version]	Version of the CUDA Toolkit supported by the installed NVIDIA driver.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GPU Details [gpu_details] (continued)

Field	Description
Driver version [driver_version]	Version of the NVIDIA driver installed on the system.
GSP Firmware Version [gsp_firmware_version]	Version of the GPU System Processor (GSP) firmware.
Index [index]	Index number of the GPU as reported by the system.
Inforom ECC [inforom_ecc]	Version of the error-correcting code (ECC) used to record data in the InfoROM.
Inforom IMG [inforom_img]	Global version of the infoROM image.
Inforom OEM [inforom_oem]	Version of the original equipment manufacturer (OEM) configuration data in the InfoROM.
Inforom PWR [inforom_pwr]	Version of the power management data in the InfoROM.
Memory Free Usage [memory_free_usage]	Amount of GPU memory currently free.
Memory Total Usage [memory_total_usage]	Total amount of GPU memory installed.
Memory Used Usage [memory_used_usage]	Amount of GPU memory currently in use.
Name [name]	Official product name of the GPU.
PCIE Current Link [pcie_current_link]	Current peripheral component interconnect express (PCIe) link generation.
Performance State [performance_state]	Power and performance levels of the GPU. Possible values run from P0 to P15, with lower numbers indicating higher performance.
Persistence Mode Status [persistence_mode_status]	Indicates whether persistence mode is enabled for the GPU. Possible values are Enabled or Disabled.
Serial Number [serial_number]	Serial number of the GPU.
UUID [uuid]	Universally unique identifier (UUID) of the GPU.
Vbios Version [vbios_version]	Basic Input/Output System (BIOS) of the GPU board.

CI relationships

Discovery creates these relationships to support the NVIDIA GPU discovery.

CI	Relationship	CI
Graphics Processing Unit [cmdb_ci_gpu]	Used by::Uses	Linux Server [cmdb_ci_linux_server]

Oracle Linux Virtualization Manager and Red Hat Virtualization discovery

The ServiceNow Discovery application uses patterns to find Oracle Linux Virtualization Manager (OLVM) and Red Hat Virtualization (RHV) components. Discovering some of these resources

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Starting with version 1.28.0, Discovery and Service Mapping Patterns supports both OLVM and RHV discovery. Note that the RHV patterns, tables, and related items have been renamed to oVirt.

Discovery uses the following patterns for OLVM and RHV discovery:

- oVirt Clusters and Hosts
- oVirt Discover Logical Datacenters
- oVirt Discover Manager Instance
- oVirt Virtual Machines
- oVirt Templates (starting with Discovery and Service Mapping Patterns version 1.29.0)

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

OLVM and RHV data model

The OLVM and RHV patterns introduce the following CI classes that are part of the data model.

CI class hierarchy

CI class	Extends from
oVirt LDC [cmdb_ci_rhv_ldc]	Logical Datacenter [cmdb_ci_logical_datacenter]
oVirt Cluster [cmdb_ci_rhv_cluster]	oVirt Object [cmdb_ci_rhv_object]
oVirt Server [cmdb_ci_rhv_server]	Virtualization Server [cmdb_ci_virtualization_server]
oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]	Virtual Machine Instance [cmdb_ci_vm_instance]
Cloud Mgmt Network Interface [cmdb_ci_nic]	Network Adapter [cmdb_ci_network_adapter]
Cloud Service Account [cmdb_ci_cloud_service_account]	Virtual Machine Object [cmdb_ci_vm_object]
IP Address [cmdb_ci_ip_address]	Configuration Item [cmdb_ci]
VNIC Endpoint [cmdb_ci_endpoint_vnic]	Endpoint [cmdb_ci_endpoint]

Prerequisites

Deploy the MID Server

Deploy and connect the dedicated MID Server to access either the OLVM or Red Hat Virtualization Manager (RHV-M) instances.

Create oVirt credentials

Configure the following credentials:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Navigate to **Discovery** > **Credentials**.
2. Select **New**.
3. Select **oVirt Credentials**, fill in the form, and then select **Submit**.

Field	Description
Name	Credential name.
User name	Name of the user for performing this discovery.
Password	Password for this user.
Active	Option to select to enable this credential for discovery.
Bearer Token	Option to select to enable use of a bearer token.

Verify API permissions

Verify that the user has read-only permissions to send the following queries.

- /ovirt-engine/api/
- /ovirt-engine/api/datacenters
- /ovirt-engine/api/clusters
- /ovirt-engine/api/hosts
- /ovirt-engine/api/vms
- /ovirt-engine/api/disks
- /ovirt-engine/api/templates
- /ovirt-engine/api/networks

Add CMDB CI Class Models

Add class models that extend the CMDB class hierarchy, using the CMDB CI Class Models application. The app is available from the ServiceNow Store. Discovery patterns use these class extensions to populate CIs. For more information, see Red Hat Virtualization (RHV) extension classes.

Create a cloud service account

Create an OLVM or RHV cloud service account on the ServiceNow instance.

Note: Confirm that Discovery Admin Workspace is using at least version 1.10.0. The **Discovery** > **Cloud Service Accounts** navigation module isn't available with earlier versions. To access **Cloud Service Accounts** with an earlier version, enter in the navigation filter: `cmdb_ci_cloud_service_account.list`.

1. Navigate to **All** > **Discovery** > **Cloud Service Accounts**
2. Select **New**.
3. On the form, fill in the fields.

Field	Description
Name	A unique name for the service account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Account Id	The account ID to which this credential belongs. The API url of the instance should be specified, including the port. For example: `https://my-manager-instance.com:443`.
Discovery credentials	The discovery credentials that were created previously.
Datacenter Type	The datacenter type, which should be oVirt LDC [cmdb_ci_rhv_ldc].

Create a discovery schedule

For more information, see Create an oVirt Discovery schedule in Discovery Admin Workspace.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the patterns.

oVirt LDC [cmdb_ci_rhv_ldc]

Field	Description
Name [name]	Name of the datacenter.
Object ID [object_id]	Unique object identifier.
Description [short_description]	Description of the datacenter that you provide for easy identification.
Quota Mode [quota_mode]	Quota mode policy.
Href ID [href_id]	Href ID.

oVirt Cluster [cmdb_ci_rhv_cluster]

Field	Description
Name [name]	Name of the cluster.
Object ID [object_id]	Unique object identifier.
Description [short_description]	Description of the cluster that you provide for easy identification.
CPU Type [cpu_type]	CPU type. For example, x86_64.
CPU Architecture [cpu_architecture]	CPU architecture. For example, Intel Broadwell IBRS SSBD MDS Family.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

oVirt Cluster [cmdb_ci_rhv_cluster] (continued)

Field	Description
KSM [ksm]	Enabled state of kernel same-page merging (KSM) memory policy.
Fencing Policy [fencing_policy]	Fencing policy.
Memory Overcommit [memory_overcommit]	Amount of over-commitment memory allowed on the cluster.
Transparent Huge Pages [thp]	Transparent huge memory pages policy.
Ballooning [ballooning]	Memory ballooning for guests.
Href ID [href_id]	Href ID.

oVirt Server [cmdb_ci_rhv_server]

Field	Description
Name [name]	Name of the server.
Serial number [serial_number]	Serial number.
Object ID [object_id]	Unique object identifier.
CPU core count [cpu_core_count]	Total number of cores.
CPU count [cpu_count]	Total number of CPUs.
CPU type [cpu_type]	CPU type.
CPU speed (MHz) [cpu_speed]	CPU speed (MHz).
Host name [host_name]	Hostname.
IP Address [ip_address]	IP address of the server.
Model ID [model_id]	Model ID.
Operating System [os]	OS family.
OS Version [os_version]	OS version.
RAM (MB) [ram]	Amount of RAM (MB).

oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]

Field	Description
Name [name]	Name of the instance.
State [state]	State of the instance.
Object ID [object_id]	Unique object identifier.
IP Address [ip_address]	IP address of the instance.
Disks [disks]	Number of disks.
Description [short_description]	Description of the instance that you provide for easy identification.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance] (continued)

Field	Description
Fully qualified domain name [fqdn]	Fully Qualified Domain Name (FQDN) of the instance.
CPU Architecture [cpu_architecture]	CPU architecture. For example: Intel Broadwell IBRS SSBD MDS Family.
Delete Protected [delete_protected]	Delete protected, possible values: true or false.
CPUs [cpus]	Number of CPUs.
HA Priority [ha_priority]	High availability (HA) priority.
High Availability [high_availability]	High availability, possible values: true or false.
Memory MAX [memory_policy_max]	Maximum memory in the dynamic memory allocation policy for the virtual machine (VM).
Memory Guarantee (MB) [memory_policy_guaranteed]	Amount of memory guaranteed for the VM (MB).
Multi Queues [multi_queues]	Multi queues.
Placement Policy [placement_policy]	Placement policy.
Stateless [stateless]	Stateless, possible values: true or false.
On Storage Error [storage_error_resume_behaviour]	Behavior of a VM that is paused due to storage I/O error. For example: AUTO_RESUME, LEAVE_PAUSED, and KILL.
Start Time [start_time]	Start time of the VM.
Stop Time [stop_time]	Stop time of the VM.
Threads [threads]	Number of threads.
Time Zone [time_zone]	Time zone of the VM.
USB [usb]	USB enabled state.
Run Once [run_once]	Run once.
RHV Type [rhv_type]	OLVM or RHV type.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Account ID for this cloud service account.
Object ID [object_id]	Account ID for this cloud service account.
Datacenter Type [datacenter_type]	Type of datacenter: oVirt LDC [cmdb_ci_rhv_ldc].
Discovery credentials [discovery_credentials]	Credentials for this cloud service account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Address [cmdb_ci_ip_address]

Field	Description
Name [name]	The IP address of the network interface card (NIC) assigned to the VM.
IP Address [ip_address]	The IP address of the NIC assigned to the VM.
IP version [ip_version]	IP version: IPv4 or IPv6.
Nic [nic]	References the Cloud Mgmt Network Interface [cmdb_ci_nic] associated with this IP address.

VNIC Endpoint [cmdb_ci_endpoint_vnic]

Field	Description
Object ID [object_id]	Unique identifier for this resource.
Name [name]	Name for the virtual network interface card (vNIC) endpoint.
IP Address [ip_address]	IP address bound to the network interface.

Cloud Mgmt Network Interface [cmdb_ci_nic]

Field	Description
Object ID [object_id]	ID of the network interface.
Name [name]	Name of the network interface.
IP Address [ip_address]	IP address bound to the network interface.
Public IP [public_ip]	Public IP address bound to the network interface.
Private IP [private_ip]	Private IP address bound to the network interface.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	Name of the disk.
Object ID [object_id]	Unique ID of the disk.
Size [size]	Storage capacity of the disk.
Storage type [storage_type]	Type of storage backend the disk is using.

Cloud Network [cmdb_ci_network]

Field	Description
Name [name]	Name of the network.
Object ID [object_id]	Unique ID of the network.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Name [name]	Name of the template.
Object ID [object_id]	Unique ID of the template.
Cores [cores]	Number of CPU cores.
Install Status [install_status]	Install status of the template. Default value is Installed.
Memory MB [memory_mb]	Memory in megabytes (MB).
Provider [provider]	Provider for the template. Value is set to **oVirt**.

On the Dependency Views map, you can view all discovered OLVM or RHV resources in your organization and the relationships between them.

OLVM or RHV dependency view

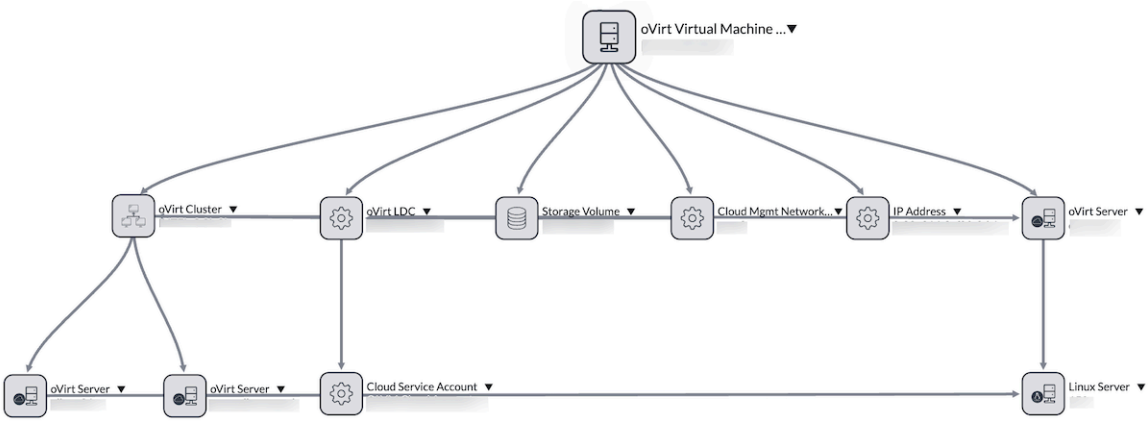

CI relationships

The OLVM and RHV patterns create the following relationships and references to support OLVM and RHV discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	oVirt LDC [cmdb_ci_rhv_ldc]
Cloud Mgmt Network Interface [cmdb_ci_nic]	Hosted on::Hosts	oVirt LDC [cmdb_ci_rhv_ldc]
Cloud Mgmt Network Interface [cmdb_ci_nic]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Cloud Network [cmdb_ci_network]	Hosted on::Hosts	oVirt LDC [cmdb_ci_rhv_ldc]
oVirt Cluster [cmdb_ci_rhv_cluster]	Members::Member of	oVirt Server [cmdb_ci_rhv_server]
oVirt LDC [cmdb_ci_rhv_ldc]	Contains::Contained by	oVirt Cluster [cmdb_ci_rhv_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
oVirt LDC [cmdb_ci_rhv_ldc]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
oVirt Server [cmdb_ci_rhv_server]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]
oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]	Hosted on::Hosts	oVirt Cluster [cmdb_ci_rhv_cluster]
oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]	Hosted on::Hosts	oVirt LDC [cmdb_ci_rhv_ldc]
oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]	Registered on::Has registered	oVirt Server [cmdb_ci_rhv_server]
oVirt Virtual Machine Instance [cmdb_ci_rhv_vm_instance]	Use End Point To::Use End Point From	VNIC Endpoint [cmdb_ci_endpoint_vnic]
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	oVirt LDC [cmdb_ci_rhv_ldc]
VNIC Endpoint [cmdb_ci_endpoint_vnic]	Implement End Point To::Implement End Point From	Cloud Mgmt Network Interface [cmdb_ci_nic]

CI references

CI	Field	Referenced CI
IP Address [cmdb_ci_ip_address]	Nic [nic]	Cloud Mgmt Network Interface [cmdb_ci_nic]

Oracle Access Manager pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Access Manager pattern to find Oracle Access Manager instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify read access to the version file

Verify you have permission for the following command: `powershell -Command "Select-String -Path '{process.environmentVariables.MW_HOME.value}\inventory\ContentsXML\comps.xml' -Pattern 'version' | Select-Object -First 5"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Access Manager pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle Access Manager@{hostname}**.
Product [product]	Product name. The value is set to **Access Manager**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Access Manager pattern creates the following relationships and references to support Oracle Access Manager discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Analytics Server pattern-based discovery

Discovery and Service Mapping Patterns finds Oracle Analytics Server (formerly Oracle Business Intelligence Enterprise Edition) components on Windows and Linux servers in your environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Oracle Analytics Server consists of integrated components that provide business intelligence capabilities. The Discovery patterns discover three core components:

- Oracle BI Server: Core analytics engine that processes queries and manages metadata
- Oracle BI Presentation Service: Web-based user interface component that delivers visualizations and dashboards
- Oracle BI Scheduler: Component that manages scheduled report generation and distribution

Prerequisites

Verify that the applications are up to date

Verify that the following applications are up to date:

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models

Verify that the Oracle Analytics Server processes are running on target systems

Verify that the following Oracle Analytics Server processes are active on target systems:

- Windows:
  - `sawserver.exe`
  - `NQScheduler.exe`
  - `nqsserver.exe`
- Linux:
  - `sawserver`
  - `nqscheduler`
  - `nqsserver`

Verify file system access permissions

Verify that the credentials used for discovery have read access to the following locations:

- Oracle BI installation directories
- `BI_ORACLE_HOME` environment variable location
- `bifoundation/version.txt` file within the Oracle BI installation

Create credentials on the ServiceNow AI Platform

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Windows: Configure Windows credentials. For more information, see Windows credentials.
- Linux: Configure SSH credentials. For more information, see SSH credentials.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle Analytics Server patterns.

Resources discovered by the Oracle BI Server pattern

Oracle BI Server [cmdb_ci_oracle_bi_server]

Field	Description
Name [name]	Name of the Oracle BI Server instance.
Version [version]	Version of the Oracle BI Server software.
Installation directory [install_directory]	File system path where the Oracle BI Server is installed.

Resources discovered by the Oracle BI Presentation Service pattern

Oracle BI Presentation Service [cmdb_ci_oracle_bi_presentation_service]

Field	Description
Name [name]	Name of the Oracle BI Presentation Service instance.
Version [version]	Version of the Oracle BI Presentation Service software.
Installation directory [install_directory]	File system path where the Oracle BI Presentation Service is installed.

Resources discovered by the Oracle BI Scheduler pattern

Oracle BI Scheduler [cmdb_ci_oracle_bi_scheduler]

Field	Description
Name [name]	Name of the Oracle BI Scheduler instance.
Version [version]	Version of the Oracle BI Scheduler software.
Installation directory [install_directory]	File system path where the Oracle BI Scheduler is installed.

CI relationships

The Oracle Analytics Server patterns create these relationships to support Oracle Analytics Server discovery.

Relationship created by the Oracle BI Server pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Oracle BI Server [cmdb_ci_oracle_bi_server]	Runs on::Runs	Windows Server [cmdb_ci_win_server] or Linux Server [cmdb_ci_linux_server]

Relationships created by the Oracle BI Presentation Service pattern

CI	Relationship	CI
Oracle BI Presentation Service [cmdb_ci_oracle_bi_presentation_service]	Runs on::Runs	Windows Server [cmdb_ci_win_server] or Linux Server [cmdb_ci_linux_server]

Relationships created by the Oracle BI Scheduler pattern

CI	Relationship	CI
Oracle BI Scheduler [cmdb_ci_oracle_bi_scheduler]	Runs on::Runs	Windows Server [cmdb_ci_win_server] or Linux Server [cmdb_ci_linux_server]

Oracle Catalog discovery

The Discovery and Service Mapping Patterns application uses the **Get Catalog info** pattern extension of the Oracle DB on Unix and Oracle DB on Windows patterns to find Oracle Catalog objects. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses the Oracle DB on Unix and Oracle DB on Windows patterns and the Get Catalog info pattern extension to run horizontal discovery.

You can use this pattern on the ServiceNow AI Platform® using the Madrid release or later.

For information on Oracle database discovery, see Oracle database discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Provide a user with permissions to run queries on ALL_TABLES

Provide a user with permissions to run queries on ALL_TABLES.

To view other users' tables in ALL_TABLES, a user must have any of the following permissions:

- Object-level privileges, for example: SELECT.
- Higher-level role, for example: SELECT ANY TABLE.

Configure applicative credentials

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure applicative credentials for the Oracle instance [cmdb_ci_db_ora_instance] CI type. For more information, see Applicative credentials.

Configure Windows credentials for Windows operating system (OS)

For more information, see Windows credentials.

Configure UNIX credentials for UNIX OS

For more information, see SSH credentials.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Get Catalog info Pattern extension.

Table and field	Description
Oracle Catalog [cmdb_ci_db_ora_catalog]	
database_instance	The name of the database instance (SID).
name	The name of the CI in the CMDB; schema and SID.
schema	The name of the catalog object.

The following is a dependency view of the relationship between an Oracle Catalog and an Oracle instance.

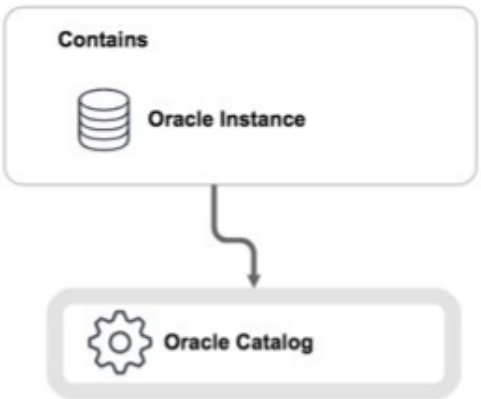

CI relationships

This relationship is created to support Oracle Catalog discovery.

CI	Relationship	CI
Oracle Catalog [cmdb_ci_db_ora_catalog]	Contains::Contained by	Oracle Instance [cmdb_ci_db_ora_instance]

Related topics

Oracle pluggable database and container database discovery

Oracle GLAS data collection using patterns

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Coherence pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Coherence pattern to find Oracle Coherence instances running on UNIX servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create SSH credentials

For more information, see SSH credentials.

Verify read access to the version file

Verify you have permission for the following command: `cat {process.environmentVariables.EMHOME.value}/inventory/ContentsXML/comps.xml | grep -i version`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Coherence pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Coherence@{hostname}**.
Product [product]	Product name. The value is set to **Coherence**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships and references

The Oracle - Coherence pattern creates the following relationships and references to support Oracle Coherence discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle Enterprise Manager Management Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Enterprise Manager Management Server pattern to find Oracle Enterprise Manager Management Server instances running on UNIX servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create SSH credentials

For more information, see SSH credentials.

Verify permission to run the version command

Verify you have permission for the following command: `{process.environmentVariables.EMHOME.value}/OPatch/opatch version`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Enterprise Manager Management Server pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Enterprise Manager Management Server@{hostname}**.
Product [product]	Product name. The value is set to **Enterprise Manager Management Server**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Enterprise Manager Management Server pattern creates the following relationships and references to support Oracle Enterprise Manager Management Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle Forms pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Forms pattern to find Forms instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the version file

Verify you have permission for the following command:
`powershell -Command "Select-String -Path '{process.environmentVariables.MW_HOME.value}\inventory\ContentsXML\comps.xml' -Pattern 'version' | Select-Object -First 5"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Forms pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Forms@{hostname}**.
Product [product]	Product name. The value is set to **Forms**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Forms pattern creates the following relationships and references to support Oracle Forms discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle GLAS data collection using patterns

The ServiceNow Discovery application uses the Oracle GLAS Data Collection pattern extensions to discover Oracle GLAS data. This data includes discovered Oracle Database, Middleware, and Java configuration items.

The following patterns collect Oracle GLAS data.

Oracle GLAS patterns

Pattern	Extends pattern
Oracle GLAS Data Collection	• Oracle DB On Unix • Oracle DB On Windows
Oracle GLAS Data Collection V2	• Oracle DB On Unix • Oracle DB On Windows
Oracle GLAS - Collect Middleware Data for Tuxedo	Tuxedo
Oracle GLAS - Collect Middleware Data for Weblogic	WebLogic
Oracle GLAS Data Collection for Non-SQL	Linux Server
Oracle GLAS HW Data Collection for Non-SQL	Windows OS - Servers
Java Installation	N/A (standalone pattern)

For details on how to view and download the collected Oracle GLAS data, see Download Oracle Global License Advisory Services (GLAS) data.

To learn about the supported versions of Oracle, refer to Detailed information on products discovered by ITOM Visibility.

The Data Collection for Oracle Global Licensing and Advisory Services application supports two database data collection methods. The default method, V1, stores data across six CMDB CI tables. V2, available starting with Data Collection for Oracle Global Licensing and Advisory Services version 1.12.0, is an alternative method optimized for large-scale or high-volume Oracle database environments that consolidates data into a single non-CMDB table. For more information about V2, see the "Oracle GLAS Data Collection V2" section.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

Install the Data Collection for Oracle Global Licensing and Advisory Services store app

Navigate to **System Definition** > **Plugins** and install the Data Collection for Oracle Global Licensing and Advisory Services app.

- The minimum version for Database collection is 1.4.0
- The minimum version for Middleware collection is 1.7.1
- The minimum version for Java collection is 1.8.4

Note: The Data Collection for Oracle Global Licensing and Advisory Services app requires the Discovery (com.snc.discovery) plugin to work correctly.

Verify the following applications are installed

- Discovery and Service Mapping Patterns
- CMDB CI Class Models

Verify operating system privileges

Verify that your Windows or UNIX user can write temporary files under `/tmp`.

Verify that you can run the following operating system commands:

- ls
- cat
- ps

Create operating system credentials

Verify that you have the following credentials for your operating system, and then create credentials for them on the ServiceNow AI Platform:

- UNIX credentials using SSH. For more information, see SSH credentials.
- Windows credentials using PowerShell/WMI. For more information, see Windows credentials.

Note: PowerShell and WMI credentials are required only for Database collection.

Verify the following Oracle privileges

Verify that you have applicative credentials for the Oracle instance (cmdb_ci_db_ora_instance) that have the following privileges:

- Read permissions on the `tnsnames.ora` file.
- Permissions to execute the *`sqlplus`* command.
- (For middleware collection) Sudo permissions to run the following command: `sudo su - oracle -c "$ORACLE_INSTANCE/bin/opmnctl status"`

Verify you have the user name and password for the relevant Oracle Database, with read permissions on the following tables.

DBA_ADVISOR_TASKS	DBA_USERS	SYS.MODEL$

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

DBA_CPU_USAGE_STATISTICS	DMSYS.DM$MODEL	SYS.REGISTRY$HISTORY
DBA_ENCRYPTED_COLUMNS	DMSYS.DM$OBJECT	SYSMAN.MGMT_INV_CONTAINER
DBA_FEATURE_USAGE_STATISTICS	DMSYS.DM$P_MODEL	SYSMAN.MGMT_LICENSES
DBA_FLASHBACK_ARCHIVE_TABLES	DVSYS.DBA_DV_REALM	SYSMAN.MGMT_LICENSE_DEFINITIONS
DBA_INDEXES	GV$IM_SEGMENTS	SYSMAN.MGMT_TARGETS
DBA_LOBS	GV$INSTANCE	SYSMAN.MGMT_VERSIONS
DBA_LOB_PARTITIONS	GV$PARAMETER	V$ARCHIVE_DEST_STATUS
DBA_LOB_SUBPARTITIONS	MDSYS.SDO_FEATURE_USAGE	V$BLOCK_CHANGE_TRACKING
DBA_OBJECTS	MDSYS.SDO_GEOM_METADATA_TABLE	V$CONTAINERS
DBA_RECYCLEBIN	MGMT$TARGET	V$DATABASE
DBA_REGISTRY	ODM.ODM_MINING_MODEL	V$DATAGUARD_CONFIG
DBA_SQLSET	ODM_DOCUMENT	V$LICENSE
DBA_SQLSET_REFERENCES	ODM_RECORD	V$OPTION
DBA_SQL_PROFILES	REDACTION_POLICIES	V$PARAMETER
DBA_TABLES	SCHEMA_VERSION_REGISTRY	V$SESSION
DBA_TABLESPACES	SYS.CDC_CHANGE_SETS$	V$VERSION
DBA_TAB_PARTITIONS	SYS.CDC_CHANGE_TABLES$	dba_aws
DBA_TAB_SUBPARTITIONS	SYS.DBA_MINING_MODELS	

Create applicative credentials for Oracle Instance

Create applicative credentials on the ServiceNow AI Platform. For more information, see Applicative credentials.

Use Oracle Wallet authentication instead of applicative credentials

For more information, see Oracle Wallet authentication for database discovery.

For virtual machines: Verify virtual machine credentials

For virtual machines, verify you have the relevant credentials to collect hardware data for machines hosted on the following virtual providers:

- VMware
- Nutanix
- IBM
- Hyper-V

For virtual machines: Collect hardware data

To populate the CPUQ folder in your GLAS report download, run the Oracle-provided hardware script on each target virtual machine. The following table lists the required script and credentials for each operating system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Operating system	Script	Required credentials
Windows	cpuq.cmd	Username and password
Linux or UNIX	cpuq.sh	Root username and password

Access to the operating system command line is required. Use **Command Prompt** on Windows systems and a UNIX Shell on Linux or UNIX systems.

1. Log in to the target virtual machine using the required credentials.
2. Open the command-line interface for the operating system.
3. Run the hardware script for the operating system.
4. Repeat for each target virtual machine.

For large-scale or high-volume environments: You can enable Oracle GLAS V2 data collection

To use V2 database data collection instead of the default V1 method, set the *sn_itom_oracleglas.enable_large_env_glas_data_collection* system property to **true**. For more information about V2, see the "*Oracle* GLAS Data Collection V2" section. For instructions, see Enable the Oracle GLAS V2 data collection method.

V1 database CI relationships

CI	Relationship/Reference	CI
ORCL V SESSION [cmdb_ci_orcl_vsession]	Contains::Contained by	Oracle Instance [cmdb_ci_db_ora_instance]
ORCL DETAIL [cmdb_ci_ora_lms_detail]	Contains::Contained by	Oracle Instance [cmdb_ci_db_ora_instance]
ORCL OVERVIEW [cmdb_ci_ora_lms_overview]	Contains::Contained by	Oracle Instance [cmdb_ci_db_ora_instance]
ORCL OPTIONS [cmdb_ci_orcl_options]	Contains::Contained by	Oracle Instance [cmdb_ci_db_ora_instance]
ORCL DBA USERS [cmdb_ci_orcl_dba_users]	Contains::Contained by	Oracle Instance [cmdb_ci_db_ora_instance]
ORCL V LICENSE [cmdb_ci_orcl_vlicense]	Contains::Contained by	Oracle Instance [cmdb_ci_db_ora_instance]

Oracle GLAS Data Collection V2

Oracle GLAS Data Collection V2 is an alternative database data collection method that uses the **Oracle GLAS Data Collection V2** pattern extension. V2 optimizes database data collection for large-scale or high-volume Oracle database environments. The V2 method applies to database data collection only; middleware and Java data collection aren't affected.

In V2, the MID Server collects data from the Oracle GLAS scripts and uses the GlideRecord API to populate the Oracle GLAS Database Audit Data [sn_itom_oracleglas_glas_data_for_database] table directly. As a result, GLAS data appears in the Content field and not in the discovery payload. To verify data population, navigate to the Oracle GLAS Database Audit Data table

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

and review the **Content** field. If data is missing or incomplete, see MID Server ECC Queue for information about accessing MID Server logs.

V2 stores all database GLAS data in the Oracle GLAS Database Audit Data table. Each record directly references the Oracle Instance [cmdb_ci_db_ora_instance] table; in V1, this association is established through a CI relationship (see the "V1 database CI relationships" section).

The Oracle GLAS Data Collection V2 pattern extension populates the following fields in the Oracle GLAS Database Audit Data table:

Oracle GLAS Database Audit Data [sn_itom_oracleglas_glas_data_for_database]

Field	Description
Host Name [hostname]	Host name of the server running the Oracle instance.
DB Name [dbname]	Name of the Oracle database.
Type [type]	Identifies the type of GLAS data stored in the record. Possible values: • ORCL_OVERVIEW • ORCL_DETAILS • ORCL_LICENSE • ORCL_SESSION • ORCL_OPTIONS • ORCL_DBA_USERS
Order [order]	Sequential order of the record. If the content for a type exceeds 10 megabytes (MB), it is stored in multiple records ordered sequentially starting from 0.
Content [content]	Collected GLAS data for the record type. Maximum size is 10 MB per record.
Oracle Instance [oracle_instance]	References the Oracle Instance [cmdb_ci_db_ora_instance] table.

When V2 is enabled, the navigation menu updates automatically, hiding the six V1 database modules and displaying the **GLAS V2 - Databases** module. To enable V2 data collection, see Enable the Oracle GLAS V2 data collection method.

Tuxedo middleware extension

Tuxedo license configuration and log files collected:

- *_InstallLog.log
- *ubbc*
- bdmconfig*
- comps.xml
- dmconfig*
- JSconfig.xml
- lic.txt

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- registry.xml
- tuxconfig*
- tuxwsvr.ini

Note: The asterisk represents one or more characters in the file name.

Processes collected:

- BBL
- DBBL
- GWWS
- JSH
- JSL
- TMJAVASVR
- TMMETADATA

WebLogic middleware extension

Files collected:

- beahomelist
- biee-domain.xml
- cluster.properties
- comps.xml
- config.xml
- domain-registry.xml
- opmn.xml
- portalconfig.xml
- pthome.xml
- registry.xml
- server.xml
- sessions.xml
- setupinfo.txt

WebLogic processes and files collected for the following middleware products:

- WebLogic
- OAS
- SOA
- OBI
- WebCenter
- FormsReports

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Processes

WebLogic		OAS				OBI
startWebLogic	httpd	f90	iooomgrrmi	r30rbe32	rwrqm60	nqsserver
java	apache	frmweb	javaw	r30rqm32	rwrqm	sawserver
startWLS	java	ifbld60	java	r30run32	rwrun60	coreapplication_obijh1
nodemanager	opmn	ifbld90	jre	r30sxc32	rwserver	nqscheduler
beasvc	dis3	ifcgi60	oad	r30sxu32	rwsxc60	nqsclustercontroller
beasvc64	dis4	ifcmp60	odisrv	rwbld60	rwsxu60	essbase
beasvcX64	dis51pr	ifcmp90	ofcguard	rwbuilder	vdeserver	essvr
oraclesvc	dis51ws	ifctrl60	oidrepld	rwcgi60	webcachea	essbasestudio
opmn	dis51	ifdbg60	oidmon	rwcli60	webcached	hyperion
wlsvc	dwfde61	ifrun60	oidldapd	rwcon60	webcache	
wlsvc64	f45des	ifsrv60	opmn	rwconverter		
wlsvcX64	f45run	ifwdb60	r30cli32	rwisv60		
jmc	f50dbg32	ifweb60	r30con32	rwmts60		
httpd	f50run32	ifweb90	r30isv32	rwqmu60		
jrmc	f60	imapds	r30qmu32	rwrbe60		

Middleware CI relationships

CI	Relationship/Reference	CI
cmdb_ci_ora_lms_overview	Contains::Contained by	cmdb_ci_linux_server
cmdb_ci_config_file_tracked	Contains::Contained by	cmdb_ci_linux_server

These relationships are created by a containment rule that describes the valid objects that the Linux CIs can contain.

Properties installed with the Middleware data collection application

To access the installed application properties, navigate to **All** > **System Properties** > **All Properties**. You must have the admin role to modify these properties.

System property	Default value	Description
sn_itom_oracleglas.isMaskingIp	*false*	Determines whether to mask the IP Address in Oracle middleware tracked files by replacing the value with "text_removed."
sn_itom_oracleglas.isMaskingUserName	*false*	Determines whether to mask the user name in Oracle middleware tracked files

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

System property	Default value	Description
		by replacing the value with "text_removed."
sn_itom_oracleglas.isMaskingPort	*true*	Determines whether to mask the port in Oracle middleware tracked files by replacing the value with "text_removed."

Control GLAS data collection

The following MID Server property controls the GLAS data collection:
sn_itom_oracleglas.disable_glas_data_collection

The property is set to **False** by default.

Note: In Data Collection for Oracle Global Licensing and Advisory Services versions earlier than 1.10.2, the property name is *sn_itom_pattern.disable_glas_data_collection*.

Oracle Java process discovery

ITOM Visibility can discover Java-running processes in your infrastructure, allowing you to comply with Oracle licensing agreements and prepare for GLAS Audits. You can perform the discovery using the Java Installation pattern (IP-based), Agent Client Collector with application patterns, or Agent Client Collector with file-based discovery.

In the pattern flow, the Java processes are discovered by Application Dependency Mapping (ADM). When the processes are discovered, the Java Installation pattern triggers and collects the data.

For information about privileged user access, see KB1705845.

For more information about Agent Client Collector with patterns, see Application patterns for the Agent Client Collector.

For more information about Agent Client Collector with file-based discovery, see Discover java installation data using Agent Client Collector for Visibility Content file-based discovery.

Data collected

Java Audit [ora_java_audit]

Field	Description
Name	Name of the record. For example: Java@hostname
Physical Machine Name	Name of the host CI
Machine Type	Either Server(S) or Desktop(D)
Java Publisher/Runtime Vendor	Java Vendor name
Operating System	Operating system in which the java is run
Environment Type	Production or Development or QA environment

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Java Audit [ora_java_audit] (continued)

Field	Description
Java Release/Version/Build	Java Version
Installation Path	Absolute installation path of the java binary
Installation Date (DateTime)	The date and time of the java binary file created on the filesystem
Release File (Y / N)	Indicates if a *release.md* file exists in the installation
Processor Make/Model	Processor model name
Sockets (Integer)	Number of CPU sockets
Cores (Integer)	Number of CPU cores
V6 and earlier	Indicates if you have a 6 or earlier java version
Open Source	Indicates if you have an open-source java installation
Unlock Commercial features	Indicates if you have the commercial features (like JFR) enabled
Mission Control	Indicates if the mission control plugin is enabled
AMC Agent	Indicates if the AMC agent plugin is enabled
Usage Tracker	Indicates if the Usage Tracker feature is enabled
Virtual Machine Object ID	Object ID of the Virtual machine instance
PID	Running process PID
IP Address	IP address of the host/VM
First Discovered	The date of the first discovery
Most recent discovery	The date of the most recent discovery
Discovery Source	The discovery method.
Application	Reference to the [cmdb_ci_appl] table
Software installation	Reference to the [cmdb_sam_sw_install] table
Java evidence	Reference to the [cmdb_ora_java_evidence] table
Host	Reference to the [cmdb_ci] table

Java Evidence [ora_java_evidence]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Running process	Reference to the [cmdb_running_process] table
Host CI	Reference to the [cmdb_ci_computer] table
File information	Provides a reference to the [cmdb_file_information] table. Relevant only for file-based discovery (FBD); in pattern-based discovery, this field remains empty.

Download the Oracle GLAS report in the Discovery Admin Workspace. For more information, see Download Oracle Global License Advisory Services (GLAS) data.

Oracle GLAS Data Collection

After Discovery finishes discovering the components of your database or middleware deployment, you can view the relevant statistics in the Oracle GLAS Data Collection.

After upgrading to Discovery Admin Workspace version 1.3.1 (August 2024 Store), you can navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** to view and access the Oracle GLAS Data Collection.

Related topics

Oracle Catalog discovery

Oracle pluggable database and container database discovery

Download Oracle Global License Advisory Services (GLAS) data

Download Oracle GLAS data in CSV format to get detailed information on patterns and configuration items. You can also download reports for Oracle Database, Middleware, Oracle Java, VM, and vCenter hardware information.

Before you begin

Role required: admin or glas_admin

About this task

ServiceNow maintains regular verification with Oracle GLAS to provide you with insights on licensing data for your Oracle estate. The following Oracle GLAS reporting versions are currently verified:

- Database: 25.2
- Middleware: 23.2
- Java: 24.4

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Oracle GLAS Data Collection** > **Download GLAS Data Collection**.

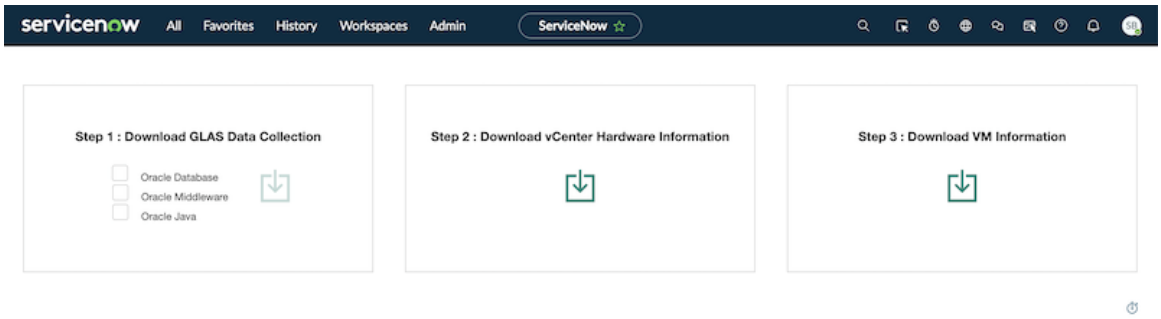

2. Download the following reports for Oracle GLAS data:
   - For downloading Oracle Database, Middleware, and Java reports, navigate to **Step 1: Download GLAS Data Collection** and select your data collection options.
   - For downloading vCenter hardware information based on your environment setup, navigate to **Step 2: Download vCenter Hardware Information**.
   - For downloading VM information, navigate to **Step 3: Download VM Information**.
3. Select the download icon ( ).
4. In the Enter email address dialog box, enter the email address where you want to receive a notification when the report is ready.
5. Select **OK**.
   After the report creation status message appears, you can check the status of the download job by selecting the link provided in the message.

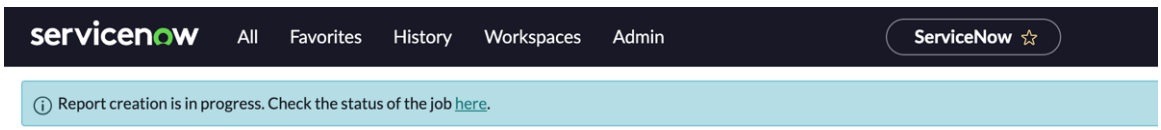

   The Oracle GLAS Report Job Log [sn_itom_oracleglas_job_log] table opens in a new window. After your report is successfully created, the status of the report changes to Complete in the Oracle GLAS Report Job Logs table.
6. Select a job log record to view its details, including the unique number, name, status, and description.
   The record also shows the date and time when the job log was created and updated.
7. On the Oracle GLAS Report Job Log form that opens, download your report by navigating to **Manage Attachments** and selecting **download**.
   - Depending on your selection in **Step 1: Download GLAS Data Collection**, the following folders are downloaded:

     Oracle Database

     - DB folder – contains the following CSV files:
       - `ORCL_V$SESSION.csv`
       - `ORCL_OPTIONS.csv`
       - `ORCL_DBA_USERS.csv`
       - `ORCL_DETAIL.csv`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- ORCL_V$LICENSE.csv
- ORCL_OVERVIEW.csv

- CPUQ folder – contains CPUQ files, vCenter Hardware, and VM information

Oracle Middleware

The following files are exported, where the file name reflects your server name and Oracle product:

- ORCL_OVERVIEW.csv
- Collection-<Server Name>_<Oracle Product>.tar.bz with the files and directories of the format FMW/u01/app/oracle/config/domains/soa_domain/config/config.xml in it

Oracle Java

- Java folder – contains Java installation information

  Note: The report that you download for Oracle Java GLAS includes only Java deployments from the Oracle publisher. Java deployments from non#Oracle publishers are excluded from the report.

- CPUQ folder – contains CPUQ files, vCenter Hardware, and VM information

- If you selected **Step 2: Download vCenter Hardware Information**, one hardware inventory CSV file is downloaded per vCenter in your environment.
- If you selected **Step 3: Download VM Information**, one VM information CSV file is downloaded per vCenter in your environment.

Note: For V1, you can also download the CSV files for Oracle Database by navigating to **Oracle GLAS Data Collection** > **Oracle GLAS Data**. From the list that opens, select the table that you want to download and then select **Download GLAS Report**.

Result

The downloaded files appear in the location that you select.

Trouble?

If the files are empty or contain an error message, check their sizes to verify that they haven't exceeded the limits set in the following properties. Navigate to **All** > **System Properties** > **All Properties** to find the following properties.

Important: Consult with ServiceNow support before changing these properties.

Property	Description
glide.discovery.file_tracking.max_file_size	The maximum file size for tracked configuration file content. Type: integer Default value: 500000
glide.discovery.file_tracking.max_files_per_ci	The maximum number of tracked configuration files per CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	Type: integer Default value: 50
glide.discovery.file_tracking.change_limit_count	Number of changes permitted on a tracked configuration file in the defined time window. Type: integer Default value: 4
glide.discovery.file_tracking.change_limit_days	Time window (in days) that limits the number of changes on a tracked configuration file. Type: integer Default value: 7

If you receive an error `Error: Skipped sensor for Horizontal Pattern because of error in payload: Payload length 5056040 exceeded the limit of 300000`, verify that the payload hasn't exceeded the system limit. Navigate to **All** > **MID Server** > **Properties** to find the following properties.

Property	Description
mid.discovery.max_payload_size	Defines the maximum overall payload size for the payload of results that come from patterns. Type: integer (bytes) Default value: 300000

If you receive an SSH timeout error, you must create and set the following properties. See "Create a MID Server property" in MID Server Properties for details.

Property	Description
mid.ssh.command_timeout_ms	Number of milliseconds an SSH command is permitted to run before timing out (default is configurable per MID Server). Type: integer
mid.sa.ssh.timeout	The SSH timeout in milliseconds. Type: integer
channel_timeout_ms	The amount of time, in milliseconds, that a MID Server waits for activity when processing an SSH command. If the MID Server doesn't

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	detect activity within the specified window, the system terminates the command. Type: integer

Enable the Oracle GLAS V2 data collection method

Enable the Oracle GLAS V2 data collection method to improve data collection performance in large-scale or high volume Oracle database environments.

Before you begin

- Verify that you have at least version 1.12.0 of Data Collection for Oracle Global Licensing and Advisory Services.
- Verify that the *sn_itom_oracleglas.disable_glas_data_collection* MID Server property is set to **false**.

Role required: glas_admin or admin

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `sn_itom_oracleglas.enable_large_env_glas_data_collection`.
3. Select the *sn_itom_oracleglas.enable_large_env_glas_data_collection* system property.
4. In the **Value** field, enter `true`.
   To switch back to V1 data collection, set the value to `false`.
5. Select **Update**.

What to do next

Run discovery again to begin data collection in V2. To verify data population, navigate to **All** > **Oracle GLAS Data Collection** > **GLAS V2 - Databases** and check the **Content** field. To download the collected data, see Download Oracle Global License Advisory Services (GLAS) data.

Oracle GoldenGate discovery

The ServiceNow Discovery and Service Mapping applications find Oracle GoldenGate version 12c components using the Oracle Golden Gate pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses the Oracle Golden Gate pattern to perform horizontal discovery to collect data into the `$report_file` file. Service Mapping performs top-down discovery on the `$report_file` file to find outgoing Oracle Golden Gate connections.

The Oracle Golden Gate pattern supports the following platforms: AIX, Linux Hewlett Packard, and Solaris.

You can use this pattern on the ServiceNow AI Platform using Kingston, London, or Madrid.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Oracle Golden Gate data model

The Oracle Golden Gate pattern introduces the following CI classes that extend an existing CMDB class.

CI classes introduced by this pattern

CI class	Extends from
Oracle Golden Gate [cmdb_ci_appl_oracle_golden_gate]	Application [cmdb_ci_appl]
Oracle Golden Gate Replicat Process [cmdb_ci_appl_ora_gg_replicat]	Application [cmdb_ci_appl]
Oracle Golden Gate Extract Process [cmdb_ci_appl_ora_gg_extract]	Application [cmdb_ci_appl]

Prerequisites

Credentials

Configure the SSH credentials, for the operating system of the server that hosts the Oracle Golden Gate Server.

User access

Give the UNIX OS user permissions to read the `$report_file`. The `$report_file` is the report file that is extracted from the Oracle Golden Gate `manager/replicat/extract` process, with the extension `.rpt`. For example:

Configuration file:

```
/base/ggs_beta/oracle/bt01pims/ggs/dirprm/mgr.prm
 REPORTFILE
```

Report file:

```
/base/ggs_beta/oracle/bt01pims/ggs/dirrpt/MGR.rpt
 PROCESSID MGR
```

Permissions to run commands

The OS user must have permissions to run the relevant OS commands for each supported platform.

Give the UNIX OS user permissions to run the following commands against the Oracle Golden Gate Server:

Command	Mandatory/Optional	Description
• Starting with Discovery and Service Mapping Patterns version 1.31.0: `$sudo + " ls " +`	Mandatory	List all report files under the Oracle Golden Gate report base directory.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Mandatory/Optional	Description
`$rpt_base_dir + " | grep '\.rpt $' \| sort "` • Before Discovery and Service Mapping Patterns version 1.31.0: `$sudo + " ls -d " + $rpt_base_dir + "*.rpt | sort "` †		
• Starting with Discovery and Service Mapping Patterns version 1.31.0: `$sudo + " ls " + $prm_base_dir + " | grep '\.prm $' \| sort "` • Before Discovery and Service Mapping Patterns version 1.31.0: `$sudo + " ls -d " + $prm_base_dir + "*.prm| sort "` †	Mandatory	List all parameter files under the Oracle Golden Gate parameter base directory.

† If this command doesn't return data, see the update set in KB3069148.

Retrieve data

Retrieve data by parsing:

- The Oracle Golden Gate manager report file (report_file) to retrieve related configuration item (CI) names and counter information.
- Variables in the process command line to retrieve the install folder, manager process name, configuration file, and parameters file.
- Variables in the Oracle Golden Gate installation folder to retrieve a list of parameter and configuration files.

EVAL closure functions

Use the following EVAL closure functions to remove duplicate entries from the extract process and the replicat process.

- var tableWithoutDuplicates = '';tableWithoutDuplicates = DuplicateRemover.removeDuplicates(${extracts}, ["name"]);CTX.setAttribute("extracts", tableWithoutDuplicates);
- var tableWithoutDuplicates = '';tableWithoutDuplicates = DuplicateRemover.removeDuplicates(${replicats}, ["name"]);CTX.setAttribute("replicats", tableWithoutDuplicates);

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the following EVAL closure functions to count the number of extract processes and replicat processes.

```
return
 ${cmdb_ci_appl_ora_gg_replicat[*].config_file}.size()
return
 ${cmdb_ci_appl_ora_gg_extract[*].config_file}.size()
```

Use the following EVAL closure function to return the privileged command.

```
return
 ${ctx}.getDiscoveryProvider(com.snc.sw.dto.ProviderTyp
e.SSH).getPrivilegedCommand();
```

Use the following EVAL closure function to return the discovery type.

```
ctx.getWork().getDiscoveryType();
```

Use the following EVAL closure function to extract the version from the installation directory path if it is empty.

```
inst_dir =
 ${install_directory}if(inst_dir.isEmpty()){return
 ${version}}if(inst_dir.startsWith('/')){return
 inst_dir.replaceAll('/.*/','')}if(!inst_dir.startsWith
('/') && !inst_dir.isEmpty()){return
 inst_dir.replaceAll('.*\\\\','')}
```

Applicative credentials

1. Navigate to **Discovery** > **Credentials**.
2. Click **New**.
3. Click **Applicative Credentials**.
4. On the form, fill in the fields.

Applicative Credentials form

Field	Description
Name	Credential name. An example is oracle_db_user.
Active	Option for enabling this credential for discovery. Select this check box to enable discovery.
Applies to	Credentials that you may or may not want to apply to **All MID servers** in your network, or to one or more **Specific MID servers**. Select **Specific MID servers**.
MID Servers	MID Servers that the credentials apply to. Select the required MID Server. This field appears when **Specific MID servers** is selected from the **Applies to** field.
Order	Order in which the platform tries this credential as it attempts to log on to devices. A smaller number indicates that the credential appears higher in the list. Establish the credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number, or none, the instance tries the credentials in a random order. The default value is **100**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
User name	Name of the user of this applicative credential. An example is oracle_db_user. You can use any user for the credential for this pattern, because the information is extracted from a local cache.
Password	Not required. You can leave this field blank or enter any value.
CI type	CI type for which this credential is used: Storage Server [cmdb_ci_storage_server]. **Note:** ServiceNow applications refer to devices and applications that comprise a service instance as configuration items (CIs).

5. Click **Submit**.

Entry point

For top-down discovery, use the Oracle Golden Gate IP address and specify the MID Server.

Data collected by Discovery during horizontal discovery

Discovery uses the Oracle Golden Gate pattern to collect the data described in the following tables.

Oracle Golden Gate [cmdb_ci_appl_oracle_golden_gate]

Field	Description
Name [name]	Name of the CI in the CMDB (<process name>@<source db>).
Version [version]	Version of the Oracle Golden Gate installation.
Source DB SID [source_db]	Manager process source database system identifier (SID).
Counter for replicat processes [count_replicat]	Counter of replicat processes that are managed by the Oracle Golden Gate manager instance.
Counter for extract processes [count_extract]	Counter of extract processes that are managed by the Oracle Golden Gate manager instance.
Configuration file [config_file]	Parameter file of the Oracle Golden Gate process. Specify the path to the configuration file and the file name, <name>.prm
Replicat File [report_file]	Report file of the Oracle Golden Gate process. Specify the path to the report file and the file name, <name>.rpt
Type [type]	Type of the Oracle Golden Gate installation. Specify the Oracle Golden Gate for Oracle technologies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Golden Gate [cmdb_ci_appl_oracle_golden_gate] (continued)

Field	Description
Installation directory [install_directory]	Folder that contains all the Oracle Golden Gate setup, configuration, libraries, and executable files.

Oracle Golden Gate Replicat Process [cmdb_ci_appl_ora_gg_replicat]

Field	Description
Name [name]	Name of the CI in the CMDB (<process name>@<source db>).
Replicat File [report_file]	Replicat process report file. Specify the path to the report file and the file name <name>.rpt
Configuration file [config_file]	Parameter file of the replicat process. Specify the path to the configuration file and the file name <name>.prm
Installation directory [install_directory]	Folder that contains all the Oracle Golden Gate setup, configuration, libraries, and executable files.
Version [version]	Version of the Oracle Golden Gate installation.
Source DB SID [source_db]	Manager process source database SID.
Operational status [operational_status]	Operational status of the CI. Select **Operational**.

Oracle Golden Gate Extract Process [cmdb_ci_appl_ora_gg_extract]

Field	Description
Name [name]	Name of the CI in the CMDB (<process name>@<source db>).
Replicat File [report_file]	Report file of the extract process. Specify the path to the report file and the file name <name>.rpt
Configuration file [config_file]	Parameter file of the extract process.
Version [version]	Version of the Oracle Golden Gate installation.
Installation directory [install_directory]	Folder that contains all the Oracle Golden Gate setup, configuration, libraries, and executable files.
Source DB SID [source_db]	Manager process source database SID.
Operational status [operational_status]	Operational status of the CI. Select **Operational**.

The Dependency Views map shows discovered load balancer CIs and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

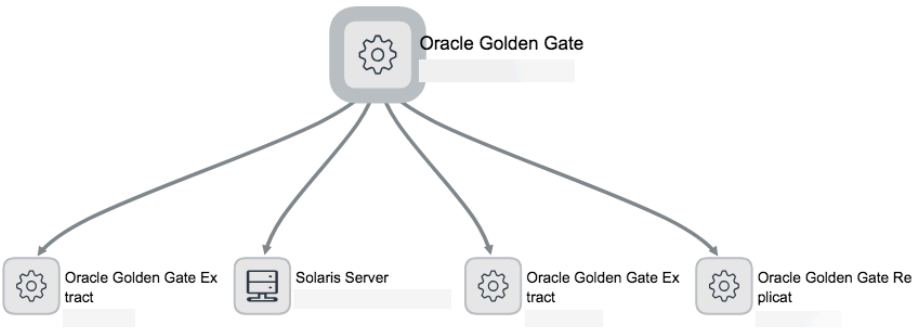

CI relationships

The Oracle Golden Gate pattern creates the following relationships to support Oracle Golden Gate discovery.

CI relationships

CI	Relationship	CI
Oracle Golden Gate [cmdb_ci_appl_oracle_golden_gate]	Managed by::Manages	Oracle Golden Gate Replicat Process [cmdb_ci_appl_ora_gg_replicat]
Oracle Golden Gate [cmdb_ci_appl_oracle_golden_gate]	Managed by::Manages	Oracle Golden Gate Extract Process [cmdb_ci_appl_ora_gg_extract]
Oracle Golden Gate [cmdb_ci_appl_oracle_golden_gate]	Runs on::Runs	Hardware [cmdb_ci_hardware]
Oracle Golden Gate Replicat Process [cmdb_ci_appl_ora_gg_replicat]	Runs on::Runs	Hardware [cmdb_ci_hardware]
Oracle Golden Gate Extract Process [cmdb_ci_appl_ora_gg_extract]	Runs on::Runs	Hardware [cmdb_ci_hardware]

Data collected by Service Mapping during top-down discovery

To discover the Oracle Golden Gate process, use the TCP entry point with the proper host and port of the Oracle Golden Gate process.

Field	Description
Name [name]	Name of the CI in the CMDB (<process name>@<source db>).
Version [version]	Version of the Oracle Golden Gate installation.
Installation directory [install_directory]	Folder that contains all the Oracle Golden Gate setup, configuration, libraries, and executable files.
Configuration file [config_file]	Parameter file of the Oracle Golden Gate process.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Replicat File [report_file]	Report file of the Oracle Golden Gate replicat process.
Source DB SID [source_db]	Manager process source database SID.
Counter for extract processes [count_extract]	Counter of extract processes that are managed by the Oracle Golden Gate manager instance.
Counter for replicat processes [count_replicat]	Counter of replicat processes that are managed by the Oracle Golden Gate manager instance.

Oracle HTTP Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - HTTP Server pattern to find Oracle HTTP Server instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the application executable

Verify you have permission for the following command: `powershell -Command "Select-String -Path '{executablePath}' -Pattern 'Server version' -SimpleMatch"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - HTTP Server pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle HTTP Server@{hostname}**.
Product [product]	Product name. The value is set to **HTTP Server**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application [cmdb_ci_appl] (continued)

Field	Description
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - HTTP Server pattern creates the following relationships and references to support Oracle HTTP Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle Internet Directory Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Internet Directory Server pattern to find Oracle Internet Directory Server instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the application executable

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify you have permission for the following command: `powershell -Command "(Get-Item '{executablePath}').VersionInfo"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Internet Directory Server pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle Internet Directory Server@{hostname}**.
Product [product]	Product name. The value is set to **Internet Directory Server**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Internet Directory Server pattern creates the following relationships and references to support Oracle Internet Directory Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Listener HD discovery

ServiceNow Discovery uses the Oracle Listener HD pattern to find and map Oracle real application cluster components. To discover these resources, install the CMDB CI Class Models and the Discovery and Service Mapping Patterns from the ServiceNow Store.

Oracle RAC enables you to cluster an Oracle database. Oracle RAC uses Oracle Clusterware for the infrastructure to bind multiple servers so they operate as a single system. Oracle Clusterware enables you to create a clustered pool of storage that can be used by any combination of non-cluster and Oracle RAC databases. Discovery uses Oracle Listener HD patterns to discover information about the Oracle Database Listener and the components in the Oracle RAC infrastructure, when such a setup is configured.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Ensure that the following plugins are installed:
  - CMDB CI Class Models plugin (Version 1.29.0 or later).
  - Discovery and Service Mapping Patterns plugin (Version 1.0.82 or later).
- Ensure that the Discovery User is in the Oracle DBA group or has the correct Permissions.
- Ensure that the MID Server has access to the target device.
- Ensure that the Oracle Clusterware/Grid infrastructure is installed and the Oracle RAC Database is configured. If there is an Oracle Clusterware setup.
- Ensure that the Oracle Listener process is running on a UNIX OS host.
- Use Oracle Clusterware minimum version 12.2.X.X (Oracle Database 12c Release 2) or higher, If there is an Oracle Clusterware setup.

  Note: This pattern is supported from the Paris release on. More UNIX OS types (AIX, Solaris, HP-UX) are supported from version 1.0.87 of the Discovery and Service Mapping Patterns plugins.

  The following commands are executed during the pattern execution. Executing the commands enables the pattern to populate allthe relevant CMDB classes with the necessary information:

  - **export TNS_ADMIN=<oracle_home_directory>/network/admin**: Sets the ORACLE_HOME to the TNS_ADMIN environment variable.
  - **export ORACLE_BASE==<oracle_home_directory>**: Sets the ORACLE_HOME to the ORACLE_BASE environment variable.
  - **lsnrctl status LISTENER_NAME**: Checks the Database Listener status.
  - **ps -ef | grep 'd.bin'**: Searches for the GRID_HOME path.
  - **crsctl query cluster site -all**: Lists the nodes and disks that the sites contain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **lsnrctl services LISTENER_NAME**: Lists the services that the listener identifies.
- **crsctl stat res -f -w '((TYPE = ora.database.type))'**: Lists the database name, associated node, type, and active status.

> **Note:** If no output is obtained from the **lsnrctl status LISTENER_NAME**, then the **Oracle Listener HD** pattern is terminated.

Limitations

The Oracle Clusterware setup is only supported for the UNIX Operating Systems. The base system **SSH Command-Oracle - Listener Details classifier probe** is inactive (Active is set to False), when installing the CMDB CI Class Models version 1.29.0 or higher.

Classifier Probes- Oracle Listener Details

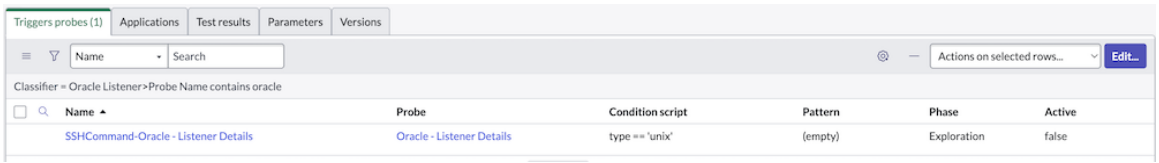

The **Oracle Listener process classification** triggers the Oracle Listener HD pattern.

Classifier Probes- Horizontal Pattern

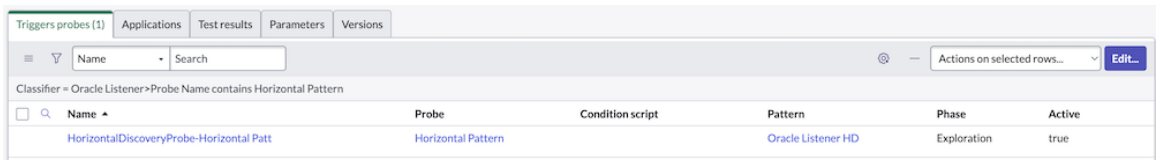

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle Listener HD pattern.

Oracle Database Listener [cmdb_ci_db_ora_listener]

Field	Description
Name [name]	The name of the Oracle Database Listener.
Oracle Home [oracle_home]	The Oracle home directory where software is installed.
Configuration [configuration]	The listener.ora network configuration file.
Scan Listener Name [scan_listener_name]	The name of the SCAN LISTENER.
Version [version]	The version of the Oracle product.
TCP port(s) [tcp_port]	The listening port for Oracle client connections.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Instance [cmdb_ci_db_ora_instance]

Field	Description
SID [sid]	The specific name that specifically identifies the instance.
Name [name]	The name of the Oracle instance combination of SID@HOST_NAME.

Oracle Database [cmdb_ci_oracle_database]

Field	Description
Name [name]	The Oracle database unique name.
Type [type]	The Oracle database type.
Operational status [operational_status]	The active status of the Oracle database.
Node [node]	The host name of the node hosting the Oracle database.

Oracle Cluster [cmdb_ci_oracle_cluster]

Field	Description
Name [name]	The name of the Oracle database cluster.
Cluster ID [cluster_id]	The GUID of the Oracle database cluster.
Cluster Status [cluster_status]	The active state of the Oracle database cluster.
IP Address [ip_address]	The IP addresses of the nodes participating in the cluster.

Oracle Cluster Node [cmdb_ci_oracle_cluster_node]

Field	Description
Name [name]	The hostname of the cluster node combination of HOSTNAME@CLUSTER_NAME.
IP Address [ip_address]	The IP address of the cluster node, if resolvable by DNS.
Server [server]	The hostname of the host hosting the cluster node.
Cluster [cluster]	The name of the cluster to which the cluster node is associated.

Oracle DB Service [cmdb_ci_db_ora_service]

Field	Description
Name [name]	The name of the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle DB Service [cmdb_ci_db_ora_service] (continued)

Field	Description
TCP port(s) [tcp_port]	The port of the service.
Listener name [listener_name]	The name of the Listener managing the service.
Version [version]	The version of the service.

Linux Server [cmdb_ci_linux_server]

Field	Description
Name [name]	The hostname of the Linux server that is hosting the Oracle Listener process.

AIX Server [cmdb_ci_aix_server]

Field	Description
Name [name]	The hostname of the AIX server that is hosting the Oracle Listener process.

HPUX Server [cmdb_ci_hpux_server]

Field	Description
Name [name]	The hostname of the HP-UX server that is hosting the Oracle Listener process.

Solaris Server [cmdb_ci_solaris_server]

Field	Description
Name [name]	The hostname of the Solaris server that is hosting the Oracle Listener process.

Dependency Views map

The following dependency views display the relationships between Oracle Listener resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Database Listener View Map

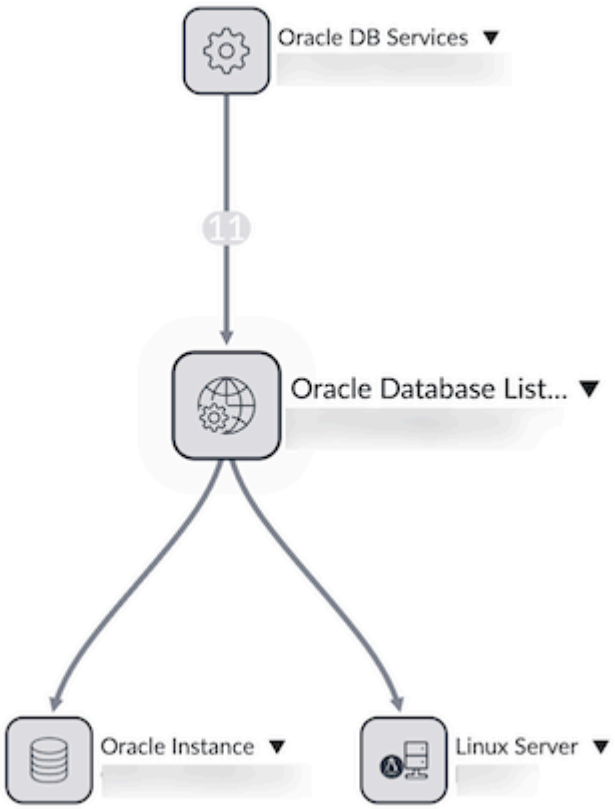

Oracle Instance View Map

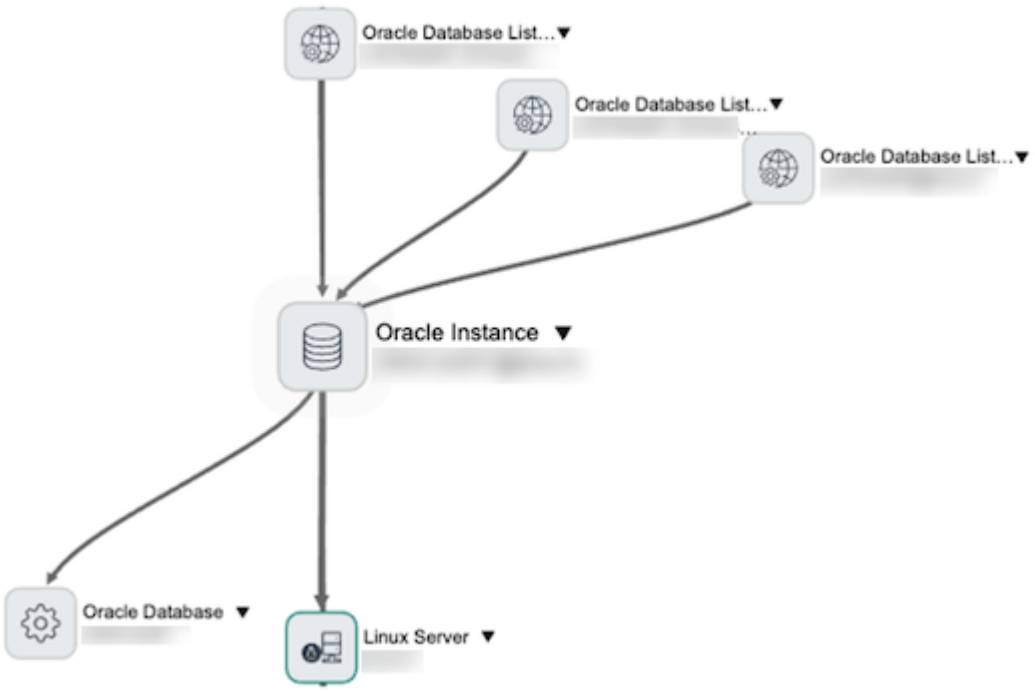

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Database View Map

Oracle Cluster View Map

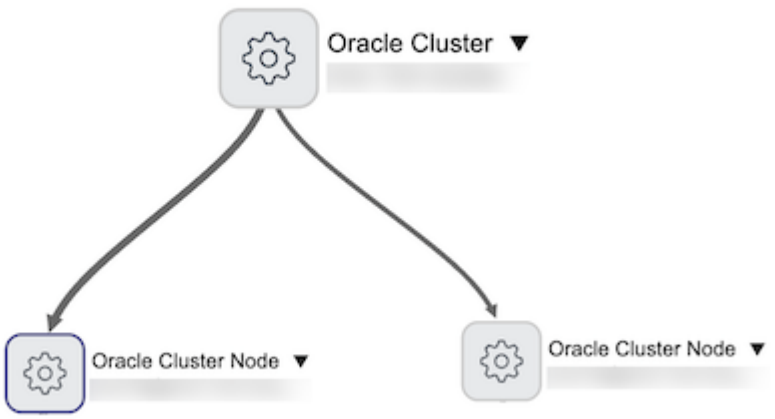

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Cluster Node View Map

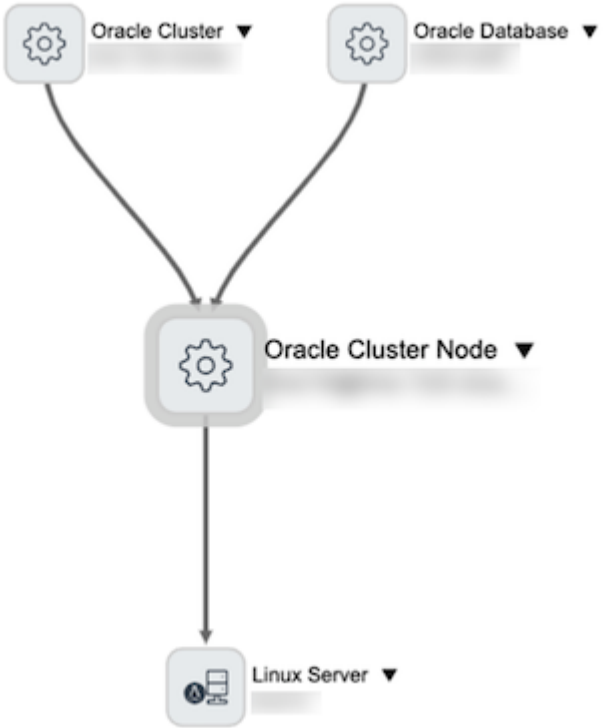

Linux Server View Map

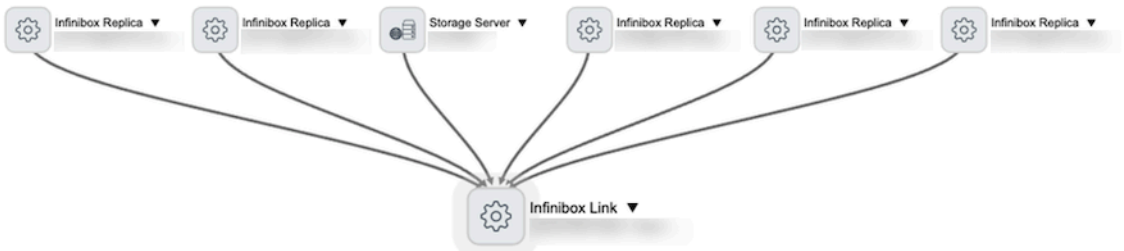

CI relationships

The Listener HD pattern creates the following relationships to support Oracle Listener discovery.

CI	Relationship	CI
Oracle Cluster [cmdb_ci_oracle_cluster]	Cluster of::Cluster	Oracle Cluster Node [cmdb_ci_oracle_cluster_node]
Oracle Cluster Node [cmdb_ci_oracle_cluster_node]	Hosted on::Hosts	Linux Server [cmdb_ci_linux_server]
Oracle Instance [cmdb_ci_db_ora_instance]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]
Oracle Database Listener [cmdb_ci_db_ora_listener]	Depends on::Used by	Oracle Instance [cmdb_ci_db_ora_instance]
Oracle Instance [cmdb_ci_db_ora_instance]	Uses::Used by	Oracle Database [cmdb_ci_db_oracle_database]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Oracle Database [cmdb_ci_db_oracle_database]	Hosted on::Hosts	Oracle Cluster Node [cmdb_ci_oracle_cluster_node]
Oracle Cluster Node [cmdb_ci_oracle_cluster_node]	Hosted on::Hosts	AIX Server [cmdb_ci_aix_server]
Oracle Instance [cmdb_ci_db_ora_instance]	Runs on::Runs	AIX Server [cmdb_ci_aix_server]
Oracle Cluster Node [cmdb_ci_oracle_cluster_node]	Hosted on::Hosts	Solaris Server [cmdb_ci_solaris_server]
Oracle Instance [cmdb_ci_db_ora_instance]	Runs on::Runs	Solaris Server [cmdb_ci_solaris_server]
Oracle Cluster Node [cmdb_ci_oracle_cluster_node]	Hosted on::Hosts	HPUX Server [cmdb_ci_hpux_server]
Oracle Instance [cmdb_ci_db_ora_instance]	Runs on::Runs	HPUX Server [cmdb_ci_hpux_server]
Oracle DB Service [cmdb_ci_db_ora_service]	Managed by::Manages	Oracle Database Listener [cmdb_ci_db_ora_listener]
Oracle Instance [cmdb_ci_db_ora_instance]	Uses::Used by	Oracle DB Service [cmdb_ci_db_ora_service]

Oracle Node Manager pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Node Manager pattern to find Oracle Node Manager instances running on UNIX servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create SSH credentials

For more information, see SSH credentials.

Verify permission to run the version command

Verify you have permission for the following command: `{process.environmentVariables.EMHOME.value}/OPatch/opatch version`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Node Manager pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle Node Manager@{hostname}**.
Product [product]	Product name. The value is set to **Node Manager**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Node Manager pattern creates the following relationships and references to support Oracle Node Manager discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle pluggable database and container database discovery

The ServiceNow Discovery application can discover Oracle pluggable databases (PDBs) and container databases (CDBs). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

A PDB is a portable collection of schemas, schema objects, and non-schema objects that appear to an Oracle Net client as a separate database. A collection of PDBs is called a container database (CDB).

For more information about PDBs and CDBs, see What are CDBs, PDBs and Containers?

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Set the user credentials and grant permissions:

- Configure Applicative credentials or Oracle Wallet authentication
- For Windows, configure Windows credentials
- For Unix, configure SSH credentials
- For the Oracle database, configure Oracle database discovery
- Provide a user with permissions to run queries on the following tables:
  - V$DATAFILE
  - V$TEMPFILE
  - V$LOG
  - DBA_FREE_SPACE
  - V$DATABASE
  - V$PDBS
  - V$PARAMETER

Note: Set applicative credentials at the CDB level to discover relationships for PDBs and the CDB size only. Alternatively, set applicative credentials at the CDB/PDB level to discover relationships for PDBs and the CDB/PDB size. If you do not set applicative credentials at the CDB level, no PDB attributes are discovered.

Data collected during Discovery

Discovery collects the following data:

Table and field	Description
cmdb_ci_db_ora_pdb_instance: Relation Type : Managed By	
sid	Name of the PDB instance
name	Name of the CI in CMDB. This is a combination of sid and the hostname.
version	Version of the PDB as taken from the CDB instance
edition	Edition of the PDB as taken from the CDB instance
install_directory	Install directory of the PDB as taken from the CDB instance
oracle_home	Oracle home of the PDB as taken from the CDB instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
cdb_name	Name of the cdb sid that contains the PDB. This refers to the Oracle database instance configuration item (CI)
cmdb_ci_db_instance_size: Relation Type : Contained By	
sid	Name of the Oracle database CDB/PDB instance
name	Name of the CI in CMDB. This is a combination of the sid and the hostname.
db_size	Total size allocated to the database instance
used_size	Size used by the database instance
free_size	Size of the free space on the database instance

The following is a dependency view of the relationship between PDBs, CDBs, and containers.

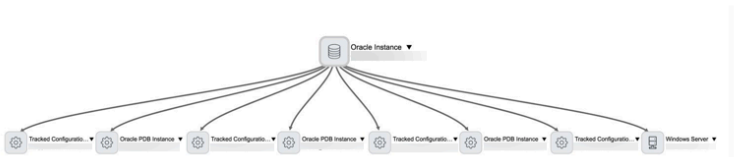

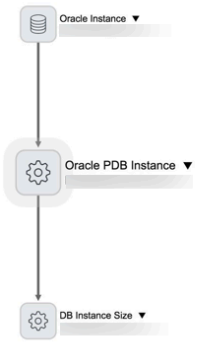

Related topics

Oracle Catalog discovery

Oracle GLAS data collection using patterns

Oracle Reports pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Reports on Windows pattern to find Oracle Reports instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify read access to the application executable

Verify you have permission for the following command:
`powershell -Command "(Get-Item '{process.environmentVariables.MW_HOME.value}\bin\rwrun.exe').VersionInfo.ProductVersion"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Reports on Windows pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle Reports@{hostname}**.
Product [product]	Product name. The value is set to **Reports**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Reports on Windows pattern creates the following relationships and references to support Oracle Reports discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Siebel CRM Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Siebel CRM Server on Windows pattern to find Siebel CRM Server instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the application executable

Verify you have permission for the following command: `powershell -Command "Get-Item '{executablePath}' | Select-Object -ExpandProperty VersionInfo | Format-List"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Siebel CRM Server on Windows pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle Siebel CRM Server@{hostname}**.
Product [product]	Product name. The value is set to **Siebel CRM Server**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships and references

The Oracle - Siebel CRM Server on Windows pattern creates the following relationships and references to support Oracle Siebel CRM Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle Siebel Gateway Name Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Siebel Gateway Name Server on Windows pattern to find Siebel Gateway Name Server instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the application executable

Verify you have permission for the following command: `powershell -Command "Get-Item '{executablePath}' | Select-Object -ExpandProperty VersionInfo | Format-List"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Siebel Gateway Name Server on Windows pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle Siebel Gateway Name Server@{hostname}**.
Product [product]	Product name. The value is set to **Siebel Gateway Name Server**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Siebel Gateway Name Server on Windows pattern creates the following relationships and references to support Oracle Siebel Gateway Name Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle Solaris LDOM discovery

Discovery uses the Solaris Logical Domain (LDOM) infrastructure pattern and Solaris LDOM shared library pattern to find all LDOM data. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses these patterns to run horizontal discovery. You can use the patterns on the ServiceNow AI Platform using the Brazil release or later.

The Solaris LDOM infrastructure pattern is triggered automatically on Solaris Server discovery. The pattern verifies if there is an LDOM controller, and stops running if it doesn't find a controller.

The Solaris LDOM shared library pattern is run as an Extension Section to the Solaris Server pattern. The pattern modifies the Serial Number (SN) on LDOM machines so that it matches the SN populated by the Solaris LDOM infrastructure pattern. For further information about Solaris Server, see Solaris discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

LDOM commands when running the Solaris LDOM infrastructure pattern

Ensure the relevant users have permissions to execute the following used commands for the Solaris LDOM infrastructure pattern:

- *"sudo /usr/sbin/virtinfo -a"*
- *"sudo /usr/sbin/ldm -V"*
- *"hostname"*
- *"svcs ldmd"*
- *"sudo /usr/sbin/ldm list-rsrc-group -a"*
- *"sudo /usr/sbin/sneep -T | grep ChassisSerialNumber 2> /dev/null"*
- *"/usr/sbin/prtdiag -v | awk '/Chassis Serial/{getline; getline; print}'"*
- *"sudo /usr/sbin/ldm list"*

Note: Some of these commands are used with conditions, and will not be executed on each discovery. For example, the main way to get the SN for servers is by using the *'sneep'* command. However, this is not available by default on Solaris servers, so *'prtdiag'* is used as an alternative method to get the SN.

Important LDOM commands when running the Solaris LDOM shared library pattern

Ensure the relevant users have permissions to execute the following used commands for the Solaris LDOM shared library pattern:

- *"sudo /usr/sbin/virtinfo -a"*
- *"sudo /usr/sbin/sneep -T | grep ChassisSerialNumber 2> /dev/null"*
- *"hostname"*

Ensure the relevant users have the permissions to execute the LDOM *virtinfo* command. This command ensures the proper LDOM role can be established.

Ensure the relevant users have the permissions to execute the LDOM *"sudo /usr/sbin/ldm list"* command for listing Guest VMs on the LDOM controller. This is required in order to set relations between the LDOM controller and the LDOM Guest VMs.

One of the *'sneep'* or *'prtdiag'* Serial Number commands has to return a correct response in order to correctly set the SN.

Data collected by Discovery during horizontal discovery

The data discovered by both patterns includes the following tables and fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Solaris Server [cmdb_ci_solaris_server]	
ldom_version	The LDOM version installed on the LDOM controller.
ldom_role	The LDOM role; "controller" or "guest".
ldom_state	The state of the LDOM implementation on the LDOM controller.
serial_number	The serial number of the machines.
virtual	Whether or not a virtual machine (for Guest VMs) is used for the LDOM controller.
ram	The memory available on the LDOM controller.
cpu_core_count	The number of CPU cores on the LDOM controller.

CI relationships

CI	Relationship	CI
LDOM controller	Hosted on::Hosts	LDOM Guest VM

Oracle TNS Listener pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - TNS Listener pattern to find Transparent Network Substrate (TNS) Listener instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the application executable

Verify you have permission for the following command: `powershell -Command "(Get-Item '{executablePath}').VersionInfo"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - TNS Listener pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **TNS Listener@{hostname}**.
Product [product]	Product name. The value is set to **TNS Listener**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - TNS Listener pattern creates the following relationships and references to support Oracle TNS Listener discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Oracle Unified Directory pattern-based discovery

Discovery and Service Mapping Patterns uses the Oracle - Unified Directory pattern to find Unified Directory instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create Windows credentials

For more information, see Windows credentials.

Verify permission to run the version command

Verify you have permission for the following command: `{process.currentDir}start-ds.bat --version`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle - Unified Directory pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Oracle Unified Directory@{hostname}**.
Product [product]	Product name. The value is set to **Unified Directory**.
Publisher [publisher]	Publisher of the application. The value is set to **Oracle**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Oracle - Unified Directory pattern creates the following relationships and references to support Oracle Unified Directory discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Palo Alto Networks firewall discovery

The ServiceNow Discovery application uses the Next-Generation Palo Alto Firewall pattern to find Palo Alto Networks firewalls. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The discovery pattern uses a set of SNMP calls to find the Palo Alto Networks firewalls. Discovery uses the pattern to run horizontal discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Palo Alto Networks firewall data model

The Next-Generation Palo Alto Firewall pattern introduces the following CI class that extends an existing CMDB class.

CI class introduced by this pattern

CI class	Extends from
Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]	Firewall Device [cmdb_ci_firewall_device]

Prerequisites

- Ensure that your network firewall device has SNMP access.
- On the ServiceNow instance, configure SNMP credentials. For more information, see SNMP credentials.
- Add the SNMP system OID record for the Palo Alto Networks device to the ServiceNow instance. Update the following:
  - Classifier: Palo Alto Firewall
  - Class: Palo Alto Firewall Device
- Deploy the pattern as follows:
  1. Download and install Firewall extension classes from the ServiceNow Store. The app adds the new CMDB classes required for network firewall discovery.
  2. Download and install the discovery pattern from the ServiceNow Store.
  3. Sync the pattern with the appropriate MID Server.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Next-Generation Palo Alto Firewall Pattern.

Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]

Field	Description
IP Address [ip_address]	IP address of the Palo Alto device.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto] (continued)

Field	Description
Serial number [serial_number]	Serial number of the Palo Alto device.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) of the Palo Alto device.
Manufacturer [manufacturer]	Palo Alto device manufacturer.
Model ID [model_id]	Model ID of the Palo Alto device.
Operational status [operational_status]	Indicates whether the Palo Alto device is in active state.
Hardware OS [hardware_os]	OS running on the hardware.
Hardware OS Version [hardware_os_version]	OS version running on the hardware.
Description [short_description]	Short description of the Palo Alto device.
Firmware version [firmware_version]	Palo Alto device firmware version.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
IP Address [ip_address]	IP address of the network adapter.
Alias [alias]	The user-assigned name for the network adapter.
Netmask [netmask]	Netmask of the network adapter.
MAC address [mac_address]	MAC address of the network adapter.
Name [name]	Name of the network adapter.
Configuration Item [cmdb_ci]	References the Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto] table.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address of the Palo Alto firewall.
Netmask [netmask]	Netmask of the Palo Alto firewall.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Domain Name System (DNS) name of the Palo Alto firewall device.
IP Address [ip_address]	Host IP address.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Next-Generation Palo Alto Firewall pattern creates the following relationships and references to support Palo Alto Networks firewall discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]	Uses::Used by	Router Interface [dscy_router_interface]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	IP Address [cmdb_ci_ip_address]

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration item [configuration_item]	Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]
Router Interface [dscy_router_interface]	Configuration Item [cmdb_ci]	Palo Alto Firewall Device [cmdb_ci_firewall_device_palo_alto]
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]

Pivotal Cloud Foundry discovery

The ServiceNow Discovery application finds Pivotal Cloud Foundry (PCF) version 3 components using the Pivotal Cloud Foundry pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

You can use this pattern on the ServiceNow AI Platform starting Kingston version.

The discovery of PCF V3 components is supported when the Discovery and Service Mapping Patterns plugin Version 1.0.99, March 2023 is installed. The Pattern continues to support the discovery of PCF V2 components.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Pivotal Cloud Foundry credentials

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create the Pivotal Cloud Foundry credentials.

1. Navigate to **Discovery** > **Credentials**.
2. Select **New**.
3. Select **CloudFoundry Credentials**.
4. On the form, fill in the fields:

Field	Description
Name	Credential name. As an example, `pcf_user credentials`
Active	Option for enabling this credential for discovery. Select this check box to enable discovery.
Applies to	Credentials that you may or may not want to apply to **All MID servers** in your network, or to one or more **Specific MID servers**. Select **Specific MID servers**.
Order	Order in which the platform tries this credential as it tries to log in to devices. A smaller number indicates that the credential seems higher in the list. Establish the credential order when using large numbers of credentials or when security locks out users after three failed logins tries. If all the credentials have the same order number, or none, the instance tries the credentials in a random order. The default value is 100.
User name	Name of the user of these credentials. Avoid leading or trailing spaces in user names. A warning seems if the platform detects leading or trailing spaces in the user name. An example is `pcf_user`. Use the `client_id` that is used to generate the JWT token.
Password	Password for the user of these credentials. Use the `client_secret` that is used to generate the JWT token.
Credential alias	Alternate name for these credentials. As an example, `cf` If this field is closed, select the unlock icon to unlock the field.

5. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

JSON web token credentials

To verify the JSON web token (JWT), perform the following token validation steps.

The format of the JWT token is: https://docs.cloudfoundry.org/api/uaa/version/4.27.0/index.html#jwt-bearer-token-grant

1. Run the following cURL script against the User Account and Authentication (UAA) server API and verify the results:

```
curl -s -k https://ip_address:port/oauth/token
-X; POST -H "Content-Type: application/x-www-
form-urlencoded" -H "Accept: application/json" -d
"client_id=<INSERT_USER_HERE>&client_secret=<INSERT_PASS_HERE>&grant_ty
```

   If this command doesn't generate a token, verify that the credentials in the ServiceNow repository are correct, otherwise go to step 2.

2. Run the following cURL script against Pivotal Cloud Foundry API:

```
curl https://ip_address:port/v3/spaces -k -v -H
"Accept: application/json" -H "Authorization:
bearer<INSERT TOKEN HERE>"
```

   If this command doesn't retrieve the PCF spaces information, verify the credentials on the ServiceNow credentials repository.

Credentials for V3 API elements

Provide a read-only user with permission to use the following API elements:

- –/v3/organizations
- –/v3/spaces
- –/v3/domains
- –/v3/routes
- –/v3/organization_quotas
- –/v3/space_quotas
- –/v3/apps
- -/v3/service_offerings
- -/v3/service_plans
- -/v3/service_instances
- -/v3/droplets

Note: The Pivotal Cloud Foundry Pattern also supports the PCF V2 APIs.

Create a serverless Discovery schedule

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Select **New**.
3. On the form, fill in the fields:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule form

Field	Description
Name	Enter a unique and descriptive name for this Discovery schedule, for example, `pcf`
Discover	For the Discover type, select **Serverlesss**
MID Server	MID Servers that the credentials apply to. Specify the required MID Server, for example, `Discovery_Server`
Active	Option for enabling this credential for discovery. Select this check box to enable discovery.
Credential alias	Alternate name for these credentials. Use the alias related with the Pivotal Cloud Foundry credentials that you created earlier.

4. Right-click the header of the Discovery Schedule form and select **Save**.

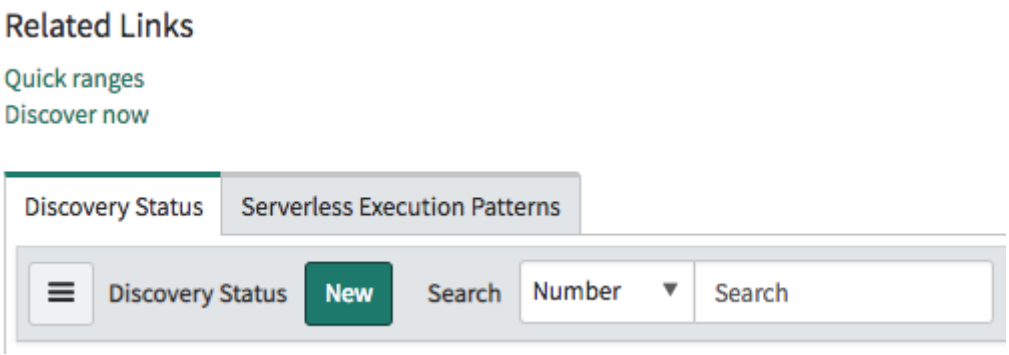

5. Select the Serverless Execution Patterns tab.
6. Select **New**.
7. On the form, fill in the fields:

Serverless Execution Pattern form

Field	Description
Name	Enter a unique and descriptive name for this Serverless Execution Pattern, for example, `pcf1`
Pattern	Select **CloudFoundry**.
Proxy Host	Fully qualified domain name of the machine on which you're installing the proxy server. Specify **Global**.
Active	Check box to enable this schedule for discovery. Select this check box.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Right-click the header of the Serverless Execution Patterns form and select **Save**.
9. Select the Serverless Execution Patterns tab
10. Create and define the serverless execution pattern as described in the product documentation.

    Configure the parameters required by the Pivotal Cloud Foundry Pattern as follows:

Discovery Pattern Launcher parameters

Parameter	Description
Organization	Name of the organization that is required to be discovered, or all where the input is "*".
CredentialAlias	Alternate name for these credentials. Use the alias related with the Pivotal Cloud Foundry credentials that you created earlier.
URL_CF	API URL of the Cloud Controller that is used to retrieve the information from Pivotal Cloud Foundry.
URL_CF_AUTH	UAA server APIthat is used to generate the token. Enter the value in the format: https://uaa.sys.dev.phx.pcf.example.com/oauth/token

11. **Update**.
12. To start discovery, navigate to the Serverless Discovery definition and select **Discover Now**.

EVAL closure functions

Use the following EVAL closure functions to concatenate all organizations with "," as a separator.

```
var rtrn = '';
var organizations = ${organizations[*].name};
var beforeReturn = '';
var separator = '';
for (var i = 0; i < organizations.size(); i++)
 {
 beforeReturn += separator + organizations.get(i);
 separator = ',';
 }
rtrn = beforeReturn;
```

Use the following EVAL closure functions to return 1 if the operational status is active and 2 if it isn't active.

```
var rtrn = '';
var currentOrgStatus =
 ${organizations[].operational_status};
var status = currentOrgStatus.toLowerCase();
if(status == "active"){
   rtrn = "1";
}
else {
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
   rtrn = "2";
}
```

Use the following EVAL closure function to return 1 if the operational status is started and 2 if it isn't started.

```
var rtrn = '';
var currentAppStatus = ${apps[].state};
var status = currentAppStatus.toLowerCase();
if(status == "started"){
   rtrn = "1";
}
else {
   rtrn = "2";
}
```

CloudFoundry Get Call

The CloudFoundry Get Call custom operation handles the generation of tokens. It uses the credentials and URL_CF_AUTH context variable to generate a token that is used to create API calls.

When a token is generated, the URL_CF is used to create the API calls defined in the **Resource** field. The custom operation also handles pagination where the response has more than one page.

CloudFoundry Get Call parses the input JSON to populate the tables and variables.

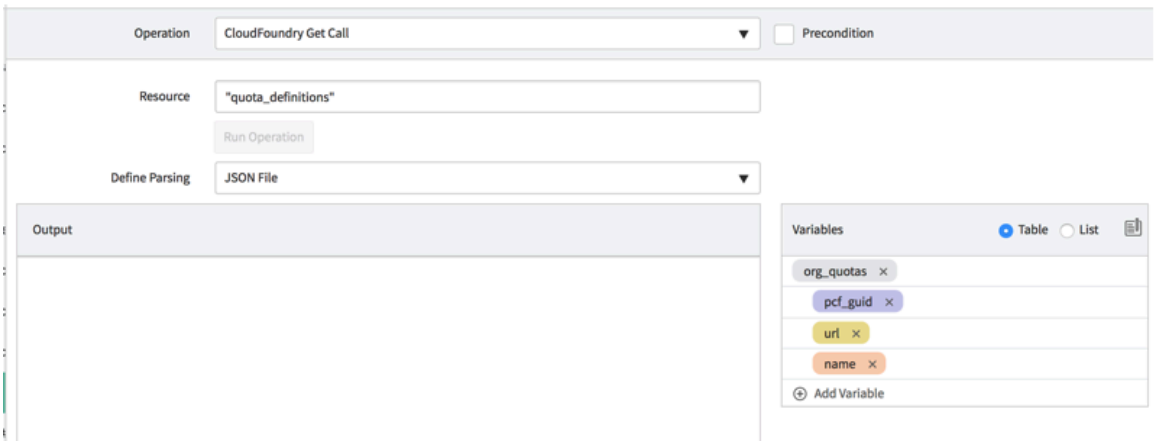

Pivotal Cloud Foundry BOSH extension section

The Pivotal Cloud Foundry BOSH is an extension section of the Pivotal Cloud Foundry Pattern.

The extension is available on the ServiceNow® Store for the June 2023 release of Discovery and Service Mapping Patterns 1.6.0.

Requirements for Pivotal Cloud Foundry BOSH

Note: The following requirements are relevant only to the BOSH extension section.

The Opsman authentication URL has to be provided as a pattern parameter with the name *URL_OPSMAN_AUTH*. For example: `https://<opsman entry point>/uaa/oauth/token`

The added credentials must have permission to execute the following APIs:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `/uaa/oauth/token`
- `/api/v0/deployed/products`

Note: Custom APIs are added as *`Tanzu Opsman Get API`* step to the Pattern.

Data collected by Discovery during horizontal discovery

Collected information

Field	Description
Main CI: PCF Foundation [cmdb_ci_pcf_foundation]	
Name	Pivotal Cloud Foundry cluster entry name `<name>-<hostCf:<portCf>`
ip_address	Pivotal Cloud Foundry host IP.
port	Pivotal Cloud Foundry service port.
short_description	Pivotal Cloud Foundry service build number.
Organization	Organization list, managed by PCF cluster. The entries are separated using a "," comma separator.
PCF Organization [cmdb_ci_pcf_organization]	
Name	Organization name
URL	Organization unique URL
pcf_guid	Organization UID in CF
operational_status	Organization status
PCF Space [cmdb_ci_pcf_space]	
Name	Space name
URL	Space unique URL
pcf_guid	Space UID in CF
PCF Application [cmdb_ci_pcf_appl]	
Name	Application name
URL	Application unique URL
pcf_guid	Application UID in CF
PCF Route [cmdb_ci_pcf_route]	
Name	Route name
URL	Route unique URL
pcf_guid	Route UID in CF
port	Route port
type	Route type
PCF Quota [cmdb_ci_pcf_quota]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Collected information (continued)

Field	Description
Name	Quota name
URL	Quota unique URL
pcf_guid	Quota UID in CF
PCF Domain [cmdb_ci_pcf_domain]	
name	Domain name
url	Domain unique URL
pcf_guid	Domain UID in CF
PCF Space instance [cmdb_ci_pcf_space_instance]	
name	Space service instance name
url	Space service instance unique URL
pcf_guid	Space service instance UID in CF
PCF Service Plan [cmdb_ci_pcf_service_plan]	
name	Service plan instance name
url	Service plan unique URL
pcf_guid	Space service instance UID in CF
PCF Service [cmdb_ci_pcf_service]	
name	Service name
url	Service unique URL
pcf_guid	Service UID in CF
PCF Droplet [cmdb_ci_pcf_droplet]	
name	Resource URL
pcf_guid	Droplet ID
url	Droplet URL
install_status	Installation status: 1, "installed"
operational_status	Operational status: 1, "operational"
Key Value [cmdb_key_value]	
key	Resource Tag key **Note:** This data is only collected for V3 API calls.
value	Resource Tag value. **Note:** This data is only collected for V3 API calls.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data discovered by the Pivotal Cloud Foundry BOSH extension section

Field	Description
BOSH Deployment (cmdb_ci_bosh_deployment)	
Name	The installation name of the resource.
object_id	The ID of the deployed product.
install_status	Hard-coded to the value: *1= installed*.
operational_status	Hard-coded to the value: *1= operational*.
deployment_version	The version of the deployed product.
deployment_type	The type of the deployed product.

CI relationships

These relationships are created to support Pivotal Cloud Foundry discovery.

CI	Relationship	CI
PCF Space [cmdb_ci_pcf_space]		
cmdb_ci_pcf_space	Contains::Contained by	cmdb_ci_pcf_organization
PCF Application [cmdb_ci_pcf_appl]		
cmdb_ci_pcf_appl	Contains::Contained by	cmdb_ci_pcf_space
PCF Route [cmdb_ci_pcf_route]		
cmdb_ci_pcf_route	Contains::Contained by	cmdb_ci_pcf_foundation
PCF Quota [cmdb_ci_pcf_quota]		
cmdb_ci_pcf_quota	Uses::Used by	cmdb_ci_pcf_space cmdb_ci_pcf_organization
PCF Domain [cmdb_ci_pcf_domain]		
cmdb_ci_pcf_domain	Uses::Used by	cmdb_ci_pcf_route
	Owns::Owned by	cmdb_ci_pcf_organization
	Contains::Contained by	cmdb_ci_pcf_foundation
PCF Space instance [cmdb_ci_pcf_space_instance]		
cmdb_ci_pcf_space_instance	Contains::Contained by	cmdb_ci_pcf_space
PCF Service Plan [cmdb_ci_pcf_service_plan]		
cmdb_ci_pcf_service_plan	Uses::Used by	cmdb_ci_pcf_space_instance
PCF Service [cmdb_ci_pcf_service]		
cmdb_ci_pcf_service	Uses::Used by	cmdb_ci_pcf_service_plan
PCF Droplet [cmdb_ci_pcf_droplet]		
cmdb_ci_pcf_appl	Contains::Contained by	cmdb_ci_pcf_droplet

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Key Value [cmdb_key_value]		
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_droplet
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_organization
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_space
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_domain
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_route
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_appl
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_service
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_service_plan
cmdb_key_value	Reference only. For example: configuration_item	cmdb_ci_pcf_space_instance

Pivotal Cloud Foundry BOSH relationships

These relationships are created to support Pivotal Cloud Foundry BOSH extension section discovery.

CI	Relationship	CI
Cloud Foundry Cluster [cmdb_ci_pcf_foundation]	Contains::Contained by	BOSH Deployment [cmdb_ci_bosh_deployment]

Pure Storage FlashArray discovery

The Discovery and Service Mapping Patterns application uses the Pure Storage pattern to find Pure Storage FlashArray on your infrastructure. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Starting March 2024, the pattern execution is disabled by default. To create a CI for Pure Storage FlashArray, synchronize the pattern to the MID Server.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

Verify the configuration of the credentials

Configure basic authentication credentials to classify FlashArray. For more information, see basic authentication.

Verify the permissions to run REST API calls

The Pure Storage Flash Array pattern supports the API 2.17 version. Note that the `/api/1.17/auth/apitoken` endpoint uses an earlier API version than the rest of the API calls.

- `/api/1.17/auth/apitoken`
- `/api/2.17/auth/session`
- `/api/2.17/arrays`
- `/api/2.17/volumes`
- `/api/2.17/volume-groups`
- `/api/2.17/host-groups`
- `/api/2.17/hardware`
- `/api/2.17/dns`
- `/api/2.17/ports`

Note: The February 2024 Pure Storage pattern version supports only iSCSI protocol discovery.

Running Discovery

To run discovery using Patterns, create a discovery schedule that targets the management IP address of the FlashArray storage. For more information, see Running discoveries in your network.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Pure Storage pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server [cmdb_ci_storage_server]

Field	Description
Name [name]	The user-assigned identifier or label for the FlashArray storage.
Serial number [serial_number]	Serial number that uniquely identifies each individual FlashArray storage.
Operating System [os]	Operating system running on the FlashArray storage.
OS Version [os_version]	Software version that is installed on the FlashArray storage.
IP Address [ip_address]	IP Address of the FlashArray storage.
Manufacturer [manufacturer]	Manufacturer name: Pure Storage.
Object ID [object_id]	Credential affinity assigned during discovery to the FlashArray storage. For more information about credential affinity, see Credential affinity for Discovery and Orchestration.
Model ID [model_id]	Model ID of the FlashArray storage.
Encryption at rest [encryption_at_rest]	Encrypted data stored on a physical device or in a database. Possible values are: • active • disable
Encryption algorithm [encryption_algorithm]	Type of algorithm the storage server employs. Possible values are: • AES-256: Advanced Encryption Standard (AES) 256 Encryption • multiple methods: Pure Storage doesn't currently support multiple methods • Unknown: If during a previous discovery an algorithm type was discovered but during the next discovery it isn't discovered, the value is unknown
Total Free Capacity (GB) [total_free_capacity]	Total amount of storage space that is currently available for use on the FlashArray storage system, which is measured in gigabytes (GB).
Physical Capacity (GB) [physical_capacity]	Total amount of physical storage space available within the FlashArray system, measured in GB. This capacity includes both allocated and unallocated storage resources.
Free Physical Capacity (GB) [free_physical_capacity]	Available physical storage space within the FlashArray system, measured in GB.
Used Physical Capacity (GB) [used_physical_capacity]	Used physical storage space within the FlashArray system, measured in GB.
Virtual Capacity (GB) [virtual_capacity]	Total virtual storage capacity provisioned within the FlashArray system, measured in GB.
Free Virtual Capacity (GB) [free_virtual_capacity]	Available virtual storage capacity within the FlashArray system, measured in GB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server [cmdb_ci_storage_server] (continued)

Field	Description
Used Virtual Capacity (GB) [used_virtual_capcity]	Used virtual storage capacity within the FlashArray system, measured in GB.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Storage Cluster Node [cmdb_ci_storage_cluster_node]

Field	Description
Name [name]	Cluster node name.
Cluster [cluster]	Host group associated with this cluster node in the FlashArray environment.
Manufacturer [manufacturer]	Manufacturer name: Pure Storage.
Comments [comments]	All the logical unit numbers (LUN) connected to the cluster node.

Storage Cluster [cmdb_ci_storage_cluster]

Field	Description
Name [name]	Storage cluster name.
Correlation ID [correlation_id]	Combines the cluster name and the storage's serial number to create a unique identifier (UID).
Manufacturer [manufacturer]	Manufacturer name: Pure Storage.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	Volume name.
Serial number [serial_number]	Serial number of the volume.
Volume ID [volume_id]	Volume ID.
Installed [install_date]	Date the volume was installed.
Size bytes [size_bytes]	Volume size in bytes.
Pool ID [pool_id]	ID of the pool containing the volumes.
Attributes [attributes]	Names of the hosts connected to the volumes.
Manufacturer [manufacturer]	Manufacturer name: Pure Storage.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.
Provided by [provided_by]	References the Storage Pool [cmdb_ci_storage_pool] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Pool [cmdb_ci_storage_pool]

Field	Description
Name [name]	Name of the storage pool.
Object ID [object_id]	ID of the storage pool.
Manufacturer [manufacturer]	Manufacturer name: Pure Storage.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] table.

Storage Controller [cmdb_ci_storage_controller]

Field	Description
Name [name]	Combination of asset and asset tag.
Model ID [model_id]	Model ID of the controller.
Assigned [assigned]	Availability status.
Device ID [device_id]	A unique label for the controller.
Manufacturer [manufacturer]	Manufacturer name: Pure Storage.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

Disk [cmdb_ci_disk]

Field	Description
Name [name]	Spindle ID.
Size bytes [size_bytes]	Disk size in bytes.
Drive type [drive_type]	Type of drive installed on the disk.
Device ID [device_id]	Combination of the storage system serial number with the disk name that uniquely identifies each disk within the FlashArray environment.
Manufacturer [manufacturer]	Disk manufacturer.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	Name of the network adapter.
IP Address [ip_address]	IP address of the network adapter.
MAC Address [mac_address]	MAC address of the network adapter.
IP default gateway [ip_default_gateway]	IP address of the default gateway.
Alias [alias]	The user-assigned name for the network adapter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network Adapter [cmdb_ci_network_adapter] (continued)

Field	Description
Netmask [netmask]	Netmask of the network adapter.
Configuration Item [cmdb_ci]	References the Storage Server [cmdb_ci_storage_server] table.

Fibre Channel Export [cmdb_ci_fc_export]

Field	Description
Name [name]	Initiator ID in FlashArray systems, which serves as a UID for the initiator in the context of Fibre Channel exports.
Export ID [export_id]	UID for a Fibre Channel export within FlashArray systems, facilitating the management and tracking of Fibre Channel connections.
Initiator WWPN [initiator_wwpn]	UID combining the Director ID, port number, and Initiator ID.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] table.
Storage [storage]	References the Storage Cluster Node [cmdb_ci_storage_cluster_node] table.

iSCSI Export [cmdb_ci_iscsi_export]

Field	Description
Name [name]	Initiator ID in FlashArray systems, which serves as a UID for the initiator in the context of iSCSI exports.
Export ID [export_id]	UID for an iSCSI export within FlashArray systems, facilitating the management and tracking of iSCSI connections.
Initiator IQN [initiator_iqn]	iSCSI Qualified Name (IQN) assigned to an initiator in the iSCSI network.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] table.
Storage [storage]	References the Storage Cluster Node [cmdb_ci_storage_cluster_node] table.

CI relationships and references

The Pure Storage pattern creates CI relationships and references to support Pure Storage FlashArray discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Disk [cmdb_ci_disk]	Contains::Contained by	Storage Server [cmdb_ci_storage_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Fibre Channel Export [cmdb_ci_fc_export]	Hosted on::Hosts	Storage Server [cmdb_ci_storage_server]
iSCSI Export [cmdb_ci_iscsi_export]	Hosted on::Hosts	Storage Server [cmdb_ci_storage_server]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	Storage Server [cmdb_ci_storage_server]
Storage Cluster Node [cmdb_ci_storage_cluster_node]	Cluster of::Cluster	Storage Cluster [cmdb_ci_storage_cluster]
Storage Controller [cmdb_ci_storage_controller]	Controller for::Controlled by	Storage Server [cmdb_ci_storage_server]
Storage Server [cmdb_ci_storage_server]	Allocated to::Allocated from	Storage Cluster Node [cmdb_ci_storage_cluster_node]
Storage Server [cmdb_ci_storage_server]	Contains::Contained by	Storage Pool [cmdb_ci_storage_pool]
Storage Server [cmdb_ci_storage_server]	Runs on::Runs	Storage Cluster [cmdb_ci_storage_cluster]
Storage Volume [cmdb_ci_storage_volume]	Contains::Contained by	Storage Cluster Node [cmdb_ci_storage_cluster_node]
Storage Volume [cmdb_ci_storage_volume]	Contains::Contained by	Storage Server [cmdb_ci_storage_server]

CI references

CI	Field	Referenced CI
Fibre Channel Export [cmdb_ci_fc_export]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server]
Fibre Channel Export [cmdb_ci_fc_export]	Storage [storage]	Storage Cluster Node [cmdb_ci_storage_cluster_node]
iSCSI Export [cmdb_ci_iscsi_export]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server]
iSCSI Export [cmdb_ci_iscsi_export]	Storage [storage]	Storage Cluster Node [cmdb_ci_storage_cluster_node]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	Storage Server [cmdb_ci_storage_server]
Disk [cmdb_ci_disk]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Storage Controller [cmdb_ci_storage_controller]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Storage Pool [cmdb_ci_storage_pool]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
Storage Volume [cmdb_ci_storage_volume]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Storage Volume [cmdb_ci_storage_volume]	Provided by [provided_by]	Storage Pool [cmdb_ci_storage_pool]

Pure Storage FlashBlade discovery

The ServiceNow Discovery application uses the FlashBlade Pure Storage pattern to find FlashBlade components. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery can find Pure Storage FlashBlade version 4.

Note:

Only Linux MID Servers are supported for Pure Storage FlashBlade discovery. Windows MID Servers are not supported.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Configure user permissions for UNIX OS

Configure a user for accessing the UNIX OS.

- Define the OS user with permissions to run the following commands.

Command	Parameter	Mandatory/ Optional	Description
Pureman	-	Mandatory	Accesses the list of pure storage commands and uses them for classification.
Echo	$$username$$	Mandatory	Fills in the username for the temporary variable.
"pureadmin list "	+ $usr + " --api-token --expose --csv"	Mandatory	Gets the user API token to generate the x-auth for the REST call commands.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Define the OS user with the required permissions to run the following REST API calls.
  - `/api/login` – Gets the x-auth token reader from the response in Groovy.
  - `/api/1.2/file -systems` – Retrieves information about the file systems.
  - `/api/1.2/blades` – Retrieves information about blades.
  - `/api/1.2/dns` – Retrieves information about DNS.
  - `/api/1.2/hardware` – Retrieves information about the storage hardware.
  - `/api/1.4/network-interfaces` – Retrieves information about network interfaces.

Configure SSH credentials

On the ServiceNow AI Platform, configure SSH credentials for the OS user. For more information, see SSH credentials.

Configure applicative credentials

On the ServiceNow AI Platform, configure applicative credentials for the OS user. The password field is not relevant; you can enter any value for the password or leave it empty. For more information, see Applicative credentials.

Access to the port used for REST calls

On relevant Linux MID Servers, provide a user with access to the port used for REST calls.

Data collected by Discovery for Pure Storage FlashBlade

Discovery populates the data in the CMDB when running the FlashBlade Pure Storage pattern.

Table and field	Description
Storage Server [cmdb_ci_storage_server]	The attributes of the storage server provided by FileShare for PureBlade.
IP address [ip_address]	
Serial number [serial_number]	
Name [name]	
Pure Storage Processor [cmdb_ci_storage_processor_pure]	The attributes of the PureBlade storage processor.
Name [name]	
Serial number [serial_number]	
Device ID [device_id]	
Model number [model_number]	
Serial [serial]	
Slot [slot]	
Type [type]	
Status [status]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Pure Storage File Share [cmdb_ci_storage_fileshare_pure]	The attributes of the storage FileShare.
Name [name]	
Path [path]	
Protocols [protocols]	
Provisioned Space [provisioned_space]	
Unique Space [unique_space]	
Snapshots Space [snapshots_space]	
Total Physical Space [total_physical_space]	
Shared space [shared_space]	
Data Reduction [data_reduction]	
System Space [system_space]	
Storage Device [cmdb_ci_storage_device]	The attributes of the storage device contained by the storage server.
Device ID [device_id]	
Name [name]	
Serial number [serial_number]	
Model number [model_number]	
Storage type [storage_type]	

The graphic illustrates CIs that are part of FlashBlade discovery.

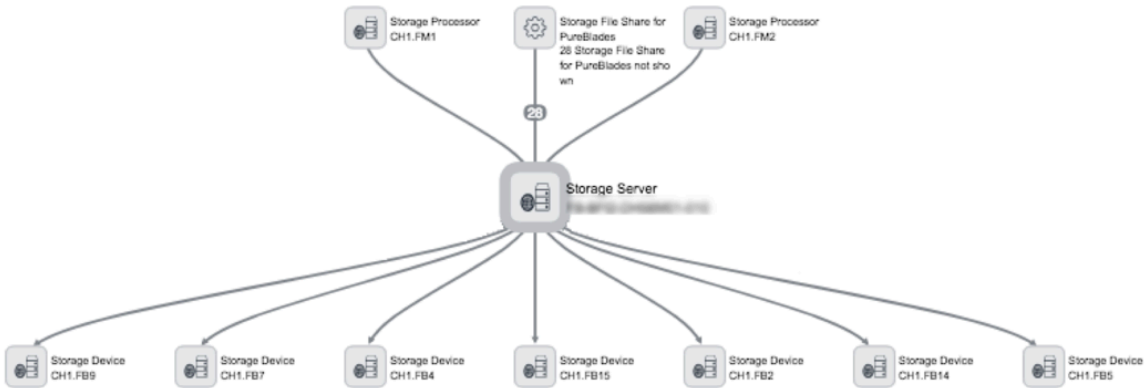

CI relationships

Discovery creates these relationships to support the Pure Storage FlashBlade discovery.

CI	Relationship	CI
Pure Storage File Share [cmdb_ci_storage_fileshare_pure]	Provided By::Provides	Storage Server [cmdb_ci_storage_server]
Storage Device [cmdb_ci_storage_device]	Contains::Contained by	Storage Server [cmdb_ci_storage_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Storage Server [cmdb_ci_storage_server]	Contains::Contained by	Pure Storage Processor [cmdb_ci_storage_processor_pure]

Red Hat 389 Directory Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Red Hat - Directory Server pattern to find 389 Directory Server instances running on UNIX servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create SSH credentials

For more information, see SSH credentials.

Verify permission to query the RPM database

Verify you have permission for the following command: `rpm -qf {process.executablePath}`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Red Hat - Directory Server pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Red Hat 389 Directory Server @{hostname}**.
Product [product]	Product name. The value is set to **389 Directory Server**.
Publisher [publisher]	Publisher of the application. The value is set to **Red Hat**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Red Hat - Directory Server pattern creates the following relationships and references to support Red Hat 389 Directory Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Red Hat JBoss EAP Host Controller pattern-based discovery

Discovery and Service Mapping Patterns uses the Red Hat JBoss - Enterprise App Platform Host Controller pattern to find JBoss Enterprise Application Platform (JBoss EAP) Host Controller instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the application module files on the application server

Verify you have permission for the following command:
`powershell -NoProfile -ExecutionPolicy Bypass -Command "Get-ChildItem '{process.environmentVariables.JBOSS_HOME.value}\modules' -Recurse -Filter MANIFEST.MF | Where-Object { $_.FullName -like '*org\jboss\as\product*' } | ForEach-Object { Select-String -Path $_.FullName -Pattern 'JBoss-Product-Release-(Name|Version)' }"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Red Hat JBoss - Enterprise App Platform Host Controller pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Red Hat JBoss Enterprise Application Platform Host Controller@{hostname}**.
Product [product]	Product name. The value is set to **JBoss Enterprise Application Platform**.
Publisher [publisher]	Publisher of the application. The value is set to **Red Hat**.
Component [component]	Component of the application. The value is set to **Host Controller**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Red Hat JBoss - Enterprise App Platform Host Controller pattern creates the following relationships and references to support Red Hat JBoss EAP Host Controller discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Red Hat JBoss EAP Server pattern-based discovery on UNIX

Discovery and Service Mapping Patterns uses the Red Hat - JBoss Enterprise App Platform Server on Unix pattern to find JBoss Enterprise Application Platform (JBoss EAP) Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

instances running on UNIX servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create SSH credentials

For more information, see SSH credentials.

Verify read access to the version file

Verify you have permission for the following command: `cat {process.currentDir}/version.txt`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Red Hat - JBoss Enterprise App Platform Server on Unix pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Red Hat JBoss Enterprise Application Platform Server@{hostname}**.
Product [product]	Product name. The value is set to **JBoss Enterprise Application Platform**.
Publisher [publisher]	Publisher of the application. The value is set to **Red Hat**.
Component [component]	Component of the application. The value is set to **Server**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships and references

The Red Hat - JBoss Enterprise App Platform Server on Unix pattern creates the following relationships and references to support Red Hat JBoss EAP Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Red Hat JBoss EAP Server pattern-based discovery on Windows

Discovery and Service Mapping Patterns uses the Red Hat JBoss Enterprise Application Platform Server pattern to find JBoss Enterprise Application Platform (JBoss EAP) Server instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the version file on the application server

Verify you have permission for the following command:
`powershell -NoProfile -Command "Get-Content '{process.environmentVariables.JBOSS_HOME.value}\version.txt'"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Red Hat JBoss Enterprise Application Platform Server pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Red Hat JBoss Enterprise Application Platform Server@{hostname}**.
Product [product]	Product name. The value is set to **JBoss Enterprise Application Platform**.
Publisher [publisher]	Publisher of the application. The value is set to **Red Hat**.
Component [component]	Component of the application. The value is set to **Server**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Red Hat JBoss Enterprise Application Platform Server pattern creates the following relationships and references to support Red Hat JBoss EAP Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Red Hat JBoss Fuse discovery

The ServiceNow Discovery and Service Mapping applications can find and map the Fuse application server using the JBoss Fuse pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Note: The Discovery and Service Mapping Patterns application supports JBoss Fuse running on Karaf only.

Prerequisites

Credentials

Configure SSH credentials in the Credentials module of the ServiceNow platform for accessing the server hosting the Fuse server.

User

Provide a user with the read rights for accessing the following directories:

- `$install_directory + "/readme.txt"`

  Access to this directory allows getting version information.

- /etc/services

  Access to this directory allows reading the hosts service definition.

Data collected by Discovery during horizontal discovery

Table and Field	Description
Jboss Fuse [cmdb_ci_appl_jboss_fuse]	
config_directory	The configuration directory of the Fuse server.
config_file	The configuration file of the Fuse server.
install_folder	The installation folder of the Fuse server.
version	The version of the Fuse server.
name	The name configured for the Fuse server during installation.

Data collected by Service Mapping during top-down discovery

Service Mapping discovers the connection between the JBoss Fuse server and the IBM Informix Dynamic server.

Red Hat Podman pattern-based discovery

Discovery and Service Mapping Patterns uses the Red Hat - Podman pattern to find Podman container engines running on UNIX servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create SSH credentials

For more information, see SSH credentials.

Verify permission to run the version command

Verify you have permission for the following command: `{process.executablePath} --version`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Red Hat - Podman pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Red Hat Podman@{hostname}**.
Product [product]	Product name. The value is set to **Podman**.
Publisher [publisher]	Publisher of the application. The value is set to **Red Hat**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Red Hat - Podman pattern creates the following relationships and references to support Red Hat Podman discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Linux Server [cmdb_ci_linux_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Red Hat Single Sign-On Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Red Hat - Single Sign-On Server pattern to find Single Sign-On (SS) Server instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify read access to the version file

Verify you have permission for the following command:
`powershell -NoProfile -Command "Get-Content '{process.environmentVariables.JBOSS_HOME.value}\version.txt'"`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Red Hat - Single Sign-On Server pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Red Hat Single Sign-On Server@{hostname}**.
Product [product]	Product name. The value is set to **Single Sign-On**.
Publisher [publisher]	Publisher of the application. The value is set to **Red Hat**.
Component [component]	Component of the application. The value is set to **Server**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application [cmdb_ci_appl] (continued)

Field	Description
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Red Hat - Single Sign-On Server pattern creates the following relationships and references to support Red Hat SSO Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

Red Hat WildFly Server pattern-based discovery

Discovery and Service Mapping Patterns uses the Red Hat - WildFly Server pattern to find WildFly Server instances running on Windows servers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- ITOM Content Service

Create Windows credentials

For more information, see Windows credentials.

Verify permission to run the WildFly version command

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify you have permission for the following command: `"{executablePath}standalone.bat" --version`.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Red Hat - WildFly Server pattern.

If you have the Software Asset Management Core (com.snc.sam.core) or Software Asset Management (com.snc.software_asset_management) plugins, Discovery also populates the Software Installation [cmdb_sam_sw_install] table.

Application [cmdb_ci_appl]

Field	Description
Name [name]	Name of the application in the following format: **Red Hat WildFly Server@{hostname}**.
Product [product]	Product name. The value is set to **WildFly**.
Publisher [publisher]	Publisher of the application. The value is set to **Red Hat**.
Component [component]	Component of the application. The value is set to **Server**.
Installation directory [install_directory]	Installation directory of the application.
Version [version]	Version of the application.
Software Install [software_install]*	References the Software Installation [cmdb_sam_sw_install] table.

* Populated only when a Software Asset Management plugin is activated.

CI relationships and references

The Red Hat - WildFly Server pattern creates the following relationships and references to support Red Hat WildFly Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Application [cmdb_ci_appl]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
Application [cmdb_ci_appl]	Software Install [software_install]	Software Installation [cmdb_sam_sw_install]*

* Populated only when a Software Asset Management plugin is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Rubrik Cluster discovery

Discovery uses multiple patterns to find all Rubrik cluster data. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

All of the information for the main CI (Rubrik Cluster) and its corresponding CI types is returned by REST API calls.

Take note of the following limitations:

- Reference fields aren't supported due to the Large Payload mechanism limitations in the Paris and Quebec releases.
- All the entities returned from the Rubrik REST API for Hosts (Servers) with IP addresses not resolved to hostnames are filtered out (excluded) from the Rubrik discovery. This is done to prevent the creation of duplicates across the CMDB.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models
- Visibility Content

Configure Basic Auth permissions

Configure Basic Auth credentials for a user with Read-Only Administrator role in Rubrik, and create a credential alias. For more information, see Create a basic authentication credential alias for Rubrik discovery.

Verify MID Server can access target

Verify that the MID Server can access the target appliance device.

Create a serverless discovery schedule

Create a serverless discovery schedule to perform targeted discovery of Rubrik cluster resources. For more information, see Create a serverless discovery schedule for Rubrik cluster discovery.

Data collected by Discovery during horizontal discovery

Resources discovered by the Rubrik AIX Host (LP) pattern

AIX Server [cmdb_ci_aix_server]

Field	Description
Name [name]	Name of the host.
Host name [host_name]	Host name of the host.
Cluster ID [cluster_id]	ID of the Rubrik cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AIX Server [cmdb_ci_aix_server] (continued)

Field	Description
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Operating System [os]	Operating system type configured on the host.
OS Version [os_version]	Operating system version, edition, or distribution configured on the host.
Operational status [operational_status]	Current operational status of the host.

Resources discovered by the Rubrik Cluster (LP) pattern

Rubrik Cluster [cmdb_ci_rubrik_cluster]

Field	Description
Cluster ID [cluster_id]	Unique ID of the Rubrik cluster.
Name [name]	Name of the Rubrik cluster.
Cluster Version [cluster_version]	Current version of the Rubrik cluster.
API Version [api_version]	API version of the Rubrik cluster.
Timezone [timezone]	Timezone used by the Rubrik cluster.

The Dependency Views map shows discovered cluster resources and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

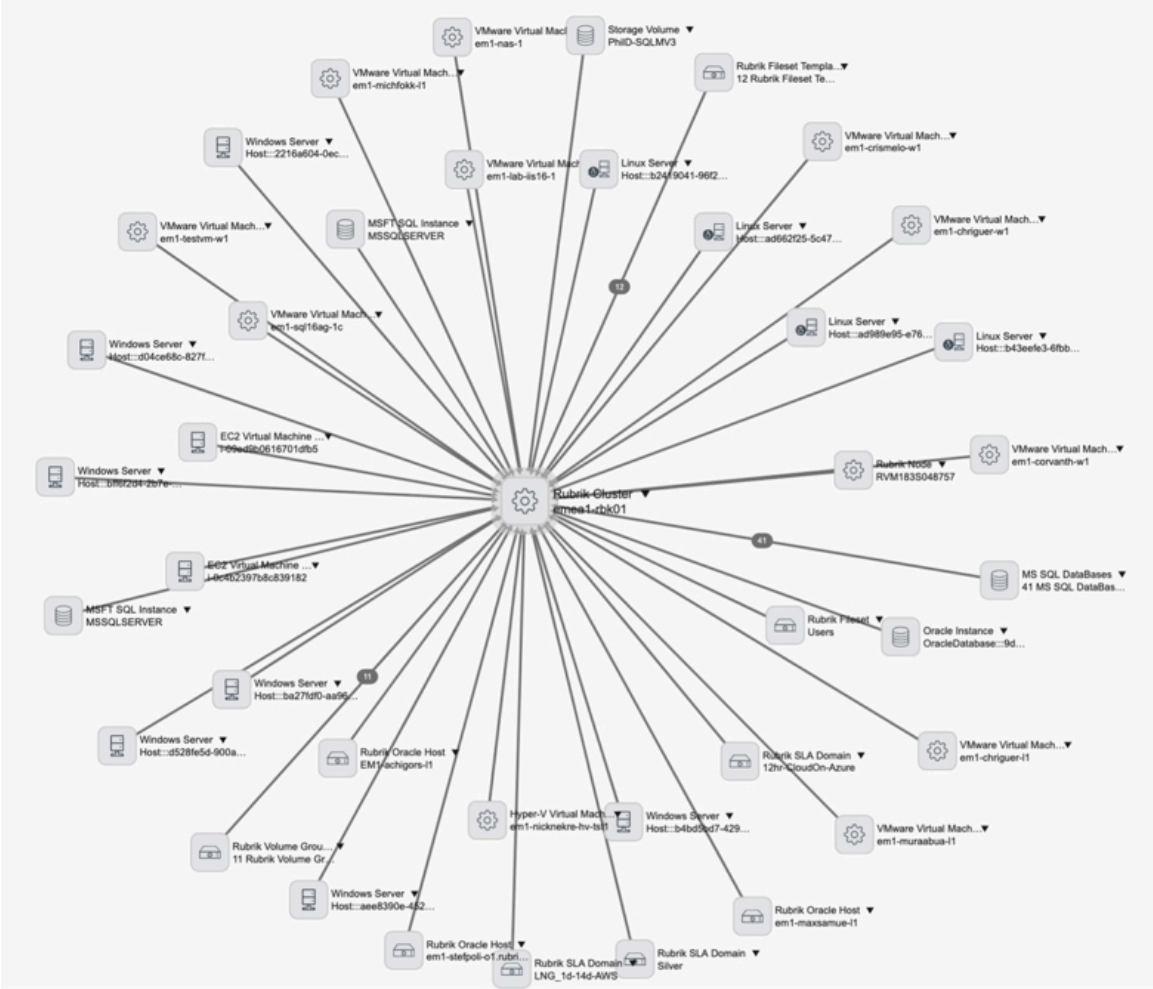

Resources discovered by the Rubrik EC2 Instance (LP) pattern

Rubrik Virtual Machine Instance [cmdb_ci_rubrik_vm_instance]

Field	Description
Object ID [object_id]	ID of the EC2 instance in the AWS context.
Name [name]	Name of the EC2 instance in the AWS context.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
Account ID [account_id]	ID of the AWS account associated with the EC2 instance.
Account name [account_name]	Name of the AWS account associated with the EC2 instance.
Instance ID [instance_id]	ID of the EC2 instance in the AWS context.
Instance name [instance_name]	Name of the EC2 instance in the AWS context.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Rubrik Virtual Machine Instance [cmdb_ci_rubrik_vm_instance] (continued)

Field	Description
Instance type [instance_type]	Template type of the EC2 instance.
Region [region]	Geographical region where the EC2 instance is located, in the AWS context.

The Dependency Views map shows discovered EC2 instances and the relationships between them.

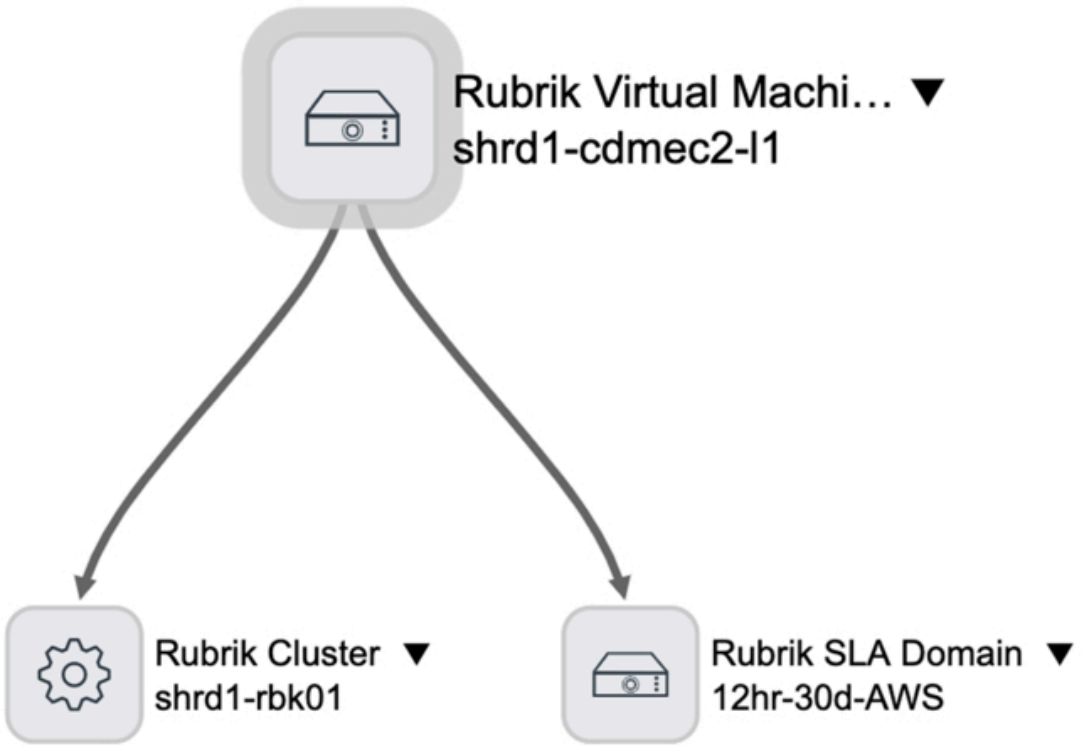

Resources discovered by the Rubrik Fileset (LP) pattern

Rubrik Fileset [cmdb_ci_rubrik_fileset]

Field	Description
Object ID [object_id]	ID of the Rubrik Fileset.
Name [name]	Name of the Rubrik Fileset.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
Host ID [host_id]	ID of the host associated with the Rubrik Fileset.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Rubrik Fileset [cmdb_ci_rubrik_fileset] (continued)

Field	Description
Template ID [template_id]	ID of the Rubrik Fileset Template from which the Rubrik Fileset defines resources.
Share ID [share_id]	ID of the Rubrik Share owning the Rubrik Fileset.

The Dependency Views map shows discovered filesets and the relationships between them.

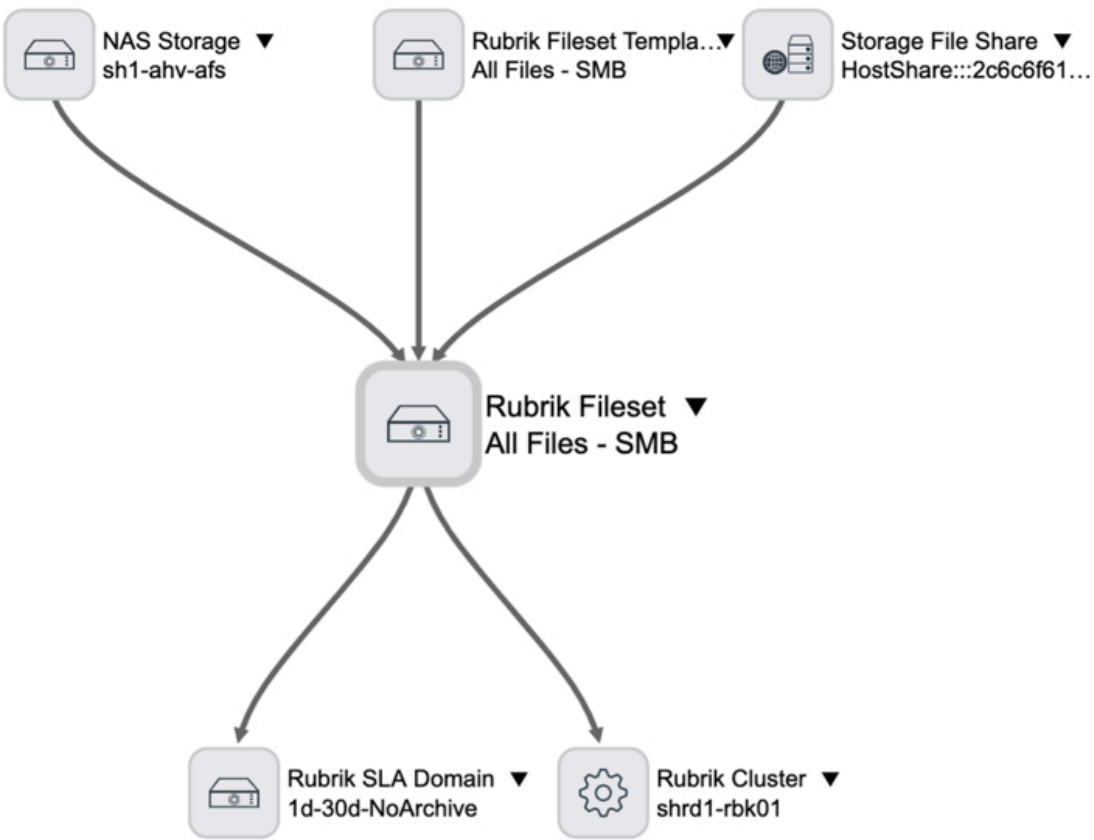

Resources discovered by the Rubrik Fileset Template (LP) pattern

Rubrik Fileset Template [cmdb_ci_rubrik_fileset_template]

Field	Description
Object ID [object_id]	ID of the Rubrik Fileset Template.
Name [name]	Name of the Rubrik Fileset Template.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Share count [share_count]	Number of shares associated with the Rubrik Fileset Template.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Rubrik Fileset Template [cmdb_ci_rubrik_fileset_template] (continued)

Field	Description
Share type [share_type]	Type of shares associated with the Rubrik Fileset Template.
Host count [host_count]	Number of hosts associated with the Rubrik Fileset Template.
OS type [os_type]	Operating system type associated with the Rubrik Fileset Template.

The Dependency Views map shows discovered fileset templates and the relationships between them.

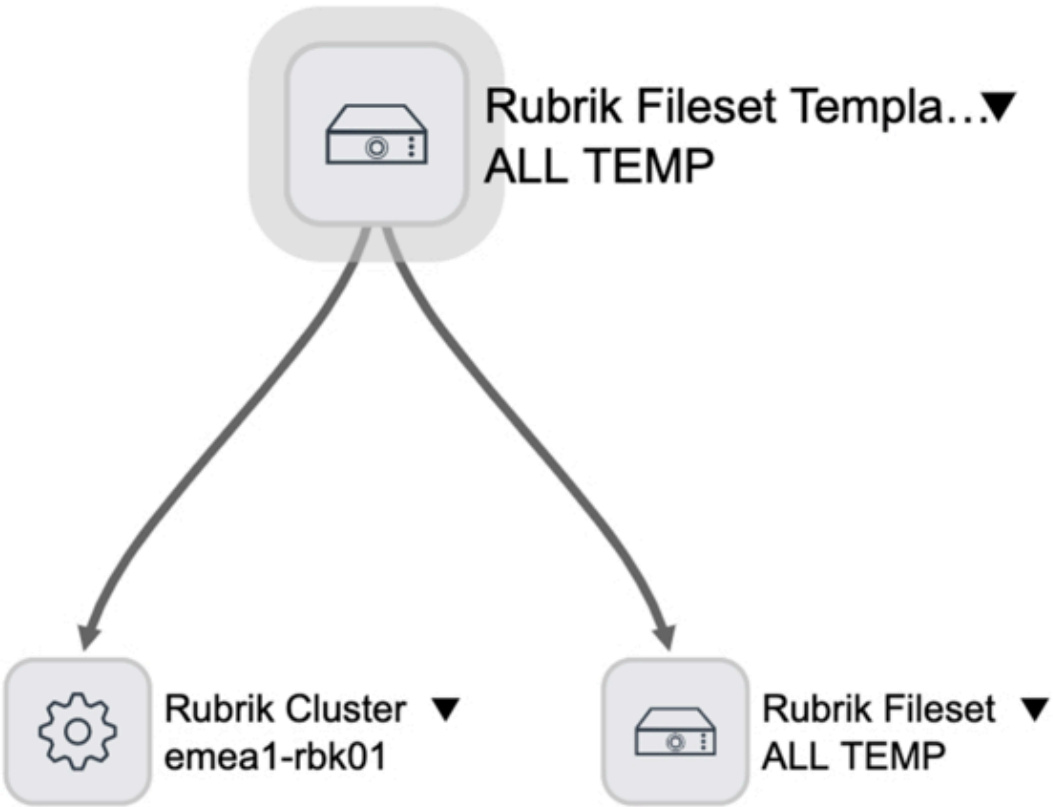

Resources discovered by the Rubrik Hyper-V Instance (LP) pattern

Note: Discovery of Hyper-V instances only updates existing CIs. This is due to a limitation in the Rubrik REST API not returning the GUID (object_id) of the Hyper-V instance. If there is no match between the data from the REST API and the CMDB, no CIs are updated. In addition, when there is no match, the Rubrik Hyper-V Instance (LP) pattern completes in an 'Error' state, which is expected behavior.

Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance]

Field	Description
Name [name]	Name of the Hyper-V instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance] (continued)

Field	Description
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
Host ID [host_id]	ID of the host associated with the Hyper-V instance.

The Dependency Views map shows discovered Hyper-V instances and the relationships between them.

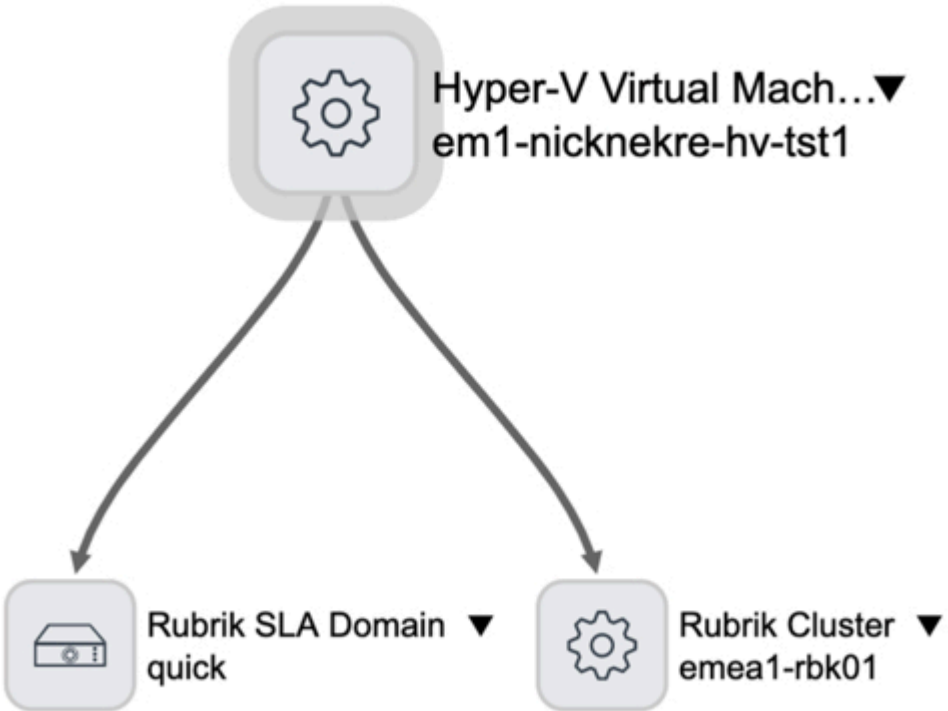

Resources discovered by the Rubrik Linux Host (LP) pattern

Linux Server [cmdb_ci_linux_server]

Field	Description
Name [name]	Name of the host.
Host name [host_name]	Host name of the host.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux Server [cmdb_ci_linux_server] (continued)

Field	Description
Operating System [os]	Operating system type configured on the host.
OS Version [os_version]	Operating system version, edition, or distribution configured on the host.
Operational status [operational_status]	Current operational status of the host.

The Dependency Views map shows discovered hosts and the relationships between them.

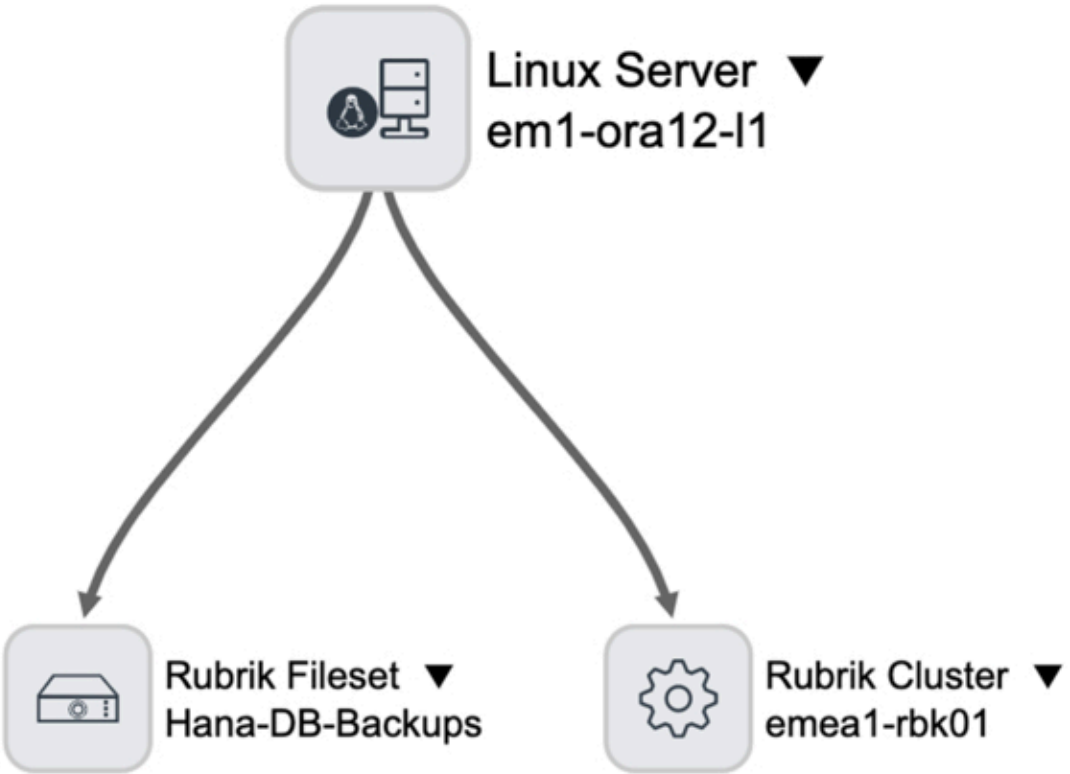

Resources discovered by the Rubrik Managed Volume (LP) pattern

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Object ID [object_id]	ID of the managed storage volume.
Volume ID [volume_id]	Volume ID of the managed storage volume.
Name [name]	Name of the managed storage volume.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume [cmdb_ci_storage_volume] (continued)

Field	Description
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
Size bytes [size_bytes]	Total capacity size in bytes of the storage volume.
Used Size bytes [used_size_bytes]	Total used capacity size in bytes.
Snapshot count [snapshot_count]	Snapshot count of the managed storage volume.
Share type [share_type]	Share type of the managed storage volume.
State [state]	State of the managed storage volume.

The Dependency Views map shows discovered managed storage instances and the relationships between them.

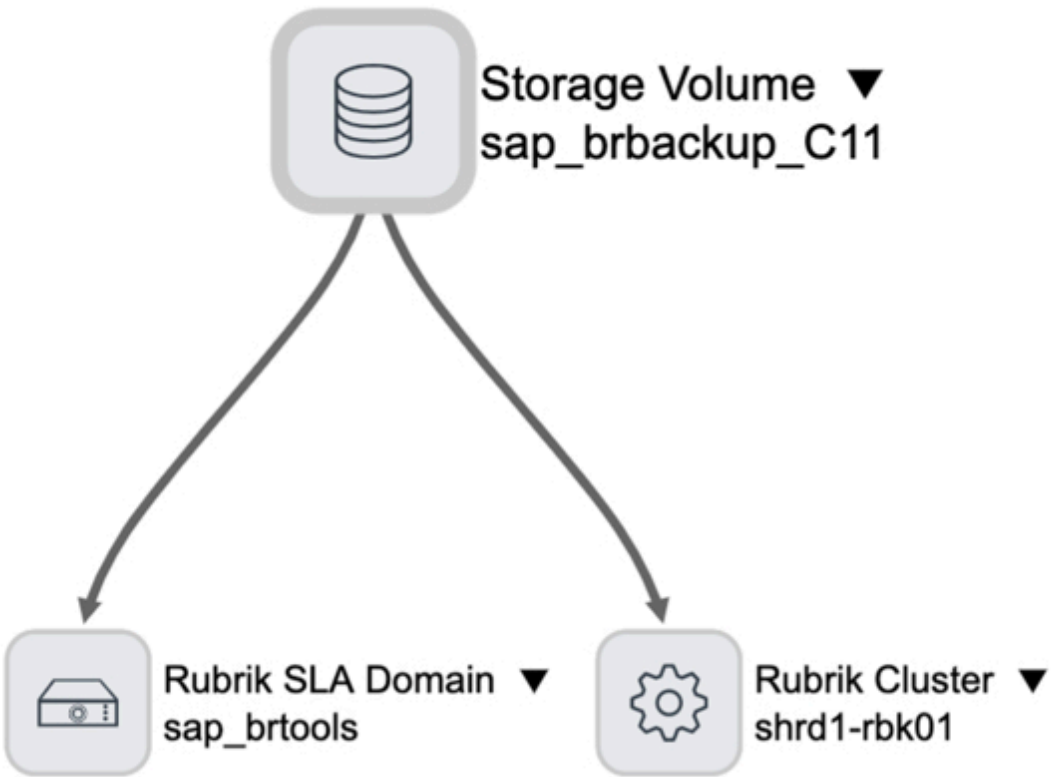

Resources discovered by the Rubrik MS SQL DB (LP) pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Relationships between MS SQL instances hosted on Windows clusters and the associated MS SQL databases are not supported from the 1.0.79 release onwards of the Discovery and Service Mapping Patterns application. This is due to insufficient information returned from the Rubrik REST API for Windows clusters. The MS SQL Instances and MS SQL databases associated to Windows clusters are filtered out from the Rubrik discovery associated patterns.

MS SQL DataBase [cmdb_ci_db_mssql_database]

Field	Description
Data Base [database]	ID of the MS SQL database.
Name [name]	Name of the MS SQL database.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
Host ID [host_id]	ID of the host where the MS SQL database is running.
Instance ID [instance_id]	ID of the database instance associated with the MS SQL database.
State [state]	State of the MS SQL database.
Recovery model [recovery_model]	Recovery model type of the MS SQL database.
Last Snapshot [last_snapshot]	Last snapshot of the MS SQL database, in the format YY-MM-DD HH-MM-SS.
Missed Snapshot count [missed_snapshot_count]	Missed snapshot count of the MS SQL database.

The Dependency Views map shows discovered MS SQL DBs and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

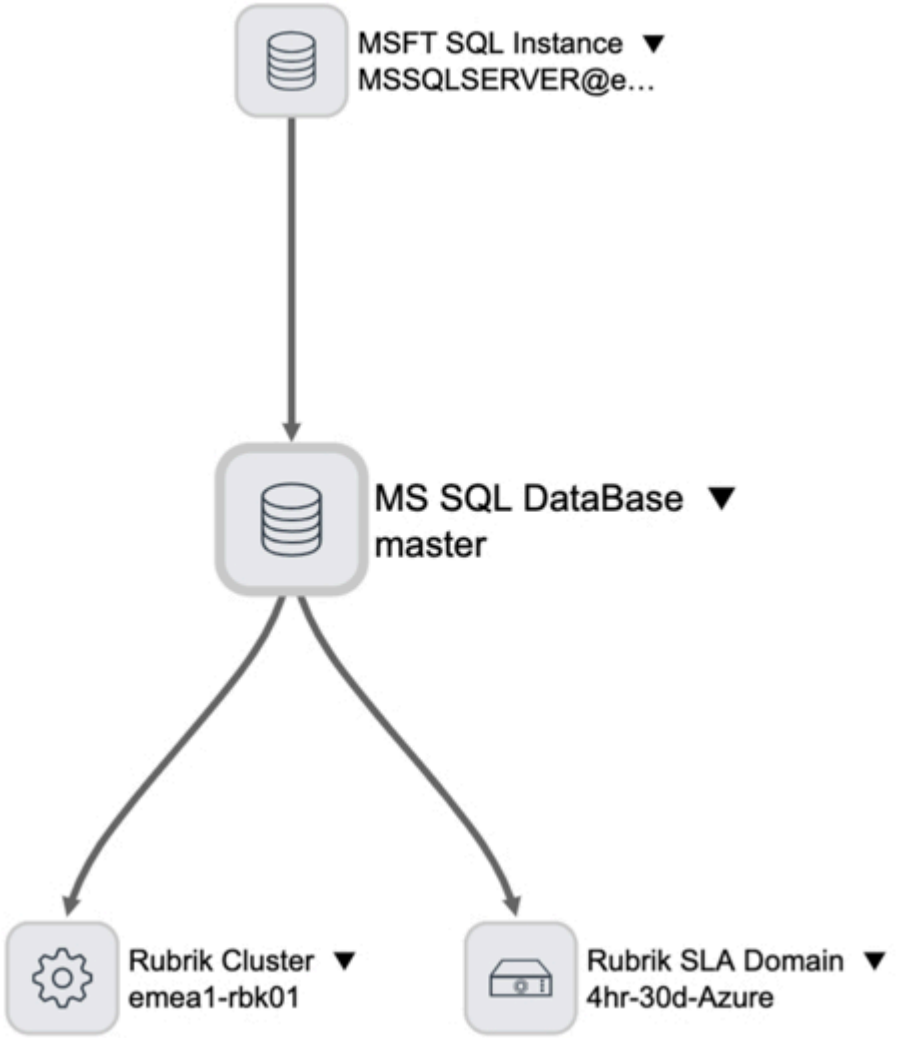

Resources discovered by the Rubrik MS SQL Instance (LP) pattern

Note: Relationships between MS SQL instances hosted on Windows clusters and the associated MS SQL databases are not supported from the 1.0.79 release onwards of the Discovery and Service Mapping Patterns application. This is due to insufficient information returned from the Rubrik REST API for Windows clusters. The MS SQL Instances and MS SQL databases associated to Windows clusters are filtered out from the Rubrik discovery associated patterns.

MSFT SQL Instance [cmdb_ci_db_mssql_instance]

Field	Description
Instance Name [instance_name]	ID of the MS SQL instance.
Name [name]	Name of the MS SQL instance.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Host ID [host_id]	ID of the host associated with the MS SQL instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MSFT SQL Instance [cmdb_ci_db_mssql_instance] (continued)

Field	Description
Version [version]	Version of the MS SQL instance.

The Dependency Views map shows discovered MS SQL instances and the relationships between them.

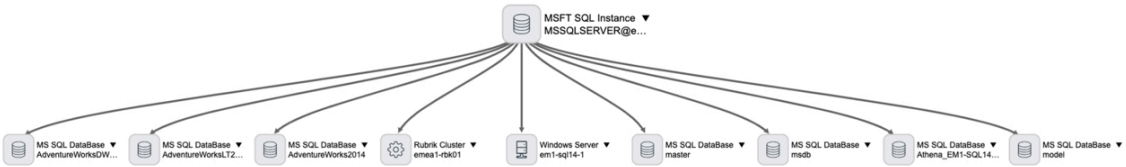

Resources discovered by the Rubrik NAS Host (LP) pattern

NAS Storage [cmdb_ci_nas_storage]

Field	Description
Name [name]	Name of the host.
Host name [host_name]	Host name of the host.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Description [short_description]	Vendor type of the NAS storage.
Operational status [operational_status]	Current operational status of the host.

The Dependency Views map shows discovered hosts and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

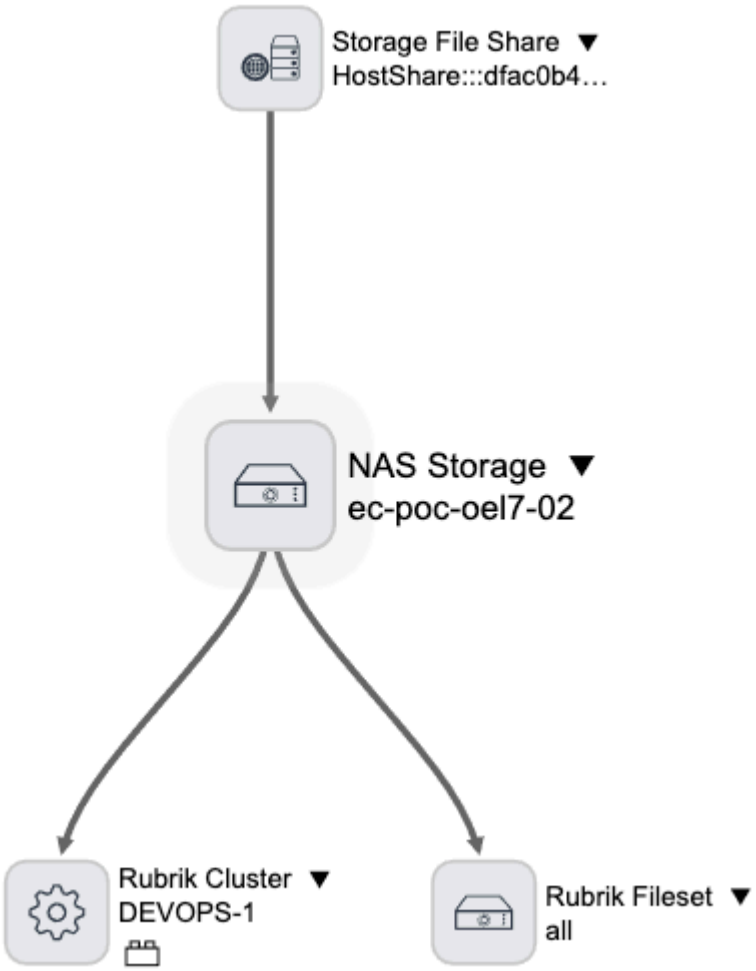

Resources discovered by the Rubrik Node (LP) pattern

Rubrik Node [cmdb_ci_rubrik_node]

Field	Description
Brik ID [brik_id]	Chassis ID of the Rubrik node.
Name [name]	Unique name of the Rubrik node, serving as the ID of the Rubrik node.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Unavailable disks [unavailable_disks]	Indicates whether any disks on the Rubrik node are unavailable. The value is true or false.
IP Address [ip_address]	IP address of the Rubrik node associated with the Rubrik cluster.
Tunnel enabled [tunnel_enabled]	Indicates whether tunneling is enabled for the Rubrik node. The value is true or false.
CPU core count [cpu_core_count]	Number of CPU cores on the Rubrik node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Rubrik Node [cmdb_ci_rubrik_node] (continued)

Field	Description
RAM (MB) [ram]	Amount of RAM in MB allocated for the Rubrik node.

Rubrik Node Disk [cmdb_ci_rubrik_node_disk]

Field	Description
Node Name [node_name]	Name of the Rubrik node associated with the disk.
Device ID [device_id]	ID of the Rubrik node disk, concatenated with the Rubrik node name in the format `ID@NODE_NAME`.
Name [name]	Name of the disk associated with the Rubrik node, in the format `NODE_NAME@DISK_ID@DISK_TYPE`.
Size bytes [size_bytes]	Total capacity size in bytes.
Unallocated size bytes [unallocated_size_bytes]	Total unallocated capacity size in bytes.
Usable size bytes [usable_size_bytes]	Total available capacity size in bytes.
Storage type [storage_type]	Type of the storage disk: SSD or HDD.
Storage path [storage_path]	Location path of the storage disk.
Operational status [operational_status]	Operational status of the storage disk.

The Dependency Views map shows discovered nodes and node disks, and the relationships between them.

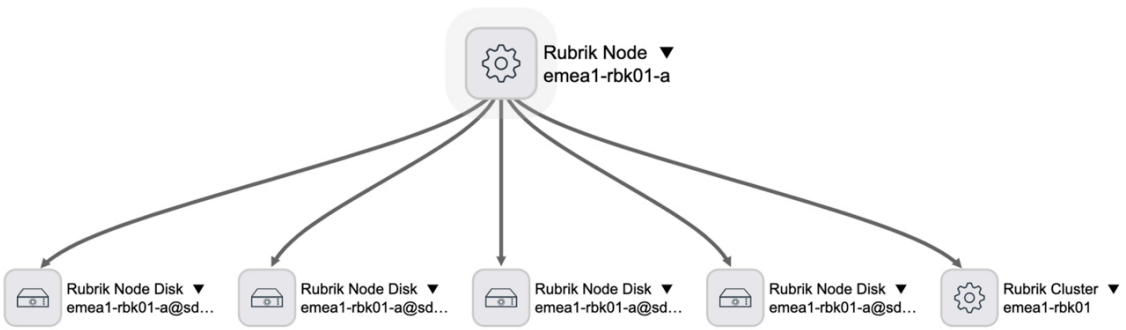

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

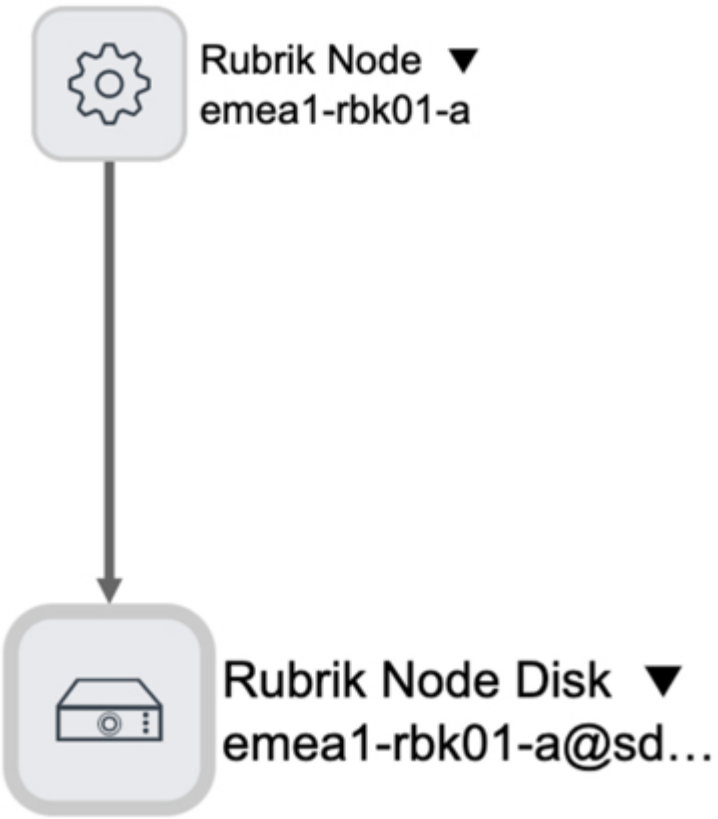

Resources discovered by the Rubrik Nutanix Instance (LP) pattern

Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]

Field	Description
Object ID [object_id]	ID of the Nutanix VM instance.
VM Instance ID [vm_inst_id]	ID of the Nutanix VM instance.
Name [name]	Name of the Nutanix VM instance.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
VM Cluster ID [vm_cluster_id]	ID of the cluster to which the Nutanix VM instance belongs.
VM Cluster name [vm_cluster_name]	Name of the cluster to which the Nutanix VM instance belongs.
OS type [os_type]	Operating system type configured on the Nutanix VM instance.

The Dependency Views map shows discovered Nutanix instances and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

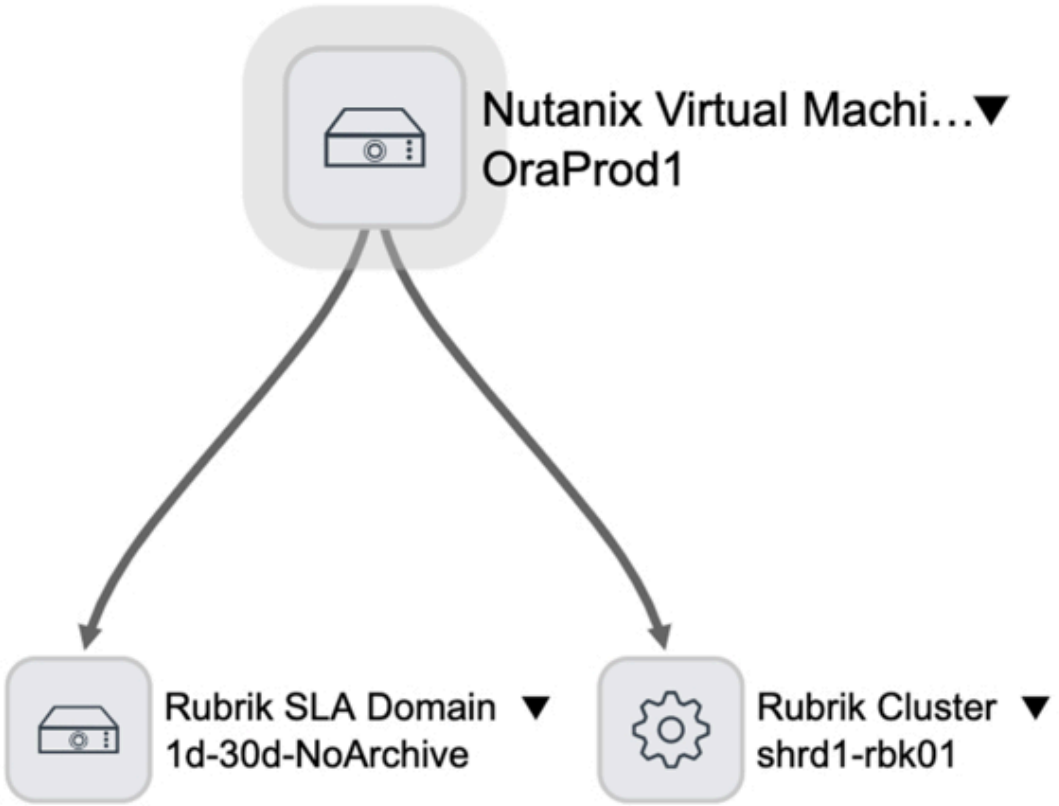

Resources discovered by the Rubrik Oracle DB (LP) pattern

Oracle Instance [cmdb_ci_db_ora_instance]

Field	Description
SID [sid]	Oracle system ID of the database.
Name [name]	ID of the Oracle database in the Rubrik context.
DB Name [database_name]	Name of the Oracle database.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
RAC_ID [rac_id]	ID of the Oracle RAC (Real Application Cluster) to which the database belongs.
Host ID [host_id]	ID of the host where the Oracle database is running.
Instance [instance]	Oracle instance ID managing the database data.
Instance count [instance_count]	Number of Oracle instances managing the database data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Instance [cmdb_ci_db_ora_instance] (continued)

Field	Description
Last Snapshot [last_snapshot]	Last snapshot of the Oracle database, in the format YY-MM-DD HH-MM-SS.
Missed Snapshot count [missed_snapshot_count]	Missed snapshot count of the Oracle database.
Tablespaces count [tablespaces_count]	Number of logical storage units associated with the Oracle database.

The Dependency Views map shows discovered Oracle DBs and the relationships between them.

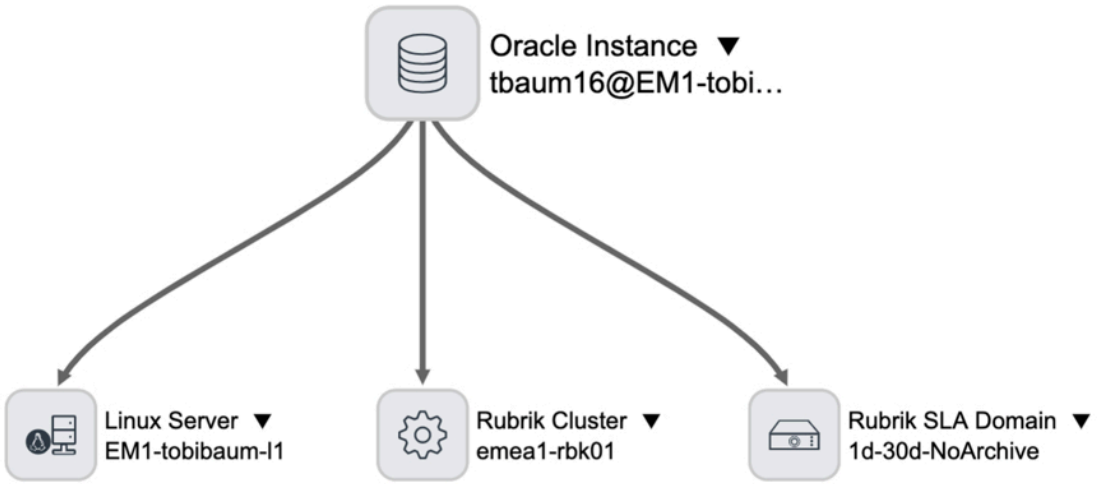

Resources discovered by the Rubrik Oracle RAC (Real Application Cluster) (LP) pattern

Rubrik Oracle RAC [cmdb_ci_rubrik_db_ora_rac]

Field	Description
Object ID [object_id]	ID of the Rubrik Oracle RAC.
Name [name]	Name of the Rubrik Oracle RAC.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
DB count [db_count]	Number of databases associated with the Rubrik Oracle RAC.
Node count [node_count]	Number of nodes participating in the Rubrik Oracle RAC.
State [state]	State of the Rubrik Oracle RAC.

The Dependency Views map shows discovered Oracle RACs and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

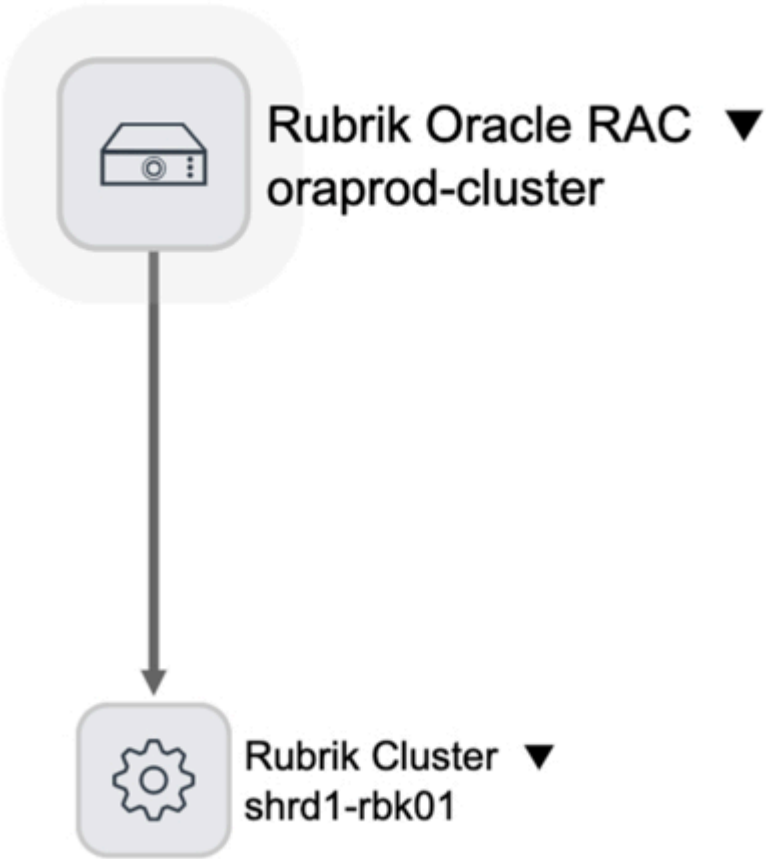

Resources discovered by the Rubrik Share (LP) pattern

Storage File Share [cmdb_ci_storage_fileshare]

Field	Description
Name [name]	ID of the host share.
Path [path]	Export point of the host share.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Host ID [host_id]	ID of the host associated with the host share.
Share type [share_type]	Share type of the host share.
Vendor type [vendor_type]	Vendor type of the host share.

The Dependency Views map shows discovered host shares and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

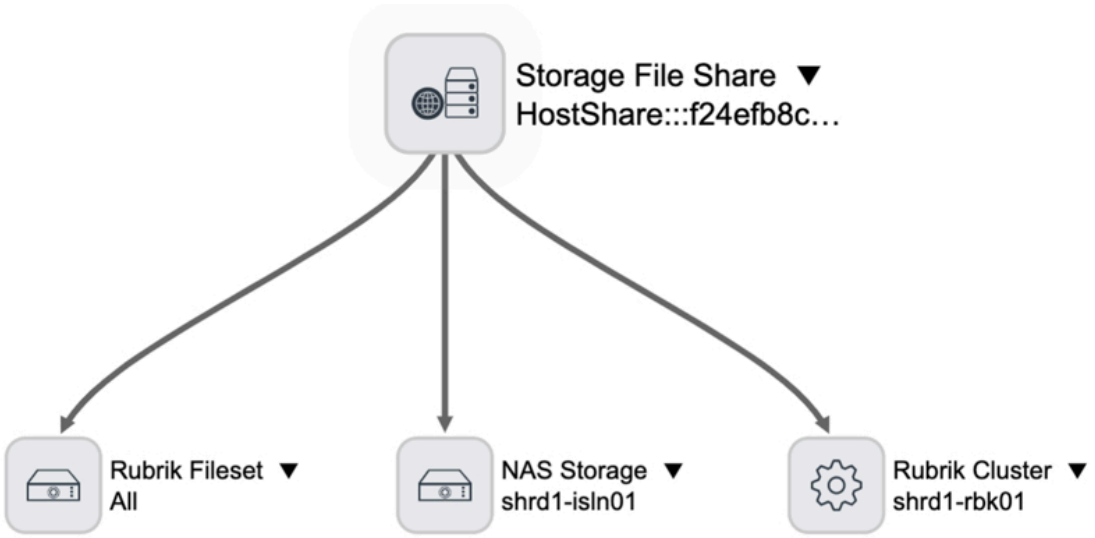

Resources discovered by the Rubrik SLA Domain (LP) pattern

Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]

Field	Description
Object ID [object_id]	ID of the Rubrik SLA domain policy.
Name [name]	Name of the Rubrik SLA domain policy.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Max Local Retention limit [max_local_retention_limit]	Maximum local retention limit of the SLA domain policy.

The Dependency Views map shows discovered SLA domains and the relationships between them.

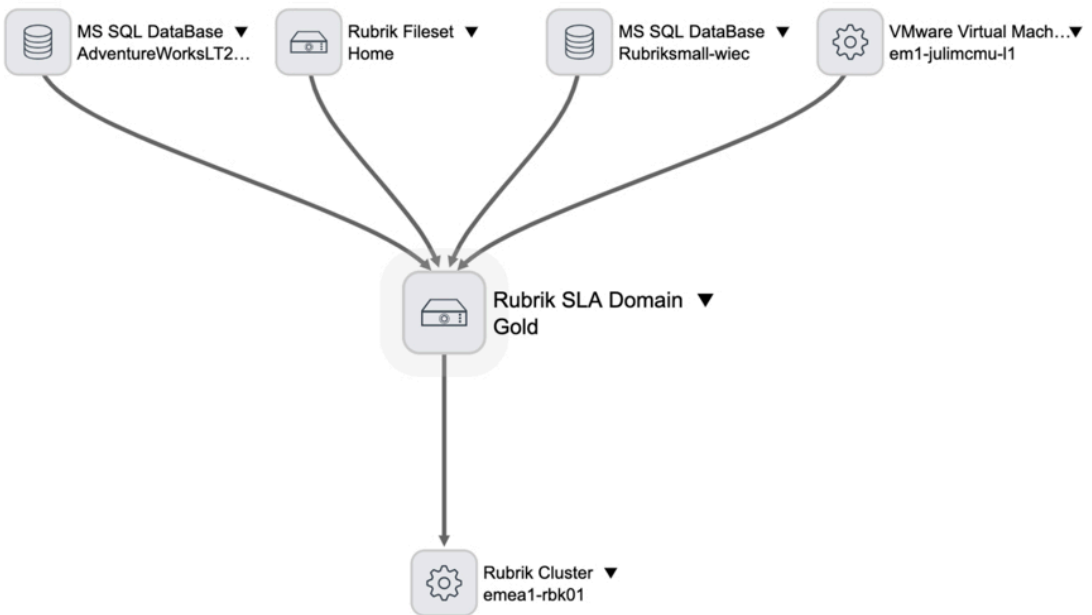

Resources discovered by the Rubrik Solaris Host (LP) pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solaris Server [cmdb_ci_solaris_server]

Field	Description
Name [name]	Name of the host.
Host name [host_name]	Host name of the host.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Operating System [os]	Operating system type configured on the host.
OS Version [os_version]	Operating system version, edition, or distribution configured on the host.
Operational status [operational_status]	Current operational status of the host.

Resources discovered by the Rubrik VMware Instance (LP) pattern

Note: Discovery of VMware instances only updates existing CIs. This is because the regular vCenter discovery does not use the IRE (Identification and Reconciliation Engine) and the Rubrik VMware pattern will create duplicates. If there is no match between the data from the REST API and the CMDB, no CIs are updated. In addition, when there is no match, the Rubrik VMware Instance (LP) pattern completes in an 'Error' state, which is expected behavior.

VMware Virtual Machine Instance [cmdb_ci_vmware_instance]

Field	Description
Object ID [object_id]	ID of the VMware instance.
Name [name]	Name of the VMware instance.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
Host ID [host_id]	ID of the VMware host.
Host name [host_name]	Name of the VMware host.
VM Cluster name [vm_cluster_name]	Name of the cluster to which the virtual machine instance belongs.
IP Address [ip_address]	IP address of the VMware instance, if assigned.
Guest OS Name [guest_os_fullname]	Guest operating system type.
State [state]	State of the VMware instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Dependency Views map shows discovered VMware instances and the relationships between them.

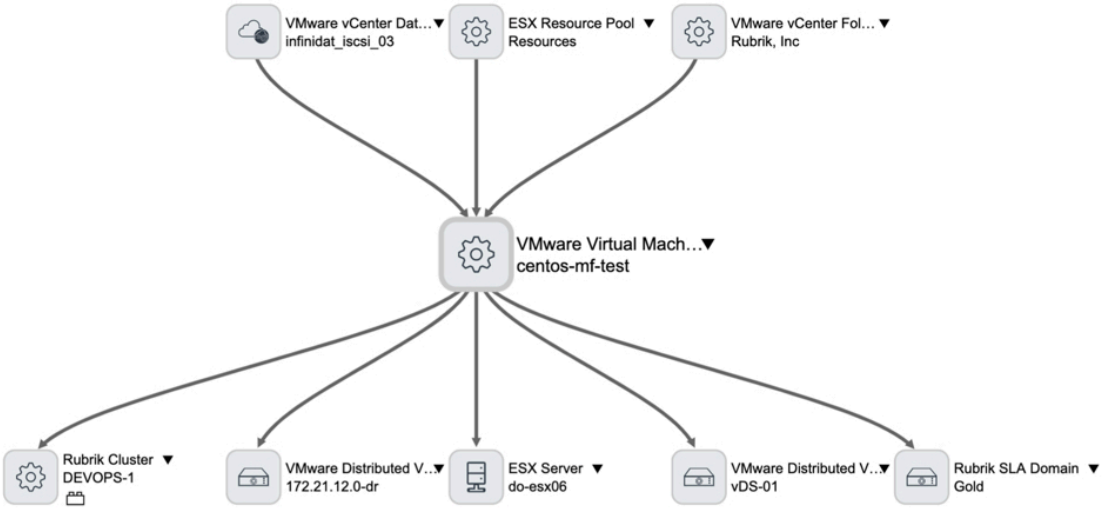

Resources discovered by the Rubrik Volume Group (LP) pattern

Rubrik Volume Group [cmdb_ci_rubrik_volume_group]

Field	Description
Object ID [object_id]	ID of the volume group.
Volume ID [volume_id]	Volume ID of the volume group.
Name [name]	Name of the volume group.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Effective SLA domain [effective_sla_domain]	ID of the Rubrik SLA Domain.
Effective SLA domain name [effective_sla_domain_name]	Name of the Rubrik SLA Domain.
Host ID [host_id]	ID of the host associated with the volume group.

The Dependency Views map shows discovered volume groups and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

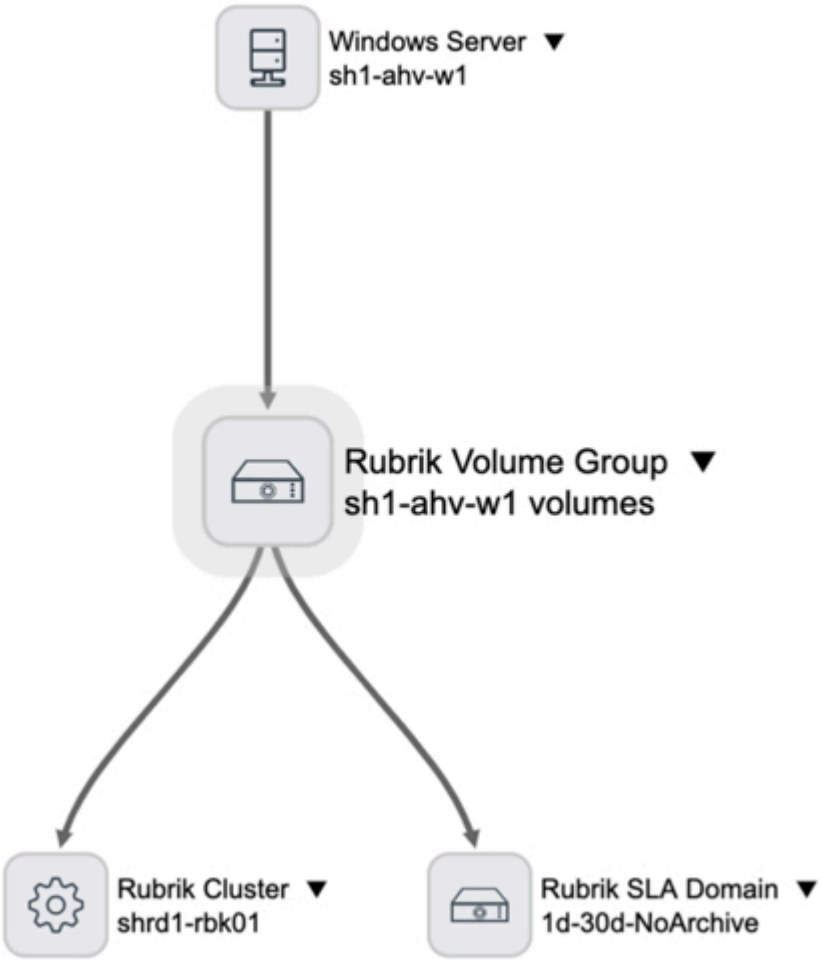

Resources discovered by the Rubrik Windows Host (LP) pattern

Windows Server [cmdb_ci_win_server]

Field	Description
Name [name]	Name of the host.
Host name [host_name]	Host name of the host.
Cluster ID [cluster_id]	ID of the Rubrik cluster.
Cluster Name [cluster_name]	Name of the Rubrik cluster.
Operating System [os]	Operating system type configured on the host.
OS Version [os_version]	Operating system version, edition, or distribution configured on the host.
Operational status [operational_status]	Current operational status of the host.

The Dependency Views map shows discovered hosts and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

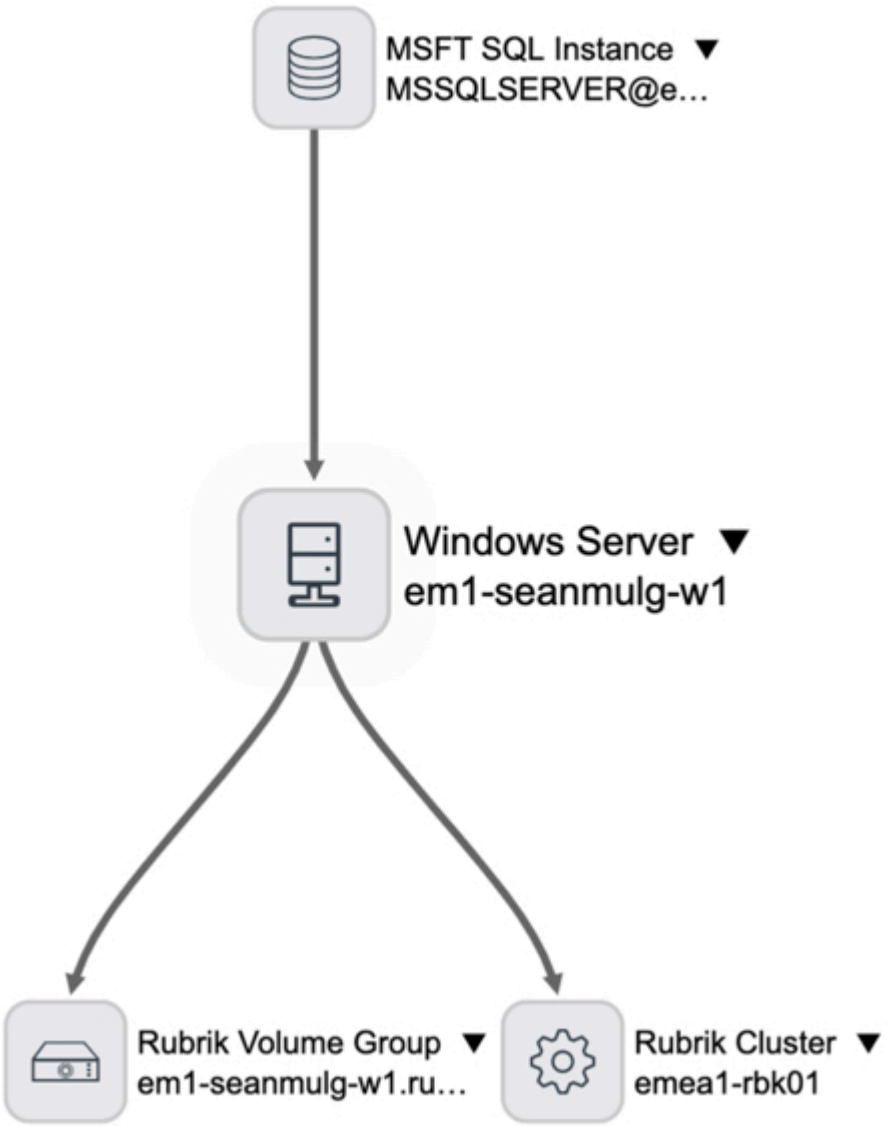

CI relationships

Relationships discovered using the Rubrik AIX Host (LP) pattern

CI	Relationship	CI
AIX Server [cmdb_ci_aix_server]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik EC2 Instance (LP) pattern

CI	Relationship	CI
Rubrik Virtual Machine Instance [cmdb_ci_rubrik_vm_instance]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Rubrik Virtual Machine Instance [cmdb_ci_rubrik_vm_instance]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships discovered using the Rubrik Fileset (LP) pattern

CI	Relationship	CI
Server [cmdb_ci_server]	Owns::Owned by	Rubrik Fileset [cmdb_ci_rubrik_fileset]
Storage File Share [cmdb_ci_storage_fileshare]	Owns::Owned by	Rubrik Fileset [cmdb_ci_rubrik_fileset]
Rubrik Fileset Template [cmdb_ci_rubrik_fileset_template]	Defines resources for::Gets resources from	Rubrik Fileset [cmdb_ci_rubrik_fileset]
Rubrik Fileset [cmdb_ci_rubrik_fileset]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Rubrik Fileset [cmdb_ci_rubrik_fileset]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]
NAS Storage [cmdb_ci_nas_storage]	Owns::Owned by	Rubrik Fileset [cmdb_ci_rubrik_fileset]

Relationships discovered using the Rubrik Fileset Template (LP) pattern

CI	Relationship	CI
Rubrik Fileset Template [cmdb_ci_rubrik_fileset_template]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik Hyper-V Instance (LP) pattern

CI	Relationship	CI
Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]

Relationships discovered using the Rubrik Linux Host (LP) pattern

CI	Relationship	CI
Linux Server [cmdb_ci_linux_server]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik Managed Volume (LP) pattern

CI	Relationship	CI
Storage Volume [cmdb_ci_storage_volume]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Storage Volume [cmdb_ci_storage_volume]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships discovered using the Rubrik MS SQL DB (LP) pattern

CI	Relationship	CI
MS SQL DataBase [cmdb_ci_db_mssql_database]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
MS SQL DataBase [cmdb_ci_db_mssql_database]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]
MSFT SQL Instance [cmdb_ci_db_mssql_instance]	Contains::Contained by	MS SQL DataBase [cmdb_ci_db_mssql_database]

Relationships discovered using the Rubrik MS SQL Instance (LP) pattern

CI	Relationship	CI
MSFT SQL Instance [cmdb_ci_db_mssql_instance]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
MSFT SQL Instance [cmdb_ci_db_mssql_instance]	Runs on::Runs	Server [cmdb_ci_server]

Relationships discovered using the Rubrik NAS Host (LP) pattern

CI	Relationship	CI
NAS Storage [cmdb_ci_nas_storage]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik Node (LP) pattern

CI	Relationship	CI
Rubrik Node [cmdb_ci_rubrik_node]	Cluster of::Cluster	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Rubrik Node [cmdb_ci_rubrik_node]	Uses::Used by	Rubrik Node Disk [cmdb_ci_rubrik_node_disk]

Relationships discovered using the Rubrik Nutanix Instance (LP) pattern

CI	Relationship	CI
Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Nutanix Virtual Machine Instance [cmdb_ci_nutanix_vm_instance]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]

Relationships discovered using the Rubrik Oracle DB (LP) pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Oracle Instance [cmdb_ci_db_ora_instance]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Oracle Instance [cmdb_ci_db_ora_instance]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]
Oracle Instance [cmdb_ci_db_ora_instance]	Runs on::Runs	Server [cmdb_ci_server]
Oracle Instance [cmdb_ci_db_ora_instance]	Cluster of::Cluster	Rubrik Oracle RAC [cmdb_ci_rubrik_db_ora_rac]

Relationships discovered using the Rubrik Oracle RAC (LP) pattern

CI	Relationship	CI
Rubrik Oracle RAC [cmdb_ci_rubrik_db_ora_rac]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik Share (LP) pattern

CI	Relationship	CI
Storage File Share [cmdb_ci_storage_fileshare]	Managed by::Manages	NAS Storage [cmdb_ci_nas_storage]
Storage File Share [cmdb_ci_storage_fileshare]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik SLA Domain (LP) pattern

CI	Relationship	CI
Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik Solaris Host (LP) pattern

CI	Relationship	CI
Solaris Server [cmdb_ci_solaris_server]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Relationships discovered using the Rubrik VMware Instance (LP) pattern

CI	Relationship	CI
VMware Virtual Machine Instance [cmdb_ci_vmware_instance]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
VMware Virtual Machine Instance [cmdb_ci_vmware_instance]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships discovered using the Rubrik Volume Group (LP) pattern

CI	Relationship	CI
Rubrik Volume Group [cmdb_ci_rubrik_volume_group]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]
Rubrik Volume Group [cmdb_ci_rubrik_volume_group]	Protected by::Protects	Rubrik SLA Domain [cmdb_ci_rubrik_sla_domain]
Server [cmdb_ci_server]	Owns::Owned by	Rubrik Volume Group [cmdb_ci_rubrik_volume_group]

Relationships discovered using the Rubrik Windows Host (LP) pattern

CI	Relationship	CI
Windows Server [cmdb_ci_win_server]	Managed by::Manages	Rubrik Cluster [cmdb_ci_rubrik_cluster]

Create a basic authentication credential alias for Rubrik discovery

Create an alias and add it to a basic authentication credential to discover Rubrik clusters.

Before you begin

Confirm that a basic authentication account is configured with the Read-Only Administrator role in Rubrik.

Role required: discovery_admin

Procedure

1. Create an alias.
   a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   b. Select **New**.
   c. Enter a unique name for the alias and select **Credential** for the alias type in the **Type** field.
   d. Select **Submit**.
2. Configure a basic authentication credential for the new alias.
   a. Navigate to **Discovery** > **Credentials**.
   b. Select **New**.
   c. Select **Basic Auth Credentials**.
   d. Unlock the **Credential alias** field and select the alias you created.
   e. Configure the rest of the **Basic Auth Credentials** form fields.
   For more information, see Basic authentication credentials.
   f. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Create a serverless discovery schedule. For more information, see Create a serverless discovery schedule for Rubrik cluster discovery.

Related topics

Rubrik Cluster discovery

Create a serverless discovery schedule for Rubrik cluster discovery

Set up a dedicated discovery schedule for each Rubrik cluster (Brik) to identify cluster resources using a serverless pattern and credential alias.

Before you begin

- Verify the MID Server is set to Active and can reach the target Rubrik cluster.
- Create an alias for the basic authentication credential. For more information, see Create a basic authentication credential alias for Rubrik discovery.
- Obtain the IP address of one of the Rubrik nodes in the target cluster.

Role required: discovery_admin

About this task

Each Rubrik cluster (Brik) requires a separate serverless discovery schedule. Each schedule must include the IP address of one of the Rubrik nodes in that cluster, and the credential alias associated with the Rubrik basic authentication credential record.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule form

Field	Description
Name	Unique name for this discovery schedule. For example, `Discover Rubrik`.
Discover	Scan type. Select **Serverless**.
MID server	Name of the MID Server to use for this schedule.
Active	When selected, activates this schedule for discovery runs.

   c. Select **Submit**.
3. Create the execution pattern.
   a. In the Discovery Schedules page, select the record you created.
   b. In the **Serverless Execution Patterns** tab, select **New**.
   c. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Serverless Execution Pattern form

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern to use for this schedule. Select the Rubrik cluster (LP) pattern.
Run Child Patterns	Set to **true** for each Rubrik discovery pattern. All child patterns associated with the Rubrik discovery pattern run automatically. This behavior is required by the Large Payload (LP) mechanism, which discovers the Rubrik cluster and all corresponding CI types.
Active	When selected, activates this schedule for discovery runs.

d. Select **Submit**.

4. Set the pattern launcher parameters.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Pattern launcher parameters

Parameter	Value
ip_address	IP address of one of the Rubrik nodes in the cluster.
credential_alias	ID of the credential alias you created for the Rubrik Basic Auth credential record.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

Rubrik Cluster discovery

Veritas Cluster Server discovery

The ServiceNow Discovery application uses the Unix Cluster – VERITAS Cluster pattern to find Veritas Cluster Server components. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

You can use this pattern on the ServiceNow AI Platform using Kingston Patch 8, London, or Madrid.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

SSH credentials

On the ServiceNow AI Platform®, configure the SSH credentials for a user with elevated rights for running the *cat* command. For more information about commands, see Service Mapping commands requiring a privileged user under Service Mapping documentation.

Permissions to read the Veritas Cluster Server configuration file

Give the Unix OS user permissions to read the following configuration file: `/etc/VRTSvcs/conf/config/main.cf`.

Permissions to run Veritas Cluster commands

Give the Unix OS user permissions to run the following commands against the Veritas Cluster Server:

Command	Parameter	Mandatory/Optional	Description
“privilege user” + cat /etc/VRTSvcs/conf/config/main.cf	-	Mandatory	Read the Veritas Cluster Server configuration file.
“privilege user” + /opt/VRTSvcs/bin/haclus -value ClusterUUID 2>/dev/null	-	Mandatory	Get the Veritas Cluster Server cluster UUID.
“privilege user” + /opt/VRTSvcs/bin/haclus -value ClusterName 2>/dev/null	-	Mandatory	Get the Veritas Cluster Server cluster name.
“privilege user” + /opt/VRTSvcs/bin/haclus -value EngineVersion	-	Mandatory	Get the Veritas Cluster Server cluster version.
“privilege user” + /opt/VRTSvcs/bin/haclus -value ClusterAddress	-	Mandatory	Get the Veritas Cluster Server cluster address.
“privilege user” + /opt/VRTSvcs/bin/haclus -value ClusState	-	Mandatory	Get the Veritas Cluster Server cluster status.
“privilege user” + /opt/VRTSvcs/bin/hasys -state	-	Mandatory	Get the Veritas Cluster Server cluster nodes.
“privilege user” + /opt/VRTSvcs/bin/hares -state	-	Mandatory	Get the cluster resources of the Veritas Cluster Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Parameter	Mandatory/Optional	Description
“privilege user” + /opt/VRTSvcs/bin/hares -display | grep -w 'Type' |grep 'global' 2>/dev/null	-	Mandatory	Get the cluster resource type of the Veritas Cluster Server.
“privilege user” + /opt/VRTSvcs/bin/hares -display | grep Group 2>/dev/null	-	Mandatory	Get the cluster resource group of the Veritas Cluster Server.
“privilege user” + /opt/VRTSvcs/bin/hagrp -state 2>/dev/null	-	Mandatory	Get the resource group name and status of the Veritas Cluster Server.

Classification probe for triggering the UNIX Cluster – VERITAS Cluster pattern

Verify that the classification probe is set to trigger the UNIX Cluster – VERITAS Cluster pattern:

1. Navigate to **Discovery Definition** > **CI Classification** > **CI Classification** > **UNIX**.
2. In the **UNIX Classification** list, click **Solaris** or **Linux**.
3. On the **Triggers probes** tab, check that the *HorizontalDiscoveryProbe-HorizontalPatt* probe is assigned to the *UNIX Cluster – VERITAS Cluster* pattern.

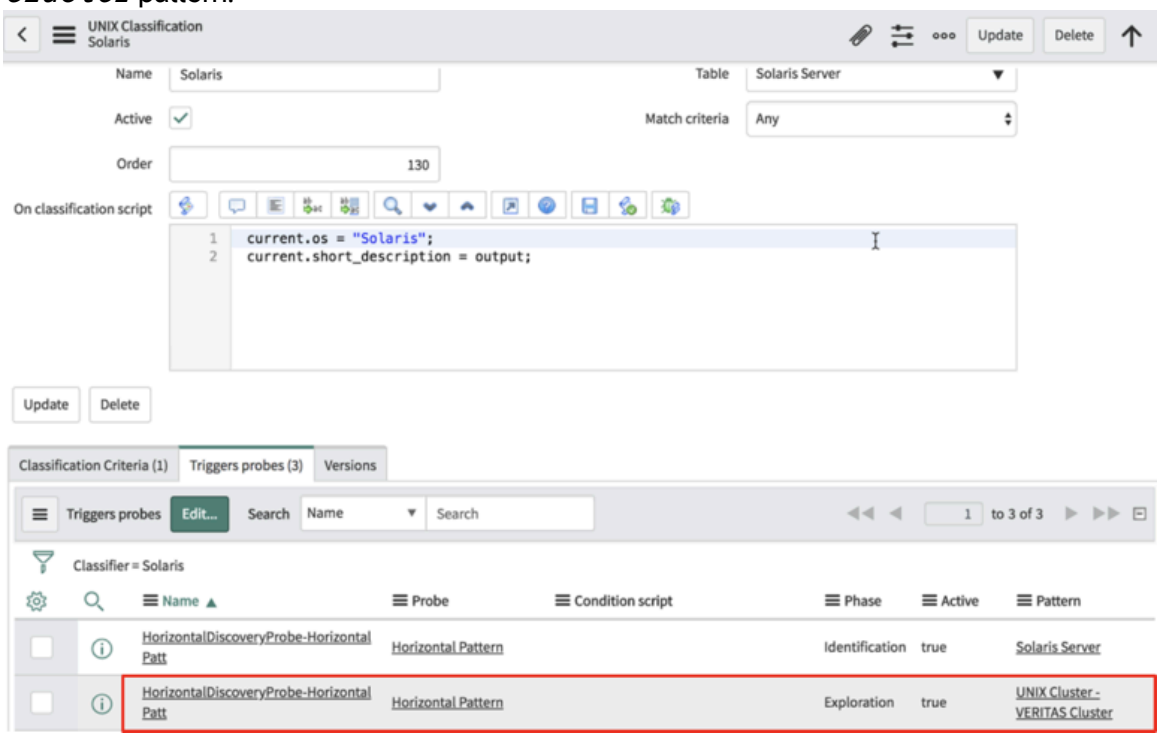

Note: The discovery log shows the error for OS discovery step even if the discovery finished successfully.

System property for the new host class

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add a system property [sys_property] *sa.host_classes* and set the value to cmdb_ci_unix_cluster. Creating a new host class for Veritas Cluster servers helps to identify this type of hosts correctly.

Limitations

You cannot customize the Unix Cluster – VERITAS Cluster pattern in the Debug mode in the Pattern Designer.

Data collected by Discovery during horizontal discovery

Table and field	Description
Unix Cluster [cmdb_ci_unix_cluster]	The attributes of the Unix Cluster.
IP address [ip_address]	
Cluster type [cluster_type]	
Name [name]	
Cluster version [cluster_version]	
Cluster status [cluster_status]	
Unix Cluster Node [cmdb_ci_unix_cluster_node]	The attributes of the Unix Cluster Node.
Name [name]	
Cluster [cluster]	
Server [server]	
Node state [node_status]	
IP address [ip_address]	
UNIX Cluster Resource Group [cmdb_ci_unix_cluster_resource_group]	The attributes of the UNIX Cluster Resource Group.
Name [name]	
Node [node]	
Server [server]	
Cluster [cluster]	
Resource group status [resource_group_status]	
UNIX Cluster Resource [cmdb_ci_unix_cluster_resource]	The attributes of the UNIX Cluster Resource.
Resource type [resource_type]	
Name [name]	
Resource status [resource_status]	
Cluster Virtual IPs [cmdb_ci_cluster_vip]	The attributes of the Cluster Virtual IP addresses.
IP address [ip_address]	
name [name]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Cluster [cluster]	
Cluster status [cluster_status]	
Node [node]	

The graphic illustrates CIs that are part of Veritas Cluster Server discovery.

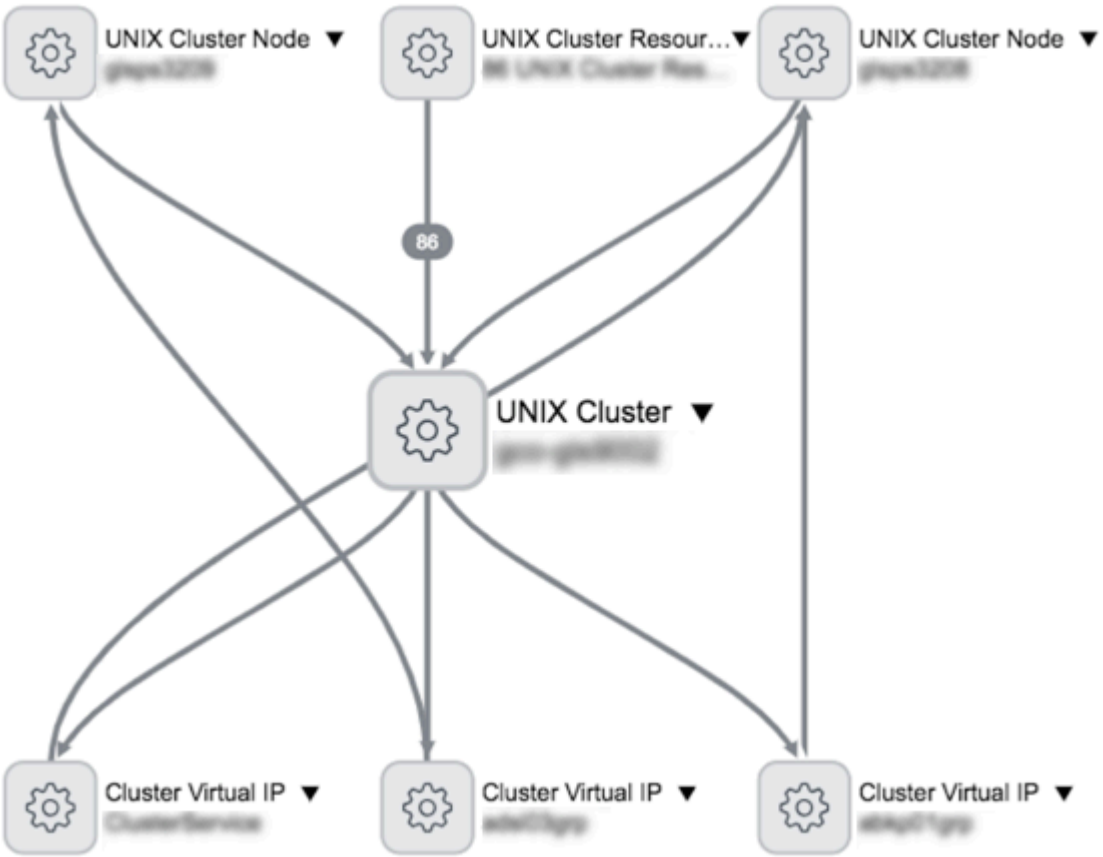

CI relationships

The Unix Cluster – VERITAS Cluster pattern creates the following CI relationships:

CI	Relationship	CI
Unix Cluster [cmdb_ci_unix_cluster]	Hosts: Hosted on	Linux Server [cmdb_ci_linux_server]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Hosts: Hosted on	Linux Server [cmdb_ci_linux_server]
Cluster Virtual IPs [cmdb_ci_cluster_vip]	Virtualized by: Virtualized	Unix Cluster [cmdb_ci_unix_cluster]
Cluster Virtual IPs [cmdb_ci_cluster_vip]	Uses: Used by	Unix Cluster Node [cmdb_ci_unix_cluster_node]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Cluster of: Cluster	Unix Cluster [cmdb_ci_unix_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Unix Cluster resource [cmdb_ci_unix_cluster_resource]	Defines resources for: Gets resources from	Unix Cluster Node [cmdb_ci_unix_cluster_node]
Unix Cluster resource [cmdb_ci_unix_cluster_resource]	Defines resources for: Gets resources from	The Cluster field on [cmdb_ci_unix_cluster]
Unix Cluster resource group [cmdb_ci_unix_cluster_resource_group]	Contains: Contained by	The Cluster field on UNIX Cluster [cmdb_ci_unix_cluster]
Unix Cluster resource group [cmdb_ci_unix_cluster_resource_group]	Contains: Contained by	The Node field on Unix Cluster Node [cmdb_ci_unix_cluster_node]
Unix Cluster resource group [cmdb_ci_unix_cluster_resource_group]	Contains: Contained by	Unix Cluster resource [cmdb_ci_unix_cluster_resource]

VMware NSX Advanced load balancer discovery

The ServiceNow Discovery application uses the VMware NSX Advanced load balancer discovery pattern to find VMware NSX load balancers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The pattern uses a set of REST APIs to find the load balancers. It is triggered by the NSX HTTP classifier that responds to the **HTTP - Classify** probe. For more information, see HTTP device discovery.

Note: There are two patterns available. The legacy VMware NSX Load Balancer pattern is for the S series. The new VMware NSX-T Load Balancer pattern is for the T series. The functionality and collected information remain the same.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Enable NSX Basic Authentication

In VMware NSX, enable NSX Basic Authentication.

Verify that the applications are up to date

- CMDB CI Class Models
- Visibility Content
- Discovery and Service Mapping Patterns

Configure basic authentication credentials

On the ServiceNow AI Platform, configure basic authentication credentials. The credentials must have the following API permissions:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- NSX:
  - /api/4.0/edges
  - /api/4.0/edges/<edge_id>loadbalancer/config
- NSX-T: /api/v1/edge-clusters

For more information about basic authentication credentials, see Basic authentication credentials.

Verify HTTP discovery is triggered against the NSX Main Controller IP

The HTTP discovery is triggered only if SSH or SNMP discovery is unsuccessful.

Create a discovery schedule

Create the discovery schedule for VMware NSX load balancer discovery or run a quick discovery on the IP address of the NSX Main Controller. For more information about running a quick discovery, see Run a Quick Discovery.

Data collected by Discovery for VMware NSX Advanced load balancer

Discovery populates the data in the CMDB when running the VMware NSX Advanced load balancer patterns.

NSX Load Balancer [cmdb_ci_nsx_lb]

Field	Description
Name [name]	Edge name.
Object ID [object_id]	Edge ID.
IP Address [ip_address]	IP address of the NSX Main Controller.
Operational status [operational_status]	Load balancer status: Active or inactive.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Name [name]	Virtual server name.
IP Address [ip_address]	Virtual server IP address.
Object ID [object_id]	Virtual server object ID.
Port [port]	Virtual server port with which to communicate.
Listener Protocol [listener_protocol]	Listener protocol. For example: HTTP, HTTPS.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Name [name]	Pool name.
Object ID [object_id]	Pool object ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Pool Member [cmdb_ci_lb_pool_member]

Field	Description
Name [name]	Pool member name.
IP Address [ip_address]	Pool member IP address.
Object ID [object_id]	Pool member object ID.
Service port [service_port]	Pool member port with which to communicate.

Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]

Field	Description
Name [name]	Monitor name.
Object ID [object_id]	Monitor object ID.
Timeout in seconds [timeout_sec]	Time in seconds to wait for a response.
Interval in seconds [check_interval_sec]	Interval in seconds after which the monitor sends the request again.
Monitor type protocol [monitor_type]	Monitor type. For example: HTTP, HTTPS, TCP.
Request path [request_path]	Path where the monitor sends the request.
Max retries [max_retries]	The maximum number of times the resource is tested before it is declared "Down".

The Dependency Views map shows all discovered load balancer CIs and the relationships between them.

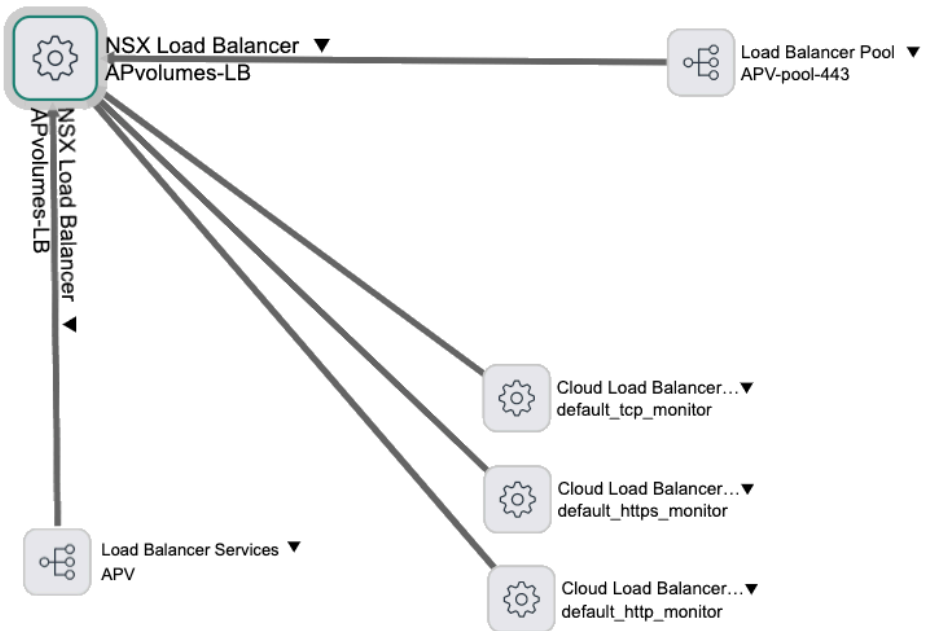

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

Discovery creates these relationships to support the VMware NSX load balancer discovery.

CI	Relationship	CI
Load Balancer Service [cmdb_ci_lb_service]	Runs on::Runs	NSX Load Balancer [cmdb_ci_nsx_lb]
Load Balancer Service [cmdb_ci_lb_service]	Uses::Used by	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Pool [cmdb_ci_lb_pool]	Hosted on::Hosts	NSX Load Balancer [cmdb_ci_nsx_lb]
Load Balancer Pool [cmdb_ci_lb_pool]	Owns::Owned by	Load Balancer Pool Member [cmdb_ci_lb_pool_member]
Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]	Hosted on::Hosts	NSX Load Balancer [cmdb_ci_nsx_lb]
Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]	Used by::Uses	Load Balancer Pool [cmdb_ci_lb_pool]

Troubleshooting

If the mapping process does not proceed as you expected, follow the following suggestions.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again. **Note:** If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware NSX-T cluster pattern-based discovery

Discovery and Service Mapping Patterns uses the NSX Cluster pattern to find VMware NSX-T infrastructure. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

NSX Cluster data model

The following diagram illustrates the tables and relationships that the Discovery and Service Mapping Patterns application creates when discovering NSX-T infrastructure resources.

Tables and relationships populated during NSX-T cluster discovery

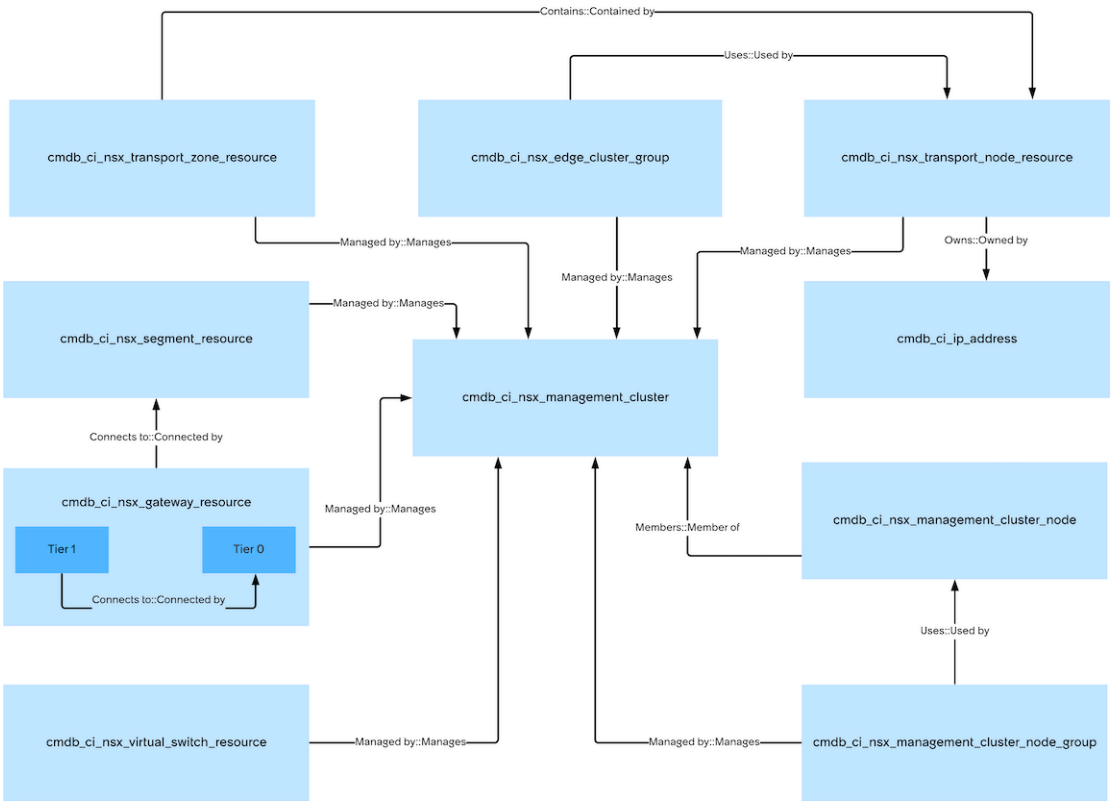

Prerequisites

Verify that the following applications are up to date

- CMDB CI Class Models
- Discovery and Service Mapping Patterns
- Visibility Content

Configure a credential alias

For more information, see Create an applicative credential alias for NSX-T cluster discovery.

Verify NSX-T API access and permissions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify that the credential has sufficient permissions to execute the following NSX-T API calls:

- `POST /api/session/create`
- `GET /api/v1/cluster/status`
- `GET /api/v1/node`
- `GET /policy/api/v1/infra/segments`
- `GET /policy/api/v1/infra/tier-0s`
- `GET /policy/api/v1/infra/tier-1s`
- `GET /api/v1/transport-nodes`
- `GET /api/v1/transport-zones`
- `GET /api/v1/fabric/virtual-switches/`
- `GET /api/v1/edge-clusters`

Create a serverless discovery schedule

For more information, see Create a serverless schedule for NSX-T cluster discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the NSX Cluster pattern.

NSX Management Cluster [cmdb_ci_nsx_management_cluster]

Field	Description
Name [name]	IP address of the NSX Management Cluster.
IP Address [ip_address]	IP address of the NSX Management Cluster.
Cluster Status [cluster_status]	Status of the management cluster as reported by the NSX-T API.
Operational status [operational_status]	Operational status of the resource, derived from the cluster status.
Install Status [install_status]	Install status of the resource. Default value is Installed.

NSX Management Cluster Node [cmdb_ci_nsx_management_cluster_node]

Field	Description
Name [name]	Display name of the cluster node, derived from the node's fully qualified domain name (FQDN).
ID [nsx_node_id]	Unique identifier of the NSX Manager node.
Fully qualified domain name [fqdn]	FQDN of the node.
Node Status [node_status]	Status of the node as reported by the NSX-T API.
Operational status [operational_status]	Operational status of the resource, derived from the node status.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NSX Management Cluster Node Group [cmdb_ci_nsx_management_cluster_node_group]

Field	Description
Name [name]	Concatenation of the cluster IP address and the group type. For example: `10.0.0.1 MANAGER`.
Object ID [object_id]	Group ID assigned by NSX-T.
Operational status [operational_status]	Operational status of the resource, derived from the group status.
Install Status [install_status]	Install status of the resource. Default value is Installed.

NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resource]

Field	Description
Name [name]	Display name of the transport node, derived from the host FQDN or virtual machine (VM) display name.
Object ID [object_id]	Unique identifier of the transport node.
Host Type [host_type]	Type of the transport node host. For example: ESXi, KVM, or Edge.
Host Name [host_name]	Host name of the transport node.
Host ID [host_id]	Identifier of the host running the transport node.
Fully qualified domain name [fqdn]	FQDN of the transport node.
IP Address [ip_address]	IP address of the transport node.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Description [short_description]	User-provided description of the transport node.

NSX Transport Zone Resource [cmdb_ci_nsx_transport_zone_resource]

Field	Description
Name [name]	Display name of the transport zone.
Object ID [object_id]	Unique identifier of the transport zone.
Transport Type [transport_type]	Transport type of the zone. For example: OVERLAY or VLAN.
Forwarding Mode [forwarding_mode]	Forwarding mode of the transport zone.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NSX Transport Zone Resource [cmdb_ci_nsx_transport_zone_resource] (continued)

Field	Description
Is Default [is_default]	Indicates whether this is the default transport zone.
Operational status [operational_status]	Operational status of the resource. The value is set to Operational.
Install Status [install_status]	Install status of the resource. The value is set to Installed.
Description [short_description]	User-provided description of the transport zone.

NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group]

Field	Description
Name [name]	Display name of the edge cluster group.
Object ID [object_id]	Unique identifier of the edge cluster.
Node Type [node_type]	Type of edge node in the cluster.
Inter Site Forwarding [enable_inter_site_forwarding]	Indicates whether inter-site forwarding is enabled for this edge cluster.

NSX Virtual Switch Resource [cmdb_ci_nsx_virtual_switch_resource]

Field	Description
Name [name]	Display name of the virtual switch.
Object ID [object_id]	Unique identifier of the virtual switch.
CM Local ID [cm_local_id]	Local identifier assigned by the compute manager.
Origin ID [origin_id]	Origin identifier of the virtual switch resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

NSX Gateway Resource [cmdb_ci_nsx_gateway_resource]

Field	Description
Name [name]	Display name of the gateway.
Object ID [object_id]	Unique identifier of the gateway.
Tier Type [tier_type]	Tier type of the gateway: Tier-0 or Tier-1.
HA Mode [ha_mode]	High-availability mode of the gateway.
Failover Mode [failover_mode]	Failover mode of the gateway.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Description [short_description]	User-provided description of the gateway resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NSX Segment Resource [cmdb_ci_nsx_segment_resource]

Field	Description
Name [name]	Display name of the segment.
Object ID [object_id]	Unique identifier of the segment.
Operational status [operational_status]	Operational status of the resource, derived from the resource state.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Description [short_description]	User-provided description of the segment resource.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address of the transport node.

Dependency Views map

On the Dependency Views map, you can view discovered NSX-T cluster resources and the relationships between them.

NSX Management Cluster CI dependency view

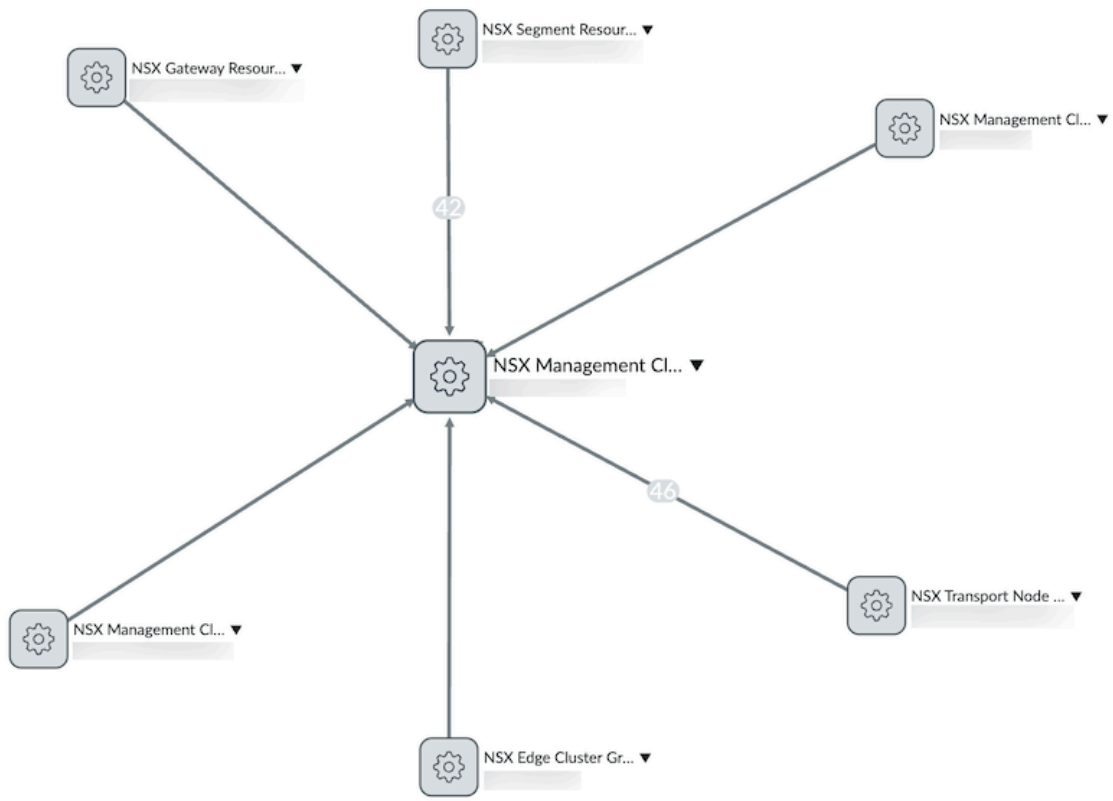

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships and references

The NSX Cluster pattern creates the following relationships and references to support NSX-T cluster discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
NSX Management Cluster Node [cmdb_ci_nsx_management_cluster_node]	Members::Member of	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Gateway Resource [cmdb_ci_nsx_gateway_resource]	Managed by::Manages	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Gateway Resource [cmdb_ci_nsx_gateway_resource]	Connects to::Connected by	NSX Segment Resource [cmdb_ci_nsx_segment_resource]
NSX Gateway Resource [cmdb_ci_nsx_gateway_resource] (Tier 1)	Connects to::Connected by	NSX Gateway Resource [cmdb_ci_nsx_gateway_resource] (Tie
NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resource]	Managed by::Manages	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resource]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
NSX Management Cluster Node Group [cmdb_ci_nsx_management_cluster_node_group]	Managed by::Manages	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Management Cluster Node Group [cmdb_ci_nsx_management_cluster_node_group]	Uses::Used by	NSX Management Cluster Node [cmdb_ci_nsx_management_cluster_
NSX Segment Resource [cmdb_ci_nsx_segment_resource]	Managed by::Manages	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Transport Zone Resource [cmdb_ci_nsx_transport_zone_resource]	Managed by::Manages	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Transport Zone Resource [cmdb_ci_nsx_transport_zone_resource]	Contains::Contained by	NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resou
NSX Virtual Switch Resource [cmdb_ci_nsx_virtual_switch_resource]	Managed by::Manages	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group]	Managed by::Manages	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group]	Uses::Used by	NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resou

CI references

CI	Field	Referenced CI
NSX Management Cluster Node [cmdb_ci_nsx_management_cluster_node]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Management Cluster Node Group [cmdb_ci_nsx_management_cluster_node_group]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
NSX Segment Resource [cmdb_ci_nsx_segment_resource]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Gateway Resource [cmdb_ci_nsx_gateway_resource]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resource]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Transport Zone Resource [cmdb_ci_nsx_transport_zone_resource]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Virtual Switch Resource [cmdb_ci_nsx_virtual_switch_resource]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group]	Cluster [cluster]	NSX Management Cluster [cmdb_ci_nsx_management_cluster]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	References one of the following NSX resource tables: • NSX Segment Resource [cmdb_ci_nsx_segment_resource] • NSX Gateway Resource [cmdb_ci_nsx_gateway_resource] • NSX Transport Node Resource [cmdb_ci_nsx_transport_node_reso • NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group]
External system metadata [cmdb_key_value_v2]	Configuration item [configuration_item]	References one of the following NSX resource tables: • NSX Segment Resource [cmdb_ci_nsx_segment_resource] • NSX Gateway Resource [cmdb_ci_nsx_gateway_resource] • NSX Transport Node Resource [cmdb_ci_nsx_transport_node_reso • NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group] • NSX Transport Zone Resource [cmdb_ci_nsx_transport_zone_resou • NSX Virtual Switch Resource [cmdb_ci_nsx_virtual_switch_resour

Tag discovery and external system metadata

The NSX Cluster pattern collects resource tags and external system metadata. Tags are stored in the Key Value [cmdb_key_value] table. External system metadata is stored in the External system metadata [cmdb_key_value_v2] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the NSX resource table associated with the tag. Possible tables: • NSX Segment Resource [cmdb_ci_nsx_segment_resource] • NSX Gateway Resource [cmdb_ci_nsx_gateway_resource] • NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resource] • NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group]

External system metadata [cmdb_key_value_v2]

Field	Description
Key [key]	Metadata property name.
String value [string_value]	Value of the metadata property.
Value type [value_type]	Data type of the metadata value.
Configuration item [configuration_item]	References the NSX resource table associated with the metadata entry. Possible tables: • NSX Segment Resource [cmdb_ci_nsx_segment_resource] • NSX Gateway Resource [cmdb_ci_nsx_gateway_resource] • NSX Transport Node Resource [cmdb_ci_nsx_transport_node_resource] • NSX Edge Cluster Group [cmdb_ci_nsx_edge_cluster_group] • NSX Transport Zone Resource [cmdb_ci_nsx_transport_zone_resource] • NSX Virtual Switch Resource [cmdb_ci_nsx_virtual_switch_resource]

Create an applicative credential alias for NSX-T cluster discovery

Create a credential alias and configure an applicative credential to enable the NSX Cluster pattern to authenticate with the NSX-T management cluster.

Before you begin

Role required: discovery_admin

About this task

The NSX Cluster pattern requires a credential alias that contains exactly one applicative credential record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create an alias.
   a. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
   b. Select **New**.
   c. Enter a unique name for the alias and select **Credential** for the alias type in the **Type** field.
   d. Select **Submit**.
2. Configure an applicative credential for the alias.
   a. Navigate to **Discovery** > **Credentials**.
   b. Select **New**.
   c. Select **Applicative Credentials**.
   d. Fill in the fields.

Applicative Credentials form fields

Field	Value
Name	Descriptive name for the credential. For example: NSX-T admin.
User name	NSX-T user name with permissions to execute the required API calls.
Password	Password for the NSX-T user.
CI type	Cluster [cmdb_ci_cluster].

   e. Unlock the **Credential alias** field and select the alias you created.
   f. Select **Submit**.

What to do next

Create a serverless discovery schedule for the NSX Cluster pattern. For more information, see Create a serverless schedule for NSX-T cluster discovery.

Create a serverless schedule for NSX-T cluster discovery

Create a serverless discovery schedule to run the NSX Cluster pattern against a VMware NSX-T management cluster.

Before you begin

- Verify that you have at least version 1.31.0 of Discovery and Service Mapping Patterns.
- Verify that you have the IP address of the NSX-T management cluster.
- Create a credential alias with an applicative credential. For more information, see Create an applicative credential alias for NSX-T cluster discovery.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule New record form

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type for this schedule. Select **Serverless**.
MID server	Name of the MID Server to use for this schedule.
Active	Setting that activates this schedule for discovery.

   c. Select **Submit**.
3. Create the execution pattern.
   a. In the Discovery Schedules page, select the record you created.
   b. In the **Serverless Execution Patterns** tab, select **New**.
   c. On the form, fill in the fields.

Serverless Execution Pattern New record form

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern for this schedule. Select the NSX Cluster pattern.
Proxy Host	Fully qualified domain name of the proxy server host. Select **Global**.
Active	Setting that activates this schedule for discovery.

   d. Select **Submit**.
4. Set the pattern launcher parameters.
   a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.
   b. On the form, fill in the fields.

Pattern launcher parameters

Parameter	Value
ip_address	IP address of the NSX-T management cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Value
cred_alias	Name of the credential alias you created.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

VMware NSX-T cluster pattern-based discovery

Available cloud discovery patterns

ITOM Visibility comes with an extensive library of patterns that can discover your cloud environment: AWS, Azure, GCP, IBM, and OCI.

ServiceNow releases new discovery patterns on the ServiceNow Store on a monthly basis to ensure that your organization can discover the latest industry-standard devices and applications. Major ServiceNow versions incorporate patterns previously released on the ServiceNow Store.

For a complete list of devices and applications that ITOM Visibility can discover, see Detailed information on products discovered by ITOM Visibility.

Alibaba Cloud discovery using patterns

Discovery and Service Mapping Patterns uses patterns to discover components of the Alibaba Cloud deployment during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models
- Visibility Content

Create Alibaba Cloud credentials

For more information, see Create Alibaba Cloud API Credentials.

Create an Alibaba Cloud service account

For more information, see Set up Alibaba Cloud service accounts.

Schedule an Alibaba Cloud discovery

For more information, see Create an Alibaba Cloud Discovery schedule in Discovery Admin Workspace.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Related topics

Discovery for Alibaba Cloud

Alibaba Cloud availability zone pattern-based discovery

Discovery and Service Mapping Patterns uses the Alibaba - Availability Zone (LP) pattern to discover availability zones during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Alibaba Cloud discovery prerequisites section in Alibaba Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Alibaba - Availability Zone (LP) pattern.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Zone ID of the Alibaba Cloud availability zone.
Object ID [object_id]	Zone ID of the Alibaba Cloud availability zone.
Install Status [install_status]	Install status of the availability zone. Default value is Installed.

CI relationships

Discovery creates these relationships to support the availability zone discovery.

CI	Relationship	CI
Alibaba Datacenter [cmdb_ci_alibaba_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

Alibaba Cloud cloud hardware type pattern-based discovery

Discovery and Service Mapping Patterns uses the Alibaba - Cloud Hardware Type (LP) pattern to discover cloud hardware types during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the Alibaba Cloud discovery prerequisites section in Alibaba Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Alibaba - Cloud Hardware Type (LP) pattern.

Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Object ID [object_id]	ID of the instance type.
vCPUs [vcpus]	Number of virtual central processing units (vCPUs).
Memory MB [memory_mb]	Memory size in megabytes (MB).
Local Storage GB [local_storage_gb]	Capacity of each local disk in gigabytes (GB).
Name [name]	Image name.
Provider [provider]	Image cloud provider: Alibaba.
Cores [cores]	Number of CPU cores.

CI relationships

Discovery creates these relationships to support the cloud hardware type discovery.

CI	Relationship	CI
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	Alibaba Datacenter [cmdb_ci_alibaba_datacenter]

Alibaba Cloud cloud OS image pattern-based discovery

Discovery and Service Mapping Patterns uses the Alibaba - Cloud OS Image (LP) pattern to discover cloud OS images during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Alibaba Cloud discovery prerequisites section in Alibaba Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Alibaba - Cloud OS Image (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Name [name]	Image name.
Object ID [object_id]	ID of the image. • IDs of public images contain the operating system version, architecture, language, and release date. For example: **win2008r2_64_ent_sp1_en-us_40G_alibase_20190318.vhd**. • IDs of custom, shared, and community images start with m. For example: **m-0a1b2c3defghijklm4n5o**.
Guest OS [guest_os]	Operating system (OS) type of the image. Possible values are Windows or Linux.
OS Status [status]	Status of the image. The default value is Available and this field can contain multiple values separated with a comma. For example: Creating, Available, or Unavailable. If you don't specify a parameter, only Available status images are returned.
Image Type [image_type]	Architecture of the image. Possible values: • i386 • x86_64 • arm64
Install Status [install_status]	Install status of the image. Default value is Installed.
Operational status [operational_status]	Operational status of the image. Default value is Operational.
Provider [provider]	Image cloud provider: Alibaba.

CI relationships

Discovery creates these relationships to support the Alibaba Cloud cloud OS image discovery.

CI	Relationship	CI
Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	Alibaba Datacenter [cmdb_ci_alibaba_datacenter]
Key Value [cmdb_key_value]	References	Cloud Image [cmdb_ci_cloud_os_image]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag discovery

The Alibaba - Cloud OS Image (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag key.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Image [cmdb_ci_cloud_os_image] table.

Alibaba Cloud datacenter pattern-based discovery

Discovery and Service Mapping Patterns uses the Alibaba - Datacenter (LP) pattern to discover datacenters during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Alibaba Cloud discovery prerequisites section in Alibaba Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Alibaba - Datacenter (LP) pattern.

Alibaba Datacenter [cmdb_ci_alibaba_datacenter]

Field	Description
Name [name]	**LocalName** field value of the Alibaba Cloud region.
Object ID [object_id]	**RegionId** field value of the Alibaba Cloud region.
Install Status [install_status]	Install status of the datacenter. Default value is Installed.
Operational status [operational_status]	Operational status of the datacenter. Default value is Operational.
Region [region]	**RegionId** field value of the Alibaba Cloud region.
Datacenter URL [datacenter_url]	Datacenter URL. Example URLs: • For GovCloud: **https://ecs.aliyun-inc.com** • For commercial cloud: **htpps://ecs.aliyuncs.com**

CI relationships

Discovery creates these relationships to support the Alibaba Cloud datacenter discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Alibaba Datacenter [cmdb_ci_alibaba_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Alibaba Cloud service account pattern-based discovery

Discovery and Service Mapping Patterns uses the Alibaba - Service Account Validation pattern to discover service accounts during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Alibaba Cloud discovery prerequisites section in Alibaba Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Alibaba - Service Account Validation pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Name [name]	Alibaba Cloud Resource Name (ARN) of the current requester.
Object ID [object_id]	The ID of the Alibaba Cloud account to which the current requester belongs.
Datacenter Type [datacenter_type]	Type of the datacenter where the account is hosted: Alibaba Datacenter [cmdb_ci_alibaba_datacenter].
Account Id [account_id]	The ID of the Alibaba Cloud account to which the current requester belongs.
Discovery credentials [discovery_credentials]	Alibaba Cloud API credentials.
Datacenter URL [datacenter_url]	Datacenter URL. Example URLs: • For GovCloud: **https://ecs.aliyun-inc.com** • For commercial cloud: **htpps://ecs.aliyuncs.com**

CI relationships

Discovery creates this reference to support the Alibaba Cloud service account discovery.

CI	Relationship	CI
Key Value [cmdb_key_value]	References	Cloud Service Account [cmdb_ci_cloud_service_account]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag discovery

The Alibaba - Service Account Validation pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag key.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table.

Alibaba Cloud storage volume pattern-based discovery

Discovery and Service Mapping Patterns uses the Alibaba - Storage Volume (LP) pattern to discover storage volumes during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Alibaba Cloud discovery prerequisites section in Alibaba Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Alibaba - Storage Volume (LP) pattern.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	Name of the disk.
Object ID [object_id]	ID of the disk.
State [state]	Status of the disk. Default value is All. For example: In_use, Available, or Attaching.
Size bytes [size_bytes]	Disk size in bytes.
Storage type [storage_type]	Type of the disk. Possible values are system or data.
Install Status [install_status]	Install status of the disk. Default value is Installed.
Operational status [operational_status]	Operational status of the disk. Default value is Operational.

CI relationships

Discovery creates these relationships to support the Alibaba Cloud storage volume discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	Alibaba Datacenter [cmdb_ci_alibaba_datacenter]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Virtual Machine Instance [cmdb_ci_vm_instance]	Uses::Used by	Storage Volume [cmdb_ci_storage_volume]
Key Value [cmdb_key_value]	References	Storage Volume [cmdb_ci_storage_volume]

Tag discovery

The Alibaba - Storage Volume (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag key.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Storage Volume [cmdb_ci_storage_volume] table.

Alibaba Cloud virtual machine (VM) pattern-based discovery

Discovery and Service Mapping Patterns uses the Alibaba - Virtual Machine (LP) pattern to discover VMs during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Alibaba Cloud discovery prerequisites section in Alibaba Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Alibaba - Virtual Machine (LP) pattern.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Name [name]	Name of the instance.
Object ID [object_id]	Instance ID.
State [state]	Status of the instance. For example: Running or Stopped.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance] (continued)

Field	Description
CPUs [cpus]	Number of Virtual Central Processing Units (vCPUs).
Memory (MB) [memory]	Memory size in megabytes (MB).
Disks [disks]	Number of local disks attached to the instance.
Guest OS Name [guest_os_fullname]	Type of the operating system (OS) of the instance. Possible values are Windows or Linux.
Install Status [install_status]	Install status of the instance. Default value is Installed.
Operational status [operational_status]	Operational status of the instance. Default value is Operational.
Network adapters [nics]	Elastic Network Interfaces (ENIs) attached to the instance.

Cloud Mgmt Network Interface [cmdb_ci_nic]

Field	Description
Name [name]	Name of the ENI.
Object ID [object_id]	ID of the ENI.
MAC Address [mac_address]	MAC address of the ENI.
Private IP [private_ip]	Private IP address of the ENI.
IP Address [ip_address]	Public IP addresses of instances.
Primary [primary]	A Boolean indicating whether a network interface is primary or secondary.
Node ID [node_id]	Instance ID.

Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Object ID [object_id]	ID of the image. • IDs of public images contain the operating system version, architecture, language, and release date. For example: **win2008r2_64_ent_sp1_en-us_40G_alibase_20190318.vhd**. • IDs of custom, shared, and community images start with m. For example: **m-0a1b2c3defghijklm4n5o**.
Provider [provider]	Image cloud provider: Alibaba.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Object ID [object_id]	ID of the instance type.
Provider [provider]	Image cloud provider: Alibaba.

CI relationships

Discovery creates these relationships to support the Alibaba Cloud virtual machine discovery.

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Implement End Point To::Implement End Point From	Cloud Mgmt Network Interface [cmdb_ci_nic]
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	Alibaba Datacenter [cmdb_ci_alibaba_datacenter]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Virtual Machine Instance [cmdb_ci_vm_instance]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Cloud Image [cmdb_ci_cloud_os_image]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	Alibaba Datacenter [cmdb_ci_alibaba_datacenter]
Key Value [cmdb_key_value]	References	Virtual Machine Instance [cmdb_ci_vm_instance]

Tag discovery

The Alibaba - Virtual Machine (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag key.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

AWS discovery using patterns

Discovery and Service Mapping Patterns uses patterns to discover components of the Amazon AWS Cloud deployment during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date:

- Discovery and Service Mapping Patterns
- CMDB CI Class Models
- Visibility Content

Update the method used for pointed discovery for the AWS CloudFormation Template (CFT) stack

If you use Cloud Provisioning and Governance, you must update the *getOperationGR(type)* method. This update enables the pointed discovery to list the resources correctly for the AWS CFT stack after provisioning. For further information about the steps required to update this method, see the Knowledge Base article KB0858437.

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Set up service accounts on the AWS Management Console

An AWS Organization is a collection of AWS accounts under a single account. In AWS Organizations, parent accounts are called management accounts. The sub-accounts that belong to a management account are called member accounts.

Note: Cloud Discovery for AWS Organizations isn't fully supported in a GovCloud isolated region.

The advantages of using management accounts in Discovery are:

Easy population of member accounts

After you configure the management account and supply the necessary credentials, you can test the connection to the account. If the test succeeds, Discovery returns a list of the member accounts in that management account. From this list, you can choose one or more member account to include in the Discovery of the management account.

(Optional) Discover member resources using dynamically acquired credentials

When you run Discovery on your cloud resources, you don't need separate credentials for each member account. The Cloud Discovery process handles credentials automatically by acquiring a temporary credential for each member via an AWS API. You can elect to use the default configuration or customize the MID Server to assume other roles for additional controls and security.

For more information, see Setup AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use IAM user policy on the AWS Management Console

To use the IAM user policy instead of credentials during discovery, configure the MID Server for AWS IAM roles. For more information, see configure the MID Server for AWS IAM roles.

To create the IAM user policy for provisioning AWS resources, see Control AWS access and permissions using policies.

Configure access to the AWS resources

To discover a single account, create an IAM account in the AWS Management Console, and verify that it has the "ReadOnlyAccess" policy applied. To discover several member or child accounts, configure the credentials as described in Access setup for AWS service accounts.

Configure a discovery schedule

Create an AWS Discovery schedule in Discovery Admin Workspace.

(Optional) Discover datacenters only for new members

Starting with Zurich Patch 2, you can discover datacenters only for new members added since the last discovery. For more information, see Discover datacenters only for new cloud accounts.

Optimize discovery by including only datacenters with resources

You can optimize discovery by limiting it to only AWS datacenters with resources.

- Verify that your service account has the following IAM permissions to access the Resource Explorer API (starting with Discovery and Service Mapping Patterns version 1.35.0).

  Resource Explorer API permissions and coverage

IAM permission	Coverage
`resource-explorer-2:Search`	Partial
`resource-explorer-2:Search` + `iam:CreateServiceLinkedRole`	Full

  For more information, go to the AWS Documentation and search for the "Getting started with Resource Explorer" article.

- Discover only datacenters with resources by setting the *mid.cloud.discovery.sonar.discover_all_aws_datacenters* MID Server property to **false**. For more information, see Limit AWS discovery to datacenters with resources.
- To exclude specific resource types from AWS datacenter discovery, configure the *sn_itom_pattern.discovery.aws.ldc.excluded_resource_types* system property with a comma-separated list of resource types to exclude (starting with Discovery and Service Mapping Patterns version 1.35.0). For more information, see Exclude AWS resource types from datacenter discovery.

Note: Discovery and Service Mapping Patterns versions 1.29.0 through 1.32.0 used the AWS Config service instead of the Resource Explorer API to determine datacenter activity. For instructions on configuring AWS Config recorder, go to the AWS Documentation and search for the "Recording resources in the AWS Config console" article.

For more information, see the **AWS resources discovery by datacenters** section.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Support for AWS services in the China region

The latest version of Discovery and Service Mapping Patterns supports discovering AWS services in the China region. You can discover these services on the ServiceNow AI Platform, starting from Xanadu Patch 3 and Washington DC Patch 9 instances.

Discovering AWS services in the China region requires using a datacenter URL when setting up an AWS service account. For example: `https://organizations.cn-northwest-1.amazonaws.com.cn`.

- To learn more about AWS master account and sub-account support in the China region, see KB1704526.
- To identify AWS patterns supported in the China region, refer to the Cloud Discovery patterns spreadsheet. The **AWS China Region Support** column has a Yes value for supported patterns.

AWS resources discovery by datacenters

Starting with version 1.29.0, Discovery and Service Mapping Patterns introduces a new AWS datacenter discovery model. The previous model discovered all datacenters, regardless of whether they contained relevant resources. The new model improves the AWS discovery performance by focusing only on datacenters that contain resources.

AWS has multiple datacenters around the world, but resources like load balancers and virtual machines are typically deployed in only some of them. The **Amazon AWS Datacenter Discovery** pattern runs before all other AWS patterns to identify datacenters with resources related to your service account ("active") and those without ("passive"). A datacenter is classified as "empty" when the Resource Explorer API is unavailable or lacks permissions, and patterns continue to run as normal. You can check the discovery log for the exact cause — for more information, see Logs for horizontal discovery.

After identifying "active", "passive", or "empty" datacenters, the discovery schedule continues to execute all AWS patterns only for "active" or "empty" datacenters, to discover your AWS cloud resources. "Passive" datacenters are ignored during the schedule. The **Refresh Datacenters** flow continues to display all regions, not just active ones. You don't need to create another schedule when a resource is added or a datacenter switches from passive to active.

The *mid.cloud.discovery.sonar.discover_all_aws_datacenters* MID Server property is set to **true** by default, which discovers all datacenters. To limit discovery to the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

"active" or "empty" datacenters, set this property to **false**. Starting with Discovery and Service Mapping Patterns version 1.35.0, you can exclude specific resource types from AWS datacenter discovery to avoid marking regions with only default or unneeded resources as active. For information on setting up active datacenter discovery, see the **Optimize discovery by including only datacenters with resources** prerequisite.

Datacenters discovered before upgrading to version 1.29.0 remain in the **Amazon AWS Datacenters** table. Depending on the service account and MID Server property settings, you might notice differences in discovery time and CMDB data.

Differences in datacenter discovery by mid.cloud.discovery.sonar.discover_all_aws_datacenters MID Server property setting

MID Server property setting	Flow	Discovered/displayed datacenters
False	New schedule	All datacenters except passive
False	Refresh Datacenters	All datacenters
True (default)	New schedule	All datacenters
True (default)	Refresh Datacenters	All datacenters

AWS datacenter model flow by mid.cloud.discovery.sonar.discover_all_aws_datacenters MID Server property setting

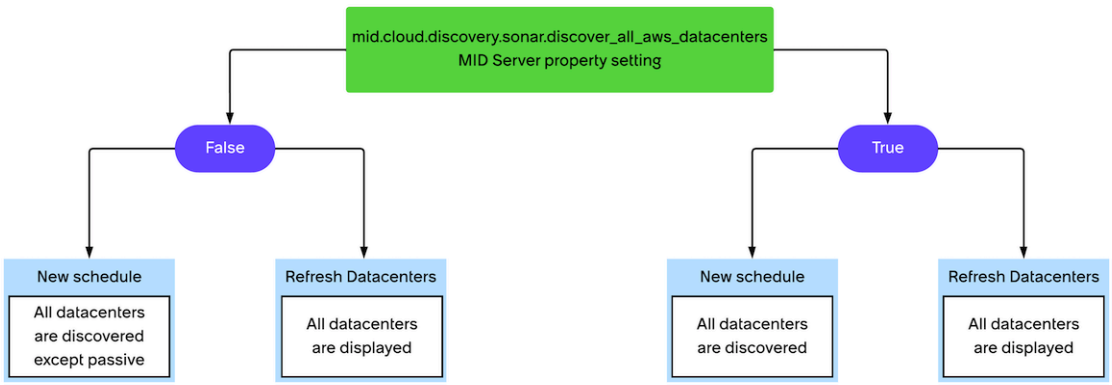

Note: Discovery and Service Mapping Patterns versions 1.29.0 through 1.32.0 used the AWS Config service instead of the Resource Explorer API to determine datacenter activity. For instructions on configuring AWS Config recorder, go to the AWS Documentation and search for the "Recording resources in the AWS Config console" article.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS discovery Dependency Views

Dependency Views displaying the cloud load balancer and connected components

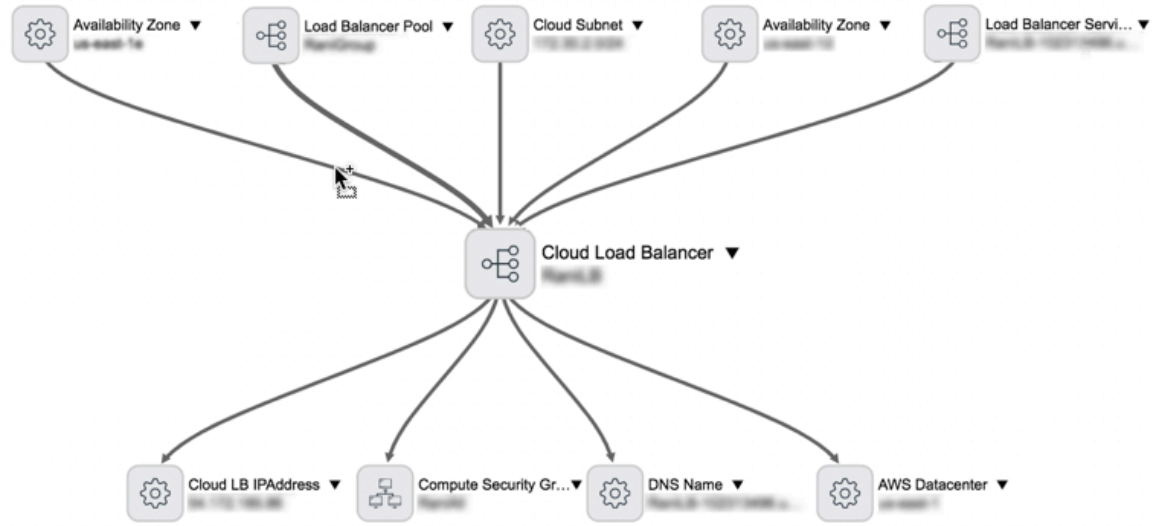

Dependency Views displaying components connected to the cloud network in the AWS environment

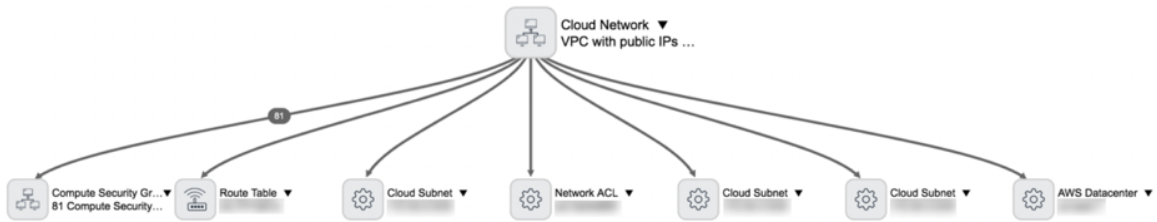

Dependency Views showing Virtual Machine and connected components in the AWS environment

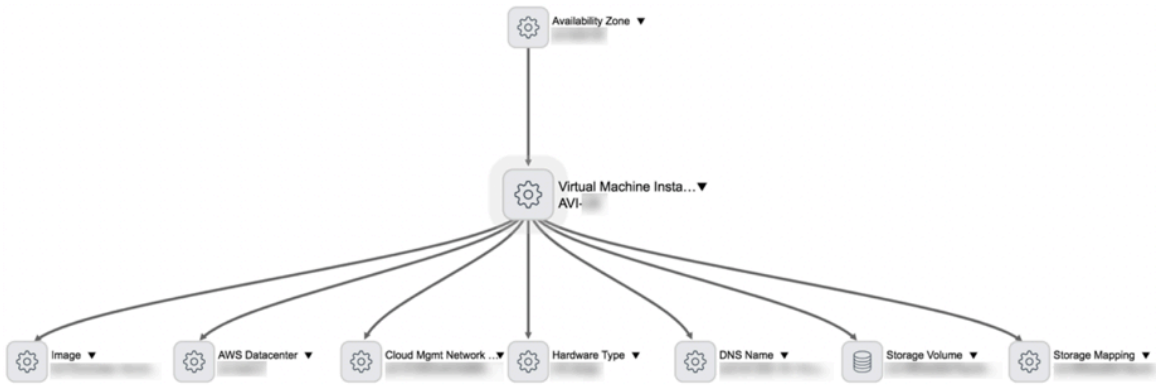

Services discovered by patterns

Horizontal discovery finds EC2 and VPC services running on AWS resources.

Services discovered by Discovery using patterns

Service name	CI class	Pattern
AWS::EC2::SecurityGroup	Compute Security Group [cmdb_ci_compute_security_group]	Amazon AWS Security Group Ports

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Services discovered by Discovery using patterns (continued)

Service name	CI class	Pattern
AWS::EC2::Subnet	Cloud Subnet [cmdb_ci_cloud_subnet]	Amazon AWS Subnet Events
AWS::EC2::VPC	Cloud Network [cmdb_ci_network]	Amazon AWS Network Events
AWS::EC2::Instance	Virtual Machine Instance [cmdb_ci_vm_instance]	Amazon AWS Virtual Server Events
EQS::EC2::Volume	Storage Volume [cmdb_ci_storage_volume]	Amazon AWS Elastic Block Storage
AWS::ElasticLoadBalancingV2::LoadBalancer	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Amazon AWS Application and Network LBs Events
AWS::ElasticLoadBalancing::LoadBalancer	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Amazon AWS Classic LBs Events

Data collected by Service Mapping during tag-based discovery

Service Mapping uses tag-based discovery to create service instance maps including the Cloud components. The Service Mapping application comes with the following preconfigured CI relationships used for tag-based discovery. These CI relationships are available from the 1.0.68 release on the ServiceNow Store.

CI	Relationship	CI
Configuration Item [cmdb_ci]	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]
Logical Datacenter [cmdb_ci_logical_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Related topics

- Kubernetes discovery using patterns
- AWS Certificate Manager discovery
- Amazon Bedrock pattern-based discovery

Amazon API Gateway discovery with Patterns

The ServiceNow Discovery and Service Mapping applications use the Amazon AWS API Gateway pattern to find Amazon API Gateways and connections to other entities. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Starting with Discovery and Service Mapping Patterns version 1.29.0, the Amazon AWS API Gateway pattern data model has changed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon AWS API Gateway pattern data model (starting with Discovery and Service Mapping Patterns version 1.29.0)

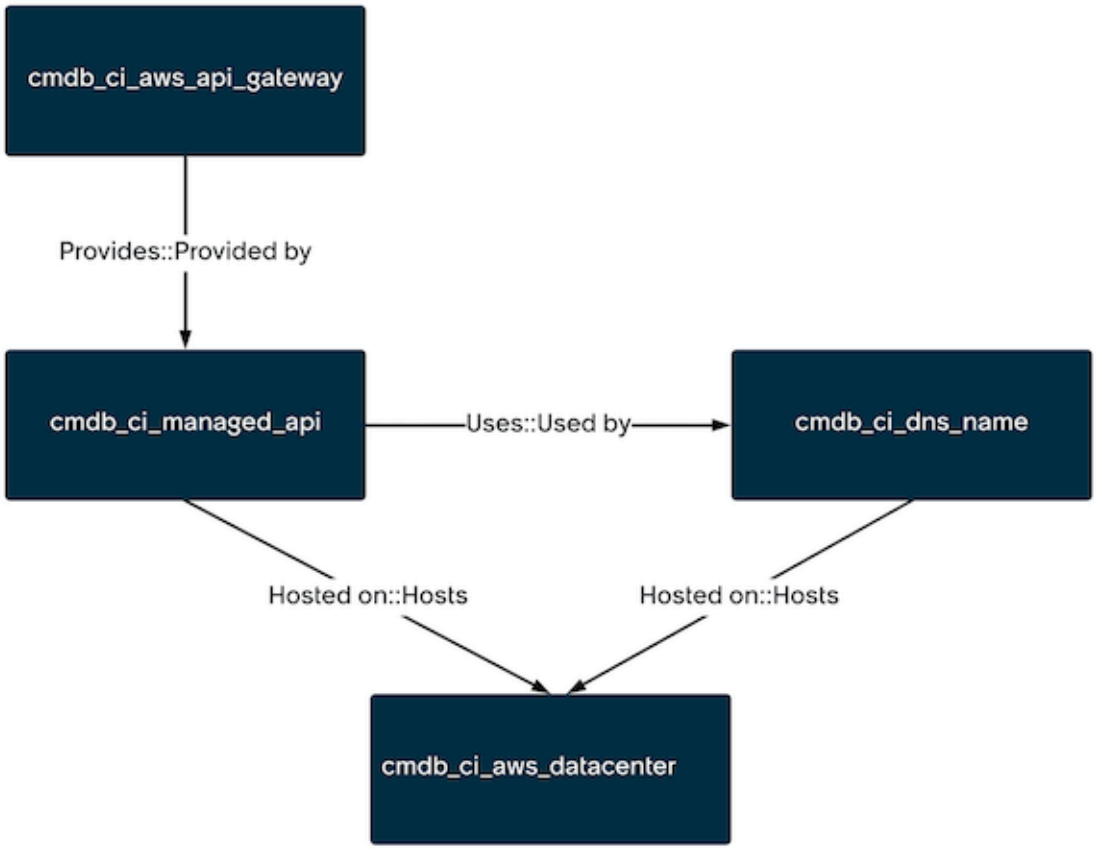

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the AWS prerequisites

For more information, see the prerequisite section in AWS discovery using patterns.

Retire legacy data model CIs to avoid duplicates

To avoid duplicate CIs, activate a schedule job to mark legacy data model CIs as retired. For more information, see Retire legacy AWS API Gateway data model CIs.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS API Gateway pattern.

AWS API Gateway [cmdb_ci_aws_api_gateway]

Field	Description
Name [name]	Concatenate of "aws-gateway-" + account ID + region. For example: **aws-gateway-000000000000-us-east-1**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS API Gateway [cmdb_ci_aws_api_gateway] (continued)

Field	Description
ID [id]	Concatenate of "aws-gateway-" + account ID + region. For example: **aws-gateway-000000000000-us-east-1**.
Install Status [install_status]	Install status of the AWS gateway. Default value is Installed.
Operational status [operational_status]	Operational status of the AWS gateway. Default value is Operational.

Managed API [cmdb_ci_managed_api]

Field	Description
ID [id]	Concatenate of "https://" + API ID + ".execute-api." + region + ".amazonaws.com". For example: **https://a1b2c3d4e5.execute-api.us-east-1.amazonaws.com**.
Name [name]	Name of the API.
Description [short_description]	Description of API.
Install Status [install_status]	Install status of the API. Default value is Installed.
Operational status [operational_status]	Operational status of the API. Default value is Operational.
Type [type]	Type of API: REST.
Version [version]	Version of API.
Fully qualified domain name [fqdn]	Concatenate of API ID + ".execute-api." + region + ".amazonaws.com". For example: **a1b2c3d4e5.execute-api.us-east-1.amazonaws.com**.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Unique name or IP address of the DNS host.
IP Address [ip_address]	IP address of the DNS host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon API gateway dependency view

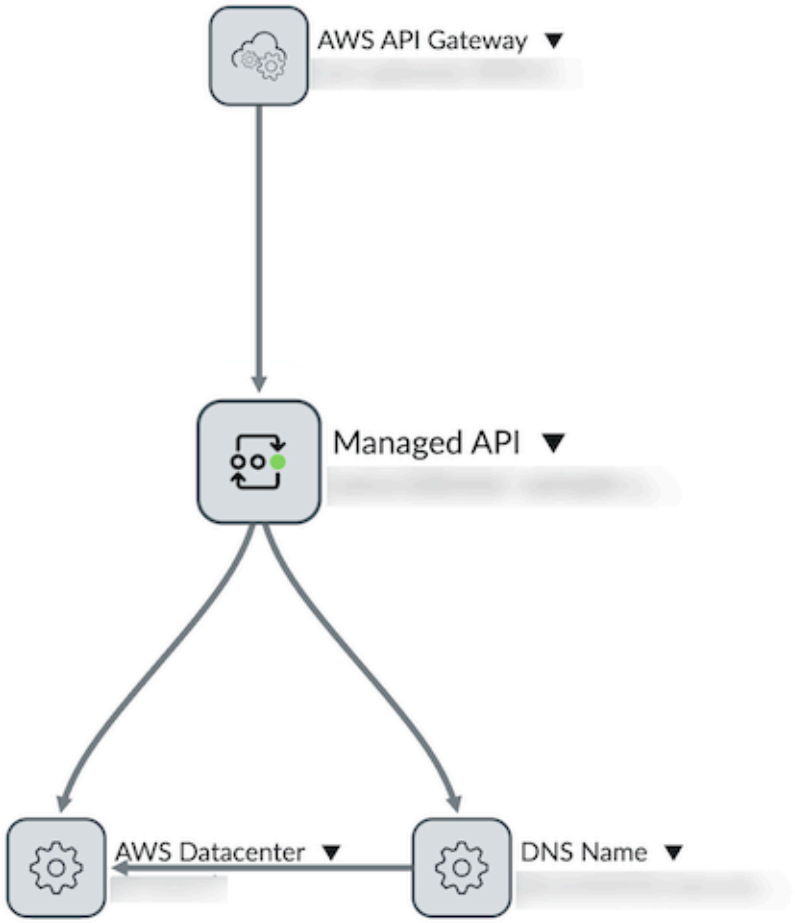

Relationships

Discovery creates these relationships to support Amazon API Gateway discovery.

CI	Relationship	CI
AWS API Gateway [cmdb_ci_aws_api_gateway]	Provides::Provided by	Managed API [cmdb_ci_managed_api]
Managed API [cmdb_ci_managed_api]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Managed API [cmdb_ci_managed_api]	Uses::Used by	DNS Name [cmdb_ci_dns_name]
DNS Name [cmdb_ci_dns_name]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Key Value [cmdb_key_value]	References	Managed API [cmdb_ci_managed_api]

Tag discovery

The Amazon AWS API Gateway pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Managed API [cmdb_ci_managed_api] table.

Connections discovered by Service Mapping during the top-down discovery

Service Mapping performs the top-down discovery of the AWS API Gateways in the context of application services. It discovers outgoing connections of the AWS API Gateways:

- AWS Lambda
- VPC Link for Virtual Private Cloud
- HTTP endpoint defined as the configured URL on the API Gateway

Test patterns for discovering AWS API Gateway and AWS Lambda

Run the horizontal and top-down discovery using Amazon AWS API Gateway [cmdb_ci_cloud_gateway] and Amazon AWS Lambda [cmdb_ci_cloud_function] patterns. Verify that the result is as expected.

Before you begin

Make sure that your deployment complies with all the prerequisites for Amazon API Gateway discovery with Patterns and AWS Lambda discovery with Patterns.

Role required: admin

Procedure

1. Create an application service using the following properties for its entry point.

Field	Description
CI type	HTTPS
URL	The URL for invoking the API Gateway. Typically, the API GW has the following format: https://{restapi_id}.execute-api.{region}.amazonaws.com/{stage_name}/ For example, https://k40o0l3af6.execute-api.us-east-1.amazonaws.com/prod/. For information about the URL format, see https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-call-api.html. **Note:** The default stage_name parameter is *prod*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Click **Discover**.
   Service Mapping runs the top-down discovery using the Amazon AWS API Gateway [cmdb_ci_cloud_gateway] pattern.

   After Service Mapping completes identifying the API Gateway configuration item (CI), it runs the **Connection to APIs** connection section and discovers connection to the AWS Lambda. It then runs the Amazon AWS Lambda [cmdb_ci_cloud_function] pattern and identifies the AWS Lambda CI.

3. After Service Mapping finishes to discover this application service, verify that it discovered and identified AWS components correctly.
   Check that the following actions happen:

   - Service Mapping correctly resolves the URL, finds all IP addresses behind it and creates an entry point for each IP address.

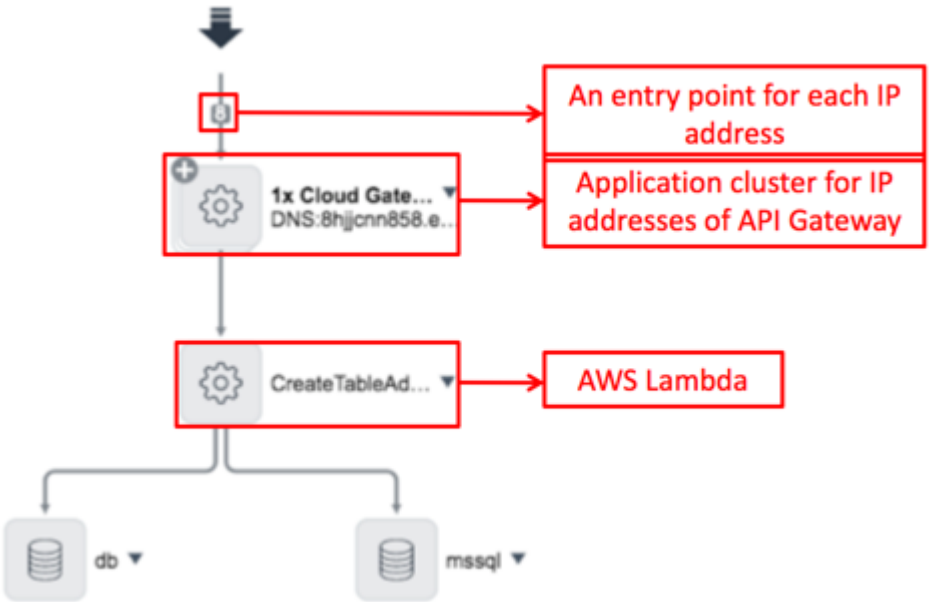

   - Service Mapping identifies the AWS API Gateway and creates an application cluster for IP address.
   - The **Properties** pane shows the properties for the AWS API Gateway: Cloud Gateway [cmdb_ci_cloud_gateway].
   - Service Mapping discovers the AWS Lambda connected to the application cluster for API Gateway.
   - Service Mapping discovers the databases connected to AWS Lambda.

Retire legacy AWS API Gateway data model CIs

Avoid duplicate AWS API Gateway configuration items (CIs) by retiring legacy data model CIs through a scheduled job.

Before you begin

- Verify you have at least version 1.29.0 of Discovery and Service Mapping Patterns.
- Run an AWS discovery to populate the CIs in the managed API data model. For more information, see Create Discovery schedules for cloud resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin

About this task

Starting with Discovery and Service Mapping Patterns December release, the Amazon AWS API Gateway pattern data model has changed. To avoid duplicate CIs, activate a schedule job to mark legacy data model CIs as retired.

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. In the **Name** column, search for `Retire Cloud Gateways migrated to Managed API`.
3. Select the **Retire Cloud Gateways migrated to Managed API** scheduled job.
4. Select **Execute Now**.

Amazon API Gateway Domain Name pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - API Gateway Domain Name - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS API Gateway - Domain Name [cmdb_aws_api_gateway_domain_name]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the domain.
Name [name]	The custom domain name associated with an API Gateway.
Certificate Name [certificate_name]	The name of the secure sockets layer and transport layer security (SSL/TLS) certificate used for the domain.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Domain Name Status [domain_name_status]	The status of the custom domain name.
Endpoint Type [endpoint_type]	The type of API endpoint. For example: EDGE or REGIONAL.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - API Gateway Domain Name - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the domain.
Name [name]	The custom domain name associated with an API Gateway.
Resource type [resource_type]	Type of resource. The value is set to **AWS::API Gateway::Domain Name**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships and references

The Amazon AWS - API Gateway Domain Name - Extended Inventory (LP) pattern creates these relationships and references to support Amazon API Gateway Domain Name discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS API Gateway - Domain Name [cmdb_aws_api_gateway_domain_name]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon Athena Workgroup pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon Athena Workgroups on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Athena Workgroup - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Athena Workgroup [cmdb_aws_athena_workgroup]

Field	Description
Name [name]	Name of the Amazon Athena workgroup.
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the workgroup.
State [state]	Current state of the workgroup. For example: ENABLED or DISABLED.
Effective Engine Version [effective_engine_version]	Engine version used by the workgroup to process queries.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Athena Workgroup - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the workgroup.
Name [name]	Name of the AmazonAmazon Athena workgroup.
Resource type [resource_type]	Resource type. The value is set to **AWS::Athena::Workgroup**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships and references

The Amazon AWS - Athena Workgroup - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon Athena Workgroup discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS Athena Workgroup [cmdb_aws_athena_workgroup]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Athena Workgroup - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon CloudFront Distribution pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon CloudFront Distributions on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - CloudFront Distribution - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS CloudFront Distribution [cmdb_aws_cloudfront_distribution]

Field	Description
Name [name]	Domain name of the CloudFront distribution.
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the CloudFront distribution.
Id [id]	Unique identifier assigned to the CloudFront distribution.
Status [status]	Deployment status of the CloudFront distribution.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - CloudFront Distribution - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the CloudFront distribution.
Name [name]	Domain name of the CloudFront distribution.
Resource type [resource_type]	Type of resource. The value is set to **AWS::CloudFront::Distribution**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Amazon AWS - CloudFront Distribution - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon CloudFront Distribution discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

CI references

CI	Field	Referenced CI
AWS CloudFront Distribution [cmdb_aws_cloudfront_distribution]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - CloudFront Distribution - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon CloudWatch Log Group pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon CloudWatch Log Groups on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - CloudWatch Log Group - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS CloudWatch Log [cmdb_aws_cloudwatch_log]

Field	Description
Name [name]	Name of the Amazon CloudWatch log group.
Object ID [object_id]	Amazon Resource Name (ARN) of the log group.
Log Group Arn [log_group_arn]	ARN of the log group.
Log Group Class [log_group_class]	Storage class of the log group.
Stored Bytes [stored_bytes]	Number of bytes stored in the log group.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - CloudWatch Log Group - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN of the log group.
Name [name]	Name of the Amazon CloudWatch log group.
Resource type [resource_type]	Type of resource. The value is set to **AWS::CloudWatch::Log**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Amazon AWS - CloudWatch Log Group - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon CloudWatch Log Group discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS CloudWatch Log [cmdb_aws_cloudwatch_log]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - CloudWatch Log Group - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon Cognito discovery

The ServiceNow Discovery and Service Mapping applications use the Amazon AWS Cognito pattern to provide authentication, authorization, and user management functions for AWS customers. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

You can use this pattern on the ServiceNow AI Platform using London Patch 8, Madrid Patch 2, or later releases.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

User permissions

Provide user with read-only permission to run the following API:

- https://cognito-idp.<region>.amazonaws.com
- Method: POST

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Body: {\"MaxResults\": 10}
- Headers: X-Amz-Target:AWSCognitoIdentityProviderService.ListUserPools,Content-Type:application/x-amz-json-1.0

AWS Credentials

On your instance, configure credentials of type **AWS Credentials** and set to **Active**.

Cloud service account

On your instance, configure the cloud service account of type **AWS Datacenter** and set to **AWS account ID**. Use the credentials defined in the preceding **AWS Credentials**.

Discovery schedule

Create a cloud application schedule for discovering AWS Cognito and configure the attributes. Set Discovery to **Cloud application**.

Execution pattern

Create and define the serverless execution pattern for cloud application discovery.

1. Create new **Cloud Execution Patterns**.
2. Define **Name**.
3. Verify that **Active** is true.
4. Verify that **Domain** is global.
5. Choose the **AWS pattern** you want to run.
6. Create multiple records if you want to run more than one pattern.

Discovery schedule for full AWS discovery

Create a discovery schedule from your **Cloud service account** created in the earlier procedure.

1. Click on **Discover Datacenter** and wait for it to finish.
2. Click **Create Discovery Schedule**.
3. This new schedule is created under the Discovery Schedule and runs all AWS patterns.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery and Service Mapping during horizontal and top-down discovery

The AWS Cognito pattern collects data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Collected information from the AWS Cognito pattern

Field	Description
Main CI: cmdb_ci_cloud_authentication	
name	A descriptive name used to identify the user pool.
object_id	This is equal to the account_id and used by IRE identification rules.
Fqdn	Example of an ARN: arn:aws:cognito-idp:eu-west-1:751200741520:userpool/eu-west-1_fim5E2mix

Tags are also being collected by an extension section that runs following the pattern. The tagging API for AWS specifies the resource type Cognito.

Collected information from the AWS Cognito tags

Field	Description
cmdb_key_value	
key	The actual tag key.
value	The tag value.
configuration_item	The unique resource ID (ARN) that identifies the resource in the AWS console.

CI relationships

The AWS Cognito pattern creates the following CI relationship.

CI	Relationship	CI
Cloud authentication [cmdb_ci_cloud_authentication]	Hosts:Hosted on	Logical datacenter [cmdb_ci_logical_datacenter]

Troubleshooting

If the mapping process does not proceed as you expected, follow the following suggestions.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
		Note: If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

Amazon DB cluster discovery with Patterns

ServiceNow Discovery uses the AWS DB cluster discovery pattern to find and map Aurora DB clusters, Aurora Serverless resources, Amazon Neptune DB instances, and Amazon DocumentDB instances. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The discovered DB clusters form new relationships with the DB Instances discovered by the AWS Relational Database Service pattern. Therefore, the AWS DB cluster pattern is triggered after the AWS Relational Database Service pattern during discovery. For further information, see Amazon Relational Database Service (Amazon RDS) discovery with Patterns.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Install the Discovery and Service Mapping Patterns application (version 1.0.75) from the ServiceNow Store.
- On the ServiceNow AI Platform, configure AWS credentials, using a secret key and an access key.
- Create a service account. Set the **Account ID** to the Amazon account ID to which RDS belongs. Use the Account ID as it appears in the AWS Management Console.
- Set read-only permissions for these REST APIs:
  - https://rds.<REGION>.amazonaws.com/?Action=DescribeDBInstances
  - https://rds.<REGION>.amazonaws.com/?Action=DescribeDBClusters
  - https://rds.<REGION>.amazonaws.com/?Action=DescribeDBInstances&DBInstanceIdentifier=<input_object_id> for events discovery
  - https://tagging.<REGION>.amazonaws.com/ for tags discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For Cloud Discovery, download the Discovery and Service Mapping Patterns from the ServiceNow Store.
- When installing the MID Server, ensure that the host machine meets or exceeds the MID Server system requirements published on the ServiceNow documentation site.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

Table and field	Description
Cloud Database Cluster [cmdb_ci_cloud_db_cluster]	
Cluster ID [cluster_id]	Contains a user-supplied DB cluster identifier. This identifier is the unique key that identifies a DB cluster.
Name [name]	Contains the name of the initial database of the DB cluster that was provided on creation. If a name was not specified, the cluster identifier is used as the name. This name is returned for the life of the DB cluster.
FQDN [fqdn]	Specifies the connection endpoint for the primary instance of the DB cluster.
TCP Port [tcp_port]	Port of the database endpoint.
Vendor [vendor]	The vendor, Amazon.
Cluster Status [cluster_status]	Specifies the current state of this DB cluster.
Operational Status [operational_status]	Status of the database, ”available” if operational.
Cloud Database [cmdb_ci_cloud_database]	
Object ID [object_id]	The instance identifier for this member of the DB cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
	Note: For Database Instances, this table is populated with only the object_id attribute. This attribute is required for identification of the Cloud Database CIs, to ensure relationships between the CIs and the Cloud Database Clusters can be made.
Name [name]	Contains the name of the initial database of the DB cluster that was provided on creation. If a name was not specified, the cluster identifier is used as the name. This name is returned for the life of the DB cluster.
FQDN [fqdn]	Specifies the connection endpoint for the primary instance of the DB cluster.
TCP Port [tcp_port]	Port of the database endpoint.
Vendor [vendor]	The vendor, Amazon.
Operational Status [operational_status]	Status of the database, "available" if operational.
IP Address [cmdb_ci_ip_address]	
IP Address [ip_address]	IP address of the Cloud DB Cluster Writer endpoint.
Name [name]	The IP address.
Fully qualified domain name [fqdn]	Address of the database endpoint. The DNS name.
Netmask [netmask]	IP netmask [0.0.0.0].

CI relationships

CI	Relationship	CI
cmdb_ci_cloud_db_cluster	Cluster off::Cluster	cmdb_ci_cloud_database
cmdb_ci_cloud_database	Hosted on::Hosts	cmdb_ci_aws_datacenter
cmdb_ci_cloud_db_cluster	Owns::Owned by	cmdb_ci_ip_address
cmdb_ci_cloud_database	Owns::Owned by	cmdb_ci_ip_address
cmdb_key_value	Reference only	cmdb_ci_cloud_db_cluster
cmdb_key_value	Reference only	cmdb_ci_cloud_database

Amazon DynamoDB discovery using Patterns

The ServiceNow Discovery and Service Mapping applications use the Amazon AWS DynamoDB pattern to find components of DynamoDB. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Amazon AWS DynamoDB pattern deploys REST APIs for fetching information about DynamoDB. You can also configure Service Mapping and Discovery to discover AWS DynamoDB components using AWS events.

The ServiceNow AI Platform supports Amazon Commercial Cloud Services (C2S) for discovery and cloud governance. C2S is the government program that brings AWS into the Intelligence Community.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

User access

On the AWS Console, provide a user with the following permissions: AmazonDynamoDBReadOnly Access and AWSLambdaInvocation-DynamoDB.

Cloud Service Account

On the ServiceNow AI Platform®, set up the service account with these properties:

- An **Account ID**, which is the Amazon account ID to which DynamoDB belongs. Use the Account ID as it appears in the AWS Management Console.
- The appropriate AWS credentials required for the account ID to reach the AWS cloud.

AWS Credentials

On the ServiceNow AI Platform®, configure AWS credentials, using a secret key and an access key.

Entry point for service containing DynamoDB

Use *AWS* **Resource Name Endpoint** as the Entry Point Type when creating a service containing DynamoDB. Provide the Amazon Resource Name (ARN) and host for the entry point.

(For top-down discovery by Service Mapping) Run horizontal discovery of the Logical Datacenters

Discover Logical Datacenters hosting DynamoDBs as described in Run Discovery on an AWS cloud service account in the ServiceNow Cloud Provisioning and Governance documentation.

Cloud Discovery

For Cloud Discovery, download the Discovery and Service Mapping pattern from the ServiceNow Store.

MID Server

When installing the MID Server, ensure that the host machine meets or exceeds the MID Server system requirements.

Discovery schedule

Create a schedule for the relevant AWS service account as described in Schedule Discovery on a service account in the ServiceNow Cloud Provisioning and Governance documentation.

Configure AWS event service

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To enable alert-based discovery, configure the AWS event service as described in Set up AWS event processing for Discovery and Service Mapping in the ServiceNow Cloud Provisioning and Governance documentation.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

Table and field	Description
DynamoDB Global Table [cmdb_ci_dynamodb_global_table]	
Object ID [object_id]	The Amazon Resource Name (ARN), expressed in this format: *arn:aws:dynamodb::xxxxxxxxxxxx:table:<table name>*
DynamoDB Table [cmdb_ci_dynamodb_table]	
Name [name]	Table name.
Creation date [creation date]	Creation date of DynamoDB according to the time zone of the MID Server that discovers this DynamoDB.
Cloud provider [cloud_provider]	Amazon AWS
Object ID [object_id]	The Amazon Resource Name (ARN), expressed in this format: *arn:aws:dynamodb:<region>:xxxxxxxxxxxx:tabl name>*
Global ARN [global_arn]	The global ARN of the table for the global DynamoDB table. This field is empty if the DynamoDB is not global.
Stream ARN [stream_arn]	The choices are: • Stream ARN • Disabled
Service name [service_name]	Set to **DynamoDB** by default.
Encryption [encryption]	Encryption configuration on DynamoDB. The choices are: • Enabled • Disabled

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Replication group [replication_group]	Table ARN of the connected global DynamoDB tables.
Read Autoscaling [read_autoscaling]	The choices are: • Enabled • Disabled
Read Units [read_units]	The number of read units on the DynamoDB table.
Read Max Units [read_max_units]	The maximum allowed number of read units on the DynamoDB table.
Read Min Units [read_min_units]	The minimum allowed number of read units on the DynamoDB table.
Write Units [write_units]	The number of write units on the DynamoDB table.
Write Max Units [write_max_units]	The maximum allowed number of write units on the DynamoDB table.
Write Min Units [write_min_units]	The minimum allowed number of write units on the DynamoDB table.
Table Size [table_size]	The currently provisioned size of the DynamoDB table in bytes.
Write Autoscaling [write_autoscaling]	The choices are: • Enabled • Disabled
Point-in-Time [point_in_time]	The choices are: • Enabled • Disabled

The graphic illustrates CIs that are part of DynamoDB discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dependency Views showing DynamoDB components

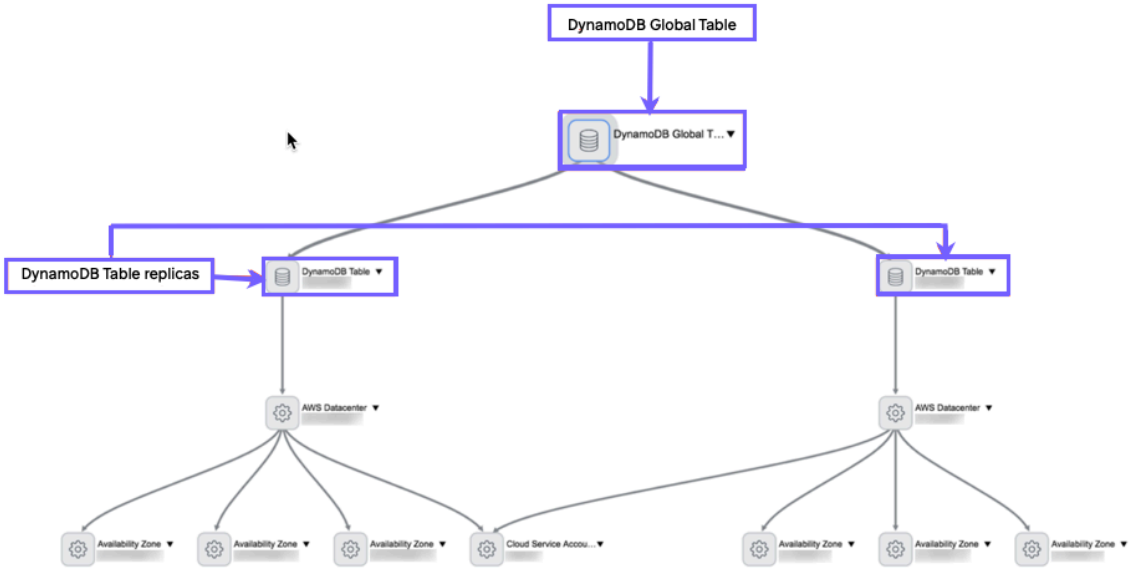

CI relationships

The Amazon AWS DynamoDB pattern creates the following CI relationships:

CI	Relationship	CI
DynamoDB Global Table [cmdb_ci_dynamodb_global_table]	Cluster of::Cluster	DynamoDB Table [cmdb_ci_dynamodb_table]

Data collected by Service Mapping during top-down discovery

The Amazon AWS DynamoDB pattern identifies connections from the DynamoDB to AWS Lambda functions.

Troubleshooting

If the mapping process does not proceed as you expected, follow the following suggestions.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
		Note: If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

Related topics

Define a Cloud REST Query for cloud computing devices

Define a Cloud REST Query for cloud computing devices

As part of creating or modifying a discovery pattern, you can use the **Cloud REST Query** operation to extract information from configuration items (CIs) of the PaaS (Platform as a Service) type, such as Microsoft Azure or Amazon Web Services.

Before you begin

- Verify that the operating system of the CIs for which you want to use the Cloud REST Query operation extends the Logical Data Center OS type [cmdb_ci_logical_datacenter]:
  1. Navigate to **Pattern Designer** > **Discovery Patterns**.
  2. Double-click the relevant pattern to open it in the Pattern Designer.
  3. Click the **Basic** tab.
  4. Note the operating system for this CI.
  5. Navigate to **System Definition** > **Tables**.
  6. Set the search field to **Label** and enter the name of the operating system as stated on the **Basic** tab of the pattern.
  7. Find the operating system in the list and verify that Logical Datacenter appears in the **Extends table** column.
- Navigate to the relevant pattern step:
  1. On the pattern form, select the relevant identification section for Discovery.

     Alternatively, select the relevant identification or connection section for Service Mapping.
  2. Select the relevant pattern step or click the **Add a step above** icon to add a step.

Basic knowledge of programming is desirable.

Role required: pd_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Deploy the Cloud REST Query operation in patterns used for discovery of PaaS CIs. This Java-script based custom operation is available only after downloading patterns version 1.0.24 or later from ServiceNow Store. Use this operation instead of the Cloud REST Call operation available in the base system.

Refer to the official API-related documentation provided by manufacturers to obtain the query syntax for the device you want to query using the Cloud REST Query operation. For example:

- http://docs.aws.amazon.com/AWSEC2/latest/APIReference/Welcome.html
- https://docs.microsoft.com/en-us/rest/api/apimanagement/

Procedure

1. Select **Cloud REST Query** from the **Operation** list.
2. Define query parameters as necessary:

Note: If you customized this operation, the query parameters may be different.

Field	Description
URL	Specify the URL as described in the official Microsoft Azure or Amazon Web Services documentation. You can use variables. You can also enter a value from the specific field in a tabular variable.
Method	Enter the method value as a string using all capital letters, for example "GET". Supported methods are: ○ GET ○ POST ○ PUT Refer to the relevant API guide for information. If the API documentation does not specifically mention the HTTP query method, use the GET method.
Body	(Optional) Enter a request body as a string. If the string contains quotation marks, use the backslash mark in front of the quotation marks to indicate that the string does not end at the quotation marks.
Headers	(Optional) If the relevant API documentation states that HTTP headers must be sent, enter these headers in the following format: Header_name1:header_value1,header_name2:header_value2 For example, `Content-Type:xml`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Formatted	Formatted JSON or not

3. To save the command output in its entirety as a variable, select NONE from the **Define Parsing** list and define the variable.
4. Select the parsing strategy from the **Define Parsing** list.
5. If working in the Debug Mode, define the parsing criteria as follows:

   a. Click **Run Operation** to see the result in the **Output** pane.

   b. In the **Output** pane, mark text or symbols that you want a variable to contain.

   c. In the variable name box, enter the name for the new variable, for example 'process_name'.

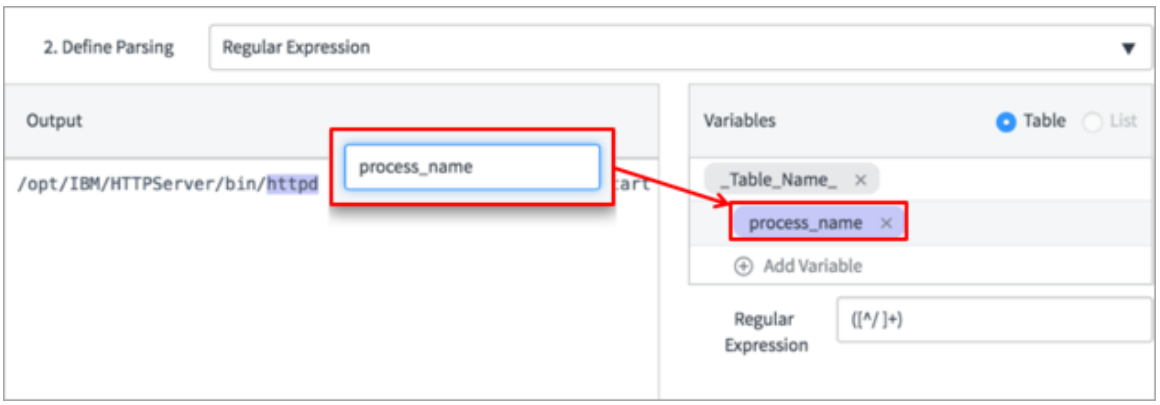

   d. Press **Enter**.
   The new variable is added in the **Variables** pane.

6. If you are not working in Debug Mode, define the parsing criteria as follows:

   a. In the Variables pane, click **Add Variable** and enter the name for the new variable.

   b. Click the **Advanced** icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

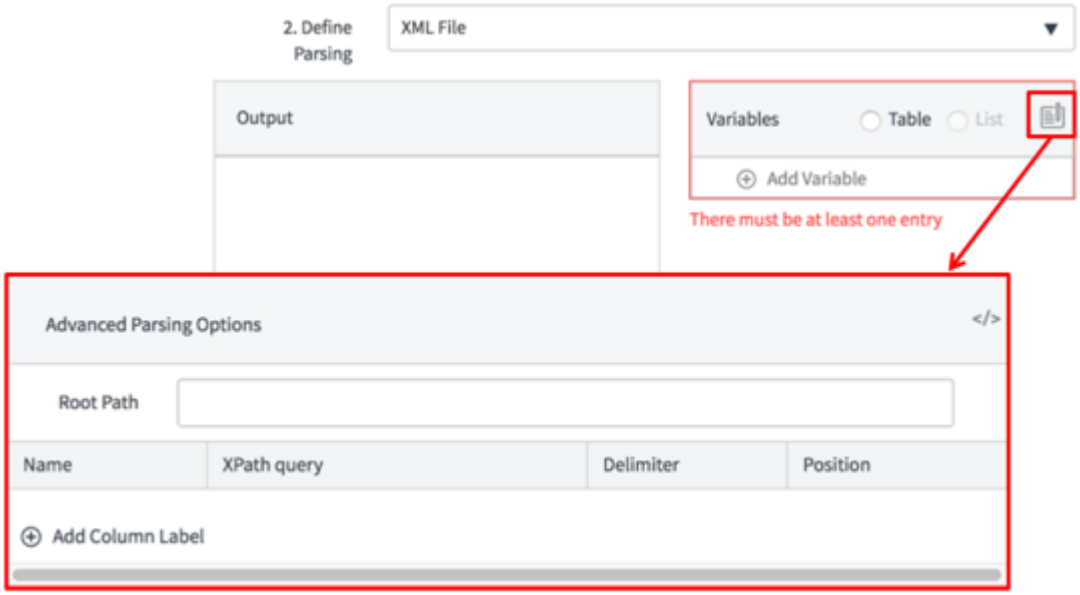

c. Click **Add Column Label**.

d. Enter the parsing query.

e. Enter a value for the delimiter.

f. Enter a value for the position.

7. Select **Terminate** to stop discovery if no results are found.

Example:

This operation is used in	This item
Hierarchy	Application
CI Type	Cloud Object Storage [cmdb_ci_cloud_object_storage]
Pattern	Amazon AWS S3
Section	Identification of AWS S3
Step number and Name	3. List buckets

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

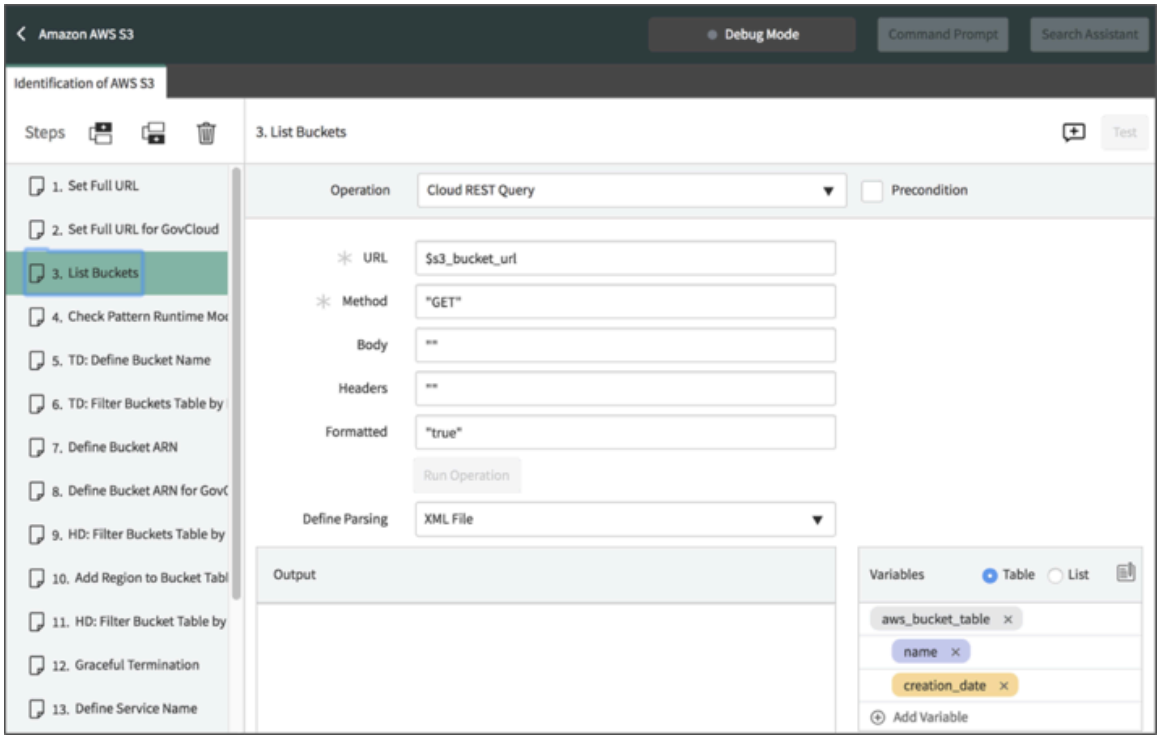

Use the Cloud REST Query operation to extract information about AWS buckets, their names and creation dates.

Amazon DynamoDB Cluster pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - DynamoDB Cluster - Extended Inventory (LP) pattern.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Name [name]	The name assigned to the DynamoDB Accelerator (DAX) cluster.
Cluster ID [cluster_id]	The Amazon Resource Name (ARN) uniquely identifying the DAX cluster.
Cluster Status [cluster_status]	The current operational status of the DAX cluster.
Cluster Type [cluster_type]	Cluster type: DynamoDB Cluster.
Fully qualified domain name [fqdn]	The endpoint address used to connect to the DAX cluster.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Node Count [node_count]	The total number of nodes in the DAX cluster.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Description of the cluster: DynamoDB Cluster.
Vendor [vendor]	Name of the vendor: Amazon.

Cluster Node [cmdb_ci_cluster_node]

Field	Description
Name [name]	The unique identifier of the DAX node within the cluster.
Serial Number [serial_number]	Unique identifier of the node in the following format: **<cluster_arn>/node/<node_id>**.
Node Status [node_status]	The current operational status of the node. For example: available, joining, or deleting.
Fully qualified domain name [fqdn]	The endpoint address used to connect to the specific DAX node.
Install Status [install_status]	Install status of the node. Default value is Installed.
Operational Status [operational_status]	Operational status of the node. Default value is Operational.
Cluster [cluster]	References the Cloud Database Cluster [cmdb_ci_cloud_db_cluster] table.
Description [short_description]	Description of the node: DynamoDB Cluster Node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - DynamoDB Cluster - Extended Inventory (LP) pattern creates these relationships to support Amazon DynamoDB Cluster discovery.

CI	Relationship	CI
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Cluster Node [cmdb_ci_cluster_node]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Cluster Node [cmdb_ci_cluster_node]	Cluster of::Cluster	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]
Cluster Node [cmdb_ci_cluster_node]	References	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon EC2 Amazon EBS Snapshot pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Elastic Compute Cloud EBS Snapshot - Extended Inventory (LP) pattern.

Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]

Field	Description
Object ID [object_id]	The Amazon Resource Name (ARN) uniquely identifying the EBS Snapshot. The ARN consists of the AWS region, account ID, and snapshot ID in the following format: **arn:aws:ec2:<region>:<accountId>:snapshot/<snapshot_id>**.
Name [name]	The unique identifier for the EC2 snapshot.
State [state]	The current state of the snapshot. For example: pending or completed.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Encrypted [is_encrypted]	Indicates whether the snapshot is encrypted.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **AWS::Elastic Compute Cloud::EBS Snapshot**.
Size (GB) [size]	The size of the resource expressed in gigabytes (GB).
Size bytes [size_bytes]	The size of the snapshot in bytes.
Volume Name [volume_name]	The identifier of the volume associated with the snapshot.

CI relationships

The Amazon AWS - Elastic Compute Cloud EBS Snapshot - Extended Inventory (LP) pattern creates these relationships to support Amazon EC2 Amazon EBS Snapshot discovery.

CI	Relationship	CI
Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon EC2 Reserved Instance pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Elastic Compute Cloud Reserved Instance - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Elastic Compute Cloud - Reserved Instance [cmdb_aws_elastic_cloud_compute_reserved_instance]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the Reserved Instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Elastic Compute Cloud - Reserved Instance [cmdb_aws_elastic_cloud_compute_reserved_instance] (continued)

Field	Description
	The ARN consists of the region, account ID, and instance ID in the following format: **arn:aws:ec2:<region>:<accountId>:reserved-instances/<instance_id>**.
Name [name]	The unique identifier assigned to the EC2 instance.
State [state]	The current status or lifecycle state of the EC2 instance.
Configuration Item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.
Instance Type [instance_type]	The type of EC2 instance. For example: t2.micro or m5.large.
Scope [scope]	The reservation scope. For example: Regional or Availability Zone.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Elastic Compute Cloud Reserved Instance - Extended Inventory (LP) pattern.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the Reserved Instance. The ARN consists of the region, account ID, and instance ID in the following format: **arn:aws:ec2:<region>:<accountId>:reserved-instances/<instance_id>**.
Name [name]	The unique identifier assigned to the EC2 instance.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **AWS::Elastic Compute Cloud::Reserved Instance**.

CI relationships

The Amazon AWS - Elastic Compute Cloud Reserved Instance - Extended Inventory (LP) pattern creates these relationships to support Amazon EC2 Reserved Instance discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Elastic Compute Cloud - Reserved Instance [cmdb_aws_elastic_cloud_compute_reserved_instance]	References	Virtual Machine Instance [cmdb_ci_vm_instance]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon EC2 VPC Endpoint Service pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Elastic Compute Cloud VPC Endpoint Service - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Elastic Compute Cloud - VPC Endpoint Service [cmdb_aws_elastic_cloud_compute_vpc_endpoint_service]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the VPC Endpoint Service. The ARN consists of the AWS region, account ID, and service ID in the following format: **arn:aws:ec2:<region>:<accountId>:vpc-endpoint-service/<service_id>.**
Name [name]	The name assigned to the VPC endpoint service.
Configuration Item [configuration_item]	References the Cloud Gateway [cmdb_ci_cloud_gateway] table.
Is Customer Managed [is_customer_managed]	Indicates whether the VPC Endpoint Service is managed by a customer or AWS.
Service ID [service_id]	The unique identifier for the VPC endpoint service.
Service Type [service_type]	The type of the VPC endpoint service. For example: Interface or Gateway.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Elastic Compute Cloud VPC Endpoint Service - Extended Inventory (LP) pattern.

Cloud Gateway [cmdb_ci_cloud_gateway]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the VPC Endpoint Service. The ARN consists of the AWS region, account ID, and service ID in the following format: **arn:aws:ec2:<region>:<accountId>:vpc-endpoint-service/<service_id>.**
Name [name]	The name assigned to the VPC endpoint service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Gateway [cmdb_ci_cloud_gateway] (continued)

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **AWS::Elastic Compute Cloud::VPC Endpoint Service**.

CI relationships

The Amazon AWS - Elastic Compute Cloud VPC Endpoint Service - Extended Inventory (LP) pattern creates these relationships to support Amazon EC2 VPC Endpoint Service discovery.

CI	Relationship	CI
Cloud Gateway [cmdb_ci_cloud_gateway]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Elastic Compute Cloud - VPC Endpoint Service [cmdb_aws_elastic_cloud_compute_vpc_endpoint_service]	References	Cloud Gateway [cmdb_ci_cloud_gateway]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon EC2 VPC Peering Connection pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Elastic Compute Cloud VPC Peering Connection - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Elastic Compute Cloud - VPC Peering Connection [cmdb_aws_elastic_compute_cloud_vpc_peering_connection]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the VPC Peering Connection. The ARN consists of the region, account ID, and VPC ID in the following format: **arn:aws:ec2:<region>:<accountId>:vpc-peering-connection/<vpc_id>**.
Name [name]	The unique identifier of the VPC involved in the peering connection.
Status [status]	The current status of the VPC peering connection.
Configuration Item [configuration_item]	References the Cloud Gateway [cmdb_ci_cloud_gateway] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Elastic Compute Cloud VPC Peering Connection - Extended Inventory (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Gateway [cmdb_ci_cloud_gateway]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the VPC Peering Connection. The ARN consists of the region, account ID, and VPC ID in the following format: **arn:aws:ec2:<region>:<accountId>:vpc-peering-connection/<vpc_id>**.
Name [name]	The unique identifier of the VPC involved in the peering connection.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **AWS::EC2::VPCPeeringConnection**.

CI relationships

The Amazon AWS - Elastic Compute Cloud VPC Peering Connection - Extended Inventory (LP) pattern creates these relationships to support Amazon EC2 VPC Peering Connection discovery.

CI	Relationship	CI
Cloud Gateway [cmdb_ci_cloud_gateway]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Elastic Compute Cloud - VPC Peering Connection [cmdb_aws_elastic_compute_cloud_vpc_peering_connection]	References	Cloud Gateway [cmdb_ci_cloud_gateway]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon ECS resource discovery with Patterns

The ServiceNow Discovery application uses the Amazon AWS - ECS pattern to find resources managed by the Amazon Elastic Container Service (Amazon ECS). This includes resources that are run on the Fargate launch type. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon ECS is a highly scalable, fast, container management service that makes it easy to run, stop, and manage Docker containers. AWS Fargate is a compute engine for Amazon ECS that enables AWS customers to run containers without having to manage servers or clusters.

Discovery uses the Amazon AWS - ECS pattern to run horizontal discovery.

Starting with Discovery and Service Mapping Patterns version 1.18.0, the Amazon AWS - ECS pattern supports the discovery of Amazon Elastic Container Registry (Amazon ECR) images. For more information on container scanning, see Scan container images.

Starting with Discovery and Service Mapping Patterns version 1.31.0, you can choose to discover Docker image CIs only, without discovering Docker container CIs. Check your entitlements to determine whether you have access to 2026 Packaging SKU. For more information, see Disable Docker container CI discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Amazon ECS data model

The Amazon AWS - ECS pattern introduces the following CI classes that extend an existing CMDB class.

CI classes introduced by this pattern

CI class	Extends from
AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]	Virtual Machine Object [cmdb_ci_vm_object]
AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]	Virtual Machine Object [cmdb_ci_vm_object]
AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]	Virtual Machine Object [cmdb_ci_vm_object]
AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]	Virtual Machine Object [cmdb_ci_vm_object]

Prerequisites

Create AWS credentials with ECS resources viewing privileges

In the AWS console, create AWS credentials with privileges to view ECS resources. For more information, see: https://docs.aws.amazon.com/en_pv/AmazonECS/latest/userguide/get-set-up-for-amazon-ecs.html.

Note: The API requests use AWS Signature Version 4 to authenticate. When the correct credentials are provided, authentication is done internally for the Cloud API Call pattern steps.

Create AWS cloud credential records or use temporary credentials

In the ServiceNow AI Platform, create AWS credential records. You can also use a temporary credential generated by the AWS Security Token Service (AWS STS) for IAM roles. For more information, see Cloud credentials

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create AWS cloud service accounts

- In the ServiceNow AI Platform, create AWS cloud service accounts using permanent or temporary AWS credentials. For more information, see Setup AWS service accounts.
- In the Cloud Service Account record, select **Discover Datacenters** to find AWS Datacenters (LDCs) for the configured AWS cloud service account.

Schedule an AWS discovery

In the ServiceNow AI Platform, schedule either a full AWS pattern discovery or only an ECS pattern discovery.

- For more information on scheduling a cloud discovery, see Create a discovery schedule in Cloud Discovery Workspace.
- For more information about creating a serverless schedule for Amazon ECS discovery, see Create a serverless schedule for Amazon ECS discovery.

Verify configuration of VM Object record to find the resource Tags

To enable the pattern to find the resource Tags, the parent ECS table (VM Object) must have a related entry for the cmdb_key_value table:

1. In the ServiceNow AI Platform navigation bar, navigate to **All** > **Configuration** > **Identification/Reconciliation**, and select **CI Identifiers**.
2. Search for and open the VM Object record.
3. Select the **Related Entries** tab and make sure it's configured as shown in the VM Object Related Entry figure. If it isn't, specify a new related entry: In the Related Entries related list, select **New**, fill out the form, and then select **Submit**.

VM Object Related Entry

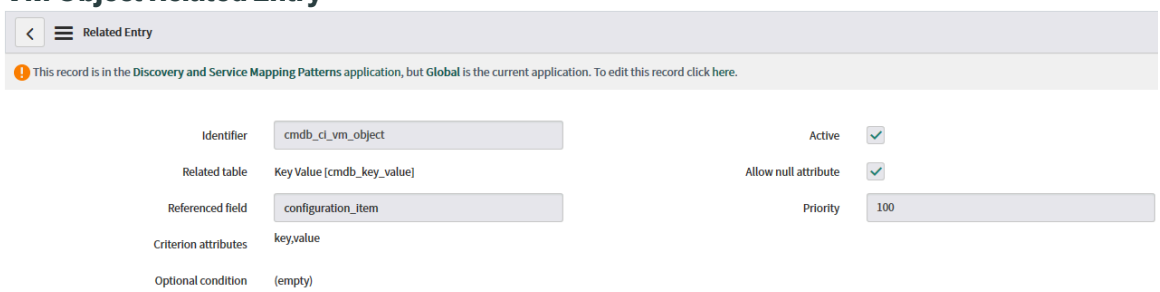

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - ECS pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]

Field	Description
Object ID [object_id]	The Amazon Resource Name (ARN) that identifies the cluster.
Name [name]	A user-generated string used to identify the cluster.
Registered Container Instances Count [registered_container_instances_count]	The number of container instances registered to the cluster. The status of these container instances can be either ACTIVE or DRAINING.

AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]

Field	Description
Object ID [object_id]	The full ARN of the task definition.
Name [name]	The name of the family to which this task definition is registered.
Revision [revision]	The revision of the task in a particular family. This is the version number of the task definition in a family.
Network Mode [network_mode]	The Docker networking mode to use for the containers in the task. Valid values: None, bridge, awsvpc, host.
CPU [cpu]	The number of CPU units used by the task.
Memory [memory]	The amount of memory, in mebibyte (MiB), used by the task.

AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]

Field	Description
Object ID [object_id]	The ARN that identifies the task.
Name [name]	A user-generated string used to identify the task.
Launch Type [launch_type]	The launch type on which the task runs.
Operational status [operational_status]	The health status of the task, determined by the health of the essential containers in it.
CPU [cpu]	The number of CPU units used by the task, as stated in the task definition.
Memory [memory]	The amount of memory, in MiB, used by the task, as stated in the task definition.

AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]

Field	Description
Object ID [object_id]	The ARN that identifies the service.
Name [name]	The name of the service.
Launch Type [launch_type]	The launch type on which the service runs.
Scheduling Strategy [scheduling_strategy]	The scheduling strategy to use for the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Docker Container [cmdb_ci_docker_container]

Field	Description
Container id [container_id]	The ARN of the container.
Name [name]	The name of the container.
Status [status]	The last-known status of the container.
Operational status [operational_status]	The health status of the container.
CPU [cpu]	The number of CPU units set for the container.
Memory [memory]	The hard limit of memory, in MiB, set for the container.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Object ID [object_id]	The ARN that identifies the virtual machine (VM) instance. **Note:** Complete Virtual Machine Instance information is discovered in a different EC2 cloud pattern. In Amazon ECS resource discovery, only the object_id is parsed to aid with the unique identification of the CI.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The Tag key.
Value [value]	The Tag value.
Configuration item [configuration_item]	References the discovered CI that the tag is associated with.

Docker Image [cmdb_ci_docker_image]

Field	Description
Name [name]	Name of the Docker image.
Image Digest [image_digest]	Unique identifier for the Docker image.
Image ID [image_id]	12-character length unique identifier for the Docker image.

Container Repository [cmdb_ci_container_repository]

Field	Description
Name [name]	Fully qualified name of the repository containing the Docker image.

Container Repository Entry [cmdb_ci_container_repository_entry]

Field	Description
Name [name]	The Docker image tag.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Container Repository Entry [cmdb_ci_container_repository_entry] (continued)

Field	Description
Category [category]	Fully qualified name of the repository containing the Docker image.

On the Dependency Views map, you can see all discovered resources in your organization that are managed by Amazon ECS, and the relationships between them.

In this example, the AWS Cloud ECS Cluster CI is hosted on an AWS Datacenter. It runs two AWS Cloud ECS Tasks and two AWS Cloud ECS Services, and uses three Docker Containers.

Dependency Views sample map

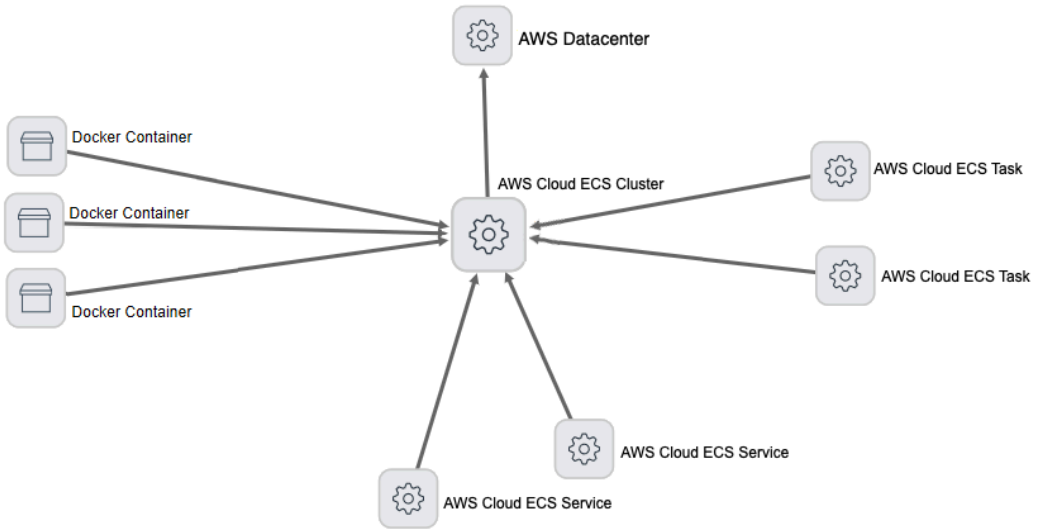

CI relationships

The Amazon AWS - ECS pattern creates the following relationships and references to support Amazon ECS resource discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Virtual Machine Instance [cmdb_ci_vm_instance]	Used by::Uses	AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]	Runs on::Runs	AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]	Uses::Used by	AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]
AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]*	Runs on::Runs	AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]
AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]*	Runs on::Runs	Virtual Machine Instance [cmdb_ci_vm_instance]
AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]	Runs on::Runs	Docker Container [cmdb_ci_docker_container]
AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]	Uses::Used by	AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]
AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]	Used by::Uses	AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]
AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]	Used by::Uses	AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]
Docker Container [cmdb_ci_docker_container]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Docker Container [cmdb_ci_docker_container]	Used by::Uses	AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]
Container Repository [cmdb_ci_container_repository]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Container Repository Entry [cmdb_ci_container_repository_entry]	Hosted on::Hosts	Container Repository [cmdb_ci_container_repository]
Docker Image [cmdb_ci_docker_image]	Provisioned From::Provisioned	Container Repository Entry [cmdb_ci_container_repository_entry]
Docker Image [cmdb_ci_docker_image]	Instantiates::Instantiated by	Docker Container [cmdb_ci_docker_container]

*The launch type determines which relationship is created. For Fargate launch type tasks, the relationship is created with the AWS Cloud ECS Cluster. For EC2 launch type tasks, the relationship is created with the Virtual Machine Instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Virtual Machine Instance [cmdb_ci_vm_instance]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]

AWS tag discovery

The Amazon AWS - ECS pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The Tag key.
Value [value]	The Tag value.
Configuration item [configuration_item]	References one of the following CIs that the tag is associated with: • AWS Cloud ECS Cluster [cmdb_ci_cloud_ecs_cluster] • Virtual Machine Instance [cmdb_ci_vm_instance] • AWS Cloud ECS TaskDefinition [cmdb_ci_cloud_ecs_task_definition] • AWS Cloud ECS Task [cmdb_ci_cloud_ecs_task] • AWS Cloud ECS Service [cmdb_ci_cloud_ecs_service]

Create a serverless schedule for Amazon ECS discovery

Create a serverless schedule to discover Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Container Registry (ECR) resources in a standalone discovery.

Before you begin

- Verify you have an AWS service account ID and have created AWS credentials. For more information, see the prerequisites section in Amazon ECS resource discovery with Patterns.
- Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Discovering ECR resources requires running the Amazon AWS - ECS pattern before using Aqua Trivy to scan container images. To discover only ECS and ECR resources without performing a full AWS pattern discovery, create a serverless discovery schedule using the Amazon AWS - ECS pattern.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule New record form

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type, which should be Serverless.
MID server	Name of the MID Server to use for this schedule.
Active	Enables this schedule for discovery.
Run	Determines the run schedule of the discovery.

   c. Select **Submit**.
3. Create the execution pattern.
   a. In the Discovery Schedules page, select the record you created.
   b. In the **Serverless Execution Patterns** tab, select **New**.
   c. On the form, fill in the fields.

Serverless Execution Pattern New record form

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern to be used for this schedule, which should be Amazon AWS - ECS.
Proxy Host	Fully qualified domain name of the machine on which you're installing the proxy server, which should be Global.
Active	Enables this schedule for discovery.

   d. Select **Submit**.
4. Set the pattern launcher parameters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.

b. On the form, fill in the fields.

Parameter	Value
cloud_account_id	AWS cloud service account ID.
cloud_cred_id	Sys_id of the AWS credentials you created on the ServiceNow AI Platform.
cloud_datacenter_type	Name of the AWS Datacenter table, which should be cmdb_ci_aws_datacenter.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

Related topics

Amazon ECS resource discovery with Patterns

Scan container images

Amazon Elastic File System (Amazon EFS) pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Elastic File System - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Elastic File System [cmdb_aws_elastic_file_system_file_system]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) for the file system.
Name [name]	The assigned name of the Amazon EFS file system.
Configuration Item [configuration_item]	References the Cloud Directory [cmdb_ci_cloud_directory] table.
File System ID [file_system_id]	The unique identifier of the Amazon EFS file system.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Elastic File System - Extended Inventory (LP) pattern.

Cloud Directory [cmdb_ci_cloud_directory]

Field	Description
Object ID [object_id]	The ARN for the file system.
Name [name]	The assigned name of the Amazon EFS file system.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **AWS::Elastic File System::File System**.

CI relationships

The Amazon AWS - Elastic File System - Extended Inventory (LP) pattern creates these relationships to support Amazon EFS discovery.

CI	Relationship	CI
Cloud Directory [cmdb_ci_cloud_directory]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
AWS Elastic File System [cmdb_aws_elastic_file_system_file_system]	References	Cloud Directory [cmdb_ci_cloud_directory]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon ElastiCache discovery with Patterns

The ServiceNow® Discovery application uses the Amazon ElastiCache discovery pattern to find Redis clusters deployed as part of the ElastiCache service. ElastiCache is compatible with both Redis and Memcached, but the pattern collects information only for Redis.

Amazon ElastiCache is a fully managed, in-memory caching service available on AWS.

Prerequisites

- Install the Discovery and Service Mapping Patterns application from the ServiceNow Store.
- Configure an AWS Service Account with valid credentials on the ServiceNow instance with permission to run the following APIs:
  - `https://elasticache.{region}.amazonaws.com?Action=DescribeReplicationGroups&Version=2015-02-02`
  - `https://tagging.{region}.amazonaws.com {"ResourceTypeFilters":[]}`

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

Replication groups are stored in the Cloud Database Clusters [cmdb_ci_cloud_db_cluster] table and node group members are stored in the Cloud Databases [cmdb_ci_cloud_database] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Amazon resource name	Name of the resource for each replication group and each node group member.
FQDN	Fully qualified domain name for each node group member.
Port	Port used by each node group member.

Amazon ElastiCache Snapshot pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - ElastiCache Snapshot - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS ElastiCache - Snapshot [cmdb_aws_elasticache_snapshot]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the ElastiCache snapshot.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS ElastiCache - Snapshot [cmdb_aws_elasticache_snapshot] (continued)

Field	Description
Name [name]	The name of the snapshot.
Status [status]	The current status of the snapshot.
Cache Cluster ID [cache_cluster_id]	The identifier of the cache cluster from which the snapshot was taken.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Engine [engine]	The type of caching engine used. For example: Redis or Memcached.
Engine Version [engine_version]	The version of the cache engine used for the snapshot.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - ElastiCache Snapshot - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the ElastiCache snapshot.
Name [name]	The name of the snapshot.
Resource type [resource_type]	Type of resource. The value is set to **AWS::ElastiCache::Snapshot**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - ElastiCache Snapshot - Extended Inventory (LP) pattern creates these relationships to support Amazon ElastiCache Snapshot discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS ElastiCache - Snapshot [cmdb_aws_elasticache_snapshot]	References	Cloud Resource [cmdb_ci_cmp_resource]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon EMR Cluster pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon EMR Clusters on your cloud environment running on EC2. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - EMR Cluster - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS EMR Cluster [cmdb_aws_emr_cluster]

Field	Description
Name [name]	Name of the EMR cluster.
Object ID [object_id]	The Amazon Resource Name (ARN) of the EMR cluster.
Cluster ID [cluster_id]	The unique identifier of the EMR cluster.
State [state]	Current state of the EMR cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS EMR Cluster [cmdb_aws_emr_cluster] (continued)

Field	Description
	For example: RUNNING, WAITING, or TERMINATING.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - EMR Cluster - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the EMR cluster.
Object ID [object_id]	The ARN of the EMR cluster.
Resource type [resource_type]	Type of resource. The value is set to **AWS::EMR::Cluster**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - EMR Cluster - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon EMR Cluster discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS EMR Cluster [cmdb_aws_emr_cluster]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Tag discovery

The Amazon AWS - EMR Cluster - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon EventBridge Event Bus pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon EventBridge Event Buses on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - EventBridge Event Bus - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS EventBridge Event Bus [cmdb_aws_eventbridge_event_bus]

Field	Description
Name [name]	Name of the event bus.
Object ID [object_id]	The Amazon Resource Name (ARN) of the event bus.
Policy [policy]	The resource-based policy that controls who can send events to the event bus.
Configuration Item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - EventBridge Event Bus - Extended Inventory (LP) pattern.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Field	Description
Name [name]	Name of the event bus.
Object ID [object_id]	The ARN of the event bus.
Type [type]	Type of resource. The value is set to **AWS::EventBridge::EventBus**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - EventBridge Event Bus - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon EventBridge Event Bus discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Messaging Service [cmdb_ci_cloud_messaging_service]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS EventBridge Event Bus [cmdb_aws_eventbridge_event_bus]	Configuration Item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

AWS Tag discovery

The Amazon AWS - EventBridge Event Bus - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Amazon FSx Backup pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon FSx Backups on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - FSx Backup - Extended Inventory (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS FSx Backup [cmdb_aws_fsx_backup]

Field	Description
Name [name]	The unique ID of the backup.
Object ID [object_id]	The Amazon Resource Name (ARN) of the backed-up resource.
Type [type]	The type of backup. For example: AUTOMATIC, USER_INITIATED, or AWS_BACKUP.
Owner ID [owner_id]	The ID of the AWS account that owns the backup.
Source Backup ID [source_backup_id]	The ID of the source backup. Populated for cross-region backup copies.
Source Backup Region [source_backup_region]	The AWS Region where the source backup was created. Populated for cross-region backup copies.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - FSx Backup - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	The unique ID of the backup.
Object ID [object_id]	The ARN of the backed-up resource.
Resource type [resource_type]	Type of resource. The value is set to **AWS::FSx::Backup**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - FSx Backup - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon FSx Backup discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS FSx Backup [cmdb_aws_fsx_backup]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - FSx Backup - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon FSx File System pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon FSx File Systems on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - FSx File System - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS FSx File System [cmdb_aws_fsx_file_system]

Field	Description
Name [name]	The ID of the file system, used as the display name.
Object ID [object_id]	The Amazon Resource Name (ARN) of the file system.
File System Id [file_system_id]	The ID of the file system.
File System Type [file_system_type]	The type of file system. For example: WINDOWS, LUSTRE, ONTAP, or OPENZFS.
File System Type Version [file_system_type_version]	The version of the file system type.
Lifecycle [lifecycle]	The lifecycle status of the file system. For example: AVAILABLE, CREATING, DELETING, FAILED, MISCONFIGURED, or UPDATING.
Owner Id [owner_id]	The ID of the AWS account that owns the file system.
DNS Name [dns_name]	The Domain Name System (DNS) name for the file system.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - FSx File System - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	The ID of the file system, used as the display name.
Object ID [object_id]	The ARN of the file system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Resource type [resource_type]	Type of resource. The value is set to **AWS::FSx::FileSystem**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - FSx File System - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon FSx File System discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS FSx File System [cmdb_aws_fsx_file_system]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - FSx File System - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon MQ Broker pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon MQ Brokers on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - MQ Broker - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS MQ Broker [cmdb_aws_mq_broker]

Field	Description
Name [name]	Name of the broker.
Object ID [object_id]	The Amazon Resource Name (ARN) of the broker.
Broker Id [broker_id]	The unique ID of the broker.
Broker State [broker_state]	The current state of the broker. For example: RUNNING, CREATION_IN_PROGRESS, DELETION_IN_PROGRESS, or REBOOT_IN_PROGRESS.
Engine Type [engine_type]	The type of broker engine. For example: ACTIVEMQ or RABBITMQ.
Engine Version [engine_version]	The version of the broker engine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS MQ Broker [cmdb_aws_mq_broker] (continued)

Field	Description
Host Instance Type [host_instance_type]	The broker's instance type, which determines the computing resources assigned to the broker.
Configuration Item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - MQ Broker - Extended Inventory (LP) pattern.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Field	Description
Name [name]	Name of the broker.
Object ID [object_id]	The ARN of the broker.
Type [type]	Type of resource. The value is set to **AWS::MQ::Broker**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - MQ Broker - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon MQ Broker discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Messaging Service [cmdb_ci_cloud_messaging_service]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS MQ Broker [cmdb_aws_mq_broker]	Configuration Item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Tag discovery

The Amazon AWS - MQ Broker - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Amazon MQ Configuration pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon MQ Configurations on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - MQ Configuration - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS MQ Configuration [cmdb_aws_mq_configuration]

Field	Description
Name [name]	Name of the MQ configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS MQ Configuration [cmdb_aws_mq_configuration] (continued)

Field	Description
Object ID [object_id]	The Amazon Resource Name (ARN) of the MQ configuration.
Id [id]	The unique ID of the MQ configuration.
Engine Type [engine_type]	The type of broker engine the configuration applies to. For example: ACTIVEMQ or RABBITMQ.
Engine Version [engine_version]	The version of the broker engine the configuration applies to.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - MQ Configuration - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the MQ configuration.
Object ID [object_id]	The ARN of the MQ configuration.
Resource type [resource_type]	Type of resource. The value is set to **AWS::MQ::Configuration**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - MQ Configuration - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon MQ Configuration discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS MQ Configuration [cmdb_aws_mq_configuration]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - MQ Configuration - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon MWAA Environment pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Managed Workflows for Apache Airflow Environment - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Managed Workflows for Apache Airflow Environment [cmdb_aws_managed_workflows_for_apache_airflow_environment]

Field	Description
Name [name]	The display name or identifier for the Airflow environment.
Airflow Version [airflow_version]	The specific version of Apache Airflow running in the environment.
Configuration Item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the Airflow environment.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Managed Workflows for Apache Airflow Environment - Extended Inventory (LP) pattern.

Cloud App Server [cmdb_ci_cloud_appserver]

Field	Description
Name [name]	The display name or identifier for the Airflow environment.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Object ID [object_id]	The ARN uniquely identifying the Airflow environment.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **AWS::MWAA::Environment**.

CI relationships

The Amazon AWS - Managed Workflows for Apache Airflow Environment - Extended Inventory (LP) pattern creates these relationships to support Amazon MWAA Environment discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud App Server [cmdb_ci_cloud_appserver]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Managed Workflows for Apache Airflow Environment [cmdb_aws_managed_workflows_for_apache_airflow_environment]	References	Cloud App Server [cmdb_ci_cloud_appserver]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon Relational Database Service (Amazon RDS) discovery with Patterns

Discovery and Service Mapping Patterns uses the Amazon AWS Relational Database Service pattern to find Amazon RDS components in your environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Amazon RDS is a managed relational database service by Amazon Web Services (AWS). Amazon RDS enables you to create and manage relational database instances in the cloud.

The ServiceNow AI Platform supports Amazon Commercial Cloud Services (C2S) for discovery and cloud governance. C2S is the government program that brings AWS into the Intelligence Community.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Support for AWS services in the China region

The latest version of Discovery and Service Mapping Patterns supports discovering AWS services in the China region. You can discover these services on the ServiceNow AI Platform, starting from Xanadu Patch 3 and Washington DC Patch 9 instances.

Discovering AWS services in the China region requires using a datacenter URL when setting up an AWS service account. For example: `https://organizations.cn-northwest-1.amazonaws.com.cn`.

- To learn more about AWS master account and sub-account support in the China region, see KB1704526.
- To identify AWS patterns supported in the China region, refer to the Cloud Discovery patterns spreadsheet. The **AWS China Region Support** column has a Yes value for supported patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS Relational Database Service pattern.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	Unique identifier for the RDS database instance, assigned by AWS.
Name [name]	Name of the database. If no database name is provided, the value is the same as the Object ID.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) endpoint used to connect to the database instance.
TCP port(s) [tcp_port]	Port number on which the database instance accepts connections.
Type [type]	Database engine type. For example: mysql, postgres, or aurora.
Version [version]	Database engine version number.
Category [category]	DB instance class that specifies the compute and memory capacity. For example: db.t3.micro or db.m5.large.
IP Address [ip_address]	IP address resolved from the FQDN.
State [state]	State of the database instance. For example: Available or Terminated.
Automated Backups [automated_backup]	Indicates whether automated backups are enabled and the retention period. Possible values are either **Disabled** or **Enabled (X Days)**, where X represents the backup retention period in days.
Deletion Protection Enabled [deletion_protection_enabled]	Indicates whether deletion protection is enabled for the database instance. Possible values are **true** or **false**. • **true**: Deletion protection is enabled and the database can't be deleted • **false**: Deletion protection is disabled
IAM Authentication Enabled [iam_authentication_enabled]	Indicates whether AWS Identity and Access Management (IAM) database authentication is enabled. Possible values are **true** or **false**. • **true**: IAM authentication is enabled and you can use IAM credentials to authenticate to the database • **false**: IAM authentication is disabled

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase [cmdb_ci_cloud_database] (continued)

Field	Description
Storage Encrypted [storage_encrypted]	Indicates whether the database storage is encrypted.
Multiple Availability Zones [multi_az]	Indicates whether the database instance is deployed across multiple Availability Zones for high availability.
Publicly Enabled [publicly_enabled]	Indicates whether the database instance is publicly accessible from the internet.
Vendor [vendor]	Database vendor. Value is set to **Amazon**.
Install Status [install_status]	Install status of the database instance. Default value is Installed.
Operational status [operational_status]	Operational status of the database instance. Default value is Operational.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IP address resolved from the database instance FQDN.
Netmask [netmask]	Network mask for the IP address. Value is set to 0.0.0.0.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Object ID [object_id]	Availability Zone identifier.
Name [name]	Availability Zone name.

Cloud Network [cmdb_ci_network]

Field	Description
Object ID [object_id]	Virtual Private Cloud (VPC) identifier for the network containing the database instance.

CI relationships

The Amazon AWS Relational Database Service pattern creates these relationships to support Amazon RDS discovery.

CI	Relationship	CI
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud DataBase [cmdb_ci_cloud_database]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
AWS Datacenter [cmdb_ci_aws_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Cloud DataBase [cmdb_ci_cloud_database]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Cloud DataBase [cmdb_ci_cloud_database]	Uses::Used by	Cloud Network [cmdb_ci_network]
Cloud Network [cmdb_ci_network]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Key Value [cmdb_key_value]	References	Cloud DataBase [cmdb_ci_cloud_database]

AWS Tag discovery

The Amazon AWS Relational Database Service pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

Amazon RDS DB Snapshot pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Relational Database Services DB Snapshot - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Relational Database Services - DB Snapshot [cmdb_aws_relational_database_services_db_snapshot]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) that uniquely identifies the database snapshot.
Name [name]	A unique identifier for the RDS database snapshot.
Availability Zone [availability_zone]	The availability zone where the database snapshot was created.
Status [status]	The current state or progress of the database snapshot.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
DB Instance Identifier [db_instance_identifier]	The identifier of the RDS DB instance associated with the snapshot.
Engine [engine]	The database engine used for the snapshot. For example: MySQL or PostgreSQL.
Engine Version [engine_version]	The version of the database engine for the snapshot.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Relational Database Services DB Snapshot - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The ARN that uniquely identifies the database snapshot.
Name [name]	A unique identifier for the RDS database snapshot.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Resource type [resource_type]	Type of resource. The value is set to **AWS::RDS::DBSnapshot**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Relational Database Services DB Snapshot - Extended Inventory (LP) pattern creates these relationships to support Amazon RDS DB Snapshot discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Relational Database Services - DB Snapshot [cmdb_aws_relational_database_services_db_snapshot]	References	Cloud Resource [cmdb_ci_cmp_resource]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon Redshift discovery with Patterns

The ServiceNow® Discovery application uses the Amazon AWS Redshift pattern to find Redshift data warehouse services available on Amazon Web Services (AWS). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Redshift runs as a cluster of compute nodes with an additional leader node. The leader node acts as an interface with client applications and coordinates communication between the compute nodes.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS Redshift pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon Redshift [cmdb_ci_aws_redshift]

Field	Description
Name [name]	Name of the Redshift cluster.
Object ID [object_id]	A unique identifier constructed from the Redshift namespace Amazon Resource Name (ARN) and the cluster name in the following format: **arn:aws:redshift:<region>:<account-id>:cluster:<cluster-name>**.
Start date [start_date]	Timestamp when the Redshift cluster was created.
Cluster Availability Status [cluster_availability_status]	Current availability status of the Redshift cluster. For example: Available, Unavailable, or Maintenance.
VPC ID [vpc_id]	ID of the virtual private cloud (VPC) in which the Redshift cluster is deployed.
Node Count [node_count]	Total number of compute nodes in the Redshift cluster.
Node Type [node_type]	Type of compute nodes used in the Redshift cluster.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) of the Redshift cluster endpoint.
TCP port(s) [tcp_port]	TCP port used by the Redshift cluster for client connections.
IP Address [ip_address]	IP address associated with the Redshift cluster endpoint.
Availability Zone [availability_zone]	References the Availability Zone [cmdb_ci_availability_zone] table.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the avaiability zone.
Object ID [object_id]	Name of the avaiability zone.

CI relationships

Discovery creates these relationships to support the Amazon Redshift discovery.

CI	Relationship	CI
Amazon Redshift [cmdb_ci_aws_redshift]	References	Availability Zone [cmdb_ci_availability_zone]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
AWS Datacenter [cmdb_ci_aws_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Key Value [cmdb_key_value]	References	Amazon Redshift [cmdb_ci_aws_redshift]

Tag discovery

The Amazon AWS Redshift pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Amazon Redshift [cmdb_ci_aws_redshift] table.

Amazon Redshift Serverless Namespace pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Redshift Serverless Namespace - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Redshift Serverless - Namespace [cmdb_aws_redshift_serverless_namespace]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the namespace.
Name [name]	The name assigned to the Redshift Serverless namespace.
Configuration Item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.
DB Name [db_name]	The name of the database associated with the namespace.
Namespace ID [namespace_id]	The unique identifier assigned to the Redshift Serverless namespace.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Redshift Serverless Namespace - Extended Inventory (LP) pattern.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the namespace.
Name [name]	The name assigned to the Redshift Serverless namespace.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **AWS::Redshift Serverless::Namespace**.

CI relationships

The Amazon AWS - Redshift Serverless Namespace - Extended Inventory (LP) pattern creates these relationships to support Amazon Redshift Serverless Namespace discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Redshift Serverless - Namespace [cmdb_aws_redshift_serverless_namespace]	References	Cloud DataBase [cmdb_ci_cloud_database]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon Redshift Serverless Snapshot pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Redshift Serverless Snapshot - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Redshift Serverless - Snapshot [cmdb_aws_redshift_serverless_snapshot]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the resource.
Name [name]	The user-defined name of the resource.
Status [status]	The current state or condition of the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Redshift Serverless Snapshot - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the resource.
Name [name]	The user-defined name of the resource.
Resource type [resource_type]	Type of resource. The value is set to **AWS::Redshift Serverless::Snapshot**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Redshift Serverless Snapshot - Extended Inventory (LP) pattern creates these relationships to support Amazon Redshift Serverless Snapshot discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Redshift Serverless - Snapshot [cmdb_aws_redshift_serverless_snapshot]	References	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon Redshift Serverless Workgroup pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Redshift Serverless Workgroup - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Redshift Serverless - Workgroup [cmdb_aws_redshift_serverless_workgroup]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) that uniquely identifies the resource.
Name [name]	The name assigned to the resource or workgroup.
Status [status]	The current operational status of the resource or workgroup.
Configuration Item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.
Workgroup ID [workgroup_id]	The unique identifier for the Redshift Serverless workgroup.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Redshift Serverless Workgroup - Extended Inventory (LP) pattern.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	The ARN that uniquely identifies the resource.
Name [name]	The name assigned to the resource or workgroup.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **AWS::Redshift Serverless::Workgroup**.

CI relationships

The Amazon AWS - Redshift Serverless Workgroup - Extended Inventory (LP) pattern creates these relationships to support Amazon Redshift Serverless Workgroup discovery.

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Redshift Serverless - Workgroup [cmdb_aws_redshift_serverless_workgroup]	References	Cloud DataBase [cmdb_ci_cloud_database]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon Route 53 pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon Route 53 domain name systems (DNS) and aliases on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS Route53 pattern.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	The DNS domain name for the hosted zone.
IP Address [ip_address]	The IP address resolved from the DNS name.

DNS Alias [cmdb_ci_dns_alias]

Field	Description
Name [name]	The alias name for the DNS record set.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Object ID [object_id]	The unique identifier (ID) for the cloud service account.
Account Id [account_id]	The unique ID for the cloud service account.

CI relationships

The Amazon AWS Route53 pattern creates these relationships to support Amazon Route 53 discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
DNS Name [cmdb_ci_dns_name]	Used by::Uses	DNS Alias [cmdb_ci_dns_alias]
DNS Name [cmdb_ci_dns_name]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Amazon Simple Storage Service (Amazon S3) discovery with Patterns

The ServiceNow Discovery application uses the Amazon AWS S3 pattern to find public and non-public storage buckets of Amazon Simple Storage Service. The pattern uses a set of REST API calls to find these resources. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

In addition to the discovery schedule, the ServiceNow instance triggers the Amazon AWS S3 pattern when AWS Config sends an event to the ServiceNow instance. Config is a service that continuously monitors AWS resources and sends events to the ServiceNow every time it senses that a resource configuration has changed.

Amazon AWS S3 pattern can also discover AWS S3 buckets located in the AWS GovCloud (US) region. For AWS S3 located in the AWS GovCloud (SU), the pattern cannot discover AWS S3 outgoing connections to other cloud resources like Lambda function. Discovery of outgoing connections for AWS S3 will be supported in future releases.

You can use this pattern on the ServiceNow platform using Jakarta Patch 10, Kingston Patch 8, or London Patch 1.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- If you use Identity and Access Management (IAM) to manage users in Amazon Web Services (AWS) platform, ensure that you created a user policy for the AWS user. For more information, refer to Create an IAM user policy for Cloud Provisioning and Governance in the ServiceNow Cloud Provisioning and Governance documentation.
- Configure the AWS service account.
- Configure AWS credentials, using a secret key and an access key.
- AmazonS3ReadOnlyAccess policy user permissions.
- Discover Logical Datacenters hosting S3 buckets as described in Run Discovery on an AWS cloud service account in the ServiceNow Cloud Provisioning and Governance documentation.
- Create a schedule for the relevant AWS service account as described in Schedule Discovery on a service account in the ServiceNow Cloud Provisioning and Governance documentation.
- To enable alert-based discovery, configure the AWS event service as described in Set up AWS event processing for Discovery and Service Mapping in the ServiceNow Cloud Provisioning and Governance documentation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For Cloud Discovery, download the Discovery and Service Mapping Patterns from the ServiceNow Store.
- When installing the MID Server, ensure that the host machine meets or exceeds the MID Server system requirements published on the ServiceNow documentation site.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

Table and field	Description
Cloud Object Storage [cmdb_ci_cloud_object_storage]	
Bucket name [bucket_name]	The name of the AWS S3 bucket. You can discover all buckets hosted on a Logical Datacenter during one discovery process.
Creation date [creation date]	Creation date of the AWS S3 bucket.
Cloud provider [cloud_provider]	Amazon AWS
Object ID [object_id]	The Amazon Resource Name (ARN), expressed in this format: *arn:aws:s3:::<bucket name>*
Service name [service_name]	Set to **S3** by default.
Owner [owner]	Account ID in which the bucket exists.
Encryption type [encryption_type]	Type of encryption for this bucket. The choices are: • None • AES-256 • AWS-KMS
ACL access type [acl_access_type]	The choices for access control types are: • Public • Not Public
Policy access type [policy_access_type]	The choices for policy access types are: • Public • Not Public

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Replication destination [replication_destination]	Bucket ARN of the Replication Destination bucket.

Note: In Discovery and Service Mapping Patterns versions before 1.30.2, the Owner field displays the bucket owner's name.

On the Dependency Views map you can see all discovered S3 buckets in your organization, including S3 buckets that you configured as replicas. In this example, Cloud Object Storage 1, Cloud Object Storage 2, and Cloud Object Storage 3 replicate Cloud Object Storage 2. In its turn, Cloud Object Storage 2 replicates Cloud Object Storage 1.

Dependency Views showing AWS S3 components

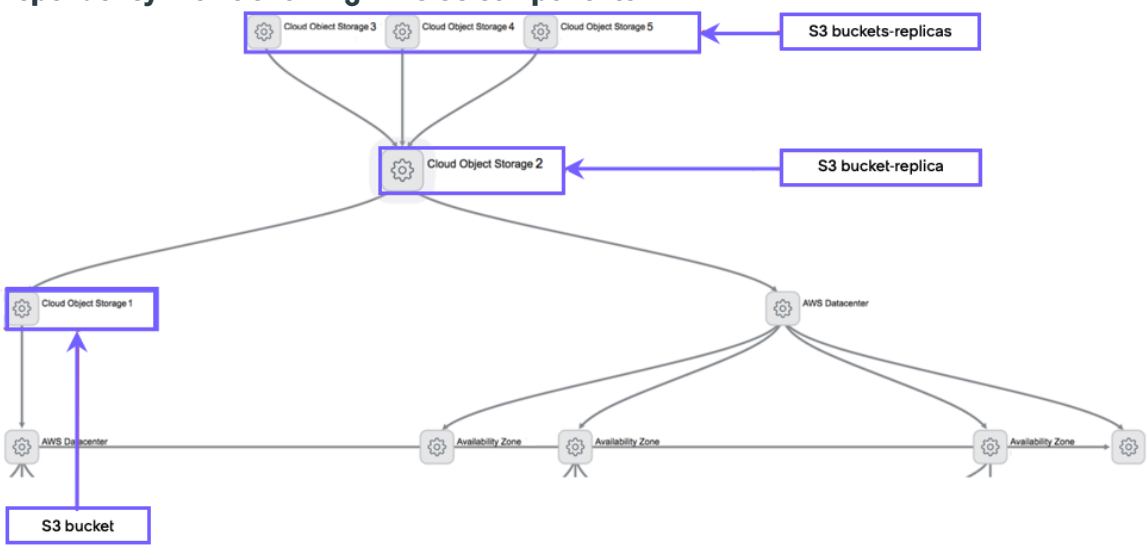

CI relationships

These relationships are created to support AWS S3 discovery:

CI	Relationship	CI
cmdb_ci_cloud_object_storage	Replicates to::Replicated By	cmdb_ci_cloud_object_storage
cmdb_ci_cloud_object_storage	Hosted on::Hosts	cmdb_ci_logical_datacenter

Data discovered by Service Mapping during top-down discovery

Service Mapping can discover AWS S3 outgoing connections to cloud resources, for example Lambda function.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application service map showing an S3 bucket CI

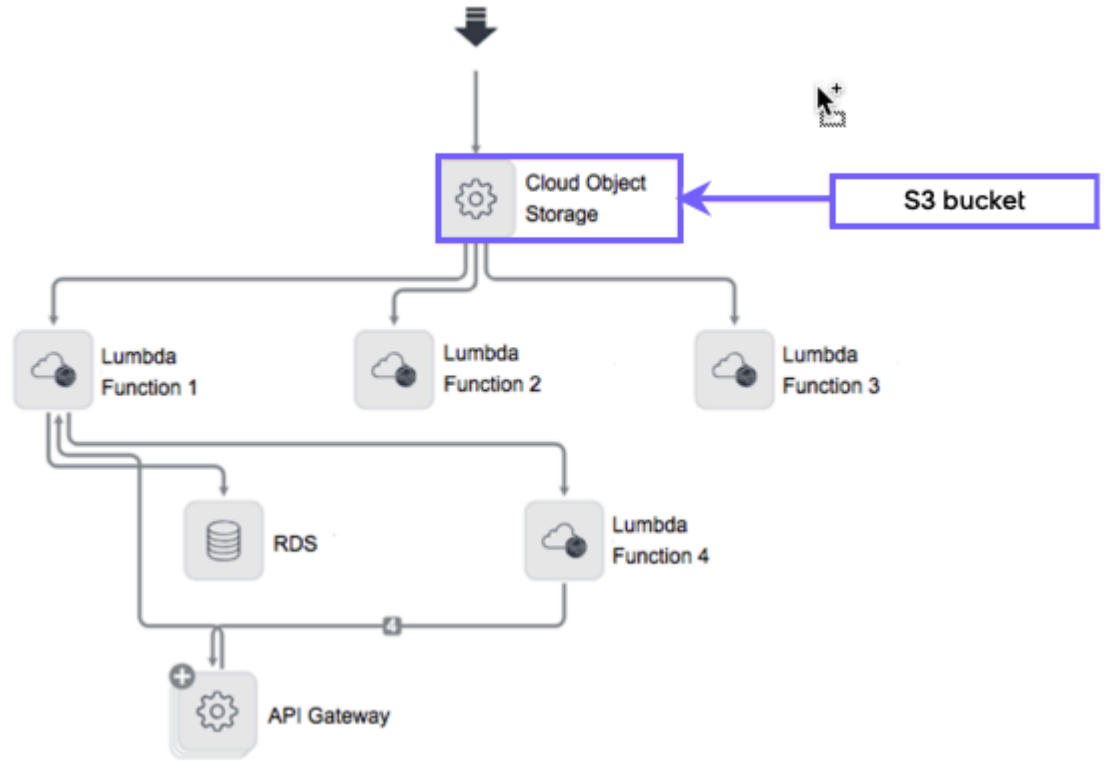

Troubleshooting

If the mapping process does not proceed as you expected, follow the following suggestions.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again. **Note:** If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon S3 Glacier Vault pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon S3 Glacier Vaults on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - S3 Glacier Vault - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS S3 Glacier Vault [cmdb_aws_s3_glacier_vault]

Field	Description
Name [name]	Name of the vault.
Object ID [object_id]	The Amazon Resource Name (ARN) of the vault.
Archives Number [archives_number]	The number of archives stored in the vault.
Size in Bytes [size_in_bytes]	The total size of all archives stored in the vault, in bytes.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - S3 Glacier Vault - Extended Inventory (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the vault.
Object ID [object_id]	The ARN of the vault.
Resource type [resource_type]	Type of resource. The value is set to **AWS::S3::GlacierVault**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - S3 Glacier Vault - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon S3 Glacier Vault discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS S3 Glacier Vault [cmdb_aws_s3_glacier_vault]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - S3 Glacier Vault - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon SageMaker Notebook Instance pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon SageMaker Notebook Instances on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - SageMaker Notebook Instance - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS SageMaker Notebook Instance [cmdb_aws_sagemaker_notebook_instance]

Field	Description
Name [name]	Name of the notebook instance.
Object ID [object_id]	The Amazon Resource Name (ARN) of the notebook instance.
Instance Status [instance_status]	The current status of the notebook instance. For example: InService, Pending, Stopped, Stopping, or Failed.
Kms Key Id [kms_key_id]	The ARN of the AWS KMS key used to encrypt the notebook instance's storage volume. This field is empty if no KMS key is configured.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS SageMaker Notebook Instance [cmdb_aws_sagemaker_notebook_instance] (continued)

Field	Description
Default Code Repository [default_code_repository]	The name or URL of the Git repository associated with the notebook instance as its default code repository.
Root Access [root_access]	Indicates whether root access is enabled for the notebook instance. The value is Enabled or Disabled.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - SageMaker Notebook Instance - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the notebook instance.
Object ID [object_id]	The ARN of the notebook instance.
Resource type [resource_type]	Type of resource. The value is set to **AWS::SageMaker::NotebookInstance**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - SageMaker Notebook Instance - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon SageMaker Notebook Instance discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS SageMaker Notebook Instance [cmdb_aws_sagemaker_notebook_instance]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - SageMaker Notebook Instance - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon SageMaker Training Job pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - SageMaker Training Job - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS SageMaker - Training Job [cmdb_aws_sagemaker_training_job]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the training job.
Name [name]	The user-defined name for the training job.
Configuration Item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.
Secondary Status [secondary_status]	The additional status information providing more detail about the training job's progress.
Training Job Status [training_job_status]	The current status of the training job.
Training Plan ARN [training_plan_arn]	The ARN of the training plan associated with the job.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - SageMaker Training Job - Extended Inventory (LP) pattern.

Cloud App Server [cmdb_ci_cloud_appserver]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the training job.
Name [name]	The user-defined name for the training job.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **AWS::SageMaker::Training Job**.

CI relationships

The Amazon AWS - SageMaker Training Job - Extended Inventory (LP) pattern creates these relationships to support Amazon SageMaker Training Job discovery.

CI	Relationship	CI
Cloud App Server [cmdb_ci_cloud_appserver]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
AWS SageMaker - Training Job [cmdb_aws_sagemaker_training_job]	References	Cloud App Server [cmdb_ci_cloud_appserver]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon Simple Email Service Identity pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon Simple Email Service (SES) Identities on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Simple Email Service Identity - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Simple Email Service Identity [cmdb_aws_simple_email_service_identity]

Field	Description
Name [name]	The email address or domain of the SES identity.
Object ID [object_id]	The email address or domain of the SES identity.
Verification Status [verification_status]	The verification status of the SES identity. For example: Success, Pending, Failed, TemporaryFailure, or NotStarted.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Simple Email Service Identity - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	The email address or domain of the SES identity.
Object ID [object_id]	The email address or domain of the SES identity.
Resource type [resource_type]	Type of resource. The value is set to **AWS::SimpleEmailService::Identity**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Simple Email Service Identity - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon SES Identity discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS Simple Email Service Identity [cmdb_aws_simple_email_service_identity]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Simple Email Service Identity - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon SQS Queue pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon SQS Queues on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - SQS Queue - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS SQS Queue [cmdb_aws_sqs_queue]

Field	Description
Name [name]	Name of the SQS queue.
Object ID [object_id]	The Amazon Resource Name (ARN) of the SQS queue.
Queue URL [queue_url]	The URL of the SQS queue.
FIFO Queue [fifo_queue]	Indicates whether the queue is a First-In-First-Out (FIFO) queue. • FIFO queue: the value is true • Standard queue: the value is false
Maximum Message Size [maximum_message_size]	The maximum size of a message that can be sent to the queue, in bytes.
Configuration Item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - SQS Queue - Extended Inventory (LP) pattern.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Field	Description
Name [name]	Name of the SQS queue.
Object ID [object_id]	The ARN of the SQS queue.
Type [type]	Type of resource. The value is set to **AWS::SQS::Queue**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - SQS Queue - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon SQS Queue discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Messaging Service [cmdb_ci_cloud_messaging_service]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS SQS Queue [cmdb_aws_sqs_queue]	Configuration Item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

AWS Tag discovery

The Amazon AWS - SQS Queue - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Amazon Timestream for InfluxDB Database Instance pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Timestream for InfluxDB Database Instance - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Timestream for InfluxDB Database Instance [cmdb_aws_timestream_for_influxdb_database_instance]

Field	Description
Name [name]	The name assigned to the InfluxDB database instance.
Status [status]	The current state or health of the InfluxDB database instance.
Configuration Item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.
DB Instance Type [db_instance_type]	The compute and memory configuration type for the InfluxDB instance.
DB Storage Type [db_storage_type]	The storage class or type used for the InfluxDB instance.
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the InfluxDB instance.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Timestream for InfluxDB Database Instance - Extended Inventory (LP) pattern.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Name [name]	The name assigned to the InfluxDB database instance.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Object ID [object_id]	The ARN uniquely identifying the InfluxDB instance.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase [cmdb_ci_cloud_database] (continued)

Field	Description
Type [type]	Type of resource. The value is set to **AWS::Timestream::InfluxDBInstance**.

CI relationships

The Amazon AWS - Timestream for InfluxDB Database Instance - Extended Inventory (LP) pattern creates these relationships to support Amazon Timestream for InfluxDB Database Instance discovery.

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Timestream for InfluxDB Database Instance [cmdb_aws_timestream_for_influxdb_database_instance]	References	Cloud DataBase [cmdb_ci_cloud_database]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Amazon VPC Flow Log pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon Virtual Private Cloud (Amazon VPC) Flow Logs on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - VPC Flow Log - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS VPC Flow Log [cmdb_aws_vpc_flow_log]

Field	Description
Name [name]	The flow log ID, used as the display name.
Object ID [object_id]	The unique identifier of the flow log.
Resource ID [resource_id]	The ID of the resource associated with the flow log, such as a VPC, subnet, or network interface.
Deliver Logs Status [deliver_logs_status]	The status of the log delivery. The value is SUCCESS or FAILED.
Flow Log Status [flow_log_status]	The status of the flow log. The value is ACTIVE when the flow log is running.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - VPC Flow Log - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	The flow log ID, used as the display name.
Object ID [object_id]	The unique identifier of the flow log.
Resource type [resource_type]	Type of resource. The value is set to **AWS::VPC::FlowLog**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - VPC Flow Log - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon VPC Flow Log discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS VPC Flow Log [cmdb_aws_vpc_flow_log]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - VPC Flow Log - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Amazon Amazon VPC Managed Prefix List pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon Virtual Private Cloud (Amazon VPC) Managed Prefix Lists on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - VPC Managed Prefix List - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS VPC Managed Prefix List [cmdb_aws_vpc_managed_prefix_list]

Field	Description
Name [name]	Name of the managed prefix list.
Object ID [object_id]	The Amazon Resource Name (ARN) of the managed prefix list.
List ID [list_id]	The unique identifier of the managed prefix list.
State [state]	The current state of the managed prefix list. For example: create-complete, modify-in-progress, or delete-complete.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - VPC Managed Prefix List - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the managed prefix list.
Object ID [object_id]	The ARN of the managed prefix list.
Resource type [resource_type]	Type of resource. The value is set to **AWS::VPC::ManagedPrefixList**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - VPC Managed Prefix List - Extended Inventory (LP) pattern creates the following relationships and references to support Amazon VPC Managed Prefix List discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS VPC Managed Prefix List [cmdb_aws_vpc_managed_prefix_list]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - VPC Managed Prefix List - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Application and Network LB pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Application Load Balancers and Network Load Balancers on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Application and Network LB (LP) pattern.

Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Field	Description
Name [name]	Name of the load balancer.
Fully Qualified Domain Name [fqdn]	IP address of the fully qualified domain name (FQDN) of the load balancer.
Object ID [object_id]	The Amazon Resource Name (ARN) of the load balancer.
DNS Name [dns_name]	The public Domain Name System (DNS) name of the load balancer.
Canonical Hosted Zone Name [canonical_hosted_zone_name]	The name of the Amazon Route 53 hosted zone associated with the load balancer.
Canonical Hosted Zone ID [canonical_hosted_zone_id]	The ID of the Amazon Route 53 hosted zone associated with the load balancer.
State [state]	The state of the load balancer.
Short Description [short_description]	A concatenation of the series of attributes for the load balancer, including LB ARN, VPC ID, type, and zone.
Comments [comments]	Identifier for internal usage (deletion strategy).

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Name of the DNS.
Object ID [object_id]	Name of the DNS.
IP Address [ip_address]	IP address of the DNS.
Comments [comments]	Identifier for internal usage (deletion strategy).

CI relationships

The Amazon AWS - Application and Network LB (LP) pattern creates the following relationships and references to support AWS Application and Network LB discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Subnet [cmdb_ci_cloud_subnet]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Contains::Contained by	DNS Name [cmdb_ci_dns_name]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Contains::Contained by	Compute Security Group [cmdb_ci_compute_security_group]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

AWS Tag discovery

The Amazon AWS - Application and Network LB (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Load Balancer [cmdb_ci_cloud_load_balancer] table.

AWS application ELB Service discovery with Patterns

Discovery uses the Amazon AWS application ELB Service discovery pattern to show all load balancers in your environment in a map.

Elastic Load Balancing (ELB) automatically distributes incoming application traffic across multiple targets.

The Amazon AWS application ELB Service discovery pattern supports top-down discovery. This pattern discovers all load balancers with these characteristics:

- Use ultra-high performance, TLS offloading at scale, and centralized certificate deployment
- Support UDP and static IP addresses for your application

For top-down discovery, the entry point is HTTP(s) Endpoint. Outgoing relations are to Amazon EC2 instances.

The ServiceNow AI Platform supports Amazon Commercial Cloud Services (C2S) for discovery and cloud governance. C2S is the government program that brings AWS into the Intelligence Community.

You can use this pattern on the ServiceNow platform using Kingston or later releases.

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

- Run a horizontal discovery, using Cloud API (CAPI). This is necessary because this discovery pattern supports top-down discovery only. For Cloud Discovery, you can run the horizontal discovery as usual.
- When the load balancers data has been stored in the Configuration Management Database (CMDB), create an application service for each load balancer.
- Set read-only permissions for these REST APIs:
  - https://elasticloadbalancing.<REGION>.amazonaws.com/?Action=DescribeLoadBalancers&LoadBalancerArns.member.1=<LB_ARN>&Version=2015-12-01
  - https://elasticloadbalancing.<REGION>.amazonaws.com/?Action=DescribeLoadBalancers&Version=2015-12-01
  - https://elasticloadbalancing.<REGION>.amazonaws.com/?Action=DescribeListeners&LoadBalancerArn=<LB_ARN>&Version=2015-12-01
  - https://elasticloadbalancing.<REGION>.amazonaws.com/?Action=DescribeRules&ListenerArn=<LISTENER_ARN>&Version=2015-12-01
  - https://elasticloadbalancing.<REGION>.amazonaws.com/?Action=DescribeTargetHealth&TargetGroupArn=<TARGET_GROUP_ARN>&Version=2015-12-01
  - https://ec2.<REGION>.amazonaws.com/?Action=DescribeInstances&Version=2016-11-15&Filter.1.Name=instance-state-name&Filter.1.Value=running&Filter.2.Name=instance-id&Filter.2.Value=<INSTANCE_ID>
- For Cloud Discovery, download the Discovery and Service Mapping Patterns from the ServiceNow Store.
- When installing the MID Server, ensure that the host machine meets or exceeds the MID Server system requirements published on the ServiceNow documentation site.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

Table and field	Description
Load Balancer Services [cmdb_ci_lb_service]	
Fully qualified domain name [fqdn]	Fully Qualified Domain Name (FQDN) of the load balancer. The DNS name.
Name [name]	Name of the load balancer, in the following format: <LB name>:<LB port>.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Serial Number [serial_number]	Serial number of the load balancer, in the following format: ##Object_ID=<LB_ARN>.
Port [port]	Port of the load balancer.
Object ID [object_id]	Unique identifier of the load balancer.
IP Address [ip_address]	IP address of the FQDN.

AWS AppSync API pattern-based discovery

Discovery and Service Mapping Patterns finds AWS AppSync APIs on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - AppSync API - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS AppSync API [cmdb_aws_appsync_api]

Field	Description
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the AppSync API.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS AppSync API [cmdb_aws_appsync_api] (continued)

Field	Description
Name [name]	Name of the AWS AppSync API.
API ID [api_id]	Unique identifier of the AppSync API.
Owner Contact [owner_contact]	Contact information for the owner of the AppSync API.
DNS Realtime [dns_realtime]	Domain Name System (DNS) endpoint for real-time subscriptions.
DNS GraphQL [dns_graphql]	DNS endpoint for GraphQL queries and mutations.
Configuration Item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - AppSync API - Extended Inventory (LP) pattern.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the AppSync API.
Name [name]	Name of the AWS AppSync API, which is the GraphQL interface.
Type [type]	Type of resource. The value is set to **AWS::AppSync::API**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships and references

The Amazon AWS - AppSync API - Extended Inventory (LP) pattern creates the following relationships and references to support AWS AppSync API discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Messaging Service [cmdb_ci_cloud_messaging_service]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS AppSync API [cmdb_aws_appsync_api]	Configuration Item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

AWS Tag discovery

The Amazon AWS - AppSync API - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

AWS Auto Scaling groups discovery with Patterns

The ® Discovery application uses the Amazon AWS - AutoScaling Groups (LP) pattern and Auto Scaling extensions to find AWS Auto Scaling groups. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Starting with Discovery and Service Mapping Patterns version 1.0.95, the pattern extension is automatically displayed and collects the Auto Scaling tags.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Configure AWS Credentials

On your instance, configure credentials of type **AWS Credentials** and set them to **Active**. For more information, see Access setup for AWS service accounts.

Configure a Cloud service account

Configure the AWS service account valid in the ServiceNow instance. For more information, see Set up AWS service accounts.

Configure user permissions to execute the list AWS Auto Scaling groups API call

Provide the user permission to run the following API: https://autoscaling.amazonaws.com/?Action=DescribeAutoScalingGroups&Version=2011-01-01

For more information, see Amazon AWS Cloud components discovery using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure Cloud Discovery schedule

Create a cloud application schedule for discovering AWS Auto Scaling groups and configure the attributes. Set **Discovery** to **Cloud application**. For more information, see Create schedules for discovering cloud resources in Cloud Discovery Workspace

Configure AWS Auto Scaling groups

Navigate to the Instance Scale Set table and verify that the configurations are set as follows:

Verify fields for table- [cmdb_ci_instance_scale_set]

Element Name	Type	Label	Max Length
desired_capacity	integer	Desired Capacity	32
min_size	integer	Minimum Size	32
max_size	integer	Maximum Size	32
health_check_type	string	Health Check Type	128

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - AutoScaling Groups (LP) pattern.

Instance Scale Set [cmdb_ci_instance_scale_set]

Field	Description
Name [name]	The Amazon Resource Name (ARN) of the Auto Scaling group.
Object ID [object_id]	The ARN of the Auto Scaling group.
Install Status [install_status]	The install status of Auto Scaling group. Default value is Installed.
Operational status [operational_status]	The operational status of Auto Scaling group. Default value is Operational.
Desired Capacity [desired_capacity]	The desired size of the group.
Maximum Size [max_size]	The maximum size of the group.
Minimum Size [min_size]	The minimum size of the group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Object ID [object_id]	The Instance ID of the Virtual Machine.

AWS EC2 Auto Scaling– Class model

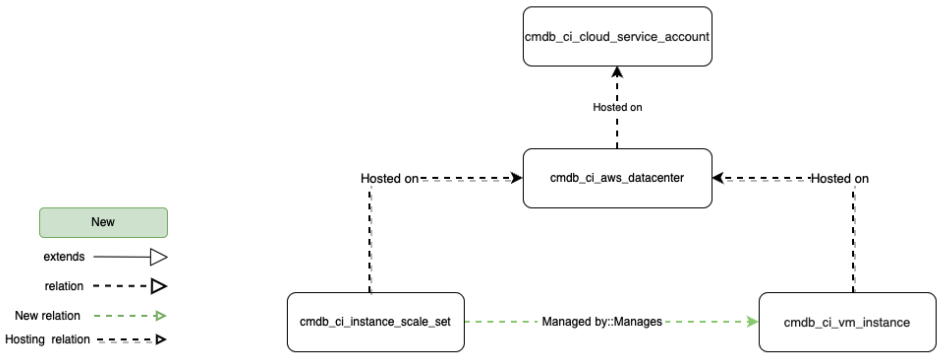

Relationships

Discovery creates these relationships to support the AWS Auto Scaling groups discovery.

CI	Relationship	CI
AWS Datacenter [cmdb_ci_aws_datacenter]	Hosted on::Hosts	Service Account [cmdb_ci_cloud_service_account]
Instance Scale Set [cmdb_ci_instance_scale_set]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Instance Scale Set [cmdb_ci_instance_scale_set]	*Managed by::Manages	VM Instance [cmdb_ci_vm_instance] **Note:** EC2 VM
Key Value [cmdb_key_value]	References	Instance Scale Set [cmdb_ci_instance_scale_set]

*Starting with Discovery and Service Mapping Patterns version 1.29.0, the relationship between Instance Scale Set and VM Instance has changed from **Members::Member of** to **Managed by::Manages**.

Tag discovery

The Amazon AWS - AutoScaling Groups (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The key of the resource tag.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Value [value]	The value of the resource tag.
Configuration item [configuration_item]	References the Instance Scale Set [cmdb_ci_instance_scale_set] table.

Related topics

Detailed information on products discovered by ITOM Visibility

AWS discovery using patterns

AWS Availability Zone pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Availability Zones in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Availability Zone (LP) pattern.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the Availability Zone.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
State [state]	The state of the Availability Zone. Possible values are: available, information, impaired, or unavailable.

CI relationships

The Amazon AWS - Availability Zone (LP) pattern creates the following relationships to support AWS Availability Zone discovery.

CI relationships

CI	Relationship	CI
AWS Datacenter [cmdb_ci_aws_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

AWS Backup Plan pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Backup Plans on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Backup Backup Plan - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Backup Backup Plan [cmdb_aws_backup_backup_plan]

Field	Description
Name [name]	Name of the AWS Backup Plan.
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the backup plan.
Plan ID [plan_id]	Unique identifier assigned to the backup plan.
Version ID [version_id]	Version identifier of the backup plan.
Creator Request ID [creator_request_id]	Unique string that identifies the request to create a backup plan.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Backup Backup Plan - Extended Inventory (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the backup plan.
Name [name]	Name of the AWS Backup Plan.
Resource type [resource_type]	Type of resource. The value is set to **AWS::Backup::BackupPlan**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships and references

The Amazon AWS - Backup Backup Plan - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Backup Plan discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Backup Backup Plan [cmdb_aws_backup_backup_plan]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Backup Backup Plan - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Backup Vault pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Backup Vaults on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Backup Backup Vault - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Backup Backup Vault [cmdb_aws_backup_backup_vault]

Field	Description
Name [name]	Name of the AWS Backup Vault.
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the backup vault.
Max Retention Days [max_retention_days]	Maximum number of days a recovery point can be retained in the backup vault.
Number of Recovery Points [number_of_recovery_points]	Number of recovery points stored in the backup vault.
Vault State [vault_state]	Current state of the backup vault.
Vault Type [vault_type]	Type of the backup vault.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Backup Backup Vault - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the backup vault.
Name [name]	Name of the AWS Backup Vault.
Resource type [resource_type]	Type of resource. The value is set to **AWS::Backup::BackupVault**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Amazon AWS - Backup Backup Vault - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Backup Vault discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Backup Backup Vault [cmdb_aws_backup_backup_vault]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Backup Backup Vault - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Batch Compute Environment pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Batch Compute Environment - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Batch - Compute Environment [cmdb_aws_batch_compute_environment]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the compute environment.
Name [name]	The name of the compute environment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Batch - Compute Environment [cmdb_aws_batch_compute_environment] (continued)

Field	Description
Configuration Item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.
ECS Cluster ARN [ecs_cluster_arn]	The ARN of the Amazon ECS cluster associated with the compute environment.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Batch Compute Environment - Extended Inventory (LP) pattern.

Cloud App Server [cmdb_ci_cloud_appserver]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the compute environment.
Name [name]	The name of the compute environment.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **AWS::Batch::ComputeEnvironment**.

CI relationships

The Amazon AWS - Batch Compute Environment - Extended Inventory (LP) pattern creates these relationships to support AWS Batch Compute Environment discovery.

CI	Relationship	CI
Cloud App Server [cmdb_ci_cloud_appserver]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS Batch - Compute Environment [cmdb_aws_batch_compute_environment]	References	Cloud App Server [cmdb_ci_cloud_appserver]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS classic ELB Service discovery with Patterns

Discovery uses the Amazon AWS classic ELB Service discovery pattern to find all load balancers that use HTTP and HTTPS traffic.

Elastic Load Balancing (ELB) automatically distributes incoming application traffic across multiple targets.

The Amazon AWS classic ELB Service discovery pattern supports top-down discovery. The two possible entry points for top-down discovery are the HTTP(s) Endpoint and TCP Endpoint. Outgoing relations are to Amazon EC2 instances.

The ServiceNow AI Platform supports Amazon Commercial Cloud Services (C2S) for discovery and cloud governance. C2S is the government program that brings AWS into the Intelligence Community.

You can use this pattern on the ServiceNow platform using Jakarta or later releases.

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Run a horizontal discovery, using Cloud API (CAPI). This is necessary because this discovery pattern supports top-down discovery only. For cloud discovery, you can run the horizontal discovery as usual.
- When the load balancers data has been stored in the Configuration Management Database (CMDB), create an application service for each load balancer.
- Set read-only permissions for these REST APIs:
  - https://elasticloadbalancing.<REGION>.amazonaws.com/?Action=DescribeLoadBalancers&Version=2012-06-01&LoadBalancerNames.member.1=<LB_NAME>
  - https://ec2.<REGION>.amazonaws.com/?Action=DescribeInstances&Version=2016-11-15&Filter.1.Name=instance-state-name&Filter.1.Value=running&Filter.2.Name=instance-id&Filter.2.Value=<INSTANCE_ID> (to create relation between LB and EC2 instance)
- For Cloud Discovery, download the Discovery and Service Mapping Patterns from the ServiceNow Store.
- When installing the MID Server, ensure that the host machine meets or exceeds the MID Server system requirements published on the ServiceNow documentation site.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

The discovered data includes the following tables and fields.

Table and field	Description
Load Balancer Services [cmdb_ci_lb_service]	
Fully qualified domain name [fqdn]	Fully Qualified Domain Name (FQDN) of the load balancer. The DNS name.
Name [name]	Name of the load balancer, in the following format: <LB name>:<LB port>.
Back End Port [back_end_port]	Port of the instance.
Server Port [server_port]	Port of the instance.
Port [port]	Port of the load balancer.
Front End Port [front_end_port]	Port of the load balancer.
Object ID [object_id]	LB port.
IP Address [ip_address]	IP address of the FQDN.
Server Protocol [server_protocol]	Protocol of the instance.
IP Address [cmdb_ci_ip_address]	
IP Address [ip_address]	IP address of the load balancer FQDN.
Name [name]	Name of the load balancer, in the following format: <LB name>:<LB port>.
Netmask [netmask]	0.0.0.0

AWS Classic Load Balancer pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Classic Load Balancers on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Classic LB (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Field	Description
Name [name]	Name of the load balancer.
Fully Qualified Domain Name [fqdn]	The DNS name of the load balancer.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
DNS Name [dns_name]	The DNS name of the load balancer.
Canonical Hosted Zone Name [canonical_hosted_zone_name]	The DNS name of the load balancer.
Canonical Hosted Zone ID [canonical_hosted_zone_id]	The ID of the Amazon Route 53 hosted zone for the load balancer.

Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]

Field	Description
Name [name]	IP address of the load balancer.
Object ID [object_id]	IP address of the load balancer.
IP Address [ip_address]	IP address of the load balancer.
Comments [comments]	Identifier for internal usage (deletion strategy).

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Name of the Domain Name System (DNS).
IP Address [ip_address]	IP address of the DNS.
Comments [comments]	Identifier for internal usage (deletion strategy).

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Name [name]	Name of the load balancer pool.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
Comments [comments]	Identifier for internal usage (deletion strategy).

Load Balancer Pool Member [cmdb_ci_lb_pool_member]

Field	Description
Name [name]	Name of the load balancer pool member (known in AWS as a target).
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Name [name]	Name of the load balancer service.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
Port [port]	The port on which the load balancer is listening.
Service port [service_port]	The port on which the load balancer is listening.
Server protocol [service_protocol]	The protocol to use for routing traffic to instances: HTTP, HTTPS, TCP, or SSL.
Comments [comments]	Identifier for internal usage (deletion strategy).

CI relationships and references

The Amazon AWS - Classic LB (LP) pattern creates the following relationships and references to support AWS Classic Load Balancer discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Load Balancer Service [cmdb_ci_lb_service]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Cloud Subnet [cmdb_ci_cloud_subnet]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Load Balancer Pool [cmdb_ci_lb_pool]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Load Balancer Pool [cmdb_ci_lb_pool]	Owns::Owned by	Load Balancer Pool Member [cmdb_ci_lb_pool_member]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Contains::Contained by	DNS Name [cmdb_ci_dns_name]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Owns::Owned by	Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Contains::Contained by	Compute Security Group [cmdb_ci_compute_security_group]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Tag discovery

The Amazon AWS - Classic LB (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Load Balancer [cmdb_ci_cloud_load_balancer] table.

AWS CloudHSM HSM pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - CloudHSM HSM - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS CloudHSM HSM [cmdb_aws_cloudhsm_hsm]

Field	Description
Name [name]	The unique identifier of the HSM within the cluster.
Availability Zone [availability_zone]	The AWS availability zone where the HSM is deployed.
Status [status]	The current operational status of the HSM.
Cluster Id [cluster_id]	The unique identifier of the HSM cluster.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Hsm Type [hsm_type]	The type or model specification of the HSM device.
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the AWS CloudHSM HSM. The ARN consists of the discovered region, account ID, and HSM ID in the following format: **arn:aws:cloudhsm:<region>:<account_id>:hsm/<hsm_id>**.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - CloudHSM HSM - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	The unique identifier of the HSM within the cluster.
Resource type [resource_type]	Type of resource. The value is set to **AWS::CloudHSM::HSM**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Object ID [object_id]	The ARN uniquely identifying the AWS CloudHSM HSM. The ARN consists of the discovered region, account ID, and HSM ID in the following format: **arn:aws:cloudhsm:<region>:<account_id>:hsm/<hsm_id>**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - CloudHSM HSM - Extended Inventory (LP) pattern creates these relationships to support AWS CloudHSM HSM discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS CloudHSM HSM [cmdb_aws_cloudhsm_hsm]	References	Cloud Resource [cmdb_ci_cmp_resource]

AWS CloudTrail Trail pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - CloudTrail Trail - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The Amazon Resource Name (ARN) uniquely identifying the CloudTrail trail.
Name [name]	The name of the CloudTrail trail.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Resource type [resource_type]	Type of resource. The value is set to **AWS::CloudTrail::Trail**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - CloudTrail Trail - Extended Inventory (LP) pattern creates these relationships to support AWS CloudTrail Trail discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

AWS CodeDeploy Deployment pattern-based discovery

Discovery and Service Mapping Patterns finds AWS CodeDeploy Deployments on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - CodeDeploy Deployment - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS CodeDeploy Deployment [cmdb_aws_codedeploy_deployment]

Field	Description
Name [name]	Name of the CodeDeploy application associated with the deployment.
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the CodeDeploy application.
Deployment Group Name [deployment_group_name]	Name of the deployment group associated with the deployment.
Deployment ID [deployment_id]	Unique identifier of the CodeDeploy deployment.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - CodeDeploy Deployment - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the CodeDeploy application.
Name [name]	Name of the CodeDeploy application associated with the deployment.
Resource type [resource_type]	Type of resource. The value is set to **AWS::CodeDeployment::Deployment**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - CodeDeploy Deployment - Extended Inventory (LP) pattern creates the following relationships and references to support AWS CodeDeploy Deployment discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS CodeDeploy Deployment [cmdb_aws_codedeploy_deployment]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - CodeDeploy Deployment - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS CodePipeline Pipeline pattern-based discovery

Discovery and Service Mapping Patterns finds AWS CodePipeline Pipelines on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - CodePipeline Pipeline - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS CodePipeline Pipeline [cmdb_aws_codepipeline_pipeline]

Field	Description
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the pipeline.
Name [name]	Name of the AWS CodePipeline pipeline.
Pipeline Type [pipeline_type]	Type of the pipeline.
Execution Mode [execution_mode]	Execution mode of the pipeline.
Version [version]	Version number of the pipeline.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - CodePipeline Pipeline - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the pipeline.
Name [name]	Name of the AWS CodePipeline pipeline.
Resource type [resource_type]	Type of resource. The value is set to **AWS::CodePipeline::Pipeline**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - CodePipeline Pipeline - Extended Inventory (LP) pattern creates the following relationships and references to support AWS CodePipeline Pipeline discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS CodePipeline Pipeline [cmdb_aws_codepipeline_pipeline]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - CodePipeline Pipeline - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Customer Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds AWS customer gateways on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Customer Gateway (LP) pattern.

Customer Gateway [cmdb_ci_customer_gateway]

Field	Description
Name [name]	Name or ID if no Name is specified for the customer gateway.
Object ID [object_id]	ID of the customer gateway.
Connection Type [connection_type]	Type of VPN connection the customer gateway supports.

Customer Gateway Endpoint [cmdb_ci_endpoint_cust_gateway]

Field	Description
Name [name]	Name or ID if no Name is specified for the customer gateway.
Object ID [object_id]	ID of the customer gateway.

CI relationships

The Amazon AWS - Customer Gateway (LP) pattern creates the following relationships and references to support AWS Customer Gateway discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Customer Gateway [cmdb_ci_customer_gateway]	Hosted on::Hosts	Virtual Machine Instance [cmdb_ci_instance]
Customer Gateway [cmdb_ci_customer_gateway]	Implement End Point To::Implement End Point From	Customer Gateway Endpoint [cmdb_ci_endpoint_cust_gateway]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Customer Gateway [cmdb_ci_customer_gateway]

AWS Tag discovery

The Amazon AWS - Customer Gateway (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Customer Gateway [cmdb_ci_customer_gateway] table.

AWS datacenter pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Regions for your AWS account on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS Datacenter discovery pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Unique identifier (ID) for the AWS service account.
Object ID [object_id]	Unique ID for the AWS service account.
Datacenter Type [datacenter_type]	Type of datacenter. The value is set to **AWS Datacenter [cmdb_ci_aws_datacenter]**.

AWS Datacenter [cmdb_ci_aws_datacenter]

Field	Description
Name [name]	The AWS Region code. For example: us-east-1 or eu-central-1.
Object ID [object_id]	The AWS Region code. For example: us-east-1 or eu-central-1.
Region [region]	The AWS Region code. For example: us-east-1 or eu-central-1.

CI relationships

The Amazon AWS Datacenter discovery pattern creates these relationships to support AWS datacenter discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
AWS Datacenter [cmdb_ci_aws_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

AWS DataSync Task pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - DataSync Task - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS DataSync - Task [cmdb_aws_datasync_task]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the DataSync task.
Name [name]	The name assigned to the AWS DataSync task.
Status [status]	The current operational status of the DataSync task.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS DataSync - Task [cmdb_aws_datasync_task] (continued)

Field	Description
Configuration Item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.
Task Mode [task_mode]	The execution mode of the DataSync task. For example: one-time or scheduled.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - DataSync Task - Extended Inventory (LP) pattern.

Cloud App Server [cmdb_ci_cloud_appserver]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the DataSync task.
Name [name]	The name assigned to the AWS DataSync task.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **AWS::DataSync::Task**.

CI relationships

The Amazon AWS - DataSync Task - Extended Inventory (LP) pattern creates these relationships to support AWS DataSync Task discovery.

CI	Relationship	CI
Cloud App Server [cmdb_ci_cloud_appserver]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AWS DataSync - Task [cmdb_aws_datasync_task]	References	Cloud App Server [cmdb_ci_cloud_appserver]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS discover Organization pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Cloud Organizations in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - discover Organization pattern.

Cloud Organizations [cmdb_ci_cloud_org]

Field	Description
Name [name]	The unique identifier (ID) of the management account of an organization.
Object ID [object_id]	The unique ID of the management account of an organization.
Root ID [root_id]	The unique ID of an organization.
Master Email [master_email]	The email address associated with the AWS account designated as the management account for the organization.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

AWS DMS Endpoints pattern-based discovery

Discovery and Service Mapping Patterns finds AWS DMS endpoints on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - DMS Endpoint - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS DMS Endpoint [cmdb_aws_dms_endpoint]

Field	Description
Name [name]	Identifier of the DMS endpoint.
Object ID [object_id]	Amazon Resource Name (ARN) that uniquely identifies the DMS endpoint.
Username [username]	Username used to connect to the endpoint database.
Database Name [database_name]	Name of the database at the endpoint.
Endpoint Type [endpoint_type]	Type of the endpoint: source or target.
Engine Name [engine_name]	Database engine used by the endpoint.
Kms Key Id [kms_key_id]	ARN of the Key Management Service (KMS) key used to encrypt the connection parameters.
Server Name [server_name]	Name or IP address of the server hosting the endpoint database.
Port [port]	Port used to connect to the endpoint database.
Status [status]	Connection status of the endpoint.
Endpoint Identifier [endpoint_identifier]	Identifier of the DMS endpoint.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - DMS Endpoint - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	ARN that uniquely identifies the DMS endpoint.
Name [name]	Identifier of the DMS endpoint.
Resource type [resource_type]	Type of resource. The value is set to **AWS::DMS::Endpoint**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Amazon AWS - DMS Endpoint - Extended Inventory (LP) pattern creates the following relationships and references to support AWS DMS endpoint discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS DMS Endpoint [cmdb_aws_dms_endpoint]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - DMS Endpoint - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Elastic Beanstalk Application pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Elastic Beanstalk Applications on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Elastic Beanstalk Application - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Elastic Beanstalk Application [cmdb_aws_elastic_beanstalk_application]

Field	Description
Name [name]	Name of the Elastic Beanstalk application.
Object ID [object_id]	The Amazon Resource Name (ARN) of the Elastic Beanstalk application.
Configuration Item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Elastic Beanstalk Application - Extended Inventory (LP) pattern.

Cloud App Server [cmdb_ci_cloud_appserver]

Field	Description
Name [name]	Name of the Elastic Beanstalk application.
Object ID [object_id]	The ARN of the Elastic Beanstalk application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud App Server [cmdb_ci_cloud_appserver] (continued)

Field	Description
Type [type]	Type of resource. The value is set to **AWS::ElasticBeanstalk::Application**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Elastic Beanstalk Application - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Elastic Beanstalk Application discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud App Server [cmdb_ci_cloud_appserver]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Elastic Beanstalk Application [cmdb_aws_elastic_beanstalk_application]	Configuration Item [configuration_item]	Cloud App Server [cmdb_ci_cloud_appserver]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud App Server [cmdb_ci_cloud_appserver]

AWS Tag discovery

The Amazon AWS - Elastic Beanstalk Application - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.

AWS Events pattern-based discovery

Discovery uses event patterns to update Amazon AWS Cloud component data in near real-time. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Events discovered by Discovery during horizontal discovery

Discovery uses patterns to find events created for Amazon AWS Cloud components. If there are events that indicate the change of state in one of the Amazon AWS Cloud components, it triggers discovery of Amazon AWS Cloud components using the patterns.

Patterns used for event discovery

Pattern	CI
Amazon AWS Application and Network LBs Events	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Amazon AWS Classic LB Events	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Amazon AWS Network Events	Cloud Network [cmdb_ci_network]
Amazon AWS Security Group Events	Compute Security Group [cmdb_ci_compute_security_group]
Amazon AWS Storage Events	Storage Volume [cmdb_ci_storage_volume]
Amazon AWS Subnet Events	Cloud Subnet [cmdb_ci_cloud_subnet]
Amazon AWS Virtual Server Events	Virtual Machine Instance [cmdb_ci_vm_instance]

Related topics

- AWS Application and Network LB pattern-based discovery
- AWS Classic Load Balancer pattern-based discovery
- AWS Network pattern-based discovery
- AWS Security Group pattern-based discovery
- AWS Storage pattern-based discovery
- AWS subnet pattern-based discovery
- AWS virtual server pattern-based discovery

AWS Global Accelerator pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Global Accelerator - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Global Accelerator [cmdb_aws_global_accelerator_accelerator]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the Global Accelerator resource.
Name [name]	The name assigned to the AWS Global Accelerator resource.
Status [status]	The current operational state of the Global Accelerator.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Is Enabled [is_enabled]	Indicates whether the Global Accelerator is enabled or disabled.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Global Accelerator - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the Global Accelerator resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Name [name]	The name assigned to the AWS Global Accelerator resource.
Resource type [resource_type]	Type of resource. The value is set to **AWS::GlobalAccelerator::Accelerator**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Global Accelerator - Extended Inventory (LP) pattern creates these relationships to support AWS Global Accelerator discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
AWS Global Accelerator [cmdb_aws_global_accelerator_accelerator]	References	Cloud Resource [cmdb_ci_cmp_resource]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

AWS Glue Database pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Glue Databases on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Glue Database - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Glue Database [cmdb_aws_glue_database]

Field	Description
Name [name]	Name of the Glue database.
Object ID [object_id]	The Amazon Resource Name (ARN) of the Glue database, constructed from the region, account ID, and database name, in the following format: `arn:aws:glue:{region}:{account-id}:database/{database-name}`.
Catalog ID [catalog_id]	The ID of the Data Catalog in which the database resides.
Location Uri [location_uri]	The location of the database, for example an HDFS path.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Glue Database - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the Glue database.
Object ID [object_id]	The ARN of the Glue database, constructed from the region, account ID, and database name, in the following format: `arn:aws:glue:{region}:{account-id}:database/{database-name}`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Resource type [resource_type]	Type of resource. The value is set to **AWS::Glue::Database**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Glue Database - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Glue Database discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Glue Database [cmdb_aws_glue_database]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Glue Database - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS hardware type pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon EC2 instance types on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery uses the Amazon AWS - Hardware Type (LP) pattern to discover AWS EC2 instance types from your AWS account.

When using the Hardware Type [cmdb_ci_compute_template] table to store the hardware types, you may notice an unusually large number of records. To avoid this issue, you can store the discovered hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. When enabled, Discovery launches the Amazon AWS - Cloud Hardware Type (LP) pattern to populate the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. For more information, see Enable the Cloud Hardware Type class extension.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the AWS hardware type patterns.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Name [name]	Name of the instance type. For example: t2.micro or m5.large
Object ID [object_id]	Name of the instance type. For example: t2.micro or m5.large
Cores [cores]	The default number of cores for the instance type.
Memory MB [memory_mb]	The size of the memory for the instance type, in megabytes (MB).
vCPUs [vcpus]	The default number of Virtual Central Processing Units (vCPUs) for the instance type.
Local Storage GB [local_storage_gb]	The total size of the local instance storage for the instance type, in gigabytes (GB).
Provider [provider]	The cloud provider, which is **AWS**. This field is only populated in Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

CI relationships

The AWS hardware type patterns create these relationships to support AWS hardware type discovery.

CI	Relationship	CI
Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Host pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Dedicated Hosts in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Host (LP) pattern.

Cloud Host [cmdb_ci_cloud_host]

Field	Description
Name [name]	Name of the host.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
CPU Core Count [cpu_core_count]	The number of host cores.
State [state]	The current state of the host.
Host Type [host_type]	Host type (instance family). For example: m5, c5.
Cloud Vendor [cloud_vendor]	The cloud vendor: AWS.
Is Virtual [virtual]	Indicates whether this is a virtual host. The value is set to **false**.

CI relationships

The Amazon AWS - Host (LP) pattern creates the following relationships and references to support AWS Host discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Host [cmdb_ci_cloud_host]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Virtual Machine Instance [cmdb_ci_vm_instance]	Runs on::Runs	Cloud Host [cmdb_ci_cloud_host]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Host [cmdb_ci_cloud_host]

AWS Tag discovery

The Amazon AWS - Host (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Host [cmdb_ci_cloud_host] table.

AWS IAM Policy pattern-based discovery

Discovery and Service Mapping Patterns finds AWS IAM Policies on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - IAM Policy - Extended Inventory (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS IAM Policy [cmdb_aws_iam_policy]

Field	Description
Name [name]	Name of the IAM policy.
Object ID [object_id]	The Amazon Resource Name (ARN) of the IAM policy.
Policy ID [policy_id]	The stable and unique string identifying the policy.
Is Customer Managed [is_customer_managed]	Indicates whether the policy is customer-managed. • Policies managed by customers: the value is true • Policies managed by AWS: the value is false
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - IAM Policy - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the IAM policy.
Object ID [object_id]	The ARN of the IAM policy.
Resource type [resource_type]	Type of resource. The value is set to **AWS::IAM::Policy**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - IAM Policy - Extended Inventory (LP) pattern creates the following relationships and references to support AWS IAM Policy discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS IAM Policy [cmdb_aws_iam_policy]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - IAM Policy - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS IAM Role pattern-based discovery

Discovery and Service Mapping Patterns finds AWS IAM Roles on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - IAM Role - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS IAM Role [cmdb_aws_iam_role]

Field	Description
Name [name]	Name of the IAM role.
Object ID [object_id]	The Amazon Resource Name (ARN) of the IAM role.
Role ID [role_id]	The stable and unique string identifying the role.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - IAM Role - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the IAM role.
Object ID [object_id]	The ARN of the IAM role.
Resource type [resource_type]	Type of resource. The value is set to **AWS::IAM::Role**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - IAM Role - Extended Inventory (LP) pattern creates the following relationships and references to support AWS IAM Role discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS IAM Role [cmdb_aws_iam_role]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - IAM Role - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS IAM User pattern-based discovery

Discovery and Service Mapping Patterns finds AWS IAM Users on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - IAM User - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS IAM User [cmdb_aws_iam_user]

Field	Description
Name [name]	Name of the IAM user.
Object ID [object_id]	The Amazon Resource Name (ARN) of the IAM user.
User ID [user_id]	The stable and unique string identifying the user.
Path [path]	The path used to organize the user within the IAM hierarchy. If no custom path is set, the value is `/`.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - IAM User - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the IAM user.
Object ID [object_id]	The ARN of the IAM user.
Resource type [resource_type]	Type of resource. The value is set to **AWS::IAM::User**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - IAM User - Extended Inventory (LP) pattern creates the following relationships and references to support AWS IAM User discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

CI references

CI	Field	Referenced CI
AWS IAM User [cmdb_aws_iam_user]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - IAM User - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Internet Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Internet Gateways in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Internet Gateway (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Internet Gateway [cmdb_ci_internet_gateway]

Field	Description
Name [name]	Name or ID if no Name is specified for the internet gateway.
Object ID [object_id]	ID of the internet gateway.

Internet Gateway Endpoint [cmdb_ci_endpoint_intgateway]

Field	Description
Name [name]	Name or ID if no Name is specified for the internet gateway.
Object ID [object_id]	ID of the internet gateway.

CI relationships

The Amazon AWS - Internet Gateway (LP) pattern creates the following relationships and references to support AWS Internet Gateway discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Internet Gateway [cmdb_ci_internet_gateway]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Internet Gateway [cmdb_ci_internet_gateway]	Implement End Point To::Implement End Point From	Internet Gateway Endpoint [cmdb_ci_endpoint_intgateway]
Cloud Network [cmdb_ci_network]	Use End Point To::Use End Point From	Internet Gateway Endpoint [cmdb_ci_endpoint_intgateway]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Internet Gateway [cmdb_ci_internet_gateway]

AWS Tag discovery

The Amazon AWS - Internet Gateway (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Internet Gateway [cmdb_ci_internet_gateway] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS IP Address pattern-based discovery

Discovery and Service Mapping Patterns finds AWS IP addresses associated with network interfaces in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - IP Address (LP) pattern.

IP Address [cmdb_ci_cloud_ip_address]

Field	Description
Name [name]	Name or ID if no Name is specified for the network interface.
IP Address [ip_address]	If available, the address of the Elastic IP address bound to the network interface. If not available, the private IP.
Object ID [object_id]	The ID of the network interface.
Public DNS [public_dns]	The public DNS name, if available.
Private IP Address [private_ip]	The IPv4 address of the network interface within the subnet.
Instance ID [instance_id]	The ID of the instance.

CI relationships

The Amazon AWS - IP Address (LP) pattern creates the following relationships to support AWS IP Address discovery.

CI relationships

CI	Relationship	CI
Cloud Key Pair [cmdb_ci_cloud_key_pair]	Contains::Contained by	IP Address [cmdb_ci_cloud_ip_address]

AWS Keyspaces discovery with Patterns

The ServiceNow® Discovery application uses the Amazon Keyspaces discovery pattern to find Amazon Keyspaces managed Apache Cassandra–compatible database service.

Amazon Keyspaces is a managed Apache Cassandra–compatible database service. A keyspace groups related tables that are relevant for one or more applications. It contains one or more tables and defines the replication strategy for all the tables it contains.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

- Install the Discovery and Service Mapping Patterns application from the ServiceNow Store.
- Configure an AWS Service Account with valid credentials on the ServiceNow instance with permission to run the following APIs:
  - `https://cassandra.{region}.amazonaws.com {Content-Type:application/x-amz-json-1.0, , X-Amz-Target: KeyspacesService.ListKeyspaces}`
  - `https://tagging.{region}.amazonaws.com {"ResourceTypeFilters": []}`

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

Keyspaces are stored in Cloud Databases [cmdb_ci_cloud_database] table.

Field	Description
Amazon resource name	Name of the keyspace.
FQDN	Fully qualified domain name of the region where the keyspace is located.
Port	Port used for incoming TLS connections (9142).

AWS Key Pair pattern-based discovery

Discovery and Service Mapping Patterns finds AWS key pairs in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Key Pair (LP) pattern.

Cloud Key Pair [cmdb_ci_cloud_key_pair]

Field	Description
Name [name]	Name of the key pair.
Object ID [object_id]	Name of the key pair.
Finger Print [finger_print]	The fingerprint of the key pair. If you used CreateKeyPair to create the key pair, this value is the SHA-1 digest of the DER encoded private key. If you used ImportKeyPair to provide AWS the public key, this value is the MD5 public key fingerprint as specified in section 4 of RFC 4716.

CI relationships

The Amazon AWS - Key Pair (LP) pattern creates the following relationships and references to support AWS Key Pair discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Servers [cmdb_ci_server]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Key Pair [cmdb_ci_cloud_key_pair]

AWS Tag discovery

The Amazon AWS - Key Pair (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Key Pair [cmdb_ci_cloud_key_pair] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Kinesis discovery with Patterns

The ServiceNow® Discovery application uses the Amazon Kinesis discovery pattern to find Kinesis services available on Amazon Web Services (AWS).

Kinesis is a managed, scalable, AWS cloud-based service that allows real-time processing of streaming large amount of data per second. AWS Kinesis contains four types of services:

- Kinesis Video Streams - capture, process, and store video streams for analytics and machine learning.
- Kinesis Data Streams - a fully-managed streaming data service.
- Kinesis Data Firehose - load data streams into AWS data stores.
- Kinesis Data Analytics- process and analyze streaming data.

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Install the Discovery and Service Mapping Patterns application from the ServiceNow Store.
- Verify the following plugin versions:
  - Discovery and Service Mapping Patterns 1.0.93 or later (sn_itom_pattern)
  - CMDB CI Class Models version 1.40.0 or later (sn_cmdb_ci_class)

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

The discovered data includes the following fields in the Cloud Database [cmdb_ci_aws_kinesis] table.

Field	Description
Object ID	Stream resource name.
Name	The name of the stream. • Kinesis Data - cmdb_ci_aws_kinesis_data • Kinesis Video - cmdb_ci_aws_kinesis_video

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• Kinesis Firehose - cmdb_ci_aws_kinesis_firehose • Kinesis Data Analytics - cmdb_ci_aws_kinesis_data_analytics
Status	The status of the stream.

Configuration item (CI) relationships

CI	Relationship	CI
AWS Kinesis Data	Hosted on::Hosts	AWS Datacenter
AWS Kinesis Video	Hosted on::Hosts	AWS Datacenter
AWS Kinesis Firehose	Hosted on::Hosts	AWS Datacenter
AWS Kinesis Data Analytics	Hosted on::Hosts	AWS Datacenter

AWS KMS Key pattern-based discovery

Discovery and Service Mapping Patterns finds AWS KMS Keys on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - KMS Key - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS KMS Key [cmdb_aws_kms_key]

Field	Description
Object ID [object_id]	The Amazon Resource Name (ARN) of the KMS Key.
Name [name]	The ARN of the KMS Key.
Arn [arn]	The ARN of the KMS Key.
Key ID [key_id]	The unique identifier of the KMS Key.
Key State [key_state]	The current state of the KMS Key. For example: Enabled, Disabled, or PendingDeletion.
Enabled [enabled]	Indicates whether the KMS Key is enabled.
Description [description]	Description of the KMS Key.
Key Algorithm [key_algorithm]	The cryptographic algorithm of the key. For example: SYMMETRIC_DEFAULT, RSA_2048, or ECC_NIST_P256.
Key Usage [key_usage]	The cryptographic operations for which the key can be used. For example: ENCRYPT_DECRYPT, SIGN_VERIFY, or GENERATE_VERIFY_MAC.
Origin [origin]	The source of the key material. For example: AWS_KMS or EXTERNAL, or AWS_CLOUDHSM.
Multi Region [multi_region]	Indicates whether the key is a multi-Region key.
Multi Region Key Type [multi_region_key_type]	The type of multi-Region key. For example: PRIMARY or REPLICA.
Primary Key [primary_key]	Details of the primary key for a replica key, including the ARN and Region.
Replica Keys [replica_keys]	Details of the replica keys for a primary key, including the ARNs and Regions.
Encryption Algorithms [encryption_algorithms]	The encryption algorithms supported by the key.
Signing Algorithms [signing_algorithms]	The signing algorithms supported by the key.
MAC Algorithms [mac_algorithms]	The MAC algorithms supported by the key.
Created Date [created_date]	Date and time when the KMS Key was created.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - KMS Key - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	The ARN of the KMS Key.
Object ID [object_id]	The ARN of the KMS Key.
Resource type [resource_type]	Type of resource. The value is set to **AWS::KMS::Key**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - KMS Key - Extended Inventory (LP) pattern creates the following relationships and references to support AWS KMS Key discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS KMS Key [cmdb_aws_kms_key]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - KMS Key - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Lambda discovery with Patterns

The ServiceNow Discovery and Service Mapping applications can find and map Lambda functions that run in your AWS cloud. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery and Service Mapping use the Amazon AWS Lambda pattern to run horizontal and top-down discovery. When you run cloud discovery, your instance queries your AWS service account using the necessary credentials, and then runs a query to obtain Lambda information. The pattern parses this information and then it is saved in the CMDB.

You can use this pattern on the ServiceNow platform using Jakarta Patch 10, Kingston Patch 8, or London Patch 1.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Update set

If your instance is on the Jakarta or Istanbul versions, import the update set provided in KB0679927: AWS Lambda top-down discovery update set. This update set is required for Service Mapping.

User policy for AWS user

If you use Identity and Access Management (IAM) to manage users in Amazon Web Services (AWS) platform, ensure that you created a user policy for the AWS user. For more information, refer to Control AWS access and permissions using policies in the ServiceNow Cloud Provisioning and Governance documentation.

AWS service account with credentials

You must set up your service account with the appropriate AWS credentials. You run discovery from the service account form.

Refer to Service accounts for Discovery and Service Mapping in product documentation.

Service account

On the ServiceNow platform, create a service account. Set the Account ID to the Amazon account IT to which AWS Lambda belongs. Use the Account ID as it appears in the AWS Management Console. Ensure to place the MID Server on the virtual machine that is part of the service account that you are planning to discover.

Discover logical datacenters using horizontal discovery

Click **Discover datacenters** under **Related Links** on the Cloud Service Account form.

User permissions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide a user with permissions to run the ListFunctions API with the Resource: arn:aws:lambda:region:account-id: function:function-name. Provide a user with permissions to run the GetFunction API with the Resource: '*'.

Amazon Resource Name (ARN)

(Optional for application services using ARN as an entry point) Ensure that you know the ARN for the AWS Lambda you need to discover. ARNs serve as the resource identifiers for Amazon Web Services.

Host

If your instance is on the Jakarta or Istanbul versions, use the following host: 127.0.0.1.

Cloud Discovery

For Cloud Discovery, download the Discovery and Service Mapping Patterns from the ServiceNow Store.

MID Server

When installing the MID Server, ensure that the host machine meets or exceeds the MID Server system requirements published on the ServiceNow documentation site.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

Horizontal discovery populates this data in the CMDB:

Table and field	Description
Cloud Function [cmdb_ci_cloud_function]	
Name [name]	The name of the function.
Object ID [object_id]	The ID (Amazon Resource Name) of the AWS Lamdba function, for example: `arn:aws:lambda:{datacenter}1:751200741520:function:{function-name}`
Version [version]	The version of the function. For example, `$Latest` represents the most recent published version.
Language [language]	The runtime environment, or language and version, that the function is written in, for example: `Node.js6.10` for JavaScript version 6.10.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Code Location URL [code_location_url]	The URL of the function that you built from the API Gateway console.
CodeSha256 [codesha256]	The cryptographic hash signature for the function.
Code Size [code_size]	The size of the function code in MB.
Function Last Modified [function_last_modified]	The date the function last changed in your AWS Lamdba environment.

Note: To see all of these fields, you need to configure the form layout as described in product documentation.

CI relationships

These relationships are created to support AWS Lambda discovery:

CI	Relationship	CI
Cloud Function [cmdb_ci_cloud_function]	Uses:Used by	Cloud DataBase [cmdb_ci_cloud_database]
	Hosted on:Hosts	AWS Datacenter [cmdb_ci_aws_datacenter] The AWS datacenter has relationships to other AWS CIs, such as availability zones, service accounts, and so on. Refer to AWS and Azure cloud discovery in product documentation.
	Provides:Is provided by	Cloud Event Source [cmdb_ci_cloud_event_source] This CI could be any event source that triggers the cloud function or the API Gateway.

Connections discovered by Service Mapping during the top-down discovery

Service Mapping performs the top-down discovery of the AWS Lambda in the context of application services. It discovers outgoing connections of the AWS Lambda:

- AWS API Gateway
- Amazon Relational Database Service (RDS)
- Other instance of AWS Lambda inside or outside the current region using ARN

Troubleshooting

If the mapping process does not proceed as you expected, follow the following suggestions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again. **Note:** If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

Related topics

Test patterns for discovering AWS API Gateway and AWS Lambda

Define a Cloud REST Query for cloud computing devices

AWS LB Pool pattern-based discovery

Discovery and Service Mapping Patterns finds AWS load balancer (LB) pools (known in AWS as target groups) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - LB Pool (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Name [name]	Name of the load balancer pool.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
Comments [comments]	Identifier for internal usage (deletion strategy).

CI relationships

The Amazon AWS - LB Pool (LP) pattern creates the following relationships and references to support AWS LB Pool discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Load Balancer Pool [cmdb_ci_lb_pool]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

AWS LB Pool Member pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Application Load Balancer (LB) pool members (known in AWS as target group members) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Enable AWS Application Load Balancer target discovery

By default, the Amazon AWS - LB Pool Member(LP) pattern doesn't execute discovery. To enable the discovery of AWS Application Load Balancer targets, set the *sn_itom_pattern.discover_aws_app_pool_members* MID Server property to **true**. For more information, see Enable AWS Application Load Balancer target discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - LB Pool Member(LP) pattern.

Load Balancer Pool Member [cmdb_ci_lb_pool_member]

Field	Description
Name [name]	Target ID, depending on the target type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Pool Member [cmdb_ci_lb_pool_member] (continued)

Field	Description
	For example: Instance ID, IP address, Lambda ARN, or Application Load Balancer ARN.
Service port [service_port]	The port on which the target is listening, if available.
Object ID [object_id]	Possible values are: • Target ID • Target ID and target port, if available, in the following format: `<target ID>#<target port>`. For example: `i-0123456789abcdef0#8080`
Comments [comments]	Comments related to the CI.
Operational status [operational_status]	Operational status of the target. Possible values are Operational or Non-Operational.
Install Status [install_status]	Installation status of the target. Possible values are Installed or Retired.
Pool [pool]	References the Load Balancer Pool [cmdb_ci_lb_pool] table.

CI relationships

The Amazon AWS - LB Pool Member(LP) pattern creates the following relationships and references to support AWS LB Pool Member discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Load Balancer Pool [cmdb_ci_lb_pool]	Owns::Owned by	Load Balancer Pool Member [cmdb_ci_lb_pool_member]

CI references

CI	Field	Referenced CI
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	Pool [pool]	Load Balancer Pool [cmdb_ci_lb_pool]

Enable AWS Application Load Balancer target discovery

Enable the *sn_itom_pattern.discover_aws_app_pool_members* MID Server property to discover AWS Application Load Balancer targets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Mid Server** > **Properties**.
2. In the **Name** field, search for `sn_itom_pattern.discover_aws_app_pool_members`.
3. Select the *sn_itom_pattern.discover_aws_app_pool_members* property.
4. In the **Value** field, enter `true`.
5. Select **Update**.

Related topics

AWS LB Pool Member pattern-based discovery

AWS discovery using patterns

AWS LB Service pattern-based discovery

Discovery and Service Mapping Patterns finds AWS load balancer (LB) listeners in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - LB Service (LP) pattern.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Name [name]	Name of the load balancer service.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
Port [port]	The port on which the load balancer is listening.
Service port [service_port]	The port on which the load balancer is listening.
Server protocol [service_protocol]	The protocol to use for routing traffic to instances: HTTP, HTTPS, TCP, or SSL.
Comments [comments]	Identifier for internal usage (deletion strategy).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - LB Service (LP) pattern creates the following relationships to support AWS LB Service discovery.

CI relationships

CI	Relationship	CI
Load Balancer Service [cmdb_ci_lb_service]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

AWS Linux Server pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Linux Server configuration items (CIs) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Verify that AWS SSM is enabled

Verify that AWS Systems Manager (AWS SSM) is enabled on the Amazon Elastic Compute Cloud (Amazon EC2) instances.

Enable Linux Server CI creation

Set the *sn_itom_pattern.aws_cloud_discovery_populate_server_ci* system property to **true**. For more information, see Configure Server CI creation during AWS cloud discovery.

Verify SSM Agent execution context

The SSM Agent must run as root to retrieve the serial number using `dmidecode`. If your security policy restricts the SSM Agent to a non-root user, the pattern can't create a Server CI for that instance.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the AWS - Linux Server (LP) pattern.

Linux Server [cmdb_ci_linux_server]

Field	Description
Name [name]	OS-level hostname of the EC2 instance. If the *glide.discovery.hostname.include_domain* property is set to false, the domain suffix is removed and only the short hostname is stored.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux Server [cmdb_ci_linux_server] (continued)

Field	Description
Object ID [object_id]	The EC2 instance ID.
Correlation ID [correlation_id]	The EC2 instance ID.
Serial number [serial_number]	BIOS serial number of the EC2 instance.
IP Address [ip_address]	IP address of the EC2 instance.
CPU core count [cpu_core_count]	Number of vCPUs available to the EC2 instance.
CPU count [cpu_count]	CPU count. The value is set to **1**.
RAM (MB) [ram]	Amount of memory available to the EC2 instance, in megabytes (MB).
Operating System [os]	Operating system of the server. The value is set to **Linux**.
Is Virtual [virtual]	Indicates that the server is a virtual machine. The value is set to **true**.
Install Status [install_status]	Install status of the resource. The value is set to **Installed**.
Operational status [operational_status]	Operational status of the resource. The value is set to **Operational**.

Serial Number [cmdb_serial_number]

Field	Description
Serial number [serial_number]	BIOS serial number of the EC2 instance.
Serial number type [serial_number_type]	Type of serial number. The value is set to **bios**.
Valid [valid]	Indicates that the serial number is active and valid. The value is set to **true**.
Configuration item [cmdb_ci]	References the Linux Server [cmdb_ci_linux_server] table.

CI relationships and references

The AWS - Linux Server (LP) pattern creates the following relationships and references to support AWS Linux Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Linux Server [cmdb_ci_linux_server]	Virtualized by::Virtualizes	Virtual Machine Instance [cmdb_ci_vm_instance]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration item [cmdb_ci]	Linux Server [cmdb_ci_linux_server]

AWS Marketplace pattern-based discovery

Discovery and Service Mapping Patterns finds active AWS Marketplace subscriptions on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The Amazon AWS - Marketplace (LP) pattern discovers the following AWS Marketplace products:

- Amazon Machine Images (AMI)
- SaaS

Amazon AWS - Marketplace (LP) pattern data model

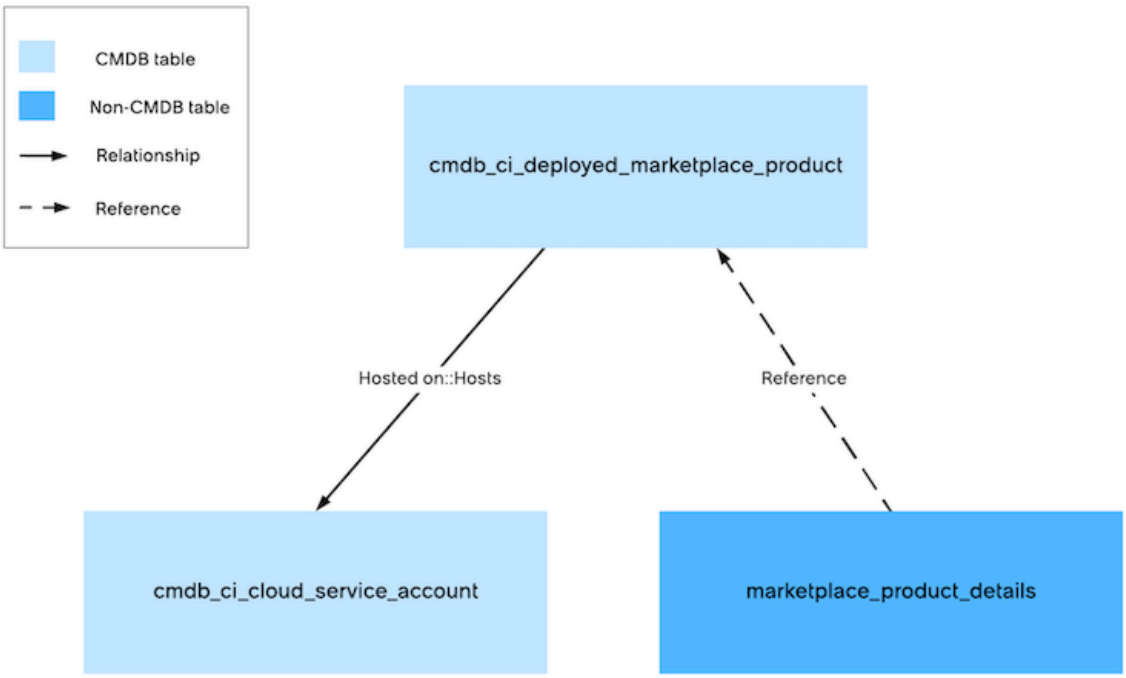

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Verify that the us-east-1 region is activated on your instance

The pattern connects to the AWS Marketplace APIs through us-east-1, regardless of your account's configured region.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Marketplace (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]

Field	Description
Name [name]	Name of the marketplace product.
Object ID [object_id]	Unique identifier of the marketplace product in the format **{productId}/Marketplace**.
Resource Type [resource_type]	Type of the marketplace product. For example: AmiProduct or SaaSProduct.
Plan Name [plan_name]	Pricing plan for the marketplace subscription. Discovery maps AWS Marketplace pricing terms to the **Plan Name** values. For example: • UsageBasedPricingTerm: Pay As You Go • ByolPricingTerm: BYOL • RecurringPaymentTerm: Recurring
Short Description [short_description]	Brief description of the marketplace product.
Organization Id [organization_id]	Unique identifier of the AWS organization associated with the subscription.
Market [market]	Geographic market where the product is sold. Discovery maps AWS Marketplace currency codes to the **Market** values. For example: • USD: US • EUR: Europe • JPY: Japan
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Marketplace Product Details [marketplace_product_details]

Field	Description
Product Code [product_code]	Unique product identifier from AWS Marketplace.
Publisher Name [publisher_name]	Name of the organization or individual that publishes the marketplace product.
Version [version]	Version of the marketplace product. The value is set to **None**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Marketplace dependency view

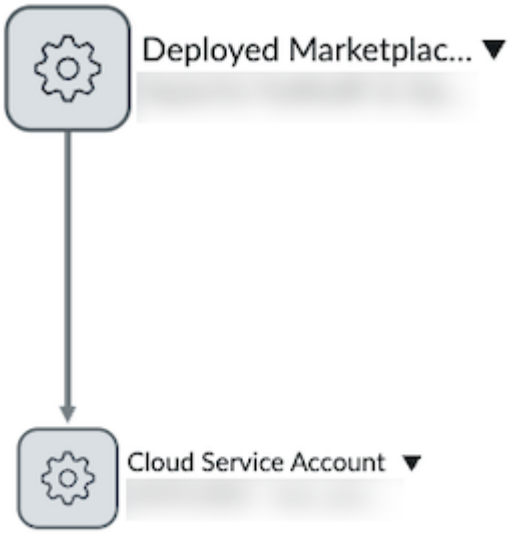

CI relationships and references

The Amazon AWS - Marketplace (LP) pattern creates the following relationships and references to support AWS Marketplace discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

CI references

CI	Field	Referenced CI
Marketplace Product Details [marketplace_product_details]	Deployed On [deployed_on]	Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]

AWS MemoryDB for Redis discovery with Patterns

The ® Discovery application uses the Amazon AWS MemoryDB pattern and extensions to find AWS MemoryDB for Redis. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Configured AWS Credentials

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

On your instance, configure credentials of type **AWS Credentials**. For more information, see Access setup for AWS service accounts.

Configured Cloud service account

Configure the AWS service account valid in the ServiceNow instance. For more information, see Set up AWS service accounts.

Configured user permissions to execute the list AWS MemoryDB for Redis API calls

Provide the user permission to run the following API:

- https://memory-db.{region}.amazonaws.com?Action=DescribeClusters&ShowShardDetails=true
- https://memory-db.{region}.amazonaws.com? Action=ListTags&ResourceArn={arn}

For more information, see Amazon AWS Cloud components discovery using patterns.

Configured Cloud Discovery schedule

Create a cloud application schedule for discovering AWS datacenters. Set **Discovery** to **Cloud discovery**. For more information, see: Create schedules for discovering cloud resources in Cloud Discovery Workspace

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected during AWS MemoryDB for Redis horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS MemoryDB pattern.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	The Amazon Resource Name of a single node in the cluster.
Name [name]	The name of a single node in the cluster.
Fully qualified domain name [fqdn]	The Fully Qualified Domain Name of each node in the cluster.
TCP port(s) [tcp_port]	The TCP port that a single node in the cluster is getting traffic from.
Type [type]	The type of database. Value is set to **memorydb**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Cluster ID [cluster_id]	The Amazon Resource Name of the cluster.
Name [name]	The name of the cluster.
Fully qualified domain name [fqdn]	The Fully Qualified Domain Name of the cluster.
TCP port(s) [tcp_port]	The TCP port that the cluster is getting traffic from.

CI relationships

The AWS MemoryDB for Redis pattern creates the following relationships and references to support AWS MemoryDB for Redis discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Cluster of::Cluster	Cloud DataBase [cmdb_ci_cloud_database]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

AWS tag discovery

The AWS MemoryDB for Redis pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] table.

Related topics

Detailed information on products discovered by ITOM Visibility

AWS discovery using patterns

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS NAT Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Network Address Translation (NAT) Gateways in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - NAT Gateway (LP) pattern.

NAT Gateway [cmdb_ci_nat_gateway]

Field	Description
Name [name]	Name of the NAT gateway.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.

NAT Endpoint [cmdb_ci_endpoint_nat]

Field	Description
Name [name]	Name of the NAT endpoint.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.

CI relationships and references

The Amazon AWS - NAT Gateway (LP) pattern creates the following relationships and references to support AWS NAT Gateway discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
NAT Gateway [cmdb_ci_nat_gateway]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
NAT Gateway [cmdb_ci_nat_gateway]	Implement End Point To::Implement End Point From	NAT Endpoint [cmdb_ci_endpoint_nat]
Cloud Network [cmdb_ci_network]	Use End Point To::Use End Point From	NAT Endpoint [cmdb_ci_endpoint_nat]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	NAT Gateway [cmdb_ci_nat_gateway]

AWS Tag discovery

The Amazon AWS - NAT Gateway (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the NAT Gateway [cmdb_ci_nat_gateway] table.

AWS Network pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Virtual Private Cloud (VPC) networks in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Network (LP) pattern.

Cloud Network [cmdb_ci_network]

Field	Description
Name [name]	Name of the VPC network.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
State [state]	The current state of the VPC. Possible values are: pending or available.
CIDR [cidr]	CIDR representation of the subnet. For example, 10.0.0.0/24.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Network [cmdb_ci_network] (continued)

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships and references

The Amazon AWS - Network (LP) pattern creates the following relationships and references to support AWS Network discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Network [cmdb_ci_network]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Network [cmdb_ci_network]

AWS Tag discovery

The Amazon AWS - Network (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Network [cmdb_ci_network] table.

AWS Network ACL pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Network access control lists (ACLs) on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - ACL (LP) pattern.

Network ACL [cmdb_ci_network_acl]

Field	Description
Name [name]	Name of the network ACL.
Object ID [object_id]	Unique identifier, allocated by AWS for this resource.

ACL Endpoint [cmdb_ci_endpoint_acl]

Field	Description
Name [name]	Name of the endpoint.
Object ID [object_id]	Unique identifier, allocated by AWS for this resource.

CI relationships

The Amazon AWS - ACL (LP) pattern creates the following relationships and references to support AWS Network ACL discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Network ACL [cmdb_ci_network_acl]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Cloud Network [cmdb_ci_network]	Contains::Contained by	Network ACL [cmdb_ci_network_acl]
Network ACL [cmdb_ci_network_acl]	Implement End Point To::Implement End Point From	ACL Endpoint [cmdb_ci_endpoint_acl]
Cloud Subnet [cmdb_ci_cloud_subnet]	Use End Point To::Use End Point From	ACL Endpoint [cmdb_ci_endpoint_acl]

AWS Network Firewall pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Network Firewall - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The Amazon Resource Name (ARN) uniquely identifying the network firewall resource.
Name [name]	The display name of the network firewall resource.
Resource type [resource_type]	Type of resource. The value is set to **AWS::NetworkFirewall::Firewall**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Network Firewall - Extended Inventory (LP) pattern creates these relationships to support AWS Network Firewall discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

AWS NIC pattern-based discovery

Discovery and Service Mapping Patterns finds AWS network interfaces (NICs) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - NIC (LP) pattern.

Cloud Mgmt Network Interface [cmdb_ci_nic]

Field	Description
Name [name]	The Name or ID if no Name is specified for the network interface.
Object ID [object_id]	The ID of the network interface.
State [state]	The status of the network interface. Possible valid values: available, associated, attaching, in-use, or detaching.
Private IP [private_ip]	The IPv4 address of the network interface within the subnet.
IP Address [ip_address]	If available, the address of the Elastic IP address bound to the network interface. If not available, the Private IP.
Public IP [public_ip]	The address of the Elastic IP address bound to the network interface.

Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]

Field	Description
Name [name]	IP address of the load balancer.
Object ID [object_id]	If available, the address of the Elastic IP address bound to the network. If not available, the Private IP.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress] (continued)

Field	Description
IP Address [ip_address]	If available, the address of the Elastic IP address bound to the network interface. If not available, the Private IP.
Comments [comments]	Identifier for internal usage (deletion strategy).

VNIC Endpoint [cmdb_ci_endpoint_vnic]

Field	Description
Name [name]	The name of the virtual network machine interface (VNIC) endpoint.
Object ID [object_id]	Unique identifier, allocated by Amazon for this resource.
IP Address [ip_address]	If available, the address of the Elastic IP address bound to the network interface. If not available, the Private IP.
Host [host]	The ID of the instance.

CI relationships and references

The Amazon AWS - NIC (LP) pattern creates the following relationships and references to support Amazon AWS NIC discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Owns::Owned by	Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	VNIC Endpoint [cmdb_ci_endpoint_vnic]
Cloud Subnet [cmdb_ci_cloud_subnet]	Contains::Contained by	Cloud Mgmt Network Interface [cmdb_ci_nic]
VNIC Endpoint [cmdb_ci_endpoint_vnic]	Implement End Point To::Implement End Point From	Cloud Mgmt Network Interface [cmdb_ci_nic]
Cloud Mgmt Network Interface [cmdb_ci_nic]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Mgmt Network Interface [cmdb_ci_nic]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Tag discovery

The Amazon AWS - NIC (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Mgmt Network Interface [cmdb_ci_nic] table.

AWS OpenSearch discovery with Patterns

The ServiceNow® Discovery application uses the Amazon OpenSearch discovery pattern to find Amazon OpenSearch service domains.

OpenSearch is an open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and clickstream analysis. OpenSearch operates as a cluster of nodes that are abstracted from the end user. An OpenSearch service domain is synonymous with an OpenSearch cluster. Domains are clusters with settings, instance types, instance counts, and storage resources.

Prerequisites

- Install the Discovery and Service Mapping Patterns application from the ServiceNow Store.
- Configure an AWS Service Account with valid credentials on the ServiceNow instance with permission to run the following APIs:
  - `https://es.{region}.amazonaws.com/2021-01-01/domain`
  - `https://es.{region}.amazonaws.com/2021-01-01/opernsearch/domain-info{"DomainNames":[]}`
  - `https://tagging.{region}.amazonaws.com {"ResourceTypeFilters":[]}`

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected by Discovery during horizontal discovery

OpenSearch service domains are stored in Cloud Databases [cmdb_ci_cloud_database] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Amazon resource name	Name of the service domain.
FQDN	Fully qualified domain name of the region where the keyspace is located.

AWS Organizations pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Organizations accounts on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS Organizations pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Unique identifier (ID) of the account.
Name [name]	Name assigned to the account.
Datacenter Type [datacenter_type]	Type of datacenter. The value is set to **AWS Datacenter [cmdb_ci_aws_datacenter]**.
Object ID [object_id]	Unique ID of the account.
Datacenter URL [datacenter_url]	URL of the AWS datacenter. This field is populated for AWS GovCloud accounts.
Is management account [is_master_account]	Indicates whether this account is the management (primary) account of the organization. • **true**: Management account • **false**: Member account
Parent account [parent_account]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table (management account).

CI relationships

The Amazon AWS Organizations pattern creates these relationships to support AWS Organizations account discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Service Account [cmdb_ci_cloud_service_account] (management account)	Members::Member of	Cloud Service Account [cmdb_ci_cloud_service_account] (member account)
Cloud Service Account [cmdb_ci_cloud_service_account] (member account)	References	Cloud Service Account [cmdb_ci_cloud_service_account] (management account)

AWS Organizational Units pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Organizational Units in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Organizational Units (LP) pattern.

AWS Organizational Unit [cmdb_ci_aws_org_unit]

Field	Description
Name [name]	The user-friendly name of the Organizational Unit (OU).
Object ID [object_id]	The unique identifier (ID) associated with this OU. The ID is unique to the organization.
Organizational ID [aws_org_id]	The unique ID associated with this OU. The ID is unique to the organization.
Org Unit Parent ID [org_unit_parent_id]	The ID of the root or the immediate parent OU.
Root ID [root_id]	The unique ID of the root of the organization.

CI relationships and references

The Amazon AWS - Organizational Units (LP) pattern creates the following relationships and references to support AWS Organizational Units discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Organizations [cmdb_ci_cloud_org]	Contains::Contained by	AWS Organizational Unit [cmdb_ci_aws_org_unit]
AWS Organizational Unit [cmdb_ci_aws_org_unit]	Contains::Contained by	Cloud Service Account [cmdb_ci_cloud_service_account]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	AWS Organizational Unit [cmdb_ci_aws_org_unit]

AWS Tag discovery

The Amazon AWS - Organizational Units (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the AWS Organizational Unit [cmdb_ci_aws_org_unit] table.

AWS OS image pattern-based discovery

Discovery and Service Mapping Patterns finds AWS OS images (both owned and executable) on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses the Amazon AWS - Owned Template (LP) and Amazon AWS - Executable Template (LP) patterns to discover AWS OS images from your AWS account. Owned images are custom images created by your AWS account, while executable images include public and shared images your account has permission to use.

When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. When enabled, Discovery launches the Amazon AWS - Owned Cloud OS Image (LP) and Amazon AWS - Executable Cloud OS Image (LP) patterns, which populate the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the AWS OS image patterns.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Name [name]	Name of the image. For example: **CentOS 7.0-64 Minimal for VSI**.
Object ID [object_id]	Unique identifier, allocated by Amazon AWS Cloud for this resource.
Image Type [image_type]	Type of the OS image.
Root Device Type [root_device_type]	The device name of the root device volume. For example: **/dev/sda1**.
Image Source [image_source]	Source of the OS image.
Guest OS [guest_os]	Operating system that is secondary to the OS.
Provider [provider]	The cloud provider, which is **AWS**. This field is only populated in the Cloud Image [cmdb_ci_cloud_os_image] table.

CI relationships

The AWS OS image patterns create these relationships to support AWS OS image discovery.

CI	Relationship	CI
Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Key value [cmdb_key_value]	References	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

AWS Tag discovery

Starting with Discovery and Service Mapping Patterns version 1.30.2, the pattern extension section discovers tags only. License discovery has moved to the Amazon AWS - Virtual Server (LP) pattern.

Key value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Public IP Address pattern-based discovery

Discovery and Service Mapping Patterns finds AWS public IP addresses (Elastic IPs) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Public IP Address (LP) pattern.

Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]

Field	Description
Name [name]	Name or allocation ID if no Name is specified for the public IP address.
Object ID [object_id]	The ID representing the allocation of the address for use with EC2-VPC.
Public IP Address [public_ip]	The Elastic IP address.

CI relationships and references

The Amazon AWS - Public IP Address (LP) pattern creates the following relationships and references to support AWS Public IP Address discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]

AWS Tag discovery

The Amazon AWS - Public IP Address (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress] table.

AWS Resource Inventory discovery with Patterns

The ServiceNow Discovery and Service Mapping applications can find and map the AWS resources available by AWS Config Service. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses the AWS Resource Inventory pattern to discover the resources available through AWS. Use this pattern for AWS resources which do not have a dedicated pattern, such as AWS Lambda or EC2. This pattern is limited to resources that are available from AWS Config Service, as described in the AWS Config Developer Guide.

You can use this pattern on the ServiceNow platform using London (Patch 8), Madrid (Patch 2), or later releases.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

- Create AWS credentials.
- Create an AWS cloud service account.

  You can use only the AWS management service account for discovery, not its sub-accounts.

  Note: The Cloud Discovery user interface refers to member accounts as sub-accounts.

- Discover AWS Datacenters (LDC)
- Schedule a full AWS discovery or use a specific inventory pattern
- Fine-tune AWS resource discovery using the Resource Inclusion List.

  If your deployment has custom patterns for AWS discovery, ensure that you do not discover AWS resources twice.

1. Ensure that the application scope is Discovery and Service Mapping Patterns:
   a. Navigate to **Settings** > **Developer**.
   b. Select `Discovery and Service Mapping Patterns` from the **Application** list.
2. Navigate to **System Definitions** > **Tables**.
3. Open the Cloud Inventory Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Under **Related Links**, click **Show List**.
5. Select resource types for which you have custom patterns, and select `Delete` from the **Actions on selected rows** list.

The Cloud Inventory Resource Inclusion List is predefined with common services. You can expand the list with additional resource types that you want the pattern to discover. The names of these resource types must conform to the appropriate vendor naming conventions.

Note: When you modify the out-of-the-box inclusion list, it is no longer updated automatically in application updates. You need to maintain the customized list on your own.

1. Open the Cloud Inventory Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table.
2. Click **New**.
3. Fill in the form, and then click **Submit**.

Field	Description
Cloud Vendor	The vendor of the resource type: AWS.
Resource Type	The AWS resource type value. For example, AWS::CloudWatch::Alarm.
Application	The application scope: Discovery and Service Mapping Patterns.

The changes are applied the next time you run the pattern.

- To discover the resources that support the AWS Config service, the pattern uses API: `https://config.[AWS region].amazonaws.com`. It specifies the resource type in the API request using the following format: `{"resourceType":"AWS::[RESOURCE]::[TYPE]"}`

  The POST method requires the following headers:

  - X-Amz-Target - StarlingDoveService.ListDiscoveredResources
  - Accept - application/json
  - Content-Type - application/x-amz-json-1.1

- To discover the resource tags for resources that support tags, the pattern uses API: `https://tagging.[AWS region].amazonaws.com`. It specifies the resource type in the API request using the following format: `{"ResourceTypeFilters":["[resource]:[type]"]}`

  The POST method requires the following headers:

  - X-Amz-Target - ResourceGroupsTaggingAPI_20170126.GetResources
  - Accept - application/json
  - Content-Type - application/x-amz-json-1.1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Table and field	Description
Main CI [cmdb_ci_cmp_resource]	
object_id	The ID of the item. This is typically the Amazon Resource Name (ARN).
name	Name of the resource.
resource_type	The asset resource type, based on the contents of the JSON file.
description	Short description of how the CI is populated.

CI relationships

The AWS Resource Inventory pattern creates the following CI relationship:

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on :: Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]

The following is a dependency view of the collected data flow.

AWS Resource Inventory dependency view

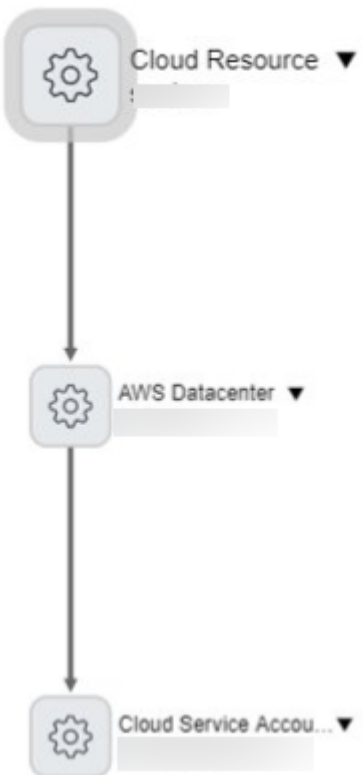

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Route Table pattern-based discovery

Discovery and Service Mapping Patterns finds AWS route tables in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Route Table (LP) pattern.

Route Table [cmdb_ci_route_table]

Field	Description
Name [name]	The Name tag of the route table, or the route table ID if no Name tag is specified.
Object ID [object_id]	The ID of the route table.
State [state]	If the route table is discoverable, the value is set to **available**.

Route Table Endpoint [cmdb_ci_endpoint_route_table]

Field	Description
Name [name]	The Name tag of the route table, or the route table ID if no Name tag is specified.
Object ID [object_id]	The ID of the route table.

CI relationships and references

The Amazon AWS - Route Table (LP) pattern creates the following relationships and references to support AWS Route Table discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Network [cmdb_ci_network]	Contains::Contained by	Route Table [cmdb_ci_route_table]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Cloud Subnet [cmdb_ci_cloud_subnet]	Use End Point To::Use End Point From	Route Table Endpoint [cmdb_ci_endpoint_route_table]
Route Table [cmdb_ci_route_table]	Implement End Point To::Implement End Point From	Route Table Endpoint [cmdb_ci_endpoint_route_table]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Route Table [cmdb_ci_route_table]

AWS Tag discovery

The Amazon AWS - Route Table (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Route Table [cmdb_ci_route_table] table.

AWS Secrets Manager Secret pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Secrets Manager Secrets on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Secrets Manager Secret - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Secrets Manager Secret [cmdb_aws_secrets_manager_secret]

Field	Description
Object ID [object_id]	The Amazon Resource Name (ARN) of the secret.
Name [name]	Name of the secret.
Primary Region [primary_region]	The primary Region for the secret, for multi-Region secrets.
Owning Service [owning_service]	The name of the service that created the secret.
Rotation Enabled [rotation_enabled]	Indicates whether automatic rotation is enabled for the secret.
Kms Key ID [kms_key_id]	The ARN or alias of the KMS key used to encrypt the secret.
Description [description]	Description of the secret.
Rotation Lambda Arn [rotation_lambda_arn]	The ARN of the Lambda function used to rotate the secret.
Rotation Days [rotation_days]	The number of days between automatic scheduled rotations of the secret.
Schedule Expression [schedule_expression]	A cron or rate expression that defines the schedule for rotating the secret.
Duration [duration]	The length of the rotation window during which Secrets Manager rotates the secret.
Secret Version To Stages [secret_version_to_stages]	A map of the secret versions and their staging labels.
Created Date [created_date]	Date and time when the secret was created.
Last Changed Date [last_changed_date]	Date and time when the secret was last changed.
Last Accessed Date [last_accessed_date]	Date and time when the secret was last accessed.
Last Rotated Date [last_rotated_date]	Date and time when the secret was last rotated.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Secrets Manager Secret - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the secret.
Object ID [object_id]	The ARN of the secret.
Resource type [resource_type]	Type of resource. The value is set to **AWS::SecretsManager::Secret**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Secrets Manager Secret - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Secrets Manager Secret discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Secrets Manager Secret [cmdb_aws_secrets_manager_secret]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Secrets Manager Secret - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Security Group pattern-based discovery

Discovery and Service Mapping Patterns finds AWS security groups in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Security Group (LP) pattern.

Compute Security Group [cmdb_ci_compute_security_group]

Field	Description
Name [name]	The name of the security group.
Object ID [object_id]	The ID of the security group.
State [state]	The current state of the security group. The value is set to **available**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships and references

The Amazon AWS - Security Group (LP) pattern creates the following relationships and references to support AWS Security Group discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Network [cmdb_ci_network]	Contains::Contained by	Compute Security Group [cmdb_ci_compute_security_group]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Compute Security Group [cmdb_ci_compute_security_group]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Compute Security Group [cmdb_ci_compute_security_group]

AWS Tag discovery

The Amazon AWS - Security Group (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Compute Security Group [cmdb_ci_compute_security_group] table.

AWS Serverless Database pattern-based discovery

Discovery and Service Mapping Patterns finds Amazon Aurora Serverless databases on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS Serverless Database pattern.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Cluster ID [cluster_id]	User-supplied identifier for the DB cluster.
Name [name]	Name of the database cluster. If a database cluster name isn't provided, the value is the same as the Cluster ID.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) for the DB cluster.
Cluster Version [cluster_version]	Database engine version number.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] (continued)

Field	Description
Vendor [vendor]	Database vendor. Value is set to **Amazon**.
Automated Backups [automated_backup]	Indicates whether automated backups are enabled and the retention period. Possible values are either **Disabled** or **Enabled (X Days)**, where X represents the backup retention period in days.
Deletion Protection Enabled [deletion_protection_enabled]	Indicates whether deletion protection is enabled for the DB cluster. Possible values are **true** or **false**. • **true**: Deletion protection is enabled and the database cannot be deleted • **false**: Deletion protection is disabled
IAM Authentication Enabled [iam_authentication_enabled]	Indicates whether AWS Identity and Access Management (IAM) database authentication is enabled. Possible values are **true** or **false**. • **true**: IAM authentication is enabled and you can use IAM credentials to authenticate to the database • **false**: IAM authentication is disabled
Storage Encrypted [storage_encrypted]	Indicates whether the DB cluster storage is encrypted.
Operational status [operational_status]	Operational status of the cluster. Default value is Operational.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	Database instance identifier or Aurora Serverless cluster identifier.
Name [name]	Name of the database instance or Aurora Serverless cluster.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) for the DB cluster.
Version [version]	Database engine version number.
Vendor [vendor]	Database vendor. Value is set to **Amazon**.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Object ID [object_id]	Availability zone name.
Name [name]	Availability zone name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IP Address [cmdb_ci_ip_address]

Field	Description
Netmask [netmask]	Netmask for the IP address. Value is set to 0.0.0.0.
Name [name]	IP address of the database cluster.

CI relationships

The Amazon AWS Serverless Database pattern creates these relationships to support AWS database discovery.

CI	Relationship	CI
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Cluster of::Cluster	Cloud DataBase [cmdb_ci_cloud_database]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud DataBase [cmdb_ci_cloud_database]
AWS Datacenter [cmdb_ci_aws_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Cloud DataBase [cmdb_ci_cloud_database]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Key Value [cmdb_key_value]	References	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]
Key Value [cmdb_key_value]	References	Cloud DataBase [cmdb_ci_cloud_database]

AWS Tag discovery

The Amazon AWS Serverless Database pattern collects tags and populates them in two entries in the Key Value [cmdb_key_value] table. One entry references the Cloud DataBase table, the other entry references the Cloud DataBase Cluster table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] table.

AWS SSM Cloud Agents pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Systems Manager (SSM) agents on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

SSM data model

The Amazon AWS - SSM Cloud Agents (LP) pattern introduces the following CI class that extends an existing CMDB class.

CI class introduced by this pattern

CI class	Extends from
Cloud System Management Agent [cmdb_ci_cloud_system_management_agent]	Virtual Machine Object [cmdb_ci_vm_object]

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - SSM Cloud Agents (LP) pattern.

Cloud System Management Agent [cmdb_ci_cloud_system_management_agent]

Field	Description
Cloud Agent Type [cloud_agent_type]	Type of cloud agent. The value is set to **AWS SSM**.
Install Status [install_status]	Install status of the SSM agent: • **Installed**: The agent is currently running. • **Absent**: The agent is not currently running.
IP Address [ip_address]	Address of the virtual machine (VM) instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud System Management Agent [cmdb_ci_cloud_system_management_agent] (continued)

Field	Description
Name [name]	Name of the VM instance that the SSM agent is running on.
Object ID [object_id]	ID of the VM instance.
Operational status [operational_status]	Operational status of the agent service. Possible values are Operational or Non-Operational.
Operating System Platform [operating_system_platform]	Operating system type of the VM instance.
Resource Type [resource_type]	Type of resource managed by SSM. Possible values are EC2Instance or ManagedInstance.
Version [version]	Version of the SSM agent.

CI relationships

The Amazon AWS - SSM Cloud Agents (LP) pattern creates the following relationships to support AWS SSM Cloud Agents discovery.

CI relationships

CI	Relationship	CI
Cloud System Management Agent [cmdb_ci_cloud_system_management_agent]	Runs on::Runs	Virtual Machine Instance [cmdb_ci_vm_instance]

AWS Step Functions State Machine pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Step Functions State Machines on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Step Functions State Machine - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Step Functions State Machine [cmdb_aws_step_functions_state_machine]

Field	Description
Name [name]	Name of the state machine.
Object ID [object_id]	The Amazon Resource Name (ARN) of the state machine.
Type [type]	Type of the state machine. The value is STANDARD or EXPRESS.
Configuration Item [configuration_item]	References the Cloud Function [cmdb_ci_cloud_function] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Step Functions State Machine - Extended Inventory (LP) pattern.

Cloud Function [cmdb_ci_cloud_function]

Field	Description
Name [name]	Name of the state machine.
Object ID [object_id]	The ARN of the state machine.
Edition [edition]	Type of the state machine. The value is STANDARD or EXPRESS.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - Step Functions State Machine - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Step Functions State Machine discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Function [cmdb_ci_cloud_function]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Step Functions State Machine [cmdb_aws_step_functions_state_machine]	Configuration Item [configuration_item]	Cloud Function [cmdb_ci_cloud_function]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Function [cmdb_ci_cloud_function]

AWS Tag discovery

The Amazon AWS - Step Functions State Machine - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Function [cmdb_ci_cloud_function] table.

AWS Storage pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Elastic Block Store (EBS) storage volumes on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Storage (LP) pattern.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	The name or ID, if no name is specified for the volume.
Object ID [object_id]	The ID of the volume.
State [state]	The volume state. The following values are valid: creating, available, in-use, deleting, deleted, or error.
Storage type [storage_type]	Type of storage. The value is set to **block**.
Volume ID [volume_id]	The ID of the EBS volume.
Size bytes [size_bytes]	The size of the volume, in bytes.
Size [size]	The size of the volume, in Gibibyte (GiB).

Block Endpoint [cmdb_ci_endpoint_block]

Field	Description
Name [name]	The name or ID, if no name is specified for the volume.
Object ID [object_id]	The ID of the volume.

CI relationships and references

The Amazon AWS - Storage (LP) pattern creates the following relationships and references to support AWS Storage discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	Block Endpoint [cmdb_ci_endpoint_block]
Block Endpoint [cmdb_ci_endpoint_block]	Implement End Point To::Implement End Point From	Storage Volume [cmdb_ci_storage_volume]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Storage Volume [cmdb_ci_storage_volume]

AWS Tag discovery

The Amazon AWS - Storage (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Storage Volume [cmdb_ci_storage_volume] table.

AWS Storage Gateway File Share pattern-based discovery

Discovery and Service Mapping Patterns finds AWS services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Storage Gateway File Share - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Storage Gateway - File Share [cmdb_aws_storage_gateway_file_share]

Field	Description
Object Id [object_id]	The Amazon Resource Name (ARN) uniquely identifying the resource.
Name [name]	The name assigned to the resource or file share.
Configuration Item [configuration_item]	References the Cloud Gateway [cmdb_ci_cloud_gateway] table.
File Share Type [file_share_type]	The type of file share provided by the Storage Gateway. For example: NFS or SMB.
Gateway Arn [gateway_arn]	The ARN of the AWS Storage Gateway associated with the resource.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Storage Gateway File Share - Extended Inventory (LP) pattern.

Cloud Gateway [cmdb_ci_cloud_gateway]

Field	Description
Object ID [object_id]	The ARN uniquely identifying the resource.
Name [name]	The name assigned to the resource or file share.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **AWS::StorageGateway::FileShare**.

CI relationships

The Amazon AWS - Storage Gateway File Share - Extended Inventory (LP) pattern creates these relationships to support AWS Storage Gateway File Share discovery.

CI	Relationship	CI
Cloud Gateway [cmdb_ci_cloud_gateway]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
AWS Storage Gateway - File Share [cmdb_aws_storage_gateway_file_share]	References	Cloud Gateway [cmdb_ci_cloud_gateway]

AWS tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

AWS Storage Gateway Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Storage Gateway Gateways on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Storage Gateway Gateway - Extended Inventory (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Storage Gateway Gateway [cmdb_aws_storage_gateway_gateway]

Field	Description
Name [name]	Name of the gateway.
Object ID [object_id]	The Amazon Resource Name (ARN) of the gateway.
Gateway ID [gateway_id]	The unique identifier of the gateway.
Gateway Type [gateway_type]	The type of the gateway. For example: STORED, CACHED, VTL, FILE_S3, or FILE_FSX_SMB.
Gateway State [gateway_state]	The current state of the gateway. The value is RUNNING or SHUTDOWN.
Host Environment [host_environment]	The host environment in which the gateway is running. For example: EC2, VMWARE, HYPER-V, KVM, SNOWBALL, or OTHER.
Configuration Item [configuration_item]	References the Cloud Gateway [cmdb_ci_cloud_gateway] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Storage Gateway Gateway - Extended Inventory (LP) pattern.

Cloud Gateway [cmdb_ci_cloud_gateway]

Field	Description
Name [name]	Name of the gateway.
Object ID [object_id]	The ARN of the gateway.
Environment [environment]	The host environment in which the gateway is running. For example: EC2, VMWARE, HYPER-V, KVM, SNOWBALL, or OTHER.
Description [short_description]	Type of resource. The value is set to **AWS::StorageGateway::Gateway**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - Storage Gateway Gateway - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Storage Gateway Gateway discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Gateway [cmdb_ci_cloud_gateway]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Storage Gateway Gateway [cmdb_aws_storage_gateway_gateway]	Configuration Item [configuration_item]	Cloud Gateway [cmdb_ci_cloud_gateway]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Gateway [cmdb_ci_cloud_gateway]

AWS Tag discovery

The Amazon AWS - Storage Gateway Gateway - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Gateway [cmdb_ci_cloud_gateway] table.

AWS sub account Organizations pattern-based discovery

Discovery and Service Mapping Patterns finds AWS sub-accounts in your organization on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Sub Account (LP) pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Unique identifier (ID) of the account.
Object ID [object_id]	Unique ID of the account.
Datacenter Type [datacenter_type]	Type of datacenter. The value is set to **AWS Datacenter [cmdb_ci_aws_datacenter]**.
Name [name]	User-friendly name of the account.
Is management account [is_master_account]	Indicates if this account is the management account of the organization.
Account Email [account_email]	Email address of the AWS service account.
Datacenter URL [datacenter_url]	URL of the datacenter used for discovery.
Install Status [install_status]	Install status of the account, derived from the account state.
Operational status [operational_status]	Operational status of the account, derived from the account state.
Parent account [parent_account]	References the management Cloud Service Account [cmdb_ci_cloud_service_account] record.

CI relationships and references

The Amazon AWS - Sub Account (LP) pattern creates the following relationships and references to support AWS sub account discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Service Account [cmdb_ci_cloud_service_account]	Members::Member of	Cloud Service Account [cmdb_ci_cloud_service_account]

CI references

CI	Field	Referenced CI
Cloud Service Account [cmdb_ci_cloud_service_account]	Parent account [parent_account]	Cloud Service Account [cmdb_ci_cloud_service_account]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Service Account [cmdb_ci_cloud_service_account]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Tag discovery

The Amazon AWS - Sub Account (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table.

AWS sub account regional pattern-based discovery

Discovery and Service Mapping Patterns finds member accounts and the primary account within an AWS Organization. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS sub account discovery pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	The unique identifier (ID) of the AWS account.
Name [name]	Name assigned to the account.
Datacenter Type [datacenter_type]	Type of datacenter. The value is set to **AWS Datacenter [cmdb_ci_aws_datacenter]**.
Datacenter URL [datacenter_url]	The URL of the AWS datacenter.
Object ID [object_id]	The unique ID of the AWS account.
Is management account [is_master_account]	Indicates whether this account is the management (primary) account of the organization. • **true**: Management account • **false**: Member account
Discovery credentials [discovery_credentials]	AWS credential from the Credentials [discovery_credentials] table.
Parent account [parent_account]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table (management account).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS sub account discovery pattern creates these relationships to support AWS sub account discovery.

CI	Relationship	CI
Cloud Service Account [cmdb_ci_cloud_service_account] (management account)	Members::Member of	Cloud Service Account [cmdb_ci_cloud_service_account] (member account)
Cloud Service Account [cmdb_ci_cloud_service_account] (member account)	References	Cloud Service Account [cmdb_ci_cloud_service_account] (management account)
Key Value [cmdb_key_value]	References	Cloud Service Account [cmdb_ci_cloud_service_account] (member account)

AWS tag discovery

The Amazon AWS sub account discovery pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table (member account).

AWS subnet pattern-based discovery

Discovery and Service Mapping Patterns finds AWS subnets on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Subnet (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Subnet [cmdb_ci_cloud_subnet]

Field	Description
Name [name]	The name or ID, if no name is specified for the subnet.
Object ID [object_id]	The ID of the subnet.
CIDR [cidr]	The IPv4 CIDR block assigned to the subnet. If an IPv6 CIDR block is also assigned, it is appended to the IPv4 value.
Available IP Count [available_ip_count]	The number of unused private IPv4 addresses in the subnet. The IPv4 addresses for any stopped instances are considered unavailable.
State [state]	The current state of the subnet. The following values are valid: pending or available.
Gateway [gateway]	ID of the internet gateway attached to the subnet, if applicable.

CI relationships and references

The Amazon AWS - Subnet (LP) pattern creates the following relationships and references to support AWS Subnet discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Network [cmdb_ci_network]	Contains::Contained by	Cloud Subnet [cmdb_ci_cloud_subnet]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud Subnet [cmdb_ci_cloud_subnet]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Subnet [cmdb_ci_cloud_subnet]

AWS Tag discovery

The Amazon AWS - Subnet (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Subnet [cmdb_ci_cloud_subnet] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Systems Manager Document pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Systems Manager Documents on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Systems Manager Document - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Systems Manager Document [cmdb_aws_systems_manager_document]

Field	Description
Name [name]	Name of the Systems Manager document.
Object ID [object_id]	The Amazon Resource Name (ARN) of the document, in the following format: `arn:aws:ssm:{region}:{account-id}:document/{document-name}`.
Document Format [document_format]	The format of the document. The value is JSON, YAML, or TEXT.
Document Type [document_type]	The type of the document. For example: Command, Automation, Session, Policy, or Package.
Document Version [document_version]	The document version.
Owner [owner]	The AWS account ID that owns the document.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Systems Manager Document [cmdb_aws_systems_manager_document] (continued)

Field	Description
Author [author]	The author of the document.
Display Name [display_name]	The display name of the document.
Review Status [review_status]	The review status of the document. The value is APPROVED, PENDING, REJECTED, or NOT_REVIEWED.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Systems Manager Document - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the Systems Manager document.
Object ID [object_id]	The ARN of the document, in the following format: `arn:aws:ssm:{region}:{account-id}:document/{document-name}`.
Resource type [resource_type]	Type of resource. The value is set to **AWS::SystemsManager::Document**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Systems Manager Document - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Systems Manager Document discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
AWS Systems Manager Document [cmdb_aws_systems_manager_document]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Systems Manager Document - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS Systems Manager Parameter Store pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Systems Manager Parameter Store parameters on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Systems Manager Parameter Store - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Systems Manager Parameter Store [cmdb_aws_systems_manager_parameter_store]

Field	Description
Name [name]	Name of the parameter.
Object ID [object_id]	The Amazon Resource Name (ARN) of the parameter.
Type [type]	The type of the parameter. The value is String, StringList, or SecureString.
Version [version]	The version number of the parameter.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Systems Manager Parameter Store - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the parameter.
Object ID [object_id]	The ARN of the parameter.
Resource type [resource_type]	Type of resource. The value is set to **AWS::SystemsManager::ParameterStore**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Systems Manager Parameter Store - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Systems Manager Parameter Store discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Systems Manager Parameter Store [cmdb_aws_systems_manager_parameter_store]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Systems Manager Parameter Store - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS tag discovery with Patterns

The ServiceNow Discovery and Service Mapping applications can use patterns to discover the tags for AWS resources and configuration item (CI) types, and then populate the CMDB with these discoveries. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

This tag-specific discovery pattern is in addition to the other AWS patterns available on the ServiceNow Store.This pattern extends the ones used for discovering the respective AWS resources.

Discovery populates the CMDB with tags for the following AWS patterns and CI types:

- Amazon AWS - ACL (LP) [cmdb_ci_network_acl]
- Amazon AWS API Gateway [cmdb_ci_cloud_gateway]
- Amazon AWS - Application and Network LB (LP) [cmdb_ci_cloud_load_balancer]
- Amazon AWS Application and Network LBs Events
- Amazon AWS - Classic LB (LP) [cmdb_ci_cloud_load_balancer]
- Amazon AWS Classic LB Events
- Amazon AWS Cognito [cmdb_ci_cloud_authentication]
- Amazon AWS - Customer Gateway (LP) [cmdb_ci_customer_gateway]
- Amazon AWS DB Cluster

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Amazon AWS DynamoDB [cmdb_ci_dynamodb_table]
- Amazon AWS ElastiCache
- Amazon AWS - ECS [cmdb_ci_cloud_service_account]
- Amazon AWS - Executable Template (LP) [cmdb_ci_os_template]
- Amazon AWS - Host (LP)
- Amazon AWS - Internet Gateway (LP)
- Amazon AWS keyspace
- Amazon AWS - Key Pair (LP)
- Amazon AWS Lambda [cmdb_ci_cloud_function]
- Amazon AWS - LB Pool (LP)
- Amazon AWS - NAT Gateway (LP)
- Amazon AWS - Network (LP)
- Amazon AWS Network Events
- Amazon AWS - NIC (LP)
- Amazon AWS - Owned Template (LP)
- Amazon AWS - Public IP Address (LP)
- Amazon AWS Redshift
- Amazon AWS Relational Database Service [cmdb_ci_cloud_database]
- Amazon AWS Resource Inventory
- Amazon AWS - Route Table (LP)
- Amazon AWS S3 [cmdb_ci_cloud_object_storage]
- Amazon AWS - Security Group (LP)
- Amazon AWS Security Group Events
- Amazon AWS Serverless Database
- Amazon AWS - Storage (LP)
- Amazon AWS Storage Events
- Amazon AWS sub account discovery
- Amazon AWS - Subnet (LP)
- Amazon AWS - Virtual Server (LP)
- Amazon AWS Virtual Server Events
- Amazon AWS - VPN Connections (LP)
- Amazon AWS - VPN Gateway (LP)

Note: For more information about the tags and patterns for each of these resources, see the documentation for the specific AWS resource.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

AWS Credentials

- On the ServiceNow AI Platform, configure AWS credentials, using a secret key and an access key.
- Have an Amazon **Account ID**. Use the Account ID as it appears in the AWS Management Console. You can run Discovery from the service account form.
- Set read-only permissions for the following URL: `https//tagging.<region>.amazonaws.com.`

Discovery schedule

Create a schedule for the relevant AWS service account. For more information, see Create a discovery schedule in Cloud Discovery Workspace.

Pattern tasks

When discovering AWS tags, the pattern completes the following high-level tasks:

- Sets the Resource Type filter: Filters the tags to only those tags related to the pattern CI type.
- Initiates the Cloud API call: The API call to the cloud environment that collects the tags and parses their attributes in the Tags table.
- Parses the resource name from the Amazon Resource Name (ARN). Each Amazon resource is assigned a name and a unique identifier.

> **Note:** Parsing the resource name is needed only for patterns where the main identification section does not work with the ARN of the AWS resource. When that is the case, a relation attribute has to be parsed from the AWS ARN.

- Includes references to cloud resource type. Establishes the relationship between the tag and the cloud resource type.

Cloud Rest Call

The pattern includes an API called "Cloud Rest Call" Post Call used to collect the following data:

- API URL used to retrieve the data — "https://tagging." + $ldc[1].object_id + ".amazonaws.com/"
- API Body — "{\"ResourceTypeFilters\":[\"" + $resource_type + "\"]}"
- Headers — "Content-Type:application/x-amz-json-1.1,X-Amz-Target:ResourceGroupsTaggingAPI_20170126.GetResources"

This data displays as shown in the following window.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data from Cloud Rest Call

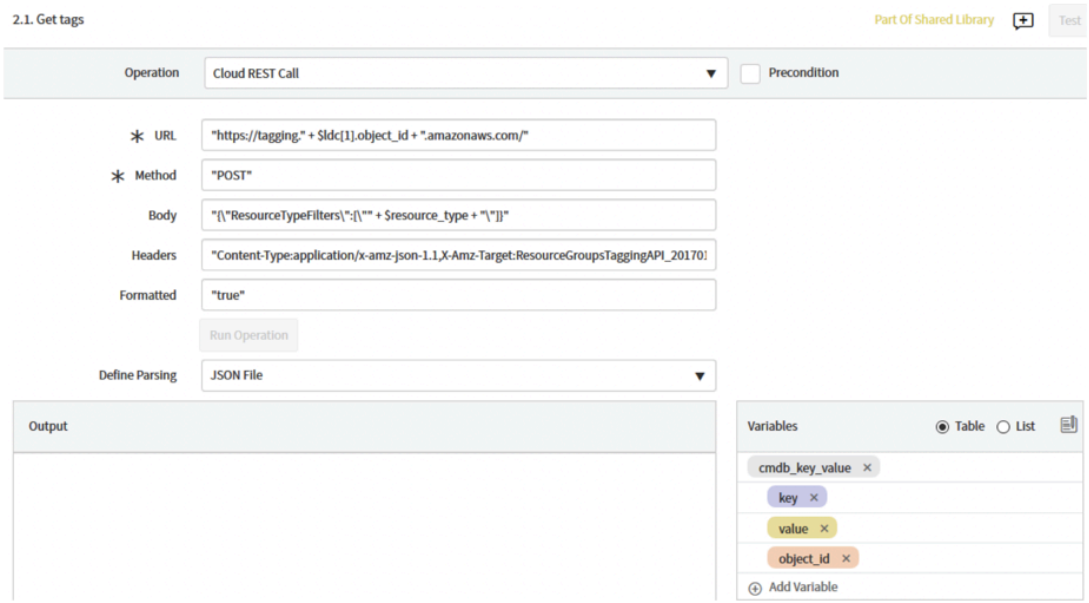

Data collected by Discovery during horizontal discovery

The pattern collects the following data during horizontal discovery:

Table and field	Additional information
CI type [cmdb_key_value]	
Key	Tag key
Value	Tag value
Object_id	Unique resource ID that identifies the resource in the AWS console.
Class model [cmdb_key_value]	
Extends table	Not applicable
Identifier – Object_id	Related entries to cloud resources
Metadata rules	Not applicable

This data is displayed as shown in the following window:

Data collected during horizontal discovery

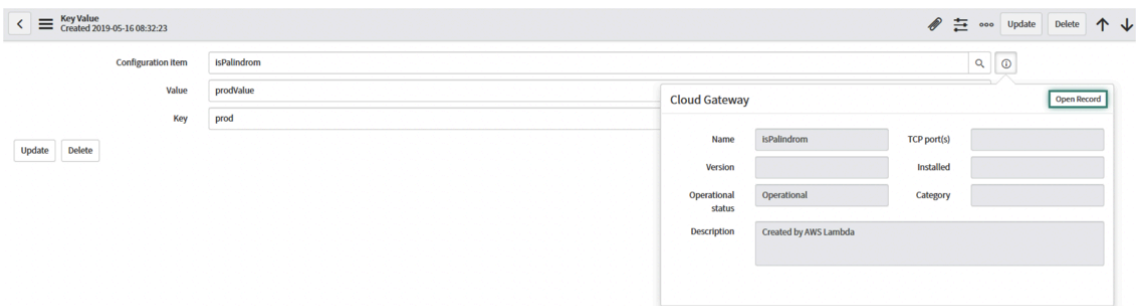

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS discovery pattern uses the following EVAL:

```
var rtrn;
var name = ${cmdb_key_value[].object_id}
name =
 name.substring(name.lastIndexOf("restapis/")+9,name.lastIndex);
name = name.substring(0,name.indexOf("/"));
rtrn = name;
```

Similar EVALS are available for patterns where the main identification section does not include ARN information in the main CI type table. When this information is not available, the Tag Extension section must parse a unique ARN used to establish the relationship to the related cloud resource.

AWS Transfer Family Server pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Transfer Family Servers on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - Transfer Family Server - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS Transfer Family Server [cmdb_aws_transfer_family_server]

Field	Description
Name [name]	The server ID, used as the display name.
Object ID [object_id]	The Amazon Resource Name (ARN) of the server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Transfer Family Server [cmdb_aws_transfer_family_server] (continued)

Field	Description
Server ID [server_id]	The unique identifier of the Transfer Family server.
Endpoint Type [endpoint_type]	The type of endpoint for the server. The value is PUBLIC, VPC, or VPC_ENDPOINT.
Identity Provider Type [identity_provider_type]	The mode of authentication for the server. For example: SERVICE_MANAGED, API_GATEWAY, AWS_DIRECTORY_SERVICE, or AWS_LAMBDA.
State [state]	The current state of the server. For example: ONLINE, OFFLINE, STARTING, STOPPING, START_FAILED, or STOP_FAILED.
User Count [user_count]	The number of users assigned to the server.
Logging Role [logging_role]	The ARN of the IAM role used to write transfer activity to Amazon CloudWatch Logs.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - Transfer Family Server - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	The server ID, used as the display name.
Object ID [object_id]	The ARN of the server.
Resource type [resource_type]	Type of resource. The value is set to **AWS::TransferFamily::Server**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - Transfer Family Server - Extended Inventory (LP) pattern creates the following relationships and references to support AWS Transfer Family Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS Transfer Family Server [cmdb_aws_transfer_family_server]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - Transfer Family Server - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

AWS virtual server pattern-based discovery

Discovery and Service Mapping Patterns finds AWS EC2 virtual machine instances on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites section in AWS discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Virtual Server (LP) pattern.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Name [name]	Name of the instance. Possible values are the tag name, if available, or the instance ID.
Object ID [object_id]	Unique identifier for the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance] (continued)

Field	Description
State [state]	Power state of the instance. Possible values: • On: Instance is running • Off: Instance is stopped
VM Instance ID [vm_inst_id]	Instance identifier from AWS.
Memory (MB) [memory]	Amount of memory allocated to the instance in megabytes (MB).
Network adapters [nics]	Number of network interface cards (NICs) attached to the instance.
Disks [disks]	Number of volumes attached to the instance.
CPUs [cpus]	Total number of virtual CPUs (vCPUs).
IP Address [ip_address]	Private IP address assigned to the instance.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) assigned to the instance.
Install Status [install_status]	Install status of the resource. For example: Installed or Retired.
Operational status [operational_status]	Operational status of the resource. For example: Operational, Non-Operational, or Retired.

Note:

If you have reactivated a retired virtual machine (VM) and want its active state to be accurately reflected in the **Operational status** and **Install Status** fields in the Virtual Machine Instance [cmdb_ci_vm_instance] table, set the *sn_itom_pattern.discover_aws_instance_statuses* MID Server property to **true**. Note that this setting involves additional API calls—one for each page of the pattern execution—which may impact performance if you're managing a large number of VMs.

Storage Mapping [cmdb_ci_storage_mapping]

Field	Description
Name [name]	Unique identifier for the storage mapping.
Object ID [object_id]	Unique identifier for the storage mapping.
Mapping Type [mapping_type]	The root device type used by the Amazon Machine Image (AMI). The AMI can use an EBS volume or an instance store volume.
Mount Point [mount_point]	Device name where the volume is mounted.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Mapping [cmdb_ci_storage_mapping] (continued)

Field	Description
	For example: `/dev/sda1`.
Host [host]	Public IP address of the instance, if available.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Subnet Endpoint [cmdb_ci_endpoint_subnet]

Field	Description
Object ID [object_id]	Subnet identifier.
Name [name]	Subnet identifier.
IP Address [ip_address]	IP address associated with the subnet endpoint.
Region [region]	AWS region where the subnet is located.

Cloud Subnet [cmdb_ci_cloud_subnet]

Field	Description
Object ID [object_id]	Subnet identifier.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	Domain Name System (DNS) name assigned to the instance.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Object ID [object_id]	AMI identifier.
Provider [provider]	Cloud provider, which is **AWS**. This field is only populated in the Cloud Image [cmdb_ci_cloud_os_image] table.

Note: When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Amazon AWS - Virtual Server (LP) pattern creates these relationships to support AWS virtual server discovery.

CI	Relationship	CI
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Virtual Machine Instance [cmdb_ci_vm_instance]
Virtual Machine Instance [cmdb_ci_vm_instance]	Contains::Contained by	Storage Mapping [cmdb_ci_storage_mapping]
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
DNS Name [cmdb_ci_dns_name]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Compute Template [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]
Virtual Machine Instance [cmdb_ci_vm_instance]	Contains::Contained by	DNS Name [cmdb_ci_dns_name]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	Subnet Endpoint [cmdb_ci_endpoint_subnet]
Cloud Subnet [cmdb_ci_cloud_subnet]	Implement End Point To::Implement End Point From	Subnet Endpoint [cmdb_ci_endpoint_subnet]
Network [cmdb_ci_network]	Contains::Contained by	Cloud Subnet [cmdb_ci_cloud_subnet]
Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Key Value [cmdb_key_value]	References	Virtual Machine Instance [cmdb_ci_vm_instance]

AWS Tag discovery

The Amazon AWS - Virtual Server (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

AWS BYOL discovery

Starting with Discovery and Service Mapping Patterns version 1.30.2, the pattern extension section discovers the following license types for Windows and Red Hat Enterprise Linux (RHEL) VMs:

- Bring Your Own License (BYOL)
- License included

The pattern stores the license type and model in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The license type key, which is one of the following: • Windows_OS_License_Type_automatic • RHEL_OS_License_Type_automatic
Value [value]	The license model, which is one of the following: • BYOL • License Included
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

AWS VPN Connections pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Virtual Private Network (VPN) connections on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - VPN Connections (LP) pattern.

Note: This pattern is not supported for CN regions.

VPN Connection [cmdb_ci_vpn_connection]

Field	Description
Name [name]	The name or ID, if no name tag is specified for the VPN connection.
Object ID [object_id]	The ID of the VPN connection.
State [state]	The current state of the VPN connection. The following values are valid: pending, available, deleting, or deleted.

Customer Gateway [cmdb_ci_customer_gateway]

Field	Description
Object ID [object_id]	The ID of the customer gateway.

Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

Field	Description
Object ID [object_id]	The ID of the virtual private gateway.

CI relationships and references

The Amazon AWS - VPN Connections (LP) pattern creates the following relationships and references to support AWS VPN Connections discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
VPN Connection [cmdb_ci_vpn_connection]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Customer Gateway [cmdb_ci_customer_gateway]	Contains::Contained by	VPN Connection [cmdb_ci_vpn_connection]
Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]	Contains::Contained by	VPN Connection [cmdb_ci_vpn_connection]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	VPN Connection [cmdb_ci_vpn_connection]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Tag discovery

The Amazon AWS - VPN Connections (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the VPN Connection [cmdb_ci_vpn_connection] table.

AWS VPN Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds AWS virtual private gateways on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - VPN Gateway (LP) pattern.

Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

Field	Description
Name [name]	The name or ID, if no name tag is specified for the VPN gateway.
Object ID [object_id]	The ID of the virtual private gateway.
Connection Type [connection_type]	The type of VPN connection the virtual private gateway supports.

Virtual Private Gateway Endpoint [cmdb_ci_endpoint_vpg]

Field	Description
Name [name]	The name or ID, if no name tag is specified for the VPN gateway endpoint.
Object ID [object_id]	The ID of the virtual private gateway.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Private Gateway Endpoint [cmdb_ci_endpoint_vpg] (continued)

Field	Description
Region [region]	The AWS region of the virtual private gateway endpoint.

CI relationships and references

The Amazon AWS - VPN Gateway (LP) pattern creates the following relationships and references to support AWS VPN Gateway discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]	Implement End Point To::Implement End Point From	Virtual Private Gateway Endpoint [cmdb_ci_endpoint_vpg]
Cloud Network [cmdb_ci_network]	Use End Point To::Use End Point From	Virtual Private Gateway Endpoint [cmdb_ci_endpoint_vpg]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

AWS Tag discovery

The Amazon AWS - VPN Gateway (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway] table.

AWS web ACL pattern-based discovery

Discovery and Service Mapping Patterns finds AWS web access control lists (web ACLs) on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Security Operations users can leverage the integration with Discovery to import web ACL rules and load balancers with attached web ACLs. For more information on setting ACL rules and using the Mitigation Controls Monitoring app, see Configure the AWS WAF integration for mitigation controls monitoring.

AWS web ACL data model

The Amazon AWS - Web ACL (LP) pattern introduces the following CI class that extends an existing CMDB class.

CI class introduced by this pattern

CI class	Extends from
Web ACL [cmdb_ci_web_acl]	Virtual Machine Object [cmdb_ci_vm_object]

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Web ACL (LP) pattern.

Web ACL [cmdb_ci_web_acl]

Field	Description
Name [name]	Name of the web ACL.
Object ID [object_id]	Unique ID for the web ACL from AWS.
Description [short_description]	Description of the web ACL provided by AWS.
Default Action [default_action]	Default action when no rules in the web ACL match. The value is **allow** or **deny**.

Note: Security Operations users can leverage the integration with Discovery to import web ACL rules and load balancers with attached web ACLs. For more information on setting ACL rules and using the Mitigation Controls Monitoring app, see Configure the AWS WAF integration for mitigation controls monitoring.

CI relationships

The Amazon AWS - Web ACL (LP) pattern creates the following relationships to support Amazon AWS Web ACL discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Web ACL [cmdb_ci_web_acl]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

AWS Windows Server pattern-based discovery

Discovery and Service Mapping Patterns finds AWS Windows Server configuration items (CIs) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Verify that AWS SSM is enabled

Verify that AWS Systems Manager (AWS SSM) is enabled on the Amazon Elastic Compute Cloud (Amazon EC2) instances.

Enable Windows Server CI creation

Set the *sn_itom_pattern.aws_cloud_discovery_populate_server_ci* system property to **true**. For more information, see Configure Server CI creation during AWS cloud discovery.

Verify SSM Agent execution context

The SSM Agent must run as SYSTEM to retrieve the serial number using `Get-CimInstance Win32_BIOS`. If your security policy restricts the SSM Agent to a non-SYSTEM user, the pattern can't create a Server CI for that instance.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Amazon AWS - Windows Server (LP) pattern.

Windows Server [cmdb_ci_win_server]

Field	Description
Name [name]	OS-level hostname of the EC2 instance. If the *glide.discovery.hostname.include_domain* property is set to false, the domain suffix is removed and only the short hostname is stored.
Object ID [object_id]	The EC2 instance ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Server [cmdb_ci_win_server] (continued)

Field	Description
Correlation ID [correlation_id]	The EC2 instance ID.
Serial number [serial_number]	BIOS serial number of the EC2 instance.
IP Address [ip_address]	IP address of the EC2 instance.
CPU core count [cpu_core_count]	Number of vCPUs available to the EC2 instance.
CPU count [cpu_count]	CPU count. The value is set to **1**.
RAM (MB) [ram]	Amount of memory available to the EC2 instance, in megabytes (MB).
Operating System [os]	Operating system of the server. The value is set to **Windows**.
Is Virtual [virtual]	Indicates that the server is a virtual machine. The value is set to **true**.
Install Status [install_status]	Install status of the resource. The value is set to **Installed**.
Operational status [operational_status]	Operational status of the resource. The value is set to **Operational**.

Serial Number [cmdb_serial_number]

Field	Description
Serial number [serial_number]	BIOS serial number of the EC2 instance.
Serial number type [serial_number_type]	Type of serial number. The value is set to **bios**.
Valid [valid]	Indicates that the serial number is active and valid. The value is set to **true**.
Configuration item [cmdb_ci]	References the Windows Server [cmdb_ci_win_server] table.

CI relationships and references

The Amazon AWS - Windows Server (LP) pattern creates the following relationships and references to support AWS Windows Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Windows Server [cmdb_ci_win_server]	Virtualized by::Virtualizes	Virtual Machine Instance [cmdb_ci_vm_instance]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration item [cmdb_ci]	Windows Server [cmdb_ci_win_server]

AWS X-Ray Sampling Rule pattern-based discovery

Discovery and Service Mapping Patterns finds AWS X-Ray Sampling Rules on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the AWS discovery prerequisites

For more information, see the prerequisites section in AWS discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see AWS Resource Inventory discovery with Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Amazon AWS - XRay Sampling Rule - Extended Inventory (LP) pattern.

You can review the non-CMDB AWS tables by navigating to **All** > **Configuration** > **AWS**. You can also search the navigation filter for the specific pattern name.

AWS XRay Sampling Rule [cmdb_aws_xray_sampling_rule]

Field	Description
Name [name]	Name of the sampling rule.
Object ID [object_id]	The Amazon Resource Name (ARN) of the sampling rule.
Priority [priority]	The priority of the sampling rule. Lower values indicate higher priority.
Resource ARN [resource_arn]	The ARN of the resource to which the sampling rule applies. An asterisk (*) matches all resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS XRay Sampling Rule [cmdb_aws_xray_sampling_rule] (continued)

Field	Description
Service Name [service_name]	The name of the service to which the sampling rule applies. An asterisk (*) matches all services.
Service Type [service_type]	The type of the service to which the sampling rule applies.
Version [version]	The version number of the sampling rule.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Amazon AWS - XRay Sampling Rule - Extended Inventory (LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the sampling rule.
Object ID [object_id]	The ARN of the sampling rule.
Resource type [resource_type]	Type of resource. The value is set to **AWS::XRay::SamplingRule**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Amazon AWS - XRay Sampling Rule - Extended Inventory (LP) pattern creates the following relationships and references to support AWS X-Ray Sampling Rule discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]

CI references

CI	Field	Referenced CI
AWS XRay Sampling Rule [cmdb_aws_xray_sampling_rule]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

AWS Tag discovery

The Amazon AWS - XRay Sampling Rule - Extended Inventory (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Limit AWS discovery to datacenters with resources

Optimize AWS discovery by limiting it to datacenters with resources.

Before you begin

- Verify that you have installed Discovery and Service Mapping Patterns, starting with version 1.29.0.
- Verify that your service account has the following IAM permissions to access the Resource Explorer API (starting with Discovery and Service Mapping Patterns version 1.35.0).

Resource Explorer API permissions and coverage

IAM permission	Coverage
`resource-explorer-2:Search`	Partial
`resource-explorer-2:Search` + `iam:CreateServiceLinkedRole`	Full

For more information, go to the AWS Documentation and search for the "Getting started with Resource Explorer" article.

Note: Discovery and Service Mapping Patterns versions 1.29.0 through 1.32.0 used the AWS Config service instead of the Resource Explorer API to determine datacenter activity. For instructions on configuring AWS Config recorder, go to the AWS Documentation and search for the "Recording resources in the AWS Config console" article.

Role required: discovery_admin

About this task

Starting with version 1.29.0, Discovery and Service Mapping Patterns introduces a new AWS datacenter discovery model that focuses discovery on datacenters with resources and excludes datacenters that don't contain resource. To limit discovery to datacenters that contain resources, set the *mid.cloud.discovery.sonar.discover_all_aws_datacenters* MID Server property to false. For more information, see AWS discovery using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** column, search for the `mid.cloud.discovery.sonar.discover_all_aws_datacenters` property.
3. Select the *`mid.cloud.discovery.sonar.discover_all_aws_datacenters`* property.
4. In the **Value** field, enter `false`.
5. Select **Update**.

What to do next

To exclude specific resource types from AWS datacenter discovery, configure the *`sn_itom_pattern.discovery.aws.ldc.excluded_resource_types`* system property with a comma-separated list of resource types to exclude (starting with Discovery and Service Mapping Patterns version 1.35.0). For more information, see Exclude AWS resource types from datacenter discovery.

Related topics

AWS discovery using patterns

Exclude AWS resource types from datacenter discovery

Exclude AWS resource types from datacenter discovery

You can configure which AWS resource types to exclude from AWS datacenter discovery, to avoid triggering discovery patterns in regions with only default resources.

Before you begin

- Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.
- Discover only datacenters with resources by setting the *`mid.cloud.discovery.sonar.discover_all_aws_datacenters`* MID Server property to **false**. For more information, see Limit AWS discovery to datacenters with resources.

Role required: discovery_admin

About this task

Some resource types, such as default Virtual Private Cloud (VPC) and security groups, exist in all regions regardless of whether workloads are deployed there. You can use the *`sn_itom_pattern.discovery.aws.ldc.excluded_resource_types`* system property to exclude these resource types. The default value of the property is empty, which includes all supported resource types in the query. Excluding these resources gives a more accurate view of which datacenters are active.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for the `sn_itom_pattern.discovery.aws.ldc.excluded_resource_types` property.
3. Select the *`sn_itom_pattern.discovery.aws.ldc.excluded_resource_types`* property.
4. In the **Value** field, enter a comma-separated list of AWS resource type identifiers to exclude. For example: `ec2:vpc,ec2:security-group`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a list of AWS resource type identifiers, go to the AWS Documentation and search for the "Resource types you can search for with Resource Explorer" article.

5. Select **Update**.

Related topics

Limit AWS discovery to datacenters with resources

AWS discovery using patterns

Configure Server CI creation during AWS cloud discovery

Create Windows Server or Linux Server configuration items (CIs) during AWS cloud discovery by enabling the *sn_itom_pattern.aws_cloud_discovery_populate_server_ci* system property.

Before you begin

- Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.
- Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns.
- Verify that AWS Systems Manager (AWS SSM) is enabled on the Amazon Elastic Compute Cloud (Amazon EC2) instances.
- Verify SSM Agent execution context.
  - For Linux: The SSM Agent must run as root to retrieve the serial number using `dmidecode`.
  - For Windows: The SSM Agent must run as SYSTEM to retrieve the serial number using `Get-CimInstance Win32_BIOS`.

Role required: admin

About this task

Before Discovery and Service Mapping Patterns version 1.35.0, Server CIs weren't created during AWS cloud discovery. Starting with version 1.35.0, the *sn_itom_pattern.aws_cloud_discovery_populate_server_ci* property controls whether Server CIs are created. The default value is false and Server CIs aren't created. You can set the property to true to create Server CIs for SSM-enabled EC2 instances.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `sn_itom_pattern.aws_cloud_discovery_populate_server_ci`.
3. Select the *sn_itom_pattern.aws_cloud_discovery_populate_server_ci* property.
4. In the **Value** field, enter `true`.
   To stop creating Server CIs during AWS cloud discovery, enter `false`.
5. Select **Update**.

What to do next

Run AWS cloud discovery or wait for the next scheduled discovery run for the changes to apply.

If your EC2 instances are slow to respond to SSM commands, increase the SSM command wait time. This provides more time before the first status check, enabling Server CIs to be created successfully. For more information, see set the SSM command wait time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

AWS discovery using patterns

AWS Linux Server pattern-based discovery

AWS Windows Server pattern-based discovery

Set the AWS SSM command wait time

Increase the wait time for AWS Systems Manager (AWS SSM) commands to support Server CI creation on slow-responding EC2 instances.

Before you begin

- Verify that the *sn_itom_pattern.aws_cloud_discovery_populate_server_ci* system property is set to `true`. For more information, see Configure Server CI creation during AWS cloud discovery.
- Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.

Role required: discovery_admin

About this task

By default, the MID Server waits 30 seconds (30000 ms) before checking whether the Server CI is created on a Windows or Linux EC2 instance. If instances in your environment are slow to respond, Server CIs may not be created. Increase the *sn_itom_pattern.discover_aws_server_command_initial_wait_ms* MID Server property value to provide more time before the first status check.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** column, search for `sn_itom_pattern.discover_aws_server_command_initial_wait_ms`.
3. Select the *sn_itom_pattern.discover_aws_server_command_initial_wait_ms* MID Server property.
4. In the **Value** field, enter the wait time in milliseconds.
   The default value is `30000` (30 seconds). The property accepts any positive integer.
5. Select **Update**.

Result

The MID Server applies the updated wait time when running the next SSM commands during AWS Cloud Discovery Server CI creation.

What to do next

Run AWS Cloud Discovery again to apply the change.

Related topics

AWS discovery using patterns

AWS Linux Server pattern-based discovery

AWS Windows Server pattern-based discovery

Configure Server CI creation during AWS cloud discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Google Cloud Platform (GCP) Cloud discovery using Patterns

Discovery and Service Mapping Patterns uses patterns to discover components of the Google Cloud Platform (GCP) deployment during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date.

- Discovery and Service Mapping Patterns starting version 1.8.0 (September 2023)
- CMDB CI Class Models starting version 1.49.0 (September 2023)
- Visibility Content starting version 6.13.0 (September 2023)

Verify the MID Server requirements

- Use a dedicated MID Server for cloud discovery. The MID Server must have 16 Cores, 8 GB with 4 GB Java Virtual Machine (JVM) memory allocation. For information on verifying MID Server specifications, see KB article KB2156327.
- Deploy and connect the dedicated MID Server with internet access.
- For more information on the MID Server configurations, see Install and configure MID Servers to access cloud environments.

Verify Google Cloud BigQuery and Google Cloud Bigtable requirements

To discover Google Cloud BigQuery and Google Cloud Bigtable, navigate to **All** > **Discovery** > **MID Server** > **Properties** and ensure that the *sn_itom_pattern.google_dbservice_table_collection* property is set to **True**.

Verify user role settings

Set the user permission for the Google Cloud Platform member to **Viewer**:

1. In the Google Cloud Platform console, navigate to **IAM**.
2. Select the relevant member from the list and click the **Edit** icon.
3. In the Edit permissions window, select `Viewer` from the **Role** list.
4. Click **Save**.

Configure Google API Credentials

For more information, see Configure the Google API credentials for Cloud Discovery.

Create a Google Cloud Platform service account

For more information, see Create GCP service accounts.

Enable collecting cloud logging events

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Google Cloud Platform patterns can collect cloud logging events generated by the native Google Cloud APIs. Configure the Google Cloud Logging service to auto-update the CMDB.

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Verify that a Discovery schedule is configured

For more information, see Create a GCP Discovery schedule in Discovery Admin Workspace.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Related topics

Kubernetes discovery using patterns

GCP Certificate Manager discovery

Google Cloud Platform (GCP) AlloyDB for PostgreSQL discovery using patterns

Discovery and Service Mapping Patterns uses the Google Cloud Platform (GCP) - AlloyDB for PostgreSQL pattern to discover AlloyDB for PostgreSQL during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - AlloyDB for PostgreSQL pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Service Account [cmdb_ci_cloud_service_account] (continued)

Field	Description
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Cluster ID [cluster_id]	The GCP system-generated unique identifier (UID) of the resource.
Name [name]	Name of the cluster resource.
Automated Backups [automated_backup]	Determines whether automated backups are enabled. Possible values are Enabled or Disabled.
Install Status [install_status]	Install status of the cluster. Default value is Installed.
Operational status [operational_status]	Operational status of the cluster. Default value is Operational.
Cluster Status [cluster_status]	Current state of this cluster. For example: READY or STOPPED.
Cluster Version [cluster_version]	Optional field that displays the database engine's major version.
Cluster Type [cluster_type]	Type of cluster. Possible values are AlloyDB – Primary cluster or AlloyDB – Secondary cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	The GCP system-generated UID of the resource.
Name [name]	Name of the instance resource.
Install Status [install_status]	Install status of the database. Default value is Installed.
Operational status [operational_status]	Operational status of the database. Default value is Operational.
State [state]	State of the database. Possible values are Available or Terminated.
CPU Count [cpu_count]	Number of central processing units (CPUs) in the virtual machine (VM) instance.
Type [type]	Type of database: AlloyDB – PostgreSQL.
Replication Type [replication_type]	Type of the instance specified at creation time. For example: Primary instance, Read pool instance, or Secondary instance.
IP Address [ip_address]	IP address for the Instance, which is the endpoint for an end-user application.

On the Dependency Views map, you can view all discovered AlloyDB for PostgreSQL resources in your organization and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AlloyDB for PostgreSQL instance dependency view

AlloyDB for PostgreSQL Cluster dependency view

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

Discovery creates these relationships to support the AlloyDB for PostgreSQL discovery.

CI	Relationship	CI
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Cloud DataBase [cmdb_ci_cloud_database]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Cluster of::Cluster	Cloud DataBase [cmdb_ci_cloud_database]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Replicates::Replicated by	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]
Key Value [cmdb_key_value]	References	Cloud DataBase [cmdb_ci_cloud_database]
Key Value [cmdb_key_value]	References	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Tag discovery

The Google Cloud Platform (GCP) - AlloyDB for PostgreSQL pattern collects tags and populates them in two entries in the Key Value [cmdb_key_value] table. One entry references the Cloud DataBase table, the other entry references the Cloud DataBase Cluster table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag key.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table
Configuration item [configuration_item]	References the Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] table

GCP BigQuery pattern-based discovery

Discovery and Service Mapping Patterns finds GCP BigQuery datasets and tables on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - BigQuery DB pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Name [name]	Name of the BigQuery dataset.
Object ID [object_id]	The fully qualified unique identifier of the dataset in the format **projectId:datasetId**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
State [state]	State of the dataset.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) for accessing the dataset.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Public DataBase [cmdb_ci_cloud_public_database]

Field	Description
Name [name]	Name of the public BigQuery dataset.
Object ID [object_id]	The fully qualified unique identifier of the dataset in the format **projectId:datasetId**.
Vendor [vendor]	Cloud vendor. Value is set to GCP.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
State [state]	State of the public dataset.
Fully qualified domain name [fqdn]	FQDN for accessing the public dataset.

Database [cmdb_ci_database]

Field	Description
Name [name]	Name of the BigQuery table.
Type [type]	Database type. Value is set to **bigquery**.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Google Cloud Platform (GCP) - BigQuery DB pattern creates these relationships to support GCP BigQuery discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud Service Account [cmdb_ci_cloud_service_account]	Uses::Used by	Cloud Public DataBase [cmdb_ci_cloud_public_database]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase [cmdb_ci_cloud_database]	Replicates to::Replicated by	Google Datacenter [cmdb_ci_google_datacenter]
Cloud Public DataBase [cmdb_ci_cloud_public_database]	Replicates to::Replicated by	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase [cmdb_ci_cloud_database]	Contains::Contained by	Database [cmdb_ci_database]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Key Value [cmdb_key_value]	References	Cloud DataBase [cmdb_ci_cloud_database]

GCP Tag discovery

The Google Cloud Platform (GCP) - BigQuery DB pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

GCP Bigtable pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Bigtable instances on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Bigtable DB pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Name [name]	Display name of the Bigtable instance.
Object ID [object_id]	Unique identifier for the Bigtable instance resource.
Location [location]	Geographic regions where the Bigtable instance clusters are deployed.
Install Status [install_status]	Install status of the Bigtable instance. Default value is Installed.
State [state]	Current operational state of the Bigtable instance. For example, Available or Terminated.
Operational status [operational_status]	Operational status of the Bigtable instance. Default value is Operational.
Type [type]	Resource type identifier. Value is set to **bigtable-instance**.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) for the Bigtable instance.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Name [name]	Name of the Bigtable cluster extracted from the cluster identifier.
Install Status [install_status]	Install status of the cluster. Default value is Installed.
Operational status [operational_status]	Operational status of the cluster. Default value is Operational.
Node Count [node_count]	Number of nodes allocated to serve the cluster.
Cluster Type [cluster_type]	Type of database cluster. Value is set to **GCP Bigtable**.
Location [location]	Geographic region where the cluster is deployed.

Database [cmdb_ci_database]

Field	Description
Name [name]	Full resource identifier of the Bigtable table.
Type [type]	Database type identifier. Value is set to **bigtable**.
Install Status [install_status]	Install status of the table. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Google Cloud Platform (GCP) - Bigtable DB pattern creates these relationships to support GCP Bigtable discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Cloud DataBase [cmdb_ci_cloud_database]	Contains::Contained by	Database [cmdb_ci_database]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud DataBase [cmdb_ci_cloud_database]	Replicates to::Replicated by	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Cluster of::Cluster	Cloud DataBase [cmdb_ci_cloud_database]
Key Value [cmdb_key_value]	References	Cloud DataBase [cmdb_ci_cloud_database]

GCP Tag discovery

The Google Cloud Platform (GCP) - Bigtable DB pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

GCP Cloud Firestore pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Cloud Firestore database instances on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Firestore DB pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	Unique identifier for the Cloud Firestore database instance.
Name [name]	Name of the Cloud Firestore database instance.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of Cloud Firestore database. Value is either Firestore Native Mode or Firestore Datastore Mode.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) for the Cloud Firestore database instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Google Cloud Platform (GCP) - Firestore DB pattern creates these relationships to support GCP Cloud Firestore discovery.

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase [cmdb_ci_cloud_database]	Replicates to::Replicated by	Google Datacenter [cmdb_ci_google_datacenter]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

Google Cloud Platform (GCP) Cloud Functions discovery using patterns

Discovery and Service Mapping Patterns uses the Google Cloud Platform (GCP) - Cloud Functions pattern to discover Cloud Functions in a Google Cloud Platform (GCP) deployment during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites

For more information, see the prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Cloud Functions pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud Function [cmdb_ci_cloud_function]

Field	Description
Version [version]	Version ID of the function, available only for Gen1 Cloud Functions.
Object ID [object_id]	Name fetched from the API response. For example: `projects/automation-226410/locations/us-central1/functions/auto-labeller`.
Name [name]	Name parsed from the API response. For example: `auto-labeller`.
Function Last Modified [function_last_modified]	Timestamp indicating the last time the function was updated.
CodeSha256 [codesha256]	SHA-256 hash algorithm of the function's code used to verify the integrity of the build. For example: `273a583f-ef1c-49c2-9bae-7d21bf3b5f7a`.
Language [language]	Runtime environment or programming language used for the function.
Install Status [install_status]	Install status of the function. Default value is Installed.
Operational status [operational_status]	Operational status of the function. Default value is Operational.

On the Dependency Views map, you can view all discovered GCP Cloud Functions resources in your organization and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP Cloud Functions dependency view

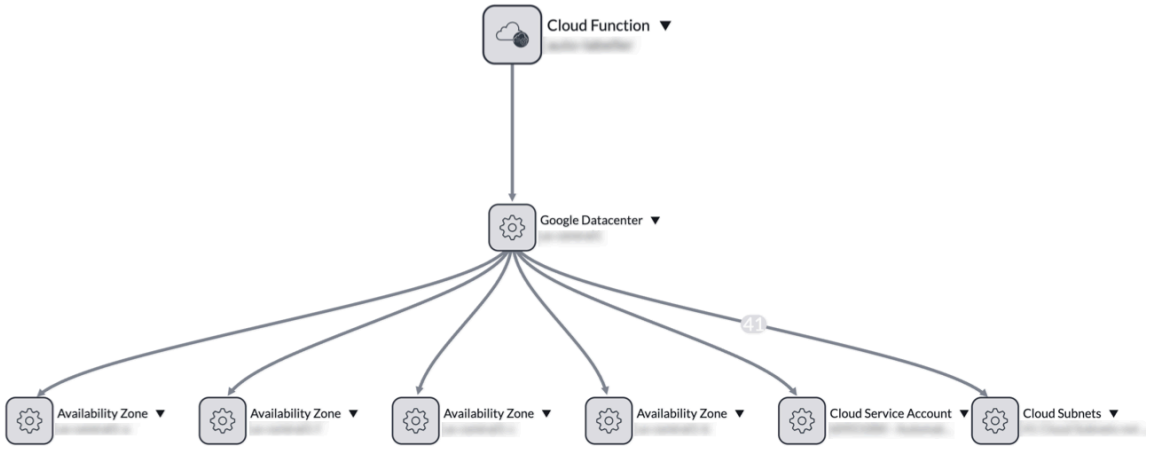

CI relationships

Discovery creates these relationships to support the GCP Cloud Functions discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud Function [cmdb_ci_cloud_function]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Key Value [cmdb_key_value]	References	Cloud Function [cmdb_ci_cloud_function]

Tag discovery

The Google Cloud Platform (GCP) - Cloud Functions pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Function [cmdb_ci_cloud_function] table.

GCP Cloud SQL pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Cloud SQL instances and their databases on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery uses the Google Cloud Platform (GCP) - Cloud SQL DB pattern to discover hardware type configurations for Cloud SQL instances from your GCP account.

When using the Hardware Type [cmdb_ci_compute_template] table to store the hardware types, you may notice an unusually large number of records. To avoid this issue, you can store the discovered hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. When enabled, Discovery populates the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. For more information, see Enable the Cloud Hardware Type class extension.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Cloud SQL DB pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	The unique resource identifier URL for the Cloud SQL instance.
Name [name]	The name of the Cloud SQL instance.
Type [type]	The database engine type for the Cloud SQL instance.
State [state]	The state of the Cloud SQL instance. For example: Available or Terminated.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Version [version]	The database version.
Fully qualified domain name [fqdn]	The unique resource identifier URL for the Cloud SQL instance.

Database [cmdb_ci_database]

Field	Description
Name [name]	The name of the database within the Cloud SQL instance.
Type [type]	The type of database.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Object ID [object_id]	The machine type for the Cloud SQL instance.
Name [name]	The machine type for the Cloud SQL instance.
vCPUs [vcpus]	The number of virtual CPUs (vCPUs) for the Cloud SQL instance.
Provider [provider]	The cloud provider, which is GCP. This field is only populated in Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

CI relationships

The Google Cloud Platform (GCP) - Cloud SQL DB pattern creates these relationships to support GCP Cloud SQL DB discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud DataBase [cmdb_ci_cloud_database]	Contains::Contained by	Database [cmdb_ci_database]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase [cmdb_ci_cloud_database]	Provisioned From::Provisioned	Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Cloud DataBase [cmdb_ci_cloud_database]	Replicates to::Replicated by	Cloud DataBase [cmdb_ci_cloud_database]
Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Key Value [cmdb_key_value]	References	Cloud DataBase [cmdb_ci_cloud_database]

GCP Tag discovery

The Google Cloud Platform (GCP) - Cloud SQL DB pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

GCP Disk Types pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Disk Types on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Disk Types pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud Disk Type [cmdb_ci_disk_type]

Field	Description
Object ID [object_id]	Unique identifier for the disk type in the format **zone:disk-type-name**.
Name [name]	Display name of the disk type in the format **disk-type-name@zone**.
Description [short_description]	Description of the disk type.
Valid Disk Size (GB) [valid_disk_size]	Valid disk size range for this disk type, in gigabytes (GB). For example **10GB-65536GB**.
Default Disk Size (GB) [default_disk_size_gb]	Default disk size in GB for this disk type.

CI relationships

The Google Cloud Platform (GCP) - Disk Types pattern creates these relationships to support GCP Disk Types discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Disk Type [cmdb_ci_disk_type]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud Disk Type [cmdb_ci_disk_type]	Hosted on::Hosts	Availability Zone [cmdb_ci_availability_zone]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

GCP Events pattern-based discovery

Discovery uses event patterns to update GCP component data in near real-time. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Events discovered by Discovery during horizontal discovery

When GCP components change state, events are created. Discovery uses event patterns to detect these changes and update the CMDB.

Patterns used for event discovery

Pattern	CI
Google Cloud Platform (GCP) - Networking - Events	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Cloud Platform (GCP) - Networking Firewall - Events	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Cloud Platform (GCP) - Subnetwork - Events	Cloud Service Account [cmdb_ci_cloud_service_account]

GCP External IP Addresses pattern-based discovery

Discovery and Service Mapping Patterns finds GCP external IP addresses on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - External IP Addresses pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]

Field	Description
Object ID [object_id]	Unique identifier for the external IP address resource.
IP Address [ip_address]	The external IP address.
Public IP Address [public_ip_address]	The public IP address.
Name [name]	Name of the external IP address resource.
State [state]	State of the IP address.
Scope [scope]	Geographic scope of the IP address.
Network Interface Owner [network_interface_owner]	The resource that is using this IP address.
Public DNS [public_dns]	Public Domain Name System (DNS) name associated with the IP address.

CI relationships

The Google Cloud Platform (GCP) - External IP Addresses pattern creates these relationships to support GCP External IP Addresses discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

GCP Firebase Realtime Database pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Firebase Realtime Database instances on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Firebase Realtime DB pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Google Datacenter [cmdb_ci_google_datacenter] (continued)

Field	Description
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	Unique identifier for the Firebase Realtime Database instance.
Name [name]	Name of the Firebase Realtime Database instance.
Type [type]	Type of Firebase database instance.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
State [state]	State of the Firebase Realtime Database instance.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) for the Firebase Realtime Database instance.

CI relationships

The Google Cloud Platform (GCP) - Firebase Realtime DB pattern creates these relationships to support GCP Firebase Realtime DB discovery.

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

GCP Load Balancer pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Load Balancers on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The Discovery patterns discover GCP Load Balancing infrastructure:

- Google Cloud Platform (GCP) - Load Balancer - HTTP pattern: Discovers Layer 7 (application layer) load balancers that use HTTP and HTTPS protocols
- Google Cloud Platform (GCP) - Load Balancer - TCP - UDP pattern: Discovers Layer 4 (transport layer) load balancers that use TCP and UDP protocols

Both patterns populate the same CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Load Balancer patterns.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Field	Description
Object ID [object_id]	Unique identifier for the load balancer.
Name [name]	Name of the load balancer.
State [state]	State of the load balancer. Default value is Available.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Object ID [object_id]	Unique identifier for the load balancer service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Service [cmdb_ci_lb_service] (continued)

Field	Description
Name [name]	Name of the load balancer service.
Service Type [service_type]	Type of service. Possible values are Frontend or Backend.
Listener Protocol [listener_protocol]	Protocol used by the service listener.
Port [port]	Port number used by the service.
IP Address [ip_address]	IP address of the frontend service.
Pool [pool]	References the Load Balancer Pool [cmdb_ci_lb_pool] table. This field is populated only by the Google Cloud Platform (GCP) - Load Balancer - HTTP pattern.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Object ID [object_id]	Unique identifier for the load balancer pool.
Name [name]	Name of the load balancer pool.
Load balancing method [load_balancing_method]	Load balancing method used by the pool. This field is populated only by the Google Cloud Platform (GCP) - Load Balancer - TCP - UDP pattern.

Load Balancer Pool Member [cmdb_ci_lb_pool_member]

Field	Description
Object ID [object_id]	Unique identifier for the pool member, known as a VM instance ID in GCP.
Name [name]	Name of the pool member, known as a VM instance name in GCP.
Operational status [operational_status]	Operational status of the pool member. • Operational: Instance status is running • Non-Operational: Other instance states

Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]

Field	Description
Object ID [object_id]	Unique identifier for the health check service.
Name [name]	Name of the health check service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Load Balancer Health Service [cmdb_ci_lb_health_service] (continued)

Field	Description
Monitor type protocol [monitor_type]	Type of health check monitor For example: HTTP or TCP.
Port [port]	Port number used for health checks.
Timeout in seconds [timeout_sec]	Timeout in seconds for health check requests.
Unhealthy threshold [unhealthy_threshold]	Number of consecutive failures before marking as unhealthy.
Healthy threshold [healthy_threshold]	Number of consecutive successes before marking as healthy.
Interval in seconds [check_interval_sec]	Interval in seconds between health checks.
Request path [request_path]	Request path for HTTP and HTTPS health checks. This field is populated only by the Google Cloud Platform (GCP) - Load Balancer - HTTP pattern.

CI relationships

Both Load Balancer patterns create these relationships to support GCP Load Balancer discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Load Balancer Service [cmdb_ci_lb_service]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Load Balancer Pool [cmdb_ci_lb_pool]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Load Balancer Pool [cmdb_ci_lb_pool]	Owns::Owned by	Load Balancer Pool Member [cmdb_ci_lb_pool_member]
Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Google Cloud Platform (GCP) - Load Balancer - HTTP pattern creates relationships to support GCP Load Balancer discovery.

CI	Relationship	CI
Load Balancer Service [cmdb_ci_lb_service]	Uses::Used by	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Service [cmdb_ci_lb_service]	References	Load Balancer Pool [cmdb_ci_lb_pool]
Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]	Used by::Uses	Load Balancer Service [cmdb_ci_lb_service]

The Google Cloud Platform (GCP) - Load Balancer - TCP - UDP pattern creates relationships to support GCP Load Balancer discovery.

CI	Relationship	CI
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Uses::Used by	Load Balancer Pool [cmdb_ci_lb_pool]
Cloud Load Balancer Health Service [cmdb_ci_lb_health_service]	Used by::Uses	Load Balancer Pool [cmdb_ci_lb_pool]

Google Cloud Platform (GCP) Memorystore discovery using patterns

Discovery and Service Mapping Patterns uses the Google Cloud Platform (GCP) - Memorystore DB pattern to discover Memorystore for Memcached and Memorystore for Redis during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

GCP Memorystore data model

The Google Cloud Platform (GCP) - Memorystore DB pattern introduces the following CI class that extends an existing CMDB class.

CI class introduced by this pattern

CI class	Extends from
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Cluster [cmdb_ci_cluster]

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Memorystore DB pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Name [name]	Name of the cluster.
Cluster ID [cluster_id]	ID of the cluster.
Install Status [install_status]	Install status of the cluster. Default value is Installed.
Operational status [operational_status]	Operational status of the cluster. Default value is Operational.
Cluster Status [cluster_status]	Current state of this cluster. For example: READY or STOPPED.
Cluster Type [cluster_type]	Type of the cluster: Redis Cluster.
Disk Space (GB) [disk_space]	Disk space in gigabytes (GB).
IAM Authentication Enabled [iam_authentication_enabled]	Determines whether the IAM authentication is enabled. Possible values are true or false.
Node Count [node_count]	Number of replica nodes per shard.
Deletion Protection Enabled [deletion_protection_enabled]	Determines whether the deletion protection is enabled. Possible values are true or false.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] (continued)

Field	Description
Asset tag [asset_tag]	NodeType of a Redis cluster node. For example: REDIS_SHARED_CORE_NANO or REDIS_HIGHMEM_MEDIUM.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	Unique name of the resource in this scope, including the project and location.
Name [name]	Unique name of the resource, parsed from the Object ID.
Type [type]	Type of database. Possible values are memcache-instance or redis-instance.
Install Status [install_status]	Install status of the database. Default value is Installed.
Operational status [operational_status]	Operational status of the database. Default value is Operational.
State [state]	State of the database. Possible values are Available or Terminated.
Fully qualified domain name [fqdn]	Fully qualified domain name (FQDN) for the resource type.
IP Address [ip_address]	Hostname or IP address of the exposed endpoint used by clients to connect to the service.

On the Dependency Views map, you can view all discovered Memorystore for Memcached or Memorystore for Redis resources in your organization and the relationships between them.

Memorystore for Memcached or Memorystore for Redis instance dependency view

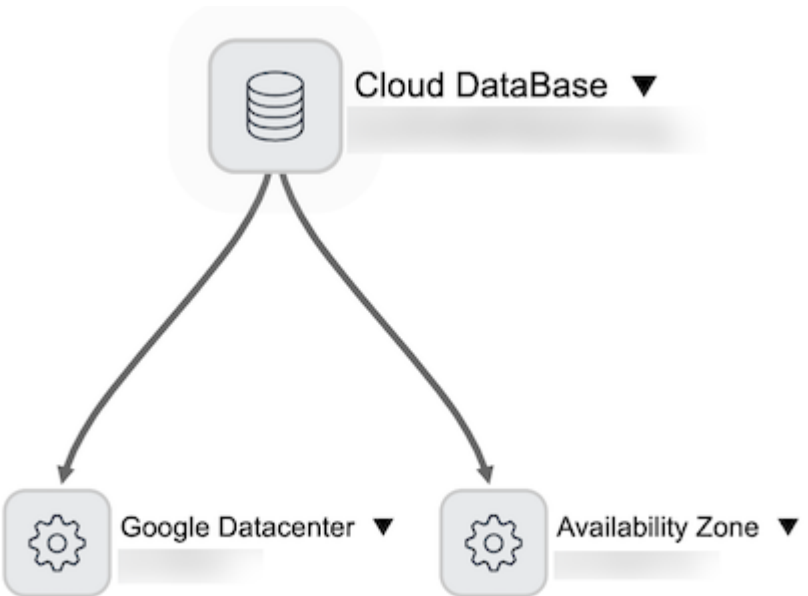

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Memorystore for Redis Cluster dependency view

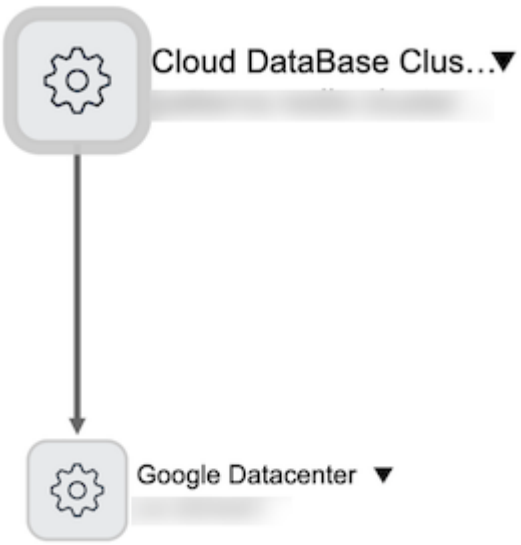

CI relationships

The Google Cloud Platform (GCP) - Memorystore DB pattern creates the following relationships and references to support Memorystore discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase [cmdb_ci_cloud_database]	Replicates to::Replicated by	Availability Zone [cmdb_ci_availability_zone]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud DataBase [cmdb_ci_cloud_database]

Tag discovery

The Google Cloud Platform (GCP) - Memorystore DB pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag key.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

GCP Networking pattern-based discovery

Discovery and Service Mapping Patterns finds GCP networking resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Networking pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Google Datacenter [cmdb_ci_google_datacenter] (continued)

Field	Description
Description [short_description]	Datacenter or region description.

Cloud Network [cmdb_ci_network]

Field	Description
Name [name]	Name of the Google Cloud Virtual Private Cloud (VPC) network.
Object ID [object_id]	Unique identifier for the VPC network resource in Google Cloud.
Description [short_description]	Description of the VPC network.
Default Gateway [default_gateway]	Default gateway IPv4 address for the network.
State [state]	Current state of the VPC network. Default value is Available.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Cloud Subnet [cmdb_ci_cloud_subnet]

Field	Description
Name [name]	Name of the Google Cloud subnetwork. If the network has auto-creation enabled, the Classless Inter-Domain Routing (CIDR) notation is used as the name.
Object ID [object_id]	Unique identifier for the subnetwork resource in GCP.
CIDR [cidr]	IP address range of the subnetwork in CIDR notation. May include primary IPv4, secondary IPv4, and IPv6 ranges as a comma-separated list.
Subnet Mask [subnet_mask]	Dotted representation of the subnet mask. For example: 255.255.240.0.
Gateway [gateway]	Gateway address for default routing out of the network.
Broadcast Address [broadcast_address]	Broadcast address of the subnet.
Available IP Count [available_ip_count]	Number of IPs that are available in the subnet. This amount does not include network and broadcast addresses.
State [state]	Current state of the subnetwork. Default value is Available.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network ACL [cmdb_ci_network_acl]

Field	Description
Name [name]	Name of the Google Cloud firewall rule.
Object ID [object_id]	Unique identifier for the firewall rule in Google Cloud.
Description [short_description]	Description of the firewall rule.
Install Status [install_status]	Install status of the firewall rule. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Network ACL Rule [cmdb_ci_network_acl_rule]

Field	Description
Name [name]	Name of the firewall rule.
Allow Deny [allow_deny]	Indicates whether traffic is allowed or denied.
Outbound [is_outbound]	Traffic direction for the firewall rule. • true: EGRESS • false: INGRESS
Source Ranges [source_ranges]	Source IP address ranges in CIDR notation for the firewall rule.
Destination Ranges [destination_ranges]	Destination IP address ranges in CIDR notation for the firewall rule.
Target Tags [target_tags]	Network tags that identify which instances the firewall rule affects.
Allowed\Denied Traffic [allowed_denied_traffic]	Protocols and ports that are allowed or denied by this rule.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Google Cloud Platform (GCP) - Networking pattern creates these relationships to support GCP networking discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Cloud Subnet [cmdb_ci_cloud_subnet]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Network [cmdb_ci_network]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud Network [cmdb_ci_network]	Contains::Contained by	Cloud Subnet [cmdb_ci_cloud_subnet]
Cloud Network [cmdb_ci_network]	Contains::Contained by	Network ACL [cmdb_ci_network_acl]
Network ACL [cmdb_ci_network_acl]	Contains::Contained by	Network ACL Rule [cmdb_ci_network_acl_rule]

Google Cloud Platform (GCP) Organization discovery with Patterns

The ServiceNow Discovery application uses the Discover Google Organization discovery pattern to find GCP Organization projects and resources. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

For details on system requirements and family compatibility, view the application listing on the ServiceNow Store website.

GCP Organization structure

Verify that a GCP Organization resource is structured correctly, with projects underneath it in a miscellaneous hierarchy.

Ensure that all projects under the GCP Organization resource have the Compute Engine API enabled.

Service Account user settings

Check that Service Account users have access to the GCP Organization resource node.

Verify that Service Account users are members of all other projects located under the GCP Organization resource node.

Ensure that Service Account users have the following Cloud Identity and Access Management credentials for projects under the GCP Organization resource:

- resourcemanager.folders.list
- resourcemanager.organizations.get
- resourcemanager.projects.get
- resourcemanager.projects.list

Credentials for API elements

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

During the discovery, the pattern uses the following API elements. The user that was added to the credentials in the instance must have permissions to send these queries.

- https://cloudresourcemanager.googleapis.com/v1/{name=organizations/*}
- https://cloudresourcemanager.googleapis.com/v1/projects
- https://cloudresourcemanager.googleapis.com/v2/folders:search

Credentials for creating a discovery schedule

Configure the following credentials:

1. Create GCP credentials.
2. Create a GCP Service Account:
   a. Credentials: The GCP credentials.
   b. Account ID: The Google Service Account ID.
   c. Datacenter type: Select **cmdb_ci_google_datacenter**.
3. Create a discovery schedule.

(Optional) Create a serverless discovery schedule

Create a discovery schedule to perform targeted discovery of GCP Organization resources.

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Click **New** and then fill in the form.

Discovery Schedule form

Field	Description
Name	Name for the Discovery schedule. For example, `Discover GCP Organization`.
Discover	Discover type. Select **Serverless**.
MID Server	Name of the MID Server.
Active	Option for enabling this schedule for discovery. Select this check box to enable discovery.

3. Right-click the header of the Discovery Schedule form and select **Save**.
4. Click the **Serverless Execution Patterns** tab, click **New**, and then fill in the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

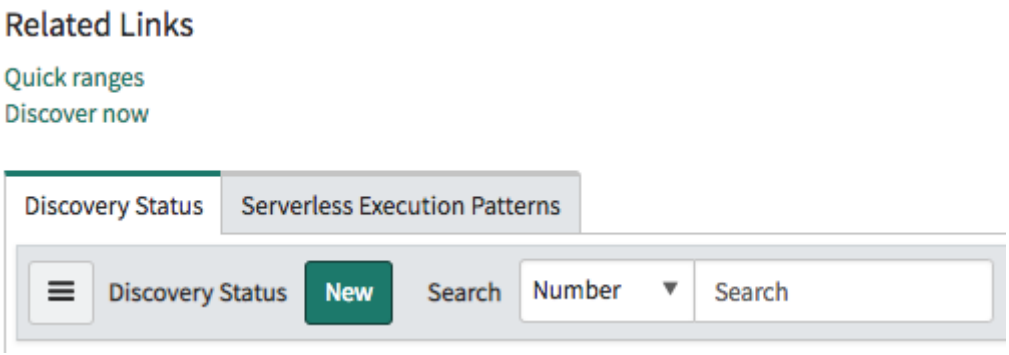

Serverless Execution Pattern form

Field	Description
Name	Name for this Serverless Execution Pattern. For example, `Discover GCP Organization`.
Pattern	Name of the pattern to run: Discover Google Organization.
Proxy Host	Fully qualified domain name of the machine on which you are installing the proxy server. Specify **Global**.
Active	Option for enabling this schedule for discovery. Select this check box to enable discovery.

5. Under **Discovery Pattern Launcher Parameters**, configure the following parameters with the relevant values:

Parameter	Value
cloud_account_id	The Project ID within GCP.
cloud_cred_id	The sysid of the GCP credentials.
cloud_datacenter_type	`cmdb_ci_google_datacenter`

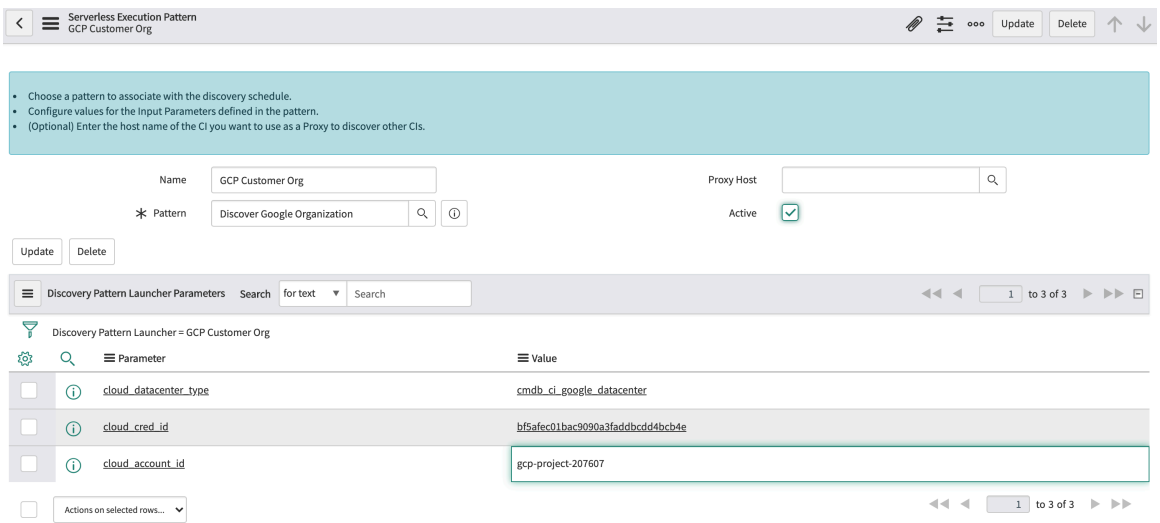

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Customers with early access to the GCP Organization pattern

Verify that the Discover-GCP-SubAccounts scheduled job is not enabled in your instance.

1. Navigate to **System Definition** > **Scheduled Jobs**.
2. Click the Discover-GCP-SubAccounts scheduled job.
3. Clear the **Active** check box, and then click **Update**.

Data collected by Discovery during horizontal discovery

Table and field	Description
Organization [cmdb_ci_cloud_org]	
name	The name of the organization resource.
object_id	The ID of the organization resource.
time	The time the organization resource was created in GCP.
operational_status	Current operational status. One of the following: • Operational • Non-Operational • Repair in Progress • DR Standby • Ready • Retired
Folder [cmdb_ci_gcp_folder]	
name	The name of the organization folder.
parent_id	The ID of the parent resource.
parent_type	The type of the parent resource. Can be `organization` or `folder`.
time	The time the resource was created in GCP.
status	The status of the folder according to the `lifecycleState` status in GCP.
object_id	The ID of the organization folder.
Project [cmdb_ci_gcp_project]	
name	The name of the project.
project_id	The ID of the project.
parent_id	The ID of the parent folder resource.
parent_type	The type of the parent resource. Can be `organization` or `folder`.
time	The time the resource was created in GCP.
operational_status	The status of the folder according to the `lifecycleState` status in GCP.
object_id	The ID of the project.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
discovery_credentials	GCP account credentials.
Resource [cmdb_key_value]	
key	The key, or label, associated with the GCP project. For example, `country`.
value	The project label value assigned to the GCP project. For example, `ca`.
configuration_item	The URL or path of the CI.

CI relationships

CI	Relationship	CI
Cloud Organization [cmdb_ci_cloud_org]	Contained by::Contains	Google Project [cmdb_ci_gcp_project]
Cloud Organization [cmdb_ci_cloud_org]	Contained by::Contains	Service Account
Cloud Organization [cmdb_ci_cloud_org]	Contained by::Contains	Google Folder [cmdb_ci_gcp_folder]
Service Account	Owns::Owned by	Google Project [cmdb_ci_gcp_project]
Google Folder [cmdb_ci_gcp_folder]	Contained by::Contains	Google Project [cmdb_ci_gcp_project]
Google Folder [cmdb_ci_gcp_folder]	Contained by::Contains	Sub Google Folder
Sub Google Folder	Contained by::Contains	Google Project [cmdb_ci_gcp_project]

Google Cloud Platform (GCP) resource inventory discovery with Patterns

The ServiceNow Discovery application uses the Google Cloud Platform (GCP) Resource Inventory pattern to find GCP resources and policies. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The pattern provides visibility for services supported by the Asset Inventory API, as well as collecting inventory data on the deployed GCP services and updating the CMDB.

The pattern collects inventory data either for all GCP-supported resources or for a preconfigured inclusion list of resources. The Cloud Inventory Resource Inclusion List contains all resource types supported by GCP Cloud Asset Inventory, except for Compute Engine resources and IAM policies. You can expand the inclusion list with additional resource types per your requirements. For more information about Google Cloud assets, see https://cloud.google.com/resource-manager/docs/cloud-asset-inventory/overview.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Starting with Discovery and Service Mapping Patterns version 1.18.1, you can discover the GCP storage by two types of discovery: Serverless and Cloud Discovery. Running Cloud Discovery schedules enables you to run one schedule per all your GCP projects without manually configuring separate schedules. You can also continue using Serverless discovery as before.

The enhanced Google Cloud Platform (GCP) asset inventory pattern runs queries to check the configurations on your system and trigger a discovery according to your configurations. The queries by order:

1. If you have the **Discovery Pattern Launcher Parameters** configured, the pattern triggers a Serverless discovery schedule.

2. If the former query carries no results, the pattern continues to query if you have the *mid.gcp_resource_inventory_bucketpath* MID Server property configured. If you do, the pattern runs a Cloud Discovery schedule by cloud accounts.

3. If the former query carries no results, the pattern continues to query if you have the *mid.gcp_resource_inventory_bucketpath.default* MID Server property configured. If you do, the pattern triggers a default Cloud Discovery schedule.

4. If none of these properties are configured, the pattern terminates gracefully.

Prerequisites

Verify the store apps are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

GCP authorization for Discovery to use the Cloud Asset API

- API endpoint: `https://cloudasset.googleapis.com/v1/projects/<account_id>:exportAssets`
- Required one or more of the following IAM permissions on the specified resource parent:
  - cloudasset.assets.exportResource
  - cloudasset.assets.exportIamPolicy

Service Account user for the cloud storage API

The ServiceNow Cloud Service Account need to have a read-only permission from GCP to access the API endpoint - https://www.googleapis.com/storage/v1.

Note: You can use the headers on the Encryption page to do the following:

- Download an object that is encrypted by a customer-supplied encryption key.
- Get object metadata with content hashes.

Permission to read and write to a Cloud Storage Bucket

- Storage Object Creator
- Storage Object Viewer
- Storage Object Admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a cloud storage bucket using Google Cloud console

1. Go to the **Google Cloud console**.
2. From the **Navigation menu**, select **Cloud Storage** > **Buckets**.
3. To create a new bucket, select **+ Create**.
4. On the **Create a bucket** page, fill in the bucket information.

Field	Description
Name your bucket	Enter a globally unique name for your bucket.
Choose where to store your data	Select a Location type and Location where the bucket data will be permanently stored. ○ Location type: Multi-region, for global storage (for example, us, eu, asia). ○ Location: List of the Cloud Storage locations available for storing your data.
Choose a storage class for your data	Select the appropriate storage class for your needs (for example, Standard, Nearline, Coldline, or Archive).
Choose how to control access to objects	Select whether or not your bucket enforces public access prevention.
Choose how to protect object data	Configure protection tools, if required.

5. Select **Create**.

Note: For more information, see Google Cloud Storage documentation.

Retention Policy for the storage bucket

Ensure that the Retention Policy for the storage bucket is not active. If the Retention Policy is active, the auto-generated inventory data file cannot be deleted by the pattern.

Create a Serverless discovery schedule

Create a discovery schedule to perform targeted discovery of GCP asset inventory.

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Click **New** and then fill in the

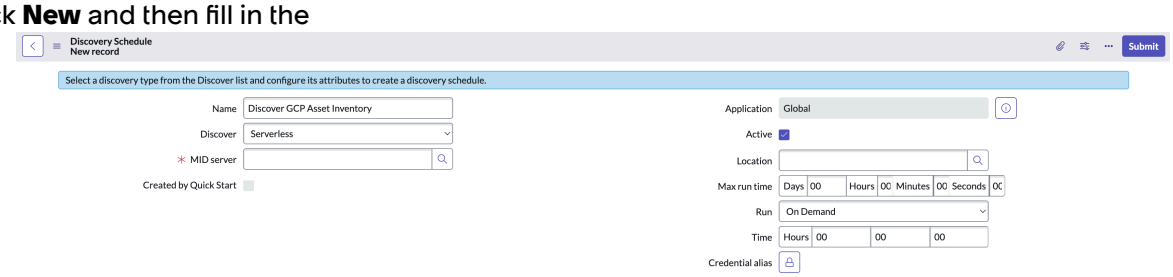

form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Schedule form

Field	Description
Name	Name for the Discovery schedule. For example, `Discover GCP Asset Inventory`.
Discover	Discover type. Select **Serverless**.
MID Server	Name of the MID Server.
Run	Option to select when should the next discovery run.

3. Right-click the header of the Discovery Schedule form and select **Save**.
4. Click the **Serverless Execution Patterns** tab, click **New**, and then fill in the form.

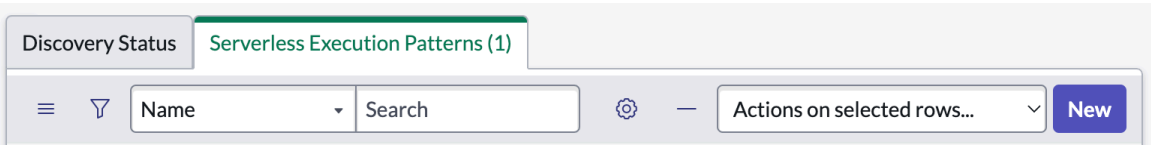

Serverless Execution Pattern form

Field	Description
Name	Name for this Serverless Execution Pattern. For example, `Discover GCP Asset Inventory`.
Pattern	Select the **Google Cloud Platform (GCP) Resource Inventory** pattern.
Proxy Host	Fully qualified domain name of the machine on which you are installing the proxy server. Specify **Global**.
Active	Option for enabling this schedule for discovery. Select this check box to enable discovery.

5. Select **Submit**.
6. In the **Discovery Pattern Launcher Parameters** tab, configure the following parameters with the relevant values:

Parameter	Value
cloud_account_id	The Project ID within GCP.
full_path_file	The complete file path of the storage bucket. For example: `gs://<bucketname>`.
cloud_cred_id	The sysid of the GCP credentials.
cloud_datacenter_type	`cmdb_ci_google_datacenter`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage discovery configurations with MID Server properties

1. Configure the *mid.gcp_resource_inventory_bucketpath* property.
   a. Navigate to **All** > **MID Server** > **Properties** and filter the list by `Name start with mid.gcp`.
   b. Select *mid.gcp_resource_inventory_bucketpath*.
   c. fill in the form.
      i. Configure the property **Name** field to include your account ID as follows: `mid.gcp_resource_inventory_bucketpath.<Cloud Account Id>`.
      ii. Fill in the **Value** field with the bucket URI, which is the complete file path of the storage bucket. For example: `gs://<bucketname>`.
      iii. In the **MID Server** field, leave it blank to set a MID Server property that affects all MID Servers. To set a MID Server property for a particularMID Server, select the preferred server.
      iv. Select **Update**.

2. Configure the *mid.gcp_resource_inventory_bucketpath.default* property.
   a. Navigate to **All** > **MID Server** > **Properties** and filter the list by `Name start with mid.gcp`.
   b. Select *mid.gcp_resource_inventory_bucketpath.default*.
   c. Fill in the **Value** field with the bucket URI, which is the complete file path of the storage bucket. For example: `gs://<bucketname>`.
   d. Select **Update**.

For more information, see Export asset metadata from one project to another

Cloud inventory resource inclusion list

- To collect inventory data for resources supported by GCP, in ServiceNow AI Platform, navigate to **Cloud Inventory Resource Inclusion List** and clear all GCP table records.

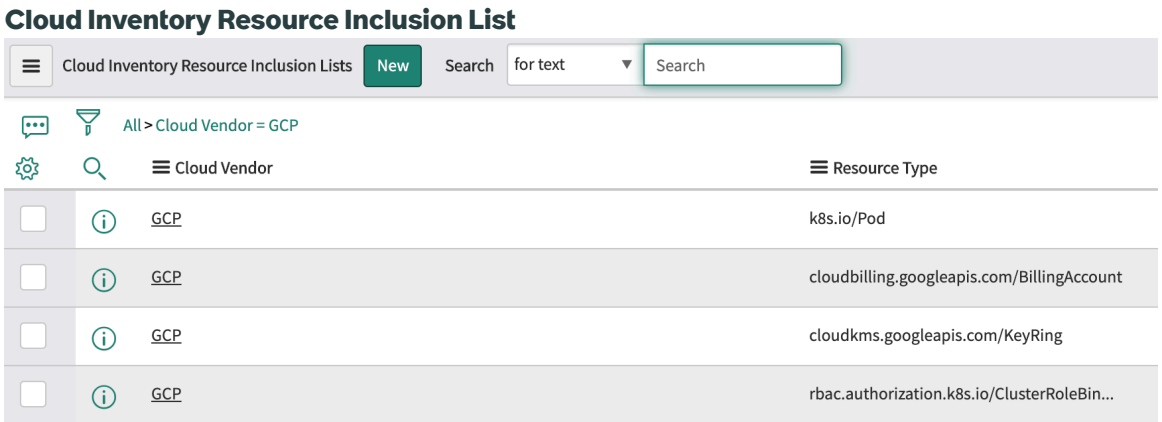

- Fine-tune GCP resource discovery using the Cloud Inventory Resource Inclusion List.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If your deployment has custom patterns for GCP discovery, ensure that you do not discover GCP resources twice:

1. Ensure that the application scope is Discovery and Service Mapping Patterns:
   a. Navigate to **Settings** > **Developer**.
   b. Select `Discovery and Service Mapping Patterns` from the **Application** list.
2. Navigate to **System Definitions** > **Tables**.
3. Open the Cloud Inventory Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table.
4. Under **Related Links**, click **Show List**.
5. Select resource types for which you have custom patterns, and select `Delete` from the **Actions on selected rows** list.

The Cloud Inventory Resource Inclusion List is predefined with common services. You can expand the list with additional resource types that you want the pattern to discover, as follows:

Note: If you modify the list provided in the base system, it is no longer updated automatically in application updates. You need to maintain customized lists yourself.

1. Open the Cloud Inventory Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table.
2. Click **New**.
3. Fill in the form, and then click **Submit**.

Note: The names of additional resource types must conform to the appropriate vendor naming conventions.

Field	Description
Cloud Vendor	The vendor of the resource type: GCP.
Resource Type	The GCP resource type value.
Application	The application scope: Discovery and Service Mapping Patterns.

The changes are applied the next time you run the pattern.

Data collected by Discovery during horizontal discovery

This pattern discovers data that provides visibility for all GCP services in your organization. The discovered data includes the following tables and fields.

Table and field	Description
Main CI [cmdb_ci_cmp_resource]	
object_id	The ID of the item. The item is accessed with this URL.
name	The name of the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
resource_type	The asset resource type, according to the data in the JSON file.
Key Value [cmdb_key_value]	
Key	The GCP tag key name.
Value	The GCP tag value name.

The Dependency Views map shows the discovered Configuration Items (CIs) in your organization and the relationships between them. Here, the only meaningful relationship between the CIs is the one that helps Discovery identify them.

Each GCP Inventory CI is related either to a Logical Datacenter (LDC) CI or to a Cloud Service Account CI. In this example, the Inventory CI is related to a Cloud Service Account CI.

Dependency Views map showing Cloud Service Account CI

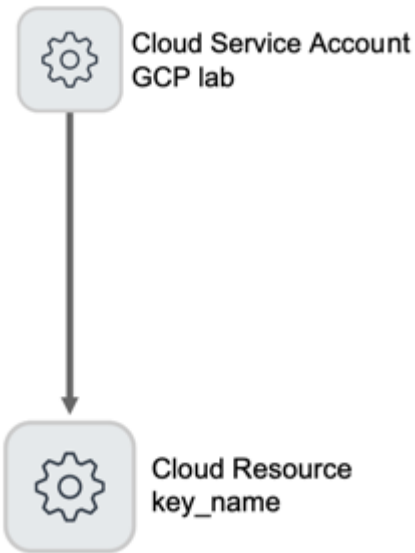

CI relationships

These relationships are created to support GCP asset inventory discovery:

CI	Relationship	CI
For Global Resources:		
Main CI [cmdb_ci_cmp_resource]	Contained by::Contains	Cloud Service Accounts
For Regional Resources:		
Main CI [cmdb_ci_cmp_resource]	HostedOn::Hosts	Logical Datacenter (LDC)

GCP Spanner pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Spanner instances and databases on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Spanner DB pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Object ID [object_id]	The full resource name of the Spanner instance.
Name [name]	The display name of the Spanner instance.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
State [state]	The state of the Spanner instance.
Type [type]	The type of cloud database resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase [cmdb_ci_cloud_database] (continued)

Field	Description
Node Count [node_count]	The number of nodes allocated to the Spanner instance.
Processing Units [processing_units]	The number of processing units allocated to the Spanner instance.
Fully qualified domain name [fqdn]	The fully qualified domain name (FQDN) for the Spanner instance.

Database [cmdb_ci_database]

Field	Description
Name [name]	The name of the Spanner database.
Type [type]	The type of database.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Google Cloud Platform (GCP) - Spanner DB pattern creates these relationships to support GCP Spanner discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud DataBase [cmdb_ci_cloud_database]	Contains::Contained by	Database [cmdb_ci_database]
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Cloud DataBase [cmdb_ci_cloud_database]	Replicates to::Replicated by	Google Datacenter [cmdb_ci_google_datacenter]
Key Value [cmdb_key_value]	References	Cloud DataBase [cmdb_ci_cloud_database]

GCP Tag discovery

The Google Cloud Platform (GCP) - Spanner DB pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

GCP SSH Keys pattern-based discovery

Discovery and Service Mapping Patterns finds GCP SSH keys stored in project metadata on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - SSH Keys pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Key Pair [cmdb_ci_cloud_key_pair]

Field	Description
Name [name]	The username associated with the SSH key.
Object ID [object_id]	The SSH public key used for authentication to Compute Engine virtual machines.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Google Cloud Platform (GCP) - SSH Keys pattern creates these relationships to support GCP SSH key discovery.

CI	Relationship	CI
Cloud Key Pair [cmdb_ci_cloud_key_pair]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

GCP Storage pattern-based discovery

Discovery and Service Mapping Patterns finds GCP persistent disks and snapshots on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Storage pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Object ID [object_id]	The unique identifier for the persistent disk.
Name [name]	The name of the persistent disk.
Volume ID [volume_id]	The unique identifier for the persistent disk.
Storage type [storage_type]	The type of persistent disk. For example: pd-standard, pd-balanced, or pd-ssd.
Size bytes [size_bytes]	The size of the persistent disk in bytes.
State [state]	The operational state of the disk. For example: Available, In Use, or Terminated.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Description [short_description]	The description of the persistent disk.

Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]

Field	Description
Object ID [object_id]	The unique identifier for the snapshot.
Name [name]	The name of the snapshot.
Volume Name [volume_name]	The name of the source persistent disk.
Capacity [capacity]	The storage size of the snapshot in bytes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Disks size (GB) [disks_size]	The total size of all attached disks in gigabytes (GB).

CI relationships

The Google Cloud Platform (GCP) - Storage pattern creates these relationships to support GCP storage discovery.

CI	Relationship	CI
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Storage Volume [cmdb_ci_storage_volume]	Replicates to::Replicated by	Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]
Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]	Provisioned From::Provisioned	Storage Volume [cmdb_ci_storage_volume]
Key Value [cmdb_key_value]	References	Storage Volume [cmdb_ci_storage_volume]
Key Value [cmdb_key_value]	References	Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

GCP Tag discovery

The Google Cloud Platform (GCP) - Storage pattern collects labels and populates them in two entries in the Key Value [cmdb_key_value] table. One entry references the Storage Volume table, the other entry references the Storage Volume Snapshot table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The label key.
Value [value]	The label value.
Configuration item [configuration_item]	References the Storage Volume [cmdb_ci_storage_volume] table.
Configuration item [configuration_item]	References the Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP virtual server pattern-based discovery

Discovery and Service Mapping Patterns finds GCP Virtual Machine (VM) instances and related resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the GCP discovery prerequisites section in Google Cloud Platform (GCP) Cloud discovery using Patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Google Cloud Platform (GCP) - Virtual Server pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Name of the project used for the discovery.
Object ID [object_id]	Name of the project used for the discovery.
Datacenter Type [datacenter_type]	Datacenter type: Google Datacenter [cmdb_ci_google_datacenter].

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Name [name]	Name of the availability zone.
Description [short_description]	Description of the availability zone.
State [state]	State of the Availability Zone. Possible values are Available or Terminated.

Google Datacenter [cmdb_ci_google_datacenter]

Field	Description
Name [name]	Datacenter or region name.
Region [region]	Datacenter or region name.
Object ID [object_id]	Unique identifier allocated by GCP for this resource.
Description [short_description]	Datacenter or region description.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Object ID [object_id]	Unique identifier for the VM instance.
VM Instance ID [vm_inst_id]	Unique identifier for the VM instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance] (continued)

Field	Description
Name [name]	Name of the VM instance.
State [state]	Operational state of the virtual machine. For example: On, Off, or Terminated.
Install Status [install_status]	Install status of the resource. Default value is Installed.
CPUs [cpus]	Number of virtual CPUs (vCPUs) allocated to the instance.
Memory (MB) [memory]	Amount of memory in megabytes (MB) allocated to the instance.
Termination Protection [termination_protection]	Indicates whether deletion protection is enabled and the instance can't be deleted. Possible values are true or false.
Disks [disks]	Number of disks attached to the VM instance.
Disks size (GB) [disks_size]	Total size of all disks attached to the instance in gigabytes (GB).
Network adapters [nics]	Number of network interfaces attached to the VM instance.
IP Address [ip_address]	Primary internal IP address of the instance.

Instance template [cmdb_ci_instance_template]

Field	Description
Object ID [object_id]	Unique identifier for the instance template.
Name [name]	Name of the instance template.
Description [short_description]	Description of the template configuration.

Cloud Mgmt Network Interface [cmdb_ci_nic]

Field	Description
Object ID [object_id]	Unique identifier combining VM ID and Network Interface Controller (NIC) name.
Name [name]	Name of the network interface.
Private IP [private_ip]	Primary internal IP address.
IP Address [ip_address]	IP address of the network interface.
Public IP [public_ip]	Public IP address of the network interface.

VNIC Endpoint [cmdb_ci_endpoint_vnic]

Field	Description
Object ID [object_id]	Unique identifier for the virtual NIC endpoint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VNIC Endpoint [cmdb_ci_endpoint_vnic] (continued)

Field	Description
Name [name]	Name of the virtual NIC endpoint.
Region [region]	Region where the endpoint is located.
IP Address [ip_address]	IP address of the endpoint.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Object ID [object_id]	Unique identifier for the persistent disk volume.
Name [name]	Name of the disk volume.
Volume ID [volume_id]	Volume identifier.
Storage type [storage_type]	Type of persistent disk. For example: **pd-balanced**, **pd-ssd**, or **pd-standard**.
Size bytes [size_bytes]	Size of the volume in bytes.
State [state]	State of the volume. For example: Available, In Use, or Terminated.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Description [short_description]	Description of the volume.

Block Endpoint [cmdb_ci_endpoint_block]

Field	Description
Object ID [object_id]	Unique identifier for the block storage endpoint.
Name [name]	Name of the storage endpoint.
Host [host]	VM instance name to which the disk is attached.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Name [name]	Name of the operating system (OS) image.
Object ID [object_id]	Unique identifier for the OS image.
Image Type [image_type]	Source type of the image. For example: RAW or machine.
Image Source [image_source]	Source of the image in the format: **disk_name/disk_id** or **snapshot_name/snapshot_id**.
FQDN [fqdn]	Fully qualified domain name (FQDN) of the image resource.
Provider [provider]	The cloud provider, which is **GCP**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image] (continued)

Field	Description
	This field is only populated in Cloud Image [cmdb_ci_cloud_os_image] table.

Note: When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Object ID [object_id]	Unique identifier for the machine type.
Name [name]	Name of the machine type. For example: **n1-standard-4**.
Short description [short_description]	Description of the machine type configuration.
vCPUs [vcpus]	Number of vCPUs in the machine type.
Memory MB [memory_mb]	Amount of memory available in the machine type in MB.
Local storage GB [local_storage_gb]	Amount of local storage space in GB.
Zone [zone]	Availability zone where the machine type is available. For example: **us-west4-b**.
Provider [provider]	The cloud provider, which is **GCP**. This field is only populated in Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

Note: When using the Hardware Type [cmdb_ci_compute_template] table to store the hardware types, you may notice an unusually large number of records. To avoid this issue, you can store the discovered hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. For more information, see Enable the Cloud Hardware Type class extension.

CI relationships

The Google Cloud Platform (GCP) - Virtual Server pattern creates these relationships to support GCP Virtual Machine discovery.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Google Datacenter [cmdb_ci_google_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Virtual Machine Instance [cmdb_ci_vm_instance]
Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Instance template [cmdb_ci_instance_template]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
Instance template [cmdb_ci_instance_template]	Uses::Used by	Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Instance template [cmdb_ci_instance_template]	Uses::Used by	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]
VNIC Endpoint [cmdb_ci_endpoint_vnic]	Implement End Point To::Implement End Point From	Cloud Mgmt Network Interface [cmdb_ci_nic]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	VNIC Endpoint [cmdb_ci_endpoint_vnic]
Cloud Mgmt Network Interface [cmdb_ci_nic]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	Google Datacenter [cmdb_ci_google_datacenter]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Block Endpoint [cmdb_ci_endpoint_block]	Implement End Point To::Implement End Point From	Storage Volume [cmdb_ci_storage_volume]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	Block Endpoint [cmdb_ci_endpoint_block]
Key Value [cmdb_key_value]	References	Virtual Machine Instance [cmdb_ci_vm_instance]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Key Value [cmdb_key_value]	References	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

GCP label discovery

The pattern extension collects labels and populates them in two entries in the Key Value [cmdb_key_value] table. One entry references the Virtual Machine Instance table, the other entry references the Image or Cloud Image table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Label name.
Value [value]	Label value.
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.
Configuration item [configuration_item]	References the Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image] table.

GCP BYOL discovery

The pattern extension section discovers the following license types for Windows and Red Hat Enterprise Linux (RHEL), and SQL Server VMs:

- Bring Your Own License (BYOL)
- License included

The pattern stores the license type and model in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The license type key, which is one of the following: • Windows_OS_License_Type_automatic • RHEL_OS_License_Type_automatic • SQL_Server_License_Type_automatic
Value [value]	The license model, which is one of the following: • BYOL • License Included
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM Cloud Platform discovery

The ServiceNow Discovery application finds IBM Cloud Platform components (Softlayer API v3 and v3.1 and Bluemix API v2) using the IBM Cloud Platform patterns. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

You can use the patterns on the ServiceNow AI Platform using London Patch 7, Madrid Patch 2, or later releases.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

MID Server

Use a dedicated MID Server for cloud discovery. The MID Server must have 16 Cores, 8 GB with 4 GB Java Memory Model (JVM) memory allocation.

Deploy and connect the dedicated MID Server with internet access.

Credentials

The cloud resources are in an IBM cloud. Create the IBM Credentials for accessing your IBM account.

1. Navigate to **Discovery** > **Credentials**.
2. Click **New**.
3. Click **IBM Credentials** and then fill in the form:

Credentials form

Field	Description
Name	Credential name. An example is `ibm-cred`.
Order	Order in which the platform tries this credential as it attempts to log on to devices. A smaller number indicates that the credential appears higher in the list. Establish the credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number, or none, the instance tries the credentials in a random order. The default value is 100.
Softlayer API User	Username for the Softlayer API. Use the Softlayer API Username defined in the Softlayer console.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Softlayer API Key	Key for the Softlayer API. Use the Softlayer API Key that is defined in the Softlayer console.
Bluemix API Key	Key for the Bluemix API. Use the Bluemix API Key defined in the IBM Cloud console.

4. Click **Submit**.

Create an IBM cloud service account

Use the credentials defined in the preceding procedure to create the IBM cloud service account.

1. Navigate to **Cloud Provisioning and Governance** > **Service Accounts**.
2. Click **New**.
3. On the form, fill in the fields:

Cloud Service Account form

Field	Description
Name	Credential name. An example is `ibm-snow`
Account Id	IBM account ID name. Enter the IBM Account ID.
Discovery credentials	Discovery credentials that you defined earlier.
Datacenter Type	Datacenter type. As an example, `IBM Datacenter [cmdb_ci_ibm_datacenter]`

4. Click **Submit**.

Create a Discovery schedule

To configure a Discovery schedule through Discovery Admin Workspace, see Create an IBM Discovery schedule in Discovery Admin Workspace.

To configure a serverless Discovery schedule:

1. On the ServiceNow AI Platform, navigate to **Discovery** > **Discovery Schedules**.
2. Click **New**.
3. On the form, fill in the fields:

Discovery Schedule form

Field	Description
Name	Name for this Discovery schedule. As an example, `Serverless Pattern`
Discover	Discover type. Select **Serverless**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
MID Server	Name of the MID Server. As an example, `Discovery_Server`
Active	Option for enabling this schedule for Discovery. Select this check box to enable Discovery.
Run	Specifies how frequently to run the Discovery schedule.

4. Right-click the header of the Discovery Schedule form and select **Save**.

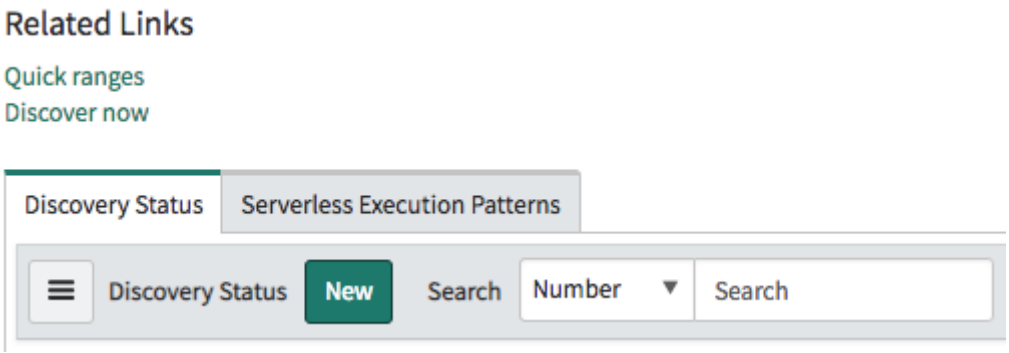

5. Click the **Serverless Execution Patterns** tab.
6. Click **New**.
7. On the form, fill in the fields:

Serverless Execution Pattern form

Field	Description
Name	Name for this Serverless Execution Pattern. As an example, `IBM1`
Pattern	Name of the pattern that you want to run. 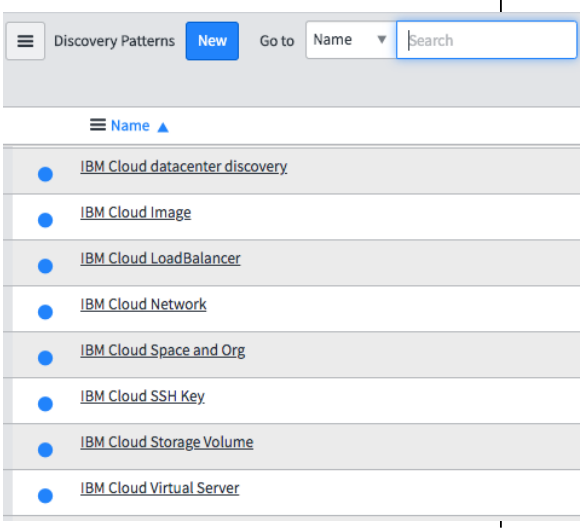

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	If you want to run multiple patterns, create multiple records.
Proxy Host	Fully qualified domain name of the machine on which you are installing the proxy server. Specify **Global**.
Active	Option for enabling this schedule for Discovery. Select this check box to enable Discovery.

8. Configure the parameters required by the IBM Cloud Platform pattern as follows:

Discovery Pattern Launcher parameters

Parameter	Description
cloud_cred_id	The sys_id of the IBM credentials specified into the configuration.
cloud_account_id	The IBM account ID.
cloud_datacenter_type	cmdb_ci_ibm_datacenter

9. **Update**.
10. To start Discovery, navigate to the Serverless Discovery definition and click **Discover Now**.

IBM flow - horizontal discovery

- To perform a discovery using a serverless schedule, the IBM Cloud Platform patterns that were selected run and discover IBM Cloud Platform resources.
- To perform a discovery using a service account, a full IBM Cloud Platform discovery schedule runs, discovering all supported IBM Cloud Platform resources.

Cloud custom operation

Cloud REST – Add Response to Context

Operation for Cloud REST queries that supports large-scale responses and better memory usage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Specify all the fields with the same values as in the Cloud REST Query custom operation, except to specify the table for the response, which is set in the **Response Variable Name** field.

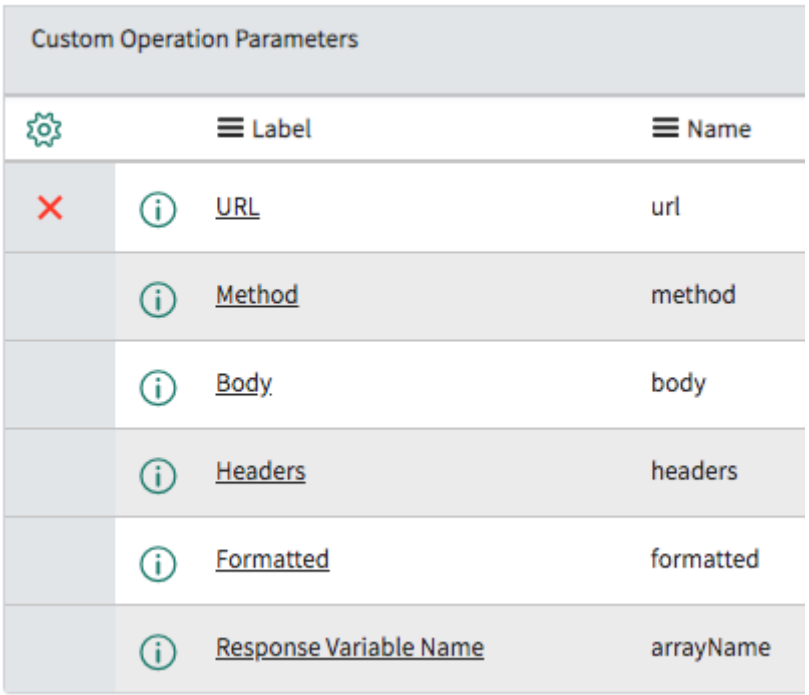

- Set **Define parsing** to None.

Methods to retrieve data for the IBM account

1. Navigate to **Pattern Designer** > **Custom Operations**.

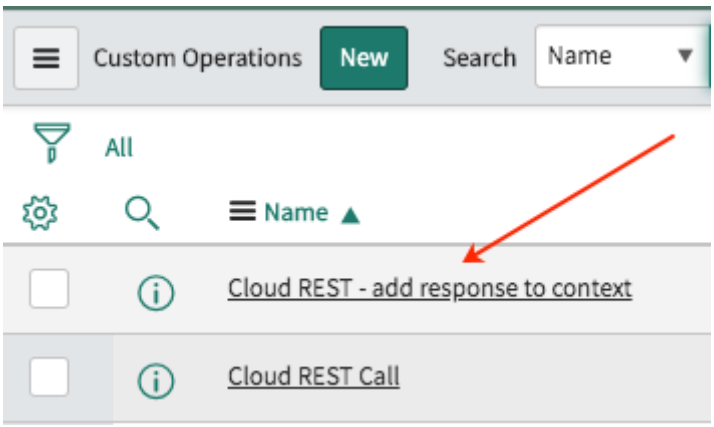

2. Click **Cloud REST – add response to context**.

Method: Get Call

API URLs used to retrieve the data:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Load Balancer *https://api.softlayer.com/rest<api_version>/SoftLayer_Network_LBaaS_LoadBalancer/getAllObjects?objectMask=mask[listeners[defaultPool],members,healthMonitors]*
- Network *https://api.softlayer.com/rest/<API_VERSION>/SoftLayer_Account/getNetworkVlans?objectMask=mask[primaryRouter.datacenter,name,id,vlanNumber,networkSpac*
- Network *https://api.softlayer.com/rest/<API_VERSION>/SoftLayer_Account/getSubnets?objectMask=mask[datacenter]*
- Network *https://api.softlayer.com/rest/<API_VERSION>/SoftLayer_Account/getSecurityGroups?objectMask=mask[rules]*
- Organizations and Spaces *https://api.<region>.bluemix.net/v2/organizations*
- *https://api.<region>.bluemix.net/v2/spaces*
- SSH Key *https://api.softlayer.com/rest/v3/SoftLayer_Account/SoftLayer_Account/getSshKeys*
- Storage *https://api.softlayer.com/rest/<api_version>/SoftLayer_Account/getNetworkStorage*
- Virtual Server *https://api.softlayer.com/rest/v3.1/SoftLayer_Account/getSecurityGroups*
- *http://api.softlayer.com/rest/v3.1/SoftLayer_Account/getSecurityGroups?objectMask=mask[id,name,networkComponentBindings[networkComponentId,net*
- *https://api.softlayer.com/rest/v3.1/SoftLayer_Account/getVirtualGuests?objectMask=mask[hostname,domain,powerState,fullyQualifiedDomainName,not*

EVAL functions

Use the following EVAL function to transform the Load Balancer API response into a full JSON format. The default response is formatted as "[<resource objects>]".

```
var cloud_lb_response = ${cloud_lb_response};
for (var i = 0; i < cloud_lb_response.size(); i++){
      cloud_lb_response.set(i, '{"result":[' +
 cloud_lb_response.get(i) + ']}');
}
CTX.setAttribute("cloud_lb_response",
 cloud_lb_response;
```

Use the following EVAL function to transform the Network API response into a full JSON format. The default response is formatted as "[<resource objects>]".

```
var network_response = ${network_response};
for (var i = 0; i < network_response.size(); i++){
      network_response.set(i, '{"result":[' +
 network_response.get(i) + ']}');
}
CTX.setAttribute("network_response", network_response);
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the following EVAL function to get the IBM region from the LB availability zone (datacenter). The same function is used for other resources inside the pattern.

```
var res = "";
res =
 IbmCloudDcToRegion.convertDctoRegion(${cloud_lb[].clou
d_lb_az_short_name});
```

Use the following EVAL function to select the required algorithm in the pattern.

```
var rtrn
var lb_algorithm =
 ${cloud_lb_pools[].lb_algorithm}.toString();
switch (lb_algorithm){
   case "ROUNDROBIN":
      lb_algorithm = "round-robin";
      break;
   case "WEIGHTED_RR":
      lb_algorithm = "round-robin";
      break;
   case "LEASTCONNECTION":
      lb_algorithm = "least-connections";
      break;
   default: break;
```

Use the following EVAL function to parse the unique region to the account that hosts the Organizations and the corresponding spaces.

```
var tableWithoutDuplicates = '';
tableWithoutDuplicates =
 DuplicateRemover.removeDuplicates(${account_region},["
region_name"]);
CTX.setAttribute("account_region_unique",
 tableWithoutDuplicates);
```

Use the following EVAL function to get information for all the organizations that are assigned to the IBM cloud account. Iteration for getting the unique region is performed.

```
var cloudcall = new CloudRestAPIClient();
var resourceTable =
 ${account_region_unique[*].region_name};
var tableSize = resourceTable.size();
var tmp_response = new Packages.java.util.ArrayList();
var all_organizations_response = new
 Packages.java.util.ArrayList();
for (var i = 0; i < tableSize ; i++){
   if  (!resourceTable.get(i)) continue;
   var url =
 "https://api."+resourceTable.get(i)+".bluemix.net/v2/o
rganizations";
   tmp_response =
 cloudcall.arrayListExecute(${service_account}, null,
 url, null, 'GET', null, 'true', CTX);

   for (var j = 0; j < tmp_response.size(); j++){
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      all_organizations_response.add('{"result":[' +
 tmp_response.get(j) + ']}');
   }
}
CTX.setAttribute("all_organizations_response",
 all_organizations_response);
```

Use the following EVAL function to get information for all the IBM cloud account assigned spaces, where the iteration for getting the unique region is performed.

```
var rtrn = '';
var cloudcall = new CloudRestAPIClient();
var resultList = new Packages.java.util.ArrayList();
var rest = '';
var resourceTable =
 ${account_region_unique[*].region_name};
var result = new Packages.java.util.ArrayList();
var all_organizations_response =
 ${all_organizations_response}
for (var i = 0; i < resourceTable.size() ; i++){
   var url =
 "https://api."+resourceTable.get(i)+".bluemix.net/v2/s
paces";
   result =
 cloudcall.arrayListExecute(${service_account}, null,
 url, null, 'GET', null, 'true', CTX);
   for (var z = 0; z <
 all_organizations_response.size(); z++){
      var jsonObj = JSON.parse(result.get(z));
      ms.log("length "+jsonObj.resources.length);
for(var j=0;j<jsonObj.resources.length;j++){
         var guid = jsonObj.resources[j].metadata.guid;
         var url = jsonObj.resources[j].metadata.url;
         var name = jsonObj.resources[j].entity.name;
         var organization_guid =
 jsonObj.resources[j].entity.organization_guid;
         var allow_ssh =
 jsonObj.resources[j].entity.allow_ssh ? "true" :
 "false";
         var row = new Packages.java.util.HashMap();
 Packages.com.snc.sw.log.DiscoLog.getLogger("IBM Cloud
 Space and Org").debug(">>>>>>>>>>>>>>|||||||||||||>>>>>>
 region alias=" + resourceTable.get(i));

 Packages.com.snc.sw.log.DiscoLog.getLogger("IBM Cloud
 Space and Org").debug(">>>>>>>>>>>>>>|||||||||||||>>>>>>
 region full name=" +
 IbmCloudDcToRegion.regionNamebyAlias(resourceTable.get
(i)));
         row.put("region",
 IbmCloudDcToRegion.regionNamebyAlias(resourceTable.get
(i)));
         row.put("space_guid",guid);
         row.put("space_url",url);
         row.put("space_name",name);

 row.put("space_organization_guid",organization_guid);
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
            row.put("space_allow_ssh",allow_ssh);
            row.put("rest", rest);
            resultList.add(row);
        }
    }
    CTX.setAttribute('spaces_table', resultList);
}
```

Use the following EVAL function to filter all the spaces by unique region.

```
var space_region = ${spaces_table[].region};

Packages.com.snc.sw.log.DiscoLog.getLogger("IBM Cloud
 Space and Org").debug("filtering space according to
 LDC. space_region=" + space_region);
var regions = ${cmdb_ci_ibm_datacenter[*].region};
Packages.com.snc.sw.log.DiscoLog.getLogger("IBM
 Cloud Space and Org").debug("input regions list=" +
 regions.toString());

var rtrn = "false";
if(regions.contains(space_region)) {
    rtrn = "true";
}
rtrn;
```

Use the following EVAL function to filter all the spaces by storage discovery.

```
var storage_response = ${storage_response};
for (var i = 0; i < storage_response.size(); i++){
      storage_response.set(i, '{"result":[' +
 storage_response.get(i) + ']}');
}
CTX.setAttribute("storage_response", storage_response);
```

Use the following EVAL function to transform the API response into full JSON format, as the untouched response is formatted as "[<resource objects>]".

```
var ssh_key = ${ssh_key_json};
for(var i = 0; i < ssh_key.size(); i++){
   ssh_key.set(i,'{"result\" : ' + ssh_key.get(i) +
 "}");
}
CTX.setAttribute("ssh_key_json", ssh_key);
```

Use the following EVAL function to populate the availability zone into the storage volume table.

```
var rtrn;
var resource_group =
 ${network_storage[].resource_name};
var regex = /.*-(\w{3}\d{2}).*/;
var az_name = regex.exec(resource_group);
if (az_name) {
   rtrn = az_name[1];
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
} else {
   rtrn;
}
```

Use the following EVAL function to populate the region into the storage volume table.

```
var res = "";
res =
 IbmCloudDcToRegion.convertDctoRegion(${network_storage
[].availability_zone});
```

Use the following EVAL function to transform the size of the storages to bytes from gigabytes.

```
var rtrn = '';

var size = ${network_storage[].size};

rtrn = (parseInt(size) * 1073741824) + '';
```

Use the following EVAL function to transform the size of the storages to bytes from gigabytes.

```
var cloudcall = new CloudRestAPIClient();
var resourceTable =
 ${network_storage[*].storage_object_id};
var tableSize = resourceTable.size();
var allResources = new Packages.java.util.ArrayList();
var response = new Packages.java.util.ArrayList();
for (var i = 0; i < tableSize ; i++){
   if  (!resourceTable.get(i)) continue;
   var url =
 "https://api.softlayer.com/rest/"+${api_version}+"/Sof
tLayer_Network_Storage/"+resourceTable.get(i)+"/getObje
ct?objectMask=mask[snapshots]";
   // the commented out code is to call the "Cloud Rest
 Call" API

   allResources =
 cloudcall.arrayListExecute(${service_account}, null,
 url, null, 'GET', null, 'true', CTX);
   for (var j = 0; j < allResources.size(); j++){
      response.add('{"result":[' + allResources.get(j)
 + ']}');
   }
}
CTX.setAttribute("all_snapshots_response", response);
```

Use the following EVAL closure that transforms the JSON that is received from the API to a format that the JSON parser can work with. The default JSON is saved in `$vm_json_array`, the new JSON is in `$vm_json`.

```
var vm_json = new Packages.java.util.ArrayList();
vm_json.addAll(${vm_json_array});
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
for (var i = 0 ; i < vm_json.size(); i++){
      vm_json.set(i, '{\"result\":[' + vm_json.get(i) +
 ']}');
}

CTX.setAttribute("vm_json", vm_json);
```

Use the following EVAL closure that returns 'on' if the power status is running. The value 'off' is returned if the power status is off.

```
var power = ${cmdb_ci_vm_instance[].state};
var rtrn = '';

if (power == 'Running'){
   rtrn = "on";
}
else {
   rtrn = "off"
}
```

Use the following EVAL closure that returns the region of the CI, based on provided availability zone name.

```
var az = ${cmdb_ci_vm_instance[].avail_zone};
var rtrn = '';
rtrn = IbmCloudDcToRegion.convertDctoRegion(az);
```

Use the following EVAL closure that builds a HashMap of HashSets, Connecting VMs with their network storages. The following key is used:

- P = page
- V = VM
- N = network storage

```
var rtrn = '';
var vmJson = ${vm_json_array};
var VMtoStorageHashMap = new
 Packages.java.util.HashMap();

for (var p = 0 ; p < vmJson.size(); p++){
   var page = JSON.parse(vmJson.get(p));
   for (var now_V in page){
      if (page[v].allowedNetworkStorage.length > 0){
         var vm_id = page[v].id + '';
         var storageSet = new
 Packages.java.util.HashSet();
         for (var n in page[v].allowedNetworkStorage){

 storageSet.add(page[v].allowedNetworkStorage[n].id +
 '');
         }
         VMtoStorageHashMap.put(vm_id,storageSet);
      }
   }
}
CTX.setAttribute("VMStorageCache", VMtoStorageHashMap);
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the following EVAL closure to remove duplicate entries from tables.

```
var tableWithoutDuplicates = '';
tableWithoutDuplicates =
 DuplicateRemover.removeDuplicates(${cmdb_ci_os_templat
e},["object_id"]);
CTX.setAttribute("cmdb_ci_os_template",
 tableWithoutDuplicates);
```

Use the following EVAL closure that builds a list of VM IDs that are related to the current security group ID.

```
var rtrn = '';
//get the Security Group ID of the current record in
 the table for comparison
var current_table_id =
 ${security_groups[].sec_group_id};
var current_table_id_string = current_table_id + '';

// get the list of instance's ID that will be then
 added to Security Group table
var needed_value_id_list =
 ${sec_group_network_components[*].vm_instance_id};
//get the list of Security Group IDs in the Network
 Component table, which will then be compared one by
 one to the current record of the security_groups table
var compare_id_list =
 ${sec_group_network_components[*].sec_group_id};

var beforeReturn = '';
var separator = '';
for (var i = 0; i <  compare_id_list.size(); i++)   {

   if   (current_table_id_string ==
 compare_id_list.get(i)){

      beforeReturn += separator +
 needed_value_id_list.get(i);
      separator = ',';
   }
}
rtrn = beforeReturn;
```

Use the following EVAL closure that builds a HashMap of HashSet between security groups and the VMs that are configured to use them.

```
rtrn = '';
var secGroupTable = ${cmdb_ci_endpoint_comp_security};
var secToVMHashSet = new Packages.java.util.HashMap();

// this loop worst case O(N*M) M = security groups, N =
 VMs
for(var i=0; i<secGroupTable.size(); i++){
   var securityTableId =
 secGroupTable.get(i).get('object_id');
   var vmListStr =
 secGroupTable.get(i).get('instance_list');
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
   var vmList = vmListStr.split(",");

   var vmSet = new Packages.java.util.HashSet();
   for (var j in vmList){
      vmSet.add(vmList[j]);
   }

   secToVMHashSet.put(securityTableId,vmSet);
}
CTX.setAttribute("securityVMCache",  secToVMHashSet);
```

Use the following EVAL closure that returns 'true' if the current VM is using the current security group ID.

```
var rtrn = '';
var vm_id = ${cmdb_ci_vm_instance[].object_id};
var vmHashSet =
 ${securityVMCache}.get(${cmdb_ci_endpoint_comp_securit
y[].object_id});
rtrn = vmHashSet.contains(vm_id);
```

Use the following EVAL closure that returns 'true' if the current VM is using the current storage volume.

```
var rtrn = '';
var vm_id = ${cmdb_ci_vm_instance[].object_id};
var vol_id = ${cmdb_ci_endpoint_iscsi[].object_id};
var vmHashSet = ${VMStorageCache}.get(vm_id);
if (vmHashSet){
   rtrn = vmHashSet.contains(vol_id);
}
else {
   rtrn = 'false'
}
```

Use the following EVAL function to transform the API response into full JSON format, as the untouched response is formatted as "[<resource objects>]".

```
var cloud_response = ${location_groups_response}
for (var i = 0; i < cloud_response.size(); i++){
    cloud_response.set(I, '{"result":[' +
 cloud_response.get(i) + ']}');
}
CTX.setAttribute("location_groups_response",
 cloud_response);
```

Parses the unique Location Groups in the "location_groups_az" table basdon the column "lg_id" and creates a new table "location_groups_az_unique" that contains the result.

```
var tableWithoutDuplicates= '';
tableWithoutDuplicates =
 DuplicateRemover.removeDuplicates(${location_groups_a
z},["lg_id"]);
CTX.setAttribute("location_groups_az_unique",
 tableWithoutDuplicates);
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Creates a list of Availability Zones per Location Group, that is stored in a column new column in the location_groups_az_uniquetable. The list records are divided by commas purposely, so that unique relations could be made in consecutive step.

```
var current_table_id=
 ${location_groups_az_unique[].lg_id};
var current_table_id_string = current_table_id+ '';
var needed_v alue_id_list =
 ${location_groups_az[*].lg_az_id};
var compare_id_list= ${location_groups_az[*].lg_id};
var beforeReturn= '';
var separator = ',';
for (var i= 0; i< compare_id_list.size(); i++)  {
    if (current_table_id_string ==
 compare_id_list.get(i)){
    beforeReturn += separator + needed_v
 alue_id_list.get(i) + separator;
  }
}
rtrn = beforeReturn;
```

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Data collected by Discovery during horizontal discovery

Data is collected using these patterns:

- IBM Cloud Load Balancer
- IBM Cloud Location Groups
- IBM Cloud Network
- IBM Cloud Organizations and Spaces
- IBM Cloud Resource Groups
- IBM Cloud SSH Key
- IBM Cloud Storage
- IBM Cloud Virtual Server

Data collected for IBM components:

IBM Cloud Load Balancer discovered resources

Field	Description
Main CI: cmdb_ci_cloud_load_balancer	
name	Load balancer name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
object_id	Load balancer unique identifier.
fqdn	Load balancer address (host name).
dns_name	Load balancer domain server name.
short_description	Load balancer description that you provide for easy identification.
operational_status	Load balancer status, one of either: • ONLINE • OFFLINE
location	Datacenter location.
state	Load balancer provisioning status.
Related CI: cmdb_ci_lb_service	
name	Pattern created name, composed of: • load balancer name • type of protocol • resource type listener
object_id	Load balancer unique identifier.
port	Port number of the listener.
front_end_port	Protocol port number of the listener.
listener_protocol	Protocol of the listener. One of the following protocols: • TCP • HTTP • HTTPS
operational_status	Provisioning status of the listener.
Related CI: cmdb_ci_lb_pool	
name	Pattern created name, composed of: • load balancer name • type of protocol • port • resource type “pool”
object_id	Load balancer pool unique identifier.
port	Protocol port of the backend.
pool_status	Load balancer pool provisioning status.
load_balancing_method	Load balancing method. One of the following algorithms:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• "ROUNDROBIN" • "WEIGHTED_RR" • "LEASTCONNECTION"
Related CI: cmdb_ci_lb_pool_member	
name	Pattern created name, composed of: • load balancer name • type of protocol • IP address of the Pool Member
object_id	Load balancer pool member unique identifier, composed of: • Member ID • Listener ID
ip_address	IP address of a load balancer member.
server_id	ID of the member in IBM cloud.
load_balancer	Name of the related Load Balancer.
Related CI: cmdb_ci_lb_health_service	
name	Pattern created name, composed of: • load balancer name • type of protocol • resource type: **Health Check** • resource ID
object_id	Unique identifier of the Health Service Check.
monitor_type	Health check type. One of the following: • TCP • HTTP • HTTPS
timeout_sec	Timeout, in seconds, to wait for health check responses.
check_interval_sec	Interval, in seconds, to perform health check.

IBM Cloud Location Groups discovered resources

Field	Description
Main CI: cmdb_ci_cloud_location_group	
name	A descriptive name used to identify a Location Group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
object_id	The ID of the Location Group record.
locationGroupType Name	The name of the Location Group Type.
short_description	A small note about a Location Group to use at your discretion.
locationGroup Type ID	The ID of the Location Group Type.

IBM Cloud Network discovered resources

Field	Description
Main CI: cmdb_ci_network	
name	VLAN number.
object_id	Network unique identifier.
access_port_vlan	VLAN number.
is_external	Network space value. The value is either: • true if the network_space is “PUBLIC” • false if the network_space is “PRIVATE”.
state	State has the value "available".
Related CI: cmdb_ci_cloud_subnet	
name	Name, composed of: “network_identifier/cidr”
object_id	Subnet unique identifier.
subnet_mask	Netmask of the subnet.
gateway	Gateway IP address.
cidr	Method for allocating IP addresses and IP routing: “network_identifier/cidr”.
category	Subnet type.
broadcast_address	Subnet broadcast address.
available_ip_count	Number of available IP addresses in the subnet.
Related CI: cmdb_ci_compute_security_group	
name	Security Group name.
object_id	Security Group unique identifier.
short_description	Security Group description that you provide for easy identification.
Related CI: cmdb_ci_compute_security_group_rule	
name	Field with the same value as the Security Group Rule object_id field.
object_id	Security Group Rule unique identifier.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
created	Date the CI was created.
direction	Direction of the information. One of the following: • inbound • outbound
Ethertype	Ethertype. One of the following: • IPv4 • IPv6
port_range_min	The first port of the range that the rule applies to.
port_range_max	The last port of the range that the rule applies to.
protocol	Protocol that the rule is applied to.
remote_group_id	Security Group ID for the Source/Destination security group. It means that this rule applies on all IP addresses that were configured in this security group.

IBM Cloud Organizations and Space discovered resources

Field	Description
Main CI: cmdb_ci_cloud_org	
name	Name of the cloud organization.
guid	Unique identifier of the organization.
url	Unique URL of the organization.
billing_enabled	Billing method in use by the organization. One of either: • Enabled / (TRUE) • Disabled / (FALSE)
quota_definition_url	Resources that are available to an organization.
spaces_url	Unique URL of the space assigned to the organization.
operational_status	Operational status of the organization.
Related CI: cmdb_ci_cloud_space	
name	Unique name of the space.
object_id	Space unique identifier (guid).
alllow_ssh	Access to the application over SSH is either allowed or not allowed. Boolean, one of either: • TRUE • FALSE
URL	URL of the space from the metadata description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM Cloud Resource Groups discovered resources

Field	Description
Main CI: cmdb_ci_resource_group	
name	Name of the resource group.
object_id	Unique identifier of the resource group.
quota_id	Value identifying the quota ID associated with the resource group.

IBM Cloud SSH Key discovered resources

Field	Description
Main CI: cmdb_ci_cloud_key_pair	
name	Name to identify an SSH key.
object_id	ID of the SSH key record.
key	SSH Key record.
finger_print	Short sequence of bytes used to authenticate or lookup a longer SSH Key. This value is automatically generated when adding or modifying the SSH Key. Read only. Changes made are ignored.
created _date	Date when the SSH Key was added. Read only. Changes made are ignored.
modify_date	Date when the SSH Key was modified. Read only. Changes made are ignored.
short_description	SSH description that you provide for easy identification.
Related CI: cmdb_ci_ibm_datacenter	
name	Name of the datacenter.
region	Region that the datacenter is in.
discovery_status	Status information about the IBM datacenter.
sys_class_name	Class of the datacenter
Related CI: cmdb_ci_cloud_service_account	
name	Name of the cloud service account.
account_id	Account ID for this cloud service account.
discovery_credentials	Credentials for this cloud service account.
datacenter_url	URL used to populate Logical Data Centers (LDCs).
datacenter_type	Type of datacenter. As an example, IBM Datacenter [cmdb_ci_ibm_datacenter]
datacenter_discovery_status	Status of last execution of Discover Datacenters.

IBM Cloud Storage discovered resources

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Main CI: cmdb_ci_storage_volume	
name	Username used to: • access a non-EVault Storage volume • register the EVault server agent with the vault backup system
object_id	Storage account unique identifier.
size	Capacity of the storage account, measured in gigabytes.
size_bytes	Capacity of the storage account, measured in bytes.
label	Name of the storage network resource.
storage_type	Description of the storage type. As an example, NAS.
fqdn	IP address of a storage resource or fully qualified domain name.
device	Unique identification number of the device associated with a Storage volume.
Related CI: cmdb_ci_storage_vol_snapshot	
name	Name of a storage network resource.
object_id	Storage unique identifier.
capacity	Capacity of the snapshot volume in gigabytes.
fqdn	IP address of a storage resource or fully qualified domain name.
location	Unique availability zone where the storage volume is located.
Related CI: cmdb_ci_ibm_datacenter	
name	Name of the datacenter.
region	Region that the datacenter is in.
discovery_status	Status information about the IBM datacenter.
sys_class_name	Class of the datacenter
Related CI: cmdb_ci_cloud_service_account	
name	Name of the cloud service account.
account_id	Account ID for this cloud service account.
discovery_credentials	Credentials for this cloud service account.
datacenter_url	URL used to populate LDCs.
datacenter_type	Type of datacenter. As an example, IBM Datacenter [cmdb_ci_ibm_datacenter]
datacenter_discovery_status	Status of last execution of Discover Datacenters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Related CI: cmdb_ci_availability_zone	
name	Name of the availability zone (AZ).
object_id	Unique identifier for this resource.
short_description	Availability zone description that you provide.

IBM Cloud Virtual Server discovered resources

Field	Description
Main CI: cmdb_ci_vm_instance	
name	Name of the VM.
object_id	Unique identifier of the VM.
dns_domain	VM domain server name.
fqdn	VM IP address.
short_description	Description in a note that is attached to the VM on the IBM console.
cpus	Number of CPU cores assigned to this VM.
memory	Memory size assigned to this VM.
disks	Number of disks assigned to this VM.
state	The running state of the VM. • On – returned as 'Running' from the API. • Off – returned as 'Halted' from the API.
placement_group_id	Placement group ID.
nics	Number of NICs.
Related CI: cmdb_ci_nic	
name	Name, composed of: • load balancer name • Port
object_id	Unique identifier of the NIC.
ip_address	IP that is configured on the NIC.
mac_address	NIC MAC details.
netmask	Mask used to divide the IP address into subnets and to specify the available hosts in the network.
status	State of the NIC.
Related CI: cmdb_ci_endpoint_vnic	
name	Name for the VNIC endpoint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
object_id	Unique identifier of the VNIC endpoint.
host	Host identifier.
Related CI: cmdb_ci_storage_volume	
name	Username used to: • access a non-EVault Storage volume • register the EVault server agent with the vault backup system
object_id	Storage account unique identifier.
size	Capacity of the storage account, measured in gigabytes.
size_bytes	Capacity of the storage account, measured in bytes.
label	Name of the storage network resource.
storage_type	Description of the storage type. As an example, NAS.
fqdn	IP address of a storage resource or fully qualified domain name.
device	Unique identification number of the guest associated with a Storage volume.
Related CI: cmdb_ci_endpoint_nfs	
name	Name for the NFS endpoint.
host	Host identifier.
protocol	Protocol that was applied.
type	NFS endpoint type.
created	Date that the NFS endpoint was created.
updated	Date that the NFS endpoint was updated.
discovery_source	Source of the discovery. As an example, ServiceWatch
host_name	Host name of the NFS endpoint
Related CI: cmdb_ci_endpoint_iscsi	
host	Host identifier for the ISCSI endpoint.
port	Port to determine how GCP iscsi contact to the VMs.
protocol	Protocol that was applied.
type	ISCSI endpoint type
created	Date that the ISCSI endpoint was created.
updated	Date that the ISCSI endpoint was updated.
discovery_source	Source of the discovery. As an example, ImportSet

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
host_name	Host name of the ISCSI endpoint
Related CI: cmdb_ci_endpoint_block	
name	Name for the block endpoint.
object_id	Unique identifier of the block endpoint.
host	Host identifier.
Related CI: cmdb_ci_os_template	
name	Name of the image. For example, 'CentOS 7.0-64 Minimal for VSI'.
object_id	Unique identifier of the image.
guest_os	OS used by the image. As an example, Linux.
root_device_type	Type of device used by root.
image_source	Source of the image.
image_type	Image type.
infuse_key	Key used to infuse the image.
update_host_name	Update host name for the image.
credentials	Credentials for the image.
version	OS version. For example, '7.0-64 Minimal for VSI'.
Related CI: cmdb_ci_cloud_key_pair	
name	Name to identify a key (key label).
object_id	ID of the key record.
key	Key record.
finger_print	Short sequence of bytes used to authenticate or look up a Key. This value is automatically generated when adding or modifying the Key. Read only. Changes made are ignored.
created _date	Date when the Key was added. Read only. Changes made are ignored.
modify_date	Date when the Key was modified. Read only. Changes made are ignored.
short_description	Description that you provide for easy identification.
Related CI: cmdb_ci_compute_security_group	
name	Security Group name.
object_id	Security Group unique identifier.
short_description	Security Group description that you provide for easy identification.
Related CI: cmdb_ci_endpoint_comp_security	
host	Host identifier of the compute security endpoint.
port	Unique identifier of the compute security endpoint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
protocol	Protocol that was applied.
type	Type of compute security endpoint.
created	Date when the compute security endpoint was added. Read only. Changes made are ignored.
updated	Date when the compute security endpoint was modified. Read only. Changes made are ignored.
discovery_source	Source of the discovery. As an example, Manual Entry.
host_name	Host name of the compute security endpoint.
Related CI: cmdb_ci_ibm_datacenter	
name	Name of the datacenter.
region	Region that the datacenter is in.
discovery_status	Status information about the IBM datacenter.
sys_class_name	Class of the datacenter
Related CI: cmdb_ci_availability_zone	
name	Name of the availability zone (AZ).
object_id	Unique identifier for this resource.
short_description	Availability zone description that you provide.
Related CI: cmdb_ci_cloud_service_account	
name	Name of the cloud service account.
account_id	Account ID for this cloud service account.
discovery_credentials	Credentials for this cloud service account.
datacenter_url	URL used to populate LDCs.
datacenter_type	Type of datacenter. As an example, IBM Datacenter [cmdb_ci_ibm_datacenter]
datacenter_discovery_status	Status of last execution of Discover Datacenters.

IBM Cloud Virtual Server – Related Items view

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

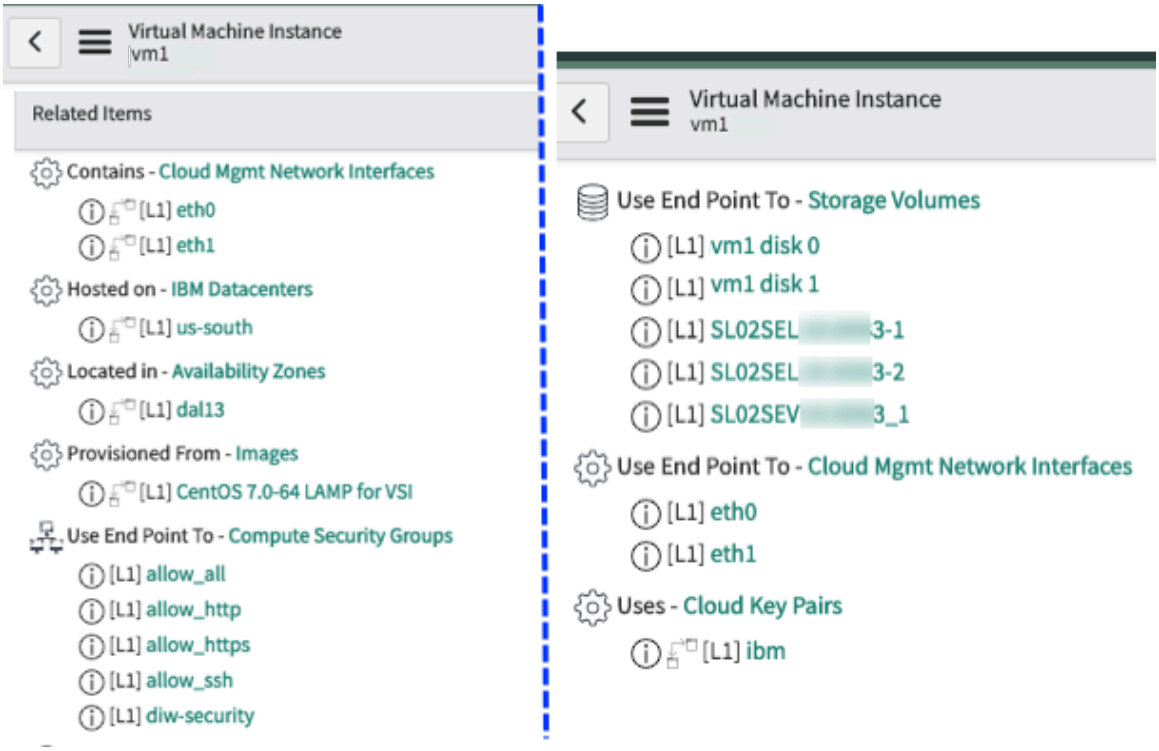

CI Relationships

These relationships are created to support IBM Cloud Platform discovery.

Relationships for IBM Cloud Load Balancer discovered resources

CI	Relationship	CI
Related CI: cmdb_ci_ibm_datacenter		
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
Related CI: cmdb_ci_cloud_load_balancer		
cmdb_ci_cloud_load_balancer	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_cloud_load_balancer	Contains::Contained by	cmdb_ci_lb_health_service
Related CI: cmdb_ci_availability_zone		
cmdb_ci_availability_zone	Hosted on::Hosts	cmdb_ci_cloud_load_balancer
cmdb_ci_availability_zone	Contains::Contained by	cmdb_ci_lb_service
Related CI: cmdb_ci_lb_service		
cmdb_ci_lb_service	Hosted on::Hosts	cmdb_ci_cloud_load_balancer
cmdb_ci_lb_service	Contains::Contained by	cmdb_ci_lb_pool
Related CI: cmdb_ci_lb_pool		
cmdb_ci_lb_pool	Hosted on::Hosts	cmdb_ci_cloud_load_balancer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
cmdb_ci_lb_pool	Owns::Owned by	cmdb_ci_lb_pool_member
Related CI: cmdb_ci_lb_health_service		
cmdb_ci_lb_health_service	Contains::Contained by	cmdb_ci_cloud_load_balancer

Relationships for IBM Cloud Location Groups discovered resources

CI	Relationship	CI
Related CI: cmdb_ci_ibm_datacenter		
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
Related CI: cmdb_ci_cloud_location_group		
cmdb_ci_cloud_location_group	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_location_group	Contains::Contained by	cmdb_ci_availability_zone

Relationships for IBM Cloud Network discovered resources

CI	Relationship	CI
Related CI: cmdb_ci_ibm_datacenter		
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
Related CI: cmdb_ci_network		
cmdb_ci_network	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_network	Contains::Contained by	cmdb_ci_cloud_subnet
Related CI: cmdb_ci_compute_security_group		
cmdb_ci_compute_security_group	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_compute_security_group	Contains::Contained by	cmdb_ci_compute_security_group_rule

Relationships for IBM Cloud Resource Groups discovered resources

CI	Relationship	CI
Related CI: cmdb_ci_resource_group		
cmdb_ci_resource_group	Hosted on::Hosts	cmdb_ci_cloud_service_account

Relationships for IBM Cloud SSH Key discovered resources

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Related CI: cmdb_ci_ibm_datacenter		
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
Related CI: cmdb_ci_cloud_key_pair		
cmdb_ci_cloud_key_pair	Hosted on::Hosts	cmdb_ci_cloud_service_account

Relationships for IBM Cloud Storage discovered resources

CI	Relationship	CI
Related CI: cmdb_ci_ibm_datacenter		
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
Related CI: cmdb_ci_storage_volume		
cmdb_ci_storage_volume	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_storage_volume	Provisioned::Provisioned From	cmdb_ci_storage_vol_snapshot
Related CI: cmdb_ci_storage_vol_snapshot		
cmdb_ci_storage_vol_snapshot	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_storage_vol_snapshot	Located in Zone::Zone contains	cmdb_ci_availability_zone
Related CI: cmdb_ci_availability_zone		
cmdb_ci_availability_zone	Contains::Contained by	cmdb_ci_storage_volume

Relationships for IBM Cloud Virtual Server discovered resources

CI	Relationship	CI
Related CI: cmdb_ci_ibm_datacenter		
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_vm_instance
Related CI: cmdb_ci_vm_instance		
cmdb_ci_vm_instance	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_vm_instance	Located in:Houses	cmdb_ci_availability_zone
cmdb_ci_vm_instance	Uses:Used By	cmdb_ci_cloud_key_pair
cmdb_ci_vm_instance	Contains:Contained By	cmdb_ci_nic
cmdb_ci_vm_instance	endpoint	cmdb_ci_nic

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Related CI: cmdb_ci_availability_zone		
cmdb_ci_availability_zone	Hosted on::Hosts	cmdb_ci_vm_instance
Related CI: cmdb_ci_nic		
cmdb_ci_nic	Hosted on::Hosts	cmdb_ci_vm_instance
Related CI: cmdb_ci_os_template		
cmdb_ci_os_template	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_os_template	Provisioned::Provisioned by	cmdb_ci_vm_instance
Related CI: cmdb_ci_compute_security_group		
cmdb_ci_compute_security_group	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_compute_security_group	Implement::Implemented from	cmdb_ci_vm_instance
Related CI: cmdb_ci_storage_volume		
cmdb_ci_storage_volume	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_storage_volume	Implement::Implemented from	cmdb_ci_vm_instance
Related CI: cmdb_ci_cloud_key_pair		
cmdb_ci_cloud_key_pair	Hosted On:Hosts	cmdb_ci_cloud_service_account

Relationships discovered for the virtual machine (VM) instance

CI	Relationship	CI
Servers [cmdb_ci_server]	Virtualized by:: Virtualizes	Virtual Machine Instance [cmdb_ci_instance]

Data collected by Service Mapping during tag-based discovery

Service Mapping uses tag-based discovery to create service instance maps including the Cloud components. The Service Mapping application comes with the following preconfigured CI relationships used for tag-based discovery. These CI relationships are available from the 1.0.68 release on the ServiceNow Store.

CI	Relationship	CI
Configuration Item [cmdb_ci]	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]
Logical Datacenter [cmdb_ci_logical_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Troubleshooting

If the mapping process does not proceed as you expected, follow the following suggestions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Symptom	Cause	Solution
Discovery fails. The discovery message contains the information about an error caused by the REST timeout.	There are many CIs sending the REST call response in the deployment. The MID Server cannot process the REST call response without exceeding the time limit controlled by the *mid.sa.cloud.request_timeout* parameter.	By default, the *mid.sa.cloud.request_timeout* parameter is set to 30000 milliseconds. Increase the value of this parameter on the relevant MID Server and run discovery again. **Note:** If the **Configuration Parameters** related list for the relevant MID Server does not show this parameter, you may need to add it.
Pattern Designer fails during a debug session. The Pattern Designer message contains information about an error caused by a timeout.	The Pattern Designer fails because of a timeout during pattern debugging (and not during discovery).	By default, the *sa.debugger.max_timeout* parameter is set to 240 seconds. Increase the value of this parameter on the relevant MID Server.

Microsoft Azure Cloud discovery using patterns

Discovery uses multiple patterns to discover components of the Microsoft Azure Cloud deployment during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date.

- Discovery and Service Mapping Patterns
- CMDB CI Class Models
- Visibility Content

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Azure Availability Set

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Wait for the **Clean-Up job for Availability zone to clear availability set record** schedule job to delete all the pre-populated Configuration Items (CIs) in the **cmdb_ci_azure_availability_set** table.

Azure Availability Zone

To run a discovery with Azure Availability Zone, register the subscription ID to the **AvailabilityZonePeering** feature with AZ CLI using `az feature register -n AvailabilityZonePeering --namespace Microsoft.Resources` to use the Check Zone Peering API. Check the status with `az feature show -n AvailabilityZonePeering --namespace Microsoft.Resources` before running discovery.

Set up Azure service accounts

Enable Cloud Discovery to access your Azure environment. For more information, see Set up Azure service accounts.

Create an Azure cloud discovery schedule

For more information, see Create an Azure Discovery schedule in Discovery Admin Workspace.

(Optional) Discover datacenters only for new subscriptions

Starting with Zurich Patch 2, you can discover datacenters only for new subscriptions added since the last discovery. For more information, see Discover datacenters only for new cloud accounts.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Azure resources discovery by datacenters

Azure has multiple datacenters around the world, but resources like load balancers and virtual machines are typically deployed in only some of them. The **Azure Datacenter Discovery** pattern executes before all other Azure patterns to identify the datacenters that have resources related to your service account ("active") and the datacenters that don't have your resources ("passive"). This model improves the performance of the Azure discovery. This execution model is more efficient than the previous one, in which all datacenters were discovered regardless of having relevant resources in them.

After identifying the "active" and "passive" datacenters, the Discovery schedule continues to execute all Azure patterns only for the "active" datacenters, to discover your Azure cloud resources. The "passive" datacenters are ignored while running the schedule.

You might notice differences in Azure discovery log, in discovery time and in the CMDB, depending on the service account and MID Server property settings.

Datacenters that have already been discovered before the upgrade to Discovery and Service Mapping Patterns version 1.15.0, remain in the **Azure Datacenters** table. However,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

the discovery runtime behavior is now determined by the value of the MID Server property *mid.cloud.discovery.sonar.discover_all_azure_datacenters*. The property is set to **false** by default, to limit the discovery execution to the "active" datacenters, rather than all datacenters. You can discover all datacenters for a service account, including "passive" ones, by setting the property to **true**. For more information, see: Create a MID Server property.

If the MID Server property is set to **false**, the Azure Datacenters table shows only active datacenters. All other behaviors remain unchanged from previous Discovery and Service Mapping Patterns versions.

Differences in datacenter discovery by Discovery and Service Mapping Patterns version

Discovery and Service Mapping Patterns version	MID Server property setting	Discovered datacenters	Datacenters contained in Azure Datacenters table	Datacenters displayed in discovery log
Discovery and Service Mapping Patterns starting with version 1.15.0	False (default)	Active only	Active only	Active only
Discovery and Service Mapping Patterns starting with version 1.15.0	True	All datacenters	All datacenters	All datacenters
Discovery and Service Mapping Patterns before version 1.15.0	False (default)	Active only	All datacenters	Active only
Discovery and Service Mapping Patterns before version 1.15.0	True	All datacenters	All datacenters	All datacenters

For management groups, Azure Cloud Discovery discovers all Azure datacenters.

Starting with Discovery and Service Mapping Patterns version 1.29.0, the **Refresh Datacenters** flow displays all regions, not just active ones. You don't need to create another schedule when a resource is added or a datacenter switches from passive to active.

Azure Hardware Type discovery

Hardware Type discovery has undergone three model changes in recent years. The 1.15.0 model triggers the **Hardware Type** pattern and the **Virtual Machine** pattern after the **Azure Datacenter Discovery** pattern. Starting Discovery and Service Mapping Patterns plugin version 1.15.0, Cloud Discovery identifies which Hardware Type model is used, and launches only one of the two patterns: **Hardware Type (LP)** or **Cloud Hardware Type (LP)**.

Migration to the Azure Hardware Type discovery model- Use cases

Discovery and Service Mapping Patterns version	Hardware Type Migration status	Which pattern executes	Discovery result
Prior to 1.0.75	Haven't migrated to the new model	**Hardware Type (LP)** pattern	The CI type created: [cmdb_ci_compute_template]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Migration to the Azure Hardware Type discovery model- Use cases (continued)

Discovery and Service Mapping Patterns version	Hardware Type Migration status	Which pattern executes	Discovery result
Discovery and Service Mapping Patterns version 1.0.75	The migration to the new model is done by migration script. See KB0955939	**Hardware Type (LP)** pattern	The CI type created: [cmdb_ci_cloud_hardware_type]
Discovery and Service Mapping Patterns version 1.6.0	The Hardware Type new model is provided OOB, enabled with the system property: *sn_itom_pattern.use a single hardware type for cloud datacenters*. For more information, seeKB1285337.	According to KB1285337 Flow Diagram	The CI type created: According to KB1285337
Discovery and Service Mapping Patterns 1.15.0	The Hardware Type new model is provided OOB enabled with the system property: *sn_itom_pattern.use a single hardware type for cloud datacenters*. For more information, seeKB1285337.	The flow is as described in KB1285337. However, only one pattern executes. The pattern that used to gracefully terminate doesn't execute.	Either **Hardware Type (LP)** pattern or **Cloud Hardware Type (LP)** pattern executes.

Tag information collected by Discovery during horizontal discovery

When running the patterns, tag information is collected to populate the cmdb_key_value table. Each tag is related to a CI that was discovered during the discovery. Tag discovery is done in the extension section of each pattern.

Data collected by Service Mapping during tag-based discovery

Service Mapping uses tag-based discovery to create service instance maps including the Cloud components. The Service Mapping application comes with the following preconfigured CI relationships used for tag-based discovery. These CI relationships are available from the 1.0.68 release on the ServiceNow Store.

CI	Relationship	CI
Configuration Item [cmdb_ci]	Hosted on::Hosts	Logical Datacenter [cmdb_ci_logical_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Logical Datacenter [cmdb_ci_logical_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Related topics

Kubernetes discovery using patterns

Azure Key Vault certificate discovery

Microsoft Foundry (classic) pattern-based discovery

Azure App Configuration store pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

The Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - App Configuration Configuration Store - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure App Configuration - Configuration Store [cmdb_azure_app_configuration_configuration_store]

Field	Description
Authentication Mode [authentication_mode]	The method or strategy used for authenticating access.
Disable Local Auth [disable_local_auth]	Indicates whether local authentication methods are disabled.
Enable Purge Protection [enable_purge_protection]	Indicates if purge protection is enabled for the resource.
Endpoint [endpoint]	The URL endpoint through which the service is accessed.
Key Vault Properties [key_vault_properties]	Properties related to the configuration of Key Vault services.
Location [location]	The Azure region where the resource is deployed.
Object Id [object_id]	The unique identifier of the resource.
Private Link Delegation [private_link_delegation]	Configurations related to the delegation of private links.
Provisioning State [provisioning_state]	The current state of the resource provisioning process.
Resource Group [resource_group]	Name of the resource group.
SKU Name [sku_name]	The stock keeping unit (SKU) representing the pricing tier of the resource.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The unique identifier of the Azure Active Directory tenant.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

The Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - App Configuration Configuration Store - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The Azure region where the resource is deployed.
Name [name]	The name assigned to the resource.
Object ID [object_id]	The unique identifier of the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.appconfiguration/configurationstores**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure App Configuration - Configuration Store [cmdb_azure_app_configuration_configuration_store]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure App Service App Service Plan pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populate data in the non-CMDB table when running the Azure - App Service App Service Plan - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure App Service - App Service Plan [cmdb_azure_app_service_app_service_plan]

Field	Description
Kind [kind]	The category or specialization of the resource, often specific to certain Azure services.
Location [location]	The Azure region where the resource is hosted.
Object Id [object_id]	The unique identifier for the resource.
Provisioning State [provisioning_state]	The current provision status of the resource.
Resource Group [resource_group]	Name of the resource group.
Status [status]	The operational status or health of the resource.
Subscription Id [subscription_id]	The subscription ID.
Tenant Id [tenant_id]	The identifier of the Azure Active Directory tenant associated with the resource.
Tier [tier]	The performance level or pricing tier assigned to the resource.
Configuration Item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - App Service App Service Plan - Extended Inventory(LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud App Server [cmdb_ci_cloud_appserver]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The Azure region where the resource is hosted.
Name [name]	The name assigned to the resource.
Object ID [object_id]	The unique identifier for the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **microsoft.web/serverfarms**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud App Server [cmdb_ci_cloud_appserver]
Cloud App Server [cmdb_ci_cloud_appserver]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure App Service - App Service Plan [cmdb_azure_app_service_app_service_plan]	References	Cloud App Server [cmdb_ci_cloud_appserver]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Microsoft Azure Application Gateway discovery with Patterns

The ServiceNow Discovery application uses the Azure Application Gateway (LP) pattern for discovering this product, while the Service Mapping application discovers Application Gateway using the Azure Application Gateway TD (LBS) pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Azure Principal credentials

On the ServiceNow platform, create a service account for cloud discovery. Set the account Id to the Azure Service Principal as it appears at the Microsoft Azure Console.

Privileges

On the Microsoft Azure Console, provide a user a user with the reader role.

Data collected by Discovery during horizontal discovery

Table and field	Additional information
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	
name	The name of the Azure Application Gateway.
Fqdn	The fully qualified domain name of the Azure Application Gateway.
Object_id	The ID of the gateway.
Install_status	The install status of the of the Azure Application Gateway.
state	The state of the Azure Application Gateway.
Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]	
name	The IP address of the Azure Application Gateway.
Object_id	The ID of the Azure Application Gateway.
ipaddress _type	The type of the IP address, which can be private or public.
ip_address	The IP address of the Azure Application Gateway.
state	The state of the IP address of the Azure Application Gateway.
operational_state	The operational state of the IP address of the Azure Application Gateway.
install_status	The install status of the IP address of the Azure Application Gateway.
Load Balancer Service [cmdb_ci_lb_service]	
serial number	The ID of the Azure Application Gateway service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Additional information
port	The TCP port that the Azure Application Gateway service listens to.
name	The FQDN with the port number of the Azure Application Gateway service, following this syntax: <fqdn>:<port>.
ip_address	The IP address of the Azure Application Gateway on which the service runs.
object_id	The ID of the Azure Application Gateway service.
operational_state	The operational state of the Azure Application Gateway service.

CI relationships

These relationships are created to support Azure Application Gateway discovery:

CI	Relationship	CI
Load Balancer Service [cmdb_ci_lb_service]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Owns::Owned by	Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Connections discovered by Service Mapping during the top-down discovery

Service Mapping performs the top-down discovery of the Azure Application Gateway in the context of application services. Service Mapping discovers the outgoing http connections from Azure Application Gateway service to pool members, such as virtual machines.

Related topics

Define a Cloud REST Query for cloud computing devices

Azure Application Insight Component pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populate data in the non-CMDB table when running the Azure - Application Insight Component - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Application Insight - Component [cmdb_azure_application_insight_component]

Field	Description
App ID [app_id]	The identifier for the application linked to this resource.
Application Type [application_type]	The type of application to which the resource is related.
Connection String [connection_string]	The string used to connect to the resource endpoint.
Instrumentation Key [instrumentation_key]	The key used for application insights to log telemetry data.
Kind [kind]	The specific subtype or variation of the resource.
Location [location]	The geographic region where the resource is deployed.
Object Id [object_id]	The unique identifier for the resource.
Provisioning State [provisioning_state]	The current provision status of the resource.
Resource Group [resource_group]	Name of the resource group.
Retention In Days [retention_in_days]	The number of days that logs or data are retained.
Subscription ID [subscription_id]	The subscription ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Application Insight - Component [cmdb_azure_application_insight_component] (continued)

Field	Description
Tenant ID [tenant_id]	The directory tenant to which the resource belongs.
Version [version]	The version number of the resource or service.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Application Insight Component - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographic region where the resource is deployed.
Name [name]	The name assigned to the resource.
Object ID [object_id]	The unique identifier for the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.insights/components**.
Description [short_description]	Description of the resource or service.
Version [version]	The version number of the resource or service.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Application Insight - Component [cmdb_azure_application_insight_component]	References	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Application Insight Data Collection Rule pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Application Insight Data Collection Rule - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Application Insight Data Collection Rule [cmdb_azure_application_insight_data_collection_rule]

Field	Description
Object Id [object_id]	The unique resource identifier for the Azure resource.
Kind [kind]	The specific resource kind or stock-keeping unit (SKU), indicating variation within the type.
DC Location [location]	The Azure region where the resource is deployed.
Resource Group [resource_group]	The name of the resource group containing the resource.
Subscription Id [subscription_id]	The unique identifier (GUID) for the Azure subscription.
Tenant Id [tenant_id]	The Azure Active Directory tenant identifier associated with the resource.
Provisioning State [provisioning_state]	The current provisioning or deployment status of the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Application Insight Data Collection Rule - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The unique resource identifier for the Azure resource.
Name [name]	The name assigned to the resource.
Location [location]	The Azure region where the resource is deployed.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.insights/datacollectionrules**.

CI relationships

The Azure - Application Insight Data Collection Rule - Extended Inventory(LP) pattern creates these relationships to support Azure Application Insight Data Collection Rule discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Application Insight Data Collection Rule [cmdb_azure_application_insight_data_collection_rule]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Application LB pattern-based discovery

Discovery and Service Mapping Patterns finds Azure application load balancers on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Application LB (LP) pattern.

Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Field	Description
Name [name]	The Name or ID if no Name is specified for the load balancer.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Location [location]	The path of the load balancer.
State [state]	The current state of the load balancer.
IP Address [ip_address]	IP address of the load balancer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Load Balancer [cmdb_ci_cloud_load_balancer] (continued)

Field	Description
Fully qualified domain name [fqdn]	The fully qualified domain name of the load balancer.
DNS Name [dns_name]	The DNS name of the load balancer.
Comments [comments]	Identifier for internal usage (deletion strategy).
Install Status [install_status]	Install status of the resource. Default value is Installed.

Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
IP Address [ip_address]	IP address of the load balancer.
IPAddress Type [ipaddress_type]	The type of the IP address. Possible values are private or public.
Fully qualified domain name [fqdn]	The fully qualified domain name of the load balancer.
Comments [comments]	Identifier for internal usage (deletion strategy).
Install Status [install_status]	Install status of the resource. Default value is Installed.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Name [name]	The Name or ID if no Name is specified for the load balancer pool.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Location [location]	Path of the load balancer pool.
Comments [comments]	Identifier for internal usage (deletion strategy).
Install Status [install_status]	Install status of the resource. Default value is Installed.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Name [name]	The Name or ID if no Name is specified for the load balancer service.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Location [location]	Path of the load balancer service.
Port [port]	The TCP port that the load balancer service listens to.
Comments [comments]	Identifier for internal usage (deletion strategy).
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Azure - Application LB (LP) pattern creates the following relationships and references to support Azure Application LB discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Owns::Owned by	Cloud LB IPAddress [cmdb_ci_cloud_lb_ipaddress]
Load Balancer Pool [cmdb_ci_lb_pool]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Load Balancer Service [cmdb_ci_lb_service]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Cloud Subnet [cmdb_ci_cloud_subnet]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Azure Tag discovery

The Azure - Application LB (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Load Balancer [cmdb_ci_cloud_load_balancer] table.

Azure Application Security Group pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Application Security Group - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Application Security Group [cmdb_azure_application_security_group_application_security_group]

Field	Description
Location [location]	The geographical region where the resource is deployed.
Object Id [object_id]	The unique identifier for the resource in Azure.
Provisioning State [provisioning_state]	Current state of the provisioning process for the resource.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The identifier for the tenant under which the resource belongs.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Application Security Group - Extended Inventory(LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical region where the resource is deployed.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier for the resource in Azure.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.network/applicationsecuritygroups**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Application Security Group [cmdb_azure_application_security_group]	References	Cloud Resource [cmdb_ci_cmp_resource] [cmdb_ci_cmp_resource_application_security_group]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Automation Account pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Automation Account - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Automation Account [cmdb_azure_automation_automation_account]

Field	Description
Disable Local Auth [disable_local_auth]	A setting indicating whether local authentication is disabled for the resource.
Identity [identity]	The managed identity associated with the resource.
Location [location]	The geographical location where the resource is hosted.
Object Id [object_id]	The unique identifier of the resource.
Public Network Access [public_network_access]	The configuration status of public network access for the resource.
Resource Group [resource_group]	Name of the resource group.
Subscription Id [subscription_id]	The subscription ID.
Tenant Id [tenant_id]	The identifier for the Azure Active Directory tenant associated with the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Automation Account - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical location where the resource is hosted.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.automation/automationaccounts**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Automation Account [cmdb_azure_automation_automation_account]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure availability sets discovery using patterns

Discovery and Service Mapping Patterns uses the Azure - Availability Set (LP) pattern to discover Azure availability sets during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Availability Set (LP) pattern.

Availability Set [cmdb_ci_availability_set]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Microsoft Azure for this resource.
Name [name]	The Name or ID if no Name is specified for the availability zone.
Install Status [install_status]	Resource provisioning status.

CI relationships

Discovery creates these relationships to support the Azure availability sets discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Availability Set [cmdb_ci_availability_set]
Azure Datacenter [cmdb_ci_azure_datacenter]	Contains::Contained by	Availability Set [cmdb_ci_availability_set]

Azure availability zones discovery using patterns

Discovery and Service Mapping Patterns uses the Azure - Availability Zones (LP) pattern to discover Azure availability zones during horizontal discovery. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Availability Zones (LP) pattern.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Object ID [object_id]	Name of the availability zone in the following format: **{Location} + "-" + {ZoneNum}**.
Name [name]	Name of the availability zone in the following format: **{Location} + "-" + {ZoneNum}**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Availability Zone [cmdb_ci_availability_zone] (continued)

Field	Description
Install Status [install_status]	Install status of the availability zone. Default value is Installed.

CI relationships

Discovery creates these relationships to support the Azure availability zone discovery.

CI	Relationship	CI
Azure Datacenter [cmdb_ci_azure_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]

Azure Blob Storage pattern-based discovery

Discovery and Service Mapping Patterns finds blob resources within Azure Blob Storage on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Storage Blobs(LP) pattern.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	The name of the storage blob.
Volume ID [volume_id]	The unique identifier for the storage volume, matching the blob name.
Object ID [object_id]	The unique Azure resource identifier combining the storage container path and blob name.
Volume container [volume_container]	The name of the storage container that holds the blob.
Storage type [storage_type]	The type of blob storage. For example: BlockBlob, PageBlob, or AppendBlob.
State [state]	The lease state of the blob. For example: Leased, Available, or Terminated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Block Endpoint [cmdb_ci_endpoint_block]

Field	Description
Name [name]	The name of the block endpoint, matching the blob name.
Object ID [object_id]	The unique Azure resource identifier for the block endpoint.

CI relationships

The Azure - Storage Blobs(LP) pattern creates these relationships to support Azure Blob Storage discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Block Endpoint [cmdb_ci_endpoint_block]	Implement End Point To::Implement End Point From	Storage Volume [cmdb_ci_storage_volume]
Cloud Storage Account [cmdb_ci_cloud_storage_account]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]

Azure Classic Load Balancer pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Classic Load Balancers on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Classic LB (LP) pattern.

Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Name [name]	The name of the load balancer.
State [state]	The current state of the load balancer.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud LB IP Address [cmdb_ci_cloud_lb_ipaddress]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Name [name]	IP address of the load balancer.
IP Address [ip_address]	IP address of the load balancer.
IPAddress Type [ipaddress_type]	The type of the IP address, which can be private or public.
Fully qualified domain name [fqdn]	The fully qualified domain name (FQDN) of the load balancer.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Name [name]	The name of the load balancing rule or inbound NAT rule.
Port [port]	The TCP port that the load balancer service listens to.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Name [name]	The name of the backend address pool.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - Classic LB (LP) pattern creates these relationships to support Azure Classic Load Balancer discovery.

CI	Relationship	CI
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Owns::Owned by	Cloud LB IP Address [cmdb_ci_cloud_lb_ipaddress]
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Load Balancer Service [cmdb_ci_lb_service]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Load Balancer Pool [cmdb_ci_lb_pool]	Hosted on::Hosts	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Load Balancer Pool [cmdb_ci_lb_pool]	Hosted on::Hosts	Load Balancer Service [cmdb_ci_lb_service]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Connections discovered by Service Mapping during top-down discovery

The Azure LoadBalancer TD pattern performs top-down discovery of Azure Classic Load Balancers. Service Mapping discovers cluster connections from the load balancer service to backend pool member IP addresses using either load balancing rules or inbound NAT rules.

Azure Compute Gallery Image Definition pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Compute Gallery Image Definition - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Compute Gallery - Image Definition [cmdb_azure_compute_gallery_image_definition]

Field	Description
Architecture [architecture]	Central processing unit (CPU) architecture of the resource, such as x86 or ARM.
Hyper-V Generation [hyper_v_generation]	The Hyper-V generation version of the virtual machine.
Kind [kind]	The specific subcategory or variation type of the resource.
Location [location]	The geographical region where the resource is deployed.
Object Id [object_id]	The unique identifier of the resource.
OS Type [os_type]	The operating system type associated with the resource.
Provisioning State [provisioning_state]	Current provisioning status of the resource.
Resource Group [resource_group]	Name of the resource group.
SKU [sku]	The stock keeping unit (SKU) that defines the pricing tier of the resource.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The unique identifier for the Azure Active Directory tenant.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Compute Gallery Image Definition - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical region where the resource is deployed.
Name [name]	The name given to the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.compute/galleries/images**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Compute Gallery - Image Definition [cmdb_azure_compute_gallery_image_definition]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Compute Snapshot pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Compute Snapshot - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Compute - Snapshot [cmdb_azure_compute_snapshot]

Field	Description
Disk State [disk_state]	The current state of the disk. For example: Attached or Unattached.
Encryption Type [encryption_type]	The encryption mechanism applied to the resource data. For example: EncryptionAtRestWithCustomerKey or EncryptionAtRestWithPlatformKey.
DC Location [location]	The geographic location where the resource is deployed.
Network Access Policy [network_access_policy]	The policy that governs network access to the resource.
Object Id [object_id]	A unique identifier for the queried resource.
OS Type [os_type]	The type of operating system associated with the resource.
Provisioning State [provisioning_state]	Current provisioning state of the resource. For example: Succeeded or Failed.
Public Network Access [public_network_access]	Indicates whether the resource can be accessed via the public network.
Resource Group [resource_group]	Name of the resource group.
SKU Tier [sku_tier]	The tier of the stock keeping unit (SKU) that determines resource performance and costs.
Subscription Id [subscription_id]	The subscription ID.
Tenant Id [tenant_id]	The unique identifier of the tenant owning the resource.
Type [type]	The type or category of the resource.
Configuration Item [configuration_item]	References the Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Compute Snapshot - Extended Inventory(LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]

Field	Description
Capacity [capacity]	The size of the disk in terms of storage capacity.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographic location where the resource is deployed.
Name [name]	The name of the resource.
Object ID [object_id]	A unique identifier for the queried resource in the following format: **/subscriptions/{subscriptionId}/resourceGroups/{resouceGroupName}/providers/Microsoft.Compute/snapshots/{snapshotName}**.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
State [state]	The current state of the disk. For example: Attached or Unattached.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]
Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Compute - Snapshot [cmdb_azure_compute_snapshot]	References	Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Container Registry pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Container Registry - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Container Registry [cmdb_azure_container_registry_registry]

Field	Description
Admin User Enabled [admin_user_enabled]	Indicates whether the administrator user account is enabled for the resource.
Encryption Status [encryption_status]	The current encryption status of the resource, specifying if data is encrypted.
Identity [identity]	The identity information for the resource, used for authentication and authorization.
Kind [kind]	The specific kind or subtype of the resource, if applicable.
Location [location]	The geographical region where the resource is deployed.
Object Id [object_id]	The unique identifier for the resource.
Provisioning State [provisioning_state]	The current provisioning state of the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Container Registry [cmdb_azure_container_registry_registry] (continued)

Field	Description
Public Network Access [public_network_access]	Indicates whether public network access is allowed for the resource.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The Azure Active Directory tenant ID associated with the resource.
Tier [tier]	The service or resource tier, indicating the level of performance or cost.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Container Registry - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical region where the resource is deployed.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier for the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.containerregistry/registries**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Container Registry [cmdb_azure_container_registry_registry]	References	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Cosmos DB for PostgreSQL Cluster pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Cosmos DB for PostgreSQL Cluster - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Cosmos DB for PostgreSQL Cluster [cmdb_azure_cosmos_db_for_postgresql_cluster_cluster]

Field	Description
Object Id [object_id]	The unique identifier of the resource.
Kind [kind]	The specific kind or stock-keeping unit (SKU) of the resource.
DC Location [location]	The Azure region where the resource is deployed.
Resource Group [resource_group]	The name of the Azure resource group containing the resource.
Subscription Id [subscription_id]	The unique identifier for the Azure subscription.
Tenant Id [tenant_id]	The universally unique identifier (UUID) of the Azure Active Directory tenant associated with the resource.
Provisioning State [provisioning_state]	The current state of the provisioning process for the resource.
Configuration Item [configuration_item]	References the Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Cosmos DB for PostgreSQL Cluster - Extended Inventory(LP) pattern.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Name [name]	The name assigned to the resource.
Location [location]	The Azure region where the resource is deployed.
Cluster ID [cluster_id]	The unique identifier of the resource.
Cluster Status [cluster_status]	The current operational status of the cluster.
Cluster Version [cluster_version]	The version number of the cluster software.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **microsoft.dbforpostgresql/servergroupsv2**.
Vendor [vendor]	Name of vendor: Microsoft.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Azure - Cosmos DB for PostgreSQL Cluster - Extended Inventory(LP) pattern creates these relationships to support Azure Cosmos DB for PostgreSQL Cluster discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Cosmos DB for PostgreSQL Cluster [cmdb_azure_cosmos_db_for_postgresql_cluster_cluster]	References	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Data Explorer Cluster pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Data Explorer Cluster - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Data Explorer Cluster [cmdb_azure_data_explorer_clusters_data_explorer_clusters]

Field	Description
Object Id [object_id]	The unique identifier of the Azure resource.
Kind [kind]	The specific kind or stock-keeping unit (SKU) of the resource.
DC Location [location]	The Azure region where the resource is deployed.
Resource Group [resource_group]	The name of the resource group that contains the resource.
Subscription Id [subscription_id]	The unique identifier of the Azure subscription.
Tenant Id [tenant_id]	The unique identifier of the Azure Active Directory tenant.
Provisioning State [provisioning_state]	The current provisioning state of the resource.
Configuration Item [configuration_item]	References the Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] table.
Enable Disk Encryption [enable_disk_encryption]	Specifies whether disk encryption is enabled on the resource.
Public Network Access [public_network_access]	Indicates if the resource is accessible over a public network.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Data Explorer Cluster - Extended Inventory(LP) pattern.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Field	Description
Name [name]	The name assigned to the Azure resource.
Location [location]	The Azure region where the resource is deployed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster] (continued)

Field	Description
Cluster ID [cluster_id]	The unique identifier of the Azure resource.
Cluster Status [cluster_status]	The current operational status of the Kusto cluster.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Type of resource. The value is set to **microsoft.kusto/clusters**.
Vendor [vendor]	Name of vendor: Microsoft.

CI relationships

The Azure - Data Explorer Cluster - Extended Inventory(LP) pattern creates these relationships to support Azure Data Explorer Cluster discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]
Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Data Explorer Cluster [cmdb_azure_data_explorer_clusters_data_explorer_clusters]	References	Cloud DataBase Cluster [cmdb_ci_cloud_db_cluster]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Data Factory pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Data Factory - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Data Factory [cmdb_azure_data_factory_data_factory]

Field	Description
Identity [identity]	Information about the identity associated with the resource.
Kind [kind]	The specific kind or stock keeping unit (SKU) of the resource.
Location [location]	The geographical location where the resource is deployed.
Object Id [object_id]	The unique identifier of the resource.
Provisioning State [provisioning_state]	The current provisioning status of the resource.
Public Network Access [public_network_access]	The status of public network accessibility for the resource.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The unique identifier of the tenant containing the resource.
Version [version]	The version of the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Data Factory [cmdb_azure_data_factory_data_factory] (continued)

Field	Description
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Data Factory - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical location where the resource is deployed.
Name [name]	The name assigned to the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.datafactory/factories**.
Description [short_description]	Description of the resource or service.
Version [version]	The version of the resource.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Data Factory [cmdb_azure_data_factory_data_factory]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Data Protection Backup Vault pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Data Protection Backup Vault - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Data Protection - Backup Vault [cmdb_azure_data_protection_backup_vault]

Field	Description
Kind [kind]	The type or category of the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Data Protection - Backup Vault [cmdb_azure_data_protection_backup_vault] (continued)

Field	Description
Location [location]	The geographic location where the resource is deployed.
Object Id [object_id]	The unique identifier for the resource.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The identifier of the Azure Active Directory tenant.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Data Protection Backup Vault - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographic location where the resource is deployed.
Name [name]	The name given to the resource.
Object ID [object_id]	The unique identifier for the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.dataprotection/backupvaults**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Data Protection - Backup Vault [cmdb_azure_data_protection_backup_vault]	References	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Database pattern-based discovery

Discovery and Service Mapping Patterns finds Azure databases, including SQL Server, MySQL, PostgreSQL, Redis, and Azure Cosmos DB, in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Verify requirements for Azure SQL Managed Instance license discovery

Required plugins and applications:

- Software Asset Management Professional for Microsoft
- Visibility Content

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure DataBase (LP) pattern.

Cloud DataBase [cmdb_ci_cloud_database]

Field	Description
Name [name]	The name of the database that you created in Azure.
Object ID [object_id]	The identification name of the database.
Fully qualified domain name [fqdn]	The FQDN that Azure assigned to your database.
Type [type]	The type of database you created.
Version [version]	The version of the database.
State [state]	The state of the database: whether it's Available or Terminated.
Vendor [vendor]	The vendor name is Azure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBase [cmdb_ci_cloud_database] (continued)

Field	Description
Operational status [operational_status]	The operational status of the database, derived from its Azure provisioning state.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Category [category]*	The stock keeping unit (SKU) family.

* Populated only by the Azure SQL Managed Instance license pattern extension.

IP Address [cmdb_ci_ip_address]

Field	Description
Netmask [netmask]	Netmask of the Azure database.

Database [cmdb_ci_database]

Field	Description
Name [name]	Name of the database in the following format: `<database name>@<server name>`.
Type [type]	Database engine type. Possible values: • Microsoft SQL Server • MySQL • Postgres SQL • Redis • DocumentDB
Serial number [serial_number]	Unique identifier of the database.
Life Cycle Stage [life_cycle_stage]	Life cycle stage of the database, derived from its Azure status. For example: Operational or End of Life.
Life Cycle Stage Status [life_cycle_stage_status]	Life cycle stage status of the database, derived from its Azure status. For example: In Use or Retired.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Name [name]	Name of the hardware type, based on the database SKU.
Object ID [object_id]	Unique identifier for the hardware type, based on the database SKU name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type] (continued)

Field	Description
vCPUs [vcpus]	Number of virtual CPUs, extracted from the SKU name.
Provider [provider]	Cloud provider of the hardware type. The value is set to **AZURE**. This field is only populated in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

Note: When using the Hardware Type [cmdb_ci_compute_template] table to store the hardware types, you may notice an unusually large number of records. To avoid this issue, you can store the discovered hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. For more information, see Enable the Cloud Hardware Type class extension.

Discovery populates the data in the CMDB when running the Azure SQL Managed Instance license pattern extension.

Serverless Hardware [cmdb_ci_serverless_hardware]

Field	Description
Object ID [object_id]	Object ID of the Azure cloud database.
Name [name]	SKU name.
Cloud Vendor [cloud_vendor]	Cloud vendor of the serverless hardware: **MS Azure**.
CPU core count [cpu_core_count]	Number of virtual cores (vCores).
CPU core thread [cpu_core_thread]	Number of vCores.
CPU count [cpu_count]	Number of vCores.
Category [category]	vCore purchasing model.
Subcategory [subcategory]	SKU tier.
Host Type [host_type]	Host type: **PaaS**.

CI relationships and references

The Azure DataBase (LP) pattern creates the following relationships and references to support Azure database discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Cloud DataBase [cmdb_ci_cloud_database]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Cloud DataBase [cmdb_ci_cloud_database]	Contains::Contained by	Database [cmdb_ci_database]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud DataBase [cmdb_ci_cloud_database]
Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Database [cmdb_ci_database]	Provisioned From::Provisioned	Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud DataBase [cmdb_ci_cloud_database]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Database [cmdb_ci_database]

The Azure SQL Managed Instance license pattern extension creates the following relationships.

CI relationships

CI	Relationship	CI
Cloud DataBase [cmdb_ci_cloud_database]	Runs on::Runs	Serverless Hardware [cmdb_ci_serverless_hardware]
Serverless Hardware [cmdb_ci_serverless_hardware]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud DataBase [cmdb_ci_cloud_database]

Azure Tag discovery

The Collect Azure DataBase Tags (LP) pattern extension collects tags from database servers and databases, and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table (for server-level tags) or the Database [cmdb_ci_database] table (for database-level tags).

Azure SQL Managed Instance license discovery

The Azure SQL Managed Instance license pattern extension populates the license type in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	**SQL_Server_PaaS_Managed_Instance_License_Type_automatic**
Value [value]	License type. The following maps the Azure portal license to ServiceNow values: • Azure Hybrid Benefit: **BYOL** • Pay as you go: **License Included** • Hybrid Failover rights: **Hybrid Failover**
Configuration item [configuration_item]	References the Cloud DataBase [cmdb_ci_cloud_database] table.

Connections found by Service Mapping during top-down discovery

The Azure DataBase TD pattern discovers the following connections during top-down discovery: SQL Server, MySQL, PostgreSQL, Redis, and Cosmos DB (DocumentDB).

Azure Databricks Workspace pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Databricks Workspace - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Databricks - Workspace [cmdb_azure_databricks_workspace]

Field	Description
Kind [kind]	The model or configuration type of the resource.
Location [location]	Geographical region where the resource is deployed.
Object Id [object_id]	The unique identifier of the resource.
Provisioning State [provisioning_state]	The current provisioning state of the resource.
Public Network Access [public_network_access]	Indicates whether public network access is enabled for the resource.
Resource Group [resource_group]	Name of the resource group.
SKU name [sku_name]	The stock keeping unit (SKU) that defines the resource's pricing tier.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The identifier for the Azure Active Directory tenant.
Workspace ID [workspace_id]	The unique identifier for the workspace associated with the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Databricks Workspace - Extended Inventory(LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	Geographical region where the resource is deployed.
Name [name]	The name assigned to the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.databricks/workspaces**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Databricks - Workspace [cmdb_azure_databricks_workspace]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Datacenter discovery pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Datacenter resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure Datacenter discovery pattern.

Azure Datacenter [cmdb_ci_azure_datacenter]

Field	Description
Object ID [object_id]	The Azure region name. For example: asia, australiacentral, or eastus.
Region [region]	The Azure region name. For example: asia, australiacentral, or eastus.
Name [name]	The display name of the Azure region. For example: Asia, Australia Central, or East US.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	The Azure subscription or management group account identifier.
Object ID [object_id]	The unique object identifier for the Azure subscription or management group.
Datacenter Type [datacenter_type]	The type of datacenter: **Azure Datacenter [cmdb_ci_azure_datacenter]**.
Is management account [is_master_account]	Indicates whether this is a management group account that manages multiple subscriptions.
Datacenter URL [datacenter_url]	The Azure management endpoint URL for the account.

CI relationships

The Azure Datacenter discovery pattern creates these relationships to support Azure Datacenter discovery.

CI	Relationship	CI
Azure Datacenter [cmdb_ci_azure_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]

Azure Dev Center pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Dev Center - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Dev Center [cmdb_azure_dev_centers_dev_center]

Field	Description
Object Id [object_id]	The unique resource identifier in Azure Resource Manager format.
Kind [kind]	The specific type or classification of the resource.
DC Location [location]	The Azure datacenter region where the resource is located.
Resource Group [resource_group]	The name of the Azure Resource Group containing the resource.
Subscription Id [subscription_id]	The unique identifier of the Azure subscription.
Tenant Id [tenant_id]	The unique identifier of the Azure Active Directory tenant.
Provisioning State [provisioning_state]	The current state of resource provisioning. For example: Succeeded or Failed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Dev Center [cmdb_azure_dev_centers_dev_center] (continued)

Field	Description
Configuration Item [configuration_item]	References the Cloud App Server [cmdb_ci_cloud_appserver] table.
Dev Center URI [dev_center_uri]	The Uniform Resource Identifier (URI) endpoint for connecting to the Dev Center resource.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Dev Center - Extended Inventory(LP) pattern.

Cloud App Server [cmdb_ci_cloud_appserver]

Field	Description
Object ID [object_id]	The unique resource identifier in Azure Resource Manager format.
Name [name]	The name of the resource.
Type [type]	Type of resource. The value is set to **microsoft.devcenter/devcenters**.
Location [location]	The Azure datacenter region where the resource is located.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Azure - Dev Center - Extended Inventory(LP) pattern creates these relationships to support Azure Dev Center discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud App Server [cmdb_ci_cloud_appserver]
Cloud App Server [cmdb_ci_cloud_appserver]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Dev Center [cmdb_azure_dev_centers_dev_center]	References	Cloud App Server [cmdb_ci_cloud_appserver]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Disk Encryption Set pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Disk Encryption Set - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Disk Encryption Set [cmdb_azure_disk_encryption_set_disk_encryption_set]

Field	Description
Encryption Type [encryption_type]	The method used to encrypt the resource data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Disk Encryption Set [cmdb_azure_disk_encryption_set_disk_encryption_set] (continued)

Field	Description
Identity type [identity_type]	The classification of the identity assigned to the resource. For example: SystemAssigned or UserAssigned.
Key Url [key_url]	The URL where the encryption key is stored or accessible.
Location [location]	The geographic region where the resource is deployed.
Object Id [object_id]	The unique identifier of the resource.
Provisioning State [provisioning_state]	The current status of the resource provisioning process.
Resource Group [resource_group]	Name of the resource group.
Tenant Id [tenant_id]	The identifier for the tenant that owns the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Disk Encryption Set - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographic region where the resource is deployed.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.compute/diskencryptionsets**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Disk Encryption Set [cmdb_azure_disk_encryption_set]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure DNS zones and record sets discovery using Patterns

The Discovery and Service Mapping Patterns application uses the Azure - DNS Zones (LP) and Azure - DNS Zone Recordsets (LP) Patterns to discover Azure Domain Name System (DNS) definitions (zones and their respective record sets). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about Azure DNS zones and record sets and their version that you can discovery, refer to Detailed information on products discovered by ITOM Visibility.

For more information on Azure cloud discovery, see Discovery for Microsoft Azure.

Prerequisites

Verify privileges

On the Microsoft Azure Console, provide a user with the API Management Service Reader role.

Create Azure Service Principal credentials

For more information see: Create Azure cloud credentials.

Configure Azure service account

For more information, see Set up a cloud account for Azure.

Create discovery schedule

Create a schedule for the relevant Azure service accounts. For more information, see Create a discovery schedule in Cloud Discovery Workspace.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - DNS Zones (LP) and Azure - DNS Zone Recordsets (LP) Patterns.

Resources discovered by the Azure - DNS Zones (LP) pattern

DNS Zone [cmdb_ci_dns_zone]

Field	Description
Name [name]	The name of the DNS Zone.
Number Of RecordSets [number_of_recordsets]	The number of the recordsets defined within the DNS Zone.
Operational status [operational_status]	The operational status of the DNS Zone.
State [state]	The state of the DNS Zone.
Install Status [install_status]	The install status of the DNS Zone.
Comments [comments]	The field for general comments.

Resources discovered by the Azure - DNS Zone Recordsets (LP) pattern

DNS Alias [cmdb_ci_dns_alias]

Field	Description
Name [name]	The name of the DNS alias record.
IP Address [ip_address]	IP address of the DNS alias record.
Comments [comments]	Unique identifier for an Azure DNS zone combining service account ID, DNS zone name, and a fixed hash suffix.
Operational status [operational_status]	The operational status of the DNS alias record.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	The name of the DNS Name record.
Operational status [operational_status]	The operational status of the DNS record.
Comments [comments]	Unique identifier for an Azure DNS zone combining service account ID, DNS zone name, and a fixed hash suffix.

CI relationships

These relationships are created to support Azure DNS zone discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	DNS Zone [cmdb_ci_dns_zone]

These relationships are created to support Azure DNS zone record set discovery.

CI	Relationship	CI
DNS Zone [cmdb_ci_dns_zone]	Contains::Contained by	DNS Name [cmdb_ci_dns_name]
DNS Zone [cmdb_ci_dns_zone]	Contains::Contained by	DNS Alias [cmdb_ci_dns_alias]
DNS Name [cmdb_ci_dns_name]	Used by::Uses	DNS Alias [cmdb_ci_dns_alias]

Azure Event Grid System Topic discovery with Patterns

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Event Grid System Topic - Extended Inventory(LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Event Grid - System Topic [cmdb_azure_event_grid_system_topic]

Field	Description
Kind [kind]	The specific kind or stock keeping unit (SKU) of the resource, detailing its capabilities or configuration.
Metric Resource Id [metric_resource_id]	The identifier used for accessing metrics related to the resource.
Object Id [object_id]	The unique identifier of the resource within Azure.
Provisioning State [provisioning_state]	The current status of the resource's provisioning process.
Resource Group [resource_group]	Name of the resource group.
Subscription Id [subscription_id]	The subscription ID.
Tenant Id [tenant_id]	The unique identifier for the Azure Active Directory tenant.
Topic Type [topic_type]	The category or type of the event topic.
Configuration Item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Event Grid System Topic - Extended Inventory(LP) pattern.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical region where the resource is hosted.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier of the resource within Azure.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **microsoft.eventgrid/systemtopics**.

CI relationships

The pattern creates these relationships to support discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]
Cloud Messaging Service [cmdb_ci_cloud_messaging_service]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Event Grid - System Topic [cmdb_azure_event_grid_system_topic]	References	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Event Hub Namespace pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Event Hub Namespace - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Event Hub - Namespace [cmdb_azure_event_hub_namespace]

Field	Description
Kind [kind]	The nature or subtype of the resource, often describing its specific role or variation.
Location [location]	The geographical location where the resource is hosted.
Minimum TLS Version [minimum_tls_version]	The minimum Transport Layer Security (TLS) version required by the resource for secure connections.
Object Id [object_id]	The unique identifier of the resource.
Provisioning State [provisioning_state]	The current status of the resource provisioning process.
Public Network #ccess [public_network_access]	Indicates whether the resource can be accessed via the public network.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The identifier of the Azure Active Directory tenant.
Configuration Item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Event Hub Namespace - Extended Inventory(LP) pattern.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical location where the resource is hosted.
Name [name]	The name given to the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service] (continued)

Field	Description
Type [type]	Type of resource. The value is set to **microsoft.eventhub/namespaces**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]
Cloud Messaging Service [cmdb_ci_cloud_messaging_service]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Event Hub - Namespace [cmdb_azure_event_hub_namespace]	References	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Events pattern-based discovery

Discovery uses event patterns to update Microsoft Azure Cloud component data in near real-time. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Events discovered by Discovery during horizontal discovery

Discovery uses patterns to find events created for Microsoft Azure Cloud components. If there are events that indicate the change of state in one of the Microsoft Azure Cloud components, it triggers the discovery of Microsoft Azure Cloud components using the patterns.

Patterns used for event discovery

Pattern	CI
Azure Application LB Event	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Azure Availability Set Event	Availability Set [cmdb_ci_availability_set]
Azure Classic LB Event	Cloud Load Balancer [cmdb_ci_cloud_load_balancer]
Azure DataBase Event	Cloud DataBase [cmdb_ci_cloud_database]
Azure Express Route Circuit Event	Cloud Direct Connect [cmdb_ci_cloud_direct_connect]
Azure Functions Event	Cloud Function [cmdb_ci_cloud_function]
Azure Local Network Gateway Event	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]
Azure NAT Gateway Event	NAT Gateway [cmdb_ci_nat_gateway]
Azure Network Event	Cloud Network [cmdb_ci_network]
Azure NIC Event	Cloud Mgmt Network Interface [cmdb_ci_nic]
Azure Private DNS Zone Event	DNS Zone [cmdb_ci_dns_zone]
Azure Public IP Event	Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]
Azure Resource Group Event	Resource Group [cmdb_ci_resource_group]
Azure Security Group Event	Compute Security Group [cmdb_ci_compute_security_group]
Azure Storage Account Event	Cloud Storage Account [cmdb_ci_cloud_storage_account]
Azure Virtual Machine Event	Virtual Machine Instance [cmdb_ci_vm_instance]
Azure Virtual Network Gateway Connection Event	Virtual Network Gateway Connection [cmdb_ci_vpc_gateway_connection]
Azure Virtual Network Peerings Event	Virtual Network Peering [cmdb_ci_vnet_peering]
Azure VM Scale Set Event	Instance Scale Set [cmdb_ci_instance_scale_set]

Related topics

- Azure Application LB pattern-based discovery
- Azure availability sets discovery using patterns
- Azure Classic Load Balancer pattern-based discovery
- Azure Database pattern-based discovery
- Azure Express Route Circuit pattern-based discovery
- Microsoft Azure Functions discovery with Patterns
- Azure Local Network Gateway pattern-based discovery
- Azure NAT Gateway pattern-based discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Azure Network and Subnet pattern-based discovery
- Azure NIC pattern-based discovery
- Azure Private DNS Zone pattern-based discovery
- Azure Public IP pattern-based discovery
- Azure Resource Group pattern-based discovery
- Azure Security Group pattern-based discovery
- Azure Storage Account pattern-based discovery
- Azure virtual machine pattern-based discovery
- Azure Virtual Network Gateway Connection pattern-based discovery
- Azure Virtual Machine Scale Sets (VMSS) Instance discovery

Azure Express Route Circuit pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Express Route Circuit resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Express Route Circuit (LP) pattern.

Cloud Direct Connect [cmdb_ci_cloud_direct_connect]

Field	Description
Object ID [object_id]	The unique Azure resource identifier for the Express Route Circuit.
Name [name]	The name of the Express Route Circuit.
State [state]	The current state of the Express Route Circuit.
VLAN ID [vlanid]	The Virtual LAN (VLAN) identifier assigned to the Express Route Circuit.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - Express Route Circuit (LP) pattern creates these relationships to support Azure Express Route Circuit discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Direct Connect [cmdb_ci_cloud_direct_connect]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Direct Connect [cmdb_ci_cloud_direct_connect]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Key Value [cmdb_key_value]	References	Cloud Direct Connect [cmdb_ci_cloud_direct_connect]

Azure Tag discovery

The Azure - Express Route Circuit (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Direct Connect [cmdb_ci_cloud_direct_connect] table.

Azure File Share pattern-based discovery

Discovery and Service Mapping Patterns finds Azure File Share resources within Storage Accounts on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - File Share (LP) pattern.

Cloud File Share [cmdb_ci_cloud_file_share]

Field	Description
Object ID [object_id]	The unique Azure resource identifier for the File Share.
Name [name]	The name of the File Share resource.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - File Share (LP) pattern creates these relationships to support Azure File Share discovery.

CI	Relationship	CI
Cloud File Service [cmdb_ci_cloud_file_service]	Contains::Contained by	Cloud File Share [cmdb_ci_cloud_file_share]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Firewall Network Security pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Network Security Azure Firewall - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Network Security - Azure Firewall [cmdb_azure_network_security_azure_firewall]

Field	Description
Object Id [object_id]	The unique identifier of the resource.
Kind [kind]	The specific resource kind or sub-type (if applicable).
DC Location [location]	The Azure region where the resource is deployed.
Resource Group [resource_group]	The name of the Azure Resource Group containing the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Network Security - Azure Firewall [cmdb_azure_network_security_azure_firewall] (continued)

Field	Description
Subscription Id [subscription_id]	The identifier of the Azure subscription hosting the resource.
Tenant Id [tenant_id]	The identifier for the Azure Active Directory tenant.
Provisioning State [provisioning_state]	The current lifecycle state of the resource. For example: Succeeded or Updating.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Firewall Policy [firewall_policy]	The identifier or reference to the firewall policy applied to the Azure Firewall resource.
Tier [tier]	The service or performance tier assigned to the resource.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Network Security Azure Firewall - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The unique identifier of the resource.
Name [name]	The name assigned to the resource.
Location [location]	The Azure region where the resource is deployed.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.network/azurefirewalls**.

CI relationships

The Azure - Network Security Azure Firewall - Extended Inventory(LP) pattern creates these relationships to support Azure Firewall Network Security discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Azure Network Security - Azure Firewall [cmdb_azure_network_security_azure_firewall]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Microsoft Azure Functions discovery with Patterns

The Discovery and Service Mapping Patterns application uses the Azure - Functions (LP) and Azure - Functions TD Patterns to discover Azure Functions apps and map them in the context of application services. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Azure Functions detailed information

Detailed information on products discovered by ITOM Visibility

Azure service account

Configure the Azure service account. For more information, see Set up a cloud account for Azure

Azure Principal credentials

On the ServiceNow AI Platform, create a service account for cloud discovery. Set the account Id to the Azure Service Principal as it appears at the Microsoft Azure Console. For more information see: Create Azure cloud credentials

Privileges

On the Microsoft Azure Console, provide a user with the API Management Service Reader role.

Run Azure Cloud discovery

For detailed information, see: Discovery for Microsoft Azure

Create discovery schedule

Create a discovery schedule in Cloud Discovery Workspace

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Table and field	Description
Azure Function [cmdb_ci_cloud_function]	
name	The name of the Function app.
Object_id	The ID of the Function app.
ip_address	The IP Address of the Function app.
fqdn	The Fully Qualified Domain Name (FQDN) of the Function app.
code_location_url	The URL of the Function app source code.
app_function_state	The state of the Function app.
operational_status	The operational status of the Function app.
install_status	The installation status of the Function app.

CI relationships

These relationships are created to support Azure Functions discovery:

CI	Relationship	CI
Azure Functions [cmdb_ci_cloud_function]	Hosted on::Hosts	Azure datacenter [cmdb_ci_azure_datacenter]

Connections found by Service Mapping during top-down discovery

Service Mapping can discover the following connections:

- Connections to the following databases: MSSql, MySql, MongoDB, and Documentum.
- HTTP connections belonging to Azure Functions environment, for example, https://scenario4-hop1-azure-function.azurewebsites.net.

> **Note:** Service Mapping does not scan java classes/jars and binary files for possible information about outgoing HTTP connections of Azure Functions. If java classes/jars and binary files are the only source of information about some outgoing HTTP connections, Service Mapping does not find such HTTP connections.

Related topics

Define a Cloud REST Query for cloud computing devices

Azure hardware type pattern-based discovery

Discovery and Service Mapping Patterns finds Azure hardware type configurations on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses the Azure - Hardware Type (LP) pattern to discover available Azure hardware type configurations (VM sizes) from your Azure account.

When using the Hardware Type [cmdb_ci_compute_template] table to store the hardware types, you may notice an unusually large number of records. To avoid this issue, you can store the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

discovered hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. When enabled, Discovery launches the Azure - Cloud Hardware Type (LP) pattern to populate the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. For more information, see Enable the Cloud Hardware Type class extension.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure hardware type patterns.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Name [name]	The name of the hardware type configuration. For example: **Standard_A1_v2** or **Standard_D2s_v3**.
Object ID [object_id]	The name of the hardware type configuration. For example: **Standard_A1_v2** or **Standard_D2s_v3**.
Cores [cores]	The number of CPU cores available for this hardware type configuration.
Memory MB [memory_mb]	The amount of memory available for this hardware type configuration, in megabytes (MB).
vCPUs [vcpus]	The number of Virtual Central Processing Units (vCPUs) available for this hardware type configuration.
Local Storage GB [local_storage_gb]	The size of the temporary disk (resource disk) available for this hardware type configuration, in gigabytes (GB).
Provider [provider]	The cloud provider, which is **AZURE**. This field is only populated in Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

CI relationships

The Azure hardware type patterns create these relationships to support Azure hardware type discovery.

CI	Relationship	CI
Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Host pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Hosts on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Host (LP) pattern.

Host [cmdb_ci_cloud_host]

Field	Description
Name [name]	The Name or ID if no Name is specified for the host.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Host Type [host_type]	The Stock Keeping Unit (SKU) name of the dedicated host.
Cloud Vendor [cloud_vendor]	The cloud vendor: **Azure**.
Is Virtual [virtual]	Indicates whether this is a virtual host. Value is set to false.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - Host (LP) pattern creates these relationships to support Azure Host discovery.

CI	Relationship	CI
Host [cmdb_ci_cloud_host]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Host [cmdb_ci_cloud_host]
Virtual Machine Instance [cmdb_ci_vm_instance]	Runs on::Runs	Host [cmdb_ci_cloud_host]
Key Value [cmdb_key_value]	References	Host [cmdb_ci_cloud_host] or Virtual Machine Instance [cmdb_ci_vm_instance]

Azure Tag discovery

The pattern extension sections discover Bring Your Own License (BYOL) or the included licenses for Windows VMs, and collect Azure resource tags from Hosts.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The key is one of the following options:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
	• Tags: Azure resource tag name • License: Windows_OS_License_Type_automatic
Value [value]	The value is one of the following options: • Tags: Azure resource tag value • Azure Hybrid Benefit: BYOL • pay-as-you-go: License Included
Configuration item [configuration_item]	• Tags: References the Host [cmdb_ci_cloud_host] table • License: References the Virtual Machine Instance [cmdb_ci_vm_instance] table

Azure Key Vault pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Key Vault - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Key Vault [cmdb_azure_key_vault_key_vault]

Field	Description
Kind [kind]	The specific form or variant of the resource type.
Location [location]	The geographical location where the resource is hosted.
Object Id [object_id]	The unique identifier of the resource.
Provisioning State [provisioning_state]	The current provisioning status of the resource.
Public Network Access [public_network_access]	Indicates whether the resource is accessible over a public network.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The unique identifier of the Azure Active Directory tenant to which the resource belongs.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Key Vault - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical location where the resource is hosted.
Name [name]	The name assigned to the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.keyvault/vaults**.
Description [short_description]	Description of the resource or service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Key Vault [cmdb_azure_key_vault_key_vault]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Key Vault Key pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Key Vault Keys on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

The Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Key Vault Key - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Key Vault - Key [cmdb_azure_key_vault_key]

Field	Description
Name [name]	Name of the Key Vault Key.
Object Id [object_id]	Unique identifier for the Key Vault Key, in the format of the Azure Resource Manager (ARM) resource ID followed by /keys/<key-name>. For example: **/subscriptions/<subscription-id>/resourceGroups/<resource-group>/providers/Microsoft.KeyVault/vaults/<vault-name>/keys/<key-name>**.
Key ID [key_id]	The versioned URI of the key in Azure Key Vault.
Key Type [key_type]	The cryptographic algorithm of the key. For example: RSA or EC.
Curve [curve]	The elliptic curve name for EC keys. For example: P-256 or P-384.
Key Ops [key_ops]	The operations permitted for the key. For example: **sign,verify** or **sign,verify,wrapKey,unwrapKey,encrypt,decrypt**.
Key Size [key_size]	The size of the key in bits.
Enabled [enabled]	Indicates whether the key is enabled.
Recovery Level [recovery_level]	The deletion recovery level for the key. For example: Recoverable, Purgeable, or Recoverable+Purgeable.
Created Date [created_date]	Date and time when the key was created.
Updated Date [updated_date]	Date and time when the key was last updated.
Expiration Date [expiration_date]	Date and time when the key expires.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in CMDB tables

The Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Key Vault Key - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Name [name]	Name of the Key Vault Key.
Object ID [object_id]	Unique identifier for the Key Vault Key, in the format of the ARM resource ID followed by /keys/<key-name>. For example: **/subscriptions/<subscription-id>/resourceGroups/<resource-group>/providers/Microsoft.KeyVault/vaults/<vault-name>/keys/<key-name>**.
Resource type [resource_type]	The Azure resource type: **microsoft.keyvault/vaults/keys**.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The Azure - Key Vault Key - Extended Inventory(LP) pattern creates the following relationships and references to support Azure Key Vault Key discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Cloud Resource [cmdb_ci_cmp_resource]	Members::Member of	Cloud Resource [cmdb_ci_cmp_resource]

CI references

CI	Field	Referenced CI
Azure Key Vault - Key [cmdb_azure_key_vault_key]	Configuration Item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Resource [cmdb_ci_cmp_resource]

Azure Tag discovery

The Azure - Key Vault Key - Extended Inventory(LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Azure Local Network Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Local Network Gateway resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Local Network Gateway (LP) pattern.

Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

Field	Description
Object ID [object_id]	The unique Azure resource identifier for the Local Network Gateway.
Name [name]	The name of the Local Network Gateway resource.
State [state]	The operational state of the Local Network Gateway.
IP Address [ip_address]	The IP address of the on-premises VPN device that the Local Network Gateway connects to.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - Local Network Gateway (LP) pattern creates these relationships to support Azure Local Network Gateway discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]
Azure Datacenter [cmdb_ci_azure_datacenter]	Hosted on::Hosts	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]
Key Value [cmdb_key_value]	References	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Tag discovery

The Azure - Local Network Gateway (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway] table.

Azure Log Analytics Workspace pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Log Analytics Workspace - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Log Analytics - Workspace [cmdb_azure_log_analytics_workspace]

Field	Description
Object Id [object_id]	The unique identifier of the Azure resource.
Kind [kind]	The specific kind or variant of the resource.
DC Location [location]	The geographical region where the resource is deployed.
Resource Group [resource_group]	The Azure resource group to which the resource belongs.
Subscription Id [subscription_id]	The unique identifier of the Azure subscription containing the resource.
Tenant Id [tenant_id]	The identifier of the Azure Active Directory tenant associated with the resource.
Provisioning State [provisioning_state]	The current lifecycle state of the resource. For example: Succeeded, Creating, or Failed.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Customer Id [customer_id]	The unique identifier of the Log Analytics workspace or customer instance.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Log Analytics Workspace - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The unique identifier of the Azure resource.
Name [name]	The name of the resource.
Location [location]	The geographical region where the resource is deployed.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.operationalinsights/workspaces**.

CI relationships

The Azure - Log Analytics Workspace - Extended Inventory(LP) pattern creates these relationships to support Azure Log Analytics Workspace discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Log Analytics - Workspace [cmdb_azure_log_analytics_workspace]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Logic App pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Logic App - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Logic App [cmdb_azure_logic_app_logic_app]

Field	Description
Access Endpoint [access_endpoint]	The endpoint URL used for accessing the resource.
Identity [identity]	The identity associated with the resource for access and security purposes.
Kind [kind]	The specific subcategory or variant of the resource.
Location [location]	The geographical region where the resource is deployed.
Object Id [object_id]	The unique identifier of the resource.
Provisioning State [provisioning_state]	The current provisioning state of the resource.
Resource Group [resource_group]	Name of the resource group.
State [state]	The overall current operational state of the resource.
Subscription Id [subscription_id]	The subscription ID.
Tenant Id [tenant_id]	The identifier of the tenant that owns the resource.
Type [type]	The type or category of the resource.
Configuration Item [configuration_item]	References the Cloud Function [cmdb_ci_cloud_function] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Logic App - Extended Inventory(LP) pattern.

Cloud Function [cmdb_ci_cloud_function]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical region where the resource is deployed.
Name [name]	The name assigned to the resource.
Object ID [object_id]	The unique identifier of the resource in the following format: **/subscriptions/{subscriptionId}/resourceGroups/**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Function [cmdb_ci_cloud_function] (continued)

Field	Description
	{resouceGroupName}/providers/ Microsoft.logic/workflows/ {workflowName}.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Function [cmdb_ci_cloud_function]
Cloud Function [cmdb_ci_cloud_function]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Logic App [cmdb_azure_logic_app_logic_app]	References	Cloud Function [cmdb_ci_cloud_function]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Managed Identity User Assigned Identity pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Managed Identity User Assigned Identity - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Managed Identity - User Assigned Identity [cmdb_azure_managed_identity_user_assigned_identity]

Field	Description
Client ID [client_id]	The unique identifier for an Azure Active Directory (Azure AD) application, also known as a client.
Location [location]	The geographical region where the resource is deployed.
Object Id [object_id]	The unique identifier of the resource.
Principal ID [principal_id]	The unique identifier of the principal service or user associated with the resource.
Resource Group [resource_group]	Name of the resource group.
Tenant ID [tenant_id]	The unique identifier of the Azure AD tenant.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Managed Identity User Assigned Identity - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Location [location]	The geographical region where the resource is deployed.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.managedidentity/ userassignedidentities**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Managed Identity - User Assigned Identity [cmdb_azure_managed_identity_user_assigned_identity]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Marketplace pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Marketplace products in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

The Azure - Marketplace LB (LP) pattern discovers the following Azure Marketplace products:

- SaaS
- Azure Application

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Virtual Machine

Note: The pattern discovers only virtual machines (VMs) created from third-party or commercial marketplace images.

Azure Marketplace data model

The Azure - Marketplace LB (LP) pattern introduces the following CI class that extends an existing CMDB class.

CI class introduced by this pattern

CI class	Extends from
Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]	Virtual Machine Object [cmdb_ci_vm_object]

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Marketplace LB (LP) pattern.

Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]

Field	Description
Name [name]	Name of the Cloud resource, usually the marketplace offering or SKU name.
Object ID [object_id]	A unique resource ID of the Cloud resource.
Resource Type [resource_type]	The type of the resource in the Cloud Marketplace. For example: `microsoft.compute/virtualmachines`.
Plan Name [plan_name]	Billing or SKU plan for a resource from the Cloud Marketplace. For example: Pay as You Go.
Market [market]	International Organization for Standardization (ISO) code of the geographical market where the resource is sold. For example: US or EU.
Organization Id [organization_id]	A unique identifier for the organization or publisher that owns the marketplace resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Marketplace Product Details [marketplace_product_details]

Field	Description
Product Code [product_code]	A unique product code of the resource within the Cloud Marketplace.
Publisher Name [publisher_name]	The organization or individual responsible for creating and offering the product or service.
Version [version]	Release number or iteration of the product.
Deployed On [deployed_on]	References the Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product] table.

CI relationships and references

The Azure - Marketplace LB (LP) pattern creates the following relationships and references to support Azure Marketplace discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]
Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]*

* This relationship is created only for resources with a global region.

CI references

CI	Field	Referenced CI
Marketplace Product Details [marketplace_product_details]	Deployed On [deployed_on]	Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product]

Azure Tag discovery

The Azure - Marketplace LB (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the Deployed Marketplace Product [cmdb_ci_deployed_marketplace_product] table.

Azure NAT Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds Azure NAT Gateway resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - NAT Gateway (LP) pattern.

NAT Gateway [cmdb_ci_nat_gateway]

Field	Description
Object ID [object_id]	The unique Azure resource identifier for the NAT Gateway.
Name [name]	The name of the NAT Gateway resource.
State [state]	The operational state of the NAT Gateway.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - NAT Gateway (LP) pattern creates these relationships to support Azure NAT Gateway discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	NAT Gateway [cmdb_ci_nat_gateway]
NAT Gateway [cmdb_ci_nat_gateway]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Cloud Subnet [cmdb_ci_cloud_subnet]	Uses::Used by	NAT Gateway [cmdb_ci_nat_gateway]
Key Value [cmdb_key_value]	References	NAT Gateway [cmdb_ci_nat_gateway]

Azure Tag discovery

The Azure - NAT Gateway (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the NAT Gateway [cmdb_ci_nat_gateway] table.

Azure Network and Subnet pattern-based discovery

Discovery and Service Mapping Patterns finds Azure virtual networks, subnets, and virtual network peerings on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Network and Subnet (LP) pattern.

Cloud Network [cmdb_ci_network]

Field	Description
Name [name]	The Name or ID if no Name is specified for the cloud network.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
State [state]	The current state of the network.
Cidr [cidr]	Classless Inter-Domain Routing (CIDR) representation of the network. For example: 10.0.0.0/24.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Cloud Subnet [cmdb_ci_cloud_subnet]

Field	Description
Name [name]	The Name or ID if no Name is specified for the cloud subnet.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
State [state]	The current state of the subnet.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Subnet [cmdb_ci_cloud_subnet] (continued)

Field	Description
CIDR [cidr]	CIDR representation of the subnet. For example, 10.0.0.0/24.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Virtual Network Peering [cmdb_ci_vnet_peering]

Field	Description
Name [name]	The Name or ID if no Name is specified for the virtual network peering.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
State [state]	The current state of the virtual network peering.
Use Remote Gateways [use_remotegateways]	Indicates whether remote gateways are used for this peering.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - Network and Subnet (LP) pattern creates the following relationships and references to support Azure Network and Subnet discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Network [cmdb_ci_network]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Network [cmdb_ci_network]
Cloud Network [cmdb_ci_network]	Contains::Contained by	Cloud Subnet [cmdb_ci_cloud_subnet]
Cloud Network [cmdb_ci_network]	Contains::Contained by	Virtual Network Peering [cmdb_ci_vnet_peering]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Network [cmdb_ci_network]

Azure Tag discovery

The Azure - Network and Subnet (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Network [cmdb_ci_network] table.

Azure Networks IP Group pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Networks IP Group - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Networks - IP Group [cmdb_azure_networks_ip_group]

Field	Description
Location [location]	The geographic region where the resource is hosted.
Object Id [object_id]	A unique identifier for the resource within Azure.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The identifier for the Azure Active Directory tenant associated with the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Networks IP Group - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographic region where the resource is hosted.
Name [name]	The name assigned to the resource.
Object ID [object_id]	A unique identifier for the resource within Azure.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.network/ipgroups**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Networks - IP Group [cmdb_azure_networks_ip_group]	References	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure NIC pattern-based discovery

Discovery and Service Mapping Patterns finds Azure network interfaces Controller (NIC) on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

(Optional) Exclude temporary Azure Databricks VMs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can reduce short-lived configuration item (CI) records by excluding temporary Azure Databricks VMs. For more information, see Exclude temporary Azure Databricks virtual machines.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - NIC (LP) pattern.

Cloud Mgmt Network Interface [cmdb_ci_nic]

Field	Description
Name [name]	The Name or ID if no Name is specified for the network interface.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
MAC Address [mac_address]	The MAC address of the network interface.
State [state]	The current state of the network interface.
NIC Status [status]	The provisioning state of the network interface as returned by Azure.
Private IP [private_ip]	The IPv4 address of the network interface within the subnet.
Public IP [public_ip]	The address of the Elastic IP address bound to the network interface.
Public DNS [public_dns]	The public DNS name of the network interface.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Mgmt Network Interface [cmdb_ci_nic] (continued)

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - NIC (LP) pattern creates the following relationships and references to support Azure NIC discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Mgmt Network Interface [cmdb_ci_nic]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Mgmt Network Interface [cmdb_ci_nic]
Cloud Subnet [cmdb_ci_cloud_subnet]	Contains::Contained by	Cloud Mgmt Network Interface [cmdb_ci_nic]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Mgmt Network Interface [cmdb_ci_nic]

Azure Tag discovery

The Azure - NIC (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Mgmt Network Interface [cmdb_ci_nic] table.

Azure OS image pattern-based discovery

Discovery and Service Mapping Patterns finds Azure OS images on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern overview

Discovery uses the Azure - Image (LP) pattern to discover Azure OS images from your Azure account.

When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. When enabled, Discovery launches the Azure - Cloud OS Image (LP) pattern, which populates the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure OS image patterns.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Name [name]	Name of the image resource.
Object ID [object_id]	Azure resource ID of the image.
Image Source [image_source]	Azure resource ID of the image.
Guest OS [guest_os]	Operating system type of the image.
Provider [provider]	The cloud provider, which is **AZURE**. This field is only populated in Cloud Image [cmdb_ci_cloud_os_image] table.

CI relationships

The Azure OS image patterns create these relationships to support Azure OS image discovery.

CI	Relationship	CI
Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]
Key Value [cmdb_key_value]	References	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Tag discovery

The Azure OS image patterns collect tags and populate them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Image [cmdb_ci_os_template] table or the Cloud Image [cmdb_ci_cloud_os_image] table.

Azure Private DNS Zone pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Private DNS Zone resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Private DNS Zone (LP) pattern.

DNS Zone [cmdb_ci_dns_zone]

Field	Description
Object ID [object_id]	The unique Azure resource identifier for the Private DNS Zone.
Name [name]	The name of the Private DNS Zone resource.
State [state]	The operational state of the Private DNS Zone.
Number Of RecordSets [number_of_recordsets]	The number of DNS record sets within the Private DNS Zone.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - Private DNS Zone (LP) pattern creates these relationships to support Azure Private DNS Zone discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	DNS Zone [cmdb_ci_dns_zone]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
DNS Zone [cmdb_ci_dns_zone]	Connected by::Connects	Cloud Network [cmdb_ci_network]
Key Value [cmdb_key_value]	References	DNS Zone [cmdb_ci_dns_zone]

Azure Tag discovery

The Azure - Private DNS Zone (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the DNS Zone [cmdb_ci_dns_zone] table.

Azure Private Gateway pattern-based discovery

Discovery and Service Mapping Patterns finds Azure virtual network gateways on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Private Gateway (LP) pattern.

Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

Field	Description
Name [name]	The Name or ID if no Name is specified for the private gateway.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Connection Type [connection_type]	Type of VPN connection the gateway supports.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway] (continued)

Field	Description
Comments [comments]	Identifier for internal usage (deletion strategy).
Install Status [install_status]	Install status of the resource. Default value is Installed.

Virtual Private Gateway Endpoint [cmdb_ci_endpoint_vpg]

Field	Description
Name [name]	The Name or ID if no Name is specified for the endpoint.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Region [region]	The Azure region in which the endpoint is located.

CI relationships

The Azure - Private Gateway (LP) pattern creates the following relationships and references to support Azure Private Gateway discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Virtual Private Gateway Endpoint [cmdb_ci_endpoint_vpg]	Implement End Point To::Implement End Point From	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]
Cloud Network [cmdb_ci_network]	Use End Point To::Use End Point From	Virtual Private Gateway Endpoint [cmdb_ci_endpoint_vpg]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]

Azure Tag discovery

The Azure - Private Gateway (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway] table.

Azure Private Link Private Endpoint pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Private Link Private Endpoint - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Private Link - Private Endpoint [cmdb_azure_private_link_private_endpoint]

Field	Description
Kind [kind]	The specific category or variant of the resource type.
Location [location]	The geographic region where the resource is deployed.
Object Id [object_id]	The unique identifier for the resource.
Provisioning State [provisioning_state]	The current provisioning or deployment state of the resource.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Tenant ID [tenant_id]	The identifier for the tenant associated with the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Private Link Private Endpoint - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographic region where the resource is deployed.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier for the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.network/privateendpoints**.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Private Link - Private Endpoint [cmdb_azure_private_link_private_endpoint]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Public IP pattern-based discovery

Discovery and Service Mapping Patterns finds Azure public IP addresses on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Public IP (LP) pattern.

Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]

Field	Description
Name [name]	The Name or ID if no Name is specified for the IP address.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
State [state]	The current state of the IP address.
Public IP Address [public_ip_address]	The address of the Elastic IP address.
Public DNS [public_dns]	The public DNS name.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Azure - Public IP (LP) pattern creates the following relationships and references to support Azure Public IP discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Cloud Mgmt Network Interface [cmdb_ci_nic]	Contains::Contained by	Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]

Azure Tag discovery

The Azure - Public IP (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress] table.

Azure Recovery Services Vault pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Recovery Services Vault - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Recovery Services - Vault [cmdb_azure_recovery_services_vault]

Field	Description
Object Id [object_id]	The unique resource identifier for the Azure resource.
Kind [kind]	The resource kind, indicating its specific variant or implementation.
DC Location [location]	The Azure region where the resource is deployed.
Resource Group [resource_group]	The name of the resource group containing the resource.
Subscription Id [subscription_id]	The Azure subscription ID under which the resource exists.
Tenant Id [tenant_id]	The Azure Active Directory tenant ID associated with the resource.
Provisioning State [provisioning_state]	The current provisioning state of the resource. For example: Succeeded, Failed, or Updating.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Recovery Services Vault - Extended Inventory(LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The unique resource identifier for the Azure resource.
Name [name]	The name assigned to the resource.
Location [location]	The Azure region where the resource is deployed.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.recoveryservices/vaults**.

CI relationships

The Azure - Recovery Services Vault - Extended Inventory(LP) pattern creates these relationships to support Azure Recovery Services Vault discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Recovery Services - Vault [cmdb_azure_recovery_services_vault]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Recovery Services Vault Backup Item pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Recovery Services Vault Backup Item - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Recovery Services - Vault Backup Item [cmdb_azure_recovery_services_vault_backup_item]

Field	Description
Object Id [object_id]	The unique identifier of the Azure resource.
Kind [kind]	The specific kind or sub-type of the resource.
DC Location [location]	The Azure region where the resource is deployed.
Resource Group [resource_group]	The name of the resource group containing the resource.
Subscription Id [subscription_id]	The unique identifier of the Azure subscription.
Tenant Id [tenant_id]	The unique identifier of the Azure Active Directory tenant.
Backup Management Type [backup_management_type]	The method or platform by which backups are managed.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Last Backup Status [last_backup_status]	The result or outcome of the most recent backup operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Recovery Services - Vault Backup Item [cmdb_azure_recovery_services_vault_backup_item] (continued)

Field	Description
Policy Name [policy_name]	The name of the backup policy associated with the resource.
Protection State [protection_state]	The operational state of backup protection on the resource.
Protection Status [protection_status]	The current protection status assigned to the resource.
Virtual Machine Id [virtual_machine_id]	The unique identifier of the associated virtual machine.
Workload Type [workload_type]	The type of workload or application being protected by the backup.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Recovery Services Vault Backup Item - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The unique identifier of the Azure resource.
Name [name]	The name of the Azure resource.
Location [location]	The Azure region where the resource is deployed.
Comments [comments]	The unique identifier of the backup vault associated with the resource.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.recoveryservices/vaults/backupfabrics/protectioncontainers/protectedItems**.

CI relationships

The Azure - Recovery Services Vault Backup Item - Extended Inventory(LP) pattern creates these relationships to support Azure Recovery Services Vault Backup Item discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Recovery Services - Vault Backup Item [cmdb_azure_recovery_services_vault_backup_item]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure Resource Group pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Resource Groups in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Resource Group (LP) pattern.

Resource Group [cmdb_ci_resource_group]

Field	Description
Name [name]	The Name or ID if no Name is specified for the resource group.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
State [state]	The current state of the resource group.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	References the Azure Datacenter [cmdb_ci_azure_datacenter] table.

CI relationships and references

The Azure - Resource Group (LP) pattern creates the following relationships and references to support Azure Resource Group discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Azure Datacenter [cmdb_ci_azure_datacenter]	Contains::Contained by	Resource Group [cmdb_ci_resource_group]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Resource Group [cmdb_ci_resource_group]
Resource Group [cmdb_ci_resource_group]	Location [location]	Azure Datacenter [cmdb_ci_azure_datacenter]

Azure Tag discovery

The Azure - Resource Group (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Resource Group [cmdb_ci_resource_group] table.

Microsoft Azure resource inventory discovery

The ServiceNow Discovery application uses the Azure Resource Inventory (LP) pattern to find resources available through Azure that don't have a dedicated pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

You can discover up to 10 thousand resources using the Azure Resource Inventory (LP) pattern.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Azure service principal

Verify the Azure service principal is configured. For more information, see Discovery for Microsoft Azure

Privileges

For versions 1.3.0 and later:

On the Microsoft Azure Console, provide a user with the API Management Service Reader role with permissions to run the following REST POST request:

```
"https://management.<resource url>/providers/
microsoft.resourcegraph/resources?api-version=2022-10-01"
```

<resource url> – The URL varies based on the Azure account type, like "azure.com" or government site address.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For versions 1.0.99 and earlier:

On the Microsoft Azure Console, provide a user with the API Management Service Reader role with permissions to run the following REST GET request:

`"https://management.azure.com/subscriptions/<subscription id>/resources"`

Create a Cloud Discovery schedule

Create a schedule for the relevant Azure service account as described in Create Discovery schedules for cloud resources

Fine-tune Azure inventory discovery using the Resource Inclusion List

If your deployment has custom patterns for Azure discovery, ensure that you don't discover Azure resources twice.

Note: Leaving the [pre_resource_whitelist] table empty causes the Pattern to populate all resources and it might cause performance issues.

1. Ensure that the application scope is Discovery and Service Mapping Patterns:
   a. Navigate to **Settings** > **Developer**.
   b. Select `Discovery and Service Mapping Patterns` from the **Application** list.
2. Navigate to **System Definitions** > **Tables**.
3. Open the Cloud Inventory Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table.
4. Under **Related Links**, click **Show List**.
5. Select resource types for which you have custom patterns, and select `Delete` from the **Actions on selected rows** list.

The Cloud Inventory Resource Inclusion List is predefined with common services. You can expand the inclusion list with additional resource types that you want the pattern to discover. The names of these resource types must conform to the appropriate vendor naming conventions.

Note: When you modify the out-of-the-box inclusion list, it is no longer updated automatically in application updates. You must maintain the customized list on your own.

1. Open the Cloud Inventory Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table.
2. Click **New**.
3. Fill in the form, and then click **Submit**.

Field	Description
Cloud Vendor	The vendor of the resource type: Azure.
Resource Type	The Azure resource type value.
Application	The application scope: Discovery and Service Mapping Patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The changes are applied the next time you run the pattern.

Modify the ServiceNow related entries

1. Navigate to **Related Entries**.
2. Click **New**.
3. For the identifier, enter **cmdb_ci_cmp_resource**.
4. For the related table, select **Key Value [cmdb_key_value]**.
5. For the referenced field, select **Configuration Item**.
6. For the criterion attributes, add **Key** and **Value**.

   You may need to click the lock icon to view and change the criterion attributes.

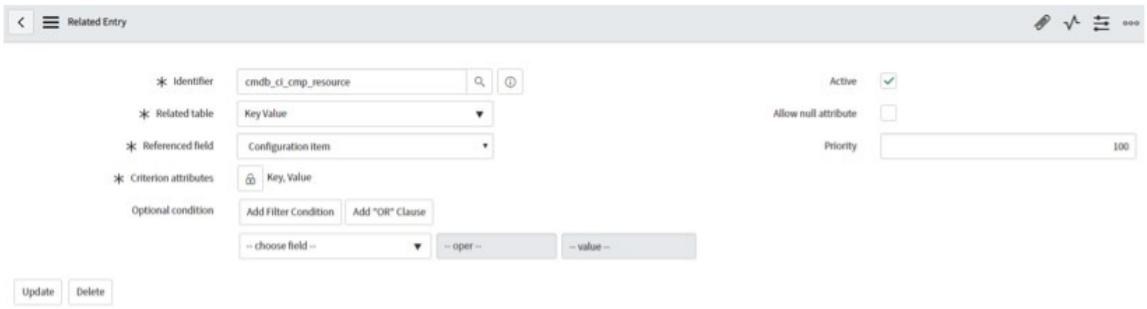

Data collected by Discovery during horizontal discovery

Table and field	Additional information
Cloud Resource [cmdb_ci_cmp_resource]	
Name	A descriptive name used to identify the Azure resource.
Resource type	The type of the resource, for example network, load balancer.
Object ID	The ID of the Azure resource.
Description	The text populated during discovery that specifies that this resource is available from Azure.
Key Value [cmdb_key_value]	
Key	The Azure tag key name.
Value	The Azure tag value name.

CI relationships

These relationships are created to support Azure resource inventory discovery:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on:Hosts	Azure datacenter [cmdb_ci_azure_datacenter]

Related topics

https://learn.microsoft.com/en-us/azure/governance/resource-graph/reference/supported-tables-resources

Azure Route Table pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Route Tables in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Route Table (LP) pattern.

Route Table [cmdb_ci_route_table]

Field	Description
Name [name]	The Name or ID if no Name is specified for the route table.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
State [state]	The current state of the route table.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Comments [comments]	Identifier for internal usage (deletion strategy).

Route [cmdb_ci_route]

Field	Description
Name [name]	The Name or ID if no Name is specified for the route.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Destination [destination]	The datacenter location of the route.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships and references

The Azure - Route Table (LP) pattern creates the following relationships and references to support Azure Route Table discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Route Table [cmdb_ci_route_table]	Contains::Contained by	Route [cmdb_ci_route]
Route Table [cmdb_ci_route_table]	Contains::Contained by	Cloud Network [cmdb_ci_network]
Azure Datacenter [cmdb_ci_azure_datacenter]	Contains::Contained by	Route Table [cmdb_ci_route_table]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Route Table [cmdb_ci_route_table]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Route Table [cmdb_ci_route_table]

Azure Tag discovery

The Azure - Route Table (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Route Table [cmdb_ci_route_table] table.

Azure Security Group pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Security Groups on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Security Group (LP) pattern.

Compute Security Group [cmdb_ci_compute_security_group]

Field	Description
Name [name]	The Name or ID if no Name is specified for the security group.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
State [state]	The current state of the security group.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships and references

The Azure - Security Group (LP) pattern creates the following relationships and references to support Azure Security Group discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Compute Security Group [cmdb_ci_compute_security_group]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Compute Security Group [cmdb_ci_compute_security_group]
Cloud Network [cmdb_ci_network]	Contains::Contained by	Compute Security Group [cmdb_ci_compute_security_group]
Cloud Subnet [cmdb_ci_cloud_subnet]	Uses::Used by	Compute Security Group [cmdb_ci_compute_security_group]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Compute Security Group [cmdb_ci_compute_security_group]

Azure Tag discovery

The Azure - Security Group (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Value [value]	Tag value.
Configuration item [configuration_item]	References the Compute Security Group [cmdb_ci_compute_security_group] table.

Azure Service Bus Namespace pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Service Bus Namespace - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

Azure Service Bus - Namespace [cmdb_azure_service_bus_namespace]

Field	Description
Identity [identity]	The identity information associated with the resource, often used for authentication or access control.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Service Bus - Namespace [cmdb_azure_service_bus_namespace] (continued)

Field	Description
Location [location]	The geographical location where the resource is hosted.
Object Id [object_id]	The unique identifier of the resource.
Resource Group [resource_group]	Name of the resource group.
Status [status]	The current status or state of the resource.
Subscription ID [subscription_id]	The subscription ID.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Service Bus Namespace - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical location where the resource is hosted.
Name [name]	The name of the resource.
Object ID [object_id]	The unique identifier of the resource.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.servicebus/namespaces**.
Description [short_description]	Description of the resource or service.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Service Bus - Namespace [cmdb_azure_service_bus_namespace]	References	Cloud Resource [cmdb_ci_cmp_resource]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Service Bus Queue pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Service Bus Queue - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Service Bus - Queue [cmdb_azure_service_bus_queue]

Field	Description
Enable Batched Operations [enable_batched_operations]	Indicates whether batched operations are enabled for the queue.
Enable Partitioning [enable_partitioning]	Indicates whether message partitioning is enabled for load distribution across multiple message brokers.
DC Location [location]	The geographic region where the Service Bus queue is deployed.
Max Message Size In KB [max_message_size_in_kb]	The maximum allowed size of a single message in kilobytes (KB).
Max Size In MB [max_size_in_mb]	The maximum storage size of the queue in megabytes (MB).
Object Id [object_id]	The unique identifier for the Service Bus queue.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Status [status]	The current status of the Service Bus queue.
Configuration Item [configuration_item]	References the Cloud Messaging Service [cmdb_ci_cloud_messaging_service] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Service Bus Queue - Extended Inventory(LP) pattern.

Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographic region where the Service Bus queue is deployed.
Name [name]	The name of the Service Bus queue.
Object ID [object_id]	The unique identifier for the Service Bus queue.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Type [type]	Type of resource. The value is set to **Microsoft.ServiceBus/namespaces/queues**.

CI relationships

The pattern creates these relationships to support discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Contains::Contained by	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]
Cloud Messaging Service [cmdb_ci_cloud_messaging_service]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Service Bus - Queue [cmdb_azure_service_bus_queue]	References	Cloud Messaging Service [cmdb_ci_cloud_messaging_service]

Azure Service Bus Topic pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Service Bus Topic - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Service Bus - Topic [cmdb_azure_service_bus_topic]

Field	Description
Enable Batched Operations [enable_batched_operations]	Indicates whether batched operations are enabled for the topic.
Enable Partitioning [enable_partitioning]	Specifies if the topic is partitioned to support scalability.
DC Location [location]	The geographical region where the topic is deployed.
Max Message Size In KB [max_message_size_in_kb]	The maximum size of any message that can be sent to the topic, measured in kilobytes (KB).
Max Size In MB [max_size_in_mb]	The maximum storage limit for the topic in megabytes (MB).
Object Id [object_id]	The unique identifier for the topic resource.
Status [status]	The status of the topic resource. For example: Active or Disabled.
Subscription Count [subscription_count]	The number of active subscriptions associated with the topic.
Resource Group [resource_group]	Name of the resource group.
Subscription ID [subscription_id]	The subscription ID.
Configuration Item [configuration_item]	References the Cloud Topic [cmdb_ci_cloud_topic] table.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Service Bus Topic - Extended Inventory(LP) pattern.

Cloud Topic [cmdb_ci_cloud_topic]

Field	Description
Install Status [install_status]	Install status of the resource. Default value is Installed.
Location [location]	The geographical region where the topic is deployed.
Name [name]	The name of the topic within the Service Bus namespace.
Object ID [object_id]	The unique identifier of the resource in the following format: **/subscriptions/{subscriptionId}/resourceGroups/{resourceGroupName}/providers/Microsoft.ServiceBus/namespaces/{namespaceName}/topics/{topicName}**
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Description [short_description]	Description of the resource or service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The pattern creates these relationships to support discovery.

CI	Relationship	CI
Cloud Resource [cmdb_ci_cmp_resource]	Contains::Contained by	Cloud Topic [cmdb_ci_cloud_topic]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Topic [cmdb_ci_cloud_topic]
Cloud Topic [cmdb_ci_cloud_topic]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Service Bus - Topic [cmdb_azure_service_bus_topic]	References	Cloud Topic [cmdb_ci_cloud_topic]

Azure Service Endpoint Policy pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB tables when running the Azure - Service Endpoint Policy - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Service Endpoint Policy [cmdb_azure_service_endpoint_policy_service_endpoint_policy]

Field	Description
Object Id [object_id]	The unique identifier of the Azure resource.
Kind [kind]	The specific kind or stock-keeping unit (SKU) of the resource.
DC Location [location]	The geographic region where the resource is deployed.
Resource Group [resource_group]	The name of the resource group containing the resource.
Subscription Id [subscription_id]	The Globally Unique Identifier (GUID) of the subscription that owns the resource.
Tenant Id [tenant_id]	The GUID of the Azure AD tenant the resource belongs to.
Provisioning State [provisioning_state]	The current provisioning status of the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.

Azure Service Endpoint Policy Definition [cmdb_azure_service_endpoint_policy_definition]

Field	Description
Object Id [object_id]	The unique identifier of the Azure resource.
Name [name]	The name assigned to the resource.
Type [type]	Type of resource. The value is set to **microsoft.network/serviceendpointpolicies**.
Provisioning State [provisioning_state]	The current provisioning status of the resource.
Configuration Item [configuration_item]	References the Cloud Resource [cmdb_ci_cmp_resource] table.
Service [service]	The specific Azure service related to the resource.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Service Endpoint Policy - Extended Inventory(LP) pattern.

Cloud Resource [cmdb_ci_cmp_resource]

Field	Description
Object ID [object_id]	The unique identifier of the Azure resource.
Name [name]	The name assigned to the resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Resource [cmdb_ci_cmp_resource] (continued)

Field	Description
Location [location]	The geographic region where the resource is deployed.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.
Resource type [resource_type]	Type of resource. The value is set to **microsoft.network/serviceendpointpolicies**.

CI relationships

The Azure - Service Endpoint Policy - Extended Inventory(LP) pattern creates these relationships to support Azure Service Endpoint Policy discovery.

CI	Relationship	CI
Azure Service Endpoint Policy Definition [cmdb_azure_service_endpoint_policy_definition]	References	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Subnet [cmdb_ci_cloud_subnet]	Uses::Used by	Cloud Resource [cmdb_ci_cmp_resource]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Resource [cmdb_ci_cmp_resource]
Cloud Resource [cmdb_ci_cmp_resource]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Service Endpoint Policy [cmdb_azure_service_endpoint_policy_service_endpoint_policy]	References	Cloud Resource [cmdb_ci_cmp_resource]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure SQL Server pattern-based discovery

Discovery and Service Mapping Patterns finds Azure SQL Server virtual machines in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - SQL Server (LP) pattern.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.

CI relationships and references

The Azure - SQL Server (LP) pattern creates the following relationships and references to support Azure SQL Server discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Virtual Machine Instance [cmdb_ci_vm_instance]

Azure License discovery

The pattern extension section discovers SQL Server license information and populates the license type in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	SQL_Server_License_Type_automatic
Value [value]	The license model, which is one of the following: • Azure Hybrid Benefit: **BYOL** • Pay-as-you-go licensing: **License Included**
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Storage Account pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Storage Accounts in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Storage Account (LP) pattern.

Cloud Storage Account [cmdb_ci_cloud_storage_account]

Field	Description
Name [name]	The Name or ID if no Name is specified for the storage account.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Sku Name [sku_name]	The name of the stock keeping unit (SKU) for backup frequency.
State [state]	The current state of the storage account. Possible values are: available, starting, error, cancelled, or off.
Blob Service [blob_service]	Indicates whether the storage account has a Blob service endpoint. If a blob endpoint exists, the value is set to true.
File Service [file_service]	Indicates whether the storage account has a File service endpoint. If a file endpoint exists, the value is set to true.
Table Service [table_service]	Indicates whether the storage account has a Table service endpoint. If a table endpoint exists, the value is set to true.
Queue Service [queue_service]	Indicates whether the storage account has a Queue service endpoint. If a queue endpoint exists, the value is set to true.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Cloud Object Service [cmdb_ci_cloud_object_service]

Field	Description
Name [name]	Name of the blob service. The value is set to **default**.
Object ID [object_id]	A unique identifier for this blob service.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud File Service [cmdb_ci_cloud_file_service]

Field	Description
Name [name]	Name of the file service. The value is set to **default**.
Object ID [object_id]	A unique identifier for this file service.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Cloud Table Service [cmdb_ci_cloud_table_service]

Field	Description
Name [name]	Name of the table service. The value is set to **default**.
Object ID [object_id]	A unique identifier for this table service.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Cloud Queue Service [cmdb_ci_cloud_queue_service]

Field	Description
Name [name]	Name of the queue service. The value is set to **default**.
Object ID [object_id]	A unique identifier for this queue service.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships and references

The Azure - Storage Account (LP) pattern creates the following relationships and references to support Azure Storage Account discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Cloud Storage Account [cmdb_ci_cloud_storage_account]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Cloud Storage Account [cmdb_ci_cloud_storage_account]
Cloud Storage Account [cmdb_ci_cloud_storage_account]	Contains::Contained by	Cloud Object Service [cmdb_ci_cloud_object_service]*
Cloud Storage Account [cmdb_ci_cloud_storage_account]	Contains::Contained by	Cloud File Service [cmdb_ci_cloud_file_service]*
Cloud Storage Account [cmdb_ci_cloud_storage_account]	Contains::Contained by	Cloud Table Service [cmdb_ci_cloud_table_service]*
Cloud Storage Account [cmdb_ci_cloud_storage_account]	Contains::Contained by	Cloud Queue Service [cmdb_ci_cloud_queue_service]*

* These relationships are created only if the corresponding service endpoint exists on the storage account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Storage Account [cmdb_ci_cloud_storage_account]

Azure Tag discovery

The Azure - Storage Account (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Storage Account [cmdb_ci_cloud_storage_account] table.

Azure Storage Containers pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Storage Containers in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Storage Container(LP) pattern.

Storage Container [cmdb_ci_storage_container]

Field	Description
Name [name]	The name of the storage container.
Object ID [object_id]	A unique identifier for the storage container.
State [state]	The lease state of the storage container.
Comments [comments]	Identifier for internal usage (deletion strategy).
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

The Azure - Storage Container(LP) pattern creates the following relationships to support Azure Storage Containers discovery.

CI relationships

CI	Relationship	CI
Cloud Object Service [cmdb_ci_cloud_object_service]	Contains::Contained by	Storage Container [cmdb_ci_storage_container]

Azure Storage Volume pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Storage Volumes (managed disks) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

(Optional) Exclude temporary Azure Databricks VMs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can reduce short-lived configuration item (CI) records by excluding temporary Azure Databricks VMs. For more information, see Exclude temporary Azure Databricks virtual machines.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Storage Volume (LP) pattern.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	The Name or ID if no Name is specified for the storage volume.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Volume ID [volume_id]	The resource ID of the storage volume.
State [state]	The current state of the storage volume. Possible values are: Available or Leased.
Storage type [storage_type]	The storage type. The value is set to **PageBlob**.
Size [size]	The size of the volume in GB.
Size bytes [size_bytes]	The size of the volume in bytes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume [cmdb_ci_storage_volume] (continued)

Field	Description
Comments [comments]	Identifier for internal usage (deletion strategy).
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships and references

The Azure - Storage Volume (LP) pattern creates the following relationships and references to support Azure Storage Volume discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Storage Volume [cmdb_ci_storage_volume]

Azure Tag discovery

The Azure - Storage Volume (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Storage Volume [cmdb_ci_storage_volume] table.

Azure sub account pattern-based discovery

Discovery and Service Mapping Patterns finds Azure subscriptions and management groups on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Management Groups data model

The Azure Management Groups pattern extension introduces the following CI class that extends an existing CMDB class.

CI class introduced by the Azure Management Groups pattern extension

CI class	Extends from
Azure Management Group [cmdb_ci_azure_management_group]	Configuration Item [cmdb_ci]

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Sub Account (LP) pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Name [name]	Name of the subscription or service account.
Object ID [object_id]	The subscription ID, allocated by Azure for this resource.
Account Id [account_id]	The subscription ID of the service account.
Datacenter Type [datacenter_type]	Type of datacenter associated with the service account. The value is set to **cmdb_ci_azure_datacenter**.
Datacenter URL [datacenter_url]	The URL of the datacenter associated with the service account.
Is management account [is_master_account]	Boolean attribute indicating whether this is the management account. The value is set to **false** for discovered sub-accounts.
Discovery credentials [discovery_credentials]	Reference field to the related Azure credentials.
Parent account [parent_account]	Reference to the primary account, if it exists.
Operational status [operational_status]	Operational status of the subscription, derived from its Azure state. For example: Operational or Non-Operational.
Life Cycle Stage [life_cycle_stage]	Life cycle stage of the subscription, derived from its Azure state. For example: Operational, End of Operation, or End of Life.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Service Account [cmdb_ci_cloud_service_account] (continued)

Field	Description
Life Cycle Stage Status [life_cycle_stage_status]	Life cycle stage status of the subscription, derived from its Azure state. For example: In Use, Action Required, On Hold, Billing Past Due, or Retired.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Discovery populates the data in the CMDB when running the Azure Management Groups pattern extension.

Azure Management Group [cmdb_ci_azure_management_group]

Field	Description
Name [name]	Management group name.
Object ID [object_id]	Management group name and tenant ID in the following format: `name+@+tenantId`. For example: itomMgmtGroup@8bcff-vdc-btrv.
Parent [parent]	References the direct parent Azure Management Group [cmdb_ci_azure_management_group] table.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Cloud Organizations [cmdb_ci_cloud_org]

Field	Description
Name [name]	Tenant ID or name. • Tenant ID: When using management-level credentials • Tenant name: When using tenant-level credentials
Object ID [object_id]	Tenant ID.
DNS Domain [dns_domain]	Domain name entered during registration, for example: servicenow.com. This field is populated only when tenant-level credentials are used.

CI relationships and references

The Azure - Sub Account (LP) pattern creates the following references to support Azure Sub Account discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud Service Account [cmdb_ci_cloud_service_account]

The Azure Management Groups pattern extension creates the following relationships and references to support Azure Sub Account discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Azure Management Group [cmdb_ci_azure_management_group]	Contains::Contained by	Cloud Service Account [cmdb_ci_cloud_service_account]
Cloud Organizations [cmdb_ci_cloud_org]	Contains::Contained by	Azure Management Group [cmdb_ci_azure_management_group]
Azure Management Group [cmdb_ci_azure_management_group]	Contains::Contained by	Azure Management Group [cmdb_ci_azure_management_group]

CI references

CI	Field	Referenced CI
Azure Management Group [cmdb_ci_azure_management_group]	Parent [parent]	Azure Management Group [cmdb_ci_azure_management_group]*

* Only references the direct parent-child management group relationship.

Azure Tag discovery

The Azure - Sub Account (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table.

Azure Subscriptions Discovery For Management Group pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Subscription entities under Management Groups on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure Subscriptions Discovery For Management Group pattern.

Cloud Service Account [cmdb_ci_cloud_service_account]

Field	Description
Account Id [account_id]	Subscription name or account ID. • Subscription account: The Azure subscription name • Management account: The management group account ID
Name [name]	Name of the account. • Subscription account: The display name of the subscription • Management account: The service account name
Discovery credentials [discovery_credentials]	Azure credential from the Credentials [discovery_credentials] table.
Datacenter URL [datacenter_url]	The Azure management endpoint URL for the account.
Datacenter Type [datacenter_type]	The type of datacenter: **Azure Datacenter [cmdb_ci_azure_datacenter]**.
Parent account [parent_account]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table (Management account).
Is management account [is_master_account]	Indicates whether this is a management group account. • **true**: Management account • **false**: Subscription account
Object ID [object_id]	Subscription name or account ID. • Subscription account: The Azure subscription name • Management account: The management group account ID

CI relationships

The Azure Subscriptions Discovery For Management Group pattern creates these references to support Azure Subscriptions Discovery For Management Group discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Cloud Service Account [cmdb_ci_cloud_service_account] (Subscription account)	References	Cloud Service Account [cmdb_ci_cloud_service_account] (Management account)
Key Value [cmdb_key_value]	References	Cloud Service Account [cmdb_ci_cloud_service_account] (Subscription account)

Azure Tag discovery

The Azure Subscriptions Discovery For Management Group pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud Service Account [cmdb_ci_cloud_service_account] table (Subscription account).

Azure virtual machine pattern-based discovery

Discovery and Service Mapping Patterns finds Azure virtual machines (VMs) on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

(Optional) Exclude temporary Azure Databricks VMs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can reduce short-lived configuration item (CI) records by excluding temporary Azure Databricks VMs. For more information, see Exclude temporary Azure Databricks virtual machines.

Create Server CIs during cloud discovery without running IP-based discovery

Starting with Discovery and Service Mapping Patterns version 1.31.0, you can create Server CIs during cloud discovery without running IP-based discovery. For more information, see Server CI population during cloud discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Virtual Machine (LP) pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Name [name]	Name of the virtual machine.
Object ID [object_id]	A unique identifier allocated by Azure for this virtual machine.
VM Instance ID [vm_inst_id]	The unique virtual machine identifier assigned by Azure.
State [state]	Power state of the virtual machine. Possible values include: On, Off, or Terminated.
Memory (MB) [memory]	Memory capacity of the virtual machine in megabytes (MB).
CPUs [cpus]	Number of virtual CPUs (vCPUs) allocated to the virtual machine.
Disks [disks]	Total number of disks attached to the virtual machine, including OS disk and data disks.
Disk size [disks_size]	Total size of all disks attached to the virtual machine in gigabytes (GB).
Network adapters [nics]	Number of network interface cards (NICs) attached to the virtual machine.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Instance Scale Set [cmdb_ci_instance_scale_set]

Field	Description
Object ID [object_id]	A unique identifier allocated by Azure for this Virtual Machine Scale Set instance.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Object ID [object_id]	A unique identifier for the managed disk.
Name [name]	Name of the managed disk.

Block Endpoint [cmdb_ci_endpoint_block]

Field	Description
Object ID [object_id]	A unique identifier for the disk endpoint.
Name [name]	Name of the disk endpoint.
Host [host]	The virtual machine object ID to which this disk endpoint is attached.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Mgmt Network Interface [cmdb_ci_nic]

Field	Description
Object ID [object_id]	A unique identifier for the NIC.
Primary [primary]	Indicates whether this NIC is the primary network interface.

VNIC Endpoint [cmdb_ci_endpoint_vnic]

Field	Description
Object ID [object_id]	A unique identifier for the virtual network interface endpoint.
Name [name]	Name of the virtual network interface endpoint.

Availability Zone [cmdb_ci_availability_zone]

Field	Description
Object ID [object_id]	A unique identifier for the availability zone in the format: region-zone.
Name [name]	A unique identifier for the availability zone in the format: region-zone.
Install Status [install_status]	Install status of the resource. Default value is Installed.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Object ID [object_id]	A unique identifier for the OS image or custom image.
Name [name]	Name of the OS image or custom image.
Provider [provider]	Cloud provider. The value is set to **AZURE**. This field is only populated in the Cloud Image [cmdb_ci_cloud_os_image] table.

Note: When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Object ID [object_id]	The VM size Stock Keeping Unit (SKU) name.
Provider [provider]	Cloud provider. The value is set to **AZURE**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type] (continued)

Field	Description
	This field is only populated in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

Note: When using the Hardware Type [cmdb_ci_compute_template] table to store the hardware types, you may notice an unusually large number of records. To avoid this issue, you can store the discovered hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. For more information, see Enable the Cloud Hardware Type class extension.

CI relationships

The Azure - Virtual Machine (LP) pattern creates these relationships to support Azure VM discovery.

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Managed by::Manages	Instance Scale Set [cmdb_ci_instance_scale_set]
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Hardware Type [cmdb_ci_compute_template] or Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Block Endpoint [cmdb_ci_endpoint_block]	Implement End Point To::Implement End Point From	Storage Volume [cmdb_ci_storage_volume]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	Block Endpoint [cmdb_ci_endpoint_block]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	VNIC Endpoint [cmdb_ci_endpoint_vnic]
VNIC Endpoint [cmdb_ci_endpoint_vnic]	Implement End Point To::Implement End Point From	Cloud Mgmt Network Interface [cmdb_ci_nic]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Virtual Machine Instance [cmdb_ci_vm_instance]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]
Image [cmdb_ci_os_template]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
or Cloud Image [cmdb_ci_cloud_os_image]		
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Virtual Machine Instance [cmdb_ci_vm_instance]
Azure Datacenter [cmdb_ci_azure_datacenter]	Contains::Contained by	Availability Zone [cmdb_ci_availability_zone]
Availability Set [cmdb_ci_availability_set]	Contains::Contained by	Virtual Machine Instance [cmdb_ci_vm_instance]
Key Value [cmdb_key_value]	References	Virtual Machine Instance [cmdb_ci_vm_instance]

Azure Tag discovery

The Azure - Virtual Machine (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

Azure BYOL discovery

The pattern extension section discovers the following license types for Windows and Red Hat Enterprise Linux (RHEL) VMs:

- Bring Your Own License (BYOL)
- License included

The pattern stores the license type and model in the Key Value [cmdb_key_value] table.

Note:

- For Windows VM using a custom image, Discovery populates key values only when the primary OS disk details are available.
- For RHEL VMs, Discovery populates information only for VMs that use Marketplace images.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The license type key, which is one of the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
	• Windows_OS_License_Type_automatic • RHEL_OS_License_Type_automatic
Value [value]	The license model, which is one of the following: • BYOL • License Included
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

Exclude temporary Azure Databricks virtual machines

Reduce short-lived configuration item (CI) records by excluding temporary Azure Databricks virtual machines (VMs).

Before you begin

Verify that you have at least version 1.30.2 of Discovery and Service Mapping Patterns.

Role required: admin

About this task

By default, Discovery includes temporary Azure Databricks VMs in discovery results. The *glide.discovery.cloud.filterout_databrick_vms* system property controls whether these VMs are excluded from the CMDB. The default value is **false**. When set to true, Discovery excludes Azure Databricks VMs during pattern-based discovery, reducing the number of short-lived CI records in your CMDB.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `glide.discovery.cloud.filterout_databrick_vms`.
3. Select the **glide.discovery.cloud.filterout_databrick_vms** system property.
4. In the **Value** field, enter **true**.
5. Select **Update**.

Related topics

Azure virtual machine pattern-based discovery

Microsoft Azure Cloud discovery using patterns

Azure Virtual Machine Scale Sets (VMSS) Instance discovery

The Discovery and Service Mapping Patterns application uses the Azure - VM Scale Set (LP) and the Azure VM Instance - Uniform Scale Set patterns to find Azure Virtual Machine Scale Sets (VMSS). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMSS let you create and manage a group of load balanced virtual machines (VMs). The number of VM instances can automatically increase or decrease in response to demand or a defined schedule. There Uniform Orchestration modes Flexible and achieve.

- VMSS with Flexible Orchestration achieves high availability at scale with identical or multiple virtual machine types. With Flexible orchestration, Azure provides a unified experience across the Azure VM ecosystem.
- VMSS with Uniform Orchestration optimized for large-scale stateless workloads with identical instances. VMSS with Uniform orchestration use a virtual machine profile or template to scale up to desired capacity. While there is some ability to manage or customize individual virtual machine instances, Uniform uses identical VM instances.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about VMSS and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

Prerequisites

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Create Server CIs during cloud discovery without running IP-based discovery

Starting with Discovery and Service Mapping Patterns version 1.31.0, you can create Server CIs during cloud discovery without running IP-based discovery. For more information, see Server CI population during cloud discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - VM Scale Set (LP) Pattern.

Instance Scale Set [cmdb_ci_instance_scale_set]

Field	Description
Install Status [install_status]	Install status of the instance scale set. Default value is Installed.
State [state]	Provisioning state of the instance scale set. Possible values are Available or Terminated.
Object ID [object_id]	A unique identifier for the instance scale set.
Name [name]	Name of the instance scale set.

Discovery populates the data in the CMDB when running the Azure VM Instance - Uniform Scale Set pattern.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Network adapters [nics]	Number of Network Interface Cards (NICs) attached to the virtual machine (VM) instance.
State [state]	Current state of the instance. For example: On, Off, or Terminated.
Install Status [install_status]	Install status of the instance. Default value is Installed.
Operational status [operational_status]	Operational status of the instance. Default value is Operational.
Disks [disks]	Number of disks attached to the instance.
Disks size (GB) [disk_size]	Total storage capacity of all disks attached to the VM, in gigabytes (GB).

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	Name or ID for the storage volume.
Object ID [object_id]	A unique identifier, allocated by Microsoft Azure for this resource.
Install Status [install_status]	Install status of the storage volume. Default value is Installed.
Operational status [operational_status]	Operational status of the storage volume. Default value is Operational.

Cloud Mgmt Network Interface [cmdb_ci_nic]

Field	Description
Name [name]	Name or ID for the network interface.
Object ID [object_id]	A unique identifier, allocated by Microsoft Azure for this resource.
Install Status [install_status]	Install status of the network interface. Default status is Installed.
Operational status [operational_status]	Operational status of the network interface. Default status is Operational.

VNIC Endpoint [cmdb_ci_endpoint_vnic]

Field	Description
Name [name]	Name or ID for the endpoint.
Object ID [object_id]	A unique identifier, allocated by Microsoft Azure for this resource.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Image [cmdb_ci_os_template]

Field	Description
Name [name]	Name or ID for the image.
Object ID [object_id]	A unique identifier, allocated by Microsoft Azure for this resource.
Install Status [install_status]	Install status of the image. Default value is Installed.
Operational status [operational_status]	Operational status of the image. Default value is operational.

Note: When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

Hardware Type [cmdb_ci_compute_template]

Field	Description
Name [name]	Name of the hardware type.

Note: When using the Hardware Type [cmdb_ci_compute_template] table to store the hardware types, you may notice an unusually large number of records. To avoid this issue, you can store the discovered hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table. For more information, see Enable the Cloud Hardware Type class extension.

CI relationships

Discovery creates these relationships to support the VMSS discovery, when running the Azure - VM Scale Set (LP) pattern.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Instance Scale Set [cmdb_ci_instance_scale_set]
Instance Scale Set [cmdb_ci_instance_scale_set]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Instance Scale Set [cmdb_ci_instance_scale_set]	Consumes::Consumed by	Load Balancer Pool [cmdb_ci_lb_pool]

Discovery creates these relationships to support the VMSS discovery, when running the when running the Azure - VM Scale Set (LP) pattern.

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Managed by::Manages	Instance Scale Set [cmdb_ci_instance_scale_set]
Storage Volume [cmdb_ci_storage_volume]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
Block Endpoint [cmdb_ci_endpoint_block]	Implement End Point To::Implement End Point From	Storage Volume [cmdb_ci_storage_volume]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	Block Endpoint [cmdb_ci_endpoint_block]
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Hardware Type [cmdb_ci_compute_template]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Hardware Type [cmdb_ci_compute_template]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Cloud Hardware Type [cmdb_ci_cloud_hardware_type]
Cloud Mgmt Network Interface [cmdb_ci_nic]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Virtual Machine Instance [cmdb_ci_vm_instance]	Implement End Point To::Implement End Point From	Cloud Mgmt Network Interface [cmdb_ci_nic]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	VNIC Endpoint [cmdb_ci_endpoint_vnic]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Image [cmdb_ci_os_template]
Image [cmdb_ci_os_template]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

Tags discovery

Both the Azure - VM Scale Set (LP) and the Azure VM Instance - Uniform Scale Set patterns collect tags and populate them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure Virtual Network Gateway Connection pattern-based discovery

Discovery and Service Mapping Patterns finds Azure Virtual Network Gateway Connection resources on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Azure discovery prerequisites section in Microsoft Azure Cloud discovery using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure - Virtual Network Gateway Connection (LP) pattern.

Virtual Network Gateway Connection [cmdb_ci_vpc_gateway_connection]

Field	Description
Object ID [object_id]	The unique Azure resource identifier for the Virtual Network Gateway Connection.
Name [name]	The name of the Virtual Network Gateway Connection resource.
State [state]	The operational state of the connection.
Protocol [protocol]	The connection protocol used by the gateway connection.
Connection Status [connection_status]	The status of the gateway connection. For example: Connected, Disconnected, or NotConnected.
Connection Type [connection_type]	The type of gateway connection. For example: IPsec, ExpressRoute, or Vnet2Vnet.
Install Status [install_status]	Install status of the resource. Default value is Installed.

CI relationships

The Azure - Virtual Network Gateway Connection (LP) pattern creates these relationships to support Azure Virtual Network Gateway Connection discovery.

CI	Relationship	CI
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Virtual Network Gateway Connection [cmdb_ci_vpc_gateway_connection]
Virtual Network Gateway Connection [cmdb_ci_vpc_gateway_connection]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Virtual Private Gateway [cmdb_ci_virtual_pvt_gateway]	Connected by::Connects	Virtual Network Gateway Connection [cmdb_ci_vpc_gateway_connection]
Key Value [cmdb_key_value]	References	Virtual Network Gateway Connection [cmdb_ci_vpc_gateway_connection]

Azure Tag discovery

The Azure - Virtual Network Gateway Connection (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Virtual Network Gateway Connection [cmdb_ci_vpc_gateway_connection] table.

Azure Web Application Firewall Policy pattern-based discovery

Discovery and Service Mapping Patterns finds Azure services on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Remove resources from the Resource Inclusion List table

Verify that the relevant resource isn't listed in the Resource Inclusion List [sa_cloud_inventory_resource_whitelist] table to avoid duplicate discovery. For more information on removing resources from the Resource Inclusion List, see Microsoft Azure resource inventory discovery.

Enable the relevant pattern

The pattern for this service is disabled by default. Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value. For more information on enabling patterns, see Activate a disabled pattern.

Configure the Discovery schedule to support GovCloud

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Discovery and Service Mapping Patterns application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

Discovery and Service Mapping Patterns application populates data in the non-CMDB table when running the Azure - Web Application Firewall Policy - Extended Inventory(LP) pattern.

You can review the non-CMDB Azure tables by navigating to **All** > **Configuration** > **Azure**. You can also search the navigation filter for the specific pattern name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Web Application Firewall - Policy [cmdb_azure_web_application_firewall_policy]

Field	Description
Object Id [object_id]	The unique identifier of the resource.
Kind [kind]	The specific kind or variant of the resource.
DC Location [location]	The Azure region where the resource is deployed.
Resource Group [resource_group]	The name of the resource group containing the resource.
Subscription Id [subscription_id]	The subscription ID associated with the resource.
Tenant Id [tenant_id]	The Azure Active Directory tenant ID associated with the resource.
Provisioning State [provisioning_state]	The latest provisioning status of the resource.
Configuration Item [configuration_item]	References the Web ACL [cmdb_ci_web_acl] table.
Policy State [policy_state]	The current state of the policy applied to the resource.

Data stored in CMDB tables

Discovery and Service Mapping Patterns application populates data in the CMDB when running the Azure - Web Application Firewall Policy - Extended Inventory(LP) pattern.

Web ACL [cmdb_ci_web_acl]

Field	Description
Object ID [object_id]	The unique identifier of the resource.
Name [name]	The name of the resource.
Location [location]	The Azure region where the resource is deployed.
Default Action [default_action]	Indicates the effective operation mode of the web application firewall (WAF) policy. Possible values are allow and detect. • allow: policy is running in **Prevention** mode and is blocking traffic • detect: policy is running in **Detection** mode and is logging only
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Web ACL [cmdb_ci_web_acl] (continued)

Field	Description
Description [short_description]	Type of resource. The value is set to **microsoft.network/applicationgatewaywebapplicationfirewallpolicies**.

CI relationships

The Azure - Web Application Firewall Policy - Extended Inventory(LP) pattern creates these relationships to support Azure Web Application Firewall Policy discovery.

CI	Relationship	CI
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	Protected by::Protects	Web ACL [cmdb_ci_web_acl]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Web ACL [cmdb_ci_web_acl]
Web ACL [cmdb_ci_web_acl]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure Web Application Firewall - Policy [cmdb_azure_web_application_firewall_policy]	References	Web ACL [cmdb_ci_web_acl]

Azure tag discovery

The pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

Azure WebSite pattern-based discovery

Discovery and Service Mapping Patterns finds Azure App Service web applications in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

This pattern discovers Azure App Service web applications only. For Azure Functions discovery, see Microsoft Azure Functions discovery with Patterns.

Pattern-based discovery and mapping requirements

Verify the Microsoft Azure discovery prerequisites

For more information, see the prerequisites section in Microsoft Azure Cloud discovery using patterns.

Configure the Discovery schedule to support GovCloud

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Azure WebSite (LP) pattern.

Cloud WebServer [cmdb_ci_cloud_webserver]

Field	Description
Name [name]	The name of the Azure web application.
Object ID [object_id]	A unique identifier, allocated by Azure for this resource.
Fully qualified domain name [fqdn]	The fully qualified domain name (FQDN) of the web application.
IP Address [ip_address]	The inbound IP address of the web application.
State [state]	The current state of the web application. Possible values are: available or terminated.
Operational status [operational_status]	The operational status of the web application. The value is 1 (Operational) when the state is available, or 2 (Non-Operational) otherwise.
Vendor [vendor]	The vendor of the web application. The value is set to **Microsoft**.
Install Status [install_status]	Install status of the resource. Default value is Installed.

IP Address [cmdb_ci_ip_address]*

Field	Description
Name [name]	The name of the web application associated with this IP address.
IP Address [ip_address]	The inbound IP address of the web application.
Fully qualified domain name [fqdn]	The fully qualified domain name (FQDN) of the web application.
Netmask [netmask]	The netmask of the IP address. The value is set to **0.0.0.0**.
Vendor [vendor]	The vendor of the web application. The value is set to **Microsoft**.

* Populated only when the web application has a non-empty inbound IP address.

CI relationships and references

The Azure WebSite (LP) pattern creates the following relationships and references to support Azure WebSite discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud WebServer [cmdb_ci_cloud_webserver]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Cloud WebServer [cmdb_ci_cloud_webserver]	Owns::Owned by	IP Address [cmdb_ci_ip_address]*

* This relationship is created only when the web application has a non-empty inbound IP address and an FQDN.

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Cloud WebServer [cmdb_ci_cloud_webserver]

Azure Tag discovery

The Azure WebSite (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Cloud WebServer [cmdb_ci_cloud_webserver] table.

Connections found by Service Mapping during top-down discovery

The Azure WebSite TD pattern discovers the following connections during top-down discovery: Microsoft SQL Server, MySQL, Redis, and NoSQL (Cosmos DB).

Oracle Cloud Infrastructure (OCI) discovery

ServiceNow Discovery uses the Oracle Cloud Infrastructure (OCI) discovery patterns to provide real-time elasticity for enterprise applications by combining Oracle autonomous services, integrated security, and cloud compute. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Starting with version 1.29.0, Discovery and Service Mapping Patterns supports the discovery of OCI GovCloud accounts.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

Verify that the following applications are installed

Verify that the following applications are installed and up to date:

- Cloud Access Interface (starting from version 1.0.0)
- Visibility Content
- CMDB CI Class Models
- Discovery and Service Mapping Patterns

Configure Oracle API credentials

Ensure that your OCI credentials are configured and that your private and public keys are converted to the RSA format. For more information, see Create Oracle API credentials.

Create OCI service account

For more information, see Create OCI service accounts.

Create a cloud discovery schedule

To run a cloud discovery at scale, see Create an OCI Discovery schedule in Discovery Admin Workspace.

Create a discovery schedule to run LDCs on.

1. Navigate to **Discovery** > **Discovery Schedules**.
2. Select **Cloud Discovery**.
3. In the **Provider** field, select **OCI**. Fill in the form to create an account or select an existing one. For more information, see Create OCI service accounts.
4. On the **Select Datacenters** page, filter which LDCs you want to run the discovery.
5. When you're finished, navigate to the **Create Schedule** page and select **Finish and Run**.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected by Discovery during horizontal discovery

The discovered Oracle configuration data includes the following table and fields.

Field	Description
Cloud Database (cmdb_ci_cloud_database)	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
name	The name of the database.
object_id	Unique identifier of the database.
operational_status	Operational status of the database.
vendor	The vendor, **Oracle**.
type	The Oracle database type.
version	Version of the database.
state	State of the database.
comments	The attribute is issued for the delete strategy.
install_status	Install status.
Image (cmdb_ci_os_template)	
name	The image name.
object_id	Unique identifier of the image.
version	The version of the guest OS.
guest_os	The name of the guest OS.
Service Account (cmdb_ci_cloud_service_account)	
name	The account name.
datacenter_type	Datacenter type.
account_id	Unique identifier of the account.
object_id	Unique identifier of the account (same value as the account id).
discovery_credentials	The credentials related to the service account.
is_master_account	**true** when it is the root compartment. **false** when it is not the root compartment.
Network (cmdb_ci_network)	
name	The network name.
object_id	Unique identifier of the network.
cidr	The cidr of the network.
domain_name	The domain name.
state	The state of the network.
NIC (cmdb_ci_nic)	
name	The NIC name.
object_id	Unique identifier of the NIC.
public_ip	The public IP address.
mac_address	The MAC address.
primary	Set **true**/**false** according to the setup.
private_ip	The private ip address.
state	The state of the NIC.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Storage Volume (cmdb_ci_storage_volume)	
name	The name of the storage volume.
object_id	Unique identifier of storage volume.
state	The state of the storage volume.
volume_id	Unique identifier of the storage volume (same value as the account id).
size	The storage volume size.
size_bytes	The storage volume size in bytes.
Subnet (cmdb_ci_cloud_subnet)	
name	The name of the subnet.
object_id	Unique identifier of subnet.
cidr	The cidr of the subnet.
domain_name	The domain name.
state	The state of the subnet.
Virtual Machine (cmdb_ci_vm_instance)	
name	The VM name.
object_id	Unique identifier of VM.
state	The state of the VM.
memory	Memory size.
cpus	Number of CPUs.

Tags are populated to the **cmdb_key_value** table and related to the relevant cmdb ci resource. In order to support it, the extension sections were created for the following resources:

- Autonomous DB
- Exadata DB
- Image
- Network
- NIC
- Storage Volume
- Account
- Subnet
- Virtual Machine

CI Relationships

Relationships discovered using the Oracle OCI - Availability Domain (LP) pattern

CI	Relationship	CI
cmdb_ci_oci_datacenter	Contains::Contained by	cmdb_ci_availablity_zone

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships discovered using the Oracle OCI - Block (LP) pattern

CI	Relationship	CI
cmdb_ci_vm_instance	Use End Point To::Use End Point From	cmdb_ci_endpoint_block
cmdb_ci_endpoint_block	Implement End Point To::Implement End Point From	cmdb_ci_storage_volume

Relationships discovered using the Oracle OCI - Fault Domain (LP) pattern

CI	Relationship	CI
cmdb_ci_oci_datacenter	Contains::Contained by	cmdb_ci_availability_zone

Note: When configuring a Cloud Discovery scheduler for part of the LDCs, the Oracle OCI - Fault Domain (LP) pattern will not trigger.

Relationships discovered using the Oracle OCI - LDC (LP) pattern

CI	Relationship	CI
cmdb_ci_oci_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account

Relationships discovered using the Oracle OCI - Network (LP) pattern

CI	Relationship	CI
cmdb_ci_network	Hosted on::Hosts	cmdb_ci_oci_datacenter

Relationships discovered using the Oracle OCI - NIC Attachment (LP) pattern

CI	Relationship	CI
cmdb_ci_vm_instance	Use End Point To::Use End Point From	cmdb_ci_endpoint_subnet
cmdb_ci_vm_instance	Use End Point To::Use End Point From	cmdb_ci_endpoint_vnic
cmdb_ci_endpoint_vnic	Implement End Point To::Implement End Point From	cmdb_ci_nic
cmdb_ci_cloud_subnet	Implement End Point To::Implement End Point From	cmdb_ci_endpoint_subnet

Relationships discovered using the Oracle OCI - NIC (LP) pattern

CI	Relationship	CI
cmdb_ci_nic	Hosted on::Hosts	cmdb_ci_oci_datacenter

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships using the Oracle OCI - Storage Volume (LP) pattern

CI	Relationship	CI
cmdb_ci_availablity_zone	Contains::Contained by	cmdb_ci_storage_volume
cmdb_ci_storage_volume	Hosted on::Hosts	cmdb_ci_oci_datacenter

Relationships discovered using the Oracle OCI - Subnet (LP) pattern

CI	Relationship	CI
cmdb_ci_network	Contains::Contained by	cmdb_ci_cloud_subnet

OCI cloud hardware type pattern-based discovery

Discovery and Service Mapping Patterns finds OCI cloud hardware types (called shapes in OCI) in your Cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the OCI discovery prerequisites

For more information, see the prerequisites section in Oracle Cloud Infrastructure (OCI) discovery.

Configure the Discovery schedule to support GovCloud

Discovering OCI GovCloud accounts requires using a datacenter URL when setting up an OCI service account. For more information, see Create OCI service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle OCI - Cloud Hardware Type (LP) pattern.

Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Name [name]	Name of the shape.
Object ID [object_id]	Name of the shape.
VCPUs [vcpus]	Number of virtual CPUs (vCPU) available to the shape.
Memory (MB) [memory_mb]	Amount of memory available to the shape, in megabytes (MB).
Cores [cores]	Number of physical cores available to the shape.
Provider [provider]	Cloud provider, which is **OCI**.

CI relationships and references

The Oracle OCI - Cloud Hardware Type (LP) pattern creates the following relationships to support OCI cloud hardware type discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Cloud Hardware Type [cmdb_ci_cloud_hardware_type]	Hosted on::Hosts	OCI Datacenter [cmdb_ci_oci_datacenter]

OCI OS image pattern-based discovery

Discovery and Service Mapping Patterns finds OCI OS images on your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern overview

Discovery uses the Oracle OCI - Image (LP) pattern to discover OCI OS images from your Oracle account.

When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. When enabled, Discovery launches the Oracle OCI - Cloud OS Image (LP) pattern, which populates the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

Pattern-based discovery and mapping requirements

Verify the OCI discovery prerequisites

For more information, see the prerequisites section in Oracle Cloud Infrastructure (OCI) discovery.

Configure the Discovery schedule to support GovCloud

Discovering OCI GovCloud accounts requires using a datacenter URL when setting up an OCI service account. For more information, see Create OCI service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the OCI OS image patterns.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Name [name]	Name of the image resource.
Object ID [object_id]	A unique identifier, allocated by Oracle Cloud Infrastructure for this resource.
Guest OS [guest_os]	Operating system type of the image.
Version [version]	Version of the operating system.
Provider [provider]	The cloud provider, which is **OCI**. This field is only populated in the Cloud Image [cmdb_ci_cloud_os_image] table.
Install Status [install_status]	Install status of the resource. Default value is Installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image] (continued)

Field	Description
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

CI relationships and references

The OCI OS image patterns create the following relationships and references to support OCI OS image discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	OCI Datacenter [cmdb_ci_oci_datacenter]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

OCI Tag discovery

The OCI OS image patterns collect tags and populate them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image] table.

OCI virtual machine pattern-based discovery

Discovery and Service Mapping Patterns finds OCI virtual machines (VMs) in your cloud environment. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Pattern-based discovery and mapping requirements

Verify the OCI discovery prerequisites

For more information, see the prerequisites section in Oracle Cloud Infrastructure (OCI) discovery.

Configure the Discovery schedule to support GovCloud

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovering OCI GovCloud accounts requires using a datacenter URL when setting up an OCI service account. For more information, see Create OCI service accounts.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Oracle OCI - Virtual Machine (LP) pattern.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Name [name]	The display name of the VM instance.
Object ID [object_id]	A unique identifier, allocated by Oracle OCI for this resource.
CPUs [cpus]	The number of Oracle CPU (OCPUs) available to the instance.
Memory (MB) [memory]	The amount of memory available to the instance, in gigabytes.
IP Address [ip_address]	The private IP address of the primary network interface attached to the instance.
State [state]	The current state of the virtual machine instance. For example: on, off, starting, stopping, scheduled, terminated, or error.
Install Status [install_status]	Install status of the resource. Default value is Installed.
Operational status [operational_status]	Operational status of the resource. Default value is Operational.

Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]

Field	Description
Object ID [object_id]	A unique identifier, allocated by Oracle OCI for this image.
Provider [provider]	The cloud provider, which is **OCI**. This field is only populated in the Cloud Image [cmdb_ci_cloud_os_image] table.

Note: When using the Image [cmdb_ci_os_template] table to store Cloud OS images, you may notice an unusually large number of records. To avoid this issue, you can store the discovered OS images in the Cloud Image [cmdb_ci_cloud_os_image] table. For more information, see Enable Cloud OS Image discovery.

CI relationships and references

The Oracle OCI - Virtual Machine (LP) pattern creates the following relationships and references to support OCI VM discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI	Relationship	CI
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on::Hosts	OCI Datacenter [cmdb_ci_oci_datacenter]
Availability Zone [cmdb_ci_availability_zone]	Contains::Contained by	Virtual Machine Instance [cmdb_ci_vm_instance]
Virtual Machine Instance [cmdb_ci_vm_instance]	DR provided by::Provides DR for	Availability Zone [cmdb_ci_availability_zone]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]
Image [cmdb_ci_os_template] or Cloud Image [cmdb_ci_cloud_os_image]	Hosted on::Hosts	OCI Datacenter [cmdb_ci_oci_datacenter]
Server [cmdb_ci_server]	Virtualized by::Virtualizes	Virtual Machine Instance [cmdb_ci_vm_instance]
Virtual Machine Instance [cmdb_ci_vm_instance]	Use End Point To::Use End Point From	VNIC Endpoint [cmdb_ci_endpoint_vnic]
VNIC Endpoint [cmdb_ci_endpoint_vnic]	Implement End Point To::Implement End Point From	Cloud Mgmt Network Interface [cmdb_ci_nic]
Virtual Machine Instance [cmdb_ci_vm_instance]	Provisioned From::Provisioned	Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

CI references

CI	Field	Referenced CI
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Virtual Machine Instance [cmdb_ci_vm_instance]

OCI Tag discovery

The Oracle OCI - Virtual Machine (LP) pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OCI BYOL discovery

Starting with Discovery and Service Mapping Patterns version 1.35.0, the pattern extension section discovers the following license types for Windows VMs:

- Bring Your Own License (BYOL)
- License included

The pattern stores the license type and model in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	The license type key, which is **Windows_OS_License_Type_automatic**.
Value [value]	The license model, which is one of the following: • BYOL • License Included
Configuration item [configuration_item]	References the Virtual Machine Instance [cmdb_ci_vm_instance] table.

OpenStack resource discovery

The ServiceNow Discovery application uses OpenStack resource discovery patterns to find OpenStack resources through REST API calls. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Discovery uses OpenStack resource discovery patterns to run horizontal discovery.

Starting with the Discovery and Service Mapping Patterns August 2024 store release, OpenStack resource discovery patterns support both private and public cloud discovery.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

OpenStack requirements

- In the OpenStack Identity service catalog, find the access URL that is used as the OpenStack identity service endpoint. For example: https://<openstack dns>:<port>.
- Configure an OpenStack user with permissions to generate a project scope token. Setting Domain ID, User Name, Project Name, and Password values is required for generating a token.

MID Server requirements

Deploy and connect a dedicated MID Server for cloud discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The dedicated MID Server must have access to OpenStack and an open port for each OpenStack service to be discovered.

Credentials

On the ServiceNow AI Platform, create OpenStack credentials for the ServiceNow cloud service account. Enter the Domain ID, User Name, Project Name, and Password values you've configured in OpenStack requirements. For more information, see Create and test your credentials.

Service account

On the ServiceNow AI Platform, configure a ServiceNow cloud service account using the created OpenStack credentials and the OpenStack datacenter details. In the **Account Id** field, provide the Project ID of the OpenStack Project to be discovered.

Create a discovery schedule

For more information, see Create an OpenStack Discovery schedule in Discovery Admin Workspace.

Create an automatic discovery schedule for OpenStack service accounts—for private cloud only

Create an automatic serverless discovery schedule for OpenStack service accounts using the OpenStack - Discover Service Accounts pattern. For more information, see the Auto Cloud discovery schedule for Open stack [KB1638685] article in the Now Support Knowledge Base.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the OpenStack resource discovery patterns.

Virtual Machine Instance [cmdb_ci_vm_instance]

Field	Description
Object ID [object_id]	Server UUID.
Name [name]	Server name.
Status [status]	Server status.
Description [short_description]	(Optional) Server description.
CPUs [cpus]	Number of virtual CPUs (vCPUs) from the Flavor that the VM instance is provisioned with.
Disks [disks]	Number of all managed volumes by Cinder that are returned from the VM instance.
Disks size (GB) [disk_size]	Total amount of all managed volumes by Cinder that are used by the VM instance measured in gigabytes (GB).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance] (continued)

Field	Description
Memory (MB) [memory]	The **ram** value from the Flavor that the VM instance is provisioned with, measured in megabytes (MB).
Network adapters [nics]	Number of unique interfaces configured in the VM instance.
State [state]	Mapped status value returned for the VM instance. The OpenStack Server statuses and their related ServiceNow **State** values are: • ACTIVE: on • RESCUED: on • PAUSED: paused • SUSPENDED: off • STOPPED: off • SOFT_DELETED: terminated • HARD_DELETED: terminated • RESIZED: off • ERROR: error • INITIALIZED: starting For any **status** value not listed, the default value is off.

Cloud Key Pairs [cmdb_ci_cloud_key_pair]

Field	Description
Object ID [object_id]	Keypair UUID.
Name [name]	Keypair name that is used to reference it.
Finger print [finger_print]	Keypair fingerprint.

Cloud Load Balancer [cmdb_ci_cloud_load_balancer]

Field	Description
Object ID [object_id]	Load balancer ID.
Name [name]	Human-readable resource name.
Description [short_description]	(Optional) Human-readable resource description.
Install Status [install_status]	Resource provisioning status.
Operational Status [operational_status]	Resource operational status.

Load Balancer Listener [cmdb_ci_lb_listener]

Field	Description
Object ID [object_id]	Listener ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Listener [cmdb_ci_lb_listener] (continued)

Field	Description
Name [name]	Human-readable name of the resource.
Install Status [install_status]	Provisioning status of the resource.
Listener Protocol [listener_protocol]	Protocol for the resource. One of HTTP, HTTPS, TCP, TERMINATED_HTTPS, or UDP.
Listener Port [listener_port]	Protocol port number for the resource.
Operational status [operational_status]	Resource operational status.
Description [short_description]	(Optional) Human-readable resource description.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Object ID [object_id]	Pool ID.
Name [name]	Human-readable resource name.
Install Status [install_status]	Resource provisioning status.
Operational Status [operational_status]	Resource operational status.
Description [short_description]	(Optional) Human-readable resource description.
Load Balancing Method [load_balancing_method]	Load balancing algorithm for the pool. One of these values: LEAST_CONNECTIONS, ROUND_ROBIN, SOURCE_IP, or SOURCE_IP_PORT.

Load Balancer Pool Member [cmdb_ci_lb_pool_member]

Field	Description
Object ID [object_id]	Pool member ID.
Name [name]	Human-readable resource name.
Install Status [install_status]	Resource provisioning status.
IP address [ip_address]	Back-end member server IP address.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Object ID [object_id]	Volume UUID.
Name [name]	Volume name.
Status [status]	Volume status.
Description [short_description]	(Optional) Human-readable description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume [cmdb_ci_storage_volume] (continued)

Field	Description
Storage type [storage_type]	Associated volume type.

Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]

Field	Description
Object ID [object_id]	Volume UUID.
Name [name]	Volume name.
Status [status]	Volume status.
Description [short_description]	(Optional) Human-readable description.

Image [cmdb_ci_os_template]

Field	Description
Object ID [object_id]	User-defined image UUID.
Name [name]	Image name. The value can be null (the JSON Null data type).
Status [status]	Image status.

Cloud Network [cmdb_ci_network]

Field	Description
Object ID [object_id]	Network ID.
Name [name]	Human-readable network name.
Operational Status [operational_status]	Network status. Possible values: ACTIVE, DOWN, BUILD, or ERROR.
Description [short_description]	(Optional) Human-readable description.

Cloud Subnet [cmdb_ci_cloud_subnet]

Field	Description
Object ID [object_id]	Subnet ID.
Name [name]	Human-readable resource name.
DHCP enabled [is_dhcp_enabled]	Whether DHCP is enabled or disabled for the subnet.
Description [short_description]	(Optional) Human-readable description.

Compute Security Group [cmdb_ci_compute_security_group]

Field	Description
Object ID [object_id]	Security group ID.
Name [name]	Human-readable resource name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Compute Security Group [cmdb_ci_compute_security_group] (continued)

Field	Description
Description [short_description]	(Optional) Human-readable description.

Security Group Rule [cmdb_ci_compute_security_group_rule]

Field	Description
Name [name]	This field has the same value as the object_id field.
Object ID [object_id]	Security group rule ID.
Description [short_description]	(Optional) Human-readable description.
Ethertype [ethertype]	Ethertype. Must be IPv4 or IPv6. Addresses represented in CIDR must match the ingress or egress rules.
Port range max [port_range_max]	The last port of the range to which the rule applies. This value is the maximum port number in the range that is matched by the security group rule. • If the protocol is TCP, UDP, DCCP, SCTP or UDP-Lite, this value must be greater than or equal to the port_range_min attribute value. • If the protocol is ICMP, this value must be an ICMP code.
Port range min [port_range_min]	The first port of the range to which the rule applies. This value is the minimum port number in the range that is matched by the security group rule. • If the protocol is TCP, UDP, DCCP, SCTP or UDP-Lite, this value must be less than or equal to the port_range_max attribute value. • If the protocol is ICMP, this value must be an ICMP type.

OpenStack Domain [cmdb_ci_cloud_openstack_domain]

Field	Description
Object ID [object_id]	Domain ID.
Name [name]	Domain name.
Description [short_description]	(Optional) Human-readable description.
Enabled [enabled]	Defines whether the domain is enabled or disabled. • True - The domain is enabled. • False - The domain is disabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OpenStack Service [cmdb_ci_cloud_openstack_service]

Field	Description
Object ID [object_id]	Service ID.
Name [name]	Service name.
Description [short_description]	(Optional) Human-readable description.
Type [type]	Service type. Describes the API implemented by the service. Possible values: Compute, ec2, identity, image, network, or volume.
Catalog enabled [enabled]	Defines whether the service and its endpoints appear in the service catalog. • True - The service and its endpoints appear in the service catalog. • False - The service and its endpoints do not appear in the service catalog.

OpenStack Endpoint [cmdb_ci_cloud_openstack_endpoint]

Field	Description
Object ID [object_id]	Endpoint ID.
Name [name]	Service name.
Interface [interface]	Interface type that describes the visibility of the endpoint. Possible values: • Public - Visible by end users on a publicly available network interface. • Internal - Visible by end users on an unmetered internal network interface. • Admin - Visible by administrative users on a secure network interface.
Catalog enabled [enabled]	Defines whether the service and its endpoints appear in the service catalog. • True - The service and its endpoints appear in the service catalog. • False - The service and its endpoints do not appear in the service catalog.

Cloud Hardware Type [cmdb_ci_cloud_hardware_type]

Field	Description
Name [name]	The **name** field of OpenStack Flavor.
Object ID [object_id]	The **id** field of OpenStack Flavor.
vCPUs [vcpus]	The **vcpus** field of OpenStack Flavor.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Hardware Type [cmdb_ci_cloud_hardware_type] (continued)

Field	Description
Memory MB [memory_mb]	The **ram** field of OpenStack Flavor.
Local Storage GB [local_storage_gb]	The **disk** field of OpenStack Flavor.

The discovered data includes the OpenStack Hypervisor View [cmdb_ci_openstack_hypervisor_view] table and its fields for private cloud only.

OpenStack Hypervisor View [cmdb_ci_openstack_hypervisor_view] table fields

Field	Description
Name [name]	Concatenated value of **hypervisor_type** and **hypervisor_hostname** fields. For example: QEMU@host.acmecorp.com.
IP Address [ip_address]	The value of **host_ip** field.
Hypervisor technology [hypervisor_technology]	The value of **hypervisor_type** field.
Object ID [object_id]	The value of **id** field.
State [state]	The value of **state** field.
Version [version]	The value of **hypervisor_version** field.

The Dependency Views map shows all discovered OpenStack resources in your organization that are available through REST API calls and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OpenStack Network Dependency

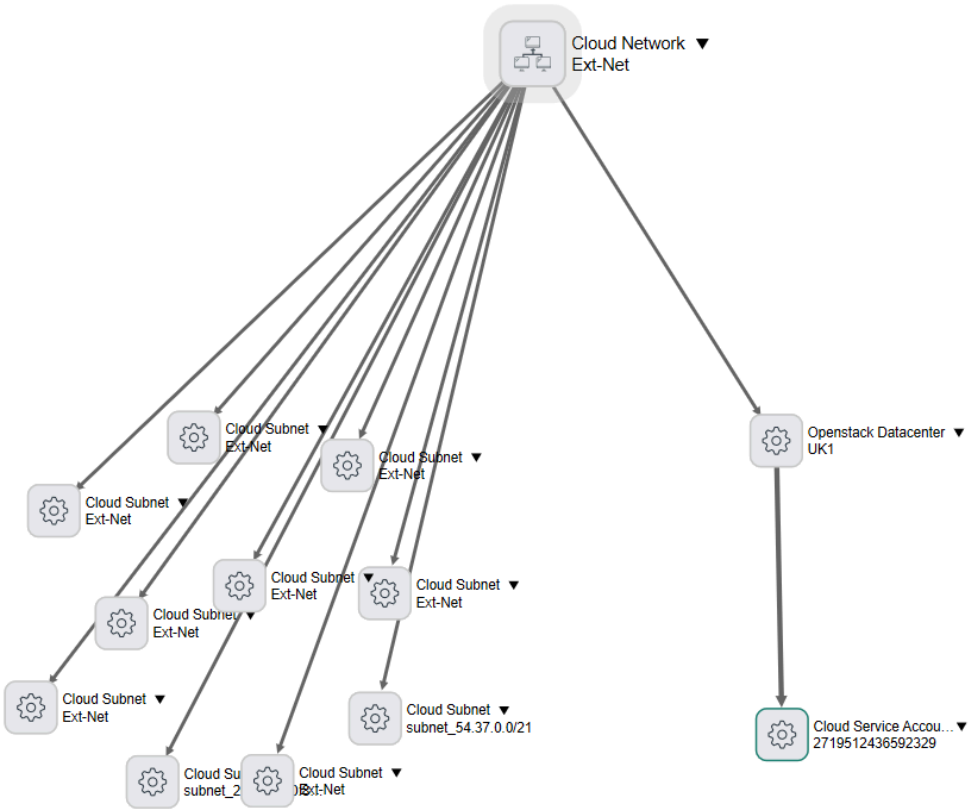

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OpenStack VM Dependency

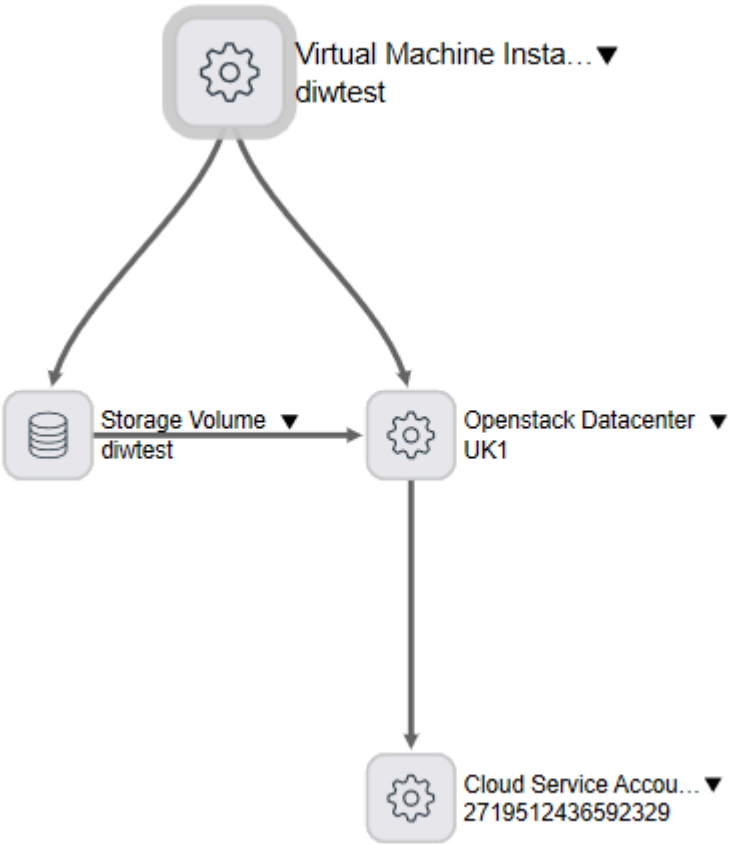

CI relationships

The relationships listed in the following table are created to support OpenStack resource discovery.

CI	Relationship	CI
cmdb_ci_openstack_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_openstack_service	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_openstack_endpoint	Use endpoint to::Use endpoint from	cmdb_ci_cloud_openstack_service
cmdb_ci_cloud_openstack_endpoint	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_cloud_openstack_domain	Owns::Owned by	cmdb_ci_cloud_service_account
cmdb_ci_network	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_network	Contains::Contained by	cmdb_ci_cloud_subnet
cmdb_ci_compute_security_group	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_compute_security_group	Contains::Contained by	cmdb_ci_compute_security_group_rule
cmdb_ci_cloud_load_balancer	Hosted on::Hosts	cmdb_ci_openstack_datacenter

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
cmdb_ci_cloud_load_balancer	Contains::Contained by	cmdb_ci_lb_listener
cmdb_ci_lb_pool	Hosted on::Hosts	cmdb_ci_cloud_load_balancer
cmdb_ci_lb_pool	Owns::Owned by	cmdb_ci_lb_pool_member
cmdb_ci_storage_volume	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_vm_instance	Uses::Used by	cmdb_ci_storage_volume
cmdb_ci_storage_vol_snapshot	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_storage_vol_snapshot	Provisioned from::Provisioned	cmdb_ci_storage_volume
cmdb_ci_vm_instance	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_cloud_key_pair	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_os_template	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_vm_instance	Provisioned from::Provisioned	cmdb_ci_os_template
cmdb_ci_vm_instance	Virtualized by::Virtualizes	cmdb_ci_server

The relationships listed in the following table are created to support OpenStack resource discovery in private cloud only.

CI	Relationship	CI
cmdb_ci_vm_instance	Registered on::Has registered	cmdb_ci_openstack_hypervisor_view
cmdb_ci_openstack_hypervisor_view	Hosted on::Hosts	cmdb_ci_openstack_datacenter
cmdb_ci_openstack_hypervisor_view	Uses::Hosts Used by	cmdb_ci_server

Configuring Pattern Designer and related applications

Install the latest available versions of the ServiceNow applications to use the latest discovery patterns.

The Pattern Designer (com.snc.pattern.designer) application provides the UI for creating or customizing discovery patterns. Pattern Designer is automatically installed with Discovery or Service Mapping. The patterns themselves are released using dedicated pattern applications.

The following user roles have access to Pattern Designer and can perform various actions. Note that customizing patterns requires basic knowledge of programming.

Pattern Designer users and access

User	Description
PD user	Has read-only access to **Discovery Pattern Log**.
PD admin	Can view, create, edit, and publish patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern Designer users and access (continued)

User	Description
PDE viewer	Starting with Pattern Designer Enhancements version 3.9.0, users can view **Command Validation Tasks**, **Command Validation Tasks Results**, and **Command List**. The pde_viewer can view the **Command Validation Tool** modules and related tables, but doesn't have permissions to modify or edit them. The pde_viewer role can view the following tables only: • Command List [pd_command_list] • Command Validation Task [pd_command_validation] • Command Validation Task Results [pd_command_validation_results] • Pattern Shared Library Mapping [pd_pattern_to_shared_library_mapping] • Temporary Variable Mappings [pd_temp_variable_value_mapping]
PD MID	Not assigned to a user directly but to the MID Server record or the user under which the MID Server runs. The role enables the MID Server to interpret and run pattern-based probes.

ServiceNow releases all discovery patterns using the following applications:

Discovery and Service Mapping Patterns (sn_itom_pattern)

This application provides the latest versions of discovery patterns the original version of which ServiceNow released on ServiceNow Store.

Visibility Content (sn_pattern_design)

This application supplies the updated version of the patterns that were part of the family releases up until Tokyo.

You can install the latest available versions of the pattern applications from the ServiceNow Store. Alternatively, you can install these applications as plugins on your ServiceNow instance.

Check which plugin contains a required pattern

Since ServiceNow releases discovery patterns using two different plugins, you may need to check which plugin to install or upgrade.

Before you begin

Role required: service_mapping_admin

About this task

ServiceNow releases all discovery patterns using the following applications:

Discovery and Service Mapping Patterns (sn_itom_pattern)

This application provides the latest versions of discovery patterns the original version of which ServiceNow released on ServiceNow Store.

Visibility Content (sn_pattern_design)

This application supplies the updated version of the patterns that were part of the family releases up until Tokyo.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Display the **Application** column.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Click **Personalize List**.

b. In the Personalize List Columns dialog box, select **Application** under **Available**, and click the right arrow.

3. **Optional:** You can filter out shared libraries to see only patterns in the list.
4. Search for the relevant pattern in the list.
5. Check to which application it belongs.

   Global stands for Visibility Content.

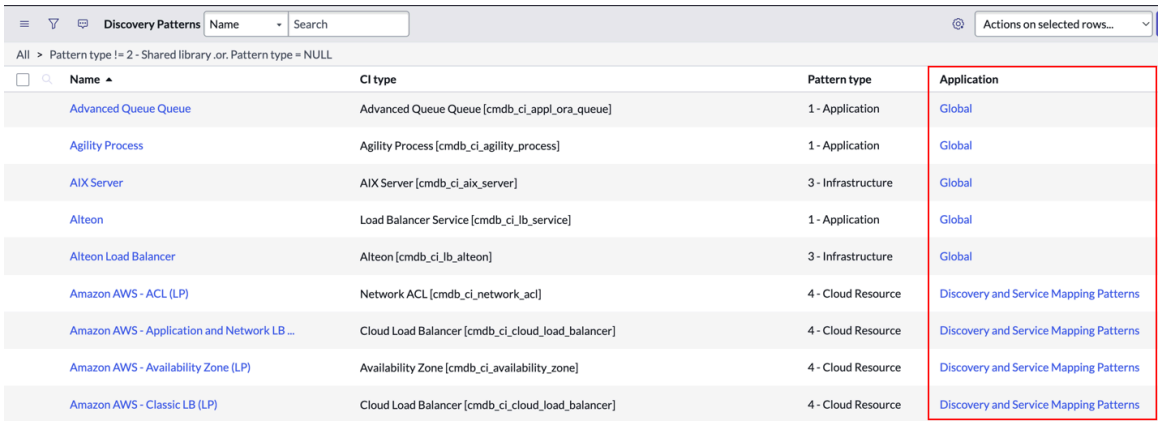

Install Discovery and Service Mapping Patterns

You can install the Discovery and Service Mapping Patterns application (sn_itom_pattern) if you have the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Discovery and Service Mapping Patterns application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

Multiple tables are installed with Discovery and Service Mapping Patterns.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Discovery and Service Mapping Patterns application (sn_itom_pattern) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.

3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

Install Visibility Content

You can install the Visibility Content application (sn_pattern_design) if you have the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Visibility Content application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

Multiple tables are installed with Visibility Content.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Visibility Content application (sn_pattern_design) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

Password2 encryption for patterns

The Password (two-way encrypted) field type is used to encrypt and decrypt data. The field type works in accordance with NIST 800-57 guidelines and provides FIPS 140-2-L3 protection.

Password2 is a text field that stores sensitive data with two-way encryption. The two-way encryption stores the data as a secure encrypted value that can be decrypted within the ServiceNow AI Platform instance.

From Quebec, Key Management Framework enables you to manage keys used for Password2 fields. However, the implementation of Password2 is required on a ServiceNow AI Platform using the Rome release or later.

Note: Only newly created records containing Password2 fields are two-way encrypted. Existing data is not automatically converted to use Password2 encryption.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify Password2 implementation

Discovery and Service Mapping Patterns store app

Verify that the Discovery and Service Mapping Patterns store app (version 1.0.76) is installed.

Verify Password2 is included in XML records

View the XML records in the ServiceNow AI Platform instance, and verify that the encryption looks similar to this example:

Create CI types for Service Mapping and Discovery

Create CI types (or CI classes) for applications and devices, which Service Mapping and Discovery do not support by default.

Before you begin

Role required: service_mapping_admin, personalize_dictionary, or admin
Verify that there is no existing CI type for this application or device.

About this task

All applications and hosts in your organization must have a corresponding configuration item (CI) type, which is necessary for discovering and processing applications and hosts correctly.

Note: The Discovery and the CMDB interface refers to CI types as CI classes.

A wide range of preconfigured CI types that cover most commonly used applications are available to you. If your organization uses a less known or proprietary application that does not have a corresponding CI type, you must create it.

A CI type is a generic notion that is used by several ServiceNow applications, but there are some attributes that are specific to Service Mapping and Discovery.

A CI type (or class) contains several important definitions that apply to all CIs belonging to it, such as:

- CI attributes are added as fields to the CMDB tables.
- Identifiers help Service Mapping and Discovery to differentiate between new and existing CIs. For example, if there is an Apache Web Server CI type defined in the CMDB, and Service Mapping and Discovery both discover an Apache Web Server CI, it processes it using identifiers. It then recognizes it as an updated version of the Apache Web Server CI that already exists in the system, not a new Apache Web Server CI.
- There are reconciliation rules that help the ServiceNow platform to consolidate CI attributes received from different applications correctly. These rules are necessary for organizations where more than one application participates in the discovery process. Reconciliation rules define how attributes of the same CI discovered by different discovery sources are merged. For example, Service Mapping discovers the version and home directory attributes of an Apache Web Server CI, while Discovery discovers the version and patch level attributes for the same Apache Web Server CI. The ServiceNow AI Platform applies the reconciliation rule and as a result Service Mapping does not overwrite the attributes found by Discovery.
- Related items, both CIs and non-CIs, serve as a reference and provide additional information about the CI. For example, you can define a serial number, a port and an network APR table as related items for a hardware server. The system does not use these items for identification. After you configure related items for the CI type, you can use horizontal discovery for CIs of this

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

type together with their related items. For information, see Discover related items together with the main CI.

Note: Non-CIs are items that do not extend the Configuration Item [cmdb_ci] table. Examples of non-CIs are ports, serial numbers, or Network ARP tables.

- CMDB stores CI class in the form of a hierarchy where some CI classes are parents to other CI classes, who automatically derive their parent's attributes in addition to attributes you configure specifically for child CI classes. CI class hierarchy is used widely for configuring CI behavior, relationships, and display. In this example, the Apache Web Server CI is a child of the Web Server CI and derives many attributes from its parent, such as name, version, model ID. By default, child CI classes inherit identifiers from their parent CI classes.

In addition to these CI type definitions, the horizontal discovery process uses a CI classification to define to which CI type a CI belongs. Create a device CI classification if you create a CI type for devices using SNMP and a process CI classification for an application CI type.

Procedure

1. Navigate to **All** > **Configuration** > **CI Class Manager**.
2. To use an existing CI type as a parent for the new CI type, select **Hierarchy** to display the CI Classes list.
3. Select the context menu for the required CI type and select **Add Child Class**.
4. Create a table to store the CI type attributes:

Table form fields

Field	Description
Label	CI type name. For example, Apache Web Server.
Name	The table name. For example, cmdb_ci_apache_web_server. Use a name similar to the other CMDB classes (for example, `u_cmdb_ci_laptop`).
Extends table	The table name of the parent CI type. All CI types are extensions of the Configuration Item [cmdb_ci] table or its child tables. For example, if the new class is Laptops, which is a subclass of Computers, select the cmdb_ci_computer table. If the new class is a top-level class, select the cmdb_ci table. The most commonly used parent CI types are: ○ cmdb_ci - basic ○ cmdb_ci_database - for databases ○ cmdb_ci_app_server - for application servers ○ cmdb_ci_infra_service - for infrastructure services ○ cmdb_ci_endpoint_inclusion - for entry points of the inclusion type ○ cmdb_ci_appl - for applications ○ cmdb_ci_web_server - for web servers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ cmdb_ci_lb - for load balancers ○ cmdb_ci_endpoint - for entry points

5. Configure how the instance determines if a discovered CI is an upgraded CI existing in the instance or a brand new CI.
See Create or edit a CI identification rule for detailed procedure.

> **Warning:** If there is no CI identification rule for a CI type, Service Mapping discovers CIs belonging to this type, but cannot interpret the results of the discovery process. In this case, the ServiceNow AI Platform rejects the discovery results for these CIs and their information is not updated.

a. On the **Identifiers** page, configure the parameters as follows:

Identifier form

Field	Description
Name	Use a meaningful name.
Applies to	Enter the CI type.
Active	Select the check box.
Independent	Clear the check box.

b. Configure the identifier entry parameters as follows:

Identifier entry form

Field	Description
Criterion attributes	Select the attribute that you want to use for identifying the CI. For example, **Class**, **Installation directory**, or **Serial number**.
Allow null attribute	Select the check box.
Priority	If there is more than one identifier entry, add priority to determine the order in which Service Mapping applies these entries. If none of the identification criteria matches the discovered CI, this CI is new.

6. **Optional:** Configure references to other CIs or non-CIs as related items for the CI type.

Related Entry form

Field	Description
Active	Check box that specifies that the related entry is active.
Related table	A related table that references the CI that is being matched.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Referenced field	A referenced field in **Related table** that should store the referenced CI. This field always references the cmdb_ci table, or a descendent of the cmdb_ci table.
Priority	Priority of the related entry for the specified **Related table**. Rules with lower priority numbers are given higher priority while matching a related item for specific related table. Related entries for the specified related table with identical priorities are applied randomly. You can keep gaps between the priority numbers, so you can assign the unused priority numbers to new entries without modifying the existing priority order.
Criterion attributes	The set of attributes to uniquely identify the related item. Attributes can belong to the current class, or to a parent class. **Note:** It is possible to add reference fields as a criterion attribute. However, such fields might not always be effective: ○ Reference fields store *sys_id*s that point to a record in another table, and thus is considered a weak criterion attribute (in terms of uniqueness) for the current table. ○ The system detects and then replaces invalid values in a reference field with 'Unknown'. For example, an invalid *Model ID* value is replaced with the value 'Unknown'. Also, if several CIs end up having that same reference field set to 'Unknown', then these CIs become duplicate CIs. Select the lock icon to view, add, or remove attributes from the identification rule.
Allow null attribute	If at least one criterion attribute in the related table is not null, allow to attempt matching with an identifier entry even if there are criterion attributes which are null.
Filter conditions	Add conditions to construct a filter to narrow the set of records that will be searched for a matching related item.

7. **Optional:** Configure the instance to consolidate CI attributes received from different data sources correctly.
   See Create or edit a CI reconciliation rule.

   Configure the following Service Mapping related parameters correctly:

   Reconciliation Definition form

Field	Description
Data source	Select `ServiceWatch`.
Applies to	Select the relevant CI type.
Optional condition	Set a condition if necessary.

   Note: If you do not create a CI reconciliation rule, data discovered by patterns is used to update CI attributes.

8. Configure the CI classification for your CI type:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For CI types representing applications, perform configuration as described in Create a Discovery process classification.
- For CI types representing SNMP devices, perform configuration as described in Create a Discovery CI classification.

Note: There is no need to create CI classifications for hosts because these classifications are included in the base system.

9. For CI types that represent inclusions, define the hierarchy for the new CI type.

   Clear the **Reverse Relationship Direction** check box while performing this configuration.

   See Create dependent relationships.

10. If necessary, customize icons that represent CIs in maps.
    See Create or modify map icons.

Create entry point types for Service Mapping

An entry point is a point where clients access a service instance. Service Mapping starts the mapping process for every application service from the entry point you define for it. Service Mapping includes a wide range of preconfigured entry point types that cover most commonly used applications. If your organization uses a less known or proprietary application that does not have a corresponding entry point type in Service Mapping, you must create it.

Before you begin

If your ServiceNow instance uses domain separation and you have access to the global domain, log in to the relevant domain. The selected domain must be a domain without any child domains.

Role required: service_mapping_admin or admin

About this task

In addition to starting the mapping process from entry points, Service Mapping patterns use entry points to discover CI outbound connections.

Entry points are modeled in the ServiceNow CMDB as CIs of endpoint type. Entry points are stored as records in the Endpoint [cmdb_ci_endpoint] tables.

Like any other CI type, an entry point contains several important definitions that apply to all CIs belonging to it:

- CI attributes are added as fields to the CMDB tables.
- Identifiers help Service Mapping and Discovery to differentiate between new and existing CIs. For example, if there is an endpoint CI type for Active Directory Forest in the CMDB 360, and Service Mapping discovers an Active Directory Forest CI. Service Mapping processes this CI using identifiers and recognizes it as an updated version of the Active Directory Forest CI that exists in the system, not a new Active Directory Forest CI.
- CI type hierarchy. CMDB stores CI class in the form of a hierarchy where some CI classes are parents to other CI classes, who automatically derive their parent's attributes in addition to attributes you configure specifically for child CI classes. CI class hierarchy is used widely for configuring CI behavior, relationships, and display. Create standard entry points as child CIs for the endpoint CI type, which creates an extension for the cmdb_ci_endpoint table. For entry points of inclusion type create child CIs for the inclusion endpoint CI type extending the cmdb_ci_endpoint_inclusion table. In an inclusion, a server hosts applications that are treated as independent objects.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Configuration** > **CI Class Manager**.
2. Select **Open Hierarchy**.
3. To create a standard entry point, right-click **Endpoint** from the **Class Hierarchy** pane and select **Add Child Class**.
4. To create an entry point of the inclusion type, right-click **Inclusion Endpoint** from the **Class Hierarchy** pane and select **Add Child Class**.
5. Define the parameters for the new entry point type as described in Create a CI class.

Related topics

Map a single application service using classic Service Mapping

Create CI types for Service Mapping and Discovery

Create or customize patterns

Create or modify a discovery pattern and define its basic attributes.

Before you begin

Make sure that the application for which you want to create a pattern, has a corresponding configuration item (CI) type and a CI classification. If the CI type you require is not in the list, create it as described in Create CI types for Service Mapping and Discovery.

If your ServiceNow instance uses domain separation and you have access to the global domain, log in to the relevant domain. The selected domain must be a domain without any child domains.

Role required: pd_admin

Basic knowledge of programming is desirable.

About this task

Patterns can be of the infrastructure or application type. Infrastructure patterns are used only by Discovery for creating lists of devices. Application patterns serve both Service Mapping and Discovery, which use the same application patterns for their purposes.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Click **New** or select the relevant pattern from the list.
3. Define the basic pattern attributes on the **Basic** tab.

Field	Description
Pattern type	○ Select **Application** for an application pattern. It can be used both for top-down discovery performed by Service Mapping and horizontal discovery performed by Discovery. ○ Select **Infrastructure** for an infrastructure pattern used for the horizontal host discovery performed by Discovery.
Name	Enter the pattern name. This name must be unique to this pattern. Use self-explanatory names such as `Apache on Unix pattern`.
CI type	Select the CI type which you want this pattern to discover.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: Discovery can find multiple CIs that belong to the same CI type.
Operating system [Application patterns only]	Select the operating system that the selected CI runs: ○ Click **All** if the CI runs on more than one operating system. Or ○ Select the relevant operating systems from the list.
Run Order [Application patterns only]	For an application pattern used by Service Mapping, define when the pattern runs **a.** Select the order in which this pattern always runs: ▪ **Before** ▪ **After** **b.** Then select the other applicable pattern. This field is only relevant if a particular pattern can be confused with another pattern. For example, both IIS and MS Exchange applications have an HTTP entry point. However MS Exchange uses some of the components of IIS. Therefore, if the IIS pattern ran first, discovery might incorrectly identify MS Exchange as IIS. To prevent this error, in the **Run Order** field in the MS Exchange pattern definition, select **Before** and **IIS**.
Description	Provide a description for this pattern.

4. **Optional:** When creating an application pattern, make the MID Server run this pattern only if the process identified on a CI matches the classification criteria for this pattern, select **Enforce Process Classification**.
   All simplified patterns created from generic applications, have this attribute enabled. For more information about creating process classification, see Discovery classifiers.
5. Define a set of identification steps for every incoming connection of a configuration item (CI) as described in KB1698521: Create or customize patterns.
6. If necessary, create more identification sections.
7. Save the pattern.
8. **Optional:** To test the pattern, select **Debug Mode** and enter the IP address of the host.
9. Upload the updated set of patterns onto the MID Server:
   a. Navigate to **Discovery** > **MID Servers**.
   b. Alternatively, navigate to **Service Mapping** > **MID Servers**.
   c. Select **Pattern Sync to Mid**.

What to do next

For application type patterns, continue with creating a connection section as described in the KB article KB1698521: Create or customize patterns.

Related topics

Discovery for cloud environment

KB0747679: Pre/Post Processing Scripts for patterns

KB0854603: Pattern Debugger Session Timed Out

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discover related items together with the main CI

Add related items to the patterns to perform horizontal discovery of configuration items with all their related items: CIs or non-CIs, like ports or serial numbers.

Before you begin

1. (Optional) Check if the CI is dependent: The system uses the identification rules to identify the CI independently of other CIs. Refer to Create or edit a CI identification rule.
2. (Optional) To understand the model better, check dependent relationship rules which define the dependency structure of the CI types and the relationship types in these service definitions. For example, you can discover a Tomcat server with Tomcat WARs belonging to it. For more information, refer to Dependent relationship rules.

Example of the Containment rule for Tomcat WAR

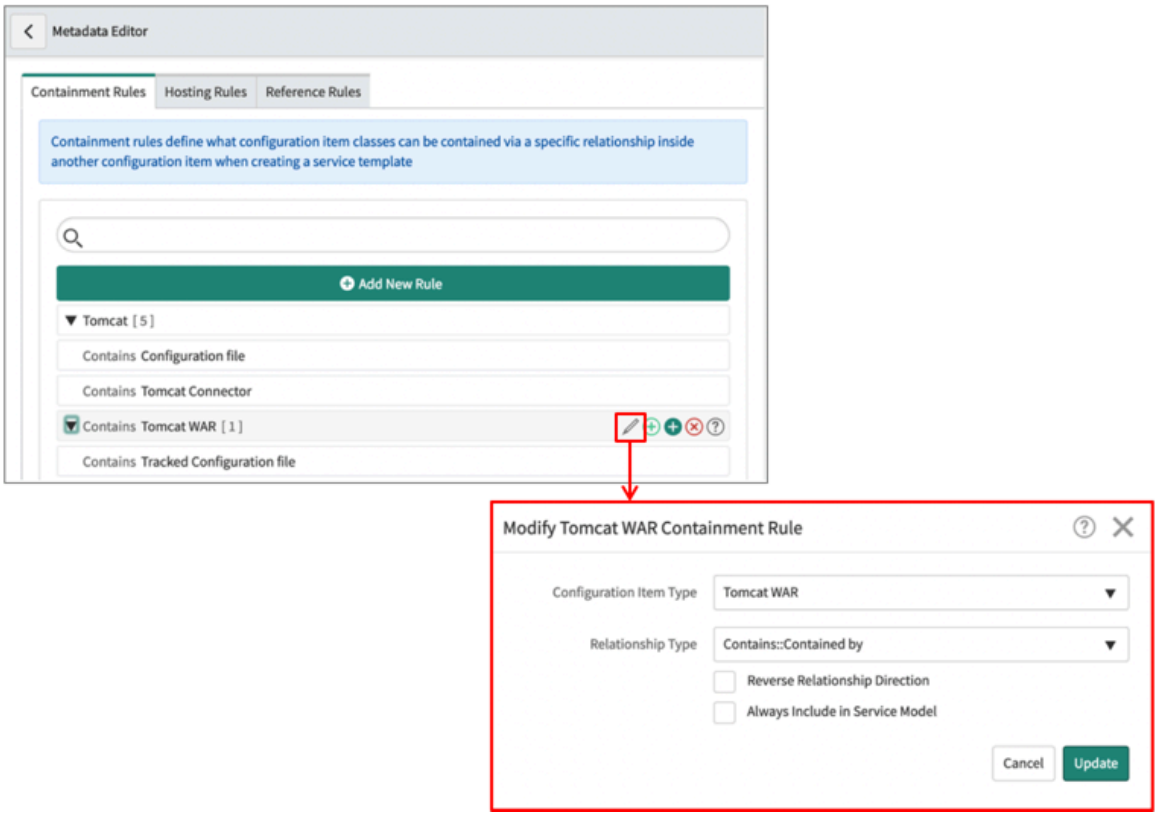

3. To add a non-CI to the pattern, add it as a related item for the CI type to which the main CI belongs. For more information on adding related items for the CI type, see Configure references to other CIs or non-CIs as related items for the CI type.

Role required: pd_admin

About this task

Related items, both CIs and non-CIs, serve as a reference and provide additional information about the CI. For example, you can define a serial number, a port and an network APR table as related items for a hardware server. The system does not use these items for identification. After you configure related items for the CI type, you can use horizontal discovery for CIs of this type together with their related items.

Note: Non-CIs are items that do not extend the Configuration Item [cmdb_ci] table. Examples of non-CIs are ports, serial numbers, or Network ARP tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pattern Designer displays attributes of related CI types, which you added to the pattern in the CI attributes pane. You can view and use these attributes while creating or modifying the main CI pattern.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Select the relevant pattern.
3. On the pattern form, select the relevant identification section.
4. Make sure that you are not in Debug mode.
5. On the **CI Attributes** pane, click the **Manage related CIs** icon.

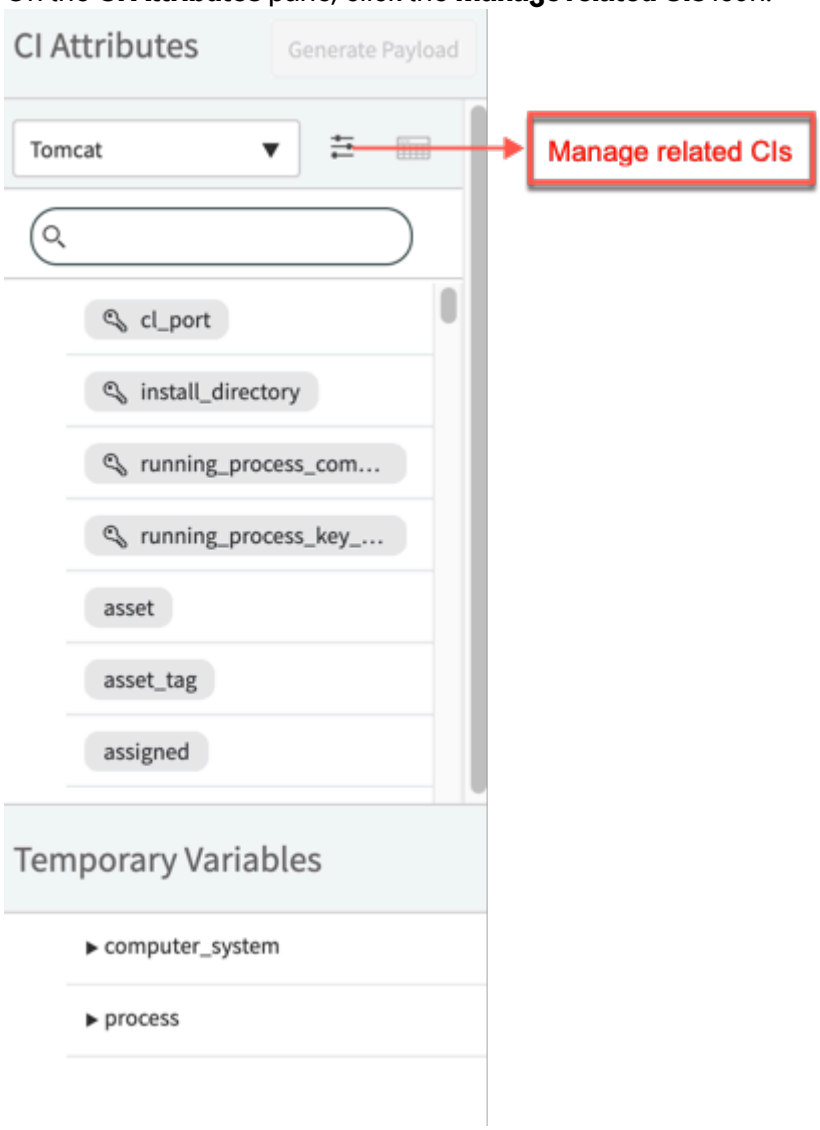

6. Select the relevant related CI type.
   You can use the search box to locate a specific CI type.
7. Click the Plus icon.
   The new related CI type is added to the list.
8. **Optional:** Add other related CI types.
9. Click **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New tabs for related CI types you added appear under the tab for the main CI:

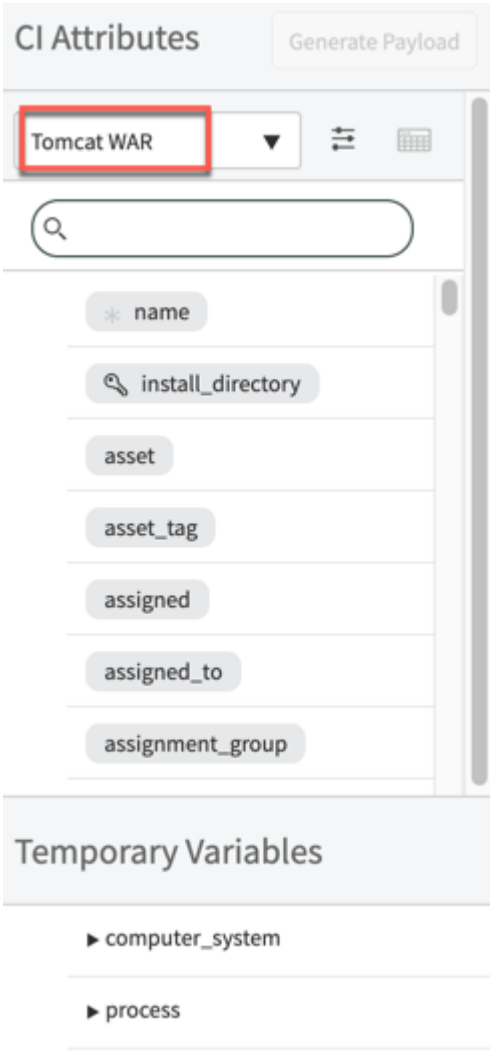

What to do next

Continue editing the pattern by adding more steps as necessary. To view or use attributes of a related CI type, click the relevant tab in the **CI Attributes** pane.

Enhance patterns without modifying identification sections

Enable patterns to search for additional attributes and modify pattern discovery logic defined in identification sections by using extension sections. Each extension section contains a preconfigured set of discovery steps referred to as a shared library.

Before you begin

If your ServiceNow instance uses domain separation and you have access to the global domain, log in to the relevant domain. The selected domain must be a domain without any child domains.

Basic knowledge of programming is desirable.

Role required: pd_admin

About this task

Patterns can be of the infrastructure or application type. Infrastructure patterns are used only by Discovery for creating lists of devices. Application patterns serve both Service Mapping and Discovery, which use the same application patterns for their purposes. Each pattern including

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

its identification and connection sections is an individual file. There are also reusable sequences of discovery steps referred to as shared libraries. Each shared library is an individual file, but it is always part of the infrastructure or application pattern and is not used as a standalone pattern. Records for both patterns and shared libraries reside in the Discovery Patterns [sa_pattern] table.

To modify the discovery logic in identification sections without changing the identification sections themselves, you can add extension sections to the pattern. Every extension section consists of a shared library, either new or existing.

Adding an extension section to a pattern

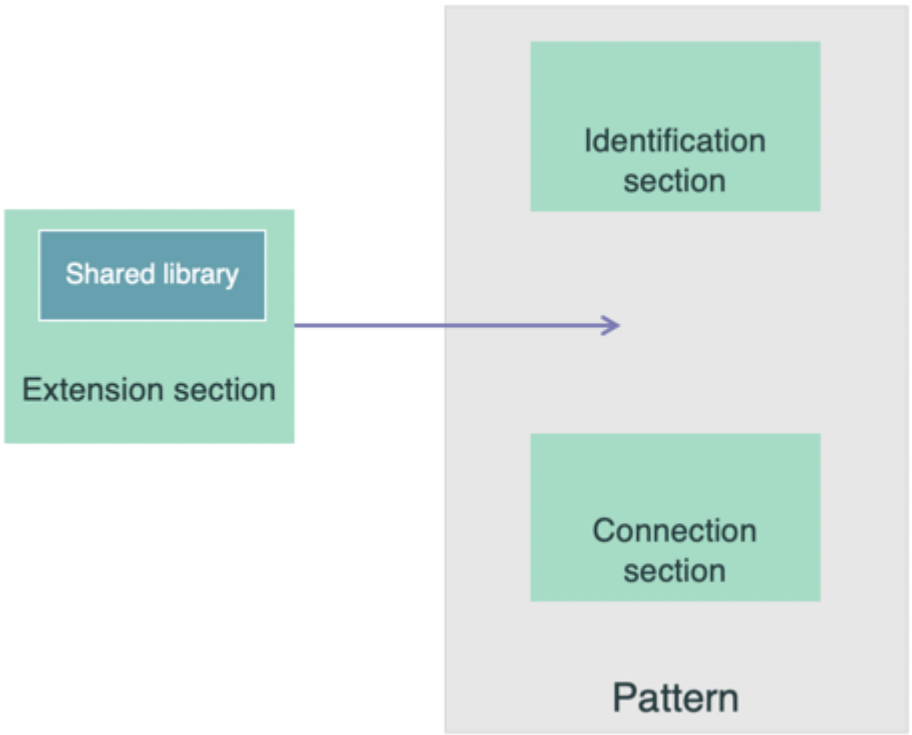

When you create or modify a shared library in the extension section, you save changes only to the shared library, not to the pattern containing this extension section.

Warning: The changes to the shared library affect all patterns using this shared library.

The MID Server runs extension sections after identification sections and only if at least one identification section completes successfully.

If there are more than one extension sections associated with this pattern, define the order in which the MID Server runs these extension sections.

You can deactivate shared libraries, so that patterns using them, skip steps in deactivated shared libraries. For more information, see Discovery Configuration Console. If a precondition defined for a shared library step is not fulfilled, the pattern skips all steps belonging to this shared library. For example, if the file, to which the variable points, does not contain "test", the pattern does not run the shared library steps 1-1 to 1-5. The precondition exists only for the pattern within which you created it. The precondition does not affect other patterns using this shared library.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Library Reference operation form

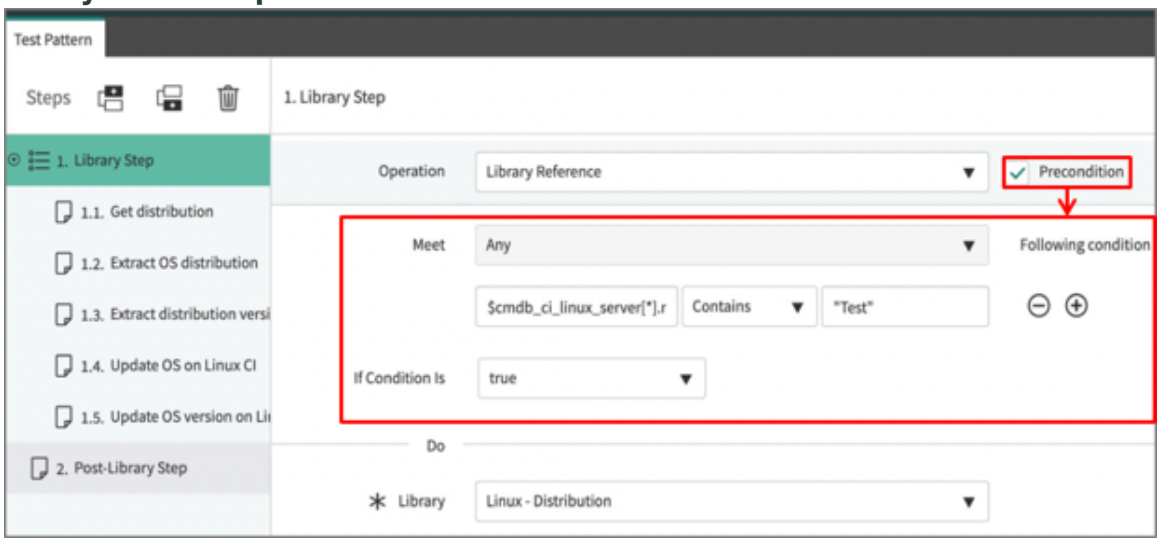

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns** and open the required pattern from the pattern list.

   The pattern can be of the Infrastructure or Application type.

2. To create an extension section and define discovery steps for its shared library.

   a. Click **New** under **Extension section**.

   b. Enter the name of this extension section.

   c. Click **Done**.
   The pattern form displays the new section under **Extension section**.

   d. Click the name of the new extension section.

   The Pattern Designer opens showing the Steps tree on the left.

   If no discovery steps have been identified for this pattern, the **Untitled Step** appears in the Steps tree in the left pane of the window.

   e. Define discovery steps.

   f. Click **Save** to save the extension section and make Service Mapping and Discovery use it. You do not publish extension sections.

3. To use an existing shared library for the new extension section:

   a. Click **Add** under **Extension section**.

   b. Select the shared library from the name list.

4. If there are more than one extension sections, define the order in which the MID Server runs the extension section.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select the check box next to the extension section.

b. Click **Edit** under **Extension section**.

c. Enter a number that determines the order in the **Order** field.
If there is only one extension section, its order is **1** by default.

The section with the lowest order number is used first.

d. Click **Done**.

What to do next
Click **Debug** to run the newly created extension section. Pattern Designer runs extension sections after running identification sections.

Fine-tune patterns using traffic-based discovery

Improve existing patterns so that Service Mapping can find configuration item (CI) connections previously found using traffic-based discovery.

Before you begin
Role required: pd_admin

Basic knowledge of programming is desirable.

About this task
Service Mapping can discover and map configuration items (CIs) following their traffic-based connections. This method is referred to as traffic-based mapping and complements pattern-based mapping.

Using traffic-based discovery, Service Mapping may discover an important CI connection. In this case, you can make sure that this connection is always found for all CIs belonging to the same CI type. You do that by modifying the pattern used to discover CIs of that CI type. A pattern is a sequence of commands designed to detect attributes of a CI and its outbound connections.

You can search CI configuration files to locate data discovered using traffic-based discovery, for example IP address and port. If there are configuration files containing this data, you create a pattern step to parse this file. Next time Service Mapping discovers this CI using the modified pattern, it extracts the information about this connection directly from the configuration file.

Procedure
1. Navigate to **Service Mapping** > **Services** > **Application Services**.
2. Select **View Map** next to the relevant service instance.
3. Right-click the relevant CI.
4. Select **Show traffic based connections**.
The Traffic Based Connections List displays the following information:

Traffic Based Connections List (when discovery based on Predictive Intelligence is disabled)

Field	Description
IP	The IP address of the application connected to the selected CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Port	The port on the selected CI that is used to communicate to the other application.
Process	The ID of the process in the selected CI.
Already on map	○ **Yes** — if this connection shows on the map. ○ **No** — if this connection is not part of the service instance and not on the map.
System decision	The setting defines if Service Mapping keeps the discovered traffic-based connection or removes it. The value comes from the algorithm that Service Mapping uses.
User decision	(Optional) The setting overrides the System decision setting, which defines if Service Mapping keeps the discovered traffic-based connection or removes it. For example, if the System decision setting for a connection is Keep, and you want to remove this connection, select `Remove`.

Note: There may be a case when you can see traffic-based connections on the map, but the Traffic Based Connections list does not display them. It happens for the connections that have been removed from the TCP Connection [cmdb_tcp] table less than three days.

5. Double-click the connection you want to use to create a pattern step.
6. If there are multiple patterns used to discover this CI, select the pattern for which you want to create the step, and click **OK**.

   Pattern Designer creates a temporary connection section for this pattern with a new step.

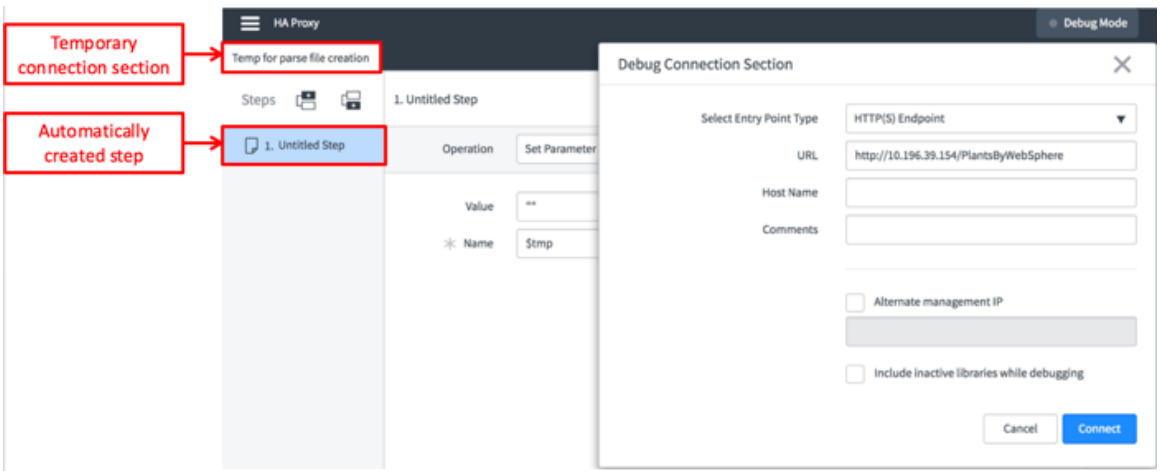

   The Debug Connection Section window is displayed with some fields automatically populated.

7. Click **Connect**.

   The Debug mode is activated.

   The Search Assistant window opens.

8. Search CI configuration files for data received using traffic-based discovery:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Under **Free Text**, review data automatically filled in from the Traffic Based Connections List.

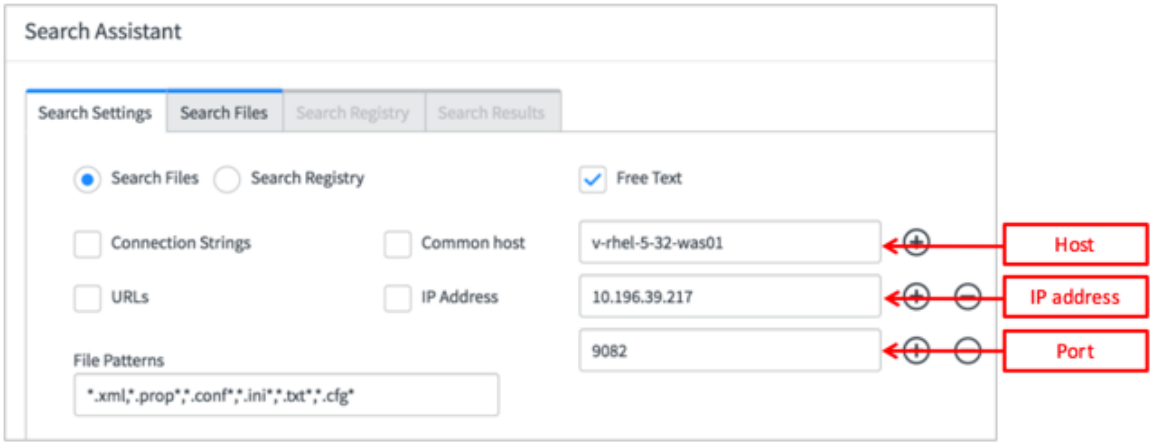

The File Patterns field displays file types in which Pattern Designer performs the search.

b. Click the **Search Files** tab and review files.
If necessary, you can remove irrelevant files from the search by clicking the Minus icon.

c. Click **Search**.

After search completes, the **Search Result** tab opens with the list of files in which any of the search criteria was found.

9. Create a pattern step for parsing the configuration file containing information about the connection:

a. On the Search Result tab, click one of the files and review connection data found in this file.

b. Wait for the file content to load and select the relevant entry under the file name:

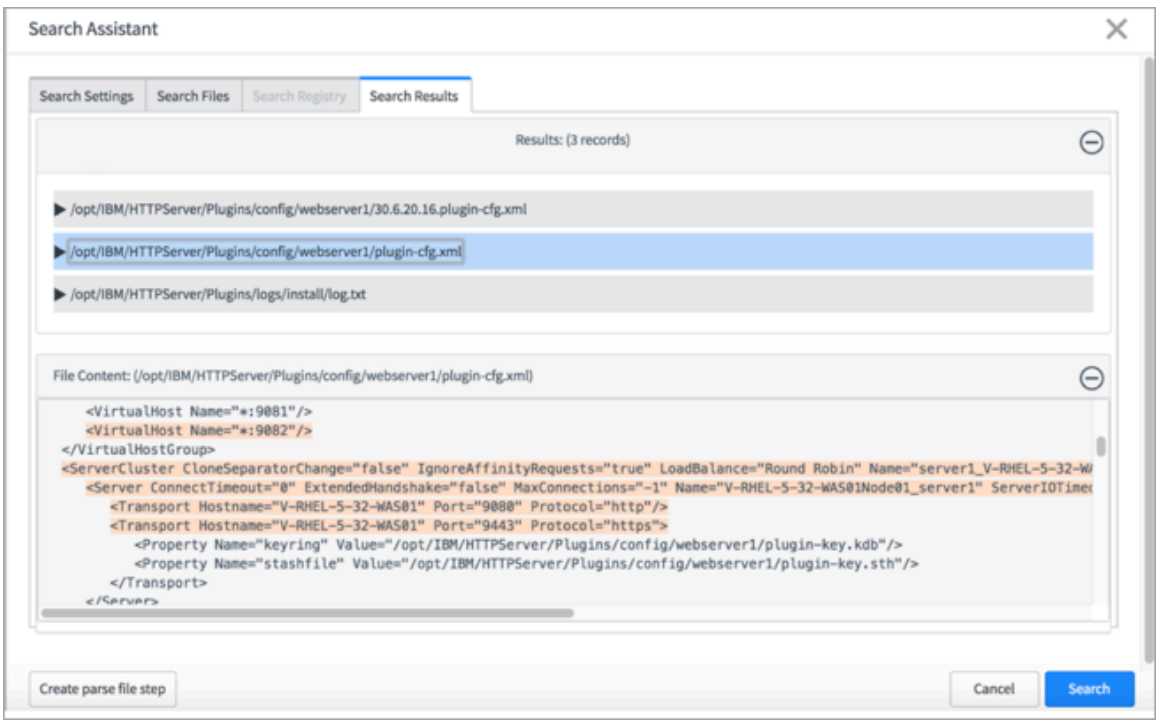

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. To use this data to create the pattern step, click **Create parse file step**.

Pattern Designer shows the new step with the Parse File operation. It automatically assigns this step the name in the following format: Parsing <file name>.

Pattern Designer automatically sets the file name, file path, and parsing strategy for the step. It also creates variables that contain data you searched for. If there is a variable, which contains the file path, Pattern Designer enters this variable for the path.

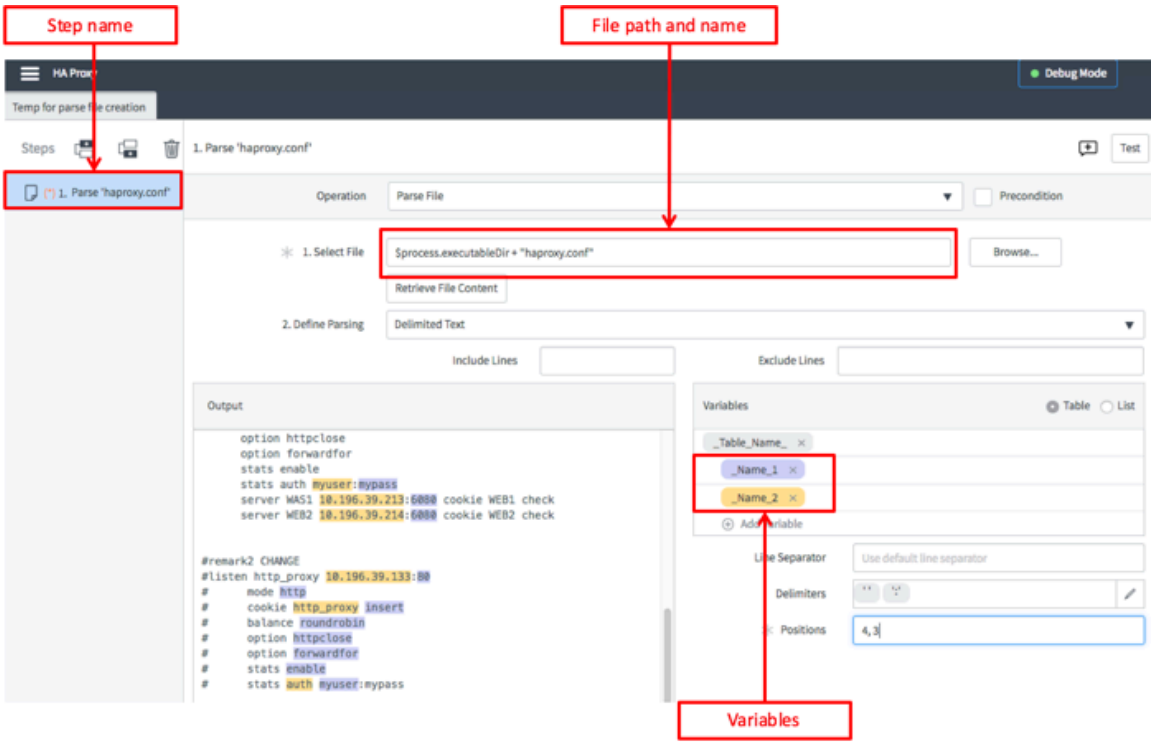

d. Optional: If necessary, you can modify parsing strategy.

e. Rename the section to provide a meaningful name, for example Discover HA Proxy client connection.
Or

Copy and past the step into the relevant connection section.

10. Complete modifying the pattern as described in Finalize a pattern.

Related topics

Traffic-based discovery in Service Mapping

Enable traffic-based discovery for CI types or specific CIs

Discovery patterns used by ITOM Visibility

Finalize a pattern

After you finish defining your pattern, make it ready for use by Service Mapping and Discovery.

Before you begin

Role required: pd_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

There are several ways to finalize a pattern after creating or modifying it:

- Saving — Save all the information you defined for the pattern creating a draft. You can continue working on the pattern draft at any time. Service Mapping and Discovery do not use the pattern draft for discovery and mapping.
- Discarding changes — If you realized that the changes you made to a pattern are incorrect, you can discard them.
- Checking — Actually run the pattern on the relevant configuration item (CI) to check that the operations that you defined for it are correct. Check your pattern before activating it. This option is available only for application patterns used by Service Mapping.
- Publishing — Make Service Mapping and Discovery use this pattern for discovery and mapping. A new version of this pattern is created. The state of the pattern version used previously changes to History.

Procedure

1. On the pattern definition form, complete your pattern definition by performing one or more of the following actions.

To accomplish this	Do this
Save the pattern	Click **Save**.
Discard unwanted changes	Click **Discard**.
Check the pattern **(only for application patterns used by Service Mapping)**	**a.** Click **Check Pattern**. **b.** Enter parameters for connecting to this CI. **c.** Click **Connect**. **d.** Wait for Pattern Designer to run the pattern and check that there are no errors. **e.** Troubleshoot any errors. If there are no errors, you can publish or save your pattern.
Publish the pattern	Click **Publish**.

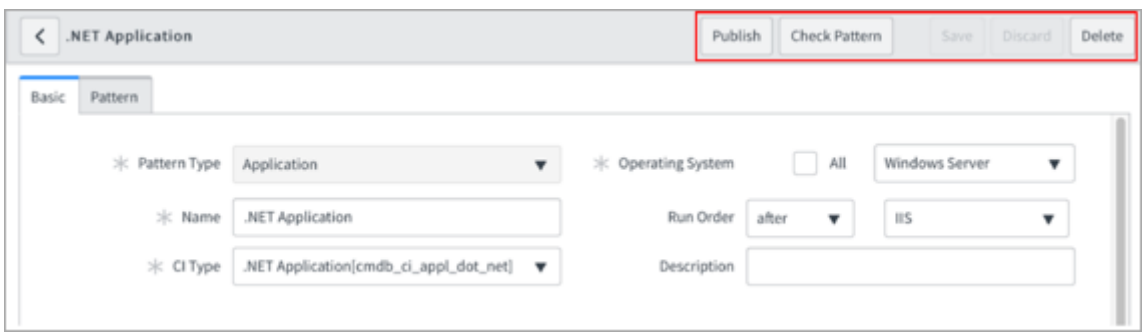

2. Close the Pattern Designer tab.

Copy patterns from one instance to another

To copy patterns from one instance to another, create an update set containing new or modified patterns with their related items in your development instance. Then, import the update set into your production instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: pd_admin

About this task
Typically, you maintain two ServiceNow instances in your organization: for production and for development. Create or modify patterns, test them, and verify results in the development instance. When you are satisfied with the discovery results, export relevant patterns from the development instance to create an update set.

Update sets is a ServiceNow feature for transferring customizations from one instance to another. Each update set is stored in the Update Set [sys_update_set] table. Specific customizations associated with the update set are the entries in the Customer Update [sys_update_xml] table, and they appear as a related list on the update set record. The method of creating update sets for patterns is different from creating update sets for other types of customizations.

Warning: Always follow this procedure to create update sets for patterns. Do not use the standard platform procedure for creating update sets.

Depending on the extent of the changes, you can use the following pattern export methods:

Export patterns

This option creates a new, customized version of the pattern and its related items necessary for the patterns to function properly. It then links them to the update set. The related items include:

- Extension sections
- Tracked files
- Related CI types
- Classification records
- Associated scripts

Note: Update sets do not include changes in script includes used in patterns.

Export patterns with CMDB items

This option creates a new, customized version of the pattern and its related items from the previous option, including any associated CMDB items. It then links all of these to the update set. The CMDB items include:

- Main CI type
- Entry point types
- Lookup tables
- Identifiers
- Hosting and reconciliation rules

The update set is an XML file with the following name convention: sys_remote_update_set<random alphanumeric string>.xml.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the development instance, perform these steps to export patterns:
   a. Navigate to **Pattern Designer** > **Discovery Patterns**.
   b. Select application or infrastructure patterns that you want to export.
      Note: You cannot export shared libraries.
   c. At the bottom of the window, click the **Actions on selected rows** button.
      Note: You cannot use more than 20 patterns at a time for the **Export patterns with CMDB items** method.
   d. Click **Export patterns** or **Export patterns with CMDB items**.
   e. When the system finishes exporting patterns, it creates the update set.
   f. To view the list of items included in the update set, click **View Logs** in the success message.
   g. Click **Close**.
      The update set file in XML format is automatically downloaded on the computer.
   h. Save the update set file locally for future use.
2. In the production instance, perform these steps to import the customized patterns:
   a. Elevate privileges to the security_admin role.
   b. Navigate to **System Update Sets** > **Retrieved Update Sets**.
   c. Select the link **Import Update Set from XML**.
   d. Select **Choose File** and select an XML file.
   e. Select **Upload**.
      The customization is now available as a retrieved update set with state *Loaded*.
3. Preview the imported update set.
   The system scans and validates the update set and displays its contents on the **Customer Updates** tab.
4. Commit the update set.
5. Click **Close** in the success message.
   The production instance contains all new or modified patterns with related items from the update set.
6. Upload the updated set of patterns and related items onto the MID Server:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **Discovery** > **MID Servers**.

b. Alternatively, navigate to **Service Mapping** > **MID Servers**.

c. Click **Pattern Sync to Mid**.

7. If the discovery fails or the result is not as expected, restart the MID Server to make sure that the MID Server runs the latest set of patterns.

Related topics

System update sets

Choose the pattern version

Every time you modify and save a pattern, you create a version of this pattern. Choose which pattern version Service Mapping and Discovery use for discovery.

Before you begin

Role required: admin

About this task

By default, the latest version is used for discovery, but you can choose any other version to use.

If necessary, you can compare pattern versions.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Click the required pattern.
   The form for this pattern opens.
3. Click the **Pattern** tab.
4. Scroll down to the **Versions** section.
   The state of the active pattern version that Service Mapping and Discovery use is **Current**.
5. Click the version you want to use.
6. Click **Revert to this version** under **Related Links**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

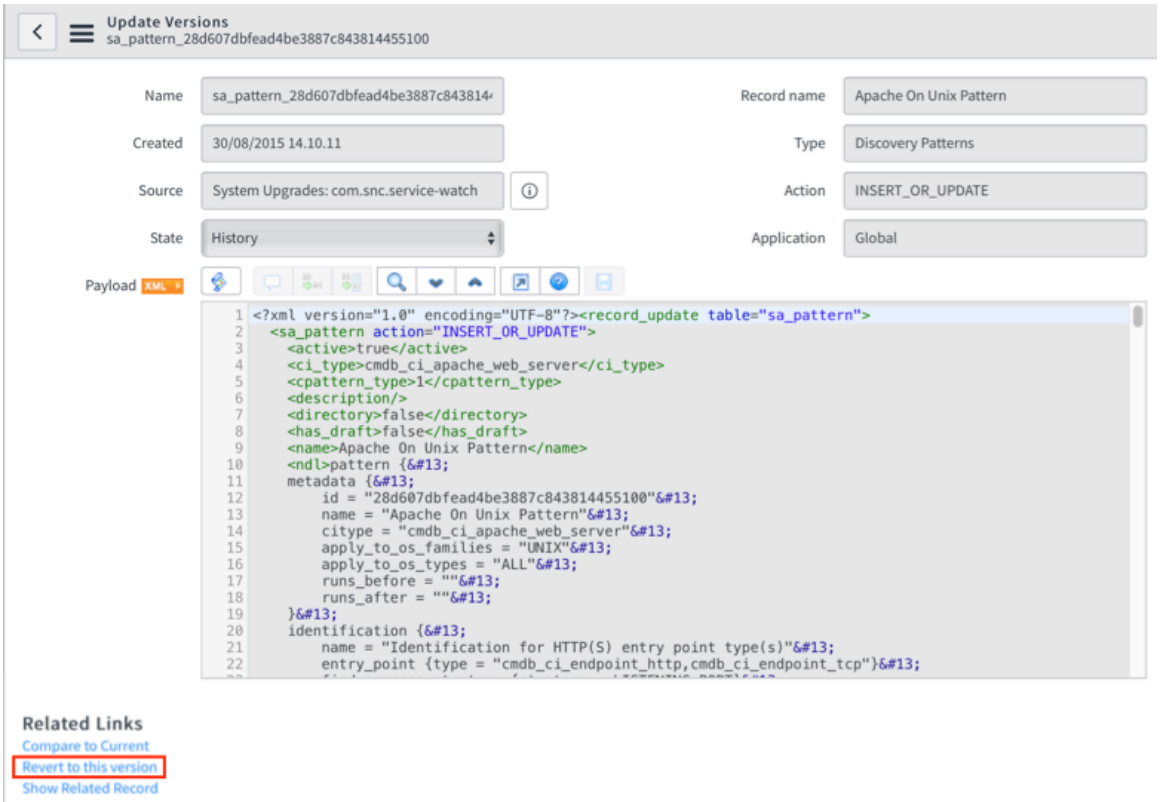

This pattern version becomes the active version that Service Mapping and Discovery use. Its state changes to Current.

Compare pattern versions

If you have multiple versions of the same pattern, you can compare them to decide which pattern version to use for discovery.

Before you begin

Role required: pd_admin

About this task

Service Mapping and Discovery use patterns in their discovery process. Every time you modify and save a pattern, you create a version of this pattern.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. Click the required pattern.
   The form for this pattern opens.
3. Click the **Pattern** tab.
4. Scroll down to the **Versions** section.
   The state of the active pattern version that Service Mapping and Discovery use is **Current**.
5. Click the previous pattern version, which you want to compare to the current version.
6. Click **Compare to Current** under **Related Links**.
7. In the Compare to Current window, click the **Pattern text** pane either under **Selected Version** or **Current Version**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

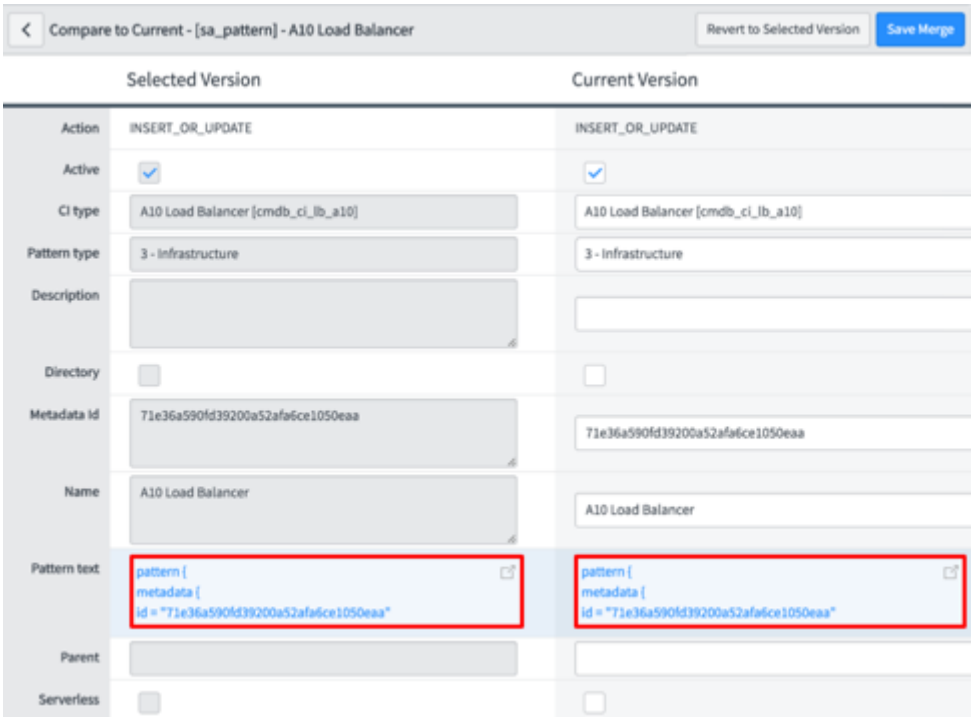

The text comparison view opens showing the text of the selected pattern version and the current pattern version side by side. The step containing a difference appears highlighted.

Highlight colors indicate the type of change:

- Updated step — purple
- New line or characters — green
- Deleted line or characters — red

8. Review the differences in the pattern versions.
9. If necessary, click **Revert chunk** to replace the step in the current version with the step from the previous version.

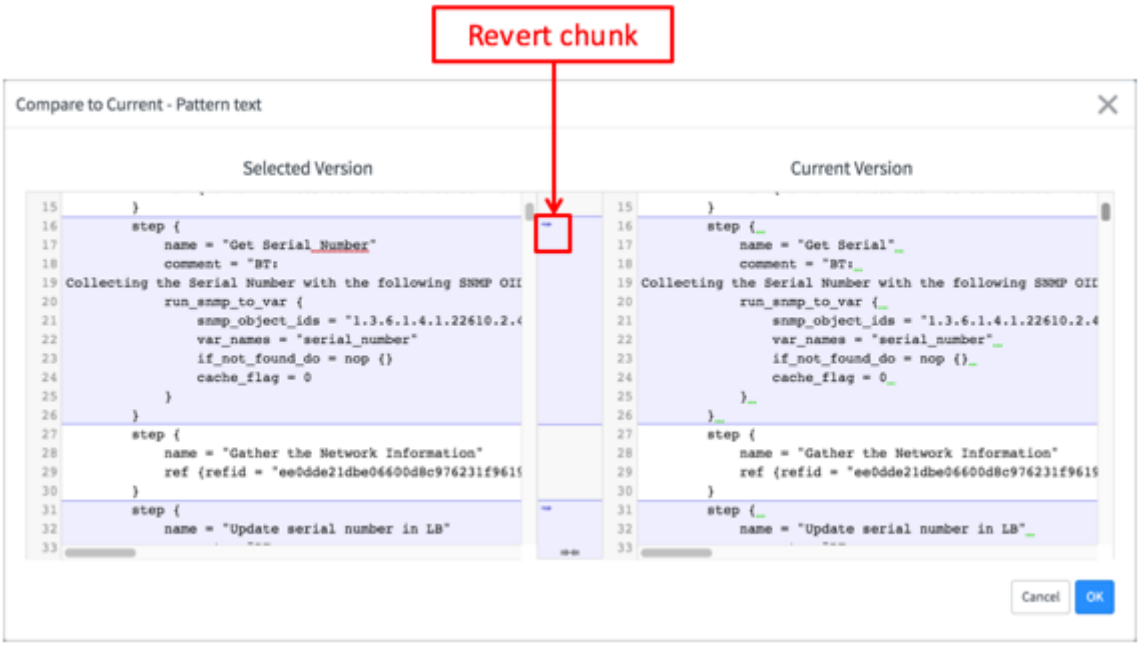

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Click **OK** when done.
11. If you copied any steps from the previous version and want to save the change, click **Save Merge**.

Activate a disabled pattern

If you want to use a pattern for discovery that's disabled by default, activate it manually.

Before you begin

Role required: discovery_admin

About this task

Starting with Visibility Content version 6.28.0, activating or deactivating a pattern won't be considered a customization, and it will continue to receive updates. Patterns that were previously activated or deactivated will reset to the latest predefined version after upgrading while retaining the last active field value.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Discovery Patterns**.
2. In the **Name** field, search for the relevant pattern.
3. Select the pattern.
4. Select the **Active** check box.
5. Select **Update**.

Enable the Cloud Hardware Type class extension

Enable the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table to store the predefined virtual machine (VM) hardware type records for Amazon AWS Cloud Microsoft Azure Cloud and Google Cloud Platform (GCP). This class extension helps prevent duplication of the VM hardware type records in the Hardware Type [cmdb_ci_compute_template] table, which improves performance for Discovery and related flows.

Before you begin

- Veirfy that the latest versions of the following applications are installed:
  - CMDB CI Class Models
  - Discovery and Service Mapping Patterns

    If the version of Discovery and Service Mapping Patterns is earlier than 1.6.0, see the Knowledge Base article KB0955939.
- Ensure that the application scope is the Discovery and Service Mapping Patterns application.

Role required: admin

About this task

Each cloud service provider offers several predefined VM hardware types, such as t1.micro (AWS), and so on. You can also create your own. By default, Discovery stores the discovered VM hardware types in the Hardware Type [cmdb_ci_compute_template] table. This table is a "child" of the cmdb_ci_logical_datacenter class, so the attributes in the Hardware Type table includes the attributes in the cmdb_ci_logical_datacenter class.

This model might create duplicate records for each discovered VM hardware type:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The count of records for each discovered VM hardware type is determined by the number of service accounts multipled by the Number of logical datacenters.

For example, let's assume an organization has an AWS account of t1.micro VM. The organization has 3000 service accounts and 10 Logical Data Centers (LDCs), such as West US. Your next discovery can detect the VM hardware type and create 30,000 records in the Hardware Type [cmdb_ci_compute_template] table.

This approach could slow down your discovery and related flows.

Avoid these issues by storing the predefined AWS and GCP VM hardware types in the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

This table extends the Hardware Type [cmdb_ci_compute_template] table. Because it's not tied to any LDC or service account, Discovery avoids adding many VM hardware type records to the Hardware Type [cmdb_ci_compute_template] table.

The Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table uses the object_id and the provider fields in the identification rule.

Note: In AWS environments, the discovery pattern discovers hardware types only to the master account's datacenter. It then creates relationships of the discovered hardware types to the corresponding regions in member accounts.

- If a discovery schedule is run for both a master account and member accounts, the pattern discovers only the master account's datacenter hardware type.
- If a discovery schedule is run for a standalone member account, the pattern discovers that member account's datacenter and its hardware type.
- If a discovery schedule is run for a master account after a standalone member account has already been discovered, the pattern doesn't rediscover the hardware type for the member account. Instead, it discovers the hardware type for the master account and maps the relationships between the master and member accounts.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** field, search for `sn_itom_pattern.use a single hardware type for cloud data centers`.
3. Select the *`sn_itom_pattern.use a single hardware type for cloud data centers`* property.
4. In the **Value** field, enter `true`.
5. Select **Update**.

What to do next

1. Clean up the Hardware Type records. For more information, see the Knowledge Base article KB1285337.

   Important: The cleanup removes all the existing compute profiles from the CMDB. Therefore, if you have manually added a hardware template to the Hardware Type [cmdb_ci_compute_template] table, you must manually add it to the Cloud Hardware Type [cmdb_ci_cloud_hardware_type] table.

2. Rerun the Cloud Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

AWS discovery using patterns

Google Cloud Platform (GCP) Cloud discovery using Patterns

Microsoft Azure Cloud discovery using patterns

Enable Cloud OS Image discovery

Enable the Cloud OS Image [cmdb_ci_cloud_os_image] table to store the cloud OS image records for Amazon AWS Cloud, Microsoft Azure Cloud, Google Cloud Platform (GCP), Oracle Cloud Infrastructure (OCI), and IBM Cloud.

Before you begin

- Ensure that the latest versions of the following applications are installed:
  - Discovery and Service Mapping Patterns
  - CMDB CI Class Models
- Ensure that the application scope is the CMDB CI Class Models application.

Note: Cloud Provisioning and Governance (CPG) doesn't support the New Cloud Image (cmdb_ci_cloud_os_image) table. If you're currently utilizing CPG Provisioning flows or have intentions to do so, don't migrate to the new Image Model.

Role required: admin

About this task

Each cloud service provider offers OS images to launch a Compute/VM instance. Cloud Discovery stores the discovered Image [cmdb_ci_os_template] table. This table is a dependent class on [cmdb_ci_logical_datacenter] or [cmdb_ci_cloud_service_account]. The new Image model creates one record for each public Image (object id) irrespective of the service account and populates them in the Cloud Image [cmdb_ci_cloud_os_image] table. There for, the new model reduces the number of records in the CMDB

Note: If you have Discovery and Service Mapping Patterns version 1.10.0- November 2023, proceed to the following procedure. If you have an earlier version, see the Knowledge Base article KB1000183.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** field, search for `sn_cmdb_ci_class.use_single_cloud_os_image`.
3. Select the *`sn_cmdb_ci_class.use_single_cloud_os_image`* property.
4. In the **Value** field, enter `true`.
5. Select **Update**.

What to do next

1. Clean up the Image records. For more information, see the Knowledge Base article KB1537376.
2. Rerun the Cloud Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discover datacenters only for new cloud accounts

If you have multiple cloud accounts and datacenters in AWS and Azure, you can discover datacenters for new cloud accounts only, instead of refreshing the entire list.

Before you begin

- Verify that you're using at least Zurich Patch 2.
- Verify that the **glide.discovery.cdu.auto_refresh_sub_accounts_and_ldcs** system property is set to true. For more information, see Create Discovery schedules for cloud resources.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. In the **Name** column, search for `cloud.discovery.refresh_datacenter_for_new_member_account`.
3. Select the **cloud.discovery.refresh_datacenter_for_new_member_account** property.
4. In the **Value** field, enter **true**.
5. Select **Update**.

What to do next

Either execute discovery immediately or wait until the predefined schedule triggers the discovery. For more information, see Create a discovery schedule in Cloud Discovery Workspace.

Related topics

AWS discovery using patterns

Microsoft Azure Cloud discovery using patterns

Improved query performance with direct field population in CI tables

The **Populate Service Account and LDC IN CMDB** scheduled job populates the Service Account and Logical Datacenter fields in cloud configuration item (CI) tables, and the Virtual Machine Object field in the Hardware [cmdb_ci_hardware] table. This direct population reduces query complexity and improves query performance.

Cloud Service Account and Logical Datacenter information is stored separately from cloud CIs. To retrieve account or datacenter information for a cloud resource, queries require joins across multiple tables, including the relationship table that connects CIs. Similarly, Virtual Machine Object information is stored separately from the Hardware CI, requiring additional joins to retrieve data from the Virtual Machine Object [cmdb_ci_vm_object] table. With potentially millions of records in the relationship table, these joins could increase query time for teams working with cloud infrastructure data.

Starting with the Discovery and Service Mapping Patterns version 1.30.2, you can enable a feature that denormalizes cloud CI tables by populating the **Service Account** [cloud_service_account] and **Logical Datacenter** [logical_datacenter] fields directly in cloud CI tables. The scheduled job also references tables extended from the Virtual Machine Object [cmdb_ci_vm_object] table via the Virtual Machine Object reference field in the Hardware [cmdb_ci_hardware] table. Populating these fields directly reduces the need for complex joins, supporting more efficient queries for reporting, analytics, and operational workflows involving cloud resources. After you enable this feature, Discovery populates these fields for both existing and newly discovered CIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query process with and without field population

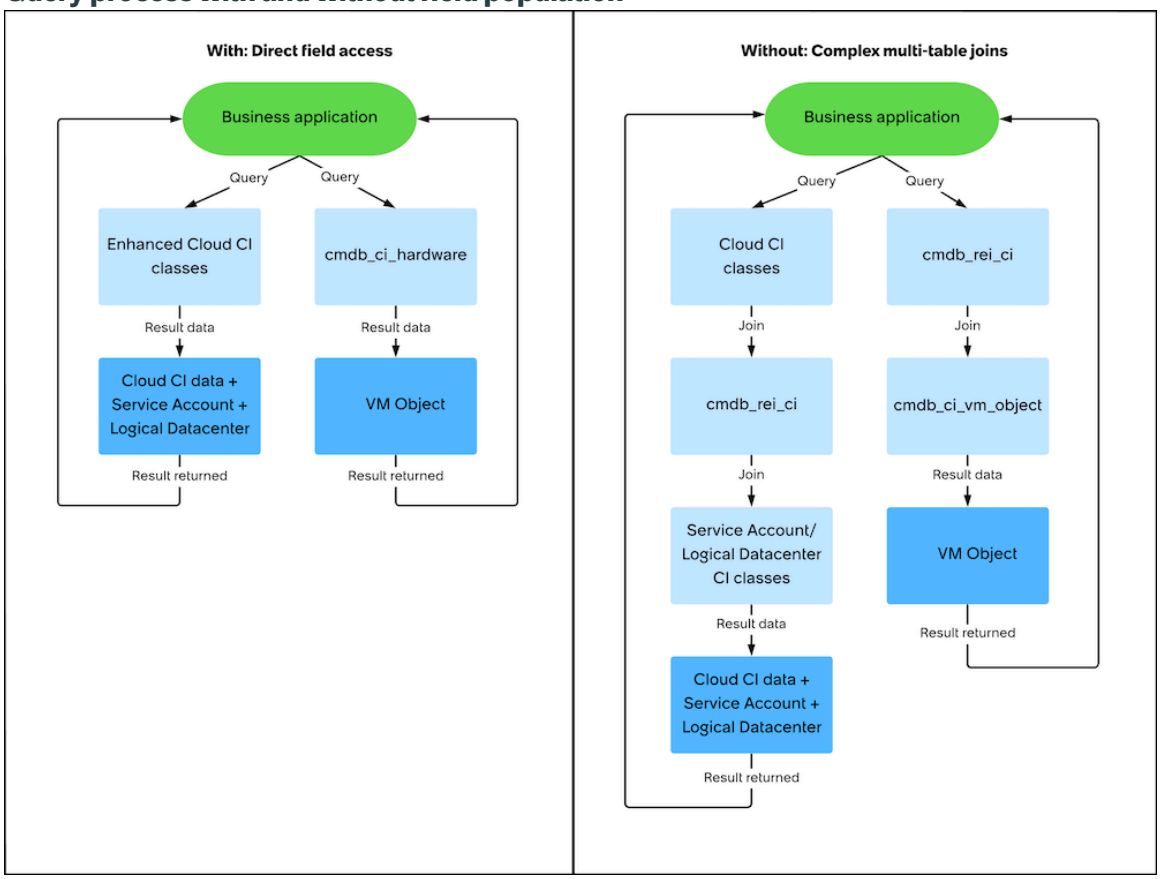

For information about enabling this feature, see Enable direct field population for query performance.

System properties

Populate Service Account and LDC IN CMDB system properties

Property	Description
sn_itom_pattern.populate_saldc.stale_job_threshold_days	Sets the number of days after which a running job is considered stale and automatically canceled. For more information, see Set the Populate Service Account and LDC job stale threshold.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Populate Service Account and LDC IN CMDB system properties (continued)

Property	Description
	• **Type**: Integer • **Default value**: 2
sn_itom_pattern.populate_saldc_full_resync	Starting with Discovery and Service Mapping Patterns version 1.35.0, determines if all CI records are reprocessed on the next job run. For more information, see Trigger a full CI table resync for direct field population. • **Type**: Boolean • **Default value**: False

Supported Cloud CI tables

Service Account and Logical Datacenter fields are added to the following cloud CI tables:

Cloud CI tables with Service Account and Logical Datacenter fields

Table label	Table name
Cloud Subnet	cmdb_ci_cloud_subnet
Cloud Network	cmdb_ci_network
Virtual Machine Instance	cmdb_ci_vm_instance
Hardware Type	cmdb_ci_compute_template
Cloud Hardware Type	cmdb_ci_cloud_hardware_type
Compute Security Group	cmdb_ci_compute_security_group

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud CI tables with Service Account and Logical Datacenter fields (continued)

Table label	Table name
Image	cmdb_ci_os_template
Cloud Image	cmdb_ci_cloud_os_image
Cloud Load Balancer	cmdb_ci_cloud_load_balancer
Cloud Resource	cmdb_ci_cmp_resource
Cloud File Service	cmdb_ci_cloud_file_service
Cloud Mgmt Network Interface	cmdb_ci_nic
Endpoint	cmdb_ci_endpoint
Cloud Function	cmdb_ci_cloud_function
Cloud Object Storage	cmdb_ci_cloud_object_storage
Cloud Gateway	cmdb_ci_cloud_gateway
Database Instance	cmdb_ci_db_instance
Cloud DataBase	cmdb_ci_cloud_database
DynamoDB Table	cmdb_ci_dynamodb_table

Supported virtual machine tables

The following virtual machine tables are referenced in the Virtual Machine Object [cmdb_ci_vm_object] field in the Hardware table:

Referenced virtual machine tables

Table label	Table name
Windows Server	cmdb_ci_win_server
Linux Server	cmdb_ci_linux_server

Related topics

- Available cloud discovery patterns
- Linux discovery
- Windows discovery

Enable direct field population for query performance

Populate Service Account, Logical Datacenter, and Virtual Machine Object fields in configuration item (CI) tables to improve query performance.

Before you begin

Verify that you have at least version 1.30.2 of Discovery and Service Mapping Patterns.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. In the **Name** column, search for `Populate Service Account and LDC IN CMDB`.
3. Select the **Populate Service Account and LDC IN CMDB** job.
4. Select the **Active** check box.
5. Select an option to run the job:
   - Select **Update** to enable the job to run on its configured schedule.
   - Select **Execute Now** to run the batch job immediately.

Related topics

Improved query performance with direct field population in CI tables

Available cloud discovery patterns

Set the Populate Service Account and LDC job stale threshold

Trigger a full CI table resync for direct field population

Set the Populate Service Account and LDC job stale threshold

Configure the number of days before a running *`Populate Service Account and LDC IN CMDB`* job is considered stale and canceled. Increase this value if a large configuration item (CI) dataset causes the job to be canceled before a full run completes.

Before you begin

Verify that you have at least version 1.30.2 of Discovery and Service Mapping Patterns.

Role required: admin

About this task

The *`sn_itom_pattern.populate_saldc.stale_job_threshold_days`* system property controls how long the *`Populate Service Account and LDC IN CMDB`* scheduled job can run before the system considers it stale. If the job exceeds this threshold, the system automatically cancels it and starts a new run. The default value is two days. If your instance has many CI records, you can increase this value to give the job more time to complete.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** field, search for `sn_itom_pattern.populate_saldc.stale_job_threshold_days`.
3. Select the *`sn_itom_pattern.populate_saldc.stale_job_threshold_days`* property.
4. In the **Value** field, enter the number of days after which the job is considered stale.
5. Select **Update**.

What to do next

Wait for the next scheduled run of the *`Populate Service Account and LDC IN CMDB`* job, or execute it immediately. For more information, see Enable direct field population for query performance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Improved query performance with direct field population in CI tables
- Enable direct field population for query performance
- Trigger a full CI table resync for direct field population

Trigger a full CI table resync for direct field population

Trigger the *Populate Service Account and LDC IN CMDB* scheduled job to reprocess all configuration item (CI) records. Configure the *sn_itom_pattern.populate_saldc_full_resync* system property when service accounts or logical datacenters have incorrect or corrupted values for a CI.

Before you begin

Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.

Role required: admin

About this task

The scheduled job normally processes only CI records created since its last run. Setting the *sn_itom_pattern.populate_saldc_full_resync* property to true configures the next job run to reprocess all CI records. The system automatically resets the property to false after the resync completes.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** field, search for `sn_itom_pattern.populate_saldc_full_resync`.
3. Select the *sn_itom_pattern.populate_saldc_full_resync* property.
4. In the **Value** field, enter `true`.
5. Select **Update**.

What to do next

Wait for the next scheduled run of the *Populate Service Account and LDC IN CMDB* job, or execute it immediately. For more information, see Enable direct field population for query performance.

Related topics

- Improved query performance with direct field population in CI tables
- Enable direct field population for query performance
- Set the Populate Service Account and LDC job stale threshold

Server CI population during cloud discovery

Cloud discovery can populate Server CIs without running IP-based discovery, reducing discovery time in large environments.

Server CI cloud discovery overview

By default, cloud discovery creates Virtual Machine (VM) Instance CIs only. To populate Server CIs, organizations run IP-based discovery, which triggers a separate schedule that discovers servers alongside applications and other data. In large environments, IP-based discovery schedules can take a long time. Some organizations require only server records from their cloud environment, without discovering application data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Starting with Discovery and Service Mapping Patterns version 1.31.0, cloud discovery can create Windows Server CIs, Linux Server CIs, or both alongside Virtual Machine Instance CIs, without running IP-based discovery. The *sn_itom_pattern.cloud_discovery_populate_server_ci* system property controls whether Server CIs are created. By default, the property is set to none and Server CIs aren't created. For more information, see Configure Server CI creation during cloud discovery.

The following patterns support Server CI population during cloud discovery:

- Azure – Virtual Machine (LP)
- Azure VM Instance – Uniform Scale Set

Note: Use either Server CI population during cloud discovery or IP-based discovery for the same VMs, but not both. Running both may result in conflicting attribute values.

For AWS Server CI creation, see AWS Linux Server pattern-based discovery and AWS Windows Server pattern-based discovery.

Data collected by Discovery during horizontal discovery

Discovery populates the following data when running the Azure – Create Server CI pattern extension. Virtual Desktop Infrastructure (VDI) instances can't be distinguished from other VMs and are registered as Windows Server or Linux Server CIs.

Windows Server [cmdb_ci_win_server]/Linux Server [cmdb_ci_linux_server]

Field	Description
Name [name]	Hostname of the server. The name is sourced from one of the following options: • If the VM is running: the VM computer name • If the VM is stopped: the OS profile computer name • If the VM is stopped and the Server CI already exists in the CMDB: the existing name is retained
Object ID [object_id]	A unique identifier for the server, matching the Azure VM object ID.
CPU Core Count [cpu_core_count]	Number of CPU cores, derived from the VM size.
CPU Count [cpu_count]	Number of CPUs. Default value is 1.
Disk Space (GB) [disk_space]	Total storage capacity of all disks attached to the VM, in gigabytes (GB).
Install Status [install_status]	Install status of the server. Default value is Installed.
IP Address [ip_address]	IP address of the server, sourced from the primary VM network interface.
Is Virtual [virtual]	Indicates the server is a virtual machine. Default value is true.
Operating System [os]	Operating system name.
Operational Status [operational_status]	Operational status of the server. Default value is Operational.
OS Version [os_version]	Operating system version, if available.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows Server [cmdb_ci_win_server]/Linux Server [cmdb_ci_linux_server] (continued)

Field	Description
RAM (MB) [ram]	Memory capacity in megabytes (MB), derived from the VM size.
Serial Number [serial_number]	Unique identifier for the server, sourced from the Azure VM ID.

Serial Number [cmdb_serial_number]

Field	Description
Serial Number [serial_number]	Unique identifier, sourced from the Azure VM ID.
Serial Number Type [serial_number_type]	Type of serial number. Default value is uuid.
Valid [valid]	Indicates the serial number is valid. Default value is true.
Configuration Item [cmdb_ci]	References the Windows Server [cmdb_ci_win_server] or Linux Server [cmdb_ci_linux_server] table.

CI relationships

The Azure – Create Server CI pattern extension creates the following relationships and references to support Server CI discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Windows Server [cmdb_ci_win_server]	Virtualized by::Virtualizes	Virtual Machine Instance [cmdb_ci_vm_instance]
Linux Server [cmdb_ci_linux_server]	Virtualized by::Virtualizes	Virtual Machine Instance [cmdb_ci_vm_instance]

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	Windows Server [cmdb_ci_win_server]
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	Linux Server [cmdb_ci_linux_server]

Related topics

- Azure virtual machine pattern-based discovery
- Azure Virtual Machine Scale Sets (VMSS) Instance discovery

Configure Server CI creation during cloud discovery

Create Server CIs during cloud discovery without running IP-based discovery, reducing discovery time in large environments.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Verify that you have at least version 1.31.0 of Discovery and Service Mapping Patterns.

Role required: admin

About this task
The *sn_itom_pattern.cloud_discovery_populate_server_ci* property controls whether Server CIs are created during cloud discovery. By default, the property is set to none and Server CIs aren't created.

Procedure
1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `sn_itom_pattern.cloud_discovery_populate_server_ci`.
3. Select the *sn_itom_pattern.cloud_discovery_populate_server_ci* property.
4. In the **Value** field, enter one of the following values.

Value	Description
none	Server CIs aren't created. Default value.
windows_only	Windows Server CIs are created.
linux_only	Linux Server CIs are created.
all	Windows Server CIs and Linux Server CIs are created for all VMs.

5. Select **Update**.

What to do next
Run cloud discovery or wait for the next scheduled discovery run for the changes to apply.

Related topics
Azure virtual machine pattern-based discovery

Azure Virtual Machine Scale Sets (VMSS) Instance discovery

Oracle Wallet authentication for database discovery

With Oracle Wallet authentication, Discovery reads database credentials directly from the target server. This approach keeps credentials within the target environment instead of storing them on the ServiceNow AI Platform.

How Oracle Wallet authentication works

Oracle database discovery uses SQL queries that require database credentials to authenticate against the target Oracle instance. By default, Discovery uses applicative credentials, which uses a database username and password stored on the ServiceNow AI Platform. Oracle Wallet authentication is supported on UNIX and Windows as an alternative to applicative credentials.

Oracle Wallet stores credentials in a wallet on the target server. Discovery reads credentials from the wallet at runtime rather than passing them from the ServiceNow AI Platform. Credentials remain in the target environment and are managed through the native Oracle mechanism.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Minimum application versions for Oracle Wallet authentication

OS	Minimum Visibility Content version	Minimum Discovery and Service Mapping Patterns version	Minimum Data Collection for Oracle Global Licensing and Advisory Services version
UNIX	6.32.0	1.31.0	1.12.0
Windows	6.35.0	1.35.0	1.13.0

For information on enabling Oracle Wallet authentication, see Enable Oracle Wallet authentication for discovery.

Supported discovery types

Oracle Wallet authentication is supported for the following Oracle database discovery types:

- Standard Oracle database discovery. For more information, see Oracle database discovery.
- Container and pluggable database (CDB/PDB) discovery. For more information, see Oracle pluggable database and container database discovery.
- Oracle Global Licensing and Advisory Services (GLAS) discovery. For more information, see Oracle GLAS data collection using patterns.

Enable Oracle Wallet authentication for discovery

Enable Oracle Wallet authentication to use credentials stored on the target server during Oracle database discovery.

Before you begin

For Oracle database discovery on UNIX systems, verify the following requirements:

- The configured Discovery user can connect via SSH to the host.
- Oracle Wallet is configured on the target server for the Discovery user. For more information, see Using Oracle Wallet authentication for Oracle Discovery on Unix/Linux [KB3059862].
- The following applications are up to date:
  - Visibility Content starting with version 6.32.0
  - Discovery and Service Mapping Patterns starting with version 1.31.0
  - For Oracle GLAS discovery: Data Collection for Oracle Global Licensing and Advisory Services starting with version 1.12.0

For Oracle database discovery on Windows systems, verify the following requirements:

- The configured Discovery user can reach and authenticate to the Windows host.
- Oracle Wallet is configured on the target server for the Discovery user. For more information, see Oracle Wallet Authentication for Discovery (Windows) [KB3150013].
- The following applications are up to date:
  - Visibility Content starting with version 6.35.0
  - Discovery and Service Mapping Patterns starting with version 1.35.0
  - For Oracle GLAS discovery: Data Collection for Oracle Global Licensing and Advisory Services starting with version 1.13.0

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Oracle Wallet stores database credentials directly on the target server. When the *glide.discovery.oracle_wallet_authentication* property is set to true, discovery reads credentials from Oracle Wallet instead of from applicative credentials stored on the instance. The property is set to false by default.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `glide.discovery.oracle_wallet_authentication`.
3. Select the *glide.discovery.oracle_wallet_authentication* system property.
4. In the **Value** field, enter `true`.
   To switch back to applicative credentials, set the value to `false`.
5. Select **Update**.

What to do next

To apply the changes, run Discovery.

Related topics

- Oracle Wallet authentication for database discovery
- Oracle database discovery
- Oracle pluggable database and container database discovery
- Oracle GLAS data collection using patterns

Reference the main CI on discovered IP addresses

Populate the **Owned By Configuration Item** field on discovered IP address records with a direct reference to the main CI, instead of requiring a dot-walk through the network adapter. Populating the reference is supported only for AIX Server discovery.

Before you begin

Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.

Role required: discovery_admin

About this task

By default, the relationship between a discovered IP address and its main CI is available only through a dot-walk via the Network Adapter [cmdb_ci_network_adapter] table (nic.configuration_item). When the *add.owned.by.attribute* system property is set to true, the **Owned By Configuration Item** field on the IP Address [cmdb_ci_ip_address] table is populated with a direct reference to the main CI. The default value is false. When enabled, any existing **Owned By Configuration Item** references are cleared and recreated for the discovered IP addresses.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `add.owned.by.attribute`.
3. Select the *add.owned.by.attribute* system property.
4. In the **Value** field, enter `true`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To revert to the default behavior and leave the **Owned By Configuration Item** field unpopulated, set the value to `false`.

5. Select **Update**.

What to do next

Run Discovery again to apply the change.

Related topics

AIX server discovery

Cryptographic Asset Compliance

ServiceNow Cryptographic Asset Compliance helps you inventory, assess, and manage cryptographic assets across cloud and on-premises environments to support post-quantum cryptography (PQC) readiness and help maintain security compliance.

Cryptographic Asset Compliance helps organizations in their goals to future-proof their cryptographic assets like digital certificates and cloud keys. Because quantum computing advances are anticipated to threaten traditional cryptographic algorithms, this application provides visibility and control for PQC migration preparation.

Key capabilities

Description	Capability
Inventory the cryptographic assets that ServiceNow Certificate Inventory and Management and ServiceNow Discovery find, including certificates (on-premises and URL-based) and keys (from AWS KMS, and Azure Key Vault).	Multi-source inventory
Track cryptographic assets at risk with policy-based risk indicators such as outdated algorithms, quantum vulnerability, and expiring certificates.	Risk indicator monitoring
Track PQC compliance and quantum vulnerability through dashboards so you can make data-driven decisions and plan migration to quantum-resistant algorithms.	Post-quantum readiness
Leverage artificial intelligence to analyze cryptographic assets and identify necessary actions.	AI-powered risk assessment

Get started

Explore

Learn about PQC, risk indicators, and how policies evaluate your cryptographic assets.

Configure

Set up the discovery of cryptographic assets and configure the certificate authorities your organization trusts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Monitor

Use the Home page, the inventory, and asset details to monitor asset health, investigate risk, and track PQC readiness.

Reference

Access field descriptions, risk indicator definitions, and technical reference information for cryptographic assets.

Exploring Cryptographic Asset Compliance

Cryptographic Asset Compliance provides centralized visibility and control of cryptographic assets for proactive risk assessment and post-quantum cryptography (PQC) readiness.

Cryptographic Asset Compliance overview

As quantum computing advances are anticipated to threaten traditional cryptographic algorithms, organizations need comprehensive visibility into their cryptographic assets. Cryptographic Asset Compliance addresses this challenge by automatically inventorying cryptographic assets like certificates, AWS KMS keys, and Azure Key Vault keys across cloud and on-premises environments. This visibility helps you assess your PQC readiness and plan migration strategies.

Cryptographic Asset Compliance workflow

The typical workflow when implementing Cryptographic Asset Compliance includes the following activities:

1. Set up discovery: Set up certificate discovery in ServiceNow Certificate Inventory and Management and cloud discovery in ServiceNow Discovery (for AWS and Azure) so cryptographic assets can be found. A scheduled job then syncs the discovered assets into Cryptographic Asset Compliance.
2. Review inventory: Analyze discovered cryptographic assets in the inventory, including certificates and keys across all connected sources.
3. Assess risks: Review the risk indicators identified for your assets, such as weak or quantum-vulnerable algorithms and certificate authority trust or ownership gaps. Use the AI-generated summary and recommendations to understand each asset's risk.
4. Monitor and maintain: Use dashboards to track cryptographic asset health, PQC compliance status, and quantum vulnerable algorithms.

Cryptographic Asset Compliance users

User	Description
SecOperations Engineer	Need visibility into cryptographic asset inventory to identify keys and certificates for PQC migration planning. They monitor asset health, identify vulnerable algorithms, and prioritize remediation.
Compliance and Governance Manager	Require audit-ready reports of cryptographic assets and quantum risk exposure to meet regulatory requirements. They

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

User	Description
	generate compliance reports, track policy adherence, monitor risk indicators, and document remediation efforts.
PKI Engineer	Manage certificate life cycle, configure discovery sources, define risk indicators, and monitor certificate expiration. They contribute to proper cryptographic asset governance and operational continuity.

Cryptographic Asset Compliance benefits

Benefit	Method	Users
Establish a trusted, up-to-date inventory of cryptographic assets and monitor its risk and PQC readiness by configuring asset discovery and the certificate authorities your organization trusts.	• Configuring the discovery of cryptographic assets • Configure trusted certificate authorities • Cryptographic Asset Compliance Home page • View and manage cryptographic assets • View cryptographic asset details	Cryptographic Asset admin sn_itom_cac.admin
Monitor the health of cryptographic assets, investigate risk indicators, and track PQC readiness to support remediation and compliance planning, using read access that doesn't change configuration.	• Cryptographic Asset Compliance Home page • View and manage cryptographic assets • View cryptographic asset details	Cryptographic Asset user sn_itom_cac.user

What to explore next

- Post-quantum cryptography overview
- Cryptographic risk indicators
- About Policy as Code Engine policies

Post-quantum cryptography overview

Post-quantum cryptography (PQC) represents new encryption algorithms that can strongly resist quantum computer attacks, addressing the urgent need for quantum-resistant security.

PQC is a new generation of cryptographic algorithms that resists attacks from quantum computers. Unlike current encryption methods such as Rivest-Shamir-Adleman (RSA) and Elliptic Curve Cryptography (ECC), PQC algorithms can strongly resist computational power of quantum computers using Shor's algorithm.

Quantum computer threat

Quantum computers pose a significant threat to current cryptographic standards with their capability to break traditional encryption methods that protect most digital communications. This

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vulnerability creates a security concern related to attacks that collect encrypted data with the intention of decrypting it after quantum computers are available.

PQC regulations

Cryptographic Asset Compliance checks PQC compliance against the following regulation standards:

1. NIST FIPS 204 (ML-DSA): Module-Lattice-Based Digital Signature Algorithm
2. NIST FIPS 203 (ML-KEM): Module-Lattice-Based Key-Encapsulation Mechanism
3. ISO/IEC 27001 Annex A 8.24: Information security control for the use of cryptography

Cryptographic agility

Cryptographic Asset Compliance supports the ability to quickly and efficiently transition between cryptographic algorithms without significant changes to system architecture. It provides comprehensive visibility into your organization's cryptographic landscape, which enables you to understand your current cryptographic dependencies before transitioning to quantum-resistant algorithms.

With Cryptographic Asset Compliance, your organization can identify which systems and applications rely on vulnerable algorithms, prioritize migration efforts based on risk assessment, and track progress throughout the PQC transition.

Cryptographic risk indicators

Risk indicators highlight cryptographic assets that need attention so you can prioritize remediation and track compliance.

Cryptographic Asset Compliance evaluates assets against policies and raises a risk indicator when an asset meets a risk condition. Risk indicators are based on policies, and each indicator points to a specific weakness, such as a weak algorithm or a missing owner. For more information, see About Policy as Code Engine policies. For more information about risk detection criteria, see Risk indicator definitions.

Supported risk indicators

For certificates:

- Weak algorithm: Uses a deprecated or quantum-vulnerable algorithm.
- Trusted CA risk: Issued by an untrusted or unrecognized certificate authority. For more information, see Configure trusted certificate authorities.
- No owner: Has no assigned owner or responsible party.
- No environment: Is not classified by environment.
- No renewal process: Has no defined renewal workflow.

For AWS KMS keys and Azure Key Vault keys:

- Weak algorithm: Uses a weak or insufficient algorithm or key size.
- Not PQC ready: Does not use a post-quantum-safe algorithm.
- Not rotated (AWS KMS keys): Has not been rotated within the recommended period.
- Pending deletion (AWS KMS keys): Is scheduled for deletion.
- Expired (Azure Key Vault keys): Is expired or lacks a proper expiration date.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Risk levels

Each asset is assigned an overall risk level, calculated from its active indicators. The available risk levels are:

- Critical: Severe vulnerabilities requiring immediate attention.
- High: Significant weaknesses that should be addressed promptly.
- Medium: Moderate concerns that require monitoring and planning.
- Protected: Quantum-safe, or no active risk indicators.

About Policy as Code Engine policies

Cryptographic Asset Compliance uses Policy as Code Engine (PaCE) policies to evaluate your cryptographic assets and flag risk.

Policies evaluate cryptographic assets and raises a risk indicator when an asset meets a risk condition. Separate policies apply to certificates, AWS KMS keys, and Azure Key Vault keys. Note that some policies calculate the overall risk level for an asset by combining the results of the other policies rather than raising an indicator of their own. For more information, see Cryptographic Asset Compliance policies.

Most policies are active by default and run automatically, so you don't have to set them up. However, the certificate authority trust policy is inactive by default because it depends on the certificate authorities that your organization trusts. You activate it after you add those authorities. For more information, see Configure trusted certificate authorities.

Managing policies

You can manage policies using PaCE. You can activate or deactivate a policy and edit its configuration, including the risk criticality that it assigns and review or revert your changes. For more information, see Manage PaCE policies.

Note: Because a change affects how risk is calculated for the assets that the policy evaluates, avoid changing a policy, other than adding trusted certificate authorities.

Configuring Cryptographic Asset Compliance

Set up cryptographic asset discovery and configure trusted certificate authorities (CAs) to prepare for your organization's cryptographic asset monitoring needs.

Configuration overview

Configuring Cryptographic Asset Compliance involves the following actions:

- Configuring the discovery of cryptographic assets- Configuring the discovery of cryptographic assets involves setting up certificate and cloud discovery.
- Configure trusted certificate authorities- Configure certificate authorities (CAs) that your organization trusts and activate the certificate authority trust policy so that certificates from other authorities are flagged as untrusted.

Configuring the discovery of cryptographic assets

Configuring the discovery of cryptographic assets involves setting up certificate and cloud discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cryptographic Asset Compliance inventories the assets that ServiceNow Certificate Inventory and Management and ServiceNow Discovery discovers rather than running discovery of cryptographic assets itself.

You must configure the discovery of cryptographic assets:

- Certificate discovery in Certificate Inventory and Management to catalog certificates (on-premises and URL-based). For more information, see Visibility to TLS certificates.
- Cloud discovery in Discovery to catalog keys from AWS KMS and Azure Key Vault. For more information, see Discovery for cloud environment.

After discovery is set up, a scheduled job syncs discovered assets into Cryptographic Asset Compliance periodically.

Onboarding from the Home page

The Home page includes a guided onboarding that helps you configure cloud and certificate discovery. The banner is collapsible and is displayed until the discovery of all cryptographic assets is configured.

The Certificates card indicates whether certificate discovery is configured by displaying either the Setup required or Configured states.

The Keys card indicates whether cloud discovery is configured by displaying the following states.

- Setup required: Cloud discovery is not configured for AWS KMS and Azure Key Vault.
- Partially configured: Cloud discovery is configured for either AWS KMS or Azure Key Vault.
- Configured: Cloud discovery is configured for both AWS KMS and Azure Key Vault.

Configure trusted certificate authorities

Configure certificate authorities (CAs) that your organization trusts and activate the certificate authority trust policy so that certificates from other authorities are flagged as untrusted.

Before you begin

Role required: Cryptographic Asset admin (sn_itom_cac.admin)

Procedure

1. Navigate to **All** > **Cryptographic Asset Compliance**.
2. Select the Settings icon .
3. In the **Settings** tab, select **Policies** if it is not already selected.
4. In the Policy page, select **Certificate authority trust**.
5. Select the **Policy builder** tab.
6. In the Data sources panel, select **Add**.
7. In the Add API Variable dialog box, fill in the fields.

Field	Description
Label	Display name for the variable, for example`Trusted CA List`.
Name	Internal name that the policy uses to reference the variable.
Type	Data type of the variable, which should be Data Array.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Mandatory	Option that determines whether a value is required for the variable.
Description	Optional notes about the variable.
Default value	The CAs that your organization trusts, as a list of quoted names. For example, `["Test CA 1","Test CA 2"]`.

8. Select **Save**.
9. Publish the policy version by selecting **Publish**.
10. Activate the policy by selecting **Activate**.
11. Confirm the activation by selecting **Activate** in the **Activate Policy** dialog box.

Result
The Certificate authority trust policy is active with the list of CAs that your organization trusts. Certificates issued by an authority not in your trusted list are flagged with the trusted CA risk indicator.

Related topics

Manage PaCE policies

Monitoring cryptographic assets

Monitor your cryptographic posture, investigate risk, and view and manage the assets that Cryptographic Asset Compliance inventories.

- Cryptographic Asset Compliance Home page

  The Cryptographic Asset Compliance Home page provides centralized visibility into cryptographic asset health, risk indicators, key metrics, and actionable insights for proactive compliance management.

- View and manage cryptographic assets

  View details of cryptographic assets and filter the data to focus on specific assets, analyze risk indicators, and export details from the Inventory page.

- View cryptographic asset details

  View detailed information about an individual cryptographic asset, including its risk indicators, cryptographic properties, AI-generated insights, and dependencies.

- AI insights and dependency details

  Analyze an individual cryptographic asset with AI insights powered by ServiceNow Otto and a dependency graph, which shows where the asset is used.

Cryptographic Asset Compliance Home page

The Cryptographic Asset Compliance Home page provides centralized visibility into cryptographic asset health, risk indicators, key metrics, and actionable insights for proactive compliance management.

The Home page presents information to help you assess your post-quantum cryptography (PQC) readiness and plan migration strategies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Metrics

- Cryptographic assets: The total number of inventoried assets.
- PQC compliant: Assets that use post-quantum-safe algorithms.
- At critical risk: Assets that need immediate attention.
- Asset Type: A breakdown of assets by type (Certificate, AWS KMS Key, and Azure Key Vault Key), with colors indicating the risk level: Protected, Medium, High, or Critical.

Recent risk indicators

Recent risk indicator cards summarize current risk conditions such as assets using weak or outdated algorithms or assets with no assigned owners. Each indicator provides an actionable recommendation.

Compliance status

The compliance status is the percentage of PQC compliance for each supported regulation: NIST FIPS 204/ML-DSA, NIST FIPS 203/ML-KEM, and ISO/IEC 27001 Annex A 8.24. The available compliance statuses are: Not compliant, Partially compliant, and Compliant.

Quantum-vulnerable algorithms

Quantum-vulnerable algorithms are ranked by the number of affected assets, with guidance to plan replacement by 2035 per NIST IR 8547. You can select the number of assets to view a filtered list of the assets on the **Inventory** page. For more information, see View and manage cryptographic assets.

Cryptographic asset breakdown

Per-asset-type counts with key risk indicators.

- Certificates: The total, at critical risk, weak algorithm, and trusted CA risk assets.
- AWK KMS Keys: The total, weak algorithms, and pending deletion assets.
- Azure Key Vault Keys: The total, critical risk, weak algorithm, and expired assets.

View and manage cryptographic assets

View details of cryptographic assets and filter the data to focus on specific assets, analyze risk indicators, and export details from the Inventory page.

Before you begin

- Discovery of cryptographic assets must have been set up. For more information, see Configuring the discovery of cryptographic assets.
- Role required: Cryptographic Asset admin (sn_itom_cac.admin) or Cryptographic Asset user (sn_itom_cac.user)

Procedure

1. Navigate to **All** > **Cryptographic Asset Compliance**.
2. In the main navigation, select the Inventory icon .
3. Review the assets' key metrics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric	Description
Cryptographic assets	Total number of discovered cryptographic assets.
PQC compliant	Total number of PQC-compliant and quantum-safe assets.
Not compliant	Total number of cryptographic assets that aren't PQC-compliant.
Critical risk	Total number of cryptographic assets requiring immediate attention.

4. **Optional:** View the assets related to the metric by selecting the metric.
5. **Optional:** Focus on specific assets by applying filters in the **Quick filters** section.

Filter	Description
Asset Type	Filter by asset type: Certificate, AWS KMS Key, Azure Key Vault Key.
Risk level	Filter by risk severity: Protected, Medium, High, and Critical.
Risk indicators	Filter by specific risk conditions, such as outdated algorithm and non-compliant.
Algorithm	Filter by cryptographic algorithm, such as RSA-2048 or ECDSA P-384.
PQC compliant	Filter by whether the assets are PQC compliant.

6. **Optional:** Apply complex filters using a condition builder by selecting **Filter**.
7. **Optional:** Sort the list of assets by a specific category in an ascending or descending order by selecting **Sort by**.
8. **Optional:** Group the list of assets by a specific category by selecting **Group by**.
9. **Optional:** Export data for external analysis or reporting.
   a. Select **Export**.
   b. Select the export format in the **File Type** field.
   The available options are Excel, CSV, JSON, or PDF.
   c. Select whether you want to download or email the file in the **Delivery Type** field.
   d. If you selected the email delivery type, in the **Email** field, enter the email address you want to send the email to.
   e. Select **Export**.
10. View more details about the asset by selecting an asset name.
    For more information, see View cryptographic asset details.

View cryptographic asset details

View detailed information about an individual cryptographic asset, including its risk indicators, cryptographic properties, AI-generated insights, and dependencies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- The discovery of cryptographic assets is set up. For more information, see Configuring the discovery of cryptographic assets.
- Role required: Cryptographic Asset admin (sn_itom_cac.admin) or Cryptographic Asset user (sn_itom_cac.user)

Procedure

1. Navigate to **All** > **Cryptographic Asset Compliance**.
2. In the main navigation, select the Inventory icon .
3. Select the asset name.
4. Select the **Details** tab and review the following information.

Metric	Description
PQC compliant	Whether the asset uses post-quantum-safe algorithms.
Risk indicators	The active risk indicators detected for the asset, such as a weak algorithm or trusted CA risk.
AI Insights	An AI-generated analysis of the asset's risk. For more information, see AI insights and dependency details.
Details	Key attributes of the asset, such as validity dates, owner, team, and environment.
Cryptographic summary	The cryptographic properties of the asset. The displayed fields vary by asset type. For more details, see Cryptographic summary fields.
Lifecycle metrics	The creation, expiry, and last attested date of a certificate. This field is displayed only if the cryptographic asset is a certificate.

5. Review where the asset is used by selecting the **Dependency graph** tab.
   For more information, see AI insights and dependency details.

AI insights and dependency details

Analyze an individual cryptographic asset with AI insights powered by ServiceNow Otto and a dependency graph, which shows where the asset is used.

AI Insights

The AI Insights section in the **Details** tab of an asset displays an AI-generated analysis. It includes an overview of the risk on the asset, the probable cause of each risk indicator, and recommended actions to address it. You can refresh these insights to reflect the latest state of an asset.

You can copy these insights to the clipboard for reuse.

Dependency graph

The **Dependency graph** tab of an asset displays a graph that shows where a cryptographic asset is used. For example, the graph shows the infrastructure on which the asset is deployed on and the applications that use it. You can search the nodes in the graph by their label. The relationships displayed depend on the asset type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For Azure Key Vault keys, the relationships displayed are:

- Key Vault: The key vault the key is stored in.
- Datacenter: The datacenter the key is hosted in.
- Applications: The applications that use the key.
- Resource Group: The resource group the key belongs to.
- Service Account: The cloud service account the key is managed by.

For AWS KMS keys, the relationships displayed are:

- Datacenter: The datacenter the key is hosted in.
- Applications: The applications that use the key.
- Cloud Service Account: The cloud service account the key is managed by.

For certificates, the relationships displayed are:

- Infrastructure: The infrastructure that the certificate is deployed on.
- Applications: The applications that use the certificate.
- Other relations: Other CIs the certificate is related to, for example, a manual endpoint.

Note: URL-based certificates have no dependent infrastructure, so the graph does not display any infrastructure details.

Cryptographic Asset Compliance reference

Reference information for Cryptographic Asset Compliance includes risk indicator definitions, policy descriptions, and field descriptions.

Reference information overview

This reference section provides detailed technical information to support your use of Cryptographic Asset Compliance. Use these references to understand risk indicator definitions, policy descriptions, and field descriptions.

Risk indicators

Reference information for how cryptographic risks are identified and categorized within your environment.

Policy descriptions

A list of the policies available to evaluate cryptographic assets and flag risk.

Field descriptions

Detailed descriptions of fields, data types, and specifications help you understand the information displayed in Cryptographic Asset Compliance interfaces and reports.

Reference topics

- Risk indicator definitions
- Cryptographic Asset Compliance policies
- Cryptographic summary fields

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Risk indicator definitions

Risk indicators are used to identify security vulnerabilities and compliance issues in cryptographic assets.

Certificate risk indicators

Risk Indicator	Description	Detection Criteria
Weak algorithm	Certificate uses cryptographic algorithms that are deprecated or considered weak by current security standards	MD5, SHA-1, RSA < 2048 bits, DSA < 2048 bits, DES, 3DES, RC4
Trusted CA risk	Certificate is not issued by a trusted certificate authority or has CA-related trust issues	Self-signed certificate, CA not in trusted root store, or CA not in organizational trust policy
No owner	Certificate has no identified owner or responsible party	Owner field is empty or not assigned
No environment	Certificate lacks environment classification	Environment field is empty or not specified
No renewal process	Certificate has no defined renewal process or workflow	No renewal process documented or automated workflow configured

Key risk indicators

Risk Indicator	Description	Applies to
Weak algorithm	Key uses a weak or insufficient algorithm or key size	AWS KMS keys, Azure Key Vault keys
Not PQC ready	Key does not use a post-quantum-safe algorithm	AWS KMS keys, Azure Key Vault keys
Not rotated	Key has not been rotated within the recommended period	AWS KMS keys
Pending deletion	Key is scheduled for deletion	AWS KMS keys
Expired	Key is expired or lacks a proper expiration date	Azure Key Vault keys

Cryptographic Asset Compliance policies

Cryptographic Asset Compliance uses policies to evaluate cryptographic assets and flag risk. The policies are available on the Settings page.

Policy name	Description
AWS KMS cumulative risk level calculation	Calculates the risk level for AWS KMS keys based on all risk indicators.
AWS KMS key rotation check	Checks whether AWS KMS keys have been rotated within the recommended period.
AWS KMS PQC readiness	Checks whether an AWS KMS key uses post-quantum-safe algorithms.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Policy name	Description
AWS KMS weak algorithm detection	Identifies AWS KMS keys that use weak or insufficient key algorithms.
Azure Key Vault cumulative risk level calculation	Calculates the risk level for Azure Key Vault keys based on all risk indicators.
Azure Key Vault key expiration check	Checks whether Azure Key Vault keys have proper expiration dates and are not expired.
Azure Key Vault PQC readiness	Checks whether an Azure Key Vault key uses post-quantum-safe algorithms.
Azure Key Vault weak algorithm detection	Identifies Azure Key Vault keys that use weak or insufficient key algorithms and sizes.
Certificate authority trust	Verifies that certificates are issued by trusted, PQC-capable authorities. For more information, see Configure trusted certificate authorities.
Cryptographic inventory & Ownership policy	Identifies unknown or unmanaged certificates.
Cumulative risk level calculation	Calculates the risk level for certificates based on all risk indicators.
PQC readiness	Checks whether a certificate is PQC compliant.
Quantum-vulnerable algorithm detection	Identifies certificates that use cryptographic algorithms that are potentially breakable by quantum computers.

Cryptographic summary fields

The fields displayed in the Cryptographic summary of an asset vary by asset type.

Certificate

Field	Description
Quantum resistant	Whether the certificate uses a quantum-resistant algorithm.
Algorithm	Signing algorithm of the certificate, such as SHA1withRSA.
Key strength	Size of the key in bits.
Hash function	Hash function used by the certificate, such as SHA-384.

AWS KMS key

Field	Description
Object ID	Unique identifier of the key.
Key usage	Cryptographic operations the key supports.
Key state	Current state of the key in AWS KMS.
Origin	Source of the key material.
MAC algorithms	Message Authentication Code (MAC) algorithms associated with the key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Key ID	Identifier of the key in AWS KMS.
Multi region	Whether the key is a multi-region key.
Multi region key type	Type of multi-region key, such as primary or replica.

Azure Key Vault key

Field	Description
Object ID	Unique identifier (URI) of the key in Azure Key Vault.
Key type	Type of cryptographic key, such as RSA.
Key size	Size of the key in bits.
Curve	For elliptic curve fields the elliptic curve of the key.
Algorithm	Cryptographic algorithm associated with the key.
Key operations	Operations the key can perform, such as encryption and decryption.
Enabled	Whether the key is enabled in the source.
Recovery level	Azure Key Vault recovery level, which indicates whether the key can be recovered or is purge able after deletion.
Vault URI	URI of the Azure key vault that stores the key.

AI Agent Topology Mapping

With AI Agent Topology Mapping, you can use patterns to identify AI infrastructure components across cloud platforms in your organization. Discover your AI agents, models, and prompts, and get centralized CMDB visibility to track security compliance and vulnerabilities across your AI deployments.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Get started

Explore	Configure	Reference
Learn more about AI Agent Topology Mapping	Install AI Agent Topology Mapping	Get details about CI tables, fields, and relationships

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Exploring AI Agent Topology Mapping

Learn how AI Agent Topology Mapping discovers AI infrastructure components across cloud platforms using patterns.

AI Agent Topology Mapping overview

AI Agent Topology Mapping extends the pattern-based discovery framework to identify and track AI-specific components in your environment. The application uses patterns to discover AI components from cloud platforms, populating the CMDB with configuration items (CIs). This approach provides centralized visibility into your AI infrastructure alongside traditional IT assets. For more information about how patterns work, see Discovery patterns used by ITOM Visibility.

AI Agent Topology Mapping discovers the following AI components:

- AI Agents: Intelligent entities that perform tasks and orchestrate AI workflows
- AI Models: Foundational models that power AI capabilities, such as large language models
- AI Prompts: Instructions and configurations that guide agent behavior and responses

AI Agent Topology Mapping users

The following user roles have access to patterns or pattern-related modules and can perform various actions. Note that customizing patterns requires basic knowledge of programming.

AI Agent Topology Mapping users and access

User	Description
Discovery admin	Can view, create, edit, and publish patterns. The role enables users to run discovery, migrate probes or CAPI to patterns, and access discovery logs and dashboards.
PD user	Has read-only access to **Discovery Pattern Log**.
PD admin	Can view, create, edit, and publish patterns.
PDE viewer	Starting with Pattern Designer Enhancements version 3.9.0, users can view **Command Validation Tasks**, **Command Validation Tasks Results**, and **Command List**. The pde_viewer can view the **Command Validation Tool** modules and related tables, but doesn't have permissions to modify or edit them. The pde_viewer role can view the following tables only: • Command List [pd_command_list] • Command Validation Task [pd_command_validation] • Command Validation Task Results [pd_command_validation_results] • Pattern Shared Library Mapping [pd_pattern_to_shared_library_mapping] • Temporary Variable Mappings [pd_temp_variable_value_mapping]
PD MID	Not assigned to a user directly but to the MID Server record or the user under which the MID Server runs. The role enables the MID Server to interpret and run pattern-based probes.
MID Server	Can grant the MID Server access to the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Agent Topology Mapping workflow

The following workflow describes how a discovery administrator uses AI Agent Topology Mapping to discover and track AI infrastructure components.

1. Install AI Agent Topology Mapping patterns from the ServiceNow Store.
2. Configure cloud credentials with appropriate permissions for AI platforms.
3. Create or update discovery schedules for environments containing AI resources.
4. Run discovery to identify AI components and populate the CMDB.
5. View discovered AI agents, models, and prompts in CMDB and non-CMDB tables.
6. Monitor discovery logs to verify successful pattern execution.
7. Schedule recurring discovery to maintain up-to-date AI infrastructure inventory.

AI Agent Topology Mapping benefits

AI Agent Topology Mapping benefits

Benefit	Description
Centralized AI infrastructure visibility	Visibility into AI agents, models, and prompts alongside other IT assets in the CMDB, supporting AI Control Tower outcomes.
Automated discovery	Patterns automatically discover AI agents from supported cloud platforms during scheduled runs and populate the CMDB.
Multi-cloud support	Discovery of AI resources across multiple cloud platforms using a single application.
Security and compliance	Tracking of AI component versions and configurations to help maintain governance and compliance requirements for AI deployments.
Vulnerability management	Tracking of AI component versions and dependencies to identify security risks.
Change impact analysis	Near real-time visibility into AI agent dependency and topology relationships to support ITSM and AIOps use cases.
Business context and service mapping	Discovered AI assets serve as the foundation for building tag-based service maps

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Agent Topology Mapping benefits (continued)

Benefit	Description
	aligned with the Common Service Data Model (CSDM).

What to explore next

To learn more about configuring and using AI Agent Topology Mapping, see:

- Configuring AI Agent Topology Mapping
- AI Agent Topology Mapping reference

Configuring AI Agent Topology Mapping

Install AI Agent Topology Mapping from the ServiceNow Store and configure discovery schedules to identify AI infrastructure components.

Configuration overview

1. Install AI Agent Topology Mapping.

   Install the application from the ServiceNow Store to enable AI agent discovery patterns.

2. Review required credentials and permissions.

   See pattern reference topics in AI Agent Topology Mapping reference for credential and permission requirements for your AI platforms.

3. Configure discovery schedules.

   Create discovery schedules for your AI platforms. Patterns run during scheduled discovery to identify AI components and populate CMDB and non-CMDB tables. For more information, see cloud discovery schedule topics in Discovery for cloud environment.

4. View discovery results.

   After discovery runs, view discovered AI components in CMDB and non-CMDB tables. See pattern reference topics in AI Agent Topology Mapping reference for details on data collected.

Install AI Agent Topology Mapping

You can install the AI Agent Topology Mapping application (sn_itom_agnt_ptrn) if you have the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the AI Agent Topology Mapping application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Verify that you have at least version 1.30.2 of Discovery and Service Mapping Patterns.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Multiple tables are installed with AI Agent Topology Mapping.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the AI Agent Topology Mapping application (sn_itom_agnt_ptrn) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   A list of the versions available to you are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

AI Agent Topology Mapping reference

Reference topics provide pattern information for AI Agent Topology Mapping, including prerequisites, tables, fields, and relationships.

Amazon Bedrock pattern-based discovery

AI Agent Topology Mapping discovers Amazon Bedrock AI services, agents, and models during horizontal discovery.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Amazon Bedrock data model

The following diagram illustrates the tables and relationships that the AI Agent Topology Mapping application creates when discovering Amazon Bedrock resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables and relationships populated during Amazon Bedrock discovery

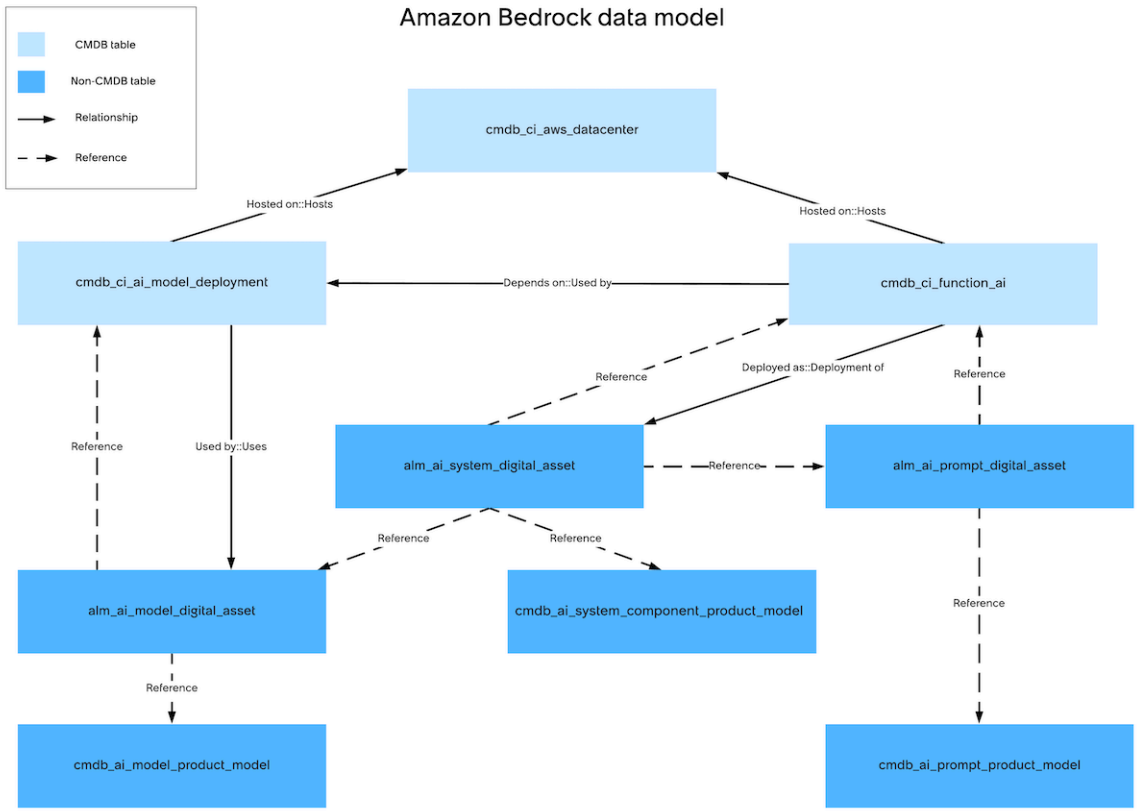

Discovery requirements

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Verify that the applications are up to date

Discovering these resources may require updating to the latest version of the following applications from the ServiceNow Store:

- AI Agent Topology Mapping
- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- Data Foundation Model (at least version 1.13.2)

Update the method used for pointed discovery for the AWS CloudFormation Template (CFT) stack

If you use Cloud Provisioning and Governance, you must update the *`getOperationGR(type)`* method. This update enables the pointed discovery to list the resources correctly for the AWS CFT stack after provisioning. For further information about the steps required to update this method, see the Knowledge Base article KB0858437.

Activate the cloud-related CI relationships

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Set up service accounts on the AWS Management Console

An AWS Organization is a collection of AWS accounts under a single account. In AWS Organizations, parent accounts are called management accounts. The sub-accounts that belong to a management account are called member accounts.

The advantages of using management accounts in Discovery are:

Easy population of member accounts

After you configure the management account and supply the necessary credentials, you can test the connection to the account. If the test succeeds, Discovery returns a list of the member accounts in that management account. From this list, you can choose one or more member account to include in the Discovery of the management account.

(Optional) Discover member resources using dynamically acquired credentials

When you run Discovery on your cloud resources, you don't need separate credentials for each member account. The Cloud Discovery process handles credentials automatically by acquiring a temporary credential for each member via an AWS API. You can elect to use the default configuration or customize the MID Server to assume other roles for additional controls and security.

For more information, see Setup AWS service accounts.

Use IAM user policy on the AWS Management Console

To use the IAM user policy instead of credentials during discovery, configure the MID Server for AWS IAM roles. For more information, see configure the MID Server for AWS IAM roles.

To create the IAM user policy for provisioning AWS resources, see Control AWS access and permissions using policies.

Configure access to the AWS resources

To discover a single account, create an IAM account in the AWS Management Console, and verify that it has the "ReadOnlyAccess" policy applied. To discover several member or child accounts, configure the credentials as described in Access setup for AWS service accounts.

Configure the Discovery schedule to support GovCloud

Discovering AWS GovCloud (US) accounts requires using a datacenter URL when setting up an AWS service account. For more information, see Create AWS service accounts.

Configure a discovery schedule

Create an AWS Discovery schedule in Discovery Admin Workspace.

(Optional) Discover datacenters only for new members

Starting with Zurich Patch 2, you can discover datacenters only for new members added since the last discovery. For more information, see Discover datacenters only for new cloud accounts.

Optimize discovery by including only datacenters with resources

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can optimize discovery by limiting it to only AWS datacenters with resources.

- Verify that your service account has the following IAM permissions to access the Resource Explorer API (starting with Discovery and Service Mapping Patterns version 1.35.0).

  Resource Explorer API permissions and coverage

IAM permission	Coverage
`resource-explorer-2:Search`	Partial
`resource-explorer-2:Search` + `iam:CreateServiceLinkedRole`	Full

  For more information, go to the AWS Documentation and search for the "Getting started with Resource Explorer" article.

- Discover only datacenters with resources by setting the *mid.cloud.discovery.sonar.discover_all_aws_datacenters* MID Server property to **false**. For more information, see Limit AWS discovery to datacenters with resources.
- To exclude specific resource types from AWS datacenter discovery, configure the *sn_itom_pattern.discovery.aws.ldc.excluded_resource_types* system property with a comma-separated list of resource types to exclude (starting with Discovery and Service Mapping Patterns version 1.35.0). For more information, see Exclude AWS resource types from datacenter discovery.

Note: Discovery and Service Mapping Patterns versions 1.29.0 through 1.32.0 used the AWS Config service instead of the Resource Explorer API to determine datacenter activity. For instructions on configuring AWS Config recorder, go to the AWS Documentation and search for the "Recording resources in the AWS Config console" article.

For more information, see the **AWS resources discovery by datacenters** section in AWS discovery using patterns.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Data collection

The AI Agent Topology Mapping application includes the following patterns for Amazon Bedrock discovery:

- Amazon AWS - Bedrock Agents
- Amazon AWS - Bedrock Foundation Model

The Amazon AWS - Bedrock Agents pattern collects AI System Digital Assets, AI Prompt Digital Assets, and Key Value information. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **AWS** > **AWS Bedrock Agents** and selecting the relevant record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Amazon AWS - Bedrock Foundation Model pattern collects AI Model Digital Assets. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **AWS** > **AWS Bedrock Models** and selecting the relevant record.

AI Agent Topology Mapping application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

The AI Agent Topology Mapping application populates data in non-CMDB tables when running the Amazon Bedrock patterns.

If you have the AI Control Tower application, you can view the information in the AI Control Tower Workspace.

AI System Digital Asset [alm_ai_system_digital_asset]

Field	Description
Display name [display_name]	Display name of the Bedrock agent.
External record reference [external_ref_id]	Amazon Resource Name (ARN) of the Bedrock agent.
Model [model]	References the AI System Component Product Model [cmdb_ai_system_component_product_model] table.
Asset type [model_category]	Asset type of the agent. Value is set to **Agentic AI**.
State [install_status]	State of the agent. Value is set to **Deployed**.
Vendor [vendor]	Vendor of the agent. The value is set to **Amazon**.
AI prompts [ai_prompts]	References the AI Prompt Digital Asset [alm_ai_prompt_digital_asset] table.
AI models [ai_models]	References the AI Model Digital Asset [alm_ai_model_digital_asset] table.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

AI System Component Product Model [cmdb_ai_system_component_product_model]

Field	Description
Name [name]	Name of the Bedrock agent.
Short description [short_description]	Description of the Bedrock agent.
Status [status]	Operational status of the agent as returned by Bedrock. For example: PREPARED.
Version [version]	Version of the Bedrock agent.
Manufacturer [manufacturer]	Name of the company defined in the *glide.appcreator.company.friendly_name* system property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI System Component Product Model [cmdb_ai_system_component_product_model] (continued)

Field	Description
Model categories [cmdb_model_category]	Model category. Value is set to **Agentic AI**.

AI Prompt Digital Asset [alm_ai_prompt_digital_asset]

Field	Description
Prompt information [prompt_info]	Instruction text defined for the Bedrock agent.
Model [model]	References the AI Prompt Product Model [cmdb_ai_prompt_product_model] table.
Asset type [model_category]	Asset type of the prompt. Value is set to **AI prompt**.
Vendor [vendor]	Vendor of the prompt. The value is set to **Amazon**.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

AI Prompt Product Model [cmdb_ai_prompt_product_model]

Field	Description
Name [name]	Name of the Bedrock agent, used to identify the associated prompt model.
Manufacturer [manufacturer]	Name of the company defined in the *`glide.appcreator.company.friendly_name`* system property.

AI Model Digital Asset [alm_ai_model_digital_asset]

Field	Description
Display name [display_name]	Name of the foundation model.
External record reference [external_ref_id]	ARN of the foundation model.
Model [model]	References the AI Model Product Model [cmdb_ai_model_product_model] table.
Asset type [model_category]	Asset type of the model. Value is set to **AI model**.
Vendor [vendor]	Provider of the foundation model.
Configuration Item [ci]	References the AI Model Deployment [cmdb_ci_ai_model_deployment] table.

AI Model Product Model [cmdb_ai_model_product_model]

Field	Description
Name [name]	Name of the foundation model.
Model number [model_number]	Model ID assigned by Amazon Bedrock.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Model Product Model [cmdb_ai_model_product_model] (continued)

Field	Description
Description [description]	Description of the foundation model.
Manufacturer [manufacturer]	Provider of the foundation model.
Model categories [cmdb_model_category]	Model category. Value is set to **AI model**.

Data stored in CMDB tables

AI Agent Topology Mapping populates the data in the CMDB when running the Amazon Bedrock patterns.

AI Function [cmdb_ci_function_ai]

Field	Description
Name [name]	Name of the Bedrock agent.
Object ID [object_id]	Concatenation of the agent ARN and version. For example: **arn:aws:bedrock:us-east-1:123456789012:agent/AGENTID123/versions/DRAFT**.
Description [short_description]	Description of the Bedrock agent.
Model number [model_number]	Agent ID assigned by Amazon Bedrock.
Operational status [operational_status]	Operational status of the resource. Default value is set to **Operational**.
Install Status [install_status]	Install status of the resource. Default value is set to **Installed**.
Product instance identifier [product_instance_id]	Version of the Bedrock agent.
Comments [comments]	Resource type identifier. Value is set to **AWS::Bedrock::Agent**.
Manufacturer [manufacturer]	Manufacturer of the resource. Value is set to **AWS**.

AI Model Deployment [cmdb_ci_ai_model_deployment]

Field	Description
Name [name]	Name of the Amazon Bedrock foundation model.
Object ID [object_id]	ARN of the foundation model.
Description [short_description]	Description of the foundation model.
Model number [model_number]	Model ID assigned by Amazon Bedrock.
Manufacturer [manufacturer]	Provider of the foundation model.
Attributes [attributes]	Lifecycle status of the model as reported by the API. For example: ACTIVE or LEGACY.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Model Deployment [cmdb_ci_ai_model_deployment] (continued)

Field	Description
Comments [comments]	Resource type identifier. Value is set to **AWS::Bedrock::Foundation Model**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.

Dependency Views map

On the Dependency Views map, you can view discovered Amazon Bedrock resources and the relationships between them.

Note: The following relationships are stored in the Asset-CI Relationship [cmdb_rel_asset_ci] table and aren't represented in the Dependency View map:

- The **Deployed as::Deployment of** relationship between AI Function [cmdb_ci_function_ai] and AI System Digital Asset [alm_ai_system_digital_asset].
- The **Used by::Uses** relationship between AI Model Deployment [cmdb_ci_ai_model_deployment] and AI Model Digital Asset [alm_ai_model_digital_asset].

Amazon Bedrock dependency view

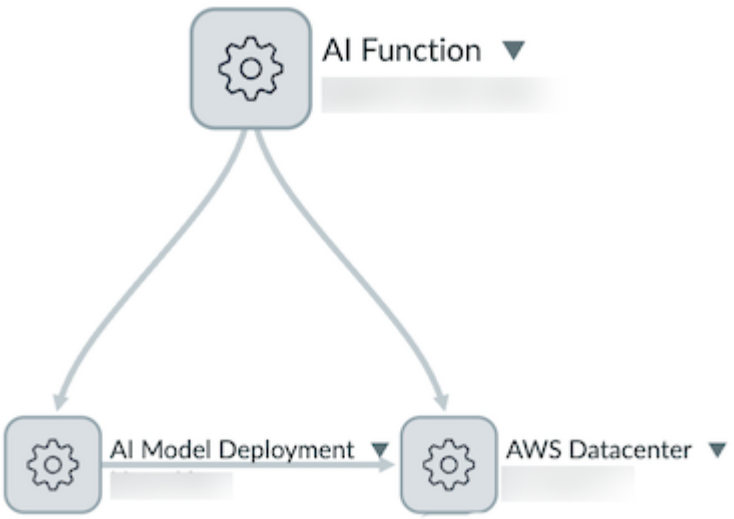

CI relationships and references

The Amazon Bedrock patterns create the following relationships and references to support Amazon Bedrock discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI/Table	Relationship	CI/Table
AI Function [cmdb_ci_function_ai]	Deployed as::Deployment of	AI System Digital Asset [alm_ai_system_digital_asset]
AI Function [cmdb_ci_function_ai]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AI Function [cmdb_ci_function_ai]	Depends on::Used by	AI Model Deployment [cmdb_ci_ai_model_deployment]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Hosted on::Hosts	AWS Datacenter [cmdb_ci_aws_datacenter]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Used by::Uses	AI Model Digital Asset [alm_ai_model_digital_asset]

CI references

CI/Table	Field	Referenced CI/Table
AI System Digital Asset [alm_ai_system_digital_asset]	Model [model]	AI System Component Product Model [cmdb_ai_system_component_product_model]
AI System Digital Asset [alm_ai_system_digital_asset]	AI prompts [ai_prompts]	AI Prompt Digital Asset [alm_ai_prompt_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	AI models [ai_models]	AI Model Digital Asset [alm_ai_model_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Model [model]	AI Prompt Product Model [cmdb_ai_prompt_product_model]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]
AI Model Digital Asset [alm_ai_model_digital_asset]	Model [model]	AI Model Product Model [cmdb_ai_model_product_model]
AI Model Digital Asset [alm_ai_model_digital_asset]	Configuration Item [ci]	AI Model Deployment [cmdb_ci_ai_model_deployment]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	AI Function [cmdb_ci_function_ai]

AWS tag discovery

The Amazon AWS - Bedrock Agents pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Value [cmdb_key_value] (continued)

Field	Description
Configuration item [configuration_item]	References the AI Function [cmdb_ci_function_ai] table.

Related topics

AWS discovery using patterns

Microsoft Foundry (classic) pattern-based discovery

AI Agent Topology Mapping discovers Microsoft Foundry (classic) services, assistants, and models during horizontal discovery.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Microsoft Foundry data model

The following diagram illustrates the tables and relationships that the AI Agent Topology Mapping application creates when discovering Microsoft Foundry resources.

Tables and relationships populated during Microsoft Foundry discovery

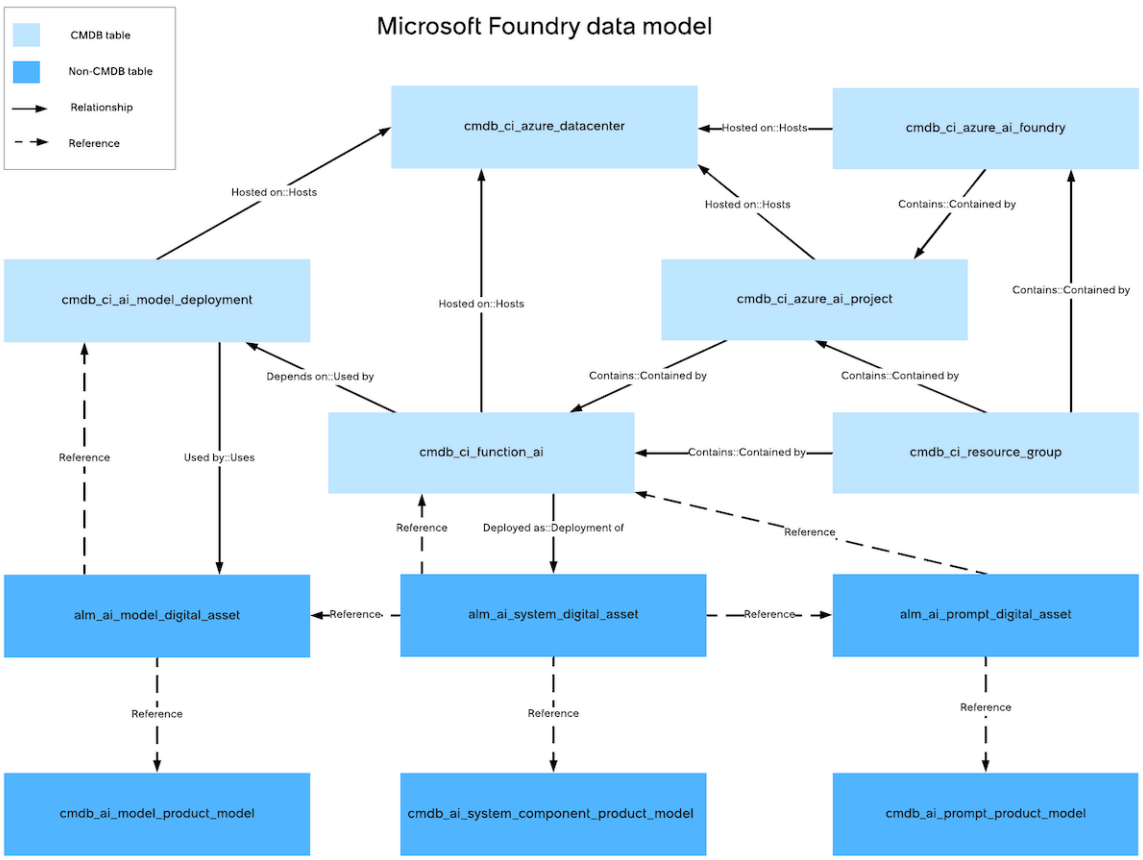

Discovery requirements

Verify the REST API permissions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Verify that the applications are up to date

Discovering these resources may require updating to the latest version of the following applications from the ServiceNow Store:

- AI Agent Topology Mapping
- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- Data Foundation Model (at least version 1.13.2)

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Azure Availability Set

Wait for the **Clean-Up job for Availability zone to clear availability set record** schedule job to delete all the pre-populated Configuration Items (CIs) in the **cmdb_ci_azure_availability_set** table.

Azure Availability Zone

To run a discovery with Azure Availability Zone, register the subscription ID to the **AvailabilityZonePeering** feature with AZ CLI using `az feature register -n AvailabilityZonePeering --namespace Microsoft.Resources` to use the Check Zone Peering API. Check the status with `az feature show -n AvailabilityZonePeering --namespace Microsoft.Resources` before running discovery.

Set up Azure service accounts

Enable Cloud Discovery to access your Azure environment.

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Create an Azure cloud discovery schedule

For more information, see Create an Azure Discovery schedule in Discovery Admin Workspace.

(Optional) Discover datacenters only for new subscriptions

Starting with Zurich Patch 2, you can discover datacenters only for new subscriptions added since the last discovery. For more information, see Discover datacenters only for new cloud accounts.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collection

The AI Agent Topology Mapping application includes the following patterns for Microsoft Foundry discovery:

- Azure - AI Foundry Agents
- Azure - AI Service Foundry
- Azure - AI Service Foundry Project
- Azure - AI Foundry Models

The Azure - AI Foundry Agents pattern collects AI System Digital Assets, AI Prompt Digital Assets, and Key Value information. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **Azure** > **Azure AI Foundry Agents** and selecting the relevant record.

The Azure - AI Foundry Models pattern collects AI Model Digital Assets. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **Azure** > **Azure AI Foundry Models** and selecting the relevant record.

The AI Agent Topology Mapping application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

The AI Agent Topology Mapping application populates data in non-CMDB tables when running the Microsoft Foundry patterns.

If you have the AI Control Tower application, you can view the information in the AI Control Tower Workspace.

AI System Digital Asset [alm_ai_system_digital_asset]

Field	Description
Display name [display_name]	Display name of the Microsoft Foundry assistant.
External record reference [external_ref_id]	Unique identifier of the Microsoft Foundry assistant.
Asset type [model_category]	Asset type of the assistant. Value is set to **Agentic AI**.
State [install_status]	State of the assistant. Value is set to **Deployed**.
Vendor [vendor]	Vendor of the assistant.Value is set to **Azure**.
Model [model]	References the AI System Component Product Model [cmdb_ai_system_component_product_model] table.
AI prompts [ai_prompts]	References the AI Prompt Digital Asset [alm_ai_prompt_digital_asset] table.
AI models [ai_models]	References the AI Model Digital Asset [alm_ai_model_digital_asset] table.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI System Component Product Model [cmdb_ai_system_component_product_model]

Field	Description
Name [name]	Name of the Microsoft Foundry assistant.
Short description [short_description]	Description of the Microsoft Foundry assistant.
Status [status]	Operational status of the assistant. Value is set to **PREPARED**.
Version [version]	Version of the Microsoft Foundry assistant.
Manufacturer [manufacturer]	Name of the company defined in the *glide.appcreator.company.friendly_name* system property.
Model categories [cmdb_model_category]	Model category. Value is set to **Agentic AI**.

AI Prompt Digital Asset [alm_ai_prompt_digital_asset]

Field	Description
Prompt information [prompt_info]	Instruction text defined for the Microsoft Foundry assistant.
Asset type [model_category]	Asset type of the prompt. Value is set to **AI prompt**.
Vendor [vendor]	Vendor of the prompt. Value is set to **Azure**.
Model [model]	References the AI Prompt Product Model [cmdb_ai_prompt_product_model] table.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

AI Prompt Product Model [cmdb_ai_prompt_product_model]

Field	Description
Name [name]	Name of the Microsoft Foundry assistant, used to identify the associated prompt model.
Manufacturer [manufacturer]	Name of the company defined in the *glide.appcreator.company.friendly_name* system property.
Model categories [cmdb_model_category]	Model category. Value is set to **AI prompt**.

AI Model Digital Asset [alm_ai_model_digital_asset]

Field	Description
Display name [display_name]	Name of the foundation model.
External record reference [external_ref_id]	Resource identifier of the foundation model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Model Digital Asset [alm_ai_model_digital_asset] (continued)

Field	Description
Asset type [model_category]	Asset type of the model. Value is set to **AI model**.
Vendor [vendor]	Provider of the foundation model.
Model [model]	References the AI Model Product Model [cmdb_ai_model_product_model] table.
Configuration Item [ci]	References the AI Model Deployment [cmdb_ci_ai_model_deployment] table.

AI Model Product Model [cmdb_ai_model_product_model]

Field	Description
Name [name]	Name of the foundation model.
Model number [model_number]	Model ID assigned by Microsoft Foundry.
Version [version]	Version of the foundation model.
Description [description]	Description of the foundation model.
Manufacturer [manufacturer]	Provider of the foundation model.
Model categories [cmdb_model_category]	Model category. Value is set to **AI model**.

Data stored in CMDB tables

The AI Agent Topology Mapping application populates data in the CMDB when running the Microsoft Foundry patterns.

AI Function [cmdb_ci_function_ai]

Field	Description
Name [name]	Name of the Microsoft Foundry assistant.
Object ID [object_id]	Unique identifier of the Microsoft Foundry assistant.
Description [short_description]	Description of the Microsoft Foundry assistant.
Manufacturer [manufacturer]	Cloud provider of the resource. Value is set to **Azure**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Comments [comments]	Resource type identifier. Value is set to **Azure::Foundry::Agent**.

Azure AI Foundry [cmdb_ci_azure_ai_foundry]

Field	Description
Name [name]	Name of the Microsoft Foundry service.
Object ID [object_id]	Unique identifier of the Microsoft Foundry service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure AI Foundry [cmdb_ci_azure_ai_foundry] (continued)

Field	Description
Description [short_description]	Description of the resource. Value is set to **Azure Foundry Service**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Azure Service Type [azure_service_type]	Azure resource type. For example: microsoft.cognitiveservices/accounts.
Azure Service Kind [azure_service_kind]	Resource kind of the Microsoft Foundry service. For example: AIServices.

Azure AI Project [cmdb_ci_azure_ai_project]

Field	Description
Name [name]	Name of the Microsoft Foundry project.
Object ID [object_id]	Unique identifier of the Microsoft Foundry project.
Comments [comments]	Azure region where the Microsoft Foundry project is deployed.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Model number [model_number]	Internal identifier of the Microsoft Foundry project.
Azure Project Kind [azure_project_kind]	Resource kind of the Microsoft Foundry project. For example: AIServices.

AI Model Deployment [cmdb_ci_ai_model_deployment]

Field	Description
Name [name]	Name of the AI model available in Microsoft Foundry.
Object ID [object_id]	Resource identifier of the AI model.
Description [short_description]	Description of the AI model.
Model number [model_number]	Model name, used as an identifier.
Manufacturer [manufacturer]	Provider of the AI model. For example: Microsoft, Anthropic, or OpenAI.
Attributes [attributes]	Lifecycle status of the model as reported by the API.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Model Deployment [cmdb_ci_ai_model_deployment] (continued)

Field	Description
	For example: Stable or Deprecated.
Comments [comments]	Azure resource type identifier. Value is set to **Microsoft.CognitiveServices/locations/models**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.

Dependency Views map

On the Dependency Views map, you can view discovered Microsoft Foundry resources and the relationships between them.

Note: The following relationships are stored in the Asset-CI Relationship [cmdb_rel_asset_ci] table and aren't represented in the Dependency View map:

- The **Deployed as::Deployment of** relationship between AI Function [cmdb_ci_function_ai] and AI System Digital Asset [alm_ai_system_digital_asset].
- The **Used by::Uses** relationship between AI Model Deployment [cmdb_ci_ai_model_deployment] and AI Model Digital Asset [alm_ai_model_digital_asset].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Foundry dependency view

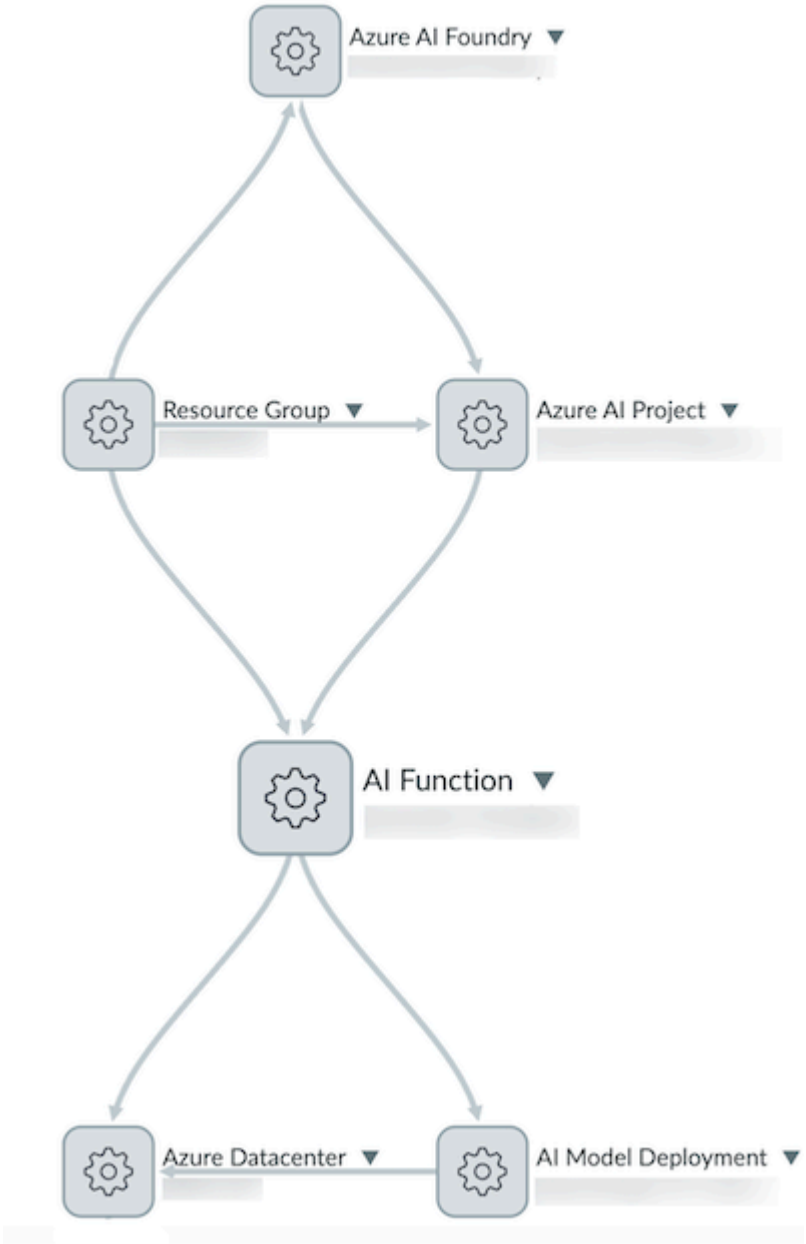

CI relationships and references

The Microsoft Foundry patterns create the following relationships and references to support Microsoft Foundry (classic) discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI/Table	Relationship	CI/Table
AI Function [cmdb_ci_function_ai]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure AI Project [cmdb_ci_azure_ai_project]	Contains::Contained by	AI Function [cmdb_ci_function_ai]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI/Table	Relationship	CI/Table
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	AI Function [cmdb_ci_function_ai]
AI Function [cmdb_ci_function_ai]	Deployed as::Deployment of	AI System Digital Asset [alm_ai_system_digital_asset]
Azure AI Foundry [cmdb_ci_azure_ai_foundry]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Azure AI Foundry [cmdb_ci_azure_ai_foundry]
Azure AI Foundry [cmdb_ci_azure_ai_foundry]	Contains::Contained by	Azure AI Project [cmdb_ci_azure_ai_project]
Azure AI Project [cmdb_ci_azure_ai_project]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Azure AI Project [cmdb_ci_azure_ai_project]
AI Function [cmdb_ci_function_ai]	Depends on::Used by	AI Model Deployment [cmdb_ci_ai_model_deployment]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Used by::Uses	AI Model Digital Asset [alm_ai_model_digital_asset]

CI references

CI/Table	Field	Referenced CI/Table
AI System Digital Asset [alm_ai_system_digital_asset]	Model [model]	AI System Component Product Model [cmdb_ai_system_component_product_model]
AI System Digital Asset [alm_ai_system_digital_asset]	AI prompts [ai_prompts]	AI Prompt Digital Asset [alm_ai_prompt_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	AI models [ai_models]	AI Model Digital Asset [alm_ai_model_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Model [model]	AI Prompt Product Model [cmdb_ai_prompt_product_model]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]
AI Model Digital Asset [alm_ai_model_digital_asset]	Model [model]	AI Model Product Model [cmdb_ai_model_product_model]
AI Model Digital Asset [alm_ai_model_digital_asset]	Configuration Item [ci]	AI Model Deployment [cmdb_ci_ai_model_deployment]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI/Table	Field	Referenced CI/Table
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Azure AI Foundry [cmdb_ci_azure_ai_foundry]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Azure AI Project [cmdb_ci_azure_ai_project]

Azure tag discovery

The Azure - AI Service Foundry pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Azure AI Foundry [cmdb_ci_azure_ai_foundry] table.

The Azure - AI Service Foundry Project pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Azure AI Project [cmdb_ci_azure_ai_project] table.

Related topics

Microsoft Azure Cloud discovery using patterns

Microsoft Foundry (new) pattern-based discovery

AI Agent Topology Mapping discovers Microsoft Foundry (new) services, agents, and models during horizontal discovery.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Microsoft Foundry data model

The following diagram illustrates the tables and relationships that the AI Agent Topology Mapping application creates when discovering Microsoft Foundry resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables and relationships populated during Microsoft Foundry discovery

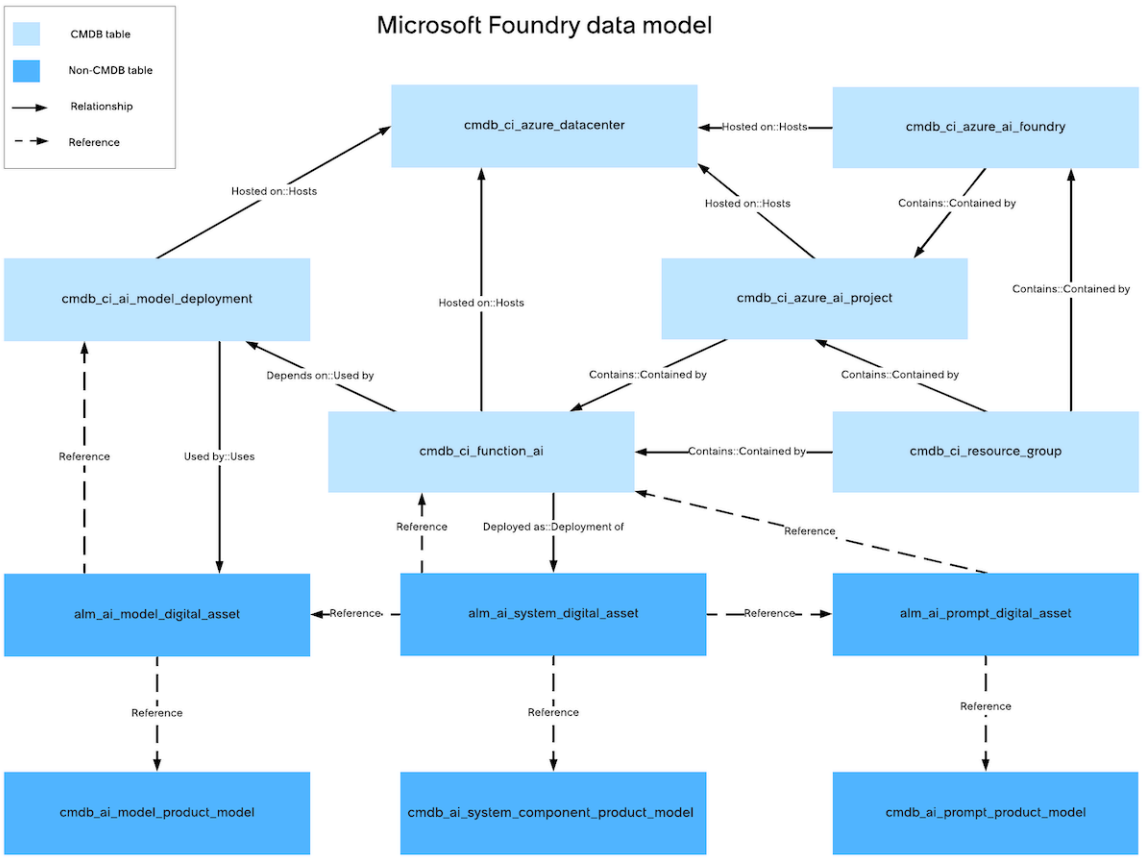

Discovery requirements

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Verify that the applications are up to date

Discovering these resources may require updating to the latest version of the following applications from the ServiceNow Store:

- AI Agent Topology Mapping
- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- Data Foundation Model (at least version 1.13.2)

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Azure Availability Set

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Wait for the **Clean-Up job for Availability zone to clear availability set record** schedule job to delete all the pre-populated Configuration Items (CIs) in the **cmdb_ci_azure_availability_set** table.

Azure Availability Zone

To run a discovery with Azure Availability Zone, register the subscription ID to the **AvailabilityZonePeering** feature with AZ CLI using `az feature register -n AvailabilityZonePeering --namespace Microsoft.Resources` to use the Check Zone Peering API. Check the status with `az feature show -n AvailabilityZonePeering --namespace Microsoft.Resources` before running discovery.

Set up Azure service accounts

Enable Cloud Discovery to access your Azure environment.

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Create an Azure cloud discovery schedule

For more information, see Create an Azure Discovery schedule in Discovery Admin Workspace.

(Optional) Discover datacenters only for new subscriptions

Starting with Zurich Patch 2, you can discover datacenters only for new subscriptions added since the last discovery. For more information, see Discover datacenters only for new cloud accounts.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Data collection

The AI Agent Topology Mapping application includes the following patterns for Microsoft Foundry discovery:

- Azure - AI New Foundry Agents
- Azure - AI Service Foundry
- Azure - AI Service Foundry Project
- Azure - AI Foundry Models

The Azure - AI New Foundry Agents pattern collects AI System Digital Assets, AI Prompt Digital Assets, and Key Value information. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **Azure** > **Azure AI Foundry Agents** and selecting the relevant record.

The Azure - AI Foundry Models pattern collects AI Model Digital Assets. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **Azure** > **Azure AI Foundry Models** and selecting the relevant record.

The AI Agent Topology Mapping application populates data in both CMDB and non-CMDB tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data stored in non-CMDB tables

The AI Agent Topology Mapping application populates data in non-CMDB tables when running the Microsoft Foundry patterns.

If you have the AI Control Tower application, you can view the information in the AI Control Tower Workspace.

AI System Digital Asset [alm_ai_system_digital_asset]

Field	Description
Display name [display_name]	Display name of the Microsoft Foundry agent.
External record reference [external_ref_id]	Unique identifier of the Microsoft Foundry agent.
Asset type [model_category]	Asset type of the agent. Value is set to **Agentic AI**.
State [install_status]	State of the agent. Value is set to **Deployed**.
Vendor [vendor]	Vendor of the agent. Value is set to **Azure**.
Model [model]	References the AI System Component Product Model [cmdb_ai_system_component_product_model] table.
AI prompts [ai_prompts]	References the AI Prompt Digital Asset [alm_ai_prompt_digital_asset] table.
AI models [ai_models]	References the AI Model Digital Asset [alm_ai_model_digital_asset] table.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

AI System Component Product Model [cmdb_ai_system_component_product_model]

Field	Description
Name [name]	Name of the Microsoft Foundry agent.
Short description [short_description]	Description of the Microsoft Foundry agent.
Status [status]	Operational status of the agent. Value is set to **PREPARED**.
Version [version]	Version of the Microsoft Foundry agent.
Manufacturer [manufacturer]	Name of the company defined in the *glide.appcreator.company.friendly_name* system property.
Model categories [cmdb_model_category]	Model category. Value is set to **Agentic AI**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Prompt Digital Asset [alm_ai_prompt_digital_asset]

Field	Description
Prompt information [prompt_info]	Instruction text defined for the Microsoft Foundry agent.
Asset type [model_category]	Asset type of the prompt. Value is set to **AI prompt**.
Vendor [vendor]	Vendor of the prompt. Value is set to **Azure**.
Model [model]	References the AI Prompt Product Model [cmdb_ai_prompt_product_model] table.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

AI Prompt Product Model [cmdb_ai_prompt_product_model]

Field	Description
Name [name]	Name of the Microsoft Foundry agent, used to identify the associated prompt model.
Manufacturer [manufacturer]	Name of the company defined in the *glide.appcreator.company.friendly_name* system property.
Model categories [cmdb_model_category]	Model category. Value is set to **AI prompt**.

AI Model Digital Asset [alm_ai_model_digital_asset]

Field	Description
Display name [display_name]	Name of the foundation model.
External record reference [external_ref_id]	Resource identifier of the foundation model.
Asset type [model_category]	Asset type of the model. Value is set to **AI model**.
Vendor [vendor]	Provider of the foundation model.
Model [model]	References the AI Model Product Model [cmdb_ai_model_product_model] table.
Configuration Item [ci]	References the AI Model Deployment [cmdb_ci_ai_model_deployment] table.

AI Model Product Model [cmdb_ai_model_product_model]

Field	Description
Name [name]	Name of the foundation model.
Model number [model_number]	Model ID assigned by Microsoft Foundry.
Version [version]	Version of the foundation model.
Description [description]	Description of the foundation model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Model Product Model [cmdb_ai_model_product_model] (continued)

Field	Description
Manufacturer [manufacturer]	Provider of the foundation model.
Model categories [cmdb_model_category]	Model category. Value is set to **AI model**.

Data stored in CMDB tables

The AI Agent Topology Mapping application populates data in the CMDB when running the Microsoft Foundry patterns.

AI Function [cmdb_ci_function_ai]

Field	Description
Name [name]	Name of the Microsoft Foundry agent.
Object ID [object_id]	Unique identifier of the Microsoft Foundry agent.
Description [short_description]	Description of the Microsoft Foundry agent.
Manufacturer [manufacturer]	Cloud provider of the resource. Value is set to **Azure**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Product instance identifier [product_instance_id]	Version of the Microsoft Foundry agent.
Comments [comments]	Resource type identifier. Value is set to **Azure::Foundry::NewAgent**.

Azure AI Foundry [cmdb_ci_azure_ai_foundry]

Field	Description
Name [name]	Name of the Microsoft Foundry service.
Object ID [object_id]	Unique identifier of the Microsoft Foundry service.
Description [short_description]	Description of the resource. Value is set to **Azure Foundry Service**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Azure Service Type [azure_service_type]	Azure resource type. For example: microsoft.cognitiveservices/accounts.
Azure Service Kind [azure_service_kind]	Resource kind of the Microsoft Foundry service. For example: AIServices.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure AI Project [cmdb_ci_azure_ai_project]

Field	Description
Name [name]	Name of the Microsoft Foundry project.
Object ID [object_id]	Unique identifier of the Microsoft Foundry project.
Comments [comments]	Azure region where the Microsoft Foundry project is deployed.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Model number [model_number]	Internal identifier of the Microsoft Foundry project.
Azure Project Kind [azure_project_kind]	Resource kind of the Microsoft Foundry project. For example: AIServices.

AI Model Deployment [cmdb_ci_ai_model_deployment]

Field	Description
Name [name]	Name of the AI model available in Microsoft Foundry.
Object ID [object_id]	Resource identifier of the AI model.
Description [short_description]	Description of the AI model.
Model number [model_number]	Model name, used as an identifier.
Manufacturer [manufacturer]	Provider of the AI model. For example: Microsoft, Anthropic, or OpenAI.
Attributes [attributes]	Lifecycle status of the model as reported by the API. For example: Stable or Deprecated.
Comments [comments]	Azure resource type identifier. Value is set to **Microsoft.CognitiveServices/locations/models**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.

Dependency Views map

On the Dependency Views map, you can view discovered Microsoft Foundry resources and the relationships between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The following relationships are stored in the Asset-CI Relationship [cmdb_rel_asset_ci] table and aren't represented in the Dependency View map:

- The **Deployed as::Deployment of** relationship between AI Function [cmdb_ci_function_ai] and AI System Digital Asset [alm_ai_system_digital_asset].
- The **Used by::Uses** relationship between AI Model Deployment [cmdb_ci_ai_model_deployment] and AI Model Digital Asset [alm_ai_model_digital_asset].

Microsoft Foundry dependency view

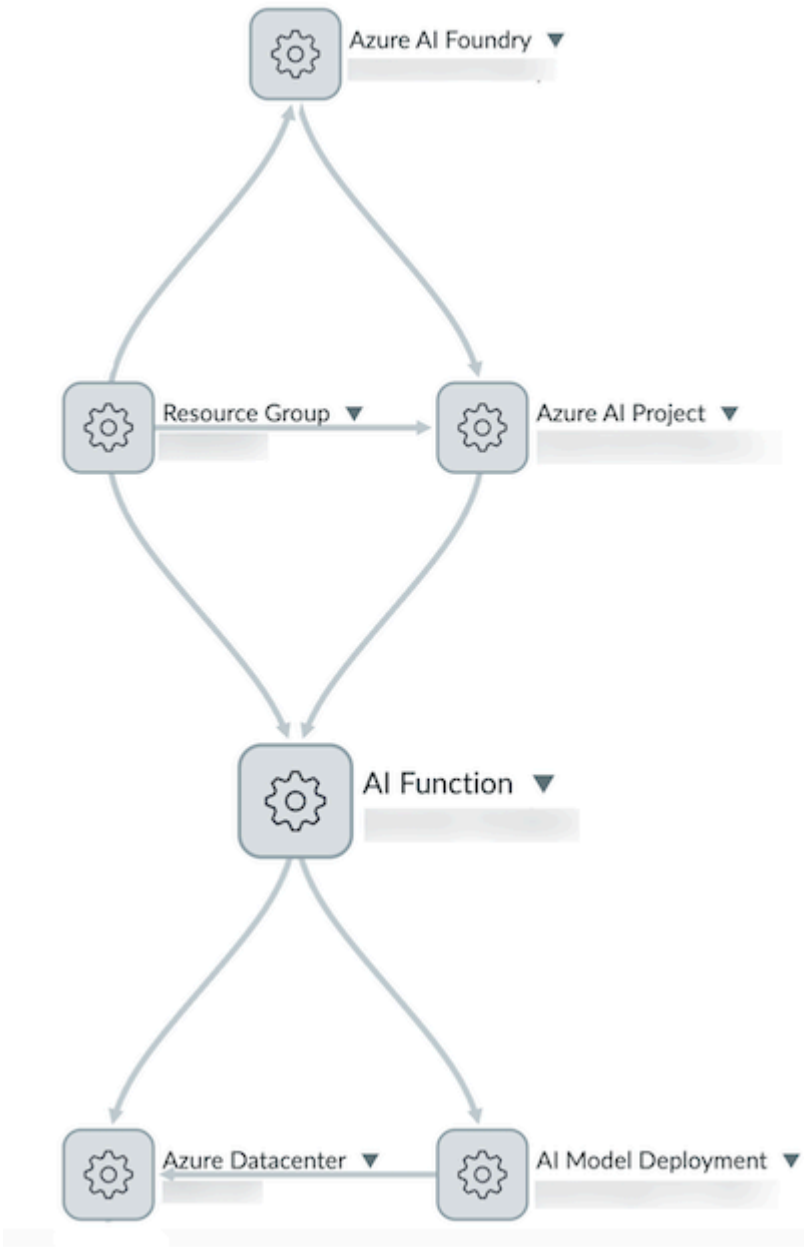

CI relationships and references

The Microsoft Foundry patterns create the following relationships and references to support Microsoft Foundry (new) discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships

CI/Table	Relationship	CI/Table
AI Function [cmdb_ci_function_ai]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure AI Project [cmdb_ci_azure_ai_project]	Contains::Contained by	AI Function [cmdb_ci_function_ai]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	AI Function [cmdb_ci_function_ai]
AI Function [cmdb_ci_function_ai]	Deployed as::Deployment of	AI System Digital Asset [alm_ai_system_digital_asset]
Azure AI Foundry [cmdb_ci_azure_ai_foundry]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Azure AI Foundry [cmdb_ci_azure_ai_foundry]
Azure AI Foundry [cmdb_ci_azure_ai_foundry]	Contains::Contained by	Azure AI Project [cmdb_ci_azure_ai_project]
Azure AI Project [cmdb_ci_azure_ai_project]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Azure AI Project [cmdb_ci_azure_ai_project]
AI Function [cmdb_ci_function_ai]	Depends on::Used by	AI Model Deployment [cmdb_ci_ai_model_deployment]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Used by::Uses	AI Model Digital Asset [alm_ai_model_digital_asset]

CI references

CI/Table	Field	Referenced CI/Table
AI System Digital Asset [alm_ai_system_digital_asset]	Model [model]	AI System Component Product Model [cmdb_ai_system_component_product_model]
AI System Digital Asset [alm_ai_system_digital_asset]	AI prompts [ai_prompts]	AI Prompt Digital Asset [alm_ai_prompt_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	AI models [ai_models]	AI Model Digital Asset [alm_ai_model_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Model [model]	AI Prompt Product Model [cmdb_ai_prompt_product_model]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI/Table	Field	Referenced CI/Table
AI Model Digital Asset [alm_ai_model_digital_asset]	Model [model]	AI Model Product Model [cmdb_ai_model_product_model]
AI Model Digital Asset [alm_ai_model_digital_asset]	Configuration Item [ci]	AI Model Deployment [cmdb_ci_ai_model_deployment]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Azure AI Foundry [cmdb_ci_azure_ai_foundry]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Azure AI Project [cmdb_ci_azure_ai_project]

Azure tag discovery

The Azure - AI Service Foundry pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Azure AI Foundry [cmdb_ci_azure_ai_foundry] table.

The Azure - AI Service Foundry Project pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Azure AI Project [cmdb_ci_azure_ai_project] table.

Related topics

Microsoft Azure Cloud discovery using patterns

Microsoft Foundry Hub pattern-based discovery

AI Agent Topology Mapping discovers Microsoft Foundry Hub assistants and models during horizontal discovery.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Foundry Hub data model

The following diagram illustrates the tables and relationships that the AI Agent Topology Mapping application creates when discovering Microsoft Foundry Hub resources.

Tables and relationships populated during Microsoft Foundry Hub discovery

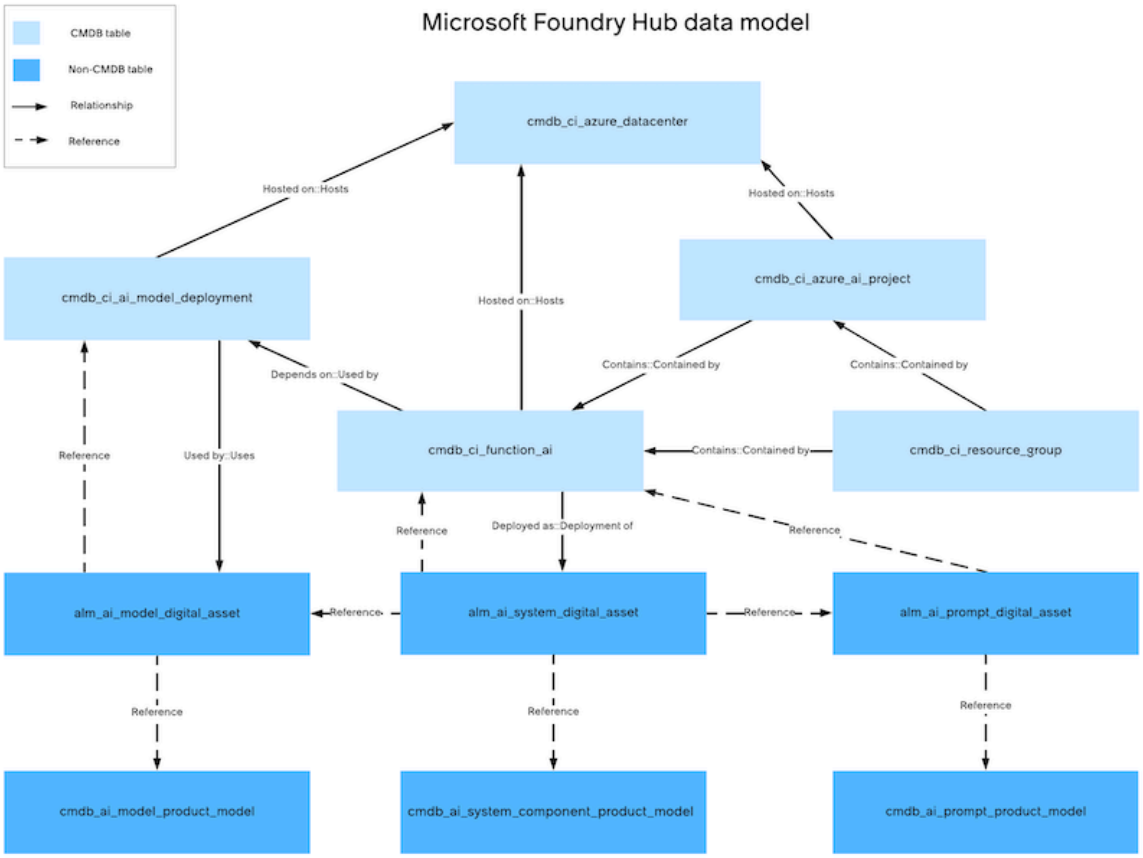

Discovery requirements

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Verify that the applications are up to date

Discovering these resources may require updating to the latest version of the following applications from the ServiceNow Store:

- AI Agent Topology Mapping
- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models
- Data Foundation Model (at least version 1.13.2)

Activate the cloud-related CI relationships

To include discovered components into service instances, enable CI relationships used in tag-based discovery by Service Mapping. These CI relationships are

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

available from the 1.0.68 release on the ServiceNow Store. For operational steps, see Tag-based discovery configuration.

Azure Availability Set

Wait for the **Clean-Up job for Availability zone to clear availability set record** schedule job to delete all the pre-populated Configuration Items (CIs) in the **cmdb_ci_azure_availability_set** table.

Azure Availability Zone

To run a discovery with Azure Availability Zone, register the subscription ID to the **AvailabilityZonePeering** feature with AZ CLI using `az feature register -n AvailabilityZonePeering --namespace Microsoft.Resources` to use the Check Zone Peering API. Check the status with `az feature show -n AvailabilityZonePeering --namespace Microsoft.Resources` before running discovery.

Set up Azure service accounts

Enable Cloud Discovery to access your Azure environment.

Discovering Azure GovCloud (US) accounts requires using a datacenter URL when setting up an Azure service account. For more information, see Set up Azure service accounts.

Create an Azure cloud discovery schedule

For more information, see Create an Azure Discovery schedule in Discovery Admin Workspace.

(Optional) Discover datacenters only for new subscriptions

Starting with Zurich Patch 2, you can discover datacenters only for new subscriptions added since the last discovery. For more information, see Discover datacenters only for new cloud accounts.

(Optional) Populate Service Account and Logical Datacenter fields in cloud CIs

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating Service Account and Logical Datacenter fields directly in cloud CIs. For more information, see Improved query performance with direct field population in CI tables.

Data collection

The AI Agent Topology Mapping application includes the following patterns for Microsoft Foundry Hub discovery:

- Azure - AI Foundry Hub Agents
- Azure - AI Foundry Hub Project
- Azure - AI Foundry Models

The Azure - AI Foundry Hub Agents pattern collects AI System Digital Assets and AI Prompt Digital Assets. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **Azure** > **Azure AI Foundry Agents** and selecting the relevant record.

The Azure - AI Foundry Models pattern collects AI Model Digital Assets. You can view this information in the **Related Links** section of the CI record by navigating to **All** > **Configuration** > **Azure** > **Azure AI Foundry Models** and selecting the relevant record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The AI Agent Topology Mapping application populates data in both CMDB and non-CMDB tables.

Data stored in non-CMDB tables

The AI Agent Topology Mapping application populates data in non-CMDB tables when running the Microsoft Foundry Hub patterns.

If you have the AI Control Tower application, you can view the information in the AI Control Tower Workspace.

AI System Digital Asset [alm_ai_system_digital_asset]

Field	Description
Display name [display_name]	Display name of the Microsoft Foundry Hub assistant.
External record reference [external_ref_id]	Unique identifier of the Microsoft Foundry Hub assistant.
Asset type [model_category]	Asset type of the assistant. Value is set to **Agentic AI**.
State [install_status]	State of the assistant. Value is set to **Deployed**.
Vendor [vendor]	Vendor of the assistant. Value is set to **Azure**.
Model [model]	References the AI System Component Product Model [cmdb_ai_system_component_product_model] table.
AI prompts [ai_prompts]	References the AI Prompt Digital Asset [alm_ai_prompt_digital_asset] table.
AI models [ai_models]	References the AI Model Digital Asset [alm_ai_model_digital_asset] table.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

AI System Component Product Model [cmdb_ai_system_component_product_model]

Field	Description
Name [name]	Name of the Microsoft Foundry Hub assistant.
Short description [short_description]	Description of the Microsoft Foundry Hub assistant.
Status [status]	Operational status of the assistant. Value is set to **PREPARED**.
Manufacturer [manufacturer]	Name of the company defined in the *glide.appcreator.company.friendly_name* system property.
Model categories [cmdb_model_category]	Model category. Value is set to **Agentic AI**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Prompt Digital Asset [alm_ai_prompt_digital_asset]

Field	Description
Prompt information [prompt_info]	Instruction text defined for the Microsoft Foundry Hub assistant.
Asset type [model_category]	Asset type of the prompt. Value is set to **AI prompt**.
Vendor [vendor]	Vendor of the prompt. Value is set to **Azure**.
Model [model]	References the AI Prompt Product Model [cmdb_ai_prompt_product_model] table.
Configuration Item [ci]	References the AI Function [cmdb_ci_function_ai] table.

AI Prompt Product Model [cmdb_ai_prompt_product_model]

Field	Description
Name [name]	Name of the Microsoft Foundry Hub assistant, used to identify the associated prompt model.
Manufacturer [manufacturer]	Name of the company defined in the *glide.appcreator.company.friendly_name* system property.
Model categories [cmdb_model_category]	Model category. Value is set to **AI prompt**.

AI Model Digital Asset [alm_ai_model_digital_asset]

Field	Description
Display name [display_name]	Name of the foundation model.
External record reference [external_ref_id]	Resource identifier of the foundation model.
Asset type [model_category]	Asset type of the model. Value is set to **AI model**.
Vendor [vendor]	Provider of the foundation model.
Model [model]	References the AI Model Product Model [cmdb_ai_model_product_model] table.
Configuration Item [ci]	References the AI Model Deployment [cmdb_ci_ai_model_deployment] table.

AI Model Product Model [cmdb_ai_model_product_model]

Field	Description
Name [name]	Name of the foundation model.
Model number [model_number]	Model ID assigned by Microsoft Foundry.
Version [version]	Version of the foundation model.
Description [description]	Description of the foundation model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Model Product Model [cmdb_ai_model_product_model] (continued)

Field	Description
Manufacturer [manufacturer]	Provider of the foundation model.
Model categories [cmdb_model_category]	Model category. Value is set to **AI model**.

Data stored in CMDB tables

The AI Agent Topology Mapping application populates data in the CMDB when running the Microsoft Foundry Hub patterns.

AI Function [cmdb_ci_function_ai]

Field	Description
Name [name]	Name of the Microsoft Foundry Hub assistant.
Object ID [object_id]	Unique identifier of the Microsoft Foundry Hub assistant.
Description [short_description]	Description of the Microsoft Foundry Hub assistant.
Manufacturer [manufacturer]	Cloud provider of the resource. Value is set to **Azure**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Comments [comments]	Resource type identifier. Value is set to **Azure::FoundryHub::Agent**.

Azure AI Project [cmdb_ci_azure_ai_project]

Field	Description
Name [name]	Name of the Microsoft Foundry Hub project.
Object ID [object_id]	Unique identifier of the Microsoft Foundry Hub project.
Comments [comments]	Azure region where the Microsoft Foundry Hub project is deployed.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.
Azure Project Kind [azure_project_kind]	Resource kind of the Microsoft Foundry Hub project. For example: Project.
Attributes [attributes]	Resource type identifier. Value is set to **Azure::FoundryHub::Project**.

AI Model Deployment [cmdb_ci_ai_model_deployment]

Field	Description
Name [name]	Name of the AI model available in Microsoft Foundry Hub.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Model Deployment [cmdb_ci_ai_model_deployment] (continued)

Field	Description
Object ID [object_id]	Resource identifier of the AI model.
Description [short_description]	Description of the AI model.
Model number [model_number]	Model name, used as an identifier.
Manufacturer [manufacturer]	Provider of the AI model. For example: Microsoft, Anthropic, or OpenAI.
Attributes [attributes]	Lifecycle status of the model as reported by the API. For example: Stable or Deprecated.
Comments [comments]	Azure resource type identifier. Value is set to **Microsoft.CognitiveServices/locations/models**.
Operational status [operational_status]	Operational status of the resource. Default value is **Operational**.
Install Status [install_status]	Install status of the resource. Default value is **Installed**.

Dependency Views map

On the Dependency Views map, you can view discovered Microsoft Foundry Hub resources and the relationships between them.

Note: The following relationships are stored in the Asset-CI Relationship [cmdb_rel_asset_ci] table and aren't represented in the Dependency View map:

- The **Deployed as::Deployment of** relationship between AI Function [cmdb_ci_function_ai] and AI System Digital Asset [alm_ai_system_digital_asset].
- The **Used by::Uses** relationship between AI Model Deployment [cmdb_ci_ai_model_deployment] and AI Model Digital Asset [alm_ai_model_digital_asset].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft Foundry Hub dependency view

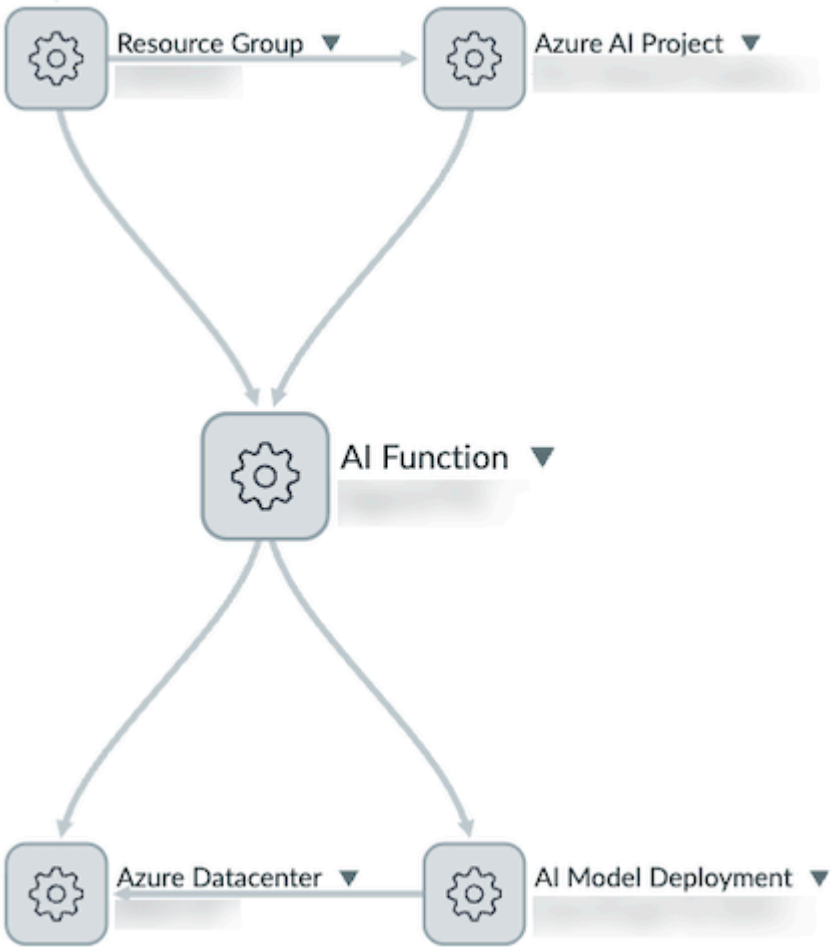

CI relationships and references

The Microsoft Foundry Hub patterns create the following relationships and references to support Microsoft Foundry Hub discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI/Table	Relationship	CI/Table
AI Function [cmdb_ci_function_ai]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
Azure AI Project [cmdb_ci_azure_ai_project]	Contains::Contained by	AI Function [cmdb_ci_function_ai]
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	AI Function [cmdb_ci_function_ai]
AI Function [cmdb_ci_function_ai]	Deployed as::Deployment of	AI System Digital Asset [alm_ai_system_digital_asset]
Azure AI Project [cmdb_ci_azure_ai_project]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI/Table	Relationship	CI/Table
Resource Group [cmdb_ci_resource_group]	Contains::Contained by	Azure AI Project [cmdb_ci_azure_ai_project]
AI Function [cmdb_ci_function_ai]	Depends on::Used by	AI Model Deployment [cmdb_ci_ai_model_deployment]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Hosted on::Hosts	Azure Datacenter [cmdb_ci_azure_datacenter]
AI Model Deployment [cmdb_ci_ai_model_deployment]	Used by::Uses	AI Model Digital Asset [alm_ai_model_digital_asset]

CI references

CI/Table	Field	Referenced CI/Table
AI System Digital Asset [alm_ai_system_digital_asset]	Model [model]	AI System Component Product Model [cmdb_ai_system_component_product_model]
AI System Digital Asset [alm_ai_system_digital_asset]	AI prompts [ai_prompts]	AI Prompt Digital Asset [alm_ai_prompt_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	AI models [ai_models]	AI Model Digital Asset [alm_ai_model_digital_asset]
AI System Digital Asset [alm_ai_system_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Model [model]	AI Prompt Product Model [cmdb_ai_prompt_product_model]
AI Prompt Digital Asset [alm_ai_prompt_digital_asset]	Configuration Item [ci]	AI Function [cmdb_ci_function_ai]
AI Model Digital Asset [alm_ai_model_digital_asset]	Model [model]	AI Model Product Model [cmdb_ai_model_product_model]
AI Model Digital Asset [alm_ai_model_digital_asset]	Configuration Item [ci]	AI Model Deployment [cmdb_ci_ai_model_deployment]
Key Value [cmdb_key_value]	Configuration item [configuration_item]	Azure AI Project [cmdb_ci_azure_ai_project]

Azure tag discovery

The Azure - AI Foundry Hub Project pattern collects tags and populates them in the Key Value [cmdb_key_value] table.

Key Value [cmdb_key_value]

Field	Description
Key [key]	Tag name.
Value [value]	Tag value.
Configuration item [configuration_item]	References the Azure AI Project [cmdb_ci_azure_ai_project] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Microsoft Azure Cloud discovery using patterns

ITOM Content Service

ServiceNow® ITOM Content Service offers extensive visibility of products in your infrastructure. The classification of processes identified by Predictive Intelligence enables wider discovery and weekly updates of new configuration items in the CMDB. Use the Discovery Admin Workspace to review and manage ITOM Content Service suggestions.

ITOM Content Service Model

About ITOM Content Service

ITOM Visibility offers multiple methods of discovering products in your infrastructure. When running discovery with Discovery and Service Mapping Patterns for example, each pattern searches for one version of a product and can offer in-depth visibility of the product's information.

While your infrastructure is growing at scale, and new IT products are rapidly released to the market, you must have wider, rapidly updating visibility. Every week, ITOM Content Service identifies products that you're currently using and delivers new Configuration Items (CI) for your usage. The latest version of Discovery Admin Workspace enables you to refresh the content service suggestions manually before reviewing them.

ITOM Content Service provides the following benefits:

- Reviews the products that you're currently using every week and delivers new configuration items (CIs).
- Based on AI, it helps you identify running processes fingerprints and SNMP System OIDs and create the CIs you need for monitoring your infrastructure.
- Identifies irrelevant processes that aren't suitable candidates for CIs.
- Discovers a higher number of products by using AI capabilities that cluster and classify running application processes.

Benefits of Software Asset Management integration with ITOM Content Service

ITAM Software Asset Management (SAM) users benefit from the integration with ITOM Content Service in multiple ways.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Gain visibility on the processes currently running on the installed software. This also enables you to find out which installed software doesn't have running processes and might be redundant.
- ITOM Content Service compliments Software Asset Management in cases of services that don't require installation, for example, an Oracle database that runs as a service and isn't installed. In cases like this, ITOM Content Service creates the installed software record based on combining the process fingerprint data with the data discovered by SAM. Without ITOM Content Service, you can't gain this kind of visibility.
- Using file-based discovery capabilities, ITOM Content Service can discover a running system file that is identified with an app. For example, ITOM Content Service can match an installed software record to an Apache Kafka running an `exe` file. For more information, see File-based Discovery.

For more information on Software Asset Management, see Software Asset Management Content Service.

Benefits of using Service Mapping with ITOM Content Service

With the latest version of ITOM Content Service, Service Mapping users can leverage Service Fingerprints suggestions to identify and select application services based on classified components. By viewing component details such as product, product category, publisher, and description, you can make informed decisions about which candidates to convert into application services. You can also see which additional processes can be integrated into the service, as certain processes collectively build a comprehensive service when combined.

Both ITOM Content Service suggestions and Service Fingerprints offer automated, weekly suggestions for candidates, but they differ in a few ways.

Comparison between ITOM Content Service suggestions and Service Fingerprints

Name	Type	Identifies	Focus	Details on publisher	Location
ITOM Content Service suggestions	Process fingerprints	Individual processes, which could be components of a service or standalone.	Specific actions of a process, offering higher granularity.	Process's publisher. If a process is part of a service but has a different publisher, the process provides details only about its own publisher.	**Content service suggestions** tab in Discovery Admin Workspace Content 360
Service Fingerprints	Service fingerprints	Services based on classified components.	The service as a whole, providing a broader perspective.	Service's publisher. If a process is part of a service but has a different publisher, the service fingerprint provides details only about the service's publisher.	**Service Fingerprints** tab in Service Mapping workspace

For more information about using Service Fingerprints in Service Mapping, see Map application services based on Automated Service Suggestions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information about Service Fingerprints candidates, see KB1706310.

Managing ITOM Content Service with the Discovery Admin Workspace

Review recommended actions and application suggestions and configure discovery through Discovery Admin Workspace.

Note that the Discovery Admin Workspace is supported starting with the Vancouver release.

For more information, see Discovery Admin Workspace Content 360.

Sharing data on the CDS

The Continuance Delivery System (CDS) is a ServiceNow® technology that synchronizes data between the Data Services instance that stores and distributes data and the customer instances.

Opting in to sharing application fingerprint process data with ITOM Content Service through the CDS enables them to be tagged and classified by AI librarians. ITOM Content Service then delivers new CIs every week. Your data is shared with your ServiceNow® instance in a safe environment. You can control access to your data by setting permissions in the Discovery Admin Workspace.

The three options for sharing your data are:

- Sharing the running processes fingerprint data, which enables access to the Process Groups [cmdb_process_groups] table.
- Sharing the SNMP System OIDs data, which enables access to the SNMP OID [discovery_snmp_oid] table.
- Sharing the Service Info data of mapped candidates, which enables access to the Application Service Candidate [application_service_candidate] table (starting with ITOM Content Service version 1.6.2 and Discovery Admin Workspace version 1.11.0).

For more information on setting the permissions, see Share data on ITOM Content Service

ITOM Content Service use case

This ITOM Content Service use case demonstrates how organizations can achieve enhanced visibility into their environment.

Use case overview

The Application Configuration Items (CIs) class is a key component of an organization's data platform within the Common Service Data Model (CSDM). Application CIs help log software activities and support critical IT operations.

Many products and applications rely on application CIs for essential tasks. A few key examples are:

- Service Mapping: Uses application CIs to create service maps
- ITOM AIOps: Monitors application health and resolves incidents using application CIs
- IT Asset Management (ITAM): Tracks and manages software applications and their associated licenses with application CIs
- IT Service Management (ITSM): Analyzes service impact and calculates blast radius with application CIs
- Enterprise Architecture (EA): Leverages application CIs to generate business outcomes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Challenge

A large organization with numerous applications requires accurate discovery. Managing application CIs is essential, yet the IT team finds it challenging to maintain visibility over its growing infrastructure.

ITOM Visibility offers several methods for discovering products within an organization's infrastructure, such as Discovery and Service Mapping Patterns. This method provides detailed visibility into a product's information. When a wide range of applications is in use, not all can have a dedicated pattern. If an application does not have a dedicated pattern, discovery populates only the Running Process [cmdb_running_process] table without identifying the product or creating a CI for it.

Without an application CI, the organization lacks accurate visibility into its infrastructure and might struggle to optimize the following essential actions:

- Bind CIs with monitoring alerts and open incidents regarding the CIs
- Schedule automated software upgrades
- Automate scripts to optimize or troubleshoot the product
- Perform impact assessments and change impact analysis
- Create security workflows

Solution

ITOM Content Service enhances the discovery process by tagging and classifying process fingerprints. This service provides details about the publisher, product, and the essential component of the CI. Unlike pattern-based discovery, ITOM Content Service identifies applications directly within the organization's environment without requiring published pattern classifiers. This approach enables the continuous delivery of content without being constrained by a scheduled release cycle.

For example: Discovery identifies the Elipse Software E3 Database Engine process (E3DBEngine.exe) and lists it in the Running Process [cmdb_running_process] table. If the E3DBEngine.exe process classification condition isn't set to trigger a pattern, it doesn't launch a pattern and a CI isn't created. However, with ITOM Content Service, the E3DBEngine.exe process is tagged and classified, enabling Discovery to turn the process into a CI during the next scan.

Result

Discovery using ITOM Content Service creates a CI record by tagging and classifying the publisher, product, and essential component for all applications in the organization's environment. This method provides enhanced and continuous visibility into the application landscape, enabling effective monitoring, proactive actions, and ongoing optimization.

Install ITOM Content Service

Install the ITOM Content Service application (sn_smart_content) to manage and monitor the creation of configuration items based on application fingerprints.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the ITOM Content Service application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin

About this task

The following plugins are activated with ITOM Content Service:

- Discovery (com.snc.discovery)
- Normalization Data Services Client (com.glide.data_services_canonicalization.client)
- File Based Discovery (com.snc.discovery.file_based_discovery)—starting with ITOM Content Service version 1.5.1

Note: The ITOM Smart Content plugin (com.snc.itom.smart.content) is available starting with XP5. This plugin is optional and can be requested through ServiceNow personnel. For more information, see Request a plugin.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the ITOM Content Service application (sn_smart_content) using the filter criteria and search bar.

   You can search for the app by its name or ID.

   In the list next to the **Install** button, the versions that are available to you are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

What to do next

Share data on ITOM Content Service

Related topics

ITOM Content Service

Discovery Admin Workspace

Share data on ITOM Content Service

Control data sharing by opting in or out of ITOM Content Service sharing options using the Discovery Admin Workspace.

Before you begin

Ensure that the following application and plugin are installed:

- ITOM Content Service (sn_smart_content)
- Discovery Admin Workspace (com.snc.itom.daw)

To share Service Info, confirm that ITOM Content Service and Discovery Admin Workspace are upgraded to the latest version.

- ITOM Content Service (at least version 1.6.2)
- Discovery Admin Workspace (at least version 1.11.0)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace**.
2. Indicate whether and what data you want to share with AI librarians.

Option	Actions
Opt in to share your data	**a.** In the **Content service setup** tile, select **Opt in**. **b.** Read the agreement and select **Yes, I have read and agree with the Opt-in Agree ment**. **c.** Select **Next**. **d.** Select your setup. ▪ Advanced - Activates a comprehensive set of classifiers for expanded visibility ▪ Basic - Activates the most common classi fiers. ▪ Custom - You select which classifiers to activate. **e.** Turn on the data sharing options that you want to use. ▪ Running process fingerprint data: Shares metadata and some samples of process es from the Process Groups table. ▪ SNMP System OIDs: Shares metadata and some samples of processes from the SNMP OID table. ▪ Service Info: Shares metadata and some samples of processes from the Applica tion Service Candidate table. **f.** Select **Save**. **g.** When the Opt-in is complete, select **Home** to return to the workspace. The **Content service suggestions** tile re places the **Content service setup** tile. In the new **Content service setup** section, a green check-mark appears indicating that you have opted in to ITOM Content Service.
Set the data sharing permissions off	**a.** In the **Content service setup** section, se lect **Configure setup**. **b.** Turn off one option or both options.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Actions
	c. Select **Save**. **d.** When the setting is complete, select **Home** to return to the workspace.
Opt out of sharing your data	**a.** In the **Content service setup** section, select **Opt out**. **b.** Select **Opt out**. **c.** Read the **Opting out of** *ITOM Content Service* system message and select **Opt out**. **d.** When the opt-out process is complete, select **Home** to return to the workspace. ITOM Content Service stops using your data and your workspace presents only CIs defined by you. For more information, see Discovery Admin Workspace.

Related topics

ITOM Content Service

Enable discovery with ITOM Content Service

Review the ITOM Content Service suggestions for application discovery and activate the discovery classifiers to enable the creation of configuration items.

Before you begin

Ensure that the following application and plugin are installed:

- ITOM Content Service (sn_smart_content)
- Discovery Admin Workspace (com.snc.itom.daw)

Ensure that ITOM Content Service is active. For more information, see Share data on ITOM Content Service.

Role required: discovery_admin

Procedure

1. Navigate to **Workspaces** > **Discovery Admin Workspace** > **Content 360**.
2. Select the **Content service suggestions** tile.
3. Manually sync with suggestions from your system by selecting **Sync candidates**.

   Note: If a message appears about the application scope, select **OK** to exit the message. Then, change the scope by selecting the application scope icon ( ) and selecting **Global**.

4. In the Sync candidates confirmation window, select **Yes**.
5. Review the Content Service suggestions for discovery by product.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To filter the list, select the products from the **Content service suggestions** tab.
- To review a product's description and details, select the link from the **Name** column.

6. Select the product you want to enable discovery for.
7. Activate discovery classifiers by product categories.

Note: The number of classifiers that will be activated is indicated in the **Activate** button.

- To enable the discovery by product categories, select the product categories check boxes, select the check box next to the **Name** column, and select **Activate**.
- To enable the discovery by individual products, select the product from the **Name** column and select **Activate**.

Result

The new configuration items are added to the CMDB the next time Discovery runs a schedule. The Discovery Admin Workspace then lists the added CIs and removes them from the suggested applications list.

Related topics

ITOM Content Service

Discovery Admin Workspace

Discover installed software data with ITOM Content Service

Use ITOM Content Service and ITSM Software Asset Management to discover version-less installed software that isn't discovered automatically with Software Asset Management Core or Software Asset Management Professional or file-based Discovery.

Before you begin

Ensure that the following application and plugins are installed:

- ITOM Content Service (sn_smart_content)
- Software Asset Management Core (com.snc.sam.core) or Software Asset Management Professional (com.snc.samp)
- File Based Discovery (com.snc.discovery.file_based_discovery)

Ensure that ITOM Content Service is active. For more information, see Share data on ITOM Content Service.

Ensure that the content service suggestions in the relevant category are activated. For more information, see Enable discovery with ITOM Content Service.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery Smart Content** > **Smart Content Categories**.
2. Activate the relevant category.
   a. In the **Name** column, search for the relevant category.
   b. Double-click (or use the keyboard shortcut) the **Create Installed Software** column.
   c. In the list, select **true**.
   d. Verify the value by selecting the green check mark.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The new installed software record is created and added to the CI Software Installations list the next time Discovery runs a schedule.

Related topics

ITOM Content Service

ITOM Content Service reference

Using ITOM Content Service enables you to have the CMDB populated with new data every week. Use the knowledge base articles to learn about the Configuration Item candidates that were released.

To learn about the Application CI candidates, see KB1584650.

To learn about the SNMP OID CI candidates, see KB1702434.

To learn about the Service Fingerprints candidates, see KB1706310.

Tag Governance

Effective tag management improves reporting and overall operations management efficiency. Use the ServiceNow® Tag Governance app to identify and remediate on-premises or cloud resources that are inconsistent and don't comply with the tag policies of your organization.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Get started

A tag is a key-value pair that is assigned to a CI or cloud resource. Your organization can use tags to categorize assets to improve visibility into cloud usage and costs.

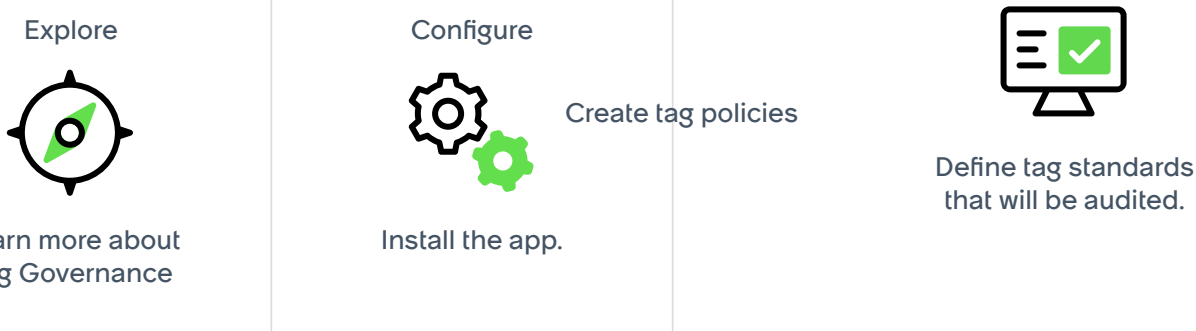

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Updating cloud tags	Remediate	Reference
	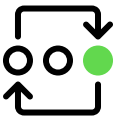	
Real-time updates to tags for cloud resources.	Remediate tag issues	Additional information on forms and settings.

Troubleshoot and get help

- Ask questions and explore other resources for Tag Governance in the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Tag Governance

Learn more about Tag Governance reports and discovery.

Tag Governance overview

Effective tag management improves reporting and overall operations management efficiency. Use the ServiceNow® Tag Governance app to identify and remediate on-premises or cloud resources that are inconsistent and don't comply with the tag policies of your organization.

Tag Governance users

Users

User	Description
Tag Governance admin	Responsible for maintaining uniform tag management policies and ensuring resources are properly tagged for visibility, compliance, and cost management.
Report viewer	View-only permissions: • Monitor reports from audit runs • Analyze the data to derive insights to guide actions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag Governance workflow

Tag Governance workflow

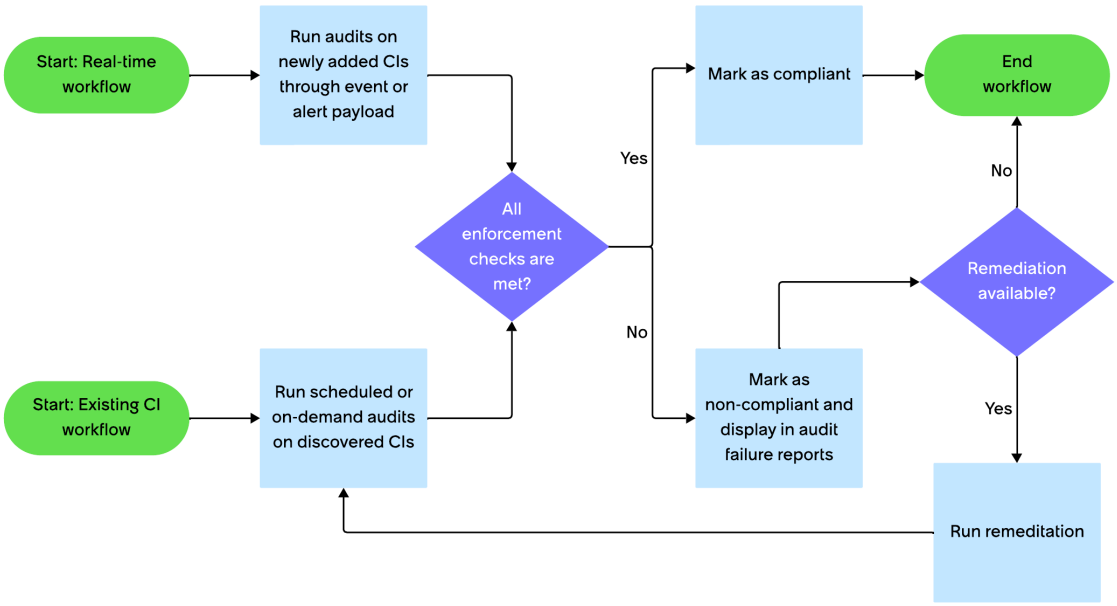

Tag Governance benefits

Tag Governance benefits

Benefit	Feature	Users
Implement cloud event processing and configure a tag policy to automatically update cloud resource tags in real time.	Performing real-time updates to tags for cloud resources	Cloud admins
Tag Governance supports domain separation, allowing you to logically group and segregate data, processes, and administrative tasks into distinct domains. This enables precise control over aspects like user access and data visibility.	Domain separation in Tag Governance	Cloud admins
Review non-compliant CIs flagged by tag audits and resolve them by adding the required tags.	Preview and remediate tag audit failures	Tag admins

Exploring

Discovering tags

The Discovery and Cloud Discovery features discover all resources in the CMDB as well as cloud resources from cloud providers such as Amazon AWS Cloud, Microsoft Azure Cloud, and Google Cloud Platform (GCP).

View Tag Governance metrics

View metrics like tag policy coverage, compliance status, and usage trends on the Tag Governance Insights dashboard.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to explore next

To learn more about configuring and using Tag Governance, see:

- Configuring Tag Governance
- Using Tag Governance
- Tag Governance reference

Discovering tags

The Discovery and Cloud Discovery features discover all resources in the CMDB as well as cloud resources from cloud providers such as Amazon AWS Cloud, Microsoft Azure Cloud, and Google Cloud Platform (GCP).

The role of discovery

Discovery and Cloud Provisioning and Governance can discover tags that are used by all major cloud providers and container ecosystems through Service Mapping patterns. Once the tags are discovered, Service Mapping can create application services that are based on the tags. Typically, organizations use tagging in virtualized, hyper-converged, or multi-cloud infrastructures. For examples of the infrastructure types and tags used for them, see Tag-based discovery in Service Mapping.

Discovery also collects cloud tags that are specific to each virtual machine (VM) and saves them to the Key Value [cmdb_key_value] table.

The discovery process also detects tags that are associated with CIs. If your organization does not use consistent tagging standards or terms (key-value pairs), however, managing CIs with tags becomes difficult.

View Tag Governance metrics

View metrics like tag policy coverage, compliance status, and usage trends on the Tag Governance Insights dashboard.

Before you begin

Verify the following:

- To be able to configure reports, log in as a Tag Governance admin (sn_itom_tag.tag_governance_admin).
- To be able to view reports, log in as a report viewer (sn_itom_tag.tag_reports_viewer).
- Tag policies must be configured and audits must be run to populate data. For more information, see Configure tag policies.

Access the dashboard

To access the dashboard, navigate to **Tag Governance** > **Dashboard**.

Dashboard details

The Tag Governance Insights dashboard displays metrics that help you make data-driven decisions, like identify tag compliance gaps, introduce tag policy adoption initiatives, assess tag compliance status, and more. For example, you can view the least used tags and decide whether they require remediation or deprecation. You can also drill down into each widget to view a list of relevant configuration items (CIs) and edit, save, or export them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The dashboard displays results from tag audit runs on discovered CIs from the CMBD CI [cmdb_ci] table.

Data visualization widgets

Widget	Type	Description
Tag policy coverage	Pie chart	Displays the total count of discovered CIs and CIs that are governed and ungoverned by a tag policy. Hover over a pie segment to view the total count and percentage of CIs for that category. Select it to view the detailed tag quality state list.
Overall compliance status across CIs	Pie chart	Displays the overall compliance status of all discovered CIs. The possible statuses are: • Non Compliant: CIs that don't meet any tag policy requirements. • Partially Compliant: CIs that meet some tag policy requirements. • Fully Compliant: CIs that meet all tag policy requirements. Hover over a pie segment to view the total count and percentage of CIs for that category. Select it to view the tag policy health list. You can filter by cloud accounts, datacenters, and policy type for specific results. For example, you can combine filters to view the tag compliance status for virtual machine instances mapped to a specific cloud account in a relevant datacenter. **Note:** The dashboard fetches cloud account and datacenter records directly from the Tag Compliance Aggregate Counts [sn_itom_tag_widget_aggregate_counts] table, displaying only cloud accounts and datacenters that have associated compliance CIs.
CI compliance by infrastructures	Stacked bar chart	Displays the compliance distribution across datacenters and cloud accounts. Each bar represents an infrastructure location with each bar segment representing their CI compliance status. Hover over a bar segment to view the total count and percentage of CIs for that category. Select it to view the tag policy run list. **Note:** The chart displays the data of top ten datacenters or cloud accounts by CI count. If there are more than ten infrastructure locations, the remaining data is combined into one bar.
Compliance trends over	Line trend report	Displays the compliance performance trends over the previous 90 days for CIs according to their compliance status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data visualization widgets (continued)

Widget	Type	Description
the past 90 days		Hover over a point on the line trend to view the total count and percentage of CIs for that category. Select it to view the tag policy run list.
Most used tags in the past 90 days	Pareto chart	Displays the most used tags across all CIs over the previous 90 days. Hover over a bar to view the total count and percentage of CIs for that category. Select it to view the tag report key counts list. **Note:** The chart displays the data of top ten tags by CI count. If there are more than ten tags, the remaining data is combined into one bar.
Least used tags	Horizontal bar chart	Displays the least used tags across all CIs. Hover over a bar to view the total count and percentage of CIs using the tag. Select it to view the tag report key counts list. **Note:** The chart displays the data of the ten least used tags by CI count. If there are more than ten tags, the remaining data is combined into one bar.

For more information about tag audit results and non-compliant CIs, see Preview and remediate tag audit failures.

Tag Categorization in Tag Governance

By automatically grouping similar tag keys into predefined categories, Tag Categorization streamlines the management of configuration items (CIs) and cloud resources while promoting clear and consistent tagging practices.

Tag categorization requirements

Note:

The Tag Categorization feature in Tag Governance leverages the CI tag category [svc_tag_categories] and the CI tag key [svc_tag_names] tables in Service Mapping. Full access to these tables requires installation of the version 1.16.3 or later of Service Mapping Plus. Without it, access to these tables might be limited, affecting your ability to recategorize tags, modify tag definitions, or fully use the feature. For more information, see Install Service Mapping Plus.

Tag Categorization overview

Tag Categorization helps streamline and improve the accuracy of tag management by providing preconfigured categories to help you group tags for your CIs. A daily scheduled job categorizes newly added tags and maps any unrecognized tag keys to the appropriate category. For more information about the predefined tag categories, see Components in Tag Governance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can modify the tag category definitions, create your own categories, or add additional tag keys as needed. For more information, see Modify tag category information for Tag Categorization.

The Tag category policy confirms that all categories are properly mapped and up-to-date. It retrieves records and checks for category mappings to identify any missing categories. CIs that are missing a category are marked as non-compliant, which helps you address and remediate any tag category issues.

The system uses the *sn_itom_tag.recategorization_required* property to determine whether recategorization is needed. By default, this property is set to false. When a change to a category occurs, a business rule sets the property to true, indicating that recategorization is needed. A daily scheduled job then checks this property and if it is set to true, begins the recategorization process. This process automatically updates the tag mappings within 24 hours of the change, but you can also initiate recategorization to apply the changes immediately. All changes and updates are stored in the CI tag category table.

Tag Categorization and domain separation

Tag Categorization supports domain-specific management. If an instance isn't domain-separated, the categories are added to the global domain. For domain-separated instances, the categories are added to each leaf domain. When you access or manage tags, the system gets domain information from the CI and adds tag mappings to the matching category in the appropriate domain. For general information about domain separation with Tag Governance, see Domain separation in Tag Governance.

Configuring Tag Governance

Install the Tag Governance app to execute tag audits, view the Tag Data Governance dashboard, and perform tag remediation.

Configuration overview

- Domain separation in Tag Governance

  Domain separation is supported in Tag Governance. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

- Install the application

  Install the Tag Governance app to execute tag audits, view the Tag Data Governance dashboard, and perform tag remediation.

  To view the list of plugins and applications installed with Tag Governance, including any dependencies, see Plugins or applications installed with ITOM Visibility.

Install the application

Install the Tag Governance app to execute tag audits, view the Tag Data Governance dashboard, and perform tag remediation.

Before you begin

- Ensure that you use the Discovery and Service Mapping Patterns for tag-based discovery.
- Install the ServiceNow Cloud Access Interface app for cloud remediation.

Role required: sn_itom_tag.tag_governance_admin, admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Applications** > **ServiceNow Store**.
2. Use the search box to find the Tag Governance application and select **Install**.

What to do next

Configure Tag Governance policies to run on discovered resources. See Configure tag policies.

Using Tag Governance

Use the ServiceNow® Tag Governance app to identify and remediate on-premises or cloud resources that are inconsistent and don't comply with the tag policies of your organization.

Tag governance tasks

- Establishing Tag Governance policies

  Configure tag policies that define the criteria for tag audits on discovered cloud resources or CIs. View audit results on tag quality and compliance on the Tag Health dashboard. You establish tag policies that support the scale and needs of the various groups in your organization so all users can benefit.

- Configure tag policies

  Configure tag policies that define the criteria for tag audits on discovered cloud resources or CIs. View audit results on tag quality and compliance on the Tag Health dashboard.

- Configure remediation policies for tag audit findings

  Configure and preview remediation options, generate keys, and perform actions to remediate non-compliance or failures that are based on tag audit reports.

- Preview and remediate tag audit failures

  Based on the tag policies that you create, you can preview remediation flows and auto-generate missing tags for non-compliant CIs.

- Performing real-time updates to tags for cloud resources

  You can set up cloud events processing and map a tag policy to update tags for cloud resources in real-time.

Establishing Tag Governance policies

Configure tag policies that define the criteria for tag audits on discovered cloud resources or CIs. View audit results on tag quality and compliance on the Tag Health dashboard. You establish tag policies that support the scale and needs of the various groups in your organization so all users can benefit.

Considerations

Follow a methodical approach when defining tagging standards to keep the count of tags low. Standardized tag policies improve the power of the Tag Governance app to monitor and manage tags across CIs in your organization.

You might need extensive discussions and research to understand the needs of all the varied teams in your organization. For example, finance teams might want CIs to be tagged based on cost-center allocations, applications assignments, and business purpose. In contrast, IT security

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

teams might want CIs categorized by OS and security patch level as part of their vulnerability management plan.

Before you finalize your tagging approach, ensure that you have considered the following issues:

- Tag filters that various teams will use.
- The level of access that different users need.
- Whether you need real-time checks.
- Whether you need to synchronize tags with cloud resources.

Related topics

Configure tag policies

Configure remediation policies for tag audit findings

Configure tag policies

Configure tag policies that define the criteria for tag audits on discovered cloud resources or CIs. View audit results on tag quality and compliance on the Tag Health dashboard.

Before you begin

Role required: sn_itom_tag.tag_governance_admin, admin.

About this task

You can specify filters that narrow the scope of an audit, for example, to audit only service accounts and datacenters.

Procedure

1. Navigate to **All** > **Tag Governance** > **Policies**.
2. Select **New** and then fill in the Tag Policies form.
   The form fields are described in Tag Policies form.
3. Select **Submit**.
   - The Tag Policy Runs related list displays the list of tag audit runs with their corresponding state, created, and updated timestamps.
   - The Latest Policy Run Findings related list displays the results of the tag audit run.

What to do next

- Select the **Show Audit Results** related link to view all audit results.
- Select the **Run Point Scan** related link to execute all applicable checks against the record.
- For more information see, Flows and Subflows.

Configure remediation policies for tag audit findings

Configure and preview remediation options, generate keys, and perform actions to remediate non-compliance or failures that are based on tag audit reports.

Before you begin

Ensure that tag policies are configured and in Active state.

Role required: sn_itom_tag.tag_governance_admin, admin

About this task

You can run remediation automatically if you map a remediation policy to a tag policy. You can also update the tags of non-compliant CIs at the cloud provider's end in real time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure remediation policies and logic and also enable updating tags on cloud resources associated with the following CMDB classes:

- Virtual Machines
- Compute Security Groups
- Storage Volumes
- Cloud Networks
- Cloud Subnet
- Network Interfaces
- Load Balancers
- Availability Zones
- Public IP addresses
- Storage Accounts
- Resource Groups

The list of CMDB classes for which you can update tags on CIs in the CMDB as well as cloud resources also displays on the Tag Remediation form. If you select the **Update Tags in Cloud** check box in the Remediation form, remediation will only preview CI failures belonging to the CMDB classes listed.

Procedure

1. Navigate to **All** > **Tag Governance** > **Remediation**.
2. Select **New** to create a new record.
3. In the **Name** field, provide a unique descriptive name.
4. In the **Policy** field, select the lookup icon.

   Note: The **Policy Type** field is automatically populated after you select a policy from the lookup list.

5. In the Tag Policies list, configure a remediation policy in one of the following ways:
   - Select an existing tag policy from the lookup list.
   - Select **New**. Fill in fields on the Tag Policies form (the fields are described in Tag Policies form.

     Note: If you are creating a new tag policy, you must run an audit before you can run remediation flows.

   You have created or selected a policy for the remediation record.
6. **Optional:** Select the **AutoRemediation** check box to automatically run remediation on non-compliant CIs.
   When an audit run finishes, remediation runs on all CIs with policy failures.
7. **Optional:** Select the **Update Tags in Cloud** check box to apply remediation and update the tags in the relevant cloud provider environment for supported CMDB CI classes.

   Important: The **Update Tags in Cloud** feature works only for AWS and Microsoft Azure Cloud resources for CMDB classes listed in the Tag Remediation form. See Performing real-time updates to tags for cloud resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. On the Tag Remediation form, select **Submit**.
   The remediation policy is configured. The Tag Remediations list opens.

Preview and remediate tag audit failures

Preview audit reports for resources that tag policies identify as non-compliant. Remediate failures by adding or updating tags.

Before you begin

Configure tag policies and remediation policies. See Configure tag policies and Configure remediation policies for tag audit findings.

Note: The update to cloud works only for AWS and Microsoft Azure Cloud resources for CMDB classes listed in the Tag Remediation form.

Role required: sn_itom_tag.tag_admin, admin

About this task

Based on the tag policies that you create, you can preview remediation flows and auto-generate missing tags for non-compliant CIs. The Tag Remediation Preview related list displays this information. See Flows and Subflows.

Note:

- If you select the **Update Tags in Cloud** check box while configuring remediation, the Preview audit filters failures belonging to the listed CI classes.
- If you selected both the **Auto-Remediate** and **Update Tags in Cloud** settings while configuring remediation, then tags are updated both in the CMDB and in cloud resources (only for supported CMDB CI classes).
- If you select only **Auto-Remediate**, the remediation runs and updates tags in the CMDB alone.

Procedure

1. Navigate to **All** > **Tag Governance** > **Remediation**.
2. From the Tag Remediations list, select the record that you want to view.
3. Select the **Preview Remediation** related link.
   The Tag Remediation Previews related list displays the non-compliant CIs.
4. Select the Tag Remediation Previews related list to identify the missing tags and their values.
5. On the Tag Remediation Keys related list, select **New** to create a tag remediation key.
   The fields on the Tag Remediation Key form are described in Tag Remediation Key form.
6. Select **Submit**.
   You have created the Tag Remediation keys associated with the policy.
7. **Optional:** Select the **Generate Keys** related link, if you have run the **Tag Key Policy**.
   The **Generate Keys** related link appears only for the *`Tag Key Policy`* policy type. Because the policy specifies the expected tags, you can select the **Generate Keys** related link to generate the keys that you need for remediation instead of manually adding them in the Tag Remediation Keys related list.
8. **Optional:** Select the **Generate Key Values** related link, if you have run the **Tag Key & Value Policy**.
   The **Generate Key Values** related link appears only for the *`Tag Key and Value Policy`* policy type. Because you have already specified the values you are looking for in the tag policy, you can select the **Generate Key Values** related link to add the expected key values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. Select the **Remediate Tags** related link.
   The remediation flow to add missing tags initiates. A relevant information message displays on the form when the Remediation is active. Select the **More Info** link on the message to navigate to the default flow context.

Result

Based on the tag audit result findings and the Tag Remediations Keys that you create, the remediation flow adds missing tags to the non-compliant CIs to achieve policy compliance.

Performing real-time updates to tags for cloud resources

You can set up cloud events processing and map a tag policy to update tags for cloud resources in real-time.

Prerequisites

- Cloud events table [sn_cmp_cloud_events] table does not have any records in the *processed* state.
- CIs impacted by tag change events are identified by either CAPI or Discovery and Service Mapping patterns.
- Configured at least one tag policy.
- Select the **Run on cloud events** check box in the tag policy you create.

Instructions depend on the cloud provider

See the following setup instructions for cloud event processing:

- AWS events-driven discovery
- Configure the Microsoft Azure Alert service to auto-update the CMDB
- Configure the Google Cloud's Operations Suite Logging service to auto-update the CMDB

Process cloud events related to changes in tags and execute tag audits on the events in real time.

Unlike regular tag audits, where the audit is performed on all CIs in the CMDB that the tag policy applies to, tag events are processed per CI. Based on the prerequisites, tag event processing occurs on events that are identified with a tag change. The impacted CI is identified and, if the associated policy is applicable to the CI, the tag policy is run on the identified CI.

If the **Run on cloud events** check box is selected during policy creation, the policy is selected to run an audit on tag change events. If the check box is cleared, the audit is not run on the CI even if there are tag change cloud events.

View the change type that applies to the tag (the change that triggered the audit) in the Tag Change Log [sn_itom_tag_change_log] table. The records in the table are stored for 90 days by default. The change log records appear whether a remediation run or an event triggered the change on the CI. You can also view tags that were added or updated as part of the change.

Optimizing performance

High event inflow or many tag policies can decrease performance. You can modify the following system properties to confirm acceptable performance:

- [sn_itom_tag.azure_window_size] - This system property is configurable (in minutes).
- [sn_itom_tag.max_audit_lag_ratio] - Specify the lag between tag audits that are already running, before the audit triggers on the records ready for audit.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- [sn_itom_tag.policy_batch_size] - Batch size on which the tag audit is run.
- [sn_itom_tag.event_batch_size] - Event batch size to be processed.

Modify tag category information for Tag Categorization

Modify tag category definitions to support Tag Categorization and ensure that tag categories align with your requirements for data management and classification.

Before you begin

Verify that you have installed version 1.16.3 of Service Mapping Plus to access to the CI tag category and CI tag key tables. For more information, see Install Service Mapping Plus.

Role required: tag_governance_admin

Procedure

1. Navigate to **All** > *Tag Governance* > **Tag Categories**.
2. Verify that you're in the leaf domain.
   a. In the page header, select the globe icon ( ).
   b. Select **Domain scope**, and choose the appropriate leaf domain.
3. Select the category definition for the tag category you want to modify.
4. In the **Regex** field, modify the definition.
   You can delete an existing value (regex) or add a new value (regex) in this field.
5. **Optional:** Enter a value in the **Priority** field to set a priority for this tag category.
   Because a tag key can belong to multiple categories. setting the priority ensures that Tag Categorization maps CIs to the correct category.
6. Select **Update** to save the changes.
7. If you want to apply your changes immediately rather than waiting for the automatic update within 24 hours, select the **Re-Categorize Tags** button.

Add or delete tag categories for Tag Categorization

Add or delete tag categories for improved tagging of your organization's configuration items (CIs) and cloud assets.

Before you begin

Verify that you have installed version 1.16.3 of Service Mapping Plus to access to the CI tag category and CI tag key tables. For more information, see Install Service Mapping Plus.

Role required: tag_governance_admin

Procedure

1. Navigate to **All** > **Tag Governance** > **Tag Categories**.
2. Verify that you're in the leaf domain.
   a. In the page header, select the globe icon ( ).
   b. Select **Domain scope**, and choose the appropriate leaf domain.
3. Add or delete a tag category.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Description
Add a tag category	**a.** Select **New**. **b.** Enter the category name in the **Tag category name** field. **c.** Select the Lookup using list icon . **d.** Select **New** in the Tag Category Definitions window. **e.** Enter the value for the category definition in the **Regex** field. **Note:** The regex identifies the tag values that are included in the category. Separate multiple values with a vertical slash (|). **f.** Update the priority value. **g.** Select **Submit**. **h.** Select **Submit** on the CI tag category - New Record page.
Delete a tag category	**a.** Select the check box next to the tag category. **b.** In the **Actions on selected rows** menu, select **Delete**.

4. If you want to apply your changes immediately rather than waiting for the automatic update within 24 hours, select the **Re-Categorize Tags** button.

Set the priority for tag category definitions for Tag Categorization

Set the priority for tag category definitions to ensure that tag keys belonging to multiple categories are properly categorized.

Before you begin

Verify that you have installed version 1.16.3 of Service Mapping Plus to access to the CI tag category and CI tag key tables. For more information, see Install Service Mapping Plus.

Role required: admin

About this task

A tag key can belong to multiple categories. Setting the priority ensures that Tag Categorization maps Configuration Items (CIs) and cloud resources to the correct category.

For example, the tag key Application might belong to both the Application and Cost Center tag categories. Set the priorities for the Cost Center and Application tag categories to 100 and 101, respectively, to prioritize mapping to the Application category. If a conflict occurs, this setting ensures that the Application category has higher priority and automatically maps the tag key to it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Tag Governance** > **Tag Categories**.
2. Verify that you're in the leaf domain.
   a. In the page header, select the globe icon ( ).
   b. Select **Domain scope**, and choose the appropriate leaf domain.
3. Select the category definition for the tag category whose priority you want to set.
4. In the Priority field, set the value.
5. Select **Update** to save the change.
6. If you want to apply your changes immediately rather than waiting for the automatic update within 24 hours, select the **Re-Categorize Tags** button.

Add or delete tag keys for Tag Categorization

Tag keys define the values available within a tag category. Add or delete tag keys to keep your tag categorization structure current as data management and classification requirements change.

Before you begin

Verify that you have installed version 1.16.3 of Service Mapping Plus to access to the CI tag category and CI tag key tables. For more information, see Install Service Mapping Plus.

Role required: tag_governance_admin

Procedure

1. Navigate to **All** > *Tag Governance* > **Tag Categories**.
2. Verify that you're in the leaf domain.
   a. In the page header, select the globe icon ( ).
   b. Select **Domain scope**, and choose the appropriate leaf domain.
3. Select the tag category to which you want to add or delete tag keys.
4. Under **CI tag keys**, either add a tag key or delete an existing tag key.

Action	Steps
Add a tag key	**a.** Double-click the empty row under **Tag key**, where you see **Insert a new row**. **b.** In the field, add a key name. **c.** Select the check mark icon ( ) to save the tag key. **d.** Repeat the previous steps to add more tag keys. **e.** Select **Update** to save your changes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Steps
Delete a tag key	**a.** Under **Tag key**, select the delete icon (⊗) next to the tag key you want to remove. **b.** Select **Update** to save your changes.

5. If you want to apply your changes immediately rather than waiting for the automatic update within 24 hours, select the **Re-Categorize Tags** button.

Tag Governance reference

Reference topics provide additional information on Tag Governance forms and other resources.

Domain separation in Tag Governance

Domain separation is supported in Tag Governance. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

Domain separation

You can identify CIs or cloud resources that are inconsistent and do not comply with the policies that you configure to govern tag data. In Tag Governance, domain separation targets tenant domain requester use cases in an application. The application has been designed to support requester activities within tenant domains. Logic is in place to route data to tenant domains, based on applicable use cases. The owner of the instance must be able to set up the application to function normally across multiple tenants. The application handles data routing to domains.

All tables in Tag Governance support data separation only; delegated domain separation is not supported.

To ensure that domain-sensitive information is not available at a global domain, ensure that you create policies at a domain level and not at a global level. While managing service providers, create policies at the service provider level. If you create policies at a global level instead, when you run the audit, the "run" findings are created in the Global domain instead of respective domains.

Related topics

Domain separation for service providers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag policies and remediation for AWS

Tag audits apply policies to discovered CIs to determine tag compliance; existence of tags, appropriate count of tags, and the presence of specified key-value pairs.

Tag policies

Run a tag audit to audit discovered CIs for the following tag policy types:

- Tag Count: Checks CIs for the tag key count that you specify
- Tag Presence: Checks CIs for the presence of the tag key values that you specify
- Tag Key & Value: Checks CIs for the presence of key-value pairs that you specify.

Specify string value or tag keys that you are looking for in a comma-separated format in the Tag Presence policy type. You can also specify a number in the Tag Count check policy type to identify CIs with one or more tags. After you run the tag audits, you can view audit results and configure remediation measures that are based on the compliance index quality. Optionally, you can also assign remediation tasks for non-compliant CIs to user groups and users for follow-on tasks.

Auto-remediation works with the Assume Role on AWS to automate the tagging of cloud resources, by following the steps in Configure temporary credential access for trusted AWS accounts.

The Assume Role must have the following permissions (policy) for remediation to work:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "tag:TagResources",
            "Resource": [
                "arn:aws:ec2:*:*:instance/*",
                "arn:aws:ec2:*:*:security-group/*",
                "arn:aws:ec2:*:*:volume/*",
                "arn:aws:ec2:*:*:vpc/*",
                "arn:aws:ec2:*:*:subnet/*",
                "arn:aws:ec2:*:*:network-interface/*",

 "arn:aws:elasticloadbalancing:*:*:loadbalancer/*",
                "arn:aws:ec2:*:*:availability-zone/*",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                "arn:aws:ec2:*:*:public-ip/*",
                "arn:aws:storagegateway:*:*:gateway/*",
                "arn:aws:resource-groups:*:*:group/*"
            ]
        }
    ]
}
```

Note:

While performing tag remediation, if it is found that Discovery is credential-less, MID Server with Tag Management capability is picked. Hence, add Tag Management capability to the correct MID Server with service account IAM role attached. This action is essential because there might be several accounts that IAM role MID Servers must choose from.

Note: You must have the ITOM Visibility entitlements to run remediation on CIs in the CMDB. For more information, contact Customer Service and Support.

Tag Policies form

You use the Tag Policy form to configure tag compliance and health policies that are used in tag audits runs on discovered cloud resources or CIs.

Tag Policies form

Field	Description
Name	Unique and descriptive name for the policy.
Short Description	A short summary that describes the policy.
Policy Type	Policy type selection for the tag compliance and health audit. • Tag Count Policy: Check for a specific count of tags for CIs. • Tag Key Policy: Check for tags with specific values (string expressions). • Tag Key & Value Policy: Check for key-value pair and value strings. • Tag Category Policy: Check for tags with a specific category. **Important:** You must work in a leaf domain to use the Tag Category Policy. Based on your selection, you will specify additional values.
Expected Tag Count	This field appears only when the **Tag Count Policy** policy type is selected.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag Policies form (continued)

Field	Description
	CIs with fewer tags than the specified number are considered as failed or non-compliant. CIs that meet the value are considered certified against this policy. For example, if you specify `4`, the policy checks whether the CIs have at least four tags specified.
Expected Tag Keys	This field appears only when the **Tag Key Policy** policy type is selected. Policy that checks discovered CIs for the tags specified in the **Expected Tag Keys** field. Specify multiple tags or string expressions in a comma-separated list. For example, if you specify `prod,finance,cost center` the policy checks whether the CIs have all the specified tags. CIs that have all the specified tags as certified or fully compliant. If CIs are missing any of the listed tags, they're considered failed or non-compliant.
Tag Key & Value Checks	Execute this policy to check discovered CIs for the tag key and values specified in the **Tag Keys & Value Checks** field. You can specify multiple tag key & value pairs or specify a script to check for a collection of values. For example, you can specify `user` in the **Key** field, and `admin` in the **Value** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag Policies form (continued)

Field	Description

	Note:
	The **Script** field appears only when **Tag Key & Value Policy** is selected in the **Policy Type** field. Use the **Script** field to specify multiple values to return a string set.
	Example code
	`// The script should return a string set to tagValues variable.`
	`// the code sets the tagValues array with name of the CI`
	`// by this the check will be done for the tag to match the CI's name`
	`tagValues = [ci.name.toString()];`

Active	Option to activate or deactivate the run schedule.
Run on cloud events	Option to enable real-time audit runs when cloud events occur. **Note:** Ensure that you have configured auto-updating the CMDB with cloud events for the relevant cloud provider.
Save cert audit result	Option to save the certification audit findings and incorporate tag audit outcomes in a CI's health calculations. **Note:** This option can be activated if there is a business need to incorporate tag audit outcomes in CI health calculations.
Run	Frequency at which the tag audit runs. • **Daily** • **Monthly** • **Weekly**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag Policies form (continued)

Field	Description
	• **Periodically** • **Once** • **On Demand** **Note:** To run an audit immediately after configuring the policy, select the **On Demand** option. • **Business Calendar: Entry Start** • **Business Calendar: Entry End**
Next scheduled run date	Date and time on which the tag audit is next scheduled to run. This field appears only for scheduled runs and not for on-demand runs.
Filters	Condition statements with a series of contextually generated fields. Use lookup lists to build conditions and filter the CIs eligible for audit. Dynamic conditions and filters you specify to fetch the desired tags based on service accounts and associated datacenters.
Table	Displays the *Configuration Item [cmdb_ci]* table to which filters are applied.
Policy Criteria	Other policy criteria that you can configure using dynamically populated reference lookup lists. The Configuration Item [cmdb_ci] table populates the data for this lookup list. For more information, see the Reference lookup.
Tag Filter	Filter CIs further based on **Key**, **Keywords**, and **Value** that you specify. Use the condition builder to customize a filter.
Task Assignment	Option to assign follow-on tasks to users or user groups.
Create Tasks	Option to create follow-on tasks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag Policies form (continued)

Field	Description
	Note: If there is an existing task mapped to a CI and tag policy, the *sn_itom_tag.allow_new_cert_follow_on_tas* system property determines whether to create a new task or update an existing one when there is a compliance failure: • To update the existing task with the compliance failure details but not create a new task, set the property to false (default) • To create a new task, set the property to true.
Group	This field appears only when you select the **Create Tasks** check box. Assignment group to which to assign the task. Use lookup lists to select a group.
User	Specific user from the assignment group to whom you want to assign the task. Use lookup lists to select a user. This field appears only when you select the **Create Tasks** check box.

Tag Remediation Key form

You use the Tag Remediation Key form to create a tag remediation key.

Tag Remediation Key form

Field	Description
Key	Tag key value that is used to remediate non-compliant CIs. For example, if the remediation preview shows the value as `cost center`, you can specify `cost center` as a new remediation key if it doesn't exist.
Order	Order of execution that the system considers when you initiate remediation. For example, you can run a tag count policy for four tags and identify CIs that are missing two tag keys. If you've configured multiple remediation keys, the system considers the two keys with the lowest values in the order field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag Remediation Key form (continued)

Field	Description
Remediation	Remediation name and record preview.
Value Type	Type of the value that you specify in the **Value** field — **Static** or **String**. Based on your selection, you must specify a script or a value for the tag key.
Value	String value for the key that you're creating. This field appears only when **String** is selected in the **Value Type** field.
Script	Script code to generate values for the key. This field appears only when **Script** is selected in the **Value Type** field.

Components in Tag Governance

Several key components are installed with, or linked to, the Tag Governance plugin to help standardize and manage your organization's tags.

Roles installed

Role title	Description	Contains roles
[tag_governance_admin]	Sets up and manages the tag governance application. Defines and enforces tag policies, manages tag categories and values, and promotes compliance with organizational standards. Performs advanced configuration and customization of tag governance features. Assign this role to application administrators.	• [itil] • [certification_admin] • [discovery_admin] • [flow_operator]
[tag_reports_viewer]	Views and analyzes tag reports related to tag usage and compliance. Assign this role to users who monitor and analyze tag data.	[cmdb_read]

Tag categories installed

Starting with version 1.7.0, the predefined categories listed in the table are installed with Tag Governance and require the installation of at least Service Mapping Plus version 1.16.3. For more information see Install Service Mapping Plus.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag category	Definition (Regex)
Application ID	applicationid|application-id|app-id|appid|application id|app id
Assignment Group	assignment-group|assignment_group|assignment group|assignmentgroup|assign group
Cost Center	costcenter|cost-center|cost_centre|costcentre|cost-centre|cost centre
Environment	env|envt|environ|environment
Owner	own|owner|owners|ownership|owned

Properties installed

Property	Description
sn_itom_tag.glide.script.block.client.globals	Controls the access client-side scripts have to global variables for enhanced security. If set to true, these scripts can't access global variables. • **Type:** boolean • **Value:** true/false • **Default value:** False
sn_itom_tag.allow_new_cert_follow_on_task	Enables the creation of new certification follow-up tasks for each audit in Tag Governance. • **Type:** boolean • **Value:** true/false • **Default value:** False
sn_itom_tag.azure_window_size_in_mins	Time window size, in minutes, to check for changes in Azure Resources. This value helps determine the start and end times for the check: **Start time:** The time the event occurred minus the time window size. **End time:** The time when the event occurred plus the time window size. • **Type:** integer • **Default value:** 10
sn_itom_tag.performance.config.bulksize	Sets the number of app data items to process at one time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Type:** integer • **Default value:** 5000
sn_itom_tag.last_event_processed	Indicates the last tag event successfully processed to verify that the system is up to date with the latest changes. • **Type:**string • **Default value:**
sn_itom_tag.max_audit_lag_ratio	Sets the maximum acceptable ratio of delay for tag audits to confirm that tag policy audits are performed within a specified time frame. • **Type:** integer • **Default value:** 2
sn_itom_tag.tag_event_batch_size	Sets the number of tag events, such as the creation, update, or deletion of tags, processed together in a single batch. • **Type:** integer • **Default value:** 500
sn_itom_tag.policy_batch_size	Sets the number of items processed together in a single batch when a tag policy audit is run. • **Type:** integer • **Default value:** 1000

Cloud Discovery Workspace

Cloud Discovery Workspace offers a comprehensive solution to manage the cloud operations of your organization.

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

Get Started

Choose one of these tiles to get started.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Install	Use
Install Cloud Discovery Workspace	Use Cloud Discovery Workspace

Troubleshoot and get help

- Ask or answer questions in the Cloud Discovery Workspace forum on the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Install Cloud Discovery Workspace

You can install the Cloud Discovery Workspace application (com.cloud_operations_workspace) if you have the admin role.

Before you begin

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

When Cloud Discovery Workspace is installed, several plugins and applications are also installed. For more information, see Plugins or applications installed with ITOM Visibility.

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Cloud Discovery Workspace requires the following plugin. Ensure that the following plugin is activated before you install Cloud Discovery Workspace.

  Required ServiceNow plugins

  Discovery (com.snc.discovery)

  ServiceNow Cloud Discovery finds applications and devices on your network, and then updates the CMDB 360 with the information it finds. For more information on activating Discovery, see Request Discovery.

Role required: admin

About this task

When you install the Cloud Discovery Workspace application, it automatically installs the following role:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Discovery Workspace role

Role name	Description and tasks	Nested roles
sn_cloud_ops_ws.cloud_ops_admin	The Cloud Discovery Workspace admin can perform the following actions: • View the Cloud Discovery Workspace dashboard • Create Cloud Discovery schedules • Monitor Cloud Discovery operations • View Cloud Discovery reports	• sn_itom_ui_comp.itom_component_user • discovery_admin • cloud_admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Cloud Discovery Workspace application (com.cloud_operations_workspace) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.
3. In the Application installation dialog box, review the application dependencies.

   Dependent plugins and applications are listed if they will be installed, are currently installed, or need to be installed. If any plugins or applications need to be installed, you must install them before you can install Cloud Discovery Workspace.
4. **Optional:** If demo data is available and you want to install it, select the **Load demo data** check box.
   Demo data are the sample records that describe application features for common use cases. Load the demo data when you first install the application on a development or test instance.
5. Select **Install**.

Cloud Discovery Workspace dashboard

The ServiceNow® Cloud Discovery Workspace dashboard provides a summary of the cloud operations of your organization and shows the ServiceNow applications that you can use to manage them.

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Required ServiceNow AI Platform roles

sn_cloud_ops_ws.cloud_ops_admin.

Access the Cloud Discovery Workspace dashboard

To open the dashboard, navigate to **All** > **Cloud Discovery Workspace**.

Use cases

User	Dashboard use
sn_cloud_ops_ws.cloud_ops_admin	• Gain insight into the trend of CI discovery, trend of Cloud Discovery events, and Cloud Discovery errors • Select the cloud provider of your organization. • Access the applications that you can use to manage the cloud operations of your organization.

Indicators

Discovery errors today: Count of all Cloud Discovery errors that occurred during the scheduled discoveries over the last 24 hours.

Reports

The following reports are available in the Cloud Discovery Workspace dashboard:

Title	Type	Source table	Description
Total configuration items	Line and area chart 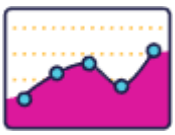	cmdb_ci	This report displays the number of CIs discovered over time. Select the report to view the list of all the configuration items discovered so far. You can select any CI record to view more information about the CI and its multi-source data if available.
Events by provider - last 24 hours	Line chart	sn_cmp_cloud_event_list	This report displays the trend of the cloud events received over the last 24 hours. The report contains a separate trend line for each cloud provider.

In addition, the dashboard cards display the applications that you can use to manage the cloud operations of your organization. Use the Cloud provider selection page to select the cloud providers of your organization. The selection helps to customize the menus of the Cloud Discovery Workspace applications and display only the cloud providers that are relevant for your organization.

Related topics

List of workspaces

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Access the UI Builder in Cloud Discovery Workspace

Access your Configurable Workspace experience in UI Builder for editing.

Before you begin

Role required: ui_builder_admin or admin

Procedure

1. In the navigation filter, navigate to **All** > **Now Experience Framework** > **UI Builder**.
2. Select **Experiences**.
3. Select your experience.

Result

Your configurable experience opens in UI Builder.

Kubernetes Explorer

Use the ServiceNow® Kubernetes Explorer to drill-down and view the Kubernetes environments and resources of your organization.

Important: Cloud Discovery Workspace is now deprecated and no longer supported or available for new activation. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

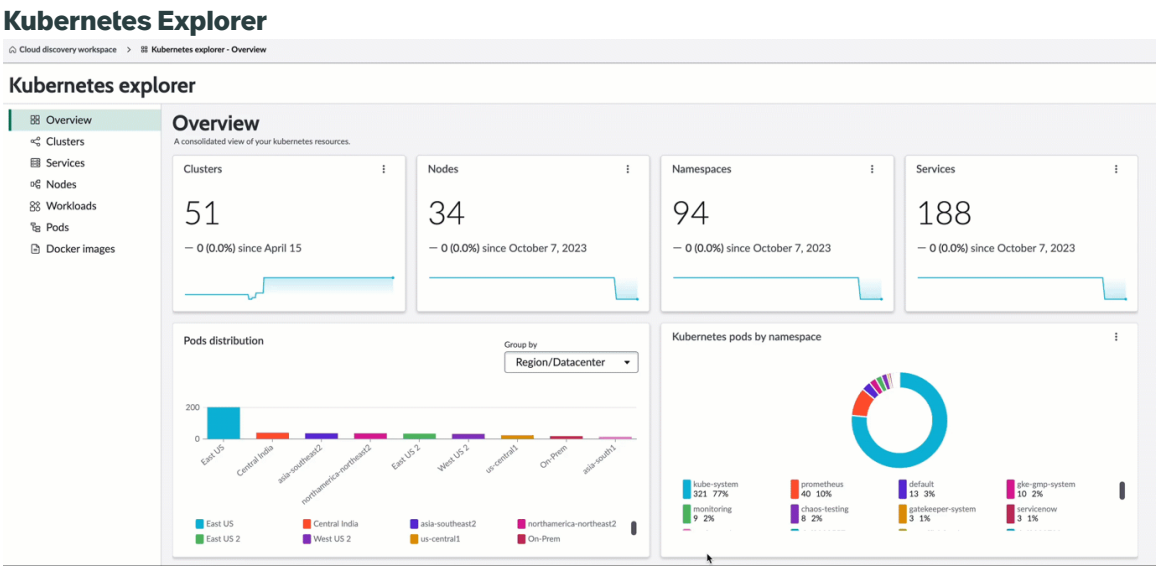

The Kubernetes Explorer supports the following cloud providers:

- Amazon Web Services (AWS): Amazon Elastic Kubernetes Service (EKS)
- Microsoft Azure Cloud: Azure Kubernetes Service
- Google Cloud Platform (GCP): Google Kubernetes Engine (GKE)
- OpenShift

Important: In the Kubernetes Explorer, any Kubernetes resource hosted on OpenShift is considered on-prem. Whereas, any Kubernetes resource hosted on Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) is considered cloud hosted.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Required ServiceNow AI Platform roles

The sn_cloud_ops_ws.cloud_ops_admin role is required to view the Kubernetes Explorer.

Access the Kubernetes Explorer

To open the dashboard, navigate to **All** > **Cloud Discovery Workspace** > **Kubernetes Explorer**.

Important: Starting with the Zurich release, Cloud Discovery Workspace is being prepared for future deprecation. It will be hidden and no longer activated on new instances, but will continue to be supported. Discovery Admin Workspace provides the latest experience for this functionality. For details, see the Application/Plugin Deprecation Process [KB0867184] article in the Now Support knowledge base.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
Application owner	Review the statistics of the applications and workloads running on the Kubernetes clusters.
IT admin, cloud admin	• Review the statistics of the Kubernetes infrastructure: Clusters, Nodes, Namespaces, Services, Pods, and Workloads. • Review the distribution of the pods across dimensions such as region/datacenter, geo location, or country. • Review the distribution of the pods across namespaces. • Review the list of Kubernetes docker images used in your organization.

Indicators

Clusters

A Kubernetes cluster is a set of nodes that run containerized applications. The Clusters indicator displays the daily total count and discovery trend of the Kubernetes clusters discovered across the Kubernetes deployments of your organization.

The Kubernetes Explorer reads the Kubernetes Cluster Analytics [sn_cow_k8s_cluster_analytics] table to compute the Clusters indicator.

Nodes

A node is a worker machine in Kubernetes and may be either a virtual or a physical machine, depending on the cluster. The Nodes indicator displays the daily total count and discovery trend of the Kubernetes nodes discovered across the Kubernetes deployments of your organization.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Kubernetes Explorer reads the Kubernetes Node Analytics [sn_cow_k8s_node_analytics] table to compute the Nodes indicator.

Namespaces

Kubernetes namespaces provides a mechanism for isolating groups of resources within the cluster. The Namespaces indicator displays the daily total count and discovery trend of the Kubernetes namespaces discovered across the Kubernetes deployments of your organization.

The Kubernetes Explorer reads the Kubernetes Namespace Analytics [sn_cow_k8s_namespace_analytics] table to compute the Namespaces indicator.

Services

Kubernetes service is an abstraction which defines a logical set of pods and a policy by which to access them (sometimes this pattern is called a micro-service). The Services indicator displays the daily total count and discovery trend of the Kubernetes services discovered across the Kubernetes deployments of your organization.

The Kubernetes Explorer reads the Kubernetes Service Analytics [sn_cow_k8s_service_analytics] table to compute the Services indicator.

The Services indicator is also available on the Services tab of the Kubernetes Explorer.

Breakdowns

- Pod distribution: This breakdown displays the distribution of Kubernetes pods hosted on the on-prem infrastructure and the cloud infrastructure.

  You can filter the breakdown across the following dimensions:

  - Region/Datacenter
  - Geo Location
  - Country

  The Pod distribution breakdown is available on the Overview tab of the Kubernetes Explorer.

- Kubernetes services by platform or Kubernetes services by type: This breakdown displays the distribution of Kubernetes services across dimensions such as platform and type.
- Nodes by region: This breakdown displays the distribution of the Kubernetes nodes across hosting regions.

Filter

You can filter the list by interacting with the data visualizations such as donut and single score card.

Data visualizations

Title	Type	Source table	Description
Kubernetes pods by namespace	Donut	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	This report displays the distribution of Kubernetes pods across the discovered namespaces.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
			This report is available on the Overview tab.
Total clusters	List	Kubernetes Cluster Analytics [sn_cow_k8s_cluster_analytics]	This report lists the Kubernetes clusters discovered across the Kubernetes deployments of your organization. This report is available on the Clusters tab. When you click the cluster name, the Dependency Views map appears for Cluster Configuration Item.
Total services	Single score	Kubernetes Workload Analytics [sn_cow_k8s_service_analytics]	This report displays the total count of the Kubernetes services. This report is available on the Services tab.
Kubernetes services by platform / type	Donut	Kubernetes Node Analytics [sn_cow_k8s_service_analytics]	This report displays the distribution of the Kubernetes services across the providers. This report is available on the Services tab.
Kubernetes services	List	Kubernetes Namespace Analytics [sn_cow_k8s_service_analytics]	This report lists the Kubernetes services discovered. This report is available on the Services tab. When you click the service name, the Dependency Views map appears for Cluster Configuration Item.
Nodes by cluster	Donut	Kubernetes Node Analytics [sn_cow_k8s_node_analytics]	This report displays the distribution of the Kubernetes nodes across the discovered clusters. This report is available on the Nodes tab.
Nodes by region	Donut	Kubernetes Node Analytics [sn_cow_k8s_node_analytics]	This report displays the distribution of the Kubernetes nodes across the discovered clusters. This report is available on the Nodes tab.
Total nodes	List	Kubernetes Node Analytics [sn_cow_k8s_node_analytics]	This report lists the Kubernetes nodes discovered across the Kubernetes deployments of your organization. This report is available on the Nodes tab. When you click the node name, the Dependency Views map appears for Cluster Configuration Item.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
Total workloads	Single score	Kubernetes Workload Analytics [sn_cow_k8s_workload_analytics]	This report displays the total count of the Kubernetes workloads discovered across the Kubernetes deployments of your organization. This report is available on the Workloads tab.
Desired/ available replicas mismatch	Single score	Kubernetes Workload Analytics [sn_cow_k8s_workload_analytics]	This report displays the count of mismatch between the desired and available replicas. If there's a mismatch, take suitable action to remediate the issue.
Workloads by class	Donut	Kubernetes Workload Analytics [sn_cow_k8s_workload_analytics]	This report displays the workloads by the workload class. This report is available on the Workloads tab.
Total workloads	List	Kubernetes Workload Analytics [sn_cow_k8s_workload_analytics]	This report lists the Kubernetes workloads discovered across the Kubernetes deployments of your organization. This report is available on the Workloads tab. When you click the workload name, the Dependency Views map appears for Cluster Configuration Item.
Total pods	Single score	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	This report displays the total count of Kubernetes pods discovered across the Kubernetes deployments of your organization. This report is available on the Pods tab.
Running pods	Single score	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	This report displays the total count of the Kubernetes pods that are in running state. This report is available on the Pods tab.
Pending pods	Single score	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	This report displays the total count of the Kubernetes pods that are in pending state. This report is available on the Pods tab.
Creating pods	Single score	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	This report displays the total count of the Kubernetes pods that are in the creating state.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Source table	Description
	42		This report is available on the Pods tab.
Crashloop backoff pods	Single score 42	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	This report displays the total count of the Kubernetes pods that are in the crashloop backoff state. This report is available on the Pods tab.
Total pods	List	Kubernetes Pod Analytics [sn_cow_k8s_pod_analytics]	This report lists the Kubernetes pods discovered across the Kubernetes deployments of your organization. This report is available on the Pods tab. When you click the pod name, the Dependency Views map appears for Cluster Configuration Item.
Docker images	List	Kubernetes Image Analytics [sn_cow_k8s_img_analytics]	This report lists the Docker images discovered across the Kubernetes deployments of your organization. This report is available on the Docker images tab. When you click the doker image name, the Dependency Views map appears for Cluster Configuration Item. You can download the SBOM. To generate the SBOM, see Download a software bill of materials (SBOM).

Dependency Views map

ServiceNow® The Dependency Views map shows an overall top-down Kubernetes perspective, starting from a top-level cluster to the container and within the container.

A Dependency Views map has one starting point, called the root CI or root node of the map. The root CI is surrounded by a darker frame that repaints itself with a pulsing effect drawing the attention to the root CI. The maps can show both upstream and downstream dependencies for the root CI. By default the Dependency Views map displays three levels, both upstream and downstream relationships. Administrators can configure the number of levels displayed. The map collapses and expands clusters to make them easier to view. By default, clusters are collapsed.

The Dependency Views module is active in all instances, and includes demo data.

Dependency Views sample map

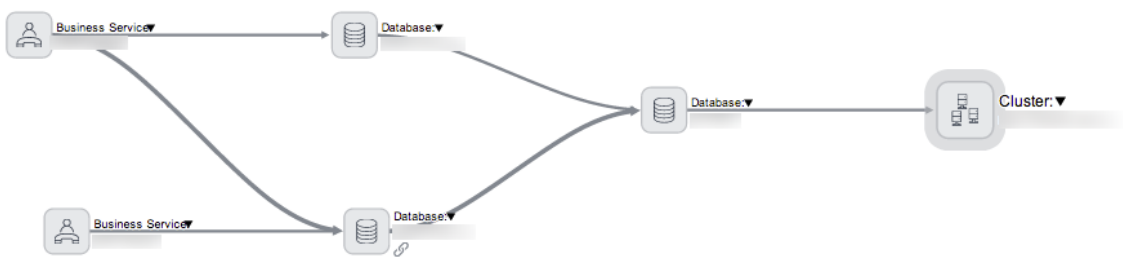

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In a Dependency Views map, map indicators indicate if a CI has any active, pending issues. You can investigate the tasks that are connected to a CI to get more details. When you return to the map from another form, the system restores the last map viewed, using the default filter and layout settings.

Many of the relationships in the map are created through the discovery process. You can also create, define, and delete CI relationships in the map. You can display the map from different perspectives and open specific records that relate to configuration items. The system refreshes the map automatically to reflect changes to the CMDB.

Note: CIs not extended from the Configuration Item [cmdb_ci] table, aren't displayed in Dependency Views maps and in CI relation formatters.

Roles

Users with the itil and ecmdb_admin roles can view maps and perform all actions in the map. Actions include access to the map views and saved filters, both from the lists in the map and from the **Saved Filters** module.

Cloud License Estimator

Cloud License Estimator estimates the resource count for all cloud resources eligible for licensing.

Explore	Configure	Generate Reports
Learn about Cloud License Estimator.	Set up credentials and configure accounts.	Generate reports with Cloud License Estimator.

Cloud License Estimator features

Learn about Cloud License Estimator features.

Use Cloud License Estimator to get the estimated resource count for all cloud resources eligible for licensing based on ITOM and CCM licensing rules. It validates the cloud account details provided by the user and estimates the resource count based on the prevailing licensing rules.

Cloud License Estimator provides the count of cloud resources along with the required license estimates. CLE supports AWS and Azure cloud environments, but not government cloud services. You can create a configuration by providing the necessary details and then run the license estimator to get the report. You can also download the report in PDF format.

Configure Cloud License Estimator

Set up the accounts and credentials needed for Cloud License Estimator.

Configure Azure Credentials for Cloud License Estimator

Configure Azure Service Principal credentials for Cloud License Estimator. If credentials are already configured in the ServiceNow instance, those credentials can be used.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: admin

Procedure

1. Navigate to **Filter Navigation** > **All** > **ITOM Cloud License Estimator** > **Home**.
2. On the Home page, select **Go to Credentials** to open the **Credentials (discovery_credentials)** page.
3. Select **New** > **Azure Service Principal** and enter the Azure Service Principal Account details.

Result
A **Credential ID** is generated, which is required when creating configurations.

Create an Azure account configuration and generate CLE reports

Create an Azure account configuration to estimate the cloud resource count for licensing. Management accounts include all subscriptions under that account.

Before you begin
Role required: admin

Procedure

1. Navigate to **Filter Navigation** > **All** > **ITOM Cloud License Estimator** > **Home**.
2. On the Home page, select **Create Configuration** to open the Create Configuration dialog.
3. Enter the following details:
   - **Name**: Enter a name for the configuration.
   - **Cloud Provider**: Choose **Azure** from the list. Fill in the additional fields that appear.
   - **Account ID**: Enter the management account ID or subscription ID.
   - **Management group**: Select this checkbox if the Account ID is a management account.
   - **Credentials**: Select the Azure Service principle credentials configured in Configure Azure Credentials for Cloud License Estimator.
   - **Activate**: Enable this option.
4. Select **Save and run license estimator**.

   The configuration is saved and the tool generates the report. You are redirected to the **Report** page where the status of the report generation can be tracked.

5. From the **Report** page, select **Download resource report** to download a PDF of the report.

Related topics

Generate Cloud License Estimator reports

Configure AWS credentials for Cloud License Estimator

Configure IAM user permanent credentials for an AWS account.

Before you begin
Role required: admin

About this task
You can configure IAM roles to generate temporary credentials using STS API for a specified account. When using IAM roles for temporary credentials, the Credential ID is optional in the configuration form. However, AWS service accounts require an IAM role to be configured in

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

service accounts. At least one service account must have credentials to generate temporary credentials.

Procedure

1. Navigate to **Filter Navigator** > **All** > **ITOM Cloud License Estimator** > **Home**.
2. On the Home page, select **Go to Credentials** to open the **Credentials (discovery_credentials)** page.
3. Select **New** > **AWS Credentials** and enter the requisite information.
   A **Credential ID** is generated, which is used when creating configurations.

Configure AWS service accounts for Cloud License Estimator

There are several ways to configure AWS Service Accounts depending on the credentials and master accounts.

Before you begin
Role required: admin

About this task
You can configure service accounts with these options:

- Configure a service account with permanent credentials. When an account is configured with permanent credentials, it is referred to as an accessor account. This account can be either a management account or member account.
- Configure a management account with the IAM role.
- Configure a member account with a custom IAM role. If the member account uses the default IAM Role (OrganizationAccountAccessRole) to generate temporary credentials, configuring a service account for the member account is not required.

Procedure

1. Navigate to **Filter Navigation** > **All** > **ITOM Cloud License Estimator** > **Home**.
2. Select the **Service Account List** icon.
3. On the **Service account** page, select **Create New**.
4. **Optional:** To configure a service account with permanent credentials, enter the following information:
   - (Optional) **Name**: Enter the AWS account name.
   - (Optional) **Cloud provider**: Select AWS.
   - (Optional) **Account ID**: Enter the AWS account ID.
   - (Optional) **Master account**: Select if this is a management account.
   - (Optional) **Credentials**: Select the AWS Credential ID created for permanent credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- (Optional) **Master account**: Leave blank if not applicable.
- (Optional) **Accessor Account and Accessor role ARN**: Not applicable for permanent credentials.

5. **Optional:** To configure a management account with the IAM role, enter the following information:
   - (Optional) **Name**: Enter the AWS account name.
   - (Optional) **Cloud provider**: Select AWS.
   - (Optional) **Account ID**: Enter the AWS account ID.
   - (Optional) **Master account**: Select if this is a management account.
   - (Optional) **Credentials**: Leave blank. No credentials are required for this account.
   - (Optional) **Master account**: Leave blank.
   - (Optional) **Accessor Account**: Select the account that has a trust relationship with this management account's IAM role.
   - **Accessor role ARN:** Enter the full ARN of the IAM role.
6. **Optional:** To configure a member account with a custom IAM role, enter the following information:
   - (Optional) **Name**: Enter the AWS account name.
   - (Optional) **Cloud provider**: Select AWS.
   - (Optional) **Account ID**: Enter the AWS account ID.
   - (Optional) **Master account**: Do not select this, as this is a member account.
   - (Optional) **Credentials**: Leave blank. No credentials are required for this account.
   - (Optional) **Master account**: Select the management account from the list.
   - (Optional) **Accessor Account**: Choose the account that has a trust relationship with this member account's IAM role. This can be either the management account or another member account.
   - **Accessor role ARN:** Enter the full ARN of the IAM role.
7. After providing the relevant information for your account type, select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create an AWS account configuration and generate CLE reports

Create an AWS account configuration to estimate the cloud resource count for licensing. Management accounts include all subscriptions under that account.

Before you begin

Role required: admin

Procedure

1. Navigate to **Filter Navigation** > **All** > **ITOM Cloud License Estimator** > **Home**.
2. On the Home page, select **Create Configuration** to open the Create Configuration dialog.
3. Enter the following details:
   - **Name**: Enter a name for the configuration.
   - **Cloud Provider**: Choose **AWS** from the list. If you choose the credential based option, additional fields appear.
   - **Account ID**: Enter the AWS Account ID.
   - **Management group**: Select this checkbox if the Account ID is a management account.
   - **Credentials**: Select the AWS credentials configured for this account.
     - **Config aggregator account**: Enter the Account ID where the aggregator is created.
     - **Config aggregator name**: Enter the name of the AWS Config Aggregator.
     - **Config aggregator region**: Enter the AWS Region where the aggregator is created.
   - **Activate**: Enable this option.
4. Select **Save and run license estimator**.

   The configuration is saved and the tool generates the report. You are redirected to the **Report** page where the status of the report generation can be tracked.

5. From the **Report** page, select **Download resource report** to download a PDF of the report.

Related topics

Generate Cloud License Estimator reports

Generate Cloud License Estimator reports

Cloud License Estimator reports provide an overview of account information.

The downloaded PDF of the resource report contains an overview of the account's information, such as the Cloud Provider, Account Configuration name, and total resource count.

The report distinguishes between CCM Licensable Resources and ITOM Licensable Resources. On the following pages, the report contains tables of the CCM and ITOM resources.

ITOM Infra Services Workspace

The dashboard integrates several tables so that issues can be monitored and resolved. The workspace provides real-time monitoring, proactive alerts, automation of common tasks, and centralized diagnostics to reduce downtime and support tickets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Explore	Configure	Workspace applications
Learn about ITOM Infra Services Workspace.	Set up your workspace.	Monitor your applications.

Explore ITOM Infra Services Workspace

The ITOM Infrastructure Services Workspace provides real-time monitoring, proactive alerts, automation of common tasks, and centralized diagnostics to reduce downtime and support tickets.

The ITOM Infrastructure Services Workspace dashboard provides an overview of applications, such as your active MID Servers and ACC agents, and their issues. The dashboard enables you to drill into issues by category and take corrective actions.

For example, selecting a MID Server card shows a table of MID Servers. Then the **Server Actions** drop-down menu enables you to perform any available action such as validating or re-keying the MID Servers. The **Server Actions** menu enables you to perform corrective actions to multiple MID Servers at once from the overview.

The ITOM Infrastructure Services dashboard also displays the top three ACC errors on the ACC agent card. You can view the error and have a quick glance at its details. Select **View all** to see all the errors.

Configure ITOM Infra Services Workspace

The dashboard provides real-time monitoring, proactive alerts, automation of common tasks, and centralized diagnostics to reduce downtime and support tickets.

The ITOM Infra Services Workspace is available out of the box with the Australia family release, and is available as a Store release in the Yokohama and Zurich family releases.

The ITOM Infra Services Workspace and MID Server workspace are provided by default. Additional applications are installed if you have the corresponding plug-ins on the instance.

Use ITOM Infra Services Workspace

The dashboard integrates different applications to provide real-time monitoring, proactive alerts, automation of common tasks, and centralized diagnostics to reduce downtime and support tickets.

MID Server Workspace

See a concise overview of MID Server activity. The workspace provides real-time monitoring and enables corrective actions across multiple MID Servers.

ITOM Infrastructure Services dashboard

The MID Server panel is displayed on the ITOM Infrastructure Services dashboard to provide an overview of MID Server issues.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Selecting any of the cards in the MID Server panel, such as **Servers down** or **Servers with issues**, takes you to that detailed table.

This page also provides a set-up link to Download a MID Server.

MID Servers tables

Navigate to the MID Servers workspace either by selecting its icon from the application list or by selecting a card from the dashboard. Each of the cards on the dashboard corresponds to the cards in this workspace view.

Selecting a card shows a table of MID Servers filtered by category. You can customize the filter by selecting the filter icon, and save the filters with the **Saved filters** drop-down menu. Group and sort the results by selecting the **Show column filter row** icon in the category row.

On the **Total MID Servers** table, a banner enables you to quickly filter for MID Servers that have been created in the past week. Use this filter to monitor new MID Servers and diagnose any issues from their deployment.

Server Actions

MID Servers can be selected by the check box on the table, then the **Server Actions** drop-down menu enables you to perform any available action. These actions include Validating or Rekeying the MID Servers. This **Server Actions** menu enables you to perform corrective actions to multiple MID Servers at once from the overview.

Detailed view

Selecting a MID Server's name navigates to a detailed view of that MID Server. On the Status, Validated, and Unresolved Issues cards, relevant correction actions are shown such as upgrading or validating the MID Server. If the MID Server has any alerts, they are shown in a banner.

If the user has the SOAP ECC queue role, then the View ECC Queue button becomes visible, which navigates to the ECC Queue table.

In the **MID Server details** view, selected from the context menu, you can edit the **Purpose** field to save any notes about the MID Server's purpose. This field is only a text note and does not change anything about the MID Server.

The **Purpose** field can be viewed from the dashboard tables by adding it to the **Personalize fields** menu, accessed through the three-dot menu beside the **Server Actions** menu.

The **MID Server configuration** and **Troobleshooting** sections of the context menu show the records assigned to this MID Server. You can access the MID Server properties or logs. Use the **Create** button to create new records, or the **Manage** button to manage existing ones. You can filter the list of records by selecting the filter icon.

Use the Agent Client Collector (ACC) admin workspace

Use the Agent Client Collector (ACC) admin workspace to monitor the status of your agents.

Before you begin

Role required: agent_client_collector_admin

About this task

For details on the information accessible on the ACC admin workspace, see Agent Client Collector (ACC) admin workspace information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **ITOM Infra Services Workspace**.
2. Select the **ACC agents** icon .
   The ACC agents page appears, displaying an overview of the agents in your system on the Dashboard tab.

ACC agents dashboard

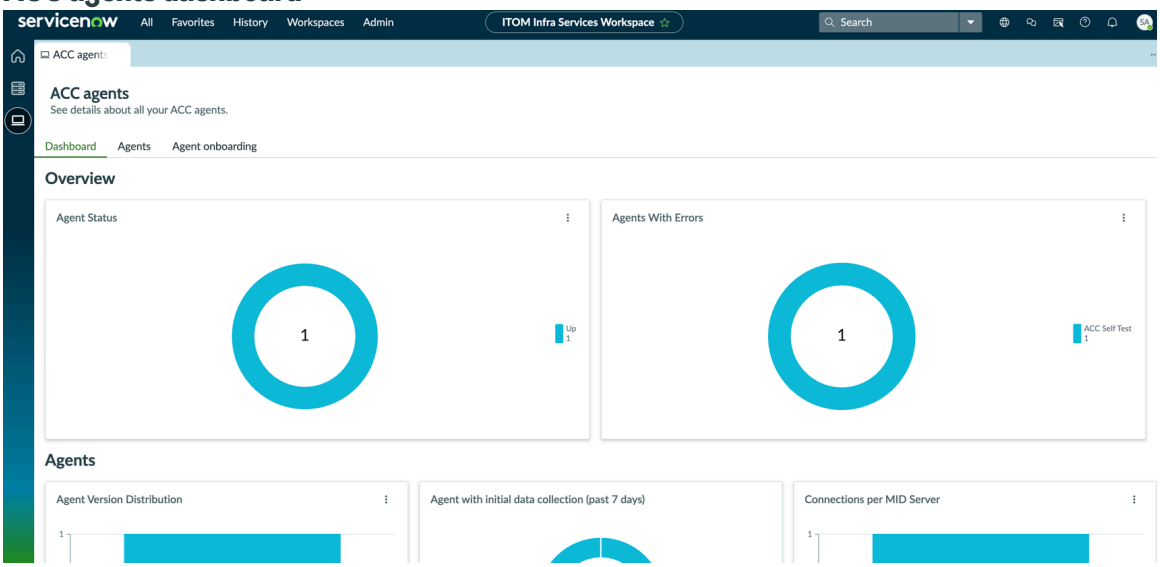

3. Select the **Agents** tab to view a list of the agents.

ACC agents page in Workspace

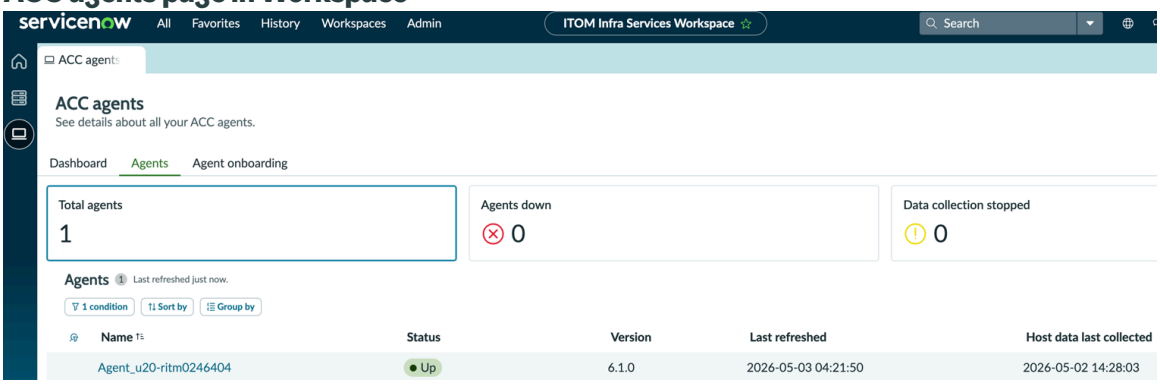

   You can filter the displayed agents by selecting one of the boxes at the top of the page. For example, select **Agents down** to filter the list by agents that are currently down.
4. Select an agent in the agents list to view details of the agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent details page

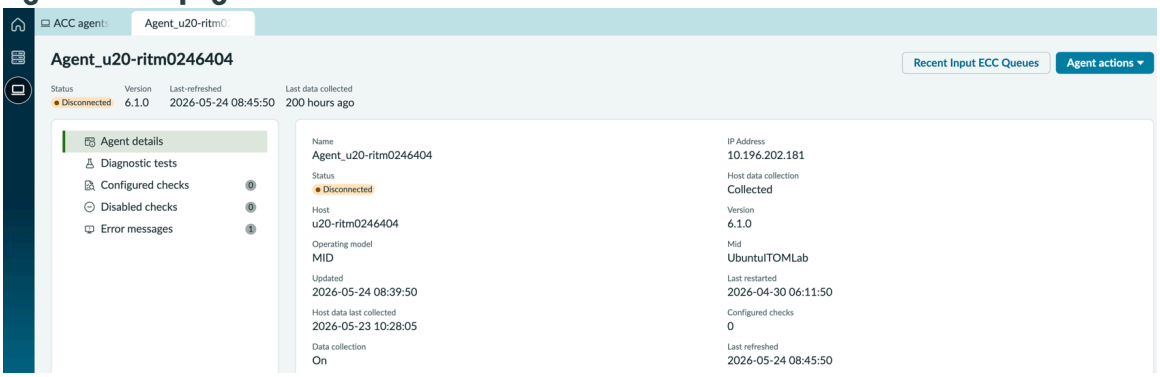

5. **Optional:** Select the **Agent actions** button and select from actions to perform on the agent. For details on the available actions, see Agent Client Collector (ACC) admin workspace agent actions.

> **Note:** The **Agent actions** button displays only when you have installed Agent Client Collector Framework version 6.0.0 or higher.

Generate an Agent Client Collector installation plan

Configure and deploy new agents in your environment, using the Agent Onboarding guide. The Agent Onboarding guide generates a customized installation plan for deploying the Agent Client Collector on your endpoints or servers.

Before you begin

Ensure that you have installed Agent Client Collector admin workspace version 1.1.0.

Role required: agent_client_collector_admin

About this task

The Agent Onboarding guide generates a customized installation plan based on your deployment type, operating system, and use case. The plan includes the commands and configuration steps required to install the Agent Client Collector on your endpoints or servers.

Procedure

1. Navigate to **Workspaces** > **ITOM Infra Services Workspace**.
2. Select the **ACC agents** icon .
3. Select the **Agent onboarding** tab.
   Page 1 of the Agent Onboarding guide opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

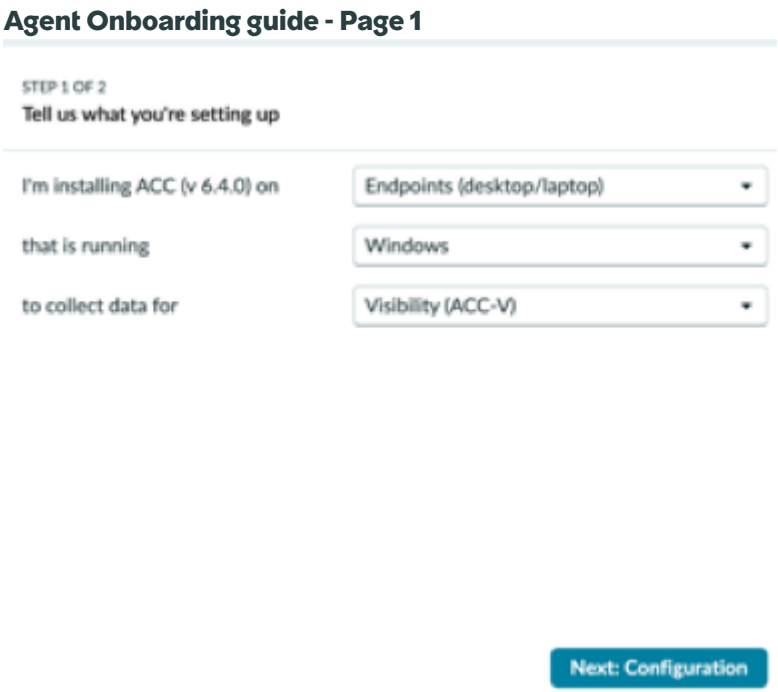

4. Configure the fields on the page.

Agent Onboarding guide fields

Field	Description
I'm installing ACC <version number> on	Place where you're installing your agent, either endpoints or servers.
that is running	Platform on which the agents are running.
to collect data for	Component for which the agent is collecting data.

5. Select **Next: Configuration**.
   Page 2 of the Agent Onboarding guide opens, displaying default configuration settings for the component for which you're collecting data.
6. Activate the toggle switch for a setting to enable customizing the value.
   For details on the available settings, see Agent onboarding configuration settings.
7. Select the **Next: Installation plan** button to view the agent installation plan, based on your configured settings.
   For details on the displayed agent installation plan steps and the commands required to be carried out for each, see Agent installation plan steps.

   The settings configured for your agent appear in the Agent configuration panel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent configuration panel

8. **Optional:** Select **Export plan** to download and print the installation plan.

Result

The installation plan is generated and ready to use. Each step in the plan includes the commands or configuration actions required to deploy the Agent Client Collector on your endpoints or servers.

What to do next

After completing the installation plan steps, verify that the agent is running and connected to the ServiceNow instance, on the ACC agents page.

Run diagnostics on an Agent Client Collector agent

Run automated self-tests on an Agent Client Collector (ACC) agent from the ITOM Infra Services Workspace to identify and address agent issues. View agent errors and invoke the relevant remediation steps.

Before you begin

- Agent Client Collector Framework version 6.2.0 or higher must be installed.
- The agent status must be Up to start a diagnostic run.
- Role required: agent_client_collector_admin

About this task

Agent diagnostics runs automated self-tests on an ACC agent directly from the agent record. After a run completes, results appear on the **All Results**, **Full Diagnostics**, and **Partial Diagnostics** tabs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **ITOM Infra Services Workspace**.
2. Select **Manage ACC agents** to view the list of ACC agents.
3. Select an agent on the **Agents** tab to open the agent record.
4. In the agent feature tree, select **Diagnostic tests**.

Agent diagnostic tests

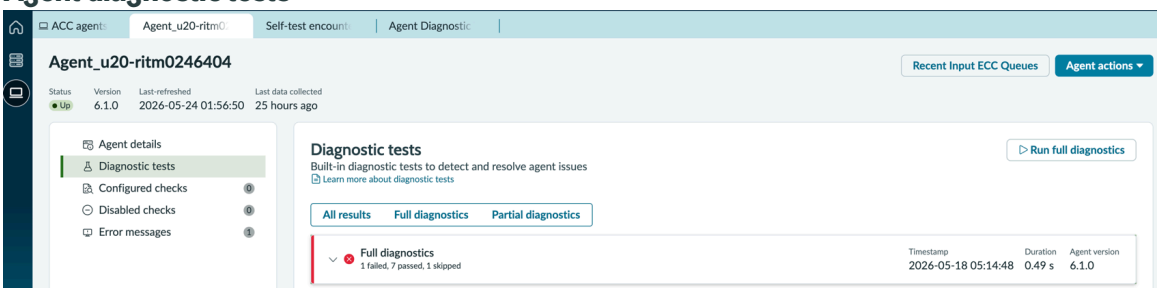

5. Select **Run full diagnostics**.
   Results populate across the All Results, Full Diagnostics, and Partial Diagnostics tabs after the run completes. Runs can take up to 60 seconds.

   If results don't appear within 60 seconds, the run timed out. Refresh the page and check for any errors.

   If the MID Server associated with the agent is not reachable, diagnostic runs cannot complete. Verify that the MID Server is running before retrying.

6. **Optional:** Select an individual test to open its detail page.
   The detail page shows the following information:

Test detail page

Area	Description
Summary fields	Severity, last result, last run timestamp, and the agent the test ran on.
Result History tab	A chronological list of past runs with status and expandable log output. Log lines can be copied to the clipboard.
Troubleshooting Guide tab	Step-by-step remediation suggestions. (Optional) Appears only when the test has a known error code with linked guidance.

7. **Optional:** Select **Run this test** to run the single test from the page.
8. **Optional:** To rerun only the failed tests, select **Rerun failed tests**.
   This button appears only when the run produced failures. Selecting it re-executes only the failed tests and leaves passing tests unchanged.

Learning Enhanced Automation Platform (LEAP)

LEAP application uses AI-driven insights and automation to help manage IT operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore Learn about LEAP benefits, features, and capabilities.	Configure Install and configure LEAP.
Use Create problem records and playbooks using LEAP.	Reference Get detailed information on important skill configuration during LEAP installation.

Supported LLMs for ServiceNow Otto for ITOM

You can use Azure OpenAI, Google Gemini or Anthropic Claude on AWS as the AI model provider for ServiceNow Otto for ITOM skills and AI agents. You can set skill-level preferences in the AI Admin Hub console. For more information, see Large language models on the ServiceNow AI Platform.

Important:

- Not all model providers are available for customers with in-country SKUs, and some AI products/features are currently unavailable for in-country customers. For more information, see the KB1584492 article in the Now Support Knowledge Base. Be sure to check for model provider availability updates in future releases.
- Some AI products/features are currently unavailable for customers in the FedRAMP, NSC DOD IL5, or Australia IRAP-Protected data centers, self-hosted customers, or in other restricted environments. For more information, see the KB0743854 article in the Now Support Knowledge Base. Be sure to check for availability updates in future releases.
- Some AI products/features are currently available only for customers in some regions. Be sure to check for availability updates in future releases.
- Some AI products and skills are not available in Regulated Markets. For more information, see KB2593939: Regulated Markets AI Products/Skills Not Available. Be sure to check for availability updates in future releases.

Troubleshoot and get help

- IT Operations Management product on the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI limitations

This application uses artificial intelligence (AI) and machine learning, which are rapidly evolving fields of study that generate predictions based on patterns in data. As a result, this application may not always produce accurate, complete, or appropriate information. Furthermore, there is no guarantee that this application has been fully trained or tested for your use case. To mitigate these issues, it is your responsibility to test and evaluate your use of this application for accuracy, harm, and appropriateness for your use case, employ human oversight of output, and refrain from relying solely on AI-generated outputs for decision-making purposes. This is especially important if you choose to deploy this application in areas with consequential impacts such as healthcare, finance, legal, employment, security, or infrastructure. You agree to abide by ServiceNow's AI Acceptable Use Policy, which may be updated by ServiceNow.

Data processing

This application requires data to be transferred from ServiceNow customers' individual instances to a centralized ServiceNow environment, which may be located in a different data center region from the one where your instance is, and potentially to a third-party cloud provider, such as Microsoft Azure. This data is handled per ServiceNow's internal policies and procedures, including our policies available through our CORE Compliance Portal.

Data collection

ServiceNow collects and uses the inputs, outputs, and edits to outputs of this application to develop and improve ServiceNow technologies including ServiceNow models and AI products. In addition, this application will collect information about scripts (and associated script records) in which ServiceNow Otto for ITOM for code generation is called. Customers can opt out of future data collection at any time, as described in the Now Assist Opt-Out page.

For more information, see the Now Assist documentation.

Exploring LEAP

LEAP categorizes similar incidents into groups and uses AI to generate resolution steps, problem records, AI-enhanced knowledge base articles, and playbooks.

LEAP overview

LEAP helps you manage incidents and operational tasks efficiently. It uses AI to gather information from existing incidents and generates resolutions that can be automated. The LEAP automation provides tools to manage incidents early, speed up resolution, use prebuilt automation, and follow clear steps to handle issues. These tools help support smooth operations.

Note:

- The plugin name has been updated from AIOps Learning Enhanced Automation Playbook to Learning Enhanced Automation Platform (LEAP) to reflect its expanded capabilities and strategic direction.
- Now Assist is renamed to ServiceNow Otto.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

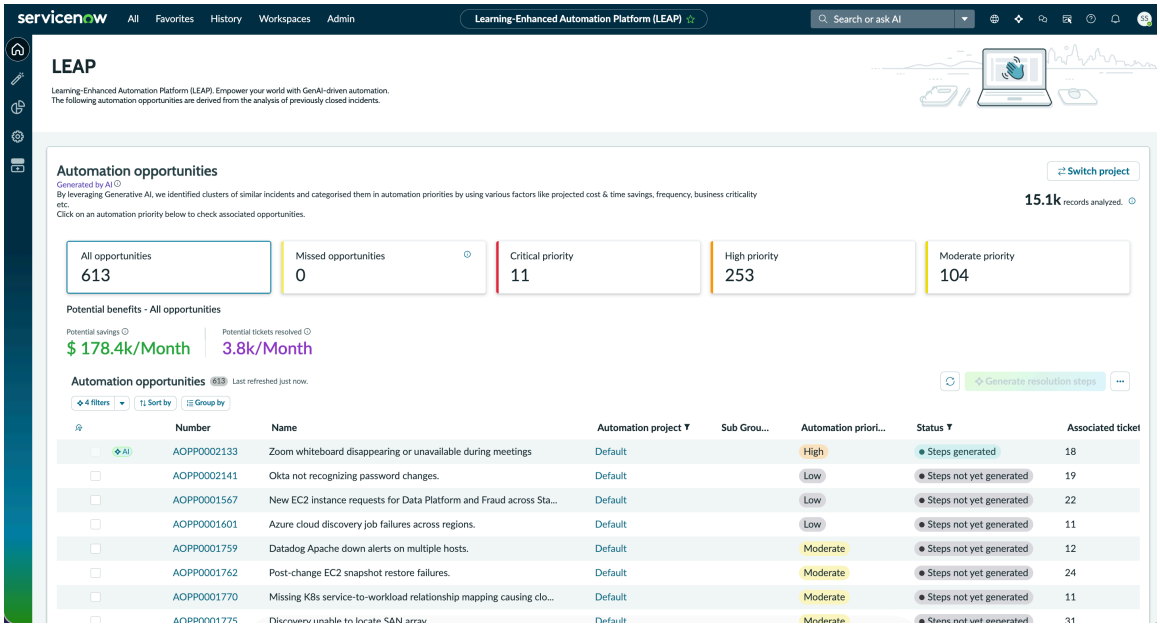

The landing page displays the number of records analyzed in the Automation opportunities section. The tooltip provides details about the duration considered for record analysis.

When the LEAP homepage loads, it displays a banner that you can use to connect Ansible Automation Platform. The banner makes the Ansible automation integration discoverable so that you can connect the automation tool and map playbooks to resolution steps. Select **Connect** to open the LEAP settings page and configure the connection. To dismiss the banner, select **Remind me later**. The banner reappears when you refresh the homepage until a connection is configured. After you connect Ansible, the banner no longer appears on the homepage.

Grouping of automation opportunities in LEAP

GAF (Group Action Framework) is the clustering engine that LEAP uses to group incidents into Automation Opportunities (AOs). GAF analyzes incident records using the short description field and applies ML-powered clustering to group incidents that share similar problem patterns. LEAP surfaces recurring issue types as discrete automation opportunities (AOs) rather than treating every incident in isolation.

GAF runs on a scheduled job, by default monthly. The schedule is configured during LEAP skill activation through the installer. On the first run, GAF processes up to six months of historical incident data using the filter set in the installer.

Before clustering begins, GAF selects the top N incidents from the eligible incident pool, where N is configurable and can be set to a maximum of 50. This selection confirms that clustering focuses on the most relevant and recent incidents within the defined scope. You can configure the fields GAF considers during analysis such as tables and columns in the LEAP settings. You can also modify the schedule run frequency in the LEAP settings.

LEAP users

User roles

Role	Description
LEAP admin	You have full control over LEAP capabilities, including the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

User roles (continued)

Role	Description
	• Turn features on and off • Execute scheduled jobs • Manage LEAP tables and the LEAP Workspace • Generate and update resolution steps • Generate AI-enhanced knowledge base articles, problem records, and playbooks from automation opportunities (AO)
LEAP viewer	You can access and view data related to LEAP, including tables, the workspace, and problems originating from AOs.
LEAP agent	You can access the LEAP menu from the Service Operations Workspace (SOW) and trigger LEAP executions directly from within SOW.

These roles work together to create a streamlined and secure approach for managing automation, resolving incidents, and sharing knowledge across teams.

Personas in LEAP

LEAP supports different personas who can have different roles assigned to enhance IT Operations Management.

LEAP personas

Persona	Description	Responsibility
Automation architect	An automation architect is a senior technical expert who designs and develops automation solutions, scaling them when required. They're a bridge between IT operations, business needs, and automation strategy. The automation architect streamlines manual, repetitive, and error-prone tasks.	• Identifies high-impact areas for automation • Develops automation scripts, integrate AI-driven insights, and maintain scalable automation frameworks • Extracts patterns from historical ticket data to optimize future resolutions • Refines automation solutions based on feedback and evolving business needs. • Collaborates with different business units to align automation with enterprise objectives.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

LEAP personas (continued)

Persona	Description	Responsibility
IT operator	An IT operator includes all L1 and L2 incident handlers who support day-to-day IT operations.	• Uses AI-generated playbooks to resolve incidents • Accesses the Knowledge Base
Buyer, Business goal owner	A buyer or a business goal owner gains strategic and operational advantages using LEAP.	• Uses LEAP to improve IT operations • Optimizes resource allocation • Automates recurring issue resolution

An automation architect can use LEAP to gather feedback and refine solutions. In a similar manner, IT operators will use LEAP to detect recurring issues and LEAP displays suggestions for preventive automation. This enables faster resolution and operators can resolve more incidents independently leading to improved efficiency and service quality.

LEAP benefits

LEAP benefits

Benefit	Feature
Promotes automation culture	Interpret data and automate records
Drives incident resolution	Measure and enhance performance
Targets Outcomes for L1 Operators	Interpret data and automate records
Improves MTTR	Measure and enhance performance
Optimize Resource Allocation	Identify and prioritize high impact areas
Provides cost predictability	Fixed pricing model for incident analysis operations

What to explore next

To learn more about configuring and using LEAP, see:

- Configuring LEAP
- Using LEAP
- LEAP reference

LEAP features

LEAP includes features such as grouping incidents, creating problem records, and generating playbooks.

Identify and prioritize high impact areas

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

LEAP uses data-driven analytics to help users identify critical issues and prioritize high-impact automation opportunities. This ensures that resources are prioritized for the areas that deliver the highest operational efficiency and business value.

Measure and enhance performance

LEAP helps improve Mean Time to Repair (MTTR) and operational efficiency by using advanced analytics and AI-driven insights to continuously monitor and optimize IT processes.

Interpret data and automate records

LEAP automatically generates problem records and executes automated playbooks by analyzing historical data patterns, enabling faster root cause identification and reducing manual intervention across complex IT environments.

Generate AI-enhanced knowledge base articles

LEAP uses AI to transform resolution steps into complete, structured knowledge base articles. Generated articles include all required sections and auto-populated metadata. Articles are published to the Knowledge Center through the standard approval workflow.

Savings metrics

LEAP tracks two types of savings metrics, projected and actual, across all automation outcome types including ServiceNow playbooks, Ansible playbooks, KB articles, and problem records.

LEAP savings metrics overview

When you use LEAP to identify and act on automation opportunities, the platform tracks savings at two stages of the workflow:

- Projected savings — displayed on the LEAP home page before you automate, based on historical incident data.
- Actual savings — recorded after automation executes, across all active outcome types including ServiceNow playbooks and Ansible playbook executions.

Both metrics are important. Projected savings help you prioritize which opportunities to pursue. Actual savings help you measure the real return on your automation investment across all outcome types.

Projected savings

Projected savings represent the estimated value that could be realized if an automation opportunity is adopted. These figures appear on the LEAP home page and on individual automation opportunity cards.

LEAP calculates projected savings using the following inputs from your historical incident data:

- Historical incident volume (number of incidents per month)
- Average manual resolution time (mean time to resolve, or MTTR)
- Estimated automation rate
- Cost per hour of IT labor

The formula is:

```
Projected savings = incident volume × Avg. manual resolution
 time × Estimated automation rate × Cost per hour
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Projected savings update dynamically as new incidents are ingested and clustered. Because they are based on estimates and historical patterns, they represent a potential value, not a guaranteed outcome.

Note: The savings number displayed on the LEAP home page is a projected value. It indicates the potential savings if all identified automation opportunities were fully adopted.

Actual savings

Actual savings represent the realized value recorded when automation executes and resolves incidents. These figures reflect real, measured outcomes and accumulate over time. LEAP tracks actual savings separately for each outcome type.

LEAP calculates actual savings per outcome type using the following inputs:

- Number of incidents resolved by the playbook or automation
- Time saved per automated resolution (manual time minus automated time)
- Cost per hour of IT labor

The formula is:

```
Actual savings = Sum of (Time saved per automated resolution ×
 Cost per hour) across all completed executions for an outcome
 type
```

Actual savings are visible in the LEAP value dashboard. The **Overview** tab shows total savings from playbook executions in the **Savings by playbooks** section. The **Ansible playbooks** tab shows cost savings and agent-hours saved specifically from Ansible-driven resolutions.

Projected and actual savings comparison

The following table summarizes the key differences between the two savings metrics.

Projected and actual savings

	Projected savings	Actual savings
Definition	Estimated value based on historical incident data, displayed on the LEAP home page before automation executes.	Realized value recorded when automation executes and resolves incidents, tracked separately per outcome type (ServiceNow playbooks, Ansible playbooks).
Data source	Historical incident volume, MTTR, and manual effort estimates.	Playbook and Ansible execution logs, time saved per automated resolution, and resolution outcomes per outcome type.
Where displayed	LEAP home page and automation opportunity cards.	**Savings by playbooks** section on the value dashboard **Overview** tab; **Cost savings via Ansible** and **Total agent-hours saved via Ansible** on the **Ansible playbooks** tab.
When calculated	Continuously, as new incidents are ingested and clustered.	Per resolution, upon successful execution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Projected and actual savings (continued)

	Projected savings	Actual savings
Formula	Incident volume × Avg. manual resolution time × Estimated automation rate × Cost per hour.	Sum of (Time saved per automated resolution × Cost per hour) across completed executions for an outcome type.

Use case: Password reset automation

This walkthrough follows a realistic scenario to show how projected and actual savings are calculated at each stage of the LEAP workflow.

The scenario: Your IT service desk receives approximately 1,200 password reset requests every month. Each request currently takes an agent about 15 minutes to resolve manually. Your organization values IT labor at $50 per hour.

Stage 1: LEAP analyzes your incident history.

LEAP ingests your historical incident data and identifies a cluster of 1,200 recurring password reset incidents per month. It recognizes this as an automation opportunity and displays it on the home page with a priority level based on the automation priority matrix.

Stage 2: LEAP calculates projected savings.

Using historical data, LEAP estimates how much time and money your team could save if this incident type were automated.

The calculation: 1,200 incidents × 15 minutes each = 18,000 minutes = 300 hours per month. At $50 per hour, that equals **$15,000 per month** or **$180,000 per year** in projected savings.

The LEAP home page displays "Projected savings: $15,000/month" on the automation opportunity card. This is an estimate based on historical patterns.

Stage 3: You create and activate a playbook.

Based on the automation opportunity, LEAP generates resolution steps for password resets. You review them, create a playbook, and activate it. From this point forward, incoming password reset incidents that match the cluster are handled by the playbook automatically.

Stage 4: Playbook runs and LEAP records actual savings.

Over the next month, the playbook executes against real incidents. LEAP measures the actual time saved for each automated resolution.

- Coverage: The playbook successfully automates 1,040 of the 1,200 incidents (87%). The remaining 160 incidents require manual handling due to edge cases such as locked accounts, expired tokens, or multi-factor authentication issues.
- Time per automated incident: Each automated resolution takes about 2 minutes (versus 15 minutes manually), saving 13 minutes per incident.
- Actual calculation: 1,040 automated incidents × 13 minutes saved = 13,520 minutes = 225 hours. At $50 per hour = **$11,267 actual savings** for the month.

Stage 5: Compare projected and actual savings.

After one month of operation, you can compare the two numbers side by side to understand the real impact of automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scenario summary: Projected and actual

Metric	Projected (before automation)	Actual (after automation)
Time saved per incident	15 minutes (estimated)	13 minutes (measured)
Monthly time savings	300 hours	225 hours
Monthly cost savings	$15,000	$11,267
Annual cost savings	$180,000	$135,204

Why projected and actual savings differ

A gap between projected and actual savings is normal and expected. The most common reasons include:

- Automation coverage: Not every incident matches the playbook. Edge cases, exceptions, and unusual incident variations still require manual handling.
- Processing overhead: Automated resolutions are fast but not instant. The playbook takes some time per incident rather than zero.
- Ramp-up period: Early in deployment, coverage is typically lower. As playbooks are refined and edge cases are addressed, actual savings trend closer to projected values over time.

Tip: Over time, as playbooks are refined and coverage expands, actual savings typically trend closer to projected values. Regularly review your resolution reports to track this improvement.

Projected savings tell you the potential value of an automation opportunity and help you prioritize which opportunities to pursue. Actual savings tell you the real value being delivered and help you measure return on investment. Both metrics are important, projected savings guide your decisions, and actual savings validate them.

LEAP prioritization logic and cost savings calculation

LEAP uses built-in prioritization logic to identify the highest-impact automation opportunities. This logic helps you focus on the most valuable automation tasks. The calculations provide an estimate of potential cost and time savings. It's best to validate these values in a non-production environment before implementing them in production.

The automation priority is calculated by combining the following key factors:

- Business impact
- Projected cost/time savings

The priority for an automation opportunity is determined using the following chart.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Automation opportunity priority

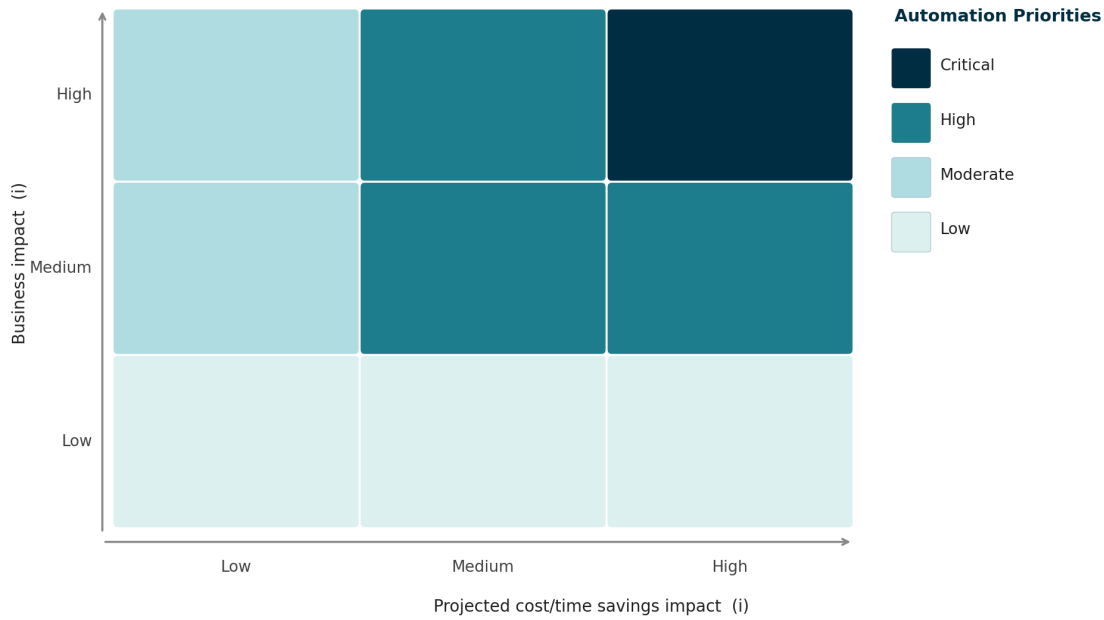

The key factors are evaluated for each group of related incidents, called#automation opportunities, and the results help teams focus on high-value areas.

Business impact

Each automation opportunity is a cluster of tickets. Every ticket is scored based on the following factors:

- Business criticality
- Major incident status
- Impact level

These factors are given values. Example values for#criticality:

- Most Critical = 100%
- Somewhat Critical = 75%
- Less Critical = 50%
- Not Critical = 25%

Properties table

Key	Default value	Description
ohf_p4	0.8	Overhead factor for P4 records.
ohf_p5	0.4	Overhead factor for P5 records.
ohf_p1	1.8	Overhead factor for P1 records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties table (continued)

Key	Default value	Description
ohf_p2	1.4	Overhead factor for P2 records.
cost_per_entry	20	Cost in dollars associated with adding a work note entry to a record.
ohf_p3	1.1	Overhead factor for P3 records.
time_per_entry	0.5	Time in hours associated with adding a work note entry to a record.
first_run_group_limit	0	Total number of groups for which resolution steps are generated in the initial run.
gaf_mapping_min_threshold	0.2	Minimum ratio of incidents pointing to an active group that is considered for remapping after clustering.

Note:

- You can modify the values in the LEAP settings page. However, the calculations will differ when you change the values.
- The values in the LEAP settings page should be modified before activating LEAP.
- If you modify the values in the LEAP settings page, you need to rerun the calculation. For this, you must deactivate and reactivate the LEAP installer skill.

Each factor is weighted. For example, say that Criticality = 0.8, Major Incident = 0.2, and Impact = 0.4.

```
Ticket Score = (Criticality × 0.8 + Major Incident × 0.2 +
 Impact × 0.4) ÷ 3
```

The average of all ticket scores in the group becomes the Business Impact Score. Based on this score, automation opportunities are labeled as high impact (the top 33%), medium impact (next 33%), and low impact (remaining 33%).

Projected cost or time savings impact

LEAP estimates how much time and cost can be saved by automating a group of incidents with historical ticket data.

The key inputs for cost and time savings are the following:

- Frequency of Incidents (FOI): Average number of similar incidents per month
- Average Manual Work Notes Count (AWNC)
- Overhead Factor (OHF): varies by incident priority based on the priorities table

LEAP uses the following formulas to calculate the projected cost or time savings:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `Projected Cost Savings = Σ (FOI × OHF) × AWNC × Cost per Work Note`
- `Projected Time Savings = Σ (FOI × OHF) × AWNC × Agent Time per Entry`

Each ticket's projected savings is calculated. The average for the group becomes the Projected Cost/Time Savings Score, which is also categorized as High, Medium, or Low. The automation opportunities are sorted in descending order based on the projected cost or time savings.

Determining overall priority

After calculating the values for the business impact and the projected cost/time savings, LEAP assigns a priority level to each automation opportunity using this matrix:

Automation priority matrix

Business impact	Projected cost/time savings	Automation priority
High	High	Very High
Medium	High	High
High	Medium	High
Medium	Medium	Medium
Medium	Low	Low
Low	Medium	Low
Low	Low	Very Low

These prioritized automation opportunities are displayed on the LEAP home page.

Tracking the actual cost and time savings

After automation is implemented, LEAP tracks actual cost and time savings and displays them on the LEAP Value Dashboard, helping you measure impact and make decisions.

Key metrics:

- Additional Work Notes (AWN): Manual entries after playbook execution
- Overhead Factor (OHF) by Priority: For example, P1 = 1.8x

Formulas:

- `Saved Work Notes (SWNs) = AWNC – AWN`
- `AWN = SWN × OHF`
- `Actual Cost Savings = AWN × Cost per Work Note`
- `Actual Time Savings = AWN × Agent Time per Entry`

Example projected cost and time savings calculations

A cluster has 20 AWNC and is mapped to a P2 incident with 6 AWN. The cost per work note is $20, and the agent time per entry is 0.5 hours.

According to the formulas, the SWN count is `(20 - 6) = 14`. According to the properties tables, the OHF = 1.4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

`SWN x OHF = AWN` so `14 x 1.4 = 19.6`

This value can be rounded off to 20 AWN. Each AWN entry represents $20 of savings.

The actual cost savings post automation = `AWN x Cost per Work Note` = `20 x 20 = $400`.

So, if you use automation, then the projected cost savings can be up to $400. Thus, the calculation provides actionable insights for ongoing process improvement and resource allocation.

AI agent

Enhance IT operations with AI-driven, autonomous artifact creation such as problem records, knowledge base articles, and playbooks using LEAP AI agent.

Important:

Generative AI may produce inaccurate or incomplete information. Always validate AI-generated content before you publish or act on it.

Generate LEAP artifacts

AI agent	AI agent role
LEAP AI agent	Uses the automation opportunities created by LEAP analysis, and creates artifacts (problem records, knowledge base articles, or playbooks) requested by users.

Important: This AI agent is turned on by default. For more information, see AI agents, skills, and agentic workflows on by default.

What the LEAP AI agent does

The LEAP AI agent autonomously scans active automation opportunities and acts on the ones that have resolution steps. IT operations teams spend less time triggering automation manually. The agent provides the following capabilities:

Create knowledge base articles

The agent is triggered when resolution steps for automation opportunities are generated successfully. It documents and generates a draft knowledge base article that includes the automation opportunity title, resolution steps, and related incident context. You review the draft and publish it. Each published article remains traceable to its source automation opportunity. The agent doesn't create a knowledge base for an automation opportunity that is already linked to an open KB.

Create problem records

The agent is triggered when resolution steps for automation opportunities are generated successfully and automatically creates a problem record. Each problem record is populated with the related incident count, the automation opportunity detail, and the associated resolution notes, which mirrors the manual **Create Problem** action. The agent doesn't create a problem record for an automation opportunity that is already linked to an open problem. Each problem record remains traceable to its source automation opportunity.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Eligibility thresholds

An automation opportunity is eligible for knowledge base article and problem record creation when it has a minimum of 15 incidents or a severity of Critical or High. You can change the incident count threshold (default: 15) and the severity threshold (default: Critical and High) on the LEAP Settings page. For the list of settings, see LEAP settings fields.

Identify AI-created artifacts

Artifacts that the LEAP AI agent creates are marked with an AI indicator on the home page and the automation opportunity details page, so that you can tell AI-created artifacts apart from those created manually. For more information, see LEAP AI indicators.

Activate or deactivate the LEAP AI agent

The LEAP AI agent is active by default. To change the activation status, navigate to **All** > **AI Agent Studio** > **Overview**, select the LEAP AI agent, then go to **Agent** > **Select channels and status** and toggle the **activation status** control.

- Active: The agent scans automation opportunities and creates artifacts automatically.
- Inactive: The agent stops scanning automation opportunities and no longer creates artifacts.

Automation projects

Use Automation projects to analyze multiple sets of incident data independently, so different teams or business units can derive separate insights and automation opportunities.

An Automation project defines the scope and configuration for LEAP analysis. Each project runs independently and can target specific assignment groups, time periods, or filtered sets of incidents. By creating multiple Automation projects, you can analyze different categories of incidents separately and generate automation opportunities tailored to each context.

Benefits of Automation projects

Automation projects provide the following benefits:

- Team-specific insights: Different teams or business units can work on separate projects concurrently without affecting each other's analysis, metrics, or automation opportunities.
- Analyze multiple data sets independently: Create separate projects to analyze different groups of incidents based on assignment groups, time periods, or other criteria. For example, analyze the last 6 months of Human Resources incidents in one project while creating a focused project for last month's spike in cases.
- Flexible taxonomy configuration: Select one or more taxonomies per project to group incidents according to your team's classification needs. Multiple projects can use the same taxonomy when analyzing different subsets of data.
- Real-time dashboard views: Switch between projects to view project-specific metrics, reports, and automation opportunities in the dashboard.
- Configurable analysis schedule: Schedule analysis jobs to run at specified frequencies for each project independently.

Each Automation project has its own independent configuration: incident, filter, grouping column, taxonomy selection, job frequency, and settings. Projects run without affecting each other.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When to use Automation projects

The following examples illustrate common scenarios where creating multiple Automation projects provides more relevant and actionable results than a single shared project.

Per-team automation

A large IT organization wants Network team incidents grouped and automated separately from Database team incidents. Because issue patterns, root causes, and resolution steps differ significantly between teams, each team gets its own project instead of a single mixed pool of recommendations.

Per-region automation

A global company wants automation tuned separately for APAC and EMEA support teams. Ticket volume, categories, and SLAs vary by region, so separate projects confirm that grouping and recommendations reflect each region's specific incident patterns.

Per-service-tier automation

A customer wants to separate automation for critical Tier-1 services, such as banking systems and core infrastructure, from general Tier-2 and Tier-3 services. Keeping high-priority incidents in a dedicated project prevents lower-priority noise from diluting the grouping and resolution logic for business-critical systems.

Controlled pilot before a wider rollout

A customer wants to try LEAP automation on one team's data first before expanding to the rest of the organization. Creating a separate project for the pilot team keeps the trial isolated from the default project that is already running in production.

Business unit isolation

A company with multiple business units, each using ServiceNow with different assignment groups, wants each unit to see automation results relevant only to their own incidents. Separate projects scoped by assignment group ensure that each business unit's recommendations are not mixed with those of other units.

Testing and sandbox configuration

An administrator wants to test grouping or filter changes safely before applying them to the production setup. A separate project scoped to a small filtered set of records acts as a sandbox, leaving the default project and its running jobs unaffected.

Default Automation project

The Default Automation project is shipped with the LEAP. The Default project represents the existing pre-multi-project setup and has no taxonomy configuration and is read-only. You can view its configuration but can't edit it from the Automation project details page. To change the default configuration, use the AI Admin Workspace.

Note:

- The initial setup experience for LEAP remains unchanged. When you first activate LEAP, you must still configure the default project through the AI Admin Workspace. Automation projects is an additional capability that is available after the default project is set up.
- Skill activation occurs in the AI Admin workspace.

Automation projects are an opt-in capability. Existing deployments continue to work without modification. You don't have to migrate your current configuration to use Automation projects. Adding a new project does not change the Default Automation project or any currently running analysis jobs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Automation project lifecycle

Managing Automation projects involves creating new projects, running analysis jobs, and switching between projects to view different automation opportunities. The following workflow describes the complete lifecycle:

1. Create the project: Define the project name, description, table, filter, grouping column, topic column, taxonomy/assignment group, and job frequency. The project is created in a **ready (editable)** state and can be viewed directly on the projects details screen or the homepage. No analysis runs at this point.
2. Start the grouping job: After the project configuration is complete, start the analysis job. Starting the job clones the required skills in the back end and activates the scheduled job. Once analysis begins, the project configuration can't be edited until the job completes or the job status is ready or failed.
3. Review automation opportunities: After the analysis completes, review automation opportunities, create resolution steps, and measure impact. The day-to-day experience is the same as the standard LEAP workflow.
4. Switch between projects: Switch between projects to view different sets of automation opportunities. Switching between projects updates data on every page and the LEAP value dashboard to reflect the selected project's data set.
5. Edit project configurations: Edit project configurations as needed. You can only edit a project when its job status is ready or failed.

Configuration constraints

When configuring an Automation project, be aware of the following constraints:

- The taxonomy field is multi-select. You can select one or more taxonomies to define which incident groups the project analyzes.
- Multiple projects can use the same taxonomy. When an incident matches multiple projects, playbooks from all matching projects are displayed.
- Filter queries are limited to 300 characters. Filters that return zero records are marked as an error before the project is saved.
- This release supports only the incident table.

Taxonomy filtering

Selected taxonomies dictate the filter query for an Automation project. For example, if you select assignment groups AG1 and AG2, the system creates a filter where the assignment group is one of AG1 or AG2.

Prediction logic

When an incident is created or updated, LEAP determines which Automation projects run predictions based on the incident's assignment group:

- The Default Automation project always runs predictions regardless of the assignment group.
- If an incident has a specific assignment group, predictions run on the Default project plus any Automation projects that match that assignment group.

Automation opportunities

LEAP groups similar incidents into automation opportunities and uses AI to generate resolution artifacts that help reduce manual effort and repeat incidents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

An automation opportunity is a group of similar incidents that LEAP identifies through AI-driven analysis. For each group, LEAP generates resolution artifacts — such as resolution steps, knowledge base articles, problem records, and playbooks — that operators can use to resolve recurring incidents faster.

Automation opportunity lifecycle

When the Group Action Framework (GAF) process re-runs, it can identify new patterns in incident data. If remapping is successful, LEAP archives automation opportunities that contain resolution steps and transfers their artifacts to newly identified opportunities. Archived automation opportunities are hidden by default in the user interface to help you focus on actionable opportunities.

The automation opportunity details page displays banner messages that indicate the archiving and remapping status:

- Archived automation opportunities display a message indicating that artifacts were remapped to a new opportunity, with a link to the new opportunity
- New automation opportunities that received artifacts display a message indicating the source archived opportunity, with a link to the archived opportunity

The Action Insights panel on the details page also shows the relationship between archived and new automation opportunities.

Managing automation opportunities

The automation opportunity details page provides tools to review and act on identified opportunities. From this page you can:

- Review the incidents grouped into the opportunity
- Generate or regenerate resolution steps, knowledge base articles, problem records, and playbooks
- Create sub-groups for large opportunities to produce more targeted resolutions
- Track missed opportunities where automation artifacts were available but not used

Automation opportunity sub-groups

Large automation opportunities in LEAP can be broken into smaller, more manageable sub-groups to achieve more granular and accurate incident resolution.

Large automation opportunities with over 150 records tend to generate overly lengthy resolution steps that lack precision and focus. These opportunities often represent a collection of related problems and issues that should be subdivided to produce more targeted and relevant solutions.

Breaking them into smaller sub-groups lets users obtain more refined resolutions for recurring issues. The Action Insights panel on the automation opportunity details page offers recommendations to create these sub-groups for opportunities containing large record volumes.

Sub-grouping example

A large automation opportunity containing 300 password reset incidents might include a mix of VPN password resets, Active Directory account lockouts, and SSO token expiry. While these are all password-related, each requires different resolution steps. LEAP displays suggestions for breaking the automation opportunity into three sub-groups. Each group receives targeted resolution steps — for example, one sub-group with VPN-specific steps, another with AD unlock procedures, and a third with SSO token refresh instructions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Automatic reorganization

During scheduled LEAP analysis reruns, clusters are automatically reorganized based on newly observed incident patterns. LEAP preserves existing artifacts and links them to updated automation opportunities rather than overwriting them.

The following table describes how each artifact type is handled during reorganization:

Artifact behavior during reorganization

Artifact	Behavior during rerun
Resolution steps	Appended to the new matching automation opportunity. Users can keep the existing steps or select **Regenerate** to create fresh ones.
Playbooks	Remain active and are linked to the new matching automation opportunity. Existing playbook configurations and execution history are preserved.
Knowledge base articles	Remain active and retain their association with the automation opportunity.
Problem records	Remain active and are linked to the new matching automation opportunity.

Automation opportunities that no longer have any associated artifacts after reorganization are marked as deactivated. Opportunities that contain resolution steps, playbooks, or other artifacts remain active and are linked to newly identified matching automation opportunities.

For example: An automation opportunity for "email delivery failures" has resolution steps and a playbook created in the previous analysis run. When LEAP reruns the analysis and reorganizes clusters, this opportunity is matched to a new cluster. The existing resolution steps are appended to the new automation opportunity. The user can review the appended steps and decide whether to keep them or regenerate new ones based on the updated cluster data.

Benefits of sub-grouping

- Provides more focused and contextual resolution steps
- Offers more targeted solutions for similar issues
- Improves incident resolution
- Enables more effective management of large automation opportunities
- Supports operational efficiency by reusing prior resolutions where appropriate

For detailed steps, see Create sub-groups for automation opportunities.

Missed automation opportunities

Track instances where users could have benefited from generated resolution steps, knowledge base articles, or playbooks but didn't have access to them.

Missed automation opportunities occur when users working on incidents could have benefited from existing LEAP automation artifacts that aren't available.

LEAP automatically logs missed opportunities when there are no artifacts recommended in Service Operations Workspace. These artifacts include resolution steps, knowledge base articles, and playbooks that could have assisted with faster incident resolution. The system

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

helps administrators identify these gaps in automation coverage and prioritize them for future improvements.

How missed opportunities are tracked

When an incident occurs and appears in SOW, and the Recommendation section of the incident does not show any artifacts, the automation opportunity is marked as 'missed automation'.

Missed opportunities are logged in a dedicated table that captures:

- incident details
- Available automation artifacts that could have helped
- User information and timing data

Types of missed opportunities

LEAP tracks several types of missed automation opportunities:

Resolution steps

Generated step-by-step procedures that could have guided incident resolution.

Knowledge base articles

Relevant documentation and solutions that could have provided immediate answers but were not consulted during incident resolution.

Playbooks

Automated workflows that could have streamlined the resolution process.

Viewing missed opportunities

The LEAP homepage displays the missed automation opportunities card.

Select the number on the missed opportunities card to filter and view only the automation opportunities that missed automation. On the details page for each automation opportunity, you can see how many times an artifact was not available when it could have assisted with incident resolution.

Missed automation opportunity details page

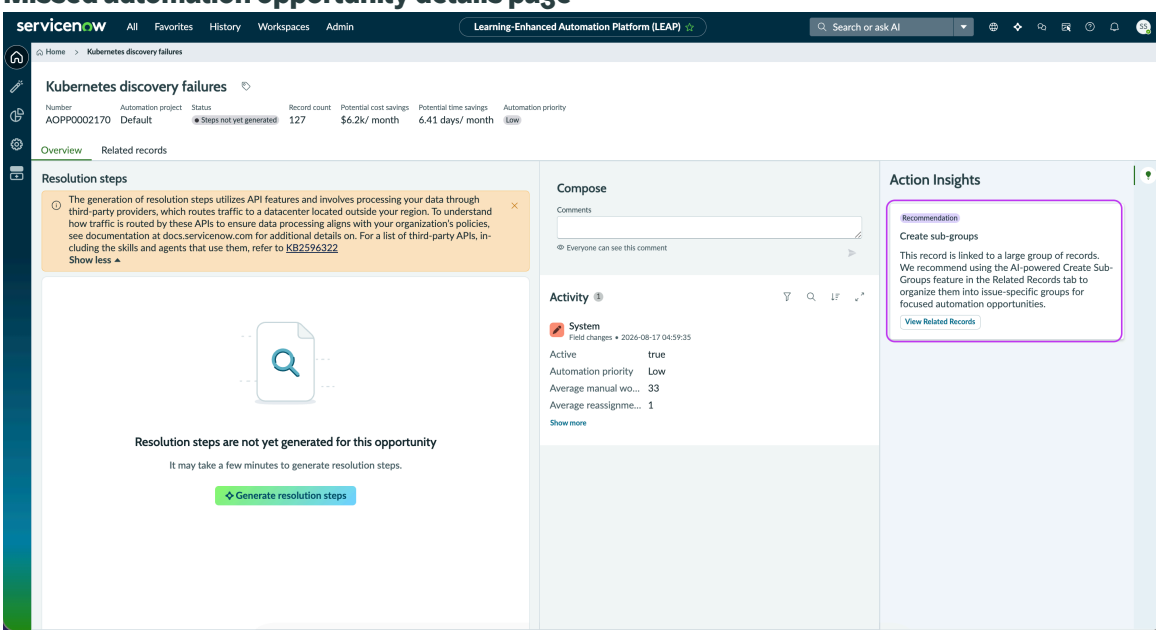

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

LEAP administrators can use missed opportunity data to identify patterns and improve automation coverage. This helps determine where additional automation artifacts might be beneficial.

Regular review of missed opportunities enables administrators to identify gaps in automation coverage for common incident types and prioritize creation of new automation artifacts based on potential impact.

Benefits of tracking missed opportunities

Tracking missed automation opportunities provides valuable insights for continuous improvement of automation strategies:

- Identifies high-impact areas where automation could provide the greatest benefit
- Helps justify investment in additional automation development
- Provides data-driven insights for improving user adoption of existing automation
- Enables measurement of automation program effectiveness and ROI
- Supports strategic planning for future automation initiatives

LEAP MCP Server

External AI clients can query LEAP automation data using Model Context Protocol (MCP) tools exposed through the ITOM MCP Server.

Overview of LEAP with MCP

LEAP exposes automation capabilities as MCP tools within the ITOM AIOps MCP Server. External AI clients such as Claude, ChatGPT, Copilot, and Gemini can invoke these tools using the Model Context Protocol (MCP), a standardized protocol for AI client integration.

MCP tools provide a governed, product-supported path for connecting external AI clients to LEAP automation workflows. Without MCP tools, the integration path is a manual process with no authentication guardrails or ServiceNow support coverage.

All MCP tool invocations use OAuth 2.0 authentication. Access to specific tools is controlled by the LEAP AI SKU tier assigned to your instance.

How MCP tools work

LEAP MCP tools are registered on the shared ITOM MCP Server (`sn_genai.itom_mcp_server`). When an external AI client connects to the server and calls `tools/list`, the LEAP tools appear alongside tools from other ITOM products. The server URL takes the form:`https://<instance>.service-now.com/sncapps/mcp-server/mcp/sn_genai_itom_mcp_server`

> ⚠ **Warning:** LEAP specific MCP tool invocations require the sn_itom_leap.leap_viewer role at minimum. Other tools on the ITOM MCP Server might still work, but LEAP specific tools fail with an unauthorized error if the user lacks this role: `User is missing required role: sn_itom_leap.leap_viewer`.

The following describes the invocation flow:

1. The external AI client authenticates to the ITOM MCP Server using an OAuth 2.0 token.
2. The client calls a LEAP MCP tool with the required input parameters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. The MCP Server routes the request to the underlying LEAP API.
4. The server returns a structured JSON response to the client.

Tool categories

LEAP MCP tools are registered as **Action** tools on the ITOM MCP Server and are organized into three categories:

LEAP MCP tools

Tools	Description
LEAP - Get App setup status	Reports LEAP setup and grouping pipeline readiness. Use this tool to check whether LEAP is configured before invoking automation opportunity tools. For example, if an AO query returns no data, `get_app_setup_status` can confirm whether setup is incomplete and return the URL of the setup page.
LEAP - Get property or setting	Retrieves specific LEAP configuration properties by key. Only properties on an explicit allow-list are accessible. For example, retrieve `missed_opportunity_threshold` (the number of missed opportunities above which an AO is flagged as high priority) or `default_resolution_steps_filter` (the encoded query that controls which AOs receive AI-generated resolution steps).
LEAP - Get automation opportunities	Provides read-only access to automation opportunity data. External clients can query for lists of AOs, retrieve full details for a specific AO, and filter results by keyword, status, or linked artifacts. For example, an AI client can ask for all AOs related to Discovery issues that have a KB article generated, or retrieve details for a specific AO number such as AOPP0000173.

Set up the MCP connection

To connect an external AI client to LEAP MCP tools, you must configure OAuth authentication on the ITOM MCP Server. A banner on the MCP Server Console confirms when OAuth setup is required. For setup instructions, see Activate the ITOM MCP Server Console.

What to explore next

- Exploring LEAP
- Automation opportunities

Ansible automation integration

The Ansible automation integration connects LEAP with Ansible Automation Platform to enable AI-driven discovery of relevant job templates and automated incident remediation from the Service Operations Workspace.

The Ansible automation integration addresses the critical operational gap of lack of automated remediation capabilities during incident resolution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key capabilities

The integration provides these core capabilities:

AI-driven automation discovery

LEAP automatically identifies relevant Ansible job templates for automation opportunities using semantic analysis.

Step-to-job mapping

The integration enables LEAP administrators to map discovered Ansible job templates to specific resolution steps through an intuitive UI.

Automated incident remediation

The integration allows incident responders to execute mapped Ansible automations directly from the Service Operations Workspace during incident resolution.

Manual step handling

The integration supports mixed workflows where some steps require manual intervention while others can be automated.

Integration architecture

The integration uses these key components:

- Ansible discovery agent: Analyzes automation opportunities and identifies candidate job templates
- Ansible execution agent: Executes mapped automation opportunities during incident remediation
- Ansible MCP server: Provides secure API access to Ansible Automation Platform
- Step-to-job mapping store: Maintains the relationship between resolution steps and Ansible job templates

Workflow overview

The integration follows this high-level workflow:

1. The Ansible discovery agent analyzes automation opportunities and discovers relevant job templates
2. LEAP administrators map discovered job templates to specific resolution steps
3. When incidents occur, LEAP predicts the automation opportunity group
4. If mappings exist, automation options are presented to incident responders
5. The Ansible execution agent launches selected automations and reports status

Benefits

The Ansible automation integration provides these benefits:

- Reduces mean time to resolution (MTTR) by automating incident remediation
- Eliminates the requirement to search for relevant job templates or automation playbooks
- Increases automation utilization across the organization
- Provides structured linkage between automation opportunities and Ansible job templates
- Enables consistent, repeatable incident resolution processes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Ansible discovery agent

The Ansible discovery agent analyzes automation opportunities and identifies relevant Ansible job templates using AI-powered semantic matching and historical execution patterns.

The Ansible discovery agent automates the process of finding relevant Ansible job templates for automation opportunities. The agent eliminates the manual effort required by automation engineers to search and identify appropriate automation resources.

Discovery process

The Ansible discovery agent follows this process to identify relevant job templates:

1. Parses the problem description of the automation opportunity for keywords and context
2. Queries Ansible Automation Platform via the Ansible MCP server to retrieve available job templates
3. Applies AI-powered matching to identify relevant job templates based on problem description similarity using semantic natural language processing
4. Presents discovered job templates to the LEAP administrator for step-to-job mapping

Matching capabilities

The agent uses these techniques to match automation opportunities with job templates:

Semantic analysis

Analyzes the meaning and context of problem descriptions to identify relevant automation patterns.

Keyword extraction

Identifies key terms and technical concepts from automation opportunity descriptions.

Integration points

The Ansible discovery agent integrates with these components:

- AI Teammate Dashboard: Entry point for automation engineers and LEAP administrators.
- Ansible MCP server: Secure API access to Ansible Automation Platform job templates.
- LEAP clustering engine: Source of automation opportunity data and problem descriptions.
- Step-to-job mapping UI: Destination for discovered job template recommendations.

AI in LEAP

AI-powered features in LEAP help administrators generate resolution steps, create problem records, publish knowledge base articles, build playbooks, and discover and execute Ansible automations.

LEAP includes several AI-powered features. All features require ServiceNow Otto for IT Operations Management (ITOM) to be installed.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI features in LEAP

AI feature	Description	Use case	Resources
AI-generated resolution steps	Uses generative AI to analyze grouped incidents and generate structured resolution steps for each automation opportunity. Resolution steps can be reviewed and modified before use.	A LEAP administrator wants to produce actionable resolution guidance for a recurring incident pattern without manually authoring steps from scratch.	Generate and modify resolution steps in LEAP
LEAP AI agent	Uses automation opportunities created by LEAP analysis to generate artifacts — problem records, AI-enhanced knowledge base articles, or playbooks — based on user requests. The agent is turned on by default.	A LEAP administrator wants to convert a recurring incident pattern into a problem record, a structured knowledge base article, or a playbook without manually authoring each artifact.	• AI agent • Create LEAP problem records • Generate LEAP knowledge base articles • Generate LEAP playbooks
Ansible discovery agent	Analyzes automation opportunities and uses AI-powered semantic matching to identify relevant Ansible job templates from Ansible Automation Platform. Presents discovered templates to the LEAP administrator for step-to-job mapping.	A LEAP administrator wants to find the most relevant Ansible job templates for an automation opportunity without manually searching the Ansible catalog.	• Ansible automation integration • Ansible discovery agent
Ansible execution agent	Automatically launches mapped Ansible job templates during incident remediation in the Service Operations Workspace. Reports execution status back to the incident responder.	An incident responder wants to trigger an Ansible automation directly from the Service Operations Workspace during active incident resolution, without switching to the Ansible platform.	• Ansible automation integration • Execute Ansible automations for incidents

Configuring LEAP

You can download and install LEAP application from the ServiceNow® Store to manage your IT operations efficiently.

Configuration overview

1. Install LEAP

   Install LEAP if you have the admin role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Setup LEAP properties

   View and modify the LEAP LEAP properties as per your requirements.

3. Activate LEAP

   Activate LEAP to enable grouping of incidents, create problem records, and generate resolution playbooks.

4. Change LEAP large language model (LLM)

   Modify the LLM model to suit your requirements.

5. Configure Ansible automation integration

   Configure the integration between LEAP and Ansible Automation Platform to enable automated discovery and execution of Ansible job templates.

Install LEAP

You can install the LEAP application using ServiceNow Otto for IT Operations Management (ITOM) (sn_itom_gen_ai) if you have the admin role. The application installs related ServiceNow® Store applications and plug-ins if they aren't already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the LEAP application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: System admin

About this task

The following items are installed with LEAP:

- Roles
- Scheduled jobs
- Tables

For more information, see Components installed with LEAP.

Procedure

1. Go to the ServiceNow Store website.
2. In the search bar, enter ServiceNow Otto for IT Operations Management (ITOM) and then select **Search**.
3. Select the ServiceNow Otto for IT Operations Management (ITOM) widget and then select **Get** to make this application available for installation on your instances.
4. Install the ServiceNow Otto for ITOM application.
   a. Navigate to **All** > **System Definition** > **Plugins**.
   b. Search for ServiceNow Otto for IT Operations Management (ITOM).
   c. Select **Install**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Setup LEAP properties

Configure LEAP properties to estimate cost and time savings calculations for your organization.

Before you begin

Role required: LEAP admin

Each knowledge base you plan to use in LEAP must have the correct Can Contribute permissions configured before you set up knowledge base routing. See Knowledge base permissions for LEAP.

About this task

If you are a first-time user of LEAP, review the default values before making changes.

Settings include three configuration areas: cost and time savings estimation, AI agent behavior, and knowledge base routing. AI agent configuration controls the conditions under which the LEAP AI agent automatically creates problem records and knowledge base articles, such as minimum incident count and severity thresholds. Knowledge base routing controls where articles are stored when created by the LEAP AI agent or manually from an automation opportunity. Selecting a default knowledge base is mandatory — you cannot save settings without one. For details on all fields, see LEAP settings fields.

Procedure

1. Select **Configure settings**.
2. On the settings page, choose one of the following options:
   - Fresh Install
     a. Select **Configure settings**.
     b. Review the default values for each property.
     c. Select **Save**.

     Note: LEAP recommends using default values initially to view and understand the cost and time savings estimates.

   - Upgrade
     a. Select **Configure settings**.
     b. Review the default values for each property.
     c. Modify values as needed.
     d. Select **Save**.

   You are redirected to the LEAP landing page. For details on each field, see LEAP settings fields
3. Activate LEAP.

Activate LEAP

Activate the LEAP skill after installing the application to assemble and categorize incidents.

Before you begin

Role required: System admin or LEAP admin

and the application scope must be LEAP.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **LEAP**.
2. Select **Admin** > **AI Admin Hub**.
3. Select **AI Skills** > **Platform** > **LEAP**.
4. On the LEAP Installer card, select **Activate**.
5. In the LEAP installer input section, enter the details for LEAP to run scans on existing records, and select **Save and continue**.
   For details on installer fields, see LEAP Installer fields
6. In the Define access section, select the role to whom you want to provide access to LEAP skills.
7. In the Select display section of *`LEAP installer`*, select the required display and select **Save and continue**.
8. In the Review and activate section of *`LEAP installer`*,review the selections, and select **Activate**.
9. Select **Return to LEAP** on the success message.
10. The grouping job starts and proceeds to group records into automation opportunities.

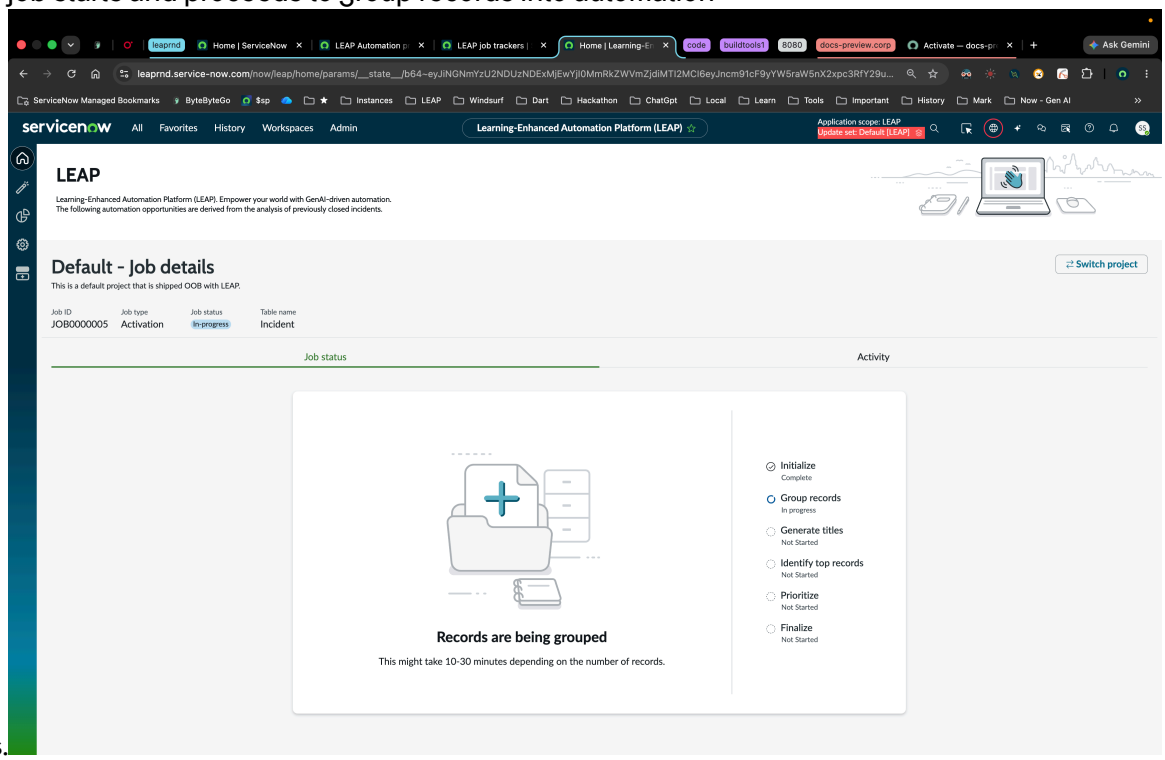

   If errors are encountered, then LEAP provides links to knowledge base articles to understand the root cause and steps to resolve issues.

Change LEAP large language model (LLM)

LEAP allows you to change the default LLM provider from the Now LLM to your required LLM.

Before you begin

Role required: admin

About this task

If you change the LLM model, the LEAP skills should be reapplied to the selected model. Because the LEAP skills aren't applied automatically to an instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Access **Admin** > **AI Admin Hub** in your workspace.
2. Select **Settings** > **Manage AI model** > **Manage model**

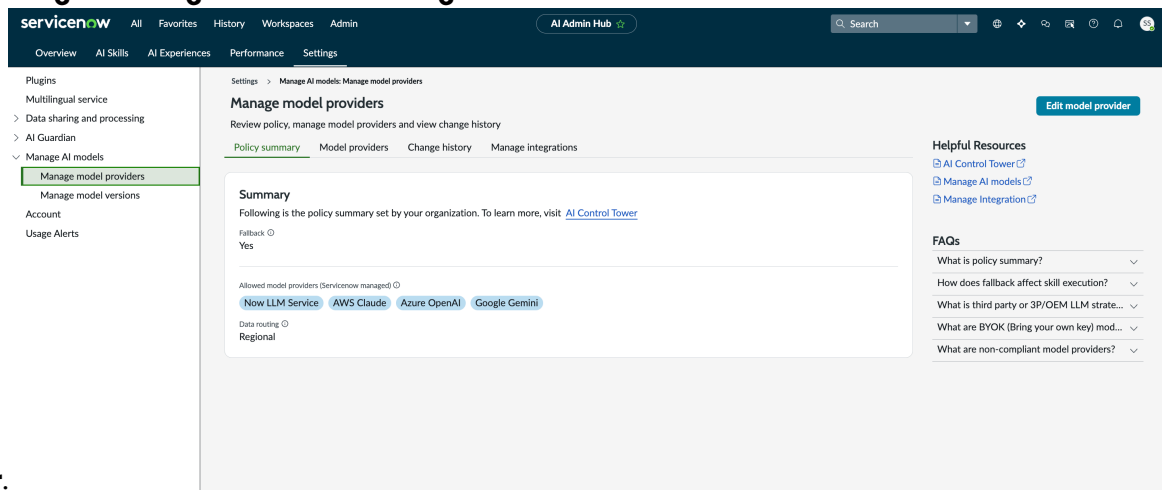

provider.

3. Select **Edit model provider**, and then select

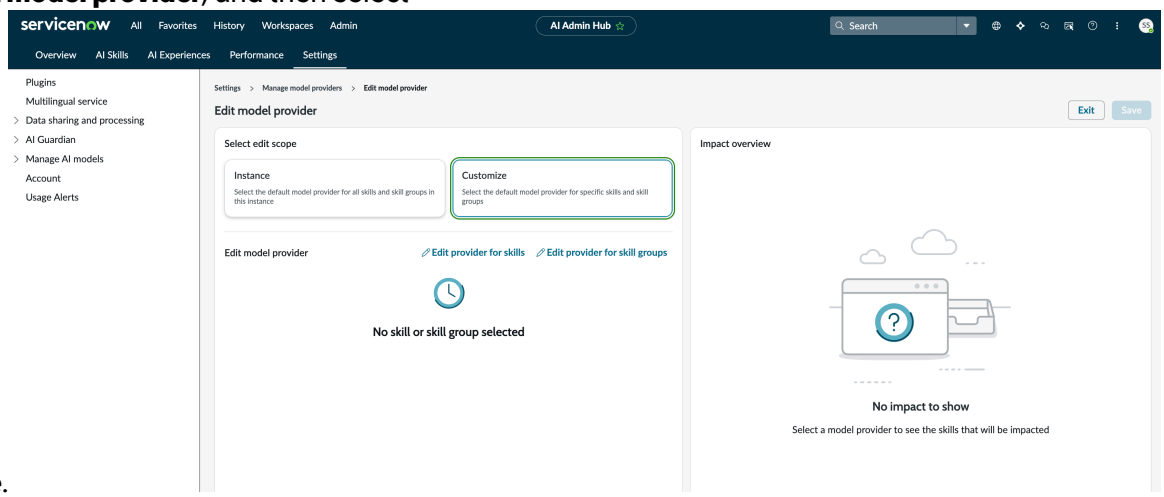

Customize.

4. In the *`Edit model provider`* section, select **Edit provider for skill groups**, select the skill group as **LEAP** and select required LLM

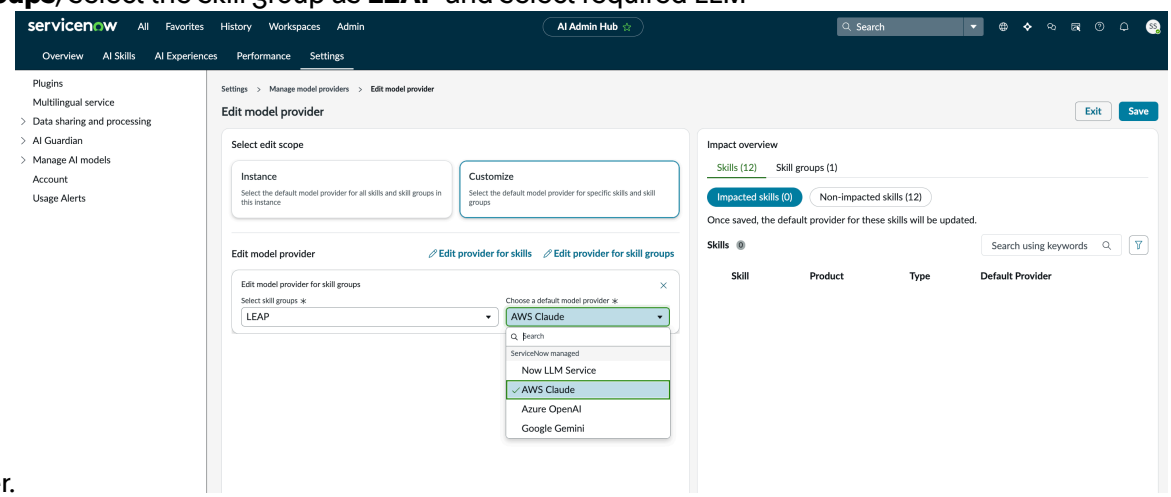

provider.

5. Select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure Ansible automation integration

Connect LEAP to Ansible Automation Platform so the Ansible Discovery and Execution Agents can automatically identify job templates and launch them during incident remediation.

Before you begin

Confirm that you have the following components installed and configured:

- LEAP application activated
- Ansible Automation Platform with MCP Server Console configured

  The MCP Server is the middle ware component that helps communication between LEAP and Ansible Automation Platform.

- Network connectivity between ServiceNow and Ansible Automation Platform
- API token for the MCP Server Console

Role required: *sn_itom_leap.leap_admin*

About this task

The Ansible Automation integration requires configuration of the MCP Server Console connection and agent to enable communication between LEAP and Ansible Automation Platform.

For more details about the integration, select **Help** on the Connectors page, and then select **Learn more** to open the LEAP and Ansible MCP server guide.

Procedure

1. On the LEAP homepage, navigate to **SettingsConnectors**, and select **Connect**.

   Note:

   Another way to start this configuration is by selecting **Connect** in the Ansible connection banner on the LEAP homepage.

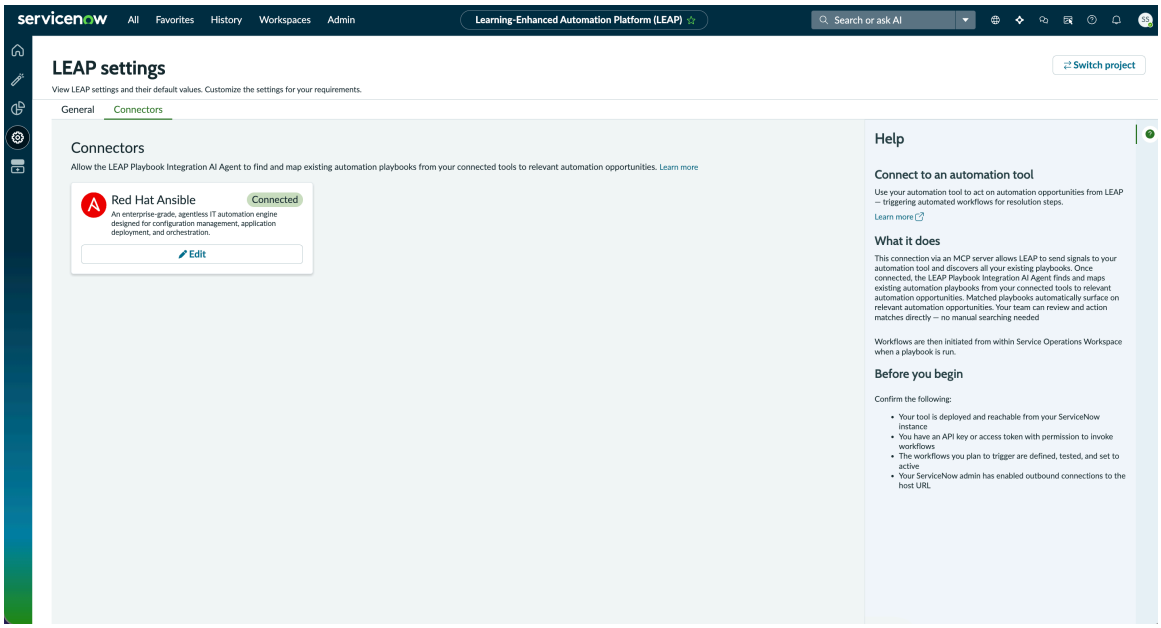

2. Configure the Ansible MCP Server Console connection by completing these fields:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Ansible MCP Server Console connection fields

Field	Description
MCP Server URL	The URL of the Ansible MCP Server Console endpoint.
Authentication Method	API key authentication method.

3. In the Token field, enter the API token used for authentication.
4. Assign the required roles to users who will work with Ansible integration:
   - sn_itom_leap.leap_admin = For LEAP administrators who configure step-to-job mappings
   - sn_itom_leap.leap_agent = For incident responders who execute Ansible automations
5. Select **Save** to apply the configuration.

Result
After you configure the Ansible Automation Integration, the Ansible discovery agent analyzes automation opportunities and identifies relevant job templates. When resolution steps are generated for a critical automation opportunity, the Ansible agent is triggered which launches mapped automation for incident remediation.

The configured connection appears in the Connectors table on the LEAP settings page.

Using LEAP

Use LEAP to analyze automation opportunities (AOs), generate resolution steps and playbooks, create problem records, build knowledge base articles, and execute Ansible automations during incident remediation.

LEAP tasks

AOs are categorized depending on how critical the services are, number of hand offs, and incident information:

- Critical
- High
- Moderate
- Low

LEAP groups incidents into AOs categorized by priority. Use AOs to research incidents, identify root causes, and develop resolutions. You can also map AOs to Ansible job templates and execute automated remediation directly from the SOW.

AOs enable you to create problem records for deeper analysis, knowledge base articles for future reference, and playbooks that give IT operators access to resolution steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Ask Otto feature of ServiceNow Otto streamlines artifact creation. Select

to open the ServiceNow Otto panel, where a workflow menu guides you through creating the artifacts you need. You can also generate artifacts using the **Actions** button.

1. Generate and modify resolution steps in LEAP

   Generate resolution steps for automation opportunity and when required get focused solutions by modifying the resolution steps.

2. Create sub-groups for automation opportunities

   Create granular automation opportunities to get targeted resolutions for similar looking issues.

3. Monitor automation outcomes with the LEAP value dashboard

   View and understand how automation helps LEAP to achieve cost and time savings.

4. Create LEAP problem records

   Create a problem record using the resolution steps of an AO. The problem record helps you perform an in-depth analysis to find the root cause of the incident.

5. Generate LEAP knowledge base articles

   Generate AI-enhanced, publication-ready knowledge base articles from resolution steps. Generated articles include structured sections, auto-populated metadata, and a link to the originating incident. Articles are routed through the standard knowledge base approval workflow and published to the Knowledge Center.

6. View LEAP resolution step recommendations in SOW

   View resolution step recommendations in SOW.

7. Generate LEAP playbooks

   Generate automated playbooks that use the resolution steps to resolve similar incidents.

8. Use LEAP playbooks in Service Operations Workspace SOW

   Use LEAP generated playbook to resolve incidents in SOW.

9. Map Ansible jobs to resolution steps

   Create mappings between automation opportunity resolution steps and Ansible job templates to enable automated incident remediation.

10. Execute Ansible automations for incidents

    Use the Ansible Execution Agent to launch mapped Ansible job templates during incident remediation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Manage Automation projects

Create and manage Automation projects to run LEAP across multiple incident taxonomies independently.

Before you begin

Role required: LEAP admin (`sn_itom_leap.leap_admin`)

The default Automation project must be activated through the AI Admin Workspace before you can create additional Automation projects. If you have not yet activated LEAP, complete the initial setup first.

About this task

After activating the default Automation project, use the following procedures to create and manage additional Automation projects:

Procedure

1. Create an Automation project.
2. Start the analysis job.
3. Switch between Automation projects.
4. Edit an Automation project.

Create an Automation project

Create an Automation project to run LEAP analysis on a specific set of incident data.

Before you begin

Role required: LEAP admin (`sn_itom_leap.leap_admin`)

About this task

Each Automation project runs LEAP analysis independently on a specific set of incident data. You can create multiple projects to analyze different incident categories, assignment groups, or time periods. Multiple projects can use the same taxonomy when analyzing different data sets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to the **Manage Automation projects** page.

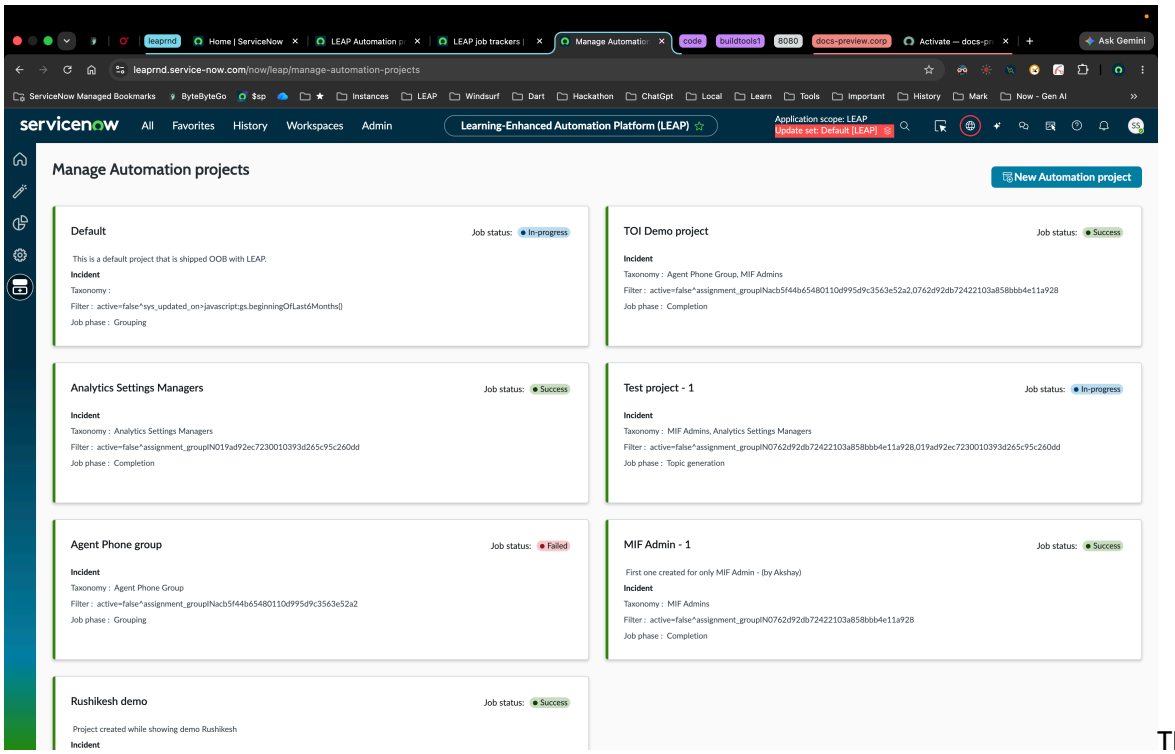

The page displays all existing Automation projects as cards, including the Default Automation project.

2. Select **New configuration**.
   The **Automation Project Details** page opens in create mode.

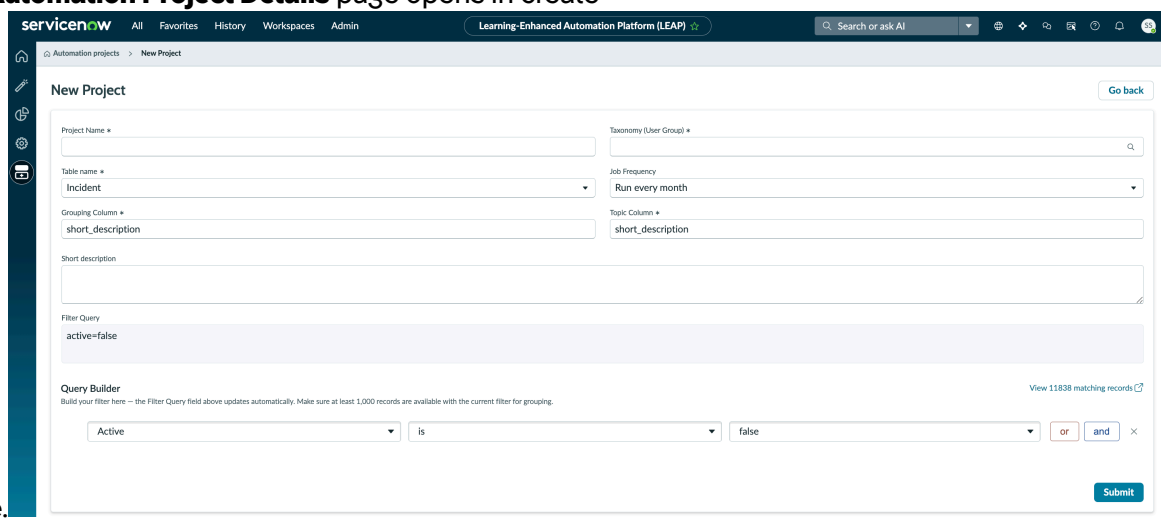

3. Complete the following fields:
   - Name: Enter a unique name for the project.
   - Description: (Optional) Describe the purpose of this project.
   - Table: Select the incident table to analyze. Only the incident table is supported in this release.
   - Filter: Define a filter to scope the incidents for this project. The filter must return at least one thousand records and can't exceed 300 characters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Grouping column: Select the column used to group incidents.
- Topic column: Select the column used to generate topics.
- Taxonomy: Select one or more taxonomies for this project. The taxonomy field is multi-select. Multiple projects can use the same taxonomy.
- Job frequency: Select how often the analysis job runs.

Mandatory fields must be filled and the filter must return at least one record before you can start the job. Validation errors are displayed inline.

4. Select **Submit**.
   The project is created in ready status. The analysis job has not started yet. You can edit the project configuration before starting the job.

What to do next
After creating the project, start the analysis job to begin generating automation opportunities.

Start the analysis job

Start the analysis job for an Automation project to begin generating automation opportunities and insights.

Before you begin
Role required: LEAP admin (`sn_itom_leap.leap_admin`)

About this task
After creating an Automation project, you must start the analysis job to activate the scheduled analysis and begin generating automation opportunities. The project configuration is locked once the job starts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Open the Automation Project in the **Automation Project Details** page.
2. Select **Start analysis**.
   LEAP generates the required grouping skills and activates the scheduled analysis job in the back end. The job status transitions from ready to running.

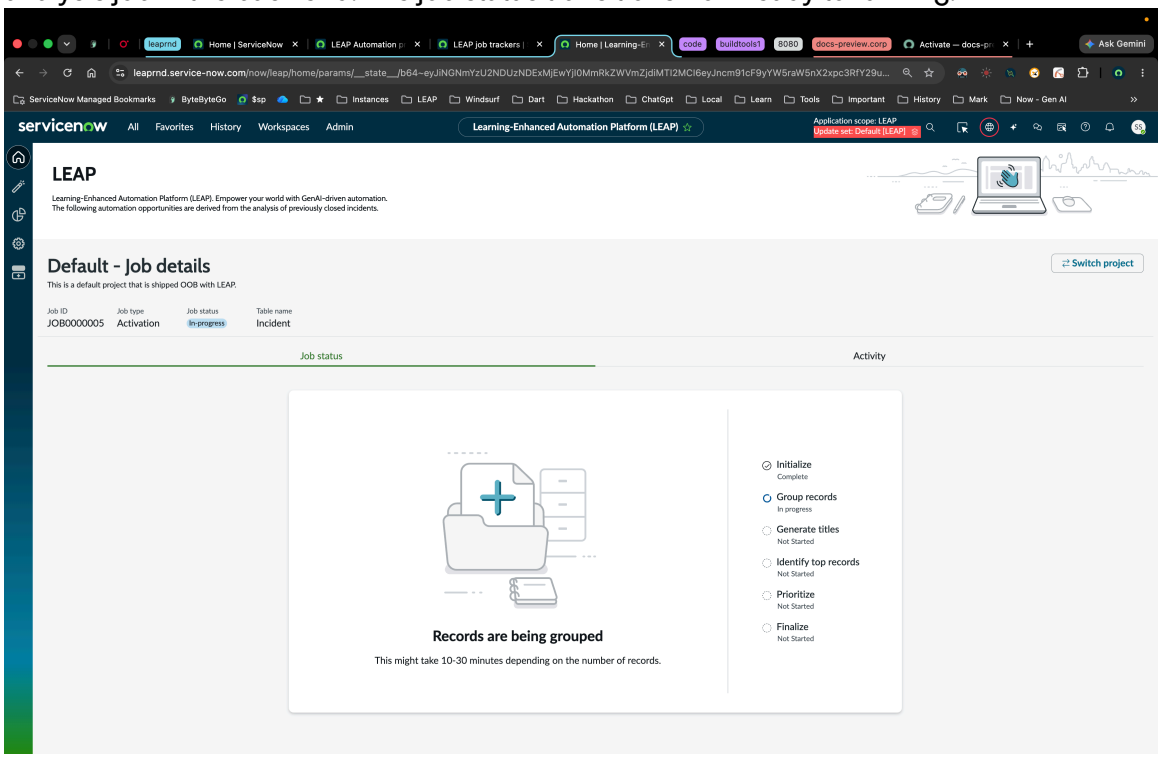

Result

When the job completes, automation opportunities and insights for this project are available in the LEAP workspace.

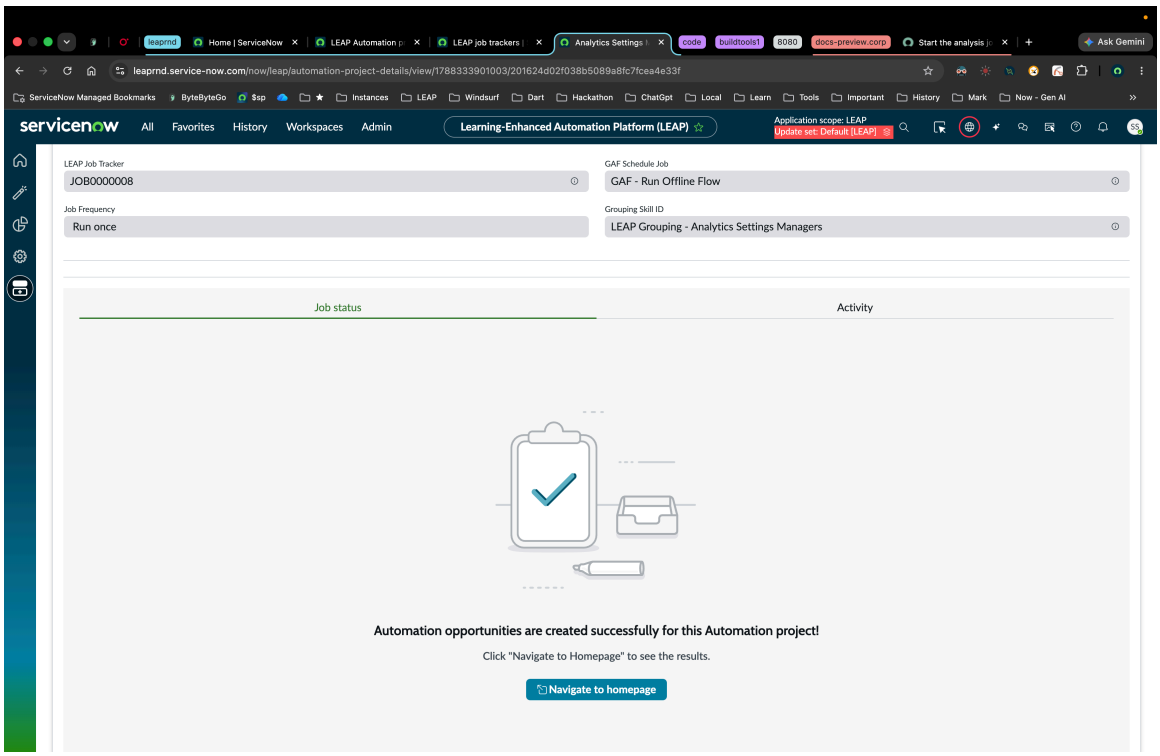

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Switch between automation projects

Switch between Automation Projects to view metrics, reports, and automation opportunities scoped to a specific project.

Before you begin

Role required:

- LEAP viewer `sn_itom_leap.leap_viewer`

About this task

When you have multiple Automation Projects configured, you can switch between them to view project-specific data.

Note: Switching projects does not affect other projects or their running analysis jobs.

Procedure

1. On any LEAP page that displays a project selector, open the **Switch project** menu.
   The project selector appears on all LEAP pages. It displays the name of the currently active Automation Project.

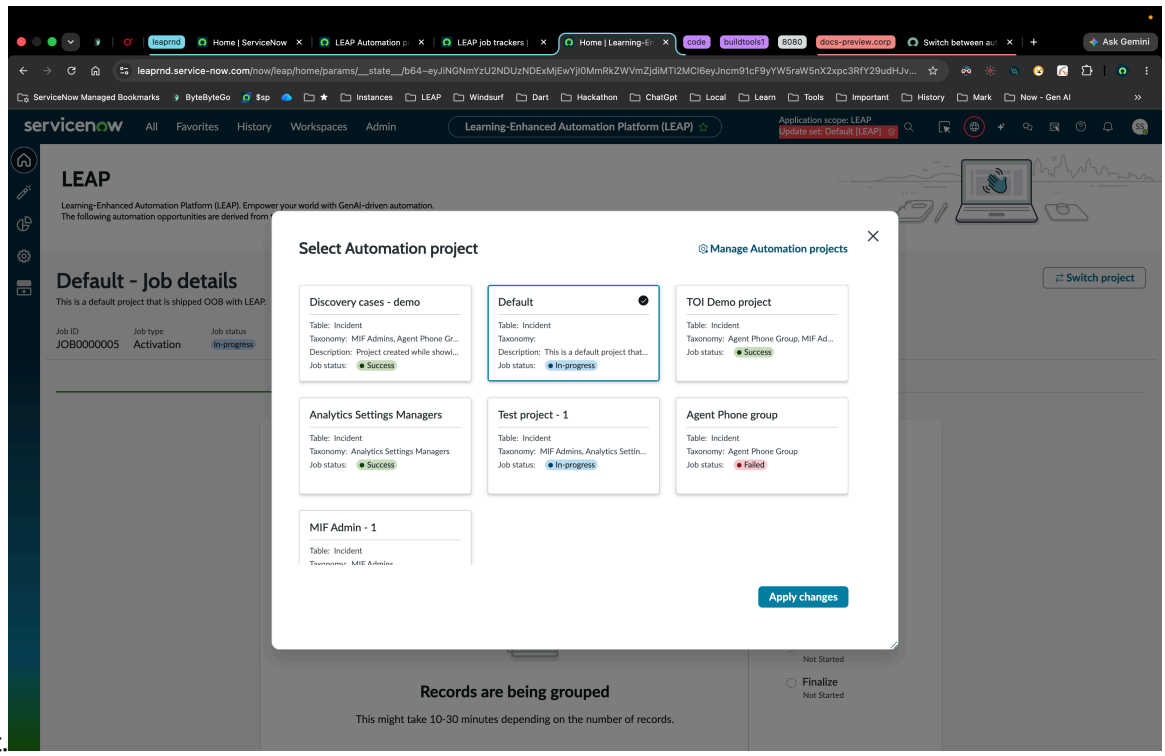

2. Select the Automation Project you want to switch to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

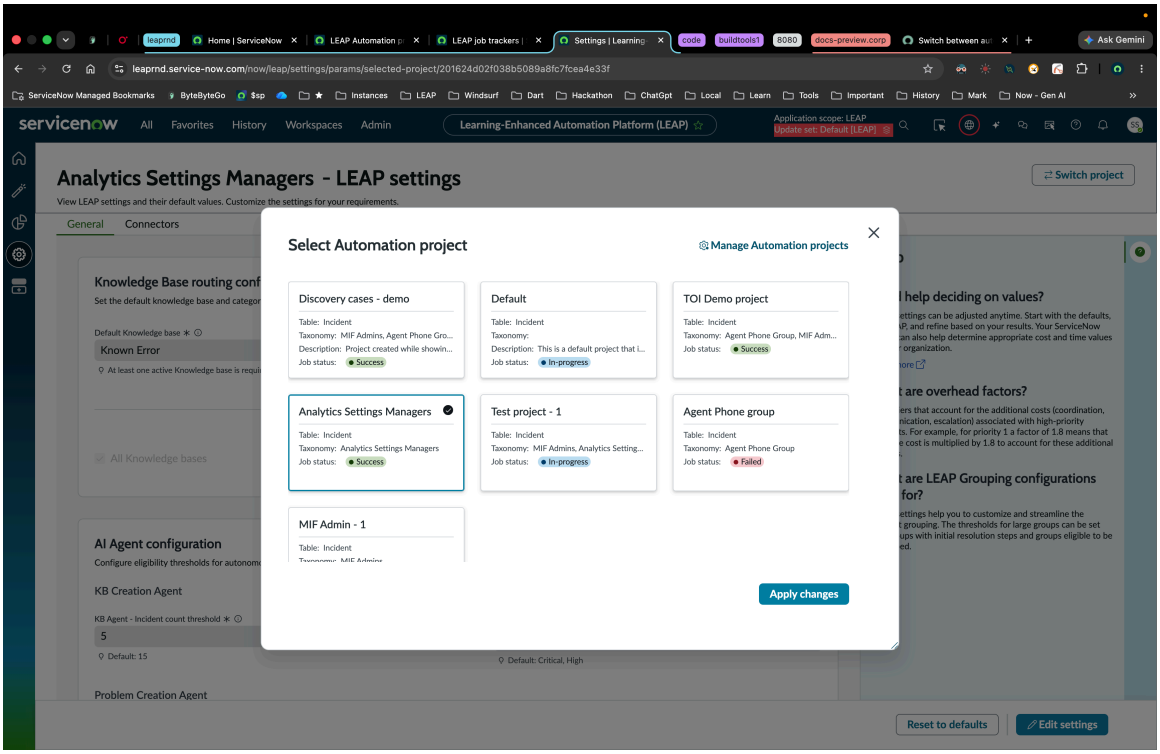

The page refreshes to display metrics or reports or automation opportunities scoped to the selected project.

Edit an Automation Project

Edit an Automation Project configuration when the analysis job is in ready status.

Before you begin

Role required: LEAP admin (`sn_itom_leap.leap_admin`)

About this task

Procedure

1. Open the Automation Project from the **Manage Automation Projects** page.
2. Select **Edit**.

   Note: You can only edit a project when its job status is ready. Projects with a running or completed analysis job can't be edited.

3. Update the fields as needed, then select **Save**.

Generate and modify resolution steps in LEAP

Learn how to generate, edit, and track AI-generated resolution steps for automation opportunities in LEAP.

About this task

Resolution steps are actions that resolve issues identified by an automation opportunity. During the initial run, GAF generates these steps by analyzing work notes and historical resolutions.

The **Generate resolution steps** button creates resolution steps as required. As your system processes new incidents and records, regenerate these steps to maintain solution accuracy. LEAP provides three search capabilities: AI search, KB search, and web search.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- AI search: Scans records within your existing instance to identify relevant resolutions.
- KB search: Searches knowledge base resources external to your organization to suggest more effective resolution steps.
- Web search: Queries web resources beyond your organization's systems to find comprehensive solutions, broadening the scope of potential resolutions.

Note:

- Regenerating resolution steps incurs chargeable transactions. Exercise caution and careful planning before using the regenerate button to avoid unexpected costs.
- Only a system administrator or a LEAP admin with system administrator access can enable these additional sources of searches. Users with agent or viewer access must contact their administrator to activate these features. Once enabled, LEAP expands its search capabilities to deliver more comprehensive and effective resolution steps.

Resolution steps are generated only for automation opportunities that match the **Default resolution steps filter** configured in LEAP settings. The default filter is set to Critical priority (`automation_priority=40`). Items are processed in batches in the background, so performance is not affected during generation. To change the priority filter, update the condition builder in LEAP settings fields.

Before you begin

Role required: LEAP admin

Procedure

1. Select **Workspaces** > **LEAP**.
2. On the LEAP landing page, select the required automation opportunity.
3. If resolution steps are unavailable in the Overview section, select **Generate resolution steps**.
4. If resolution steps are available, and you want to get alternative resolution steps, then use the **Regenerate** button.

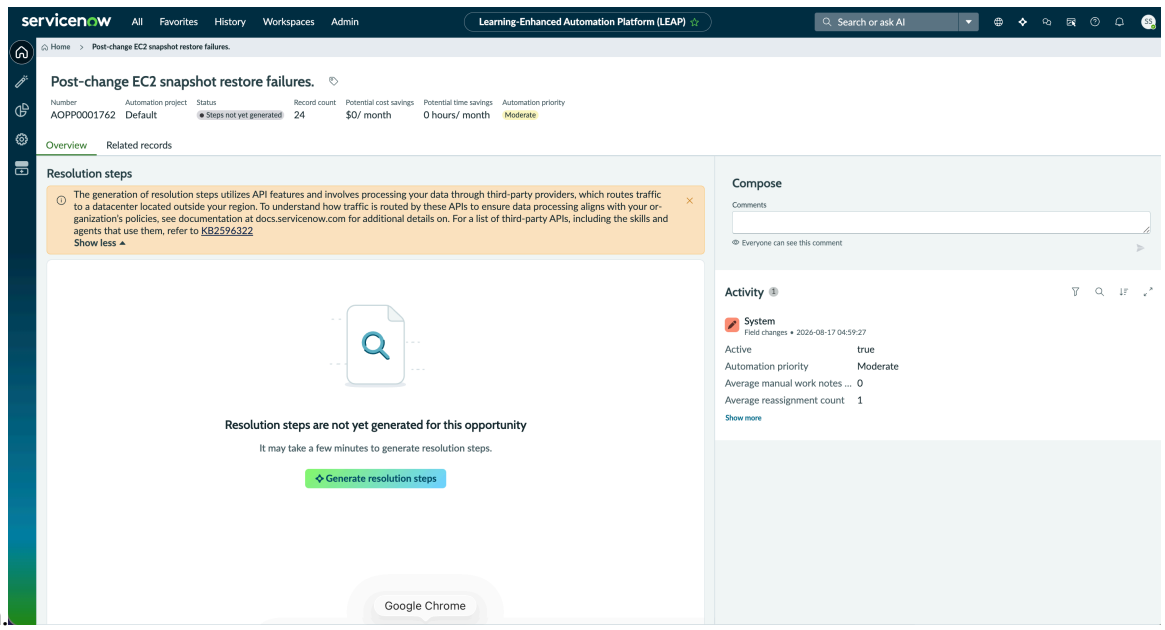

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- The Regenerate button is turned off while manual editing is in progress.
- Resolution steps are generated based on the capability enabled either using AI search or Knowledge Base or Web search.

5. To manually modify resolution steps, select the edit icon.
   Each update to the resolution steps is recorded in the LEAP Outcome Snapshots table for traceability.
6. To revert to the last generated resolution steps, use the **Revert to AI generated steps** button.

Result

Create sub-groups for automation opportunities

Large automation opportunities in LEAP are broken into smaller, more manageable sub-groups to support focused and contextual incident resolution steps.

Before you begin

Role required: LEAP admin

Procedure

1. Navigate to LEAP landing page, and select the automation opportunity that you want to refine.
2. Select **Related records** tab.
3. Select **Create Sub Groups**.
   The records in the automation opportunity are divided into sub-groups.

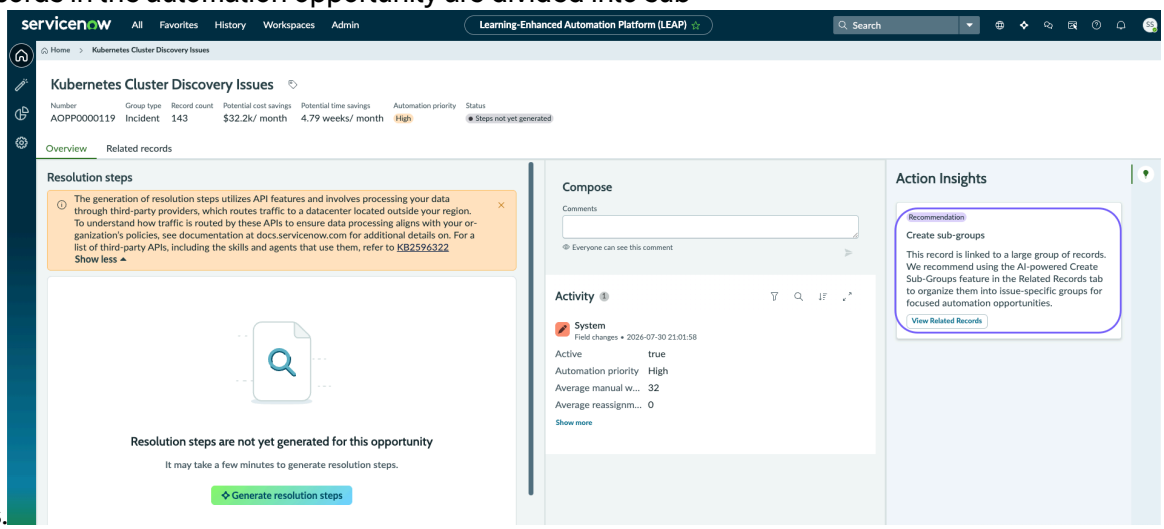

4. Select the **Sub Groups** tab to view the child sub-groups created.
   Each child sub-group has it own ID and you can generate focused resolution steps with only the records in the sub-group. On the LEAP landing page, the parent automation opportunity shows child automation opportunity in the *Sub Groups* column.

Monitor automation outcomes using the LEAP value dashboard

Use the LEAP value dashboard to monitor automation outcomes across ServiceNow playbooks, Ansible playbooks, KB articles, and problem records.

Before you begin

Role required: LEAP admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The LEAP value dashboard tracks automation outcomes across all outcome types that LEAP produces. Use the **Overview** tab to monitor cross-outcome metrics and identify high-value automation opportunities.

Procedure

1. Navigate to the LEAP dashboard by selecting the dashboard icon .

   The **Default - LEAP value dashboard** page opens. To view a different automation project dashboard, select **Switch project** in the upper-right corner.

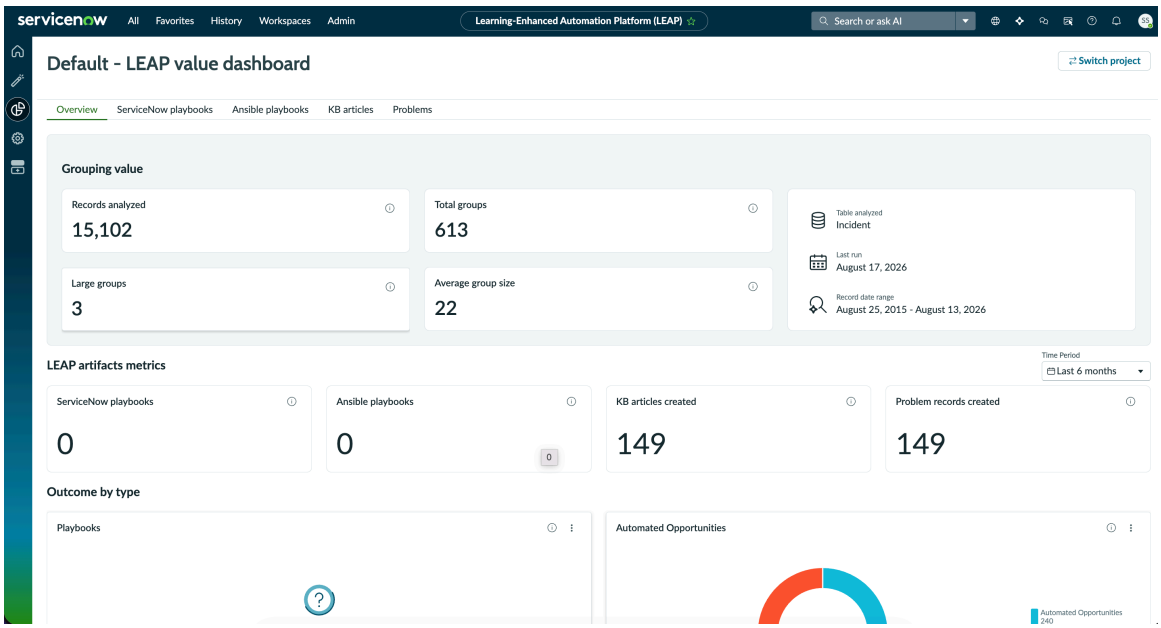

   The **Overview** tab opens by default, displaying grouping statistics and a summary of all automation outcomes.

2. Review the **Grouping value** section to understand the scope and results of the analysis.

   This section displays key grouping metrics and analysis metadata:

   - **Records analyzed** - total number of records that LEAP processed in the current analysis run.
   - **Total groups** and **Average group size** - clusters of similar incidents that LEAP identified as candidates for automation.
   - **Large groups** - groups that exceed the automation threshold and represent high-value automation opportunities.
   - **Table analyzed** - the ServiceNow table that LEAP used as the data source, for example, Incident.
   - **Last run** - the date of the most recent GAF analysis run.
   - **Record date range** - date range of records included in the analysis.

3. Review the **LEAP artifacts metrics** section.

   Use the **Time Period** list in the upper-right of this section to filter all artifact counts by a time range. The default is **Last 6 months**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Review the four artifact summary cards: **ServiceNow playbooks**, **Ansible playbooks**, **KB articles created**, and **Problem records created**.
Each card shows the total count of that artifact type that LEAP created in the selected time period.

You can see a cross-outcome summary of all artifacts that LEAP generated.

4. Review the **Outcome by type** section.

a. Review the **Playbooks** donut chart.
The chart breaks down total playbooks by type: **ServiceNow playbooks** and **Ansible playbooks**. The center value shows the combined total.

b. Review the **Automated Opportunities** donut chart.
The chart shows the proportion of automation opportunities that are automated compared to those that are not automated. The center value shows the total number of groups.

5. Review the **Savings by playbooks** section.

a. Check **Total savings by playbooks** and **Total agent-hours saved by playbooks**.
These values show realized savings from ServiceNow and Ansible playbook executions.

b. Review the **Savings distribution** chart.
The chart shows how total savings are distributed across playbook types.

You can see quantified savings from playbook-driven automation for the selected time period.

Result

The **Overview** tab provides a cross-outcome snapshot of all artifacts that LEAP generated and the total savings from playbook automation. To review detailed metrics for a specific outcome type, see the related tasks: Review ServiceNow playbook metrics, Review Ansible playbook metrics, Review KB article metrics, and Review problem record metrics.

Review ServiceNow playbook metrics

Use the **ServiceNow playbooks** tab on the LEAP value dashboard to review execution trends, playbook status, and top-performing playbooks.

Before you begin

Role required: LEAP admin

About this task

The **ServiceNow playbooks** tab shows execution volume, status breakdown, and top-performing playbooks for the selected time period. Use the **Time Period** list in the upper-right corner to filter all metrics by a date range. The default is **Last 6 months**.

Procedure

1. Open the LEAP value dashboard and select the **ServiceNow playbooks** tab.
The **ServiceNow playbooks executions overview** section opens.
2. Review the **Tickets resolved by playbooks - Monthly trend** chart.
The bar chart shows the number of tickets resolved by ServiceNow® playbooks each month over the selected time period.
3. Review the three donut charts **Playbooks by status**, **Published playbooks by automation priority**, and **Playbooks utilization overview**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Playbooks by status** — shows the split between draft and published playbooks.
- **Published playbooks by automation priority** — shows published playbooks grouped by priority level, such as, Critical.
- **Playbooks utilization overview** — shows the ratio of executed to unexecuted playbooks.

4. Review the tables **Top 10 Playbooks – by tickets resolved** and **Top 10 Playbooks – by time savings**.
   Each table lists playbooks by LEAP number and group description. Select a LEAP number to open the automation opportunity detail page.

Result
You can see execution volume, status breakdown, and the top-performing ServiceNow playbooks for the selected time period.

Review Ansible playbook metrics

Use the **Ansible playbooks** tab on the LEAP value dashboard to review Ansible execution counts, agent-hours saved, and top automation opportunities resolved by Ansible.

Before you begin
Role required: LEAP admin

About this task
The **Ansible playbooks** tab shows how Ansible automation is contributing to ticket resolution and time savings. Use the **Time Period** list in the upper-right corner to filter all metrics by a date range. The default is **Last 6 months**.

Procedure

1. Open the LEAP value dashboard and select the **Ansible playbooks** tab.
   The **Ansible playbooks executions overview** section opens.
2. Review the four metric cards at the top of the section.
   - **Ansible playbooks executed** - total Ansible playbook executions triggered by LEAP in the selected time period.
   - **Total agent-hours saved via Ansible** - cumulative agent time saved through Ansible automation.
   - **Tickets resolved by Ansible** - total tickets resolved by Ansible playbook execution.
   - **Cost savings via Ansible** - total cost savings from Ansible-driven resolutions.
3. Review the **Ansible executions - monthly trend** chart.
   The bar chart shows the number of Ansible executions per month over the selected time period.
4. Review the **Top 10 Automation Opportunities – by ticket resolved using Ansible playbooks** and **Top 10 Automation Opportunities – by time savings using Ansible playbooks** tables.
   Both tables list automation opportunities by name. Select an automation opportunity name to open its detail page.

Result
You can see how Ansible automation is contributing to ticket resolution and time savings across your automation opportunities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Review KB article metrics

Use the **KB articles** tab on the LEAP value dashboard to review the volume of knowledge base articles that LEAP created and identify which automation opportunities are generating the most knowledge content.

Before you begin
Role required: LEAP admin

About this task
The **KB articles** tab shows KB article creation trends and which automation opportunities are producing the most knowledge content. Use the **Time Period** list in the upper-right corner to filter all metrics by a date range. The default is **Last 6 months**.

Procedure

1. Open the LEAP value dashboard and select the **KB articles** tab.
   The **KB articles overview** section opens.
2. Check the **KB articles created – monthly trend** chart.
   The `bar chart` shows the number of KB articles that LEAP created each month over the selected time period.
3. Review the **KB articles created by source** donut chart.
   The chart shows the proportion of KB articles created by the LEAP AI agent compared with KB articles created manually.
4. Review the **Top 10 opportunities that generated KB articles** table.
   The table lists automation opportunities that produced KB articles, with a link to each generated knowledge article in the **Knowledge** column. Select an automation opportunity number to open its detail page.

Result
You can see the volume of AI-generated versus manually created KB articles and identify which automation opportunities are producing the most knowledge content.

Review problem record metrics

Use the **Problems** tab on the LEAP value dashboard to review the volume of problem records that LEAP created and identify which automation opportunities are driving root cause investigation.

Before you begin
Role required: LEAP admin

About this task
The **Problems** tab shows problem record creation trends and which automation opportunities are generating the most root cause investigation. Use the **Time Period** list to filter all metrics by a date range. The default is **Last 6 months**.

Procedure

1. Open the LEAP value dashboard and select the **Problems** tab.
   The **Problem records overview** section opens.
2. Review the **Problem records created – monthly trend** chart.
   The bar chart shows the number of problem records that LEAP created each month over the selected time period.
3. Review the **Problem records created by source** donut chart.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The chart shows the proportion of problem records created by the LEAP AI agent compared to those created manually.

4. Review the **Top 10 opportunities that generated Problem records** table.
   The table lists automation opportunities that produced problem records, with a link to each problem record in the **Problem** column. Select the automation opportunity number to open its detail page.

Result

You can see the volume of AI-generated versus manually created problem records and identify which automation opportunities are driving root cause investigation.

Create LEAP problem records

Create problem records for LEAP from available incident clusters, also known as automation opportunities, for further analysis.

Before you begin

Role required: LEAP admin

About this task

The problem record creation process automatically transfers critical information from incidents to problem records. When a problem record is created, LEAP includes all relevant information. This information includes:

- incident group description
- incident group link
- resolution steps

The automated data transfer is characterized with:

Root cause documentation

LEAP populates the cause notes section of the problem record with the same details as available in existing incident records.

Incident attachment

All incidents associated with the group are automatically linked to the problem record for comprehensive tracking.

Async processing

Problem records are generated efficiently without manual intervention even during high-volume scenarios.

Procedure

1. Select **Workspaces** > **LEAP**.
2. On the LEAP landing page, select an automation opportunity.
3. If resolution steps don't appear in the Overview section, select **Generate resolution steps**.
4. Select **Actions** > **Create problem record** button to create a record for further in-depth analysis.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

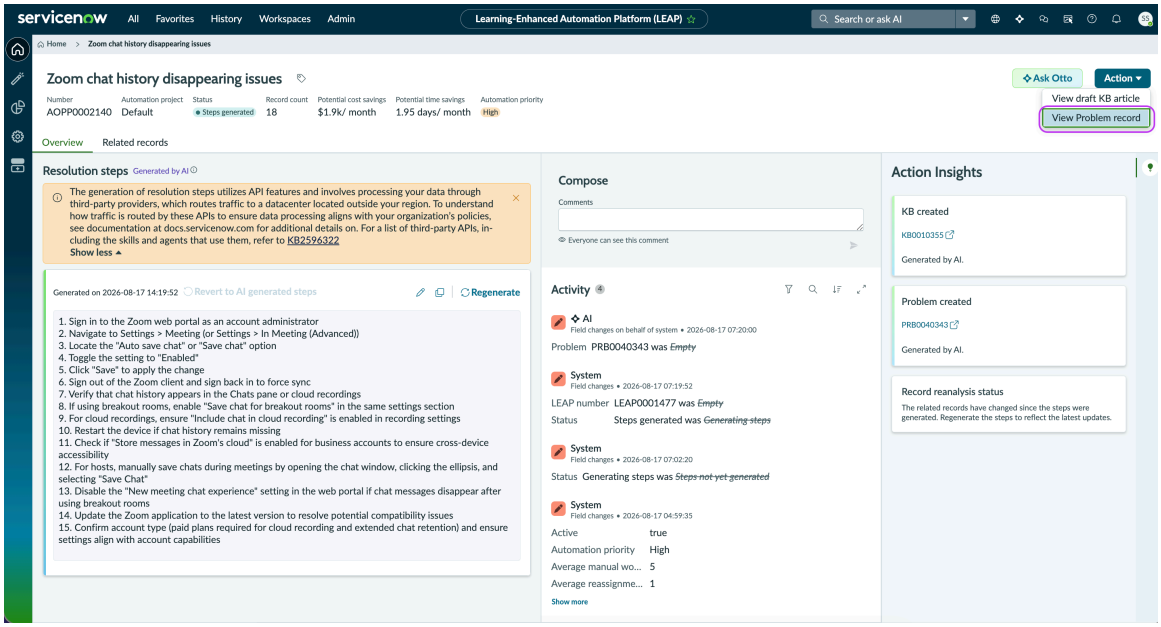

Another option is to invoke the LEAP AI agent by selecting the

Ask Otto button and creating problem record.

Generate LEAP knowledge base articles

Generate AI-enhanced knowledge base articles from automation opportunity resolution steps to share structured, publication-ready knowledge across your organization.

Before you begin

Resolution steps must exist for the automation opportunity before generating a knowledge base article.

Role required: LEAP admin

A default knowledge base must be configured in LEAP properties before LEAP AI agent creates knowledge base articles. See LEAP settings fields.

Each knowledge base used in LEAP must have the correct Can Contribute permissions configured. See Knowledge base permissions for LEAP.

About this task

LEAP uses AI to transform resolution steps into a complete, structured knowledge base article. The generated article includes these sections:

- Summary
- Symptoms — inferred from incident data and resolution context
- Resolution steps — expanded with sub-steps and code or command formatting
- Related resources
- Metadata — including Article Number, Product, Category, Visibility, and Valid To date

The article is created in Draft state and routed through the standard knowledge base approval workflow before publication. Each generated article maintains a link to the originating automation opportunity and is suggested when similar incidents occur.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When the LEAP AI agent automatically creates an article, it is routed to the default knowledge base and category configured in LEAP properties. When you create an article manually using the **Actions** menu, a dialog opens where you can select the knowledge base and category before the article is created.

Note: AI-generated content may be inaccurate. Review the generated article before submitting it for approval to confirm it reflects the correct resolution information.

Procedure

1. Navigate to **Workspaces** > **LEAP**.
2. On the LEAP landing page, select the automation opportunity for which you want to generate a knowledge base article.
3. If resolution steps are unavailable in the Overview section, select **Generate resolution steps** and wait for the steps to be generated.
4. Select **Actions** > **Draft KB article**.

   A dialog opens to configure the location and category of the knowledge base where the article created should reside.

   Tip: You can also generate a knowledge base article by selecting the button to open the ServiceNow Otto panel and selecting the **draft KB article** option from the menu. Articles created this way route to the default knowledge base configured in LEAP properties.

5. In the dialog, select a knowledge base from the **Select knowledge base** list.
   The list shows only the knowledge bases selected in the **Eligible knowledge bases** field in LEAP properties. If no eligible knowledge bases are configured, all active knowledge bases are shown.
6. Select a category from the **Category** list.
   Categories are filtered based on the knowledge base you selected.
7. In the **KB Meta** field, add or remove tags to improve the discoverability of the article for virtual agents and search.
   Meta tags help virtual agents find the correct knowledge base article. Tags are pre-populated based on the automation opportunity context. You can add new tags or remove existing ones as needed.
8. Select **Draft KB article**.
   LEAP generates a complete knowledge base article and opens it in the Knowledge Center in draft state.
9. Review the generated article, including all sections and automatically populated metadata, and make any corrections.
   Technical commands and code snippets are formatted with monospace font and include platform-specific context where applicable.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Submit the article through the knowledge base approval workflow.
    After approval, the article is published to the Knowledge Center. It is linked to the originating incident record and suggested when similar incidents occur.

View LEAP resolution step recommendations in SOW

Use LEAP to view AI-generated resolution step recommendations that match the incident description in SOW to resolve incidents more efficiently.

Before you begin
Role required: LEAP agent

About this task
When you work on an incident in SOW, LEAP compares the incident short description against existing resolution steps. Only the steps relevant to the incident description appear in the **Recommendations** tab under LEAP Resolution Steps. Steps that don't match the incident description are filtered out. If fewer steps appear than expected, this indicates successful filtering of irrelevant steps.

Note: AI-generated content may contain inaccuracies. Review all recommended resolution steps before applying them to verify their accuracy and suitability for your environment.

Procedure

1. Open an incident in SOW.
2. In the insights panel, select the **Recommendations** tab.
   LEAP displays resolution steps that match the incident's short description.
3. Review the recommended resolution steps.
   If fewer steps appear than expected, LEAP filters out steps that don't match the incident description.
4. Apply the resolution steps to resolve the incident.

Generate LEAP playbooks

Create standard, dynamic, and automated playbooks using LEAP to promote fast incident resolution and response time.

Before you begin
Role required: LEAP admin

Procedure

1. Select **Workspaces** > **LEAP**.
2. On the LEAP landing page, select the required automation opportunity to create a playbook.
3. If resolution steps are unavailable in the Overview section, select **Generate resolution steps**.
4. Select one of the following options to create playbook.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

UI option	Steps
Actions	Select Actions button , and select **Create Playbook**.
Ask Otto	**a.** Select button. **b.** In the ServiceNow Otto panel, select **Create playbook**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

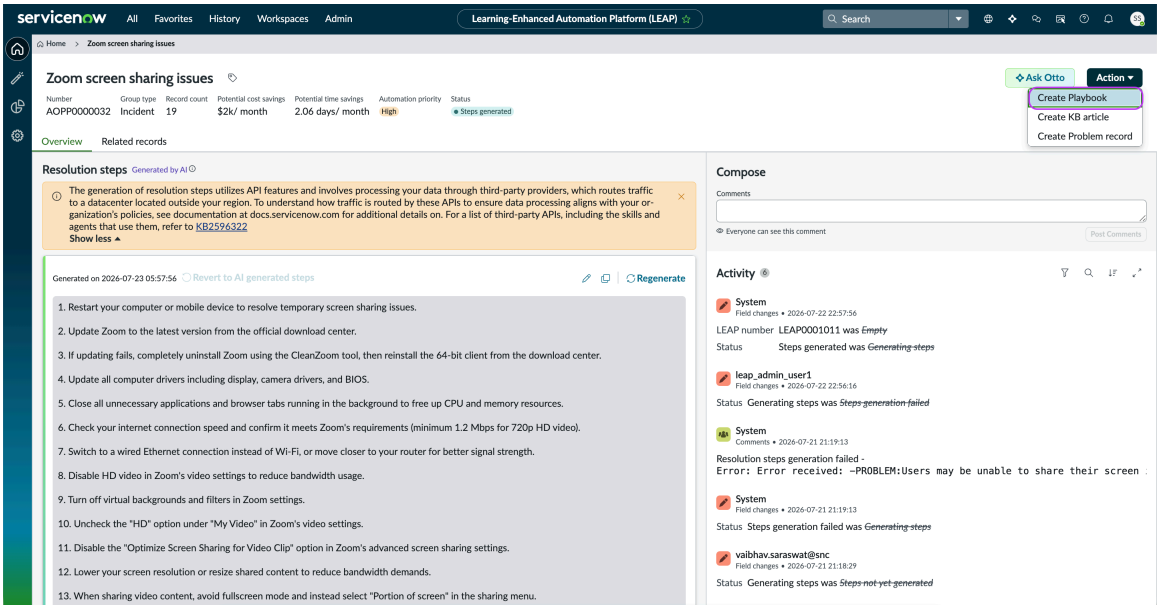

Workflow Studio opens and the playbook generation process starts. The **Playbook name** and **directions** fields are pre-filled with the problem record title and the generated resolution steps.

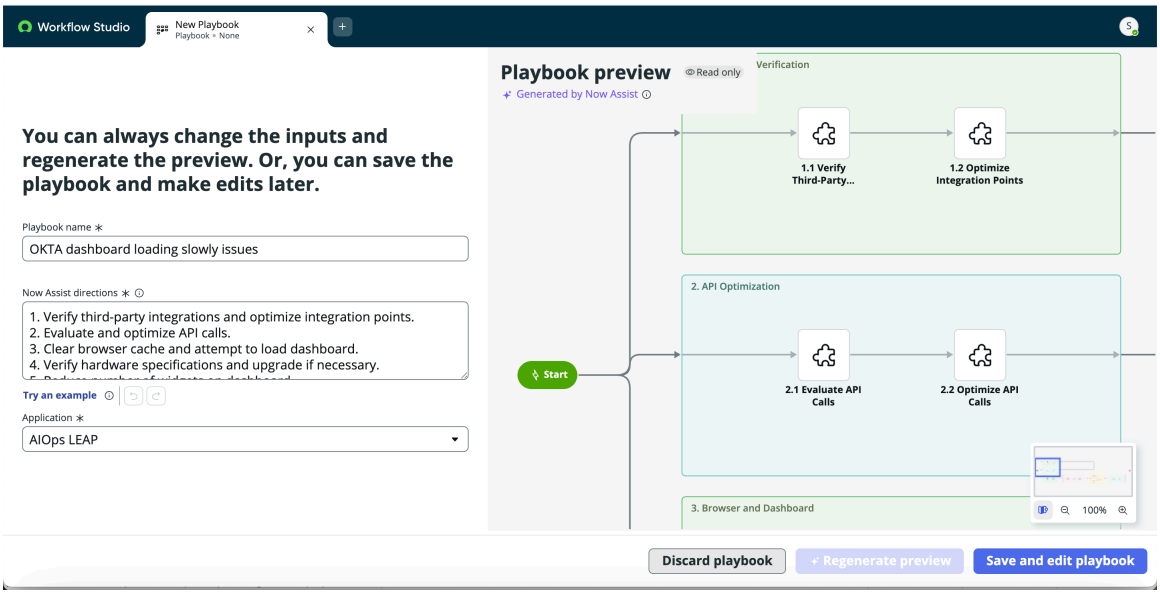

Use LEAP playbooks in Service Operations Workspace SOW

Use LEAP playbooks in the SOW.

Before you begin

Role required: LEAP admin

Procedure

1. Navigate to **All** > **LEAP**.
2. Select the required automation opportunity.
3. Generate LEAP playbooks.
4. Open the SOW interface using the Workspaces menu.
5. Select the required incident, and access the insights' panel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Select the playbook icon

.
7. If a LEAP playbook is available for the selected incident, review the AI-filtered resolution steps and use them to resolve the issue.

Map Ansible jobs to resolution steps

Create mappings between automation opportunity resolution steps and Ansible job templates to enable automated incident remediation.

Before you begin

Before mapping Ansible jobs to resolution steps:

- The automation opportunity must have generated resolution steps
- The Ansible discovery agent must have analyzed the automation opportunity
- Job templates related to the resolution steps must exist and be available in your connected Ansible Automation Platform instance.
-

Role required: LEAP admin

About this task

Step-to-job mapping creates the relationship between specific resolution steps and Ansible job templates. This mapping enables the Ansible execution agent to automatically launch the appropriate automation during incident remediation. At remediation time, the Ansible execution agent reads this mapping to determine which job template to launch for each step, in sequence.

Procedure

1. Navigate to **Learning-Enhanced Automation Platform(LEAP)**.
2. Locate the automation opportunity that you want to configure for Ansible integration.
3. In the automation opportunity details page, the *`Ready for review`* section displays whether there are Ansible playbooks detected.
4. Select **Review details** to open the mapping modal.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

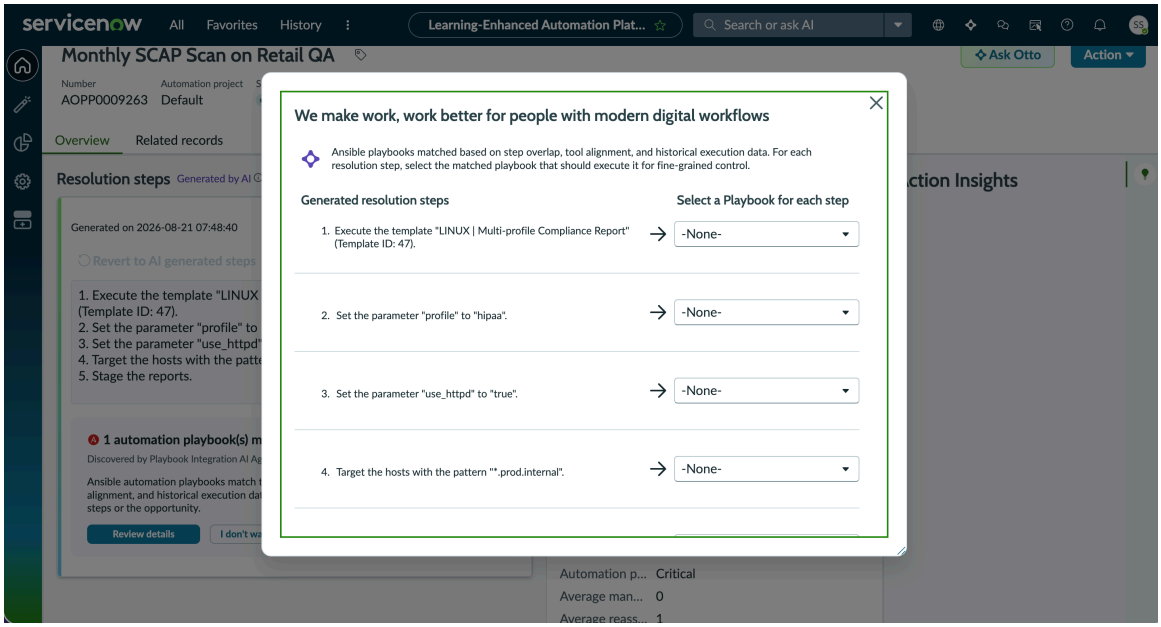

This button appears only after the Ansible discovery agent has analyzed the automation opportunity and identified candidate job templates. If the button does not appear, confirm that the Ansible discovery agent has completed analysis.

5. Review the resolution steps displayed in the mapping modal.
   The modal parses the resolution steps and displays each step with a numbered index.
6. For each resolution step, select the appropriate Ansible job templates from the drop-down list:
   - Select a job template for the step to be automated.

     A step can be automated if a matching job template exists.
   - Leave the drop-down empty if the step requires manual intervention.

   The drop-down is populated with job templates discovered by the Ansible discovery agent. Each job template shows:

   - Job template name
   - Description
7. Review your mappings to verify that:
   - Automated steps have appropriate job template selected.
   - Manual steps are left unmapped (empty drop-down).
   - The sequence of steps makes logical sense for incident remediation.
8. Select **Save** to create the mapping.
   The mapping is saved to the `sn_itom_leap_ansible_mapping` table.

Result
The step-to-job mapping is created. The automation opportunity detail page now displays each resolution step with its mapped job template. Manual steps are indicated with **(manual)**. The mapping is available for incident remediation in Service Operations Workspace.

Example: Step-to-job mapping

After saving a mapping, the automation opportunity displays resolution steps like this:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Step 1: Restart the nginx service
  ▶ nginx-restart-playbook

Step 2: Flush DNS cache
  (manual)

Step 3: Verify service health
  ▶ health-check-playbook
```

What to do next

After creating the mapping:

Note: If resolution steps are regenerated after saving, this mapping is deleted and must be recreated.

Editing replaces the existing mapping. Don't edit a mapping while it is in use during an active incident remediation.

Execute Ansible automations for incidents

Use the Ansible Execution Agent to automatically launch mapped Ansible job templates during incident remediation in the Service Operations Workspace, reducing manual effort and accelerating mean time to resolution.

Before you begin

Before executing Ansible automations:

- The Ansible Execution Agent must be configured and enabled
- Step-to-job mappings must exist for the predicted automation opportunity group

  For more information, refer to Map Ansible jobs to resolution steps.

Role required: LEAP agent

About this task

When an incident arrives in SOW, LEAP predicts the automation opportunity group using existing clustering. For critical automation opportunities where resolution steps are generated, if step-to-job mappings exist for that group, automation options are presented to incident responders for execution.

Procedure

1. In SOW, select the incident.
2. In the **Recommendations** section, review the available automation options.
   This section appears only if LEAP has predicted a critical automation opportunity group with step-to-job mappings.
3. Select the required method to complete the resolution

Method	Steps
Manual	**a.** Review the manual step description provided by the agent. **b.** Perform the required manual actions. **c.** Select **Step Completed** to confirm that the manual step is finished.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Steps
Automated	**a.** Select **Execute Ansible Playbook**. **b.** Provide the required inputs when prompted and confirm to proceed.

4. Monitor job execution status and respond to agent prompts:
5. Review the execution summary when all steps are completed.
   The agent provides a summary including total steps processed, jobs launched with their Ansible Job IDs, manual steps completed, and any failures or errors.
6. Update the incident record with the remediation results and close the incident if resolution is successful.

Result

The mapped Ansible automations are executed in sequence, with manual steps completed as needed. The incident is remediated using the automated workflow defined in the step-to-job mapping.

Example: Example execution flow

For a database connectivity incident, the execution might proceed as follows:

1. Agent: "Step 1: 'Restart the nginx service' — launching nginx-restart-playbook"
2. Agent: "Launched successfully. Ansible Job ID: 12345. Status: Running"
3. Agent: "Job completed successfully. Proceeding to next step."
4. Agent: "Step 2: 'Flush DNS cache' — this step has no playbook mapped. Please complete this step manually."
5. User: "Step completed" (after manually flushing DNS cache)
6. Agent: "Step 3: 'Verify service health' — launching health-check-playbook"
7. Agent: "All steps completed. Summary: 2 jobs launched, 1 manual step completed."

What to do next

After execution:

- Review the Ansible job logs in Ansible Automation Platform for detailed execution information.
- Document any issues or improvements for future automation opportunity mappings.
- Consider updating step-to-job mappings based on execution results and feedback.

LEAP reference

These topics provide detailed field description information required during LEAP installation.

LEAP Installer fields

Field values that you configure when installing LEAP using the LEAP Installer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

LEAP Installer skill activation form

Field	Description
Run frequency	Frequency at which LEAP runs on the configured table. This field is automatically set to Run every month.
Group records from this table	Table from which LEAP queries and filters records for grouping. This field is automatically set to Incident.
Use this column for grouping	Column from the selected table used to read records for grouping. This field is automatically set to Short description.
Use this filter	Condition filter used to select records for grouping. The default filter is `active=false^ORactive=true`. Modify this filter to change the scope of records that LEAP processes. **Note:** By default, incidents are grouped and filtered for the last 6 months.
Use this column for topic generation	Column from the selected table used to generate titles for the created groups. This field is automatically set to Short description.

Define access

Field	Description
Roles	Roles assigned to the ACL entry that determine who can access the skill. This field is automatically set to `sn_itom_leap.leap_genai_worker`. Roles selected here are available in the Select display step.

Components installed with LEAP

Several types of components are installed with activation of the plugin, including tables, user roles, and scheduled jobs.

Roles installed

Role title [name]	Description
LEAP admin [sn_itom_leap.leap_admin]	You have full control over LEAP capabilities, including the following: • Turn features on and off • Execute scheduled jobs • Manage LEAP tables and the LEAP Workspace

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description
	• Generate and update resolution steps • Create and view problems and generate LEAP playbooks from automation opportunities (AO)
LEAP viewer [sn_itom_leap.leap_viewer]	You can access and view data related to LEAP, including tables, the workspace, and problems originating from AOs.
LEAP agent [sn_itom_leap.leap_agent]	You can access the LEAP menu from the Service Operations Workspace (SOW) and trigger LEAP executions directly from within SOW.

Scheduled jobs installed

Scheduled job	Description
Populate Record Metrics Aggregate Records	Runs once per day to populate the LEAP aggregate metrics table by time periods.
Update Record Metrics Table with Calculated Metrics	Runs daily to calculate and update the records in the LEAP Record Metrics table for all the records where the metrics haven't been calculated yet.
GAF- Run Offline Flow	Runs monthly to group the tasks using the task data provided in the skill configuration while activating the LEAP skill.

Tables installed

Table	Description
LEAP aggregate metrics	Contains metrics, like cost savings, time savings, and record count, for each time period. This data is used in the LEAP report page in the workspace.
LEAP feedback	Stores the feedback given by the agent to the playbooks that are run to fix the tasks.
LEAP GAF mapping	Maintains a map between the GAF groups and the published LEAP instances to retain information in the subsequent GAF clustering runs.
LEAP group	Stores the LEAP group details, like group description, resolution steps, and whether the LEAP is automated or not.
LEAP group ranking	Maintains data, like business impact, projected cost savings, problem record, and status, that is used in the Automation Opportunities page of the workspace.
LEAP properties	Sets custom values to the properties that are used in LEAP.
LEAP job tracker	Maintains the status of background jobs init, grouping, processing, and completion phases.
LEAP outcome snapshot	Maintains snapshots of automation outcomes generated by AI or edited by users.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
LEAP Integration Project	Maintains record of integration projects that organize LEAP functionality and properties.
LEAP record metrics	Maintains data that's used in the Published LEAP page in the workspace. It contains calculated data, like time saved, tickets resolved, and rating trend.
LEAP missed opportunity	Maintains record of missed automation opportunity artifact creation of type; resolution steps, kb article, or playbook.
LEAP group strategy	Maintains the strategic data for automation opportunities, including the association of records and GAF group.
LEAP Hierarchical job	Records hierarchical job execution status with GAF group associations.
LEAP Ansible Agent status	Records Ansible agent execution and conversation status for playbook runs.
LEAP Ansible Step Mapping	Maintains record of the step to job mapping performed by LEAP.
Discovered Job Catalog	Maintains record of discovered Ansible jobs and their associations with automation opportunities.

LEAP settings fields

LEAP settings page field values help estimate cost and time savings when automation is used. These settings support cost predictability.

LEAP settings fields

Overhead factors are multipliers that are used for critical priority incidents and represent hidden costs.

Field	Description
Cost per work note ($)	Represents the cost for each work note, including labor and operational expenses for every incident. This value calculates total cost savings when automation reduces the number of work notes. Example: If your organization spends an average of $50 per work note entry, including analyst time and overhead, enter 50.
Time per work note (hrs)	Average time spent on each work note entry and helps estimate the total time impact across incidents. Example: If analysts spend an average of 30 minutes on each work note, enter 0.5. For a cluster of 100 incidents with 4 work notes each, LEAP estimates 200 hours of time savings when automation is applied.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
P1 - High priority incidents	Priority 1 incidents often require more coordination, communication, and escalation — beyond just the direct work notes. For P1 incidents, the base cost is multiplied by 1.8 to reflect the additional impact. Example: If the base cost per work note is $50 and a P1 incident has 4 work notes, the adjusted cost is $50 × 4 × 1.8 = $360.
P2 - High priority incidents	For priority 2 incidents, the base cost is multiplied by 1.4 to reflect the additional impact. Example: If the base cost per work note is $50 and a P2 incident has 4 work notes, the adjusted cost is $50 × 4 × 1.4 = $280.
P3 - Moderate priority incidents	For priority 3 incidents, the base cost is multiplied by 1.1 to reflect the additional impact. Example: If the base cost per work note is $50 and a P3 incident has 4 work notes, the adjusted cost is $50 × 4 × 1.1 = $220.
P4 - Low priority incidents	For priority 4, the base cost is multiplied by 0.8 to reflect the additional impact. Example: If the base cost per work note is $50 and a P4 incident has 4 work notes, the adjusted cost is $50 × 4 × 0.8 = $160.
P5 - Planning incidents	For priority 5, the base cost is multiplied by 0.4 to reflect the additional impact. Example: If the base cost per work note is $50 and a P5 incident has 4 work notes, the adjusted cost is $50 × 4 × 0.4 = $80.
Grouping configuration	
Large group threshold	Groups that have more incidents than the threshold set are marked as large groups. The default value set is 100 incidents. The threshold helps LEAP prioritize high-volume clusters for automation analysis. **Example:** If the threshold is set to 200, any group containing more than 200 incidents is flagged as a large group.
Initial run: Groups with resolution steps	The maximum number of groups for which resolution steps are generated in the first run of LEAP. The default is 500. Increasing this value generates broader coverage but may extend processing time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Example: If set to 500, LEAP generates resolution steps for up to 500 groups in the first run.
Remapping eligibility threshold	The minimum ratio of incidents to group size required for a group to be eligible for remapping. The default is 0.2, which represents 20% of the group size. Higher values result in more selective remapping. Increasing this value limits remapping to only well-populated groups, reducing noise from small or transient clusters. **Example:** If set to 0.2, a group must contain at least 20% of the expected group size before it's eligible for remapping.
Default resolution steps filter	A condition builder filter that determines which automation opportunities are eligible for resolution step generation. By default, this filter is set to 'Critical priority' (automation_priority=40). Therefore, resolution steps are generated only for critical automation opportunities. To include additional priorities, update the filter in the condition builder. The condition builder is accessible when you edit LEAP settings. Example: To generate resolution steps for both Critical and High priority automation opportunities, update the filter to select both Critical and High priority. 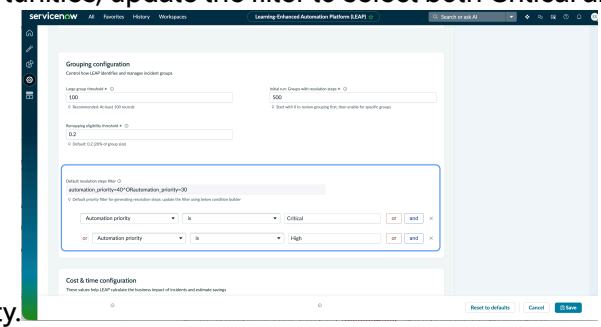

LEAP AI agent settings fields

LEAP AI agent settings

Field	Description
Problem creation agent	
Problem agent: Minimum incident count	Minimum number of incidents a cluster must contain before the LEAP Problem creation agent autonomously creates a Problem record. The default value is 15.
Problem agent: Severity	Automation opportunity severity levels that make a cluster eligible for problem record creation by the LEAP Problem creation agent. You can select multiple severity levels. The default values are Critical and High.
Knowledge base agent	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

LEAP AI agent settings (continued)

Field	Description
KB agent: Minimum incident count	Minimum number of incidents a cluster must contain before the LEAP AI agent autonomously creates a knowledge base article. The default value is 15.
KB agent: Severity	Automation opportunity severity levels that make a cluster eligible for knowledge base article creation by the LEAP AI agent. You can select multiple severity levels. The default values are Critical or High.
Knowledge base routing	
Default knowledge base	The knowledge base where articles created automatically by the LEAP AI agent are stored. This field is mandatory. You can't save settings without selecting a value. Articles created through the ServiceNow Otto panel are also routed to this knowledge base.
Default category	The category within the default knowledge base where articles created by the LEAP AI agent are stored. The available categories are filtered based on the selected default knowledge base.
Eligible knowledge bases	The knowledge bases that users can select from when manually creating a knowledge base article from an automation opportunity using **Actions** > **Draft KB article**. When specific knowledge bases are selected, only those appear in the knowledge base selection dialog. If this field is not configured, all active knowledge bases are shown.

LEAP AI indicators

AI indicators mark the artifacts that the LEAP AI agent creates. They appear on the home page list and on the automation opportunity details page.

Where AI indicators appear

The following table describes where the AI indicator appears and how it behaves.

AI indicator locations

Location	Behavior
Home page list	An AI icon AI column identifies the automation opportunities that have artifacts created by the LEAP AI agent. When you hover over the AI icon, a pop up card appears with information related to the agent-generated problem (PRB) and knowledge base (KB) records:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI indicator locations (continued)

Location	Behavior
	• Select the PRB link to open the record in the Service Operations Workspace. • Select kb link to open the article in the Knowledge Center.
Automation opportunity details page	The Action insights section for agent-generated artifact appears in a gradient style.

Automation opportunity details page

The automation opportunity details page displays resolution steps, activity, and action insights for a selected automation opportunity (AO). The page also indicates whether an AO is active or archived.

Page overview

The details page for an automation opportunity is divided into three panels:

- **Resolution steps** — displays AI-generated steps for resolving the incidents grouped in the AO. Steps carry a **Generated by AI** label and a generation timestamp. You can regenerate steps or revert to the original AI-generated version.
- **Compose** — use this panel to add comments visible to all users.
- **Action Insights** — displays contextual information about the AO, including its relationship to other AOs and any reanalysis alerts.

AO details fields

The following fields appear in the AO header.

AO details fields

Field	Description
Number	Unique identifier for the AO (for example, AOPP0000107).
Group type	The type of records grouped in the AO (for example, Incident).
Record count	The number of records associated with the AO.
Potential cost savings	The estimated monthly cost savings if the AO is actioned.
Potential time savings	The estimated monthly time savings if the AO is actioned.
Automation priority	The priority level assigned to the AO (for example, Critical or High).
Status	The current state of resolution step generation.

Resolution steps panel

The **Resolution steps** panel displays AI-generated steps to help resolve incidents in the AO. The panel includes the date and time the steps were generated.

The following actions are available on an active AO:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Regenerate** — generates a fresh set of resolution steps based on the latest record data.
- **Revert to AI generated steps** — restores the original AI-generated steps if you have manually edited them.

Warning: The generation of resolution steps involves processing your data through third-party providers, which routes traffic to a datacenter outside your region.

Resolution steps on an archived AO are read-only. The **Regenerate** and **Revert to AI generated steps** actions are not available on archived AOs.

Action Insights panel

The **Action Insights** panel displays contextual cards based on the state of the AO.

Action Insights cards

Card	Description
Archived AO	Appears on an archived AO that has no successor, or on an archived AO whose artifacts have been reassigned to a new AO. When a successor exists, links to the new active AO and includes the AO number and the timestamp of the reassignment.
New AO	Appears on a new AO that includes artifacts from an archived AO. Links to the archived AO and includes the AO number and the timestamp of the reassignment.
KB created	Appears when a knowledge base article has been created for the AO. Links to the KB article.
Problem created	Appears when a problem record has been created for the AO. Links to the problem record.
Record reanalysis status	Appears when the records associated with the AO have changed since the resolution steps were generated. Prompts you to regenerate the steps to reflect the latest data.

Banner messages

Banner messages appear at the top of the details page to indicate the AO's relationship to other AOs.

Banner messages

Message	When it appears
"This automation opportunity is archived."	Appears on an archived AO that has no successor AO yet.
"This automation opportunity is archived. Its artifacts were reassigned to: [AO link]"	Appears on an archived AO whose artifacts have been reassigned to a new active AO. The link navigates to the new AO.
"This automation opportunity includes artifacts from the archived AO: [AO link]"	Appears on a new AO that includes artifacts from an archived AO. The link navigates to the archived AO.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Knowledge base permissions for LEAP

To use a knowledge base with LEAP, the Can Contribute field must be set to allow users with the knowledge role to submit content. If this is not configured correctly, LEAP cannot route articles to the intended knowledge base.

LEAP routes knowledge base articles to the knowledge bases configured in LEAP properties. If the Can Contribute permission is not set correctly on a knowledge base, routing fails and the following error is displayed:

```
KB creation failed. Verify that the 'Can Contribute' field in the
selected knowledge base is set to 'Users with 'knowledge' role'.
```

Knowledge base permission requirements

Field	Location	Required value	Impact if not set
Can Contribute	**Knowledge** > **Administration** > **Knowledge Bases** > **Can Contribute tab** > **Collection**	Users with 'knowledge' role	LEAP cannot route articles to the knowledge base and displays a KB creation error.

Configure this permission on every knowledge base you plan to use in LEAP, including the default knowledge base and any knowledge bases configured in the **Eligible knowledge bases** field. Once configured, proceed to set up knowledge base routing in LEAP properties. See LEAP settings fields.

ITOM Visibility reference

Reference topics provide additional information about mapping and fine-tuning application services using ITOM Visibility lists and forms.

Data collected by ITOM Visibility

ITOM Visibility collects unique data for each type of device and stores it in dedicated tables, fields, and relationships.

Operating system

Operating system

Find servers and computers based on the machine's operating system:

- Windows
- Solaris
- Linux
- Mac (OS/X) discovery
- HP-UX
- AIX

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtualization

Virtualization

Collect information on deployments that include virtualization:

- Hyper-V discovery
- IBM Virtualization
- Red Hat Virtualization
- Nutanix Acropolis
- VMware vCenter
- OpenStack Virtualization

Cloud

Cloud

- AWS
- Microsoft Azure
- Google Cloud Platform (GCP)
- IBM cloud
- VMware cloud

Containers

Containers

- Kubernetes and OpenShift
- Amazon ECS resource discovery with Patterns
- Pivotal Cloud Foundry discovery

Software

Software

Find a variety of software, including:

- Apigee Edge
- ColdFusion
- HP Operations Manager

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- IBM Websphere DataPower
- JBOSS on Windows and Linux
- Web and email servers, including Microsoft IIS, Apache, Tomcat, Exchange mailbox
- Puppet automation software
- Red Hat JBoss Fuse
- SAP applications
- WebLogic and WebSphere servers
- General software

Clustered applications

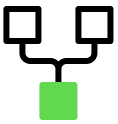

Clustered applications

- IBM PowerHA Cluster (HACMP)
- Linux Red Hat cluster
- Veritas Cluster Server discovery
- Windows server cluster

Databases

Databases

Discover databases, including:

- Oracle
- MySQL
- Microsoft SQL
- MongoDB
- SAP HANA
- PostgreSQL
- Sybase
- Amazon DynamoDB
- Oracle Database 12c

Network devices

Network devices

- Routers and switches
- Load balancers:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- A10
- F5 BIG-IP
- Cisco GSS
- Cisco CSS
- Citrix NetScaler
- HAProxy
- NGINX
- Alteon
- ACE
- Radware-appDirector
- AWS application ELB Service discovery with Patterns

- Firewalls
  - Cisco firewall
  - Fortinet firewall
  - Juniper firewall
  - Palo Alto firewall
- IP networks and specific IP addresses
- Cisco UCS devices
- Cisco Switch Wireless Access Point (WAP)
- Printers and power supplies
- IP services and daemons

You can also find devices based on TCP connections and Layer-2, SNMP-level discovery.

Storage

Storage

Discover these types of storage devices:

- Pure Storage FlashBlade
- NetApp server
- Discovery of storage area networks (SAN)
- Direct Attached Storage (DAS)
- Network Attached Storage (NAS)
- Host bus adapters (HBA)
- EMC Isilon

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Detailed information on products discovered by ITOM Visibility

Discovery and Service Mapping can discover a wide range of operating systems and applications.

Discovery finds computers, servers, printers, a variety of IP-enabled devices, and the applications that run on them. It can then update the configuration items (CIs) in your Configuration Management Database (CMDB) with the data it collects. This discovery method is referred to as horizontal discovery. Service Mapping maps dependencies, based on a connection between devices and applications. This method is referred to as top-down mapping. The top-down mapping helps you immediately see the impact of a problematic object on the rest of the service instance operation.

On top of hosts and applications supported by default, you can discover additional hosts and applications by deploying patterns available on Store. For reference information on store released patterns, see Available on-premise discovery patterns.

If your organization uses devices or applications, which are not supported by default or using patterns available at ServiceNow Store, you can configure Discovery and Service Mapping to discover them as described in Discovery patterns used by ITOM Visibility.

Cloud Discovery Patterns find the cloud resources of AWS, Azure, Google Cloud Platform (GCP), IBM, and Oracle.

If you want to validate the necessary pattern commands before running discovery, use the Command Validation Tool. For more information, see Validate commands used in pattern-based discovery.

ITOM Content Service Provides visibility to your applications by using AI capabilities that cluster and classify running application processes. For more information, see ITOM Content Service.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
.Net Framework	Windows	1.x.x, 2.x.x, 3.x.x, 4.x.x	WMI WinRM	Top-down	.NET Application
A10 load balancer	A10	N/A	REST	Horizontal	A10 Load Balancer
Adobe JRun	Windows UNIX	4.x	WMI WinRM SSH	Horizontal and top-down	Jrun

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Apache Cassandra database and DataStax	UNIX	N/A	SSH	Horizontal and top-down	Cassandra
Apache HBase	UNIX	N/A	SSH	Horizontal	Probes: • HBase - Version • HBase - Main Class • HBase - Configuration
Apache HTTP Server	Windows UNIX	2.x	WMI WinRM SSH	Horizontal and top-down	Apache
Apache Kafka and Zookeeper	Apache Kafka	N/A	SSH	Horizontal	Kafka and Zookeeper
Apache Tomcat Servlet container HTTP web server	Windows UNIX	6.x, 7.x, 8.x	WMI WinRM SSH	Horizontal and top-down	Tomcat
Apigee Edge Enterprise edition	Linux	4.x.x	SSH HTTP	Horizontal and top-down	APIGee
Avi Vantage load balancer including Avi Controller and GSLB	VMware AWS Avi Networks	N/A	REST	Horizontal and top-down	AVI LB - Avi Vantage load balancer AVI LB - AVI Top Down AVI Session - Avi Controller AVI GSLB support
Alibaba Cloud availability zone	Alibaba Cloud	N/A	REST	Horizontal	Alibaba - Availability Zone (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Alibaba Cloud cloud hardware type	Alibaba Cloud	N/A	REST	Horizontal	Alibaba - Cloud Hardware Type (LP)
Alibaba Cloud cloud OS images	Alibaba Cloud	N/A	REST	Horizontal	Alibaba - Cloud OS Image (LP)
Alibaba Cloud datacenters	Alibaba Cloud	N/A	REST	Horizontal	Alibaba - Datacenter (LP)
Alibaba Cloud service accounts	Alibaba Cloud	N/A	REST	Horizontal	Alibaba - Service Account Validation
Alibaba Cloud storage volume	Alibaba Cloud	N/A	REST	Horizontal	Alibaba - Storage Volume (LP)
Alibaba Cloud virtual machines	Alibaba Cloud	N/A	REST	Horizontal	Alibaba - Virtual Machine (LP)
Amazon API Gateway	AWS	N/A	REST	Horizontal and top-down	Amazon AWS API Gateway
Amazon API Gateway Domain Name	AWS	N/A	REST	Horizontal	Amazon AWS - API Gateway Domain Name - Extended Inventory (LP)
Amazon Application Load Balancer Service	AWS	N/A	REST	Top-down	Amazon AWS application ELB service - TD
Amazon Athena Workgroups	AWS	N/A	REST	Horizontal	Amazon AWS - Athena Workgroup - Extended Inventory (LP)
Amazon Bedrock	AWS	N/A	REST	Horizontal	Patterns: • Amazon AWS - Bedrock Agents • Amazon AWS -

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					Bedrock Foundation Model
Amazon CloudFront Distribution	AWS	N/A	REST	Horizontal	Amazon AWS - CloudFront Distribution - Extended Inventory (LP)
Amazon CloudWatch Log Group	AWS	N/A	REST	Horizontal	Amazon AWS - CloudWatch Log Group - Extended Inventory (LP)
Amazon Cognito	AWS	N/A	REST	Horizontal	Amazon AWS Cognito
Amazon DB cluster discovery with Patterns	AWS	N/A	REST	Horizontal and top-down	Amazon AWS DB Cluster
Amazon DynamoDB	AWS	N/A	REST	Horizontal and top-down	Amazon AWS DynamoDB (pattern)
Amazon DynamoDB Cluster	AWS	N/A	REST	Horizontal	Amazon AWS - DynamoDB Cluster - Extended Inventory (LP)
Amazon EC2 Amazon EBS Snapshot	AWS	N/A	REST	Horizontal	Amazon AWS - Elastic Compute Cloud EBS Snapshot - Extended Inventory (LP)
Amazon EC2 Reserved Instance	AWS	N/A	REST	Horizontal	Amazon AWS - Elastic Compute Cloud Reserved Instance - Extended Inventory (LP)
Amazon EC2 VPC Endpoint Service	AWS	N/A	REST	Horizontal	Amazon AWS - Elastic Compute Cloud VPC Endpoint

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					Service - Extended Inventory (LP)
Amazon EC2 VPC Peering Connection	AWS	N/A	REST	Horizontal	Amazon AWS - Elastic Compute Cloud VPC Peering Connection - Extended Inventory (LP)
Amazon ECS resource	AWS	N/A	REST	Horizontal	AWS ECS/ Fargate
Amazon Elastic File System (Amazon EFS)	AWS	N/A	REST	Horizontal	Amazon AWS - Elastic File System - Extended Inventory (LP)
Amazon ElastiCache	AWS	N/A	REST	Horizontal	Amazon ElastiCache discovery
Amazon ElastiCache Snapshot	AWS	N/A	REST	Horizontal	Amazon AWS - ElastiCache Snapshot - Extended Inventory (LP)
Amazon EMR Cluster	AWS	N/A	REST	Horizontal	Amazon AWS - EMR Cluster - Extended Inventory (LP)
Amazon EventBridge Event Bus	AWS	N/A	REST	Horizontal	Amazon AWS - EventBridge Event Bus - Extended Inventory (LP)
Amazon FSx Backup	AWS	N/A	REST	Horizontal	Amazon AWS - FSx Backup - Extended Inventory (LP)
Amazon FSx File System	AWS	N/A	REST	Horizontal	Amazon AWS - FSx File System - Extended Inventory (LP)
Amazon MQ Broker	AWS	N/A	REST	Horizontal	Amazon AWS - MQ Broker

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					- Extended Inventory (LP)
Amazon MQ Configuration	AWS	N/A	REST	Horizontal	Amazon AWS - MQ Configuration - Extended Inventory (LP)
Amazon MWAA Environment	AWS	N/A	REST	Horizontal	Amazon AWS - Managed Workflows for Apache Airflow Environment - Extended Inventory (LP)
Amazon Relational Database Service	AWS	N/A	REST	Horizontal and top-down	Amazon AWS Relational Database Service
Amazon RDS DB Snapshot	AWS	N/A	REST	Horizontal	Amazon AWS - Relational Database Services DB Snapshot - Extended Inventory (LP)
Amazon Redshift	AWS	N/A	REST	Horizontal	Amazon AWS Redshift
Amazon Redshift Serverless Namespace	AWS	N/A	REST	Horizontal	Amazon AWS - Redshift Serverless Namespace - Extended Inventory (LP)
Amazon Redshift Serverless Snapshot	AWS	N/A	REST	Horizontal	Amazon AWS - Redshift Serverless Snapshot - Extended Inventory (LP)
Amazon Redshift Serverless Workgroup	AWS	N/A	REST	Horizontal	Amazon AWS - Redshift Serverless Workgroup - Extended Inventory (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Amazon Route 53	AWS	N/A	REST	Horizontal	Amazon AWS Route53
Amazon S3	AWS	N/A	REST	Horizontal and top-down	AWS S3
Amazon S3 Glacier Vault	AWS	N/A	REST	Horizontal	Amazon AWS - S3 Glacier Vault - Extended Inventory (LP)
Amazon SageMaker Notebook Instance	AWS	N/A	REST	Horizontal	Amazon AWS - SageMaker Notebook Instance - Extended Inventory (LP)
Amazon SageMaker Training Job	AWS	N/A	REST	Horizontal	Amazon AWS - SageMaker Training Job - Extended Inventory (LP)
Amazon Simple Email Service Identity	AWS	N/A	REST	Horizontal	Amazon AWS - Simple Email Service Identity - Extended Inventory (LP)
Amazon SQS Queue	AWS	N/A	REST	Horizontal	Amazon AWS - SQS Queue - Extended Inventory (LP)
Amazon Timestream for InfluxDB Database Instance	AWS	N/A	REST	Horizontal	Amazon AWS - Timestream for InfluxDB Database Instance - Extended Inventory (LP)
Amazon VPC Flow Log	AWS	N/A	REST	Horizontal	Amazon AWS - VPC Flow Log - Extended Inventory (LP)
Amazon VPC Managed Prefix List	AWS	N/A	REST	Horizontal	Amazon AWS - VPC Managed Prefix List

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					- Extended Inventory (LP)
AWS Application and Network LB	AWS	N/A	REST	Horizontal	Amazon AWS - Application and Network LB (LP)
AWS application ELB Service	AWS	N/A	REST	Horizontal	Amazon AWS application ELB Service
AWS AppSync API	AWS	N/A	REST	Horizontal	Amazon AWS - AppSync API - Extended Inventory (LP)
AWS Auto Scaling groups discovery with Patterns	AWS	N/A	REST	Horizontal	AWS Auto Scaling groups (LP) (pattern)
AWS Availability Zone pattern-based discovery	AWS	N/A	REST	Horizontal	Amazon AWS - Availability Zone (LP)
AWS Backup Plan	AWS	N/A	REST	Horizontal	Amazon AWS - Backup Backup Plan - Extended Inventory (LP)
AWS Backup Vault	AWS	N/A	REST	Horizontal	Amazon AWS - Backup Backup Vault - Extended Inventory (LP)
AWS Batch Compute Environment	AWS	N/A	REST	Horizontal	Amazon AWS - Batch Compute Environment - Extended Inventory (LP)
AWS classic ELB Service	AWS	N/A	REST	Horizontal	Amazon AWS classic ELB Service
AWS Classic LB	AWS	N/A	REST	Horizontal	Amazon AWS - Classic LB (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
AWS CloudHSM HSM	AWS	N/A	REST	Horizontal	Amazon AWS - CloudHSM HSM - Extended Inventory (LP)
AWS CloudTrail Trail	AWS	N/A	REST	Horizontal	Amazon AWS - CloudTrail Trail - Extended Inventory (LP)
AWS CodeDeploy Deployment	AWS	N/A	REST	Horizontal	Amazon AWS - CodeDeploy Deployment - Extended Inventory (LP)
AWS CodePipeline Pipeline	AWS	N/A	REST	Horizontal	Amazon AWS - CodePipeline Pipeline - Extended Inventory (LP)
AWS Customer Gateway	AWS	N/A	REST	Horizontal	Amazon AWS - Customer Gateway (LP)
AWS datacenter	AWS	N/A	REST	Horizontal	Amazon AWS Datacenter discovery
AWS DataSync Task	AWS	N/A	REST	Horizontal	Amazon AWS - DataSync Task - Extended Inventory (LP)
AWS Cloud Organizations	AWS	N/A	REST	Horizontal	Amazon AWS - discover Organization
AWS DMS Endpoints	AWS	N/A	REST	Horizontal	Amazon AWS - DMS Endpoint - Extended Inventory (LP)
AWS Elastic Beanstalk Application	AWS	N/A	REST	Horizontal	Amazon AWS - Elastic Beanstalk Application - Extended Inventory (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
AWS Elastic Load Balancer Service	AWS	N/A	REST	Top-down	Amazon AWS classic ELB Service - TD
AWS Events	AWS	N/A	REST	Horizontal	Patterns: • Amazon AWS Virtual Server Events • Amazon AWS Security Group Events • Amazon AWS Subnet Events • Amazon AWS Storage Events • Amazon AWS Network Events • Amazon AWS Classic LB Events • Amazon AWS Application and Network LBs Events
AWS Global Accelerator	AWS	N/A	REST	Horizontal	Amazon AWS - Global Accelerator - Extended Inventory (LP)
AWS Glue Database	AWS	N/A	REST	Horizontal	Amazon AWS - Glue Database - Extended Inventory (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
AWS hardware type	AWS	N/A	REST	Horizontal	Patterns: • Amazon AWS - Hardware Type (LP) • Amazon AWS - Cloud Hardware Type (LP)
AWS Host	AWS	N/A	REST	Horizontal	Amazon AWS - Host (LP)
AWS IAM Policy	AWS	N/A	REST	Horizontal	Amazon AWS - IAM Policy - Extended Inventory (LP)
AWS IAM Role	AWS	N/A	REST	Horizontal	Amazon AWS - IAM Role - Extended Inventory (LP)
AWS IAM User	AWS	N/A	REST	Horizontal	Amazon AWS - IAM User - Extended Inventory (LP)
AWS Internet Gateway	AWS	N/A	REST	Horizontal	Amazon AWS - Internet Gateway (LP)
AWS IP Address	AWS	N/A	REST	Horizontal	Amazon AWS - IP Address (LP)
AWS Keyspaces	AWS	N/A	REST	Horizontal	Amazon AWS Keyspaces
AWS Key Pair	AWS	N/A	REST	Horizontal	Amazon AWS - Key Pair (LP)
AWS Kinesis Discovery	AWS	N/A	REST	Horizontal	Patterns: • Amazon AWS - Kinesis Video Stream Services (LP) • Amazon AWS - Kinesis

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					Data Stream Services (LP) • Amazon AWS - Kinesis Firehose Stream Services (LP) • Amazon AWS - Kinesis Data Analytics Stream Services (LP)
AWS KMS Key	AWS	N/A	REST	Horizontal	Amazon AWS - KMS Key - Extended Inventory (LP)
AWS Lambda discovery with Patterns	AWS	N/A	REST	Horizontal and top-down	Amazon AWS Lambda
AWS LB Pool	AWS	N/A	REST	Horizontal	Amazon AWS - LB Pool (LP)
AWS LB Pool Member	AWS	N/A	REST	Horizontal	Amazon AWS - LB Pool Member(LP)
AWS LB Service	AWS	N/A	REST	Horizontal	Amazon AWS - LB Service (LP)
AWS Linux Server	AWS	N/A	REST	Horizontal	Amazon AWS - Linux Server (LP)
AWS Marketplace	AWS	N/A	REST	Horizontal	Amazon AWS - Marketplace (LP)
AWS MemoryDB for Redis discovery with Patterns	AWS	N/A	REST	Horizontal	AWS MemoryDB (pattern)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
AWS NAT Gateway	AWS	N/A	REST	Horizontal	Amazon AWS - NAT Gateway (LP)
AWS Network	AWS	N/A	REST	Horizontal	Amazon AWS - Network (LP)
AWS Network ACL	AWS	N/A	REST	Horizontal	Amazon AWS - ACL (LP)
AWS Network Firewall	AWS	N/A	REST	Horizontal	Amazon AWS - Network Firewall - Extended Inventory (LP)
AWS NIC	AWS	N/A	REST	Horizontal	Amazon AWS - NIC (LP)
AWS OpenSearch	AWS	N/A	REST	Horizontal	Amazon AWS OpenSearch
AWS Organizations	AWS	N/A	REST	Horizontal	Amazon AWS Organizations
AWS Organizational Units	AWS	N/A	REST	Horizontal	Amazon AWS - Organizational Units (LP)
AWS OS images	AWS	N/A	REST	Horizontal	Patterns: • Amazon AWS - Owned Template (LP) • Amazon AWS - Executable Template (LP) • Amazon AWS - Owned Cloud OS Image (LP) • Amazon AWS - Executable Cloud OS Image (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
AWS Public IP Address	AWS	N/A	REST	Horizontal	Amazon AWS - Public IP Address (LP)
AWS Resource Inventory discovery with Patterns	AWS	N/A	REST	Horizontal	AWS Resource Inventory
AWS Route Table	AWS	N/A	REST	Horizontal	Amazon AWS - Route Table (LP)
AWS Secret Manager Secret	AWS	N/A	REST	Horizontal	Amazon AWS - Secrets Manager Secret - Extended Inventory (LP)
AWS Security Group	AWS	N/A	REST	Horizontal	Amazon AWS - Security Group (LP)
AWS Serverless Database	AWS	N/A	REST	Horizontal	Amazon AWS Serverless Database
AWS SSM Cloud Agents	AWS	N/A	REST	Horizontal	Amazon AWS - SSM Cloud Agents (LP)
AWS Step Functions State Machine	AWS	N/A	REST	Horizontal	Amazon AWS - Step Functions State Machine - Extended Inventory (LP)
AWS Storage	AWS	N/A	REST	Horizontal	Amazon AWS - Storage (LP)
AWS Storage Gateway File Share	AWS	N/A	REST	Horizontal	Amazon AWS - Storage Gateway File Share - Extended Inventory (LP)
AWS Storage Gateway Gateway	AWS	N/A	REST	Horizontal	Amazon AWS - Storage Gateway Gateway - Extended Inventory (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
AWS sub account Organizations	AWS	N/A	REST	Horizontal	Amazon AWS - Sub Account (LP)
AWS sub account regional	AWS	N/A	REST	Horizontal	Amazon AWS sub account discovery
AWS subnet	AWS	N/A	REST	Horizontal	Amazon AWS - Subnet (LP)
AWS Systems Manager Document	AWS	N/A	REST	Horizontal	Amazon AWS - Systems Manager Document - Extended Inventory (LP)
AWS Systems Manager Parameter Store	AWS	N/A	REST	Horizontal	Amazon AWS - Systems Manager Parameter Store - Extended Inventory (LP)
AWS Transfer Family Server	AWS	N/A	REST	Horizontal	Amazon AWS - Transfer Family Server - Extended Inventory (LP)
AWS Virtual Server	AWS	N/A	REST	Horizontal	Amazon AWS - Virtual Server (LP)
AWS VPN Connections	AWS	N/A	REST	Horizontal	Amazon AWS - VPN Connections (LP)
AWS VPN Gateway	AWS	N/A	REST	Horizontal	Amazon AWS - VPN Gateway (LP)
AWS web ACL	AWS	N/A	REST	Horizontal	Amazon AWS - Web ACL (LP)
AWS Windows Server	AWS	N/A	REST	Horizontal	Amazon AWS - Windows Server (LP)
AWS X-Ray Sampling Rule	AWS	N/A	REST	Horizontal	Amazon AWS - XRay Sampling Rule

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					- Extended Inventory (LP)
Azure App Configuration store	Microsoft Azure	N/A	REST	Horizontal	Azure - App Configuration Configuration Store - Extended Inventory(LP)
Azure App Service App Service Plan	Microsoft Azure	N/A	REST	Horizontal	Azure - App Service App Service Plan - Extended Inventory(LP)
Azure Application Gateway	Microsoft Azure	N/A	REST	Horizontal	Azure Application Gateway (LP)
Azure Application Insight Component	Microsoft Azure	N/A	REST	Horizontal	Azure - Application Insight Component - Extended Inventory(LP)
Azure Application Insight Data Collection Rule	Microsoft Azure	N/A	REST	Horizontal	Azure - Application Insight Data Collection Rule - Extended Inventory(LP)
Azure Application LB	Microsoft Azure	N/A	REST	Horizontal	Azure - Application LB (LP)
Azure Application Security Group	Microsoft Azure	N/A	REST	Horizontal	Azure - Application Security Group - Extended Inventory(LP)
Azure Automation Account	Microsoft Azure	N/A	REST	Horizontal	Azure - Automation Account - Extended Inventory(LP)
Azure availability sets	Microsoft Azure	N/A	REST	Horizontal	Azure - Availability Set (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Azure availability zones	Microsoft Azure	N/A	REST	Horizontal	Azure - Availability Zones (LP)
Azure Blob Storage	Microsoft Azure	N/A	REST	Horizontal	Azure - Storage Blobs(LP)
Azure Classic Load Balancer	Microsoft Azure	N/A	REST	Horizontal	Azure - Classic LB (LP)
Azure Compute Gallery Image Definition	Microsoft Azure	N/A	REST	Horizontal	Azure - Compute Gallery Image Definition - Extended Inventory(LP)
Azure Compute Snapshot	Microsoft Azure	N/A	REST	Horizontal	Azure - Compute Snapshot - Extended Inventory(LP)
Azure Container Registry	Microsoft Azure	N/A	REST	Horizontal	Azure - Container Registry - Extended Inventory(LP)
Azure Cosmos DB for PostgreSQL Cluster	Microsoft Azure	N/A	REST	Horizontal	Azure - Cosmos DB for PostgreSQL Cluster - Extended Inventory(LP)
Azure Data Explorer Cluster	Microsoft Azure	N/A	REST	Horizontal	Azure - Data Explorer Cluster - Extended Inventory(LP)
Azure Data Factory	Microsoft Azure	N/A	REST	Horizontal	Azure - Data Factory - Extended Inventory(LP)
Azure Data Protection Backup Vault	Microsoft Azure	N/A	REST	Horizontal	Azure - Data Protection Backup Vault - Extended Inventory(LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Azure Database	Microsoft Azure	N/A	REST	Horizontal and top-down	• Azure DataBase (LP) • Azure DataBase TD
Azure Databricks Workspace	Microsoft Azure	N/A	REST	Horizontal	Azure - Databricks Workspace - Extended Inventory(LP)
Azure Datacenter discovery	Microsoft Azure	N/A	REST	Horizontal	Azure Datacenter discovery
Azure Dev Center	Microsoft Azure	N/A	REST	Horizontal	Azure - Dev Center - Extended Inventory(LP)
Azure Disk Encryption Set	Microsoft Azure	N/A	REST	Horizontal	Azure - Disk Encryption Set - Extended Inventory(LP)
Azure DNS zones and record sets discovery using Patterns	Microsoft Azure	N/A	REST	Horizontal	Patterns: Azure DNS Zones HD or Azure DNS Zones (LP) Azure DNS Zone Recordsets (LP)
Azure Event Grid System Topic	Microsoft Azure	N/A	REST	Horizontal	Azure - Event Grid System Topic - Extended Inventory(LP)
Azure Event Hub Namespace	Microsoft Azure	N/A	REST	Horizontal	Azure - Event Hub Namespace - Extended Inventory(LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Azure Events	Microsoft Azure	N/A	REST	Horizontal	Patterns: • Azure Application LB Event • Azure Availability Set Event • Azure Classic LB Event • Azure DataBase Event • Azure Express Route Circuit Event • Azure Functions Event • Azure Local Network Gateway Event • Azure NAT Gateway Event • Azure Network Event • Azure NIC Event • Azure Private DNS Zone Event • Azure Public IP Event • Azure Resource Group Event • Azure Security Group Event

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					• Azure Storage Account Event • Azure Virtual Machine Event • Azure Virtual Network Gateway Connection Event • Azure Virtual Network Peerings Event • Azure VM Scale Set Event
Azure Express Route Circuit	Microsoft Azure	N/A	REST	Horizontal	Azure - Express Route Circuit (LP)
Azure File Share	Microsoft Azure	N/A	REST	Horizontal	Azure - File Share (LP)
Azure Firewall Network Security	Microsoft Azure	N/A	REST	Horizontal	Azure - Network Security Azure Firewall - Extended Inventory(LP)
Azure Functions	Microsoft Azure	N/A	REST	Horizontal and top-down	Patterns: Azure Functions TD or Azure Functions (LP)
Azure hardware type	Microsoft Azure	N/A	REST	Horizontal and top-down	Pattens:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					• Azure - Hardware Type (LP) • Azure - Cloud Hardware Type (LP)
Azure Host	Microsoft Azure	N/A	REST	Horizontal	Azure - Host (LP)
Azure Key Vault	Microsoft Azure	N/A	REST	Horizontal	Azure - Key Vault - Extended Inventory(LP)
Azure Key Vault Key	Microsoft Azure	N/A	REST	Horizontal	Azure - Key Vault Key - Extended Inventory(LP)
Azure LoadBalancer Service TD	Microsoft Azure	N/A	REST	Top-down	Azure LoadBalancer TD
Azure Local Network Gateway	Microsoft Azure	N/A	REST	Horizontal	Azure - Local Network Gateway (LP)
Azure Log Analytics Workspace	Microsoft Azure	N/A	REST	Horizontal	Azure - Log Analytics Workspace - Extended Inventory(LP)
Azure Logic App	Microsoft Azure	N/A	REST	Horizontal	Azure - Logic App - Extended Inventory(LP)
Azure Managed Identity User Assigned Identity	Microsoft Azure	N/A	REST	Horizontal	Azure - Managed Identity User Assigned Identity - Extended Inventory(LP)
Azure Marketplace	Microsoft Azure	N/A	REST	Horizontal	Azure - Marketplace LB (LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Azure NAT Gateway	Microsoft Azure	N/A	REST	Horizontal	Azure - NAT Gateway (LP)
Azure Network and Subnet	Microsoft Azure	N/A	REST	Horizontal	Azure - Network and Subnet (LP)
Azure Networks IP Group	Microsoft Azure	N/A	REST	Horizontal	Azure - Networks IP Group - Extended Inventory(LP)
Azure NIC	Microsoft Azure	N/A	REST	Horizontal	Azure - NIC (LP)
Azure OS image	Microsoft Azure	N/A	REST	Horizontal	Patterns: • Azure - Image (LP) • Azure - Cloud OS Image (LP)
Azure Private DNS Zone	Microsoft Azure	N/A	REST	Horizontal	Azure - Private DNS Zone (LP)
Azure Private Gateway	Microsoft Azure	N/A	REST	Horizontal	Azure - Private Gateway (LP)
Azure Private Link Private Endpoint	Microsoft Azure	N/A	REST	Horizontal	Azure - Private Link Private Endpoint - Extended Inventory(LP)
Azure Public IP	Microsoft Azure	N/A	REST	Horizontal	Azure - Public IP (LP)
Azure Recovery Services Vault	Microsoft Azure	N/A	REST	Horizontal	Azure - Recovery Services Vault - Extended Inventory(LP)
Azure Recovery Services Vault Backup Item	Microsoft Azure	N/A	REST	Horizontal	Azure - Recovery Services Vault Backup Item

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					- Extended Inventory(LP)
Azure Resource Group	Microsoft Azure	N/A	REST	Horizontal	Azure - Resource Group (LP)
Azure Resource Inventory	Microsoft Azure	N/A	REST	Horizontal	Azure Inventory and tags
Azure Route Table	Microsoft Azure	N/A	REST	Horizontal	Azure - Route Table (LP)
Azure Security Group	Microsoft Azure	N/A	REST	Horizontal	Azure - Security Group (LP)
Azure Service Bus Namespace	Microsoft Azure	N/A	REST	Horizontal	Azure - Service Bus Namespace - Extended Inventory(LP)
Azure Service Bus Queue	Microsoft Azure	N/A	REST	Horizontal	Azure - Service Bus Queue - Extended Inventory(LP)
Azure Service Bus Topic	Microsoft Azure	N/A	REST	Horizontal	Azure - Service Bus Topic - Extended Inventory(LP)
Azure Service Endpoint Policy	Microsoft Azure	N/A	REST	Horizontal	Azure - Service Endpoint Policy - Extended Inventory(LP)
Azure SQL Server	Microsoft Azure	N/A	REST	Horizontal	Azure - SQL Server (LP)
Azure Storage Account	Microsoft Azure	N/A	REST	Horizontal	Azure - Storage Account (LP)
Azure Storage Container	Microsoft Azure	N/A	REST	Horizontal	Azure - Storage Container(LP)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Azure Storage Volume	Microsoft Azure	N/A	REST	Horizontal	Azure - Storage Volume (LP)
Azure sub account	Microsoft Azure	N/A	REST	Horizontal	Azure - Sub Account (LP)
Azure Subscriptions Discovery For Management Group	Microsoft Azure	N/A	REST	Horizontal	Azure Subscriptions Discovery For Management Group
Azure Virtual Machine	Microsoft Azure	N/A	REST	Horizontal	Azure - Virtual Machine (LP)
Azure Virtual Machine Scale Sets (VMSS) Instance discovery	Microsoft Azure	N/A	REST	Horizontal	Patterns: • Azure - VM Scale Set (LP) • Azure VM Instance - Uniform Scale Set
Azure Virtual Network Gateway Connection	Microsoft Azure	N/A	REST	Horizontal	Azure - Virtual Network Gateway Connection (LP)
Azure Web Application Firewall Policy	Microsoft Azure	N/A	REST	Horizontal	Azure - Web Application Firewall Policy - Extended Inventory(LP)
Azure WebSite	Microsoft Azure	N/A	REST	Horizontal and top-down	Patterns: • Azure WebSite (LP) • Azure WebSite TD
BMC Control-M	UNIX Windows	6.x	WMI WinRM SSH	Horizontal	Cntl-M Enterprise Manager

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Bring Your Own License (BYOL)	Amazon AWS Cloud Microsoft Azure Cloud Google Cloud Platform (GCP)	N/A	REST Azure Resource Graph	Horizontal	The “Image License” shared library is an extension of the patterns: • Amazon AWS - Executable Template (LP) • Amazon AWS - Owned Template (LP) The "Azure Virtual Machine License" shared library is an extension of the "Azure - Virtual Machine (LP)" pattern. The "Google Cloud Platform (GCP) – VM license" shared library is an extension of the "Google Cloud Platform (GCP) - Virtual Server" pattern.
CA eTrust Directory server	Windows	7.x, 8.x	WMI WinRM	Top-down	CA eTrust Directory server
CA Identity Manager	Windows	11.x, 12.x	WMI	Top-down	CA Identity Manager Provisioning

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Provisioning Server			WinRM		Server for Windows
CA Introscope Enterprise Manager	Windows Unix	8.x, 9.x	WMI WinRM SSH	Top-down	CA Introscope Enterprise Manager
CA Policy Server	Windows Unix	11.x	WMI WinRM SSH	Top-down	CA Policy Server
CA Site Minder Agent	Windows Unix	11.x	WMI WinRM SSH	Horizontal and top-down	CA Site Minder Agent
Cisco Content Services Switch Load balancer	ACE	6.x	SNMP SSH	Horizontal and top-down	Cisco CSS SNMP
Cisco ACE Application Control Engine	ACE	2.x	SSH	Horizontal and top-down	ACE Load Balancer
Cisco Call Manager	Cisco	3.x, 4.x	SNMP SSH	Top-down	Cisco CallManager
Cisco Fiber Connect	Cisco	No data	SNMP SSH	Top-down	Cisco Fibre Connect
Cisco Global Site Selector Load Balancer	ACE	3.x	SNMP SSH	Horizontal and top-down	Cisco GSS
Cisco Unified Computing System	UCS	3.x	REST	Horizontal	UCS - HD
Citrix Delivery Controller	Windows	7.x>7.5, 8.x	WMI WinRM	Top-down	Delivery Controller

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Citrix Netscaler Global Server Load Balancing (GSLB)	Netscaler	10.x	SNMP SSH NS.conf	Horizontal and top-down	Netscaler GLB
Citrix Netscaler Load Balancer	Netscaler	9.x, 10.x	SNMP SSH NS.conf	Horizontal and top-down	Citrix Netscaler
Citrix Netscaler SDX	Citrix NetScaler SDX	13	SSH	Horizontal	Citrix NetScaler SDX
Citrix License Server	Windows	7.x	WMI	Horizontal	License Server
Citrix Xen Hyper-V	Xen Server	7.6	SSH	Horizontal	Citrix Xen Hyper-V
Cloudian Storage	Linux	7.x	SSH	Horizontal	Cloudian Storage discovery
Connect APK	Windows	1.x.x	WMI WinRM	Top-down	Connect-It Service
Container image scanning for software decomposition	Linux		REST HTTP	Horizontal	container image scan
Software Bill of Material (SBOM)	Linux	N/A	REST HTTP	Horizontal	SBOM is an extension section of the container image scan pattern
Docker virtualization	Linux	N/A	N/A	Horizontal	Docker (pattern)
Epic Systems Corporation	UNIX	2014.x.x, 2015.x.x	SSH	Top-down	EPIC Cache
Database Administrator (DBA) report discovery	N/A	N/A	N/A	Horizontal	DBA is an extension

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					section of the patterns: • Cassandra DB • MSSQL DB • MySQL DB • Oracle DB
Dell PowerMax storage discovery with Patterns	Dell EMC	9 and 10	N/A	Horizontal	EMC PMAX phase1 (pattern) EMC PMAX phase2 (pattern)
Dell Data Domain storage discovery using Patterns	DD OS	DD OS 7.3 or later	REST	Horizontal	DELL EMC Data Domain (pattern)
Dell EMC XtremIO storage array discovery	Linux	6.3.0 and 6.3.1	REST	Horizontal	EMC XtremIO (pattern)
EMC Isilon	OneFS	8.0	SNMP REST	Horizontal	EMC Isilon
F5 BIG-IP Device Service Clustering	F5	11.x, 12.x	SNMP SSH REST	Horizontal	F5 Cluster
F5 BIG-IP	F5	11.x, 12.x	REST	Horizontal and top-down	F5 Load Balancer, F5 Load Balancer SSH
Fortinet firewalls	Fortinet Firewall device	FortiOS 5.2.x, FortiOS 5.4.x, FortiOS 5.6.x, FortiOS 6.2.x, FortiOS 6.4.x, FortiOS 7.2.x	SNMP	Horizontal	Next Generation Fortinet Network Firewall
Fortinet firewalls and	Fortinet Firewall device	FortiOS 6.2.x, FortiOS 6.4.x, FortiOS 7.x	REST	Horizontal	Next Generation Fortinet

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
FortiGate VDOMs					Network Firewall - REST
Google Cloud Apigee API Platform	Linux	4.x.x	SSH	Horizontal and top-down	APIGee Service
Google Cloud Platform (GCP) AlloyDB for PostgreSQL	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - AlloyDB for PostgreSQL
Google Cloud BigQuery	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - BigQuery DB
Google Cloud Bigtable	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Bigtable DB
Google Cloud FireStore	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Firestore DB
Google Cloud Platform (GCP) Cloud Functions	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP)- Cloud Functions
Google Cloud SQL	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Cloud SQL DB
Google Cloud Platform (GCP) Disk Types	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Disk Types
Google Cloud Platform (GCP) Events	GCP	N/A	REST	Horizontal	• Google Cloud Platform (GCP) - Networking - Events • Google Cloud Platform (GCP) - Networking

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					Firewall - Events • Google Cloud Platform (GCP) - Subnetwork - Events
Google Cloud Platform (GCP) External IP Addresses	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - External IP Addresses
Google Firebase Realtime DB	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Firebase Realtime DB
Google Cloud Platform (GCP) Load Balancer	GCP	N/A	REST	Horizontal	• Google Cloud Platform (GCP) - Load Balancer - HTTP • Google Cloud Platform (GCP) - Load Balancer - TCP - UDP
Google Cloud Platform (GCP) Memorystore	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Memorystore DB
Google Cloud Platform (GCP) Networking	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Networking
Google Cloud Platform (GCP) Organization	GCP	N/A	REST	Horizontal	GCP Organization

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Google Cloud Platform (GCP) resource inventory	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) Resource Inventory
Google Cloud Spanner	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Spanner DB
Google Cloud Platform (GCP) SSH Keys	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - SSH Keys
Google Cloud Platform (GCP) storage	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Storage
Google Cloud Platform (GCP) Virtual Server	GCP	N/A	REST	Horizontal	Google Cloud Platform (GCP) - Virtual Server
HAProxy Community edition load balancers	UNIX	1.5, 1.6, 1.7	SSH	Horizontal and top-down	HA Proxy
HP Operations Manager	Windows Unix	9.x	WMI WinRM SSH	Horizontal and top-down	HP Operations Manager
HP Quality Center	Windows	10.x	WMI WinRM	Top-down	HP Quality Center on Windows
HP Service Manager	Windows Unix	9.x	WMI WinRM SSH	Horizontal and top-down	HP Service Manager Application Server
HP UCMDB	Windows	N/A	WMI WinRM	Horizontal and top-down	HP uCMDB on Windows
HP-UX UNIX Server	HP-UX	11i	SSH	Horizontal	HP-UX Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
IBM AIX	UNIX	6.x, 7.x	SSH	Top-down	AIX Server
IBM CICS Transaction Gateway	Windows UNIX	5.x	WMI WinRM SSH	Top-down	CTG
IBM Cloud Platform	IBM Cloud Platform	N/A	REST	Horizontal	IBM Cloud Load Balancer, IBM Cloud Location Groups, IBM Cloud Network, IBM Cloud Organizations and Spaces, IBM Cloud Resource Groups, IBM Cloud SSH Key, IBM Cloud Storage, IBM Cloud Virtual Server
IBM Customer Information Control System	Windows UNIX	4.x, 5.x	WMI WinRM	Top-down	CICS
IBM Db2 on Unix	UNIX	8.x, 9.x, 10.x, 11.x	WMI WinRM SSH	Horizontal	DB2
IBM Db2 on Windows	Windows	8.x, 9.x, 10.x, 11.x	WMI WinRM SSH	Horizontal	DB2
IBM Flash System	N/A	N/A	REST	Horizontal	IBM Flash System
IBM Virtualization and Hardware	UNIX	N/A	SSH	Horizontal	IBM HMC Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Management Console (HMC) components					
IBM PowerHA Cluster (HACMP)	IBM	N/A	REST	Horizontal	IBM PowerHA Cluster (HACMP)
IBM Security Access Manager	ISAM	8.0.1	SNMP REST	Horizontal and top-down	ISAM Server
IBM WebSEAL discovery	ISAM	9.*,10.*.	REST	Horizontal and top-down	IBM WebSEAL
IBM Tivoli Access Manager	Windows UNIX	6.x	WMI WinRM SSH	Top-down	ITAM - Asset Center
IBM Tivoli WebSEAL	Windows UNIX	5.x, 6.x	WMI WinRM SSH	Top-down	Webseal (pattern)
IBM Websphere Application server	Windows UNIX	5.x, 6.x, 7.x, 8.x, 9.x	WMI WinRM SSH	Horizontal and top-down	Websphere
IBM WebSphere DataPower SOA Appliance	DataPower	I40, I50	SNMP REST	Horizontal and top-down	DataPower
IBM WebSphere Message Queue	Windows UNIX	6.x, 7.x, 8.x, 9.x	WMI WinRM SSH	Horizontal and top-down	The following patterns: • WMQ On Unix • WMQ On Windows

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
					• WMQ Queue Unix • WMQ Queue Windows
IBM App Connect Enterprise and HTTP listener discovery	UNIX	6, 7, 8, 9, 10, 11, 12.	N/A	Horizontal and top-down	WMB on UNIX (pattern)
IBM WebSphere Portal	Windows UNIX	6.x, 7.x, 8.x	WMI WinRM SSH	Horizontal and top-down	Websphere Portal On Windows
IBM Z mainframe	Z/OS	OS/390	SSH	Horizontal	IBM zOS Server
Infini-Box		5.0.x	REST HTTP	Horizontal	InfiniBox
Interconnect Web Service	Windows	2014	WMI WinRM	Top-down	Inter Connect
Istio Service Mash	Kubernetes	1.15 and later	HTTP	Horizontal	ISTIO Service mesh
Kubernetes	Kubernetes	1.21 - 1.36.2	REST HTTP	Horizontal	• Kubernetes • Kubernetes Cluster - Per-Namespace LP
Linux Server	UNIX	Ubuntu 16, RHEL 5.x, CentOS 7.x	SSH	Horizontal	Linux Server
Linux Pacemaker Cluster	Linux	N/A	SSH	Horizontal	Pacemaker Linux cluster

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Microsoft BizTalk Orchestration Server	Windows	2004, 2006, 2007, 2009, 2010	WMI WinRM	Top-down	BizTalk server
Microsoft Certificate Authority	Microsoft	N/A	WMI	Horizontal	Microsoft Certificate Authority
Microsoft Dynamics CRM	Windows	2003, 2005, 2007, 2010	WMI WinRM	Top-down	Dynamic CRM Component
Microsoft Fast Search	Windows	2005	WMI WinRM	Horizontal and top-down	Fast Search
Microsoft Foundry (classic)	Microsoft Azure	N/A	REST	Horizontal	Patterns: • Azure - AI Foundry Agents • Azure - AI Service Foundry • Azure - AI Service Foundry Project • Azure - AI Foundry Models
Microsoft Foundry (new)	Microsoft Azure	N/A	REST	Horizontal	Patterns: • Azure - AI New Foundry Agents • Azure - AI Service Foundry • Azure - AI Service Foundry Project • Azure - AI Foundry Models

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Microsoft Foundry Hub	Microsoft Azure	N/A	REST	Horizontal	Patterns: • Azure - AI Foundry Hub Agents • Azure - AI Foundry Hub Project • Azure - AI Foundry Models
Microsoft Hyper-V Server	Windows	2008, 2012, 2016	WMI WinRM	Horizontal	Hyper-V Server
Microsoft Identity Integration Server	Windows	2003	WMI WinRM	Top-down	IIFP On Windows Pattern
Microsoft Internet Information Services	Windows	6.0, 7.x, 8.x	WMI WinRM	Horizontal and top-down	IIS
Microsoft Message Queuing	Windows	3.x, 4.x, 5.x	WMI WinRM	Top-down	MSMQ (pattern)
Microsoft SharePoint	Windows	2003, 2007, 2010, 2013	WMI WinRM	Horizontal and top-down	Microsoft SharePoint
Microsoft SQL Server and Cluster discovery	Windows	2014, 2016, 2017, 2019, 2022	WMI WinRM	Horizontal and top-down	MSSQL DB On Windows (pattern)
Microsoft SQL Server Analysis Services (SSAS)	Windows	2014, 2016, 2017, 2019, 2022	WMI WinRM	Horizontal and top-down	SSAS (pattern)
Microsoft SQL Server Integration Services (SSIS) discovery	Windows	2014, 2016, 2017, 2019, 2022	WMI WinRM	Top-down	SSIS (pattern)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Microsoft SQL Server Reporting Services	Windows	2014, 2016, 2017, 2019, 2022	WMI WinRM	Top-down	SSRS (pattern)
Microsoft Windows Desktops	Windows	2003, 2008, 2012,	WMI WinRM	Horizontal	Windows OS - Desktops
Microsoft Windows Server	Windows	2003, 2008,2012,2016, 2022	WMI WinRM	Horizontal	Windows OS - Servers
MongoDB	Windows UNIX	3.x.x	WMI WinRM SSH	Horizontal and top-down	mongos server
NetApp Server and Cluster discovery	NetApp	7.x,8.x	SNMP	Horizontal	Patterns: NetApp Storage 7-Mode NetApp Cluster SNMP
NetApp Server and Cluster discovery	NetApp	7.x,8.x,9.x	HTTP	Horizontal	NetApp Cluster HTTP (pattern)
NetApp SolidFire storage system	NetApp SolidFire	N/A	HTTP	Horizontal	NetApp SolidFire storage system
Network router	Windows UNIX	N/A	SNMP	Horizontal	Network Router
Network switch	Windows UNIX	N/A	SNMP	Horizontal	Network Switch
NGINX Web Server	Windows UNIX	1.10, 1.10.1	WMI WinRM SSH	Top-down	Nginx Pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Nutanix Acropolis (AOS)	UNIX	N/A	REST	Horizontal	Nutanix Components
Nutanix Prism Central	UNIX	N/A	REST	Horizontal	Nutanix Components
OpenText Documentum	Windows UNIX	6.x, 7.x	WMI WinRM SSH	Top-down	Pattern for license server
OpenStack resources	OpenStack	N/A	REST	Horizontal	OpenStack (pattern)
OCI cloud hardware type	OCI	N/A	REST	Horizontal	Oracle OCI - Cloud Hardware Type (LP)
OCI OS image	OCI	N/A	REST	Horizontal	• Oracle OCI - Image (LP) • Oracle OCI - Cloud OS Image (LP)
OCI virtual machine	OCI	N/A	REST	Horizontal	Oracle OCI - Virtual Machine (LP)
Oracle Access Manager	Windows	N/A	WMI WinRM	Horizontal	Oracle - Access Manager
Oracle Application Server	Windows UNIX	9.x, 10.x	WMI WinRM SSH	Top-down	Oracle iAS
Oracle Cloud Infrastructure	Oracle	N/A	REST	Horizontal	Oracle Cloud Infrastructure
Oracle Clusterware	Linux	10.x, 11.x, 12.x	SSH	Horizontal	Oracle CRS
Oracle Coherence	UNIX	N/A	SSH	Horizontal	Oracle - Coherence
Oracle Concurrent Server	Windows	10.x, 11.x, 12.x	WMI	Top-down	Concurrent Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
	UNIX		WinRM SSH		
Oracle Database	Windows UNIX	8.x.x, 9.x.x, 10.x.x, 11.x.x, 12.x.x	WMI WinRM SSH	Horizontal and top-down	Oracle DB
Oracle Database Advanced Queuing	Windows UNIX	9.x.x, 10.x.x, 11.x.x, 12.x.x	WMI WinRM SSH	Top-down	Advanced Queue Queue
Oracle Database 12c	Windows UNIX	N/A	WMI WinRM SSH	Horizontal	PDB/CDB
Oracle Discover Server	Windows UNIX	9.x, 10.x	WMI WinRM SSH	Top-down	Discoverer Engine
Oracle E-Business Suite	Windows UNIX	10.x, 11.x	WMI WinRM SSH	Top-down	E-Business Suite
Oracle Enterprise Manager Management Server	UNIX	N/A	SSH	Horizontal	Oracle - Enterprise Manager Management Server
Oracle Forms	Windows UNIX	9.x, 10.x	WMI WinRM SSH	Top-down	Forms Engine
Oracle Forms	Windows	N/A	WMI WinRM	Horizontal	Oracle - Forms

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Oracle Fulfillment Server	Windows UNIX	9.x, 10.x	WMI WinRM SSH	Top-down	Fulfillment Server
Oracle GoldenGate	Windows UNIX	N/A	WMI SSH	Horizontal and top-down	Oracle GG
Oracle Linux Virtualization Manager (OLVM) and Red Hat Virtualization (RHV)	OLVM and RHV	N/A	REST	Horizontal	Patterns: • oVirt Clusters and Hosts • oVirt Discover Logical Datacenters • oVirt Discover Manager Instance • oVirt Virtual Machines • oVirt Templates
Oracle Global License Advisory Services (GLAS)	N/A	N/A	SSH	Horizontal	Oracle GLAS Data Collection (pattern)
Oracle HTTP Server	Windows UNIX	N/A	WMI WinRM SSH	Top-down	HTTP Server
Oracle HTTP Server	Windows	N/A	WMI	Horizontal	Oracle - HTTP Server
Oracle Internet Directory Server	Windows	N/A	WMI	Horizontal	Oracle - Internet Directory Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Oracle iPlanet Web Server	Solaris	4.x, 6.x, 7.x	SSH	Horizontal and top-down	Sun iPlanet Webserver
Oracle Java Enterprise System(JES)	Solaris	7.x	SSH	Top-down	Sun JES pattern
Oracle Java processes	JDK	All	N/A	Horizontal	Java installation pattern ACC based discovery
Oracle Listener	Linux	12.2.x.x	SSH	Horizontal	Oracle Listener HD
Oracle Metric Server	Windows UNIX	9.x, 10.x	WMI WinRM SSH	Top-down	Metric Server
Oracle MySql Server	Windows UNIX	4.x, 5.x	WMI WinRM SSH	Horizontal and top-down	My SQL server On Windows and Linux (pattern)
Oracle Node Manager	UNIX	N/A	SSH	Horizontal	Oracle - Node Manager
Oracle Notification Server	Windows UNIX	9.x, 10.x	WMI WinRM SSH	Top-down	Oracle Notification Server
Oracle Peoplesoft	Windows /Unix	8.x	WMI WinRM SSH	Top-down	Peoplesoft Application Server
Oracle Process Manager	Windows /Unix	9.x.x, 10.x.x	WMI WinRM SSH	Top-down	Process Manager

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Oracle RAC	Windows UNIX	9.x.x, 10.x.x, 11.x.x, 12.x.x	WMI WinRM SSH	Top-down	App TNS Service
Oracle Reports	Windows	N/A	WMI	Horizontal	Oracle - Reports
Oracle Report Server	Windows UNIX	9.x.x, 10.x.x	WMI WinRM SSH	Top-down	Report Server
Oracle Solaris Server	Solaris	9, 10, 11, SPARC	SSH	Horizontal	Solaris Server
Oracle Solaris LDOM	Solaris	10, 11	SSH	Horizontal	Solaris Logical Domain (LDOM) infrastructure
Oracle Solaris Logical Domain infrastructure	Solaris	N/A	SSH	Horizontal	Solaris LDOM
Oracle TNS Listener	Windows	N/A	WMI	Horizontal	Oracle - TNS Listener
Oracle Siebel CRM Server	Windows	N/A	WMI WinRM	Horizontal	Oracle - Siebel CRM Server
Oracle Siebel Gateway Name Server	Windows	N/A	WMI	Horizontal	Oracle - Siebel Gateway Name Server
Oracle Unified Directory	Windows	N/A	WMI WinRM	Horizontal	Oracle - Unified Directory
Oracle Weblogic Application Server	Windows UNIX	8.x, 9.x, 10.x, 11.x, 12.x	WMI WinRM SSH	Horizontal and top-down	WebLogic

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Pivotal Cloud Foundry	PCF	N/A	REST	Horizontal and top-down	Patterns: Pivotal cloud foundry Bosh
Pivotal RabbitMQ	Windows UNIX	3.x.x	WMI WinRM SSH	Top-down	RabbitMQ Cluster Unix pattern
PostgreSQL DB	Windows UNIX	6.x, 7.x, 8.x	WMI WinRM SSH	Horizontal and top-down	PostgreSQL DB
Puppet	UNIX	N/A	SSH	Horizontal	Probes: • Puppet – Master Info • Puppet – Certificate Requests • Puppet – Manifest • Puppet – Modules
Pure Storage FlashArray discovery	Pure Storage	N/A	REST	Horizontal	Pure Flash Array Storage (pattern)
Pure Storage FlashBlade	Pure Storage	4.10.7	REST SSH	Horizontal	Pure Storage
Radware Alteon RadWare ADC	Alteon	v31, v29.5	SNMP	Horizontal and top-down	Alteon Load Balancer
RadWare AppDirector Load Balancer	Radware	N/A	SNMP	Horizontal and top-down	AppDirector Load Balancer
Red Hat 389 Directory Server	UNIX	N/A	SSH	Horizontal	Red Hat - Directory Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
Red Hat Cluster Suite	Linux	RH 5.x, 6.x, 7.x	SSH	Horizontal	RH Cluster
Red Hat JBoss Application Server	Windows UNIX	6.x, 7.x	WMI WinRM SSH	Horizontal and top-down	Jboss (pattern)
Red Hat JBoss EAP Host Controller	Windows	N/A	WMI WinRM	Horizontal	Red Hat JBoss - Enterprise App Platform Host Controller
Red Hat JBoss EAP Server on UNIX	UNIX	N/A	SSH	Horizontal	Red Hat - JBoss Enterprise App Platform Server on Unix
Red Hat JBoss EAP Server on Windows	Windows	N/A	WMI WinRM	Horizontal	Red Hat JBoss Enterprise Application Platform Server
Red Hat JBoss Fuse discovery	Windows UNIX	N/A	SSH	Horizontal and top-down	Jboss Fuse (pattern)
Red Hat OpenShift components of Kubernetes	RHEL	API 1.x	REST	Horizontal	Collect OpenShift info library used as an extension section of the Kubernetes pattern
Red Hat Podman	UNIX	N/A	SSH	Horizontal	Red Hat - Podman
Red Hat SSO Server	Windows	N/A	WMI WinRM	Horizontal	Red Hat - Single Sign-On Server
Red Hat WildFly Server	Windows	N/A	WMI WinRM	Horizontal	Red Hat - WildFly Server
Rubrik cluster		N/A	HTTP	Horizontal	Multiple patterns

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
			HTTPS		
SAP Business Object XI Scheduler	Windows	2004, 2005, 2008	WMI WinRM	Horizontal and top-down	SAP BO BOXISchedule Router on Windows Pat
SAP Crystal Management Server	Windows	2004, 2005, 2008	WinRM	Horizontal and top-down	SAP Business Objects CMS Server on Windows
SAP HANA	Windows UNIX	2015, 2016, 2017	WMI WinRM SSH	Horizontal and top-down	SAP HANA DB (pattern)
SAP HANA Catalog	UNIX	N/A	N/A	Horizontal	SAP Hana 2.0 DB Catalog (pattern)
SAP Sybase ASE discovery	Windows UNIX	16	WMI WinRM SSH	Horizontal and top-down	Sybase (pattern)
SAP Sybase ASE discovery catalog	UNIX	16	SSH	Horizontal	Sybase (pattern extension)
Cisco Switch Wireless Access Point (WAP)	Cisco Network Switch	N/A	SNMP	Horizontal	WAP
Sun Directory Proxy Server	Solaris	4.0.8	SSH	Top-down	Sun Directory Proxy Server
Sun Java System Directory Server	Solaris	6.3	SSH	Top-down	Sun Directory Server
Thunder ADC	A10	2.x, 4.x	SNMP/SSH	Horizontal and top-down	A10 Load Balancer
Tibco ActiveMatrix Adapter	UNIX	7.1.0	SSH	Top-down	Tibco Adapter
Tibco BusinessWorks	Windows	5.7.2	WMI	Horizontal	ActiveMatrix Business Works

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Platform	Version	Protocol	Discovery type	Pattern (or probe if indicated)
	UNIX		WinRM SSH		
Tibco Enterprise Message Service (EMS)	Windows UNIX	5.x, 6.x	WMI WinRM SSH	Horizontal and top-down	Enterprise Message Service
Tibco Enterprise Message Service (EMS) Queue	Windows UNIX	5.x, 6.x	WMI WinRM SSH	Horizontal and top-down	Enterprise Message Service
Tibco Hawk	UNIX	5.2.0	SSH	Top-down	TibcoHawk
Veritas Enterprise Vault	Windows	14.1	WMI WinRM	Top-down	Enterprise Vault
Veritas Cluster Server	UNIX	6.x.x	SSH	Horizontal and top-down	Veritas Cluster
VMware NSX Advanced Load Balancer	VMware	1	REST API	Horizontal and top-down	NSX
VMware NSX-T infrastructure	VMware	1	REST API	Horizontal	NSX Cluster

Cloud resource discovery references

ITOM Visibility discovers resources on cloud platforms.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Data collected for Amazon AWS Cloud Discovery

Discovery collects information about cloud resources in AWS datacenters. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date.

- Discovery and Service Mapping Patterns.
- CMDB CI Class Models.
- Visibility Content.

Configure a discovery schedule

Create a discovery schedule in Cloud Discovery Workspace

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Note: You can test the AWS REST APIs using Postman API platform. For more information, see the How to test AWS REST API using POSTMAN [KB0782183] article in the Now Support Knowledge Base.

Data collected for AWS Cloud Discovery

The following tables show examples of data collected during AWS Cloud Discovery.

AWS Datacenter [cmdb_ci_aws_datacenter]

Label	Field Name
Name	name
Region	region
Object ID	object_id

Availability Zone [cmdb_ci_availability_zone]

Label	Field Name
Name	name

Virtual Machine Instance [cmdb_ci_vm_instance]

Label	Field Name
Name	name
State	state

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Virtual Machine Instance [cmdb_ci_vm_instance] (continued)

Label	Field Name
Object ID	object_id
CPUs	cpus
Disks	disks
Memory	memory
Network adapters	nics
VM Instance ID	vm_inst_id

Compute Security Groups [cmdb_ci_compute_security_group]

Label	Field Name
Name	name
Object ID	object_id
State	state

Images [cmdb_ci_os_template]

Label	Field Name
Name	name
Object ID	object_id
Guest OS	guest_os
Root device type	root_device_type
Image source	image_source
Image type	image_type

Hardware Type [cmdb_ci_compute_template]

Label	Field Name
Name	name
vCPUs	vcpus
Memory MB	memory_mb
Local Storage GB	local_storage_gb

Storage Volume [cmdb_ci_storage_volume]

Label	Field Name
Name	name
State	state
Object ID	object_id

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume [cmdb_ci_storage_volume] (continued)

Label	Field Name
Storage type	storage_type
Size	size

Cloud Networks [cmdb_ci_network] and VMware vCenter Network [cmdb_ci_vcenter_network]

Label	Field Name
Name	name
State*	state
CIDR*	cidr

*Not found on VMware vCenter networks.

Cloud Subnets [cmdb_ci_cloud_subnet]

Label	Field Name
Name	name
Status	status
CIDR	cidr

Cloud Management Network Interfaces [cmdb_ci_nic]

Label	Field Name
Name	name
Netmask	netmask
MAC Manufacturer	mac_manufacturer
Status	install_status

Cloud Load Balancers [cmdb_ci_cloud_load_balancer]

Label	Field Name
Name	name
Object ID	object_id
State	state

Resource Groups [cmdb_ci_resource_group]

Label	Field Name
Name	name
Object ID	object_id

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource Groups [cmdb_ci_resource_group] (continued)

Label	Field Name
State	state

Public IP Addresses [cmdb_ci_cloud_public_ipaddress]

Label	Field Name
Name	name
Object ID	object_id
Public IP address	public_ip_address
Public DNS	public_dns

Storage Accounts [cmdb_ci_cloud_storage_account]

Label	Field Name
Name	name
Object ID	object_id
Sku Name	sku_name
State	state

DNS Alias [cmdb_ci_dns_alias] and DNS name [cmdb_ci_dns_name]

Label	Field name
DNS Alias [cmdb_ci_dns_alias]	
Name	name
Category	category
Status	status
DNS name [cmdb_ci_dns_name]	
Name	name
IP address	ip_address

Cloud DataBases [cmdb_ci_cloud_database]

Label	Field Name	Description
Name	name	The name of the database that you created in AWS.
Object ID	object_id	This is also the name of the database.
Type	Type	The type of database you created.
Fully qualified domain name	fqdn	The FQDN that AWS assigned to your database. An example format for AWS is as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud DataBases [cmdb_ci_cloud_database] (continued)

Label	Field Name	Description
		`database-name.{random-number}.{datacenter}.rds.amazonaws.com`
State	state	The state of the database: whether it is Available or Terminated.
TCP port(s)	tcp_port	The TCP port that the database communicates through.
Category	category	The instance class of the database, for example: db.t2.micro.
Automated Backups	automated_backup	Indicates whether automated backups are enabled and the retention period.

Cloud WebServer [cmdb_ci_cloud_webserver]

Label	Field Name
Name	name
Install status	install_status
Vendor	vendor
Fully qualified domain name	fqdn
Operational status	operational_status
State	state

Kubernetes Clusters [cmdb_ci_kubernetes_cluster]

Label	Field Name	Description
Name	name	The name of the kube-controller-manager leader.
Port	port	The identifier for the Kubernetes apiserver port.
Kubernetes UID	k8s_uid	The unique identifier for the Kubernetes Cluster.
Vendor	vendor	The vendor of the Kubernetes Cluster.
Fully qualified domain name	fqdn	The fully qualified domain name of the Kubernetes apiserver.

Related topics

- AWS discovery using patterns
- Service Graph Connector for AWS
- Agent Client Collector Discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected for Microsoft Azure Cloud Discovery

Discovery collects information about cloud resources in Microsoft Azure datacenters. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- CMDB CI Class Models
- Visibility Content

Create Azure cloud credentials

Create Microsoft Azure credentials that can access the Microsoft Azure account. For more information, see Create Azure cloud credentials.

Configure a discovery schedule

Create a discovery schedule in Cloud Discovery Workspace

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

How Discovery finds Microsoft Azure resources

Discovery uses the Cloud Provisioning and Governance API and Discovery and Service Mapping Patterns to find cloud resources.

Default patterns

Pattern	Description
Azure DataBase (LP)	Discovers database types of: • Microsoft.sql/servers • microsoft.sql/managedInstances • microsoft.dbformysql • microsoft.documentdb/databaseaccounts • microsoft.cache/redis • microsoft.dbforpostgresql/servers • microsoft.dbformysql/servers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Default patterns (continued)

Pattern	Description
	• microsoft.dbformysql/flexibleservers • microsoft.dbforpostgresql/flexibleservers Populates the [cmdb_ci_cloud_database] table.
Azure LoadBalancer TD	Retrieves Microsoft Azure load balancers and populates the cmdb_ci_lb_service table. This pattern is only used by Service Mapping for top-down discovery, not by the Discovery application for horizontal discovery.
Azure WebSite (LP)	Retrieves Microsoft Azure web servers and populates the Cloud WebServers [cmdb_ci_cloud_webserver] and IP address [cmdb_ci_ip_address] tables.

Data collected by Discovery during horizontal discovery

Azure Datacenter [cmdb_ci_azure_datacenter]

Label	Field Name
Name	name
Region	region
Object ID	object_id

Availability Zone [cmdb_ci_availability_zone]

Label	Field Name
Name	name

Virtual Machine Instance [cmdb_ci_vm_instance]

Label	Field Name
Name	name
State	state
Object ID	object_id
CPUs	cpus
Disks	disks
Disks size (GB)	disks_size
Memory	memory
Network adapters	nics
VM Instance ID	vm_inst_id

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Compute Security Groups [cmdb_ci_compute_security_group]

Label	Field Name
Name	name
Object ID	object_id
State	state

Images [cmdb_ci_os_template]

Label	Field Name
Name	name
Object ID	object_id
Guest OS	guest_os
Image source	image_source

Storage Volume [cmdb_ci_storage_volume]

Label	Field Name
Name	name
State	state
Object ID	object_id
Storage type	storage_type
Size	size

Cloud Networks [cmdb_ci_network] and VMware vCenter Network [cmdb_ci_vcenter_network]

Label	Field Name
Name	name
State*	state
CIDR*	cidr

*Not found on VMware vCenter networks.

Cloud Subnets [cmdb_ci_cloud_subnet]

Label	Field Name
Name	name
Status	status
CIDR	cidr

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Provisioning and Governance Network Interfaces [cmdb_ci_nic]

Label	Field Name
Name	name
Netmask	netmask
MAC Address	mac_address
MAC Manufacturer	mac_manufacturer
Status	install_status

Cloud Load Balancers [cmdb_ci_cloud_load_balancer]

Label	Field Name
Name	name
Object ID	object_id
State	state

Azure Deployments [cmdb_ci_azure_deployment]

Label	Field Name
Name	name
Provisioning state	provisioning_state

Note: This table is only available for Discovery when the Cloud Provisioning and Governance plugin is activated.

Resource Groups [cmdb_ci_resource_group]

Label	Field Name
Name	name
Object ID	object_id
State	state
Operational Status	operational_status
Install Status	install_status

Note:

- When a resource group is discovered for the first time, the **State** field is set to **Available**, **operational_status=Operational**, and **Status=Installed**; based on the response mapping.
- When a resource group is terminated on Microsoft Azure, the CMPReconciler scripts set the **State** field to **Terminated**, and **operational_status=Non-Operational** and **Status=Absent**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Public IP Addresses [cmdb_ci_cloud_public_ipaddress]

Label	Field Name
Name	name
Object ID	object_id
Public IP address	public_ip_address
Public DNS	public_dns

Storage Accounts [cmdb_ci_cloud_storage_account]

Label	Field Name
Name	name
Object ID	object_id
Sku Name	sku_name
State	state

Cloud Databases [cmdb_ci_cloud_database]

Label	Field Name
Fully qualified domain name	fqdn
Name	name
Object ID	object_id
Operational Status	operational_status
Type	type
State	state
Vendor	vendor
Version	version **Note:** For NoSQL and MSSQL Managed Instances, the API responses for these databases do not return version data.
Category*	category

* Populated only by the Azure SQL Managed Instance license pattern extension. For more information, see the "Azure SQL Managed Instance license discovery" section in Microsoft Azure Cloud discovery using patterns.

Cloud WebServer [cmdb_ci_cloud_webserver]

Label	Field Name
Name	name
Install status	install_status

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud WebServer [cmdb_ci_cloud_webserver] (continued)

Label	Field Name
Vendor	vendor
Fully qualified domain name	fqdn
Operational status	operational_status
State	state

IP Address [cmdb_ci_ip_address]

Label	Field Name
Fully qualified domain name	fqdn
Install status	install_status
IP Address	ip_address
Name	name
Netmask	netmask
Operational status	operational_status
Vendor	vendor

Database [cmdb_ci_database]

Label	Field Name
Life-cycle Stage Status	life_cycle_stage_status
Name	name
Status	status
Serial Number	serial_number
Type	type

Compute Template [cmdb_ci_compute_template]

Label	Field Name
Name	name
Object ID	object_id
VCPUS	vcpus

Hardware Type [cmdb_ci_hardware_type]

Label	Field Name
Name	name
Object ID	object_id
Provider	provider

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hardware Type [cmdb_ci_hardware_type] (continued)

Label	Field Name
vCPUs	vcpus

Related topics

Microsoft Azure Cloud discovery using patterns

Service Graph Connector for Microsoft Azure

Agent Client Collector Discovery

Data collected for GCP Cloud Discovery

Discovery collects information about cloud resources in Google Cloud Platform (GCP). Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date.

- Discovery and Service Mapping Patterns.
- CMDB CI Class Models.
- Visibility Content.

Configure a discovery schedule

Create a discovery schedule in Cloud Discovery Workspace

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

GCP data collected

Table	ServiceNow Field	GCP Field
Region [cmdb_ci_google_datacenter]	object_id	items.id
	name	name
	short_description	description
	status	status
	region	items.name

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	ServiceNow Field	GCP Field
Virtual Server [cmdb_ci_vm_instance]	object_id	id
	name	name
	state	status
	cpus	Not mapped
	memory	Not mapped
	disks	disks (array size)
	disks_size	Not mapped
	nics	networkInterfaces (array size)
	termination_protection	deletionProtection
	Not mapped	canIpForward
	Not mapped	selfLink
	sys_created_on	creationTimestamp
Storage Volume (cmdb_ci_storage_volume)	object_id	items.id
	name	name
	short_description	description
	state	status
	status	status
	size	sizeGb
	storage_type	storage_type
Cloud Public IP Address [cmdb_ci_cloud_public_ipaddress]	object_id	items.id
	name	name
	public_ip_address	address
	state	status
Availablility Zone [cmdb_ci_availability_zone]	object_id	items.id
	name	name
	short_description	description
	state	status
	status	status
Hardware Template [cmdb_ci_compute_template]	object_id	items.id
	name	name
	short_description	description
	vcpus	guestCpus
	memory_mb	memoryMb
	local_storage_gb	imageSpaceGb
Instance Template [cmdb_ci_instance_template]	object_id	items.id

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	ServiceNow Field	GCP Field
	name	name
	short_description	description
Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]	object_id	items.id
	name	name
	state	status
	volume_name	sourceDisk
	capacity	storageBytes
	No mapping	storageBytesStatus
	No mapping	selfLink
	No mapping	labelFingerprint
	No mapping	diskSizeGb
Disk Type [cmdb_ci_disk_type]	object_id	zone:name
	name	name@zone
	short_description	description
	valid_disk_size	validDiskSize
	default_disk_size_gb	defaultDiskSizeGb
Subnet [cmdb_ci_cloud_subnet]	object_id	id
	name	name
	cidr	ipCidrRange
	gateway	gatewayAddress
Image [cmdb_ci_os_template]	object_id	items.id
	name	name
	image_type	sourceType
	No mapping	status
	No mapping	archiveSizeBytes
	No mapping	diskSizeGb
	image_source	• sourceDisk • sourceDiskId
Load Balancer (cmdb_ci_cloud_load_balancer)	object_id	id
	name	name
	No mapping	defaultService
Load Balancer Service (cmdb_ci_lb_service)	object_id	id
	name	name
	pool	backends

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	ServiceNow Field	GCP Field
	port	port
	listener_protocol	protocol
	No mapping	portName
	No mapping	loadBalancingScheme
	No mapping	sessionAffinity
	No mapping	enableCDN
	service_type	Backend/Frontend
Load Balancer Health Service [cmdb_ci_lb_health_service]	object_id	id
	name	name
	monitor_type	type (tcp/http/https)
	request_path	requestPath
	port	port
	unhealthy_threshold	unhealthyThreshold
	healthy_threshold	healthyThreshold
	timeout_sec	timeoutSec
	check_interval_sec	checkIntervalSec
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	object_id	Instance ID from GCP
	name	instance.substring after the last'/'
	Operational status	status
Load Balancer Pool [cmdb_ci_lb_pool]	object_id	id
	name	name
Network [cmdb_ci_network]	state	No mapping
	cidr	IPv4Range
	is_external	No mapping
	name	name
	object_id	items_id
	description	description
Network ACL [cmdb_ci_network_acl]	object_id	id
	name	name
	short_description	description
Network ACL Rule [cmdb_ci_network_acl_rule]	is_outbound (True/False)	ac
	dest_from_port	No mapping
	dest_to_port	No mapping
	src_from_port	No mapping
	src_to_port	No mapping

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	ServiceNow Field	GCP Field
	source_ranges	sourceRanges
	destination_ranges	destinationRanges
	target_tags	targetTags
	allowed_traffic	allowed
Kubernetes Clusters [cmdb_ci_kubernetes_cluster]	IP address	ip_address
	Name	name
	Port	port
	Kubernetes UID	k8s_uid
	Vendor	vendor

Related topics

Google Cloud Platform (GCP) Cloud discovery using Patterns

Service Graph Connector for GCP

Agent Client Collector Discovery

Data collected for IBM Cloud Discovery

Discovery collects information about cloud resources in the IBM Cloud Platform and creates relationships between the CIs it finds. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date.

- Discovery and Service Mapping Patterns.
- CMDB CI Class Models.
- Visibility Content.

Configure a discovery schedule

Create a discovery schedule in Cloud Discovery Workspace

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load balancers

Data collected for load balancers

Table	Field Name / Description
Cloud Load Balancer [cmdb_ci_cloud_load_balancer]	**object_id**: The load balancer's unique identifier.
	name: Name of the load balancer.
	fqdn: Address (Host name) of the load balancer. This is the fully qualified domain name of the load balancer.
	dns_name: Address (Host name) of the load balancer.
	short_description: Description of the load balancer.
	operational_status: The operational status of the load balancer, either **ONLINE** or **OFFLINE**.
	location: References the load balancer's location with type.
	state: Provisioning status of the load balancer, either **Available** or **Terminated**.
Load Balancer Service [cmdb_ci_lb_service]	**object_id**: Unique identifier of the load balancer service.
	name: Name created by the pattern that contains the following: • Load balancer name • Protocol • Port • Resource type: **-listener**
	port: Listener's protocol port number.
	front_end_port: Listener's protocol port number.
	listener_protocol: Listener's protocol, either **TCP**, **HTTP**, or **HTTPS**.
	operational_status: The provisioning status of the listener.
Load Balancer Pool [cmdb_ci_lb_pool]	**object_id**: Unique identifier of the load balancer pool.
	name: Name created by the pattern that contains the following: • Load balancer name • Protocol • Port • Resource type: **-pool**
	port: Protocol port of the back end.
	pool_status: Provisioning status of a load balancer pool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected for load balancers (continued)

Table	Field Name / Description
	load_balancing_method: Method used to load balance the pool.
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	**object-id**: Unique identifier of the pool member, created from the ID of the member and the listener ID.
	name: Name created by the pattern that contains the following: • Load balancer name • Protocol • Port • Resource type: **-member** • IP address of the pool member
	ip_address: IP address of a load balancer member.
	server_id: ID of the member in the IBM Cloud.
	load_balancer: Name of the related load balancer.
Load Balancer Health Service [cmdb_ci_lb_health_service]	**object_id**: Unique identifier of the healthservice check.
	name: Name created by the pattern that contains the following: • Load balancer name • Protocol • Resource type: **Health Check** • Resource ID
	monitor_type: Type of health check. Valid values are **TCP**, **HTTP**, or **HTTPS**.
	timeout_sec: Timeout in seconds to wait for the health check response.
	check_interval_sec: Interval in seconds to perform the health check.

Relationships created for load balancers

CI	Relationship	CI
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_load_balancer	Contains::Contained by	cmdb_ci_lb_health_service
cmdb_ci_cloud_load_balancer	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_availability_zone	Contains::Contained by	cmdb_ci_cloud_load_balancer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships created for load balancers (continued)

CI	Relationship	CI
cmdb_ci_lb_service	Hosted on::Hosts	cmdb_ci_cloud_load_balancer
cmdb_ci_availability_zone	Contains::Contained by	cmdb_ci_lb_service
cmdb_ci_lb_pool	Hosted on::Hosts	cmdb_ci_cloud_load_balancer
cmdb_ci_lb_service	Contains::Contained by	cmdb_ci_lb_pool
cmdb_ci_lb_pool	Owns::Owned by	cmdb_ci_lb_pool_member

Organizations and spaces

Data collected for organizations and spaces

Table	Fields
IBM Cloud Organizations [cmdb_ci_cloud_org]	• **guid**: The organization's unique identifier. • **url**: The organization's unique URL. • **billing_enabled**: Indicates if the organization's billing method is enabled or disabled. • **name**: Unique name of the cloud organization. • **quota_definition_url**: Resources that are available to the organization. • **spaces_url**: Unique URL of the space assigned to the organization. • **operational_status**: Operational status of the organization.
IBM Cloud Space [cmdb_ci_cloud_space]	• **object_id**: Unique identifier (GUID) of the space. • **name**: Unique name of the space. • **url**: Unique URL of the space. • **allow_ssh**: The IP address of a storage resource.

Relationships created for organizations and spaces

CI	Relationship	CI
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_org	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_space	Hosted on::Hosts	cmdb_ci_cloud_org
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_cloud_space

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SSH Keys

Data collected for SSH keys

Table	Fields
Cloud Key Pairs [cmdb_ci_cloud_key_pair]	• **object_id**: ID of the SSH key record. • **name**: Descriptive name used to identify the SSH key. • **key**: The SSH key. • **finger_print**: Short sequence of bytes used to authenticate or lookup a longer SSH key. This is generated automatically when adding or modifying the SSH key. This property is read only. Changes made are ignored. • **created_date**: Date the SSH key was added. This property is read only. Changes made are ignored. • **modify_date**: Date the SSH key was last modified. This property is read only. Changes made are ignored. • **short_description**: Brief note about the SSH key.

Relationships created for SSH keys

CI	Relationship	CI
cmdb_ci_ibm_datacenter	Contained by::Contains	cmdb_ci_availability_zone
cmdb_ci_cloud_key_pair	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_key_pair	Hosted on::Hosts	cmdb_ci_cloud_service_account

IBM Cloud storage

Data collected for IBM Cloud storage

Table	Fields
Storage Volume [cmdb_ci_storage_volume]	• **object_id**: Unique identifier of the storage account. • **size**: Capacity of the storage account, measured in gigabytes. • **size_bytes**: Capacity of the storage account, measured in bytes. • **name**: Username used to access a non-EVault Storage volume. This username is

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected for IBM Cloud storage (continued)

Table	Fields
	used to register the EVault server agent with the vault backup system. • **label**: Network resource name for the storage. • **storage_type**: Description of the storage object. • **fqdn**: IP address of the storage resource. • **device**: Unique identification number of the guest associated with the storage volume.
Storage Volume Snapshot [cmdb_ci_storage_vol_snapshot]	• **object_id**: Unique identifier of the storage account. • **name**: Name of the storage's network resource. • **capacity**: The volume's configured snapshot space size. • **fqdn**: IP address of the storage resource. • **location**: Unique availability zone where the storage volume is located.

Relationships created for IBM Cloud storage

CI	Relationship	CI
cmdb_ci_ibm_datacenter	Contains::Contained by	cmdb_ci_availability_zone
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_storage_volume	Provisioned::Provisioned From	cmdb_ci_storage_vol_snapshot
cmdb_ci_storage_vol_snapshot	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_storage_volume	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_storage_vol_snapshot	Located in Zone::Zone contains	cmdb_ci_availability_zone
cmdb_ci_availability_zone	Contains::Contained by	cmdb_ci_storage_volume

IBM Cloud virtual servers

Data collected for IBM Cloud virtual servers

Table	Fields
Virtual Machine Instance [cmdb_ci_vm_instance]	• **name**: Name of the virtual machine. • **object_id**: ID of the virtual machine. • **dns_domain**: Domain of the virtual machine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected for IBM Cloud virtual servers (continued)

Table	Fields

	• **fqdn**: IP address of the virtual machine.
	• **short_description**: Notes attached to the virtual machine on the IBM console.
	• **cpus**: Number of CPUs on the virtual machine.
	• **memory**: Amount of memory on the virtual machine.
	• **disks**: Number of disks used by the virutal machine.
	• **state**: The running state of the virtual machine. The choices are:
	  ○ **Running**
	  ○ **Halted**
	• **placement_group_id**: Name of the IBM placement group containing the virtual machine.
	• **nics**: Number of network interfaces on the virtual machine.

Cloud Mgmt Network Interface [cmdb_ci_nic]	• **name**: Name and port of the NIC, in the format **<name><port>**. • **ip_address**: IP address configured on the NIC. • **mac_address**: MAC address of the NIC. • **netmask**: Netmask of the NIC. • **object_id**: ID of the NIC. • **status**: State of the NIC - AVAILABLE.
Image [cmdb_ci_os_template]	• **name**: Long description of the template. For example: **CentOS 7.0-64 Minimal for VSI**. • **version**: Operating system version. For example, **7.0-64 Minimal for VSI**. • **object_id**: ID of the image. • **guest_os**: Operating system of the virtual machine, such as Windows, Ubuntu, or CentOS.
Cloud Key Pair [cmdb_ci_cloud_key_pair]	• **name**: The key label. • **object_id**: ID of the key.
Compute Security Group [cmdb_ci_compute_security_group]	• **name**: Security group name. • **object_id**: ID of the security group • **short_description**: Security group notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected for IBM Cloud virtual servers (continued)

Table	Fields
Storage Volume [cmdb_ci_storage_volume]	**object_id**: The storage volume ID.

Relationships created for IBM Cloud virtual servers

CI	Relationship	CI
cmdb_ci_vm_instance	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_vm_instance	Located in::Houses	cmdb_ci_availability_zone
cmdb_ci_nic	Contained by::Contains	cmdb_ci_vm_instance
cmdb_ci_os_template	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_os_template	Provisioned by::Provisions	cmdb_ci_vm_instance
cmdb_ci_cloud_key_pair	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_cloud_key_pair	Used by::Uses	cmdb_ci_vm_instance
cmdb_ci_compute_security_group	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_compute_security_group	Implement End Point From::Implement End Point To	cmdb_ci_vm_instance
cmdb_ci_storage_volume	Hosted on::Hosts	cmdb_ci_ibm_datacenter
cmdb_ci_storage_volume	Implement End Point From::Implement End Point To	cmdb_ci_vm_instance

IBM Cloud service account

Data collected for IBM Cloud service account

Table	Field
Cloud Service Account [cmdb_ci_cloud_service_account]	• **name**: Display name of the cloud service account. • **account_id**: Service account ID. • **object_id**: Cloud ID of the service account. This is the same as the account_id. • **datacenter_type**: This is **cmdb_ci_ibm_datacenter**. • **discovery_credentials**: The sys_id of the credentials.
IBM Datacenter [cmdb_ci_ibm_datacenter]	• **name**: Display name of the cloud service account. • **account_id**: Service account ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected for IBM Cloud service account (continued)

Table	Field
	• **object_id**: Cloud ID of the service account. This is the same as the account_id. • **datacenter_type**: This is **cmdb_ci_ibm_datacenter**.

Relationships created for IBM Cloud service accounts

CI	Relationship	CI
cmdb_ci_cloud_service_account	Contained by::Contains	cmdb_ci_availability_zone
cmdb_ci_cloud_service_account	Contained by::Contains	cmdb_ci_ibm_datacenter
cmdb_ci_cloud_service_account	Hosts::Hosted on	cmdb_ci_cloud_key_pair
cmdb_ci_cloud_service_account	Hosts::Hosted on	cmdb_ci_vm_instance
cmdb_ci_cloud_service_account	Hosts::Hosted on	cmdb_ci_lb_service
cmdb_ci_ibm_datacenter	Contained by::Contains	cmdb_ci_availability_zone
cmdb_ci_ibm_datacenter	Hosted on::Hosts	cmdb_ci_cloud_service_account
cmdb_ci_ibm_datacenter	Hosts::Hosted on	cmdb_ci_network

Data collected for VMware Cloud Discovery

Discovery collects information about VMware resources in your cloud service accounts.

Distributed Resource Scheduler (DRS) discovery

VMware Discovery collects these DRS settings and saves them to the DRS VM Config [cmdb_ci_drs_vm_config] table:

- DRS settings of vCenter clusters.
- DRS settings of VMs configured to override cluster behavior.

DRS VM Config [cmdb_ci_drs_vm_config]

Field label and Name	Description
Cluster [cluster]	The vCenter cluster that the VM belongs to.
DRS behavior [drs_behavior]	The VM-level Distributed Resource Scheduler behavior that overrides cluster defaults.
DRS enabled [drs_enabled]	Indicates whether DRS is active for the VM.
Virtual machine [virtual_machine]	The VM instance using these DRS settings.

VMware tags

VMware Cloud Discovery finds VMware tags from vCenter and saves them to the Key Value [cmdb_key_value] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Values [cmdb_key_value]

Field label and Name	Description
Configuration item [configuration_item]	Referenced VM with tag attached in vCenter.
Key [key]	Category name of the tag applied to the resource in vCenter.
Tag [tag]	Source in vCenter from which tags are fetched. In this case, the value is always **Tags**.
Value [value]	Tag name that's applied to the resource in vCenter.

Data collected for VMware vCenter Server

Discovery identifies and classifies information and data about VMware vCenter servers.

The vCenter table schema is illustrated in the following diagram:

vCenter table schema

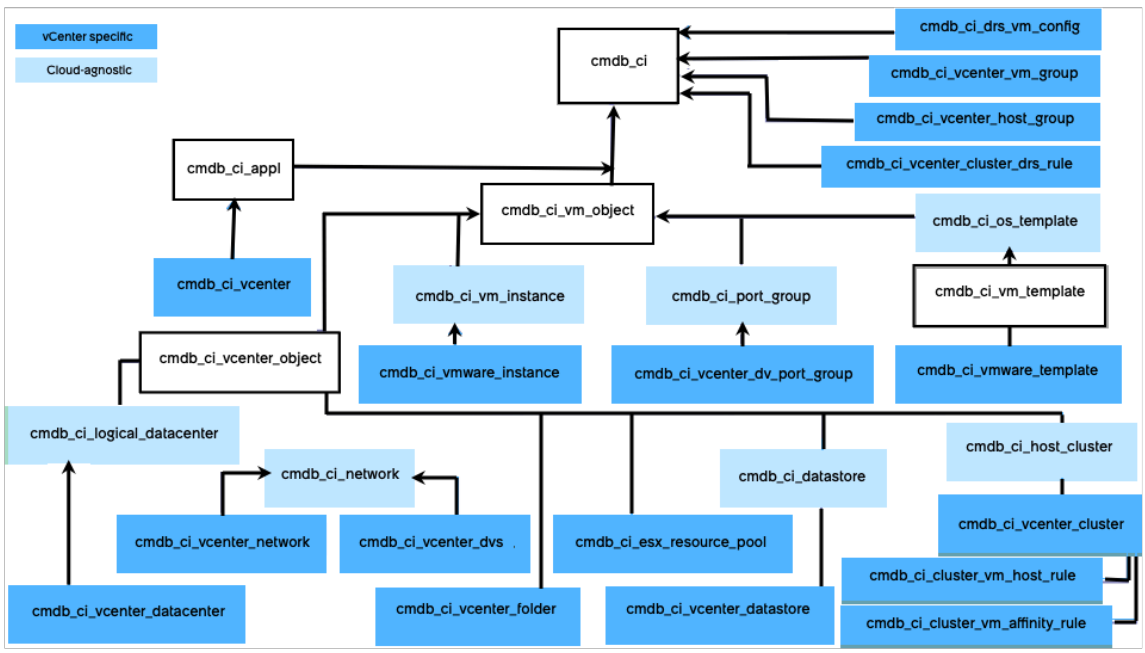

Several tables are cloud-agnostic tables, meaning that they can be populated for any cloud resource, not just vCenter resources. Look in the sub tables that extend the cloud-agnostic tables to find discovered configuration items (CIs).

Cloud-agnostic table	Look in this vCenter-specific table
Virtual Machine Instance [cmdb_ci_vm_instance]	VMware Virtual Machine Instance [cmdb_ci_vmware_instance]
Operating System Template [cmdb_ci_os_template]	VMware Virtual Machine Template [cmdb_ci_vmware_template]
Logical Datacenter [cmdb_ci_logical_datacenter]	VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud-agnostic table	Look in this vCenter-specific table
Cloud Networks [cmdb_ci_network]	VMware vCenter Network [cmdb_ci_vcenter_network]
Datastore [cmdb_ci_datastore]	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]
Host Cluster [cmdb_ci_host_cluster]	VMware vCenter Cluster [cmdb_ci_vcenter_cluster]

vCenter data

Discovery uses multiple vCenter probes to collect this data from vCenter. The data is saved in tables extend from the Configuration item [cmdb_ci] table.

VMware vCenter Instance [cmdb_ci_vcenter]

Field label and Name	Description
Name [name]	The display name of the vCenter instance.
IP Address [ip_address]	The IP address of the vCenter server.
API version [api_version]	The version of the vCenter API in use.
Full name [fullname]	The full vCenter instance name returned by VMware.
Instance UUID [instance_uuid]	A unique identifier assigned to the vCenter instance.
URL [url]	The endpoint URL used to connect to vCenter.
Fully qualified domain name [fqdn]	The fully qualified domain name of the vCenter server.

The combination of the vCenter instance in the cmdb_ci_vcenter table and the object ID in the cmdb_ci_vm_object table identifies a specific vCenter.

The tables extend from the Virtual Machine Objects [cmdb_ci_vm_object] table, which holds object IDs for all cloud-based resources, except virtualization servers.

DRS VM Config [cmdb_ci_drs_vm_config]

Field label and Name	Description
Cluster [cluster]	The vCenter cluster that the VM belongs to.
DRS behavior [drs_behavior]	The VM-level Distributed Resource Scheduler behavior that overrides cluster defaults.
DRS enabled [drs_enabled]	Indicates whether DRS is active for the VM.
Virtual machine [virtual_machine]	The VM instance using these DRS settings.

This table stores DRS behavior for the VMs that override the cluster behavior. For more information on DRS, see VMware's documentation on Distributed Resource Scheduler.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware vCenter VM Group [cmdb_ci_vcenter_vm_group]

Field label and Name	Description
Cluster [cluster]	The cluster that this VM group is associated with.

The VMware vCenter VM group [cmdb_ci_vcenter_vm_group] table has Contains::Contained by relationship with cmdb_ci_vmware_instance. This group has a set of VMs.

VMware vCenter Host Group [cmdb_ci_vcenter_host_group]

Field label and Name	Description
Cluster [cluster]	The cluster that this host group belongs to.

The VMware vCenter Host Group [cmdb_ci_vcenter_host_group] table has Contains::Contained by relationship with cmdb_ci_vmware_instance. This group has a set of hosts.

VMware vCenter Cluster DRS Rule [cmdb_ci_vcenter_cluster_drs_rule]

Field label and Name	Description
Rule UUID [rule_uuid]	Unique identifier for the DRS rule in vCenter.
Active [active]	Whether the rule is currently enforced.
Cluster [cluster]	The cluster that the rule applies to.

VMware Virtual Machine Instance [cmdb_ci_vmware_instance]

Field label and Name	Description
Name [name]	VM name as discovered from vCenter.
Template [template]	Indicates whether this VM is based on a template.
CPUs [cpus]	Number of assigned virtual CPUs.
Disks [disks]	Number of virtual disks attached to the VM.
Disks size (GB) [disks_size]	Combined size of all VM disks.
Memory (MB) [memory]	Memory allocated to the VM.
Network adapters [nics]	Number of virtual NICs attached.
Image path [image_path]	Path to the VM's image on the datastore.
State [state]	Current VM power or operational state.
vCenter Instance UUID [vcenter_uuid]	Unique ID of the vCenter instance managing this VM.
vCenter Reference [vcenter_ref]	Reference to the vCenter record in CMDB.
BIOS UUID [bios_uuid]	Hardware-level VM UUID from VMware.
VM Instance UUID [vm_instance_uuid]	VMware-provided VM unique ID.
IP address [ip_address]	Primary IP detected for the VM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware Distributed Virtual Port Groups [cmdb_ci_vcenter_dv_port_group]

Field label and Name	Description
Name [name]	Name of the port group.
vCenter Instance UUID [vcenter_uuid]	vCenter instance that manages this port group.
Distributed Virtual Switch Reference [dvs_ref]	Reference to the switch that the port group belongs to.

VMware Virtual Machine Template [cmdb_ci_vmware_template]

Field label and Name	Description
Name [The template name from VMWare.]	Template name in vCenter.
CPUs [cpus]	Number of CPUs assigned when deployed.
Disks size (GB) [disks_size]	Disk size of the template.
Guest ID [guest_id]	VMware guest OS identifier.
Memory (MB) [memory]	Allocated template memory.
Image path [image_path]	Path to the template image.
vCenter Instance UUID [vcenter_uuid]	vCenter instance where the template resides.
vCenter Reference [vcenter_ref]	CMDB reference to vCenter record.
Network adapters [nics]	Number of NICs defined in the template.
Base name [base_name]	The base name used for naming VMs created from this template.
VM namer [vm_namer]	VMware auto-naming method for derived VMs.
BIOS UUID [bios_uuid]	Template UUID.
VM Instance UUID [vm_instance_uuid]	Unique template instance ID.

VMware vCenter Object [cmdb_ci_vcenter_object]

Field label and Name	Description
Managed object reference ID [morid]	Unique Managed Object Reference ID in vCenter.
vCenter Instance UUID [vcenter_uuid]	vCenter that owns this object.
vCenter Reference [vcenter_ref]	CMDB reference to the vCenter record.

VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]*

Field label and Name	Description
Region [region]	Geographic or logical region of the datacenter.
Managed object reference ID [morid]	VMware Managed Object Reference ID.
vCenter Instance UUID [vcenter_uuid]	vCenter that manages this datacenter.
vCenter Reference [vcenter_ref]	CMDB reference back to vCenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]* (continued)

Field label and Name	Description
Top-level folder for VMs [folder_moreid]	Top-level folder where VMs reside.
Top-level folder for hosts [host_morid]	Top-level folder for hosts.

*The VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter] table is extended from the [cmdb_ci_logical_datacenter] table.

VMware vCenter Network [cmdb_ci_vcenter_network]*

Field label and Name	Description
Managed object reference ID [morid]	Object reference ID in vCenter.
vCenter Instance UUID [vcenter_uuid]	Parent vCenter instance.
vCenter Reference [vcenter_ref]	CMDB reference to vCenter.
Network accessible by either hosts or virtual machines [accessible]	Whether hosts or VMs can connect to the network.

*The VMware vCenter Network [cmdb_ci_vcenter_network] table is extended from the [cmdb_ci_network] table.

VMware Distributed Virtual Switch [cmdb_ci_vcenter_dvs]

Field label and Name	Description
VMCount [vm_count]	Number of VMs attached to the switch.
Host count [host_count]	Number of ESXi hosts connected.

VMware vCenter Folder [cmdb_ci_vcenter_folder]

Field label and Name	Description
Managed object reference ID [morid]	VMware folder identifier.
vCenter Instance UUID [vcenter_uuid]	vCenter managing the folder.
vCenter Reference [vcenter_ref]	CMDB reference to vCenter.
Full path [fullpath]	Full folder hierarchy path.

ESX Resource Pool [cmdb_ci_esx_resource_pool]

Field label and Name	Description
Managed object reference ID [morid]	VMware Managed Object Reference ID.
vCenter Instance UUID [vcenter]	vCenter that manages this pool.
vCenter Reference [vcenter_ref]	CMDB reference.
Owner [owner]	The VM or host that owns the pool.
Owner Managed Object Reference ID [owner_morid]	Managed Object Reference of the owner.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ESX Resource Pool [cmdb_ci_esx_resource_pool] (continued)

Field label and Name	Description
CPU expandable [cpu_expandable]	Whether CPU resources can exceed limits.
CPU limit (MHz) [cpu_limit_mhz]	Max CPU allocation in MHz.
CPU reserved (MHz) [cpu_reserved_mhz]	CPU reserved for the pool.
CPU shares [cpu_shares]	CPU scheduling weight.
Full path [fullpath]	Location path in vCenter hierarchy.
Memory expandable [mem_expandable]	Whether memory can exceed the configured maximum.
Memory limit (MB) [mem_limit_mb]	Memory limit in MB.
Memory reserved (MB) [mem_reserved_mb]	Memory guaranteed for the pool.
Memory shares [mem_shares]	Memory scheduling weight.

VMware vCenter Datastore [cmdb_ci_vcenter_datastore]

Field label and Name	Description
Managed object reference ID [morid]	ID of the datastore in vCenter.
vCenter Instance UUID [vcenter_uuid]	Parent vCenter ID.
Accessible [accessible]	Indicates if hosts can access it.
vCenter Reference [vcenter_ref]	CMDB reference.
Type [type]	Filesystem type (NFS, VMFS, etc.).
Capacity (GB) [capacity]	Total datastore capacity.
Free space (GB) [freespace]	Remaining space.
URL [url]	Unique locator.
Clustered [clustered]	Whether it belongs to a storage pod.

The VMware vCenter Datastore [cmdb_ci_vcenter_datastore] table is extended from the [cmdb_ci_datastore] table.

VMware vCenter Cluster [cmdb_ci_vcenter_cluster]*

Field label and Name	Description
Managed object reference ID [morid]	Cluster ID.
vCenter Instance UUID [vcenter_uuid]	Parent vCenter ID.
Effective hosts [effectivehosts]	Number of functioning hosts.
vCenter Reference [vcenter_ref]	CMDB reference.
Effective CPU [effectivecpu]	Available CPU capacity.
Effective memory [effectivememory]	Available RAM.
Number of effective hosts [effectivehosts]	Count of hosts contributing resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware vCenter Cluster [cmdb_ci_vcenter_cluster]* (continued)

Field label and Name	Description
Number of hosts [numhosts]	Total hosts.
Total CPU [totalcpu]	Total CPU capacity.
Total memory [totalmemory]	Total memory.
Number of CPU cores [numcpucores]	Total core count.
Number of CPU threads [numcputhreads]	Total thread count.
DRS Behavior [drs_behavior]	Cluster-level DRS settings.
DRS Enabled [drs_enabled]	Whether DRS is active.
DRS VMotion Rate [drs_vmotion_rate]	vMotion automation aggressiveness level.

*The VMware vCenter Cluster [cmdb_ci_vcenter_cluster] table is extended from the [cmdb_ci_host_cluster] table.

Cluster VM Host Rule Info [cmdb_ci_cluster_vm_host_rule]

Field label and Name	Description
VM Group [vm_group]	VM group to which the rule applies.
Host Group [host_group]	Host group associated with the rule.
Mandatory [mandatory]	Indicates whether the rule must be strictly enforced.
Affinity [affinity]	Whether VMs must run together or be separated.

Cluster VM Affinity Rule Info [cmdb_ci_cluster_vm_affinity_rule]

Field label and Name	Description
Affinity [affinity]	Defines whether grouped VMs must stay together.

The Cluster VM Affinity Rule Info [cmdb_ci_cluster_vm_affinity_rule] table has Contains::Contained by relationship with cmdb_ci_vmware_instance. This group has a set of VMs for which this affinity is applied.

VMWare Distributed Virtual Switch [cmdb_ci_vcenter_dvs]

Field label and Name	Description
Name [name]	The name of the distributed virtual switch as defined in vCenter.
VMCount [vm_count]	The number of virtual machines currently connected to this distributed switch.
Host count [host_count]	The number of ESXi hosts participating in this distributed switch.

Discovery also maps the relationships to VMs and to distributed virtual port groups.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware Distributed Virtual Port Group [cmdb_ci_vcenter_dv_port_group]

Field label and Name	Description
Name [name]	The name of the port group that provides network configuration for connected VMs.
vCenter Instance UUID [vcenter_uuid]	The VMware-generated UUID of the vCenter instance that manages this port group.
Distributed Virtual Switch Reference [dvs_ref]	The switch to which this port group belongs.

Discovery also maps the relationship to the distributed virtual switch.

You can add these related lists to view additional discovered data:

- **Storage Volumes**

  The virtual disks for this virtual machine. This data is saved in the Storage Volume [cmdb_ci_storage_volume] table with the value `type=v Disk`.

Storage Volume [cmdb_ci_storage_volume]

Field label and Name	Description
Name [name]	Name of the storage volume.
Size [size]	Size of the disk/volume.

- **Network Adapters**

  The virtual network adapters for the virtual disks. This data is saved in the Network Adapter [cmdb_ci_network_adapter] table.

Network Adapter [cmdb_ci_network_adapter]

Field label and Name	Description
Name [name]	NIC name.
IP Address [ip_address]	Assigned IP.
MAC address [mac_address]	NIC MAC address.
Netmask [netmask]	Subnet mask.

vCenter discovery with Software Asset Management

If Software Asset Management is active, Discovery populates these vCenter tables using the VMWare - vCenter ESX Hosts License probe.

VMWare License Key [samp_vmware_license_key]

Field label and Name	Description
vCenter Reference [vcenter_ref]	The reference link to the vCenter instance that owns this license.
Cost Unit [cost_unit]	The unit used by VMware to measure license cost.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMWare License Key [samp_vmware_license_key] (continued)

Field label and Name	Description
Edition [edition]	The edition level of the VMware product.
Features [features]	The list of capabilities enabled by this license.
License Key [license_key]	The VMware-issued alphanumeric software license key.
Product Name [product_name]	The VMware product that the license applies to.
Product Version [product_version]	The version of the licensed product.
Rights Owned [rights_owned]	Total entitlements purchased under this license.
Rights Used [rights_used]	Entitlements currently consumed by vCenter hosts.

VMWare License Key Usage [samp_vmware_license_key_usage]

Field label and Name	Description
Rights Used [rights_used]	Number of license units currently in use.
Expiration Date [expiration_date]	The date the license expires.
Used Features [used_features]	The specific product features being used.
License Key [license_key]	The license key to which this usage record applies.
Software Install [software_install]	The installation record tied to the consumed license.
Used By [used_by]	The asset or component consuming the license.

vCenter relationships

Discovery automatically creates relationships for vCenter components using data from a key class. Subsequent discoveries use the same key class to automatically validate and remove relationships that are no longer valid.

vCenter CIs can be members of folders or clusters, which affect how Discovery creates their relationships.

- If a CI is in a folder, Discovery creates a relationship between that CI and the folder. If that CI isn't in a folder, Discovery creates the relationship between the CI and the datacenter. These vCenter CIs can be in a folder:
  - VM Instance
  - VM Template
  - vCenter Network
  - Datastore
  - vCenter Folder
  - vCenter Cluster

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If an ESX server is in a cluster, Discovery creates a relationship between the ESX server and the cluster. If an ESX server isn't a member of a cluster, then Discovery creates a relationship to the datacenter.
- If a resource pool is in a cluster, Discovery creates a relationship between the resource pool and the cluster. If the resource pool isn't a member of a cluster, then Discovery creates a relationship to the ESX server.

This diagram illustrates vCenter relationships:

vCenter relationships

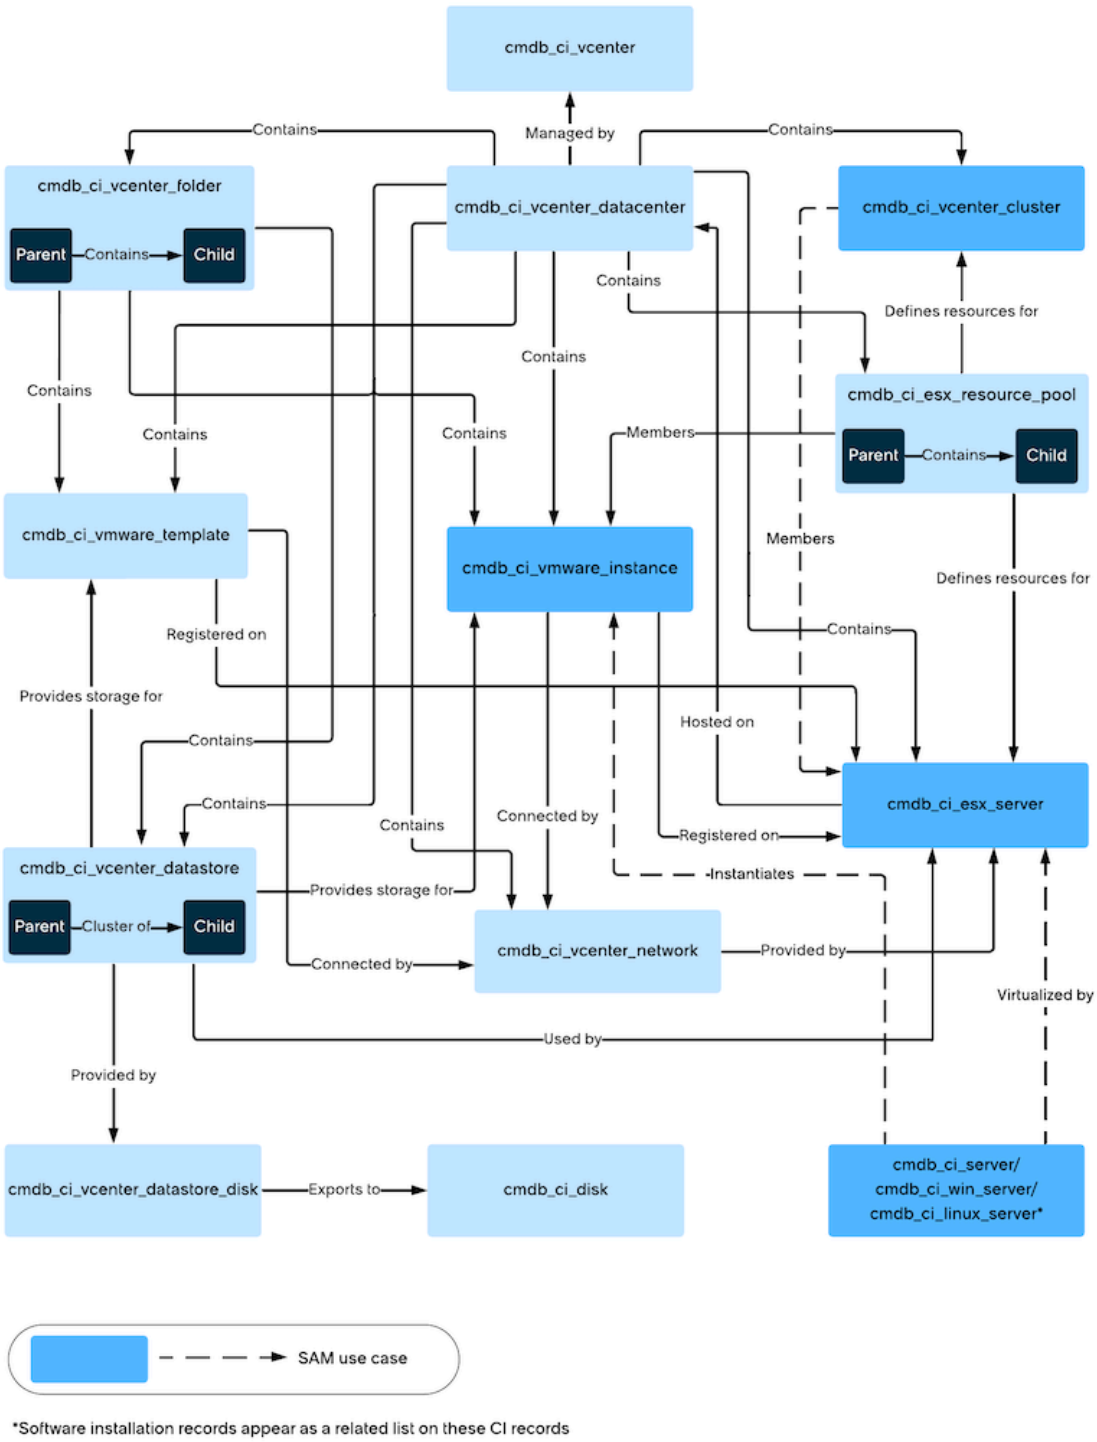

*Software installation records appear as a related list on these CI records

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
One of the following server CIs: • Server [cmdb_ci_server] • Windows Server [cmdb_ci_win_server] • Linux Server [cmdb_ci_linux_server]	Virtualized by::Virtualizes	ESX Server [cmdb_ci_esx_server]*†
One of the following server CIs: • Server [cmdb_ci_server] • Windows Server [cmdb_ci_win_server] • Linux Server [cmdb_ci_linux_server]	Instantiates::Instantiated by	VM Instance [cmdb_ci_vmware_instance]*†
ESX Server [cmdb_ci_esx_server]	Hosted on::Hosts	VMware vCenter datacenter [cmdb_ci_vcenter_datacenter]
VMware vCenter Network [cmdb_ci_vcenter_network]	Provided by::Provides	ESX Server [cmdb_ci_esx_server]
VMware vCenter Cluster [cmdb_ci_vcenter_cluster]	Members::Member of	ESX Server [cmdb_ci_esx_server]*
VMware Virtual Machine Instance [cmdb_ci_vmware_instance]	Registered on::Has registered	ESX Server [cmdb_ci_esx_server]
VMware Virtual Machine Instance [cmdb_ci_vmware_instance]	Connected by::Connects	VMware vCenter Network [cmdb_ci_vcenter_network]
Virtual Machine Template [cmdb_ci_vmware_template]	Connected by::Connects	VMware vCenter Network [cmdb_ci_vcenter_network]
Virtual Machine Template [cmdb_ci_vmware_template]	Registered on::Has registered	ESX Server [cmdb_ci_esx_server]
VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	Used by::Uses	ESX Server [cmdb_ci_esx_server]
VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	Provides storage for::Stored on	Virtual Machine Template [cmdb_ci_vmware_template]
VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	Provides storage for::Stored on	VMware Virtual Machine Instance [cmdb_ci_vmware_instance]
VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	Provided by::Provides	Datastore Disk [cmdb_ci_vcenter_datastore_disk]
VMware vCenter Datastore [cmdb_ci_vcenter_datastore] (parent)	Cluster of::Cluster	VMware vCenter Datastore [cmdb_ci_vcenter_datastore] (child)
ESX Resource Pool [cmdb_ci_esx_resource_pool]	Defines resources for::Get resources from	VMware vCenter Cluster [cmdb_ci_vcenter_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
ESX Resource Pool [cmdb_ci_esx_resource_pool]	Defines resources for::Get resources from	ESX Server [cmdb_ci_esx_server]
ESX Resource Pool [cmdb_ci_esx_resource_pool]	Members::Member of	VMware Virtual Machine Instance [cmdb_ci_vmware_instance]
ESX Resource Pool [cmdb_ci_esx_resource_pool] (parent)	Contains::Contained by	ESX Resource Pool [cmdb_ci_esx_resource_pool] (child)
VMware vCenter Folder [cmdb_ci_vcenter_folder]	Contains::Contained by	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]
VMware vCenter Folder [cmdb_ci_vcenter_folder]	Contains::Contained by	Virtual Machine Template [cmdb_ci_vmware_template]
VMware vCenter Folder [cmdb_ci_vcenter_folder]	Contains::Contained by	VMware Virtual Machine Instance [cmdb_ci_vmware_instance]
VMware vCenter Folder [cmdb_ci_vcenter_folder] (parent)	Contains::Contained by	VMware vCenter Folder [cmdb_ci_vcenter_folder] (child)
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Managed by::Manages	VMware vCenter Instance [cmdb_ci_vcenter]
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	VMware vCenter Folder [cmdb_ci_vcenter_folder]
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]
VMware vCenter datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	VMware vCenter Network [cmdb_ci_vcenter_network]
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	Virtual Machine Template [cmdb_ci_vmware_template]
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	VMware Virtual Machine Instance [cmdb_ci_vmware_instance]
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	ESX Resource Pool [cmdb_ci_esx_resource_pool]
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	ESX Server [cmdb_ci_esx_server]
VMware vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Contains::Contained by	VMware vCenter Cluster [cmdb_ci_vcenter_cluster]
Datastore Disk [cmdb_ci_vcenter_datastore_disk]	Exports to::Imports from	Disk [cmdb_ci_disk]

* These relationships are part of Software Asset Management (ServiceNow Otto for SAM) use case.

† The business rule "Virtual Computer Check" creates the relationships between the listed server CIs and the ESX Server and VM Instance. The server CI must be discovered after the vCenter is discovered. This triggers the business rule and creates the relationships.

Note: Software installation records appear as a related list on Server [cmdb_ci_server], Windows Server [cmdb_ci_win_server], and Linux Server [cmdb_ci_linux_server] CI records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Management relationships

These additional relationships are created when Cloud Management (CMP) is active.

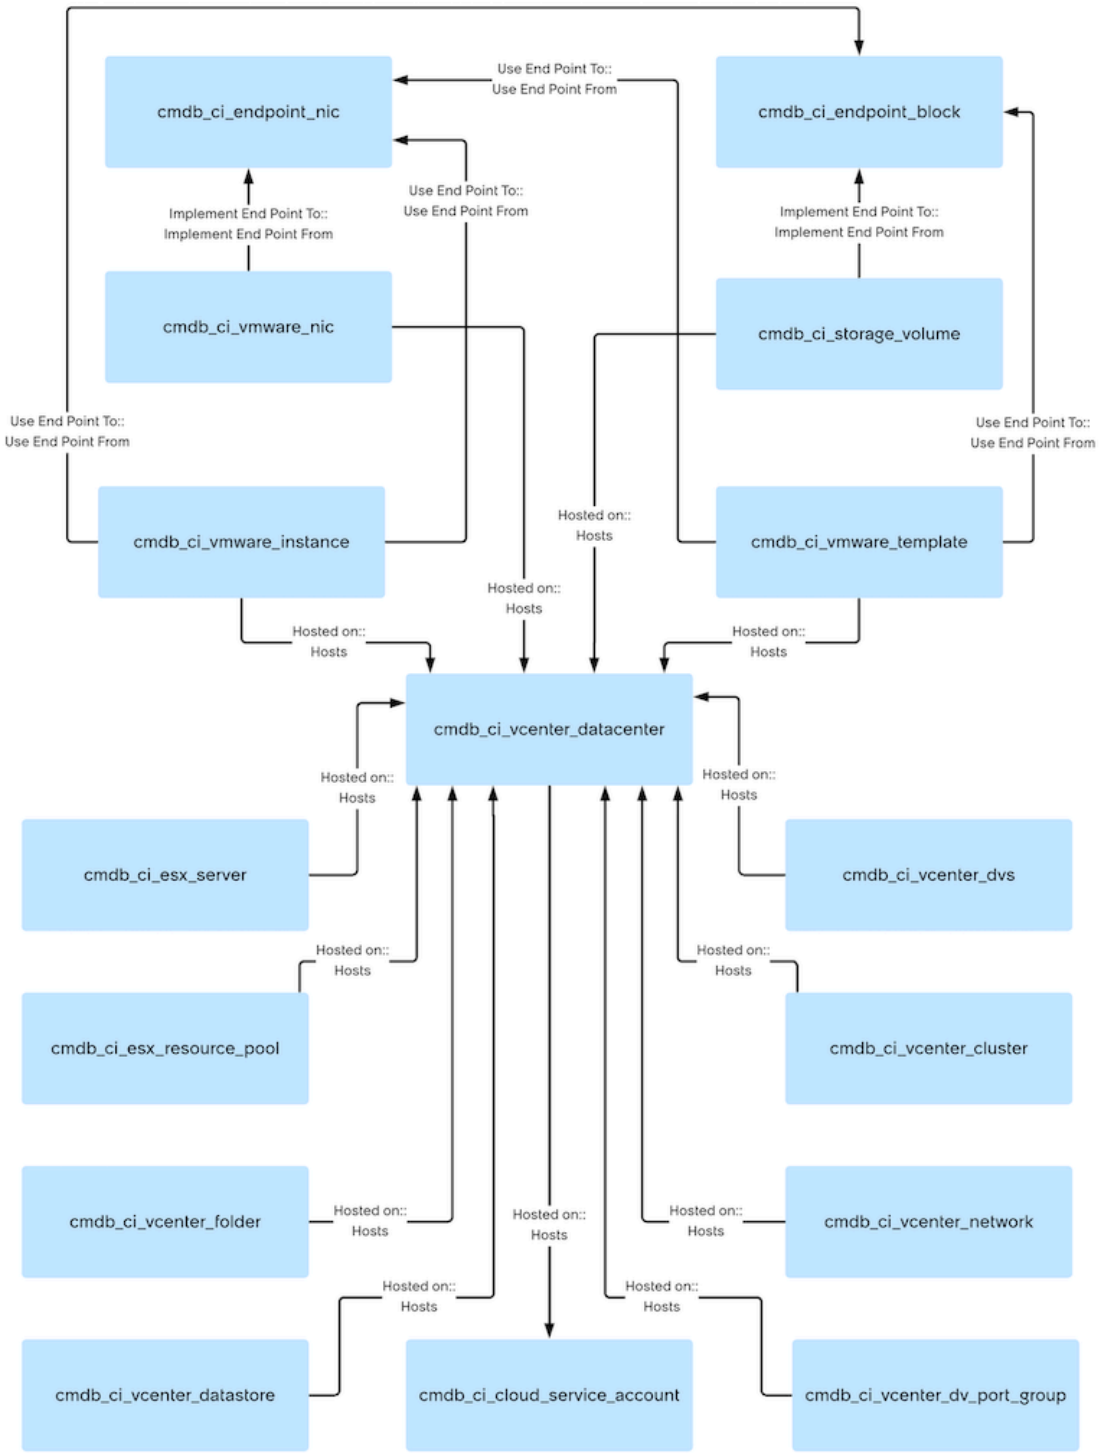

Parent class	Relationship type	Child class
ESX Server [cmdb_ci_esx_server]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parent class	Relationship type	Child class
vCenter Folder [cmdb_ci_vcenter folder]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
vCenter Datacenter [cmdb_ci_vcenter_datacenter]	Hosted on::Hosts	Cloud Service Account [cmdb_ci_cloud_service_account]
vCenter Cluster [cmdb_ci_vcenter_cluster]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
Resource Pools [cmdb_ci_esx_resource_pool]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
VM Instance [cmdb_ci_vmware_instance]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
VM Template [cmdb_ci_vmware_template]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
vCenter Network [cmdb_ci_vcenter_network]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
Distributed Virtual Switch [cmdb_ci_vcenter_dvs]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
Distributed Virtual Port Group [cmdb_ci_vcenter_dv_port_group]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
Datastore [cmdb_ci_vcenter_datastore]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
Virtual Disk [cmdb_ci_storage_volume]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
Virtual NIC [cmdb_ci_vmware_nic]	Hosted on::Hosts	vCenter Datacenter [cmdb_ci_vcenter_datacenter]
VM Template [cmdb_ci_vmware_template]	Use End Point To::Use End Point From	Block Endpoint [cmdb_ci_endpoint_block]
VM Instance cmdb_ci_vmware_instance	Use End Point To::Use End Point From	Block Endpoint [cmdb_ci_endpoint_block]
Virtual Disk [cmdb_ci_storage_volume]	Implement End Point To::Implement End Point From	Block Endpoint [cmdb_ci_endpoint_block]
Virtual NIC [cmdb_ci_vmware_nic]	Implement End Point To::Implement End Point From	NIC Endpoint [cmdb_ci_endpoint_nic]
VM Template [cmdb_ci_vmware_template]	Use End Point To::Use End Point From	NIC Endpoint [cmdb_ci_endpoint_nic]
VM Instance [cmdb_ci_vmware_instance]	Use End Point To::Use End Point From	NIC Endpoint [cmdb_ci_endpoint_nic]

VMware tags

You can attach tags to vSphere objects, such as virtual machines, through the vSphere interface. The tags can then be grouped into categories. Objects with tags are sortable and searchable based on the parameters that you give the tags and categories. The VMWare — vCenter VM Tags probe discovers these tags. You can view the discovered tags by opening the virtual machine record in **cmdb_ci_vmware_instance**. Then go to the **Key Values** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key Values [cmdb_key_value]

Field label and Name	Description
Configuration item [configuration_item]	Referenced VM with tag attached in vCenter.
Key [key]	Category name of the tag applied to the resource in vCenter.
Tag [tag]	Source in vCenter from which tags are fetched. In this case, the value is **Tags**.
Value [value]	Tag name that's applied to the resource in vCenter.

Datastore data collected by Discovery

Discovery identifies each datastore in the system and creates the relationships with the virtual machines and the ESX servers that use these datastores. Discovery uses the VMWare - vCenter Datastores probe to collect this data.

Field label and Name	Table Name	Description
Accessible [accessible]	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	Whether the datastore is collected or not
Capacity (GB) [capacity]	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	Amount of space provided by the datastore.
Clustered [clustered]	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	When the datastore is clustered (belongs to a storage pod).
Free space (GB) [freespace]	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	Amount of space still available on the datastore.
Type [type]	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	The type of file system volume, such as WFS or NFS.
URL [url]	VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	The unique URL locator for the datastore.

HostMounts

Discovery uses both the VMWare - vCenter ESX Hosts and VMWare - vCenter Datastores probes to collect datastore host mount data.

Field label and Column	Table	Description
Accessible [accessible]	VMware Datastore HostMount [vcenter_datastore_hostmount]	Whether the ESXi host can access the datastore.
Access Mode [access_mode]	VMware Datastore HostMount [vcenter_datastore_hostmount]	The mode in which the ESXi host accesses the datastore.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field label and Column	Table	Description
VMware vCenter Datastore [datastore]	VMware Datastore HostMount [vcenter_datastore_hostmount]	The datastore mounted on this ESXi host.
ESX Server [esx_server]	VMware Datastore HostMount [vcenter_datastore_hostmount]	The ESXi host that the datastore is mounted on.
vCenter Reference [vcenter_ref]	VMware Datastore HostMount [vcenter_datastore_hostmount]	The vCenter instance that manages this mount.

Datastore Discovery

A datastore is a collection of one or more physical disks, such as an iSCSI disk, and can be used by more than one ESXi host. However, a physical disk can only connect to one datastore. Because ESXi hosts can share datastores, it's easy to move virtual machines between hosts that have a common datastore.

Note: From the perspective of a virtual machine attached to a datastore, storage is provided by a single disk.

The advantages to connecting a datastore to multiple disks are:

- The ability to mirror the disks for failover protection.
- The ease of adding storage to the datastore.

For information about the tables used by Discovery to store data for physical disks and their relationship to datastores and ESXi hosts, see, ESXi server discovery.

In this example configuration, two ESXi hosts share a common datastore that uses different types of storage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Datastore Discovery with different storage types

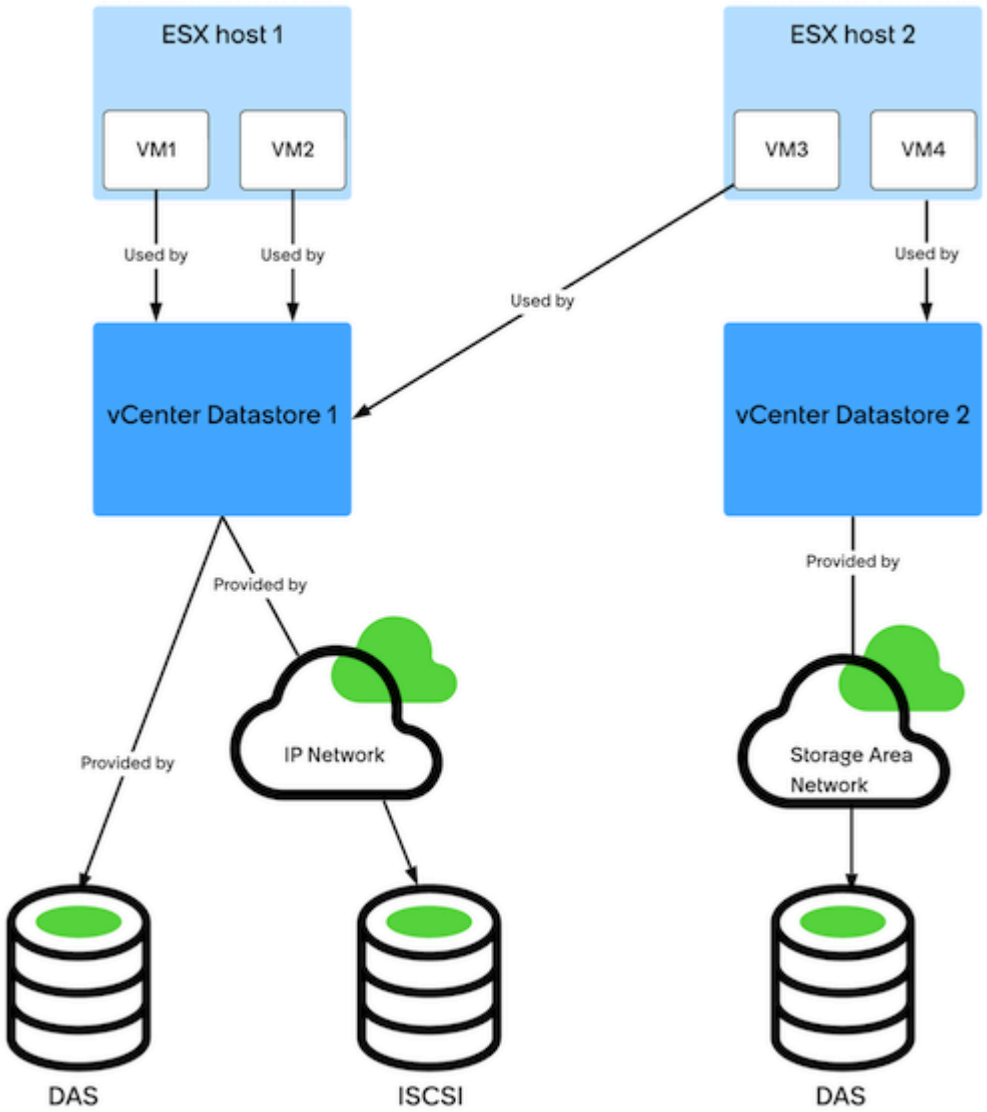

Relationships

ServiceNow provides tables that contain the relationships between an ESXi host and its datastore and the specific disks to which the datastore is connected.

ServiceNow Discovery establishes the relationships between a datastore, the disks attached to the datastore, and the ESXi server that hosts the virtual machines using that datastore. From the perspective of the ESXi host, iSCSI and fibre channel disks connected to the datastore are treated as physical disks. Discovery doesn't show the direct relationship of the storage disks to the ESXi host.

Note: Storage can be hosted on computers that aren't discovered by Discovery. ESXi hosts are discovered through vCenter, and storage is discovered separately through CIM. The system can only establish the relationship between the two when storage is discovered before ESXi.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Datastore tables

Table	Description
cmdb_ci_vcenter_datastore_disk	Stores the relationship of the physical disk to the datastore.
vcenter_datastore_hostmount	Stores the relationship between the datastore and the ESXi server to which it's connected.
cmdb_ci_disk	Contains the physical disks connected directly to the datastore. This table also contains a reference to the ESXi host.
cmdb_ci_fc_disk	Contains the fibre channel disks in a storage area network (SAN) connected to the datastore. This table also contains a reference to the ESXi host.
cmdb_ci_iscsi_disk	Contains the iSCSI disks in an IP network connected to the datastore. This table also contains a reference to the ESXi host.

VM Instance state and status fields

These tables represent the state and status of the cmdb_ci_vm_instance in various flows such as vCenter Discovery and vCenter Events and the business rules which are triggered.

vCenter Discovery that change various states and status fields related to Virtual Machine Instance

Source	Script or Event name	Field name	VM previous state	VM new state
vCenter Discovery	VCenterVMsSensor	Install Status	Retired	Installed
vCenter Discovery	VCenterDatacentersSensor	Install Status State Operational status	Installed	Retired Terminated Non-operational

vCenter Events that change various states and status fields related to Virtual Machine Instance

Source	Script or Event name	Field name	VM previous state	VM new state
vCenter events	VCenterVmStateUpdater powerOnEvtsToUse = ['VmPoweredOffEvent', 'VmPowerOffOnIsolationEvent', 'VmShutdownOnIsolationEvent', 'VmSuspendedEvent'],	Install Status State	Installed Any	Installed ON

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vCenter Events that change various states and status fields related to Virtual Machine Instance (continued)

Source	Script or Event name	Field name	VM previous state	VM new state
vCenter events	powerOffEvtsToUse = ['VmPoweredOnEvent', 'DrsVmPoweredOnEvent', 'VmRestartedOnAlternateHostEvent', 'VmSuspendedEvent'],	Install Status State	Installed Any	Installed OFF
vCenter events	suspendEvtsToUse = ['VmPoweredOnEvent', 'DrsVmPoweredOnEvent', 'VmRestartedOnAlternateHostEvent', 'VmPoweredOffEvent', 'VmPowerOffOnIsolationEvent']	Install Status State	Installed Any	Installed Paused
vCenter events	VmRemovedEvent	Install Status State	Installed Any	Retired Terminated

Business rules that change various states and status fields related to Virtual Machine Instance and related Server

Business rule	Field name	Plugin name
Cascade Operational Status to vminstance	cmdb_ci_server: Operational status	Pattern Designer
Sync Ops Status for CMDB CI	• cmdb_ci: Operational status, Install Status • cmdb_ci_hardware: Operational status, Hardware Status, Substatus	CMDB

ESXi server discovery

Discovery identifies and classifies information about ESXi servers and ESXi resource pools through the discovery of vCenter and not from the direct discovery of any ESXi servers.

Important: ESXi server discovery is done through vCenter. Don't specify the IP address of the ESXi server in a Discovery Schedule. Instead, discover the vCenter through the Discovery Schedule.

Required Roles

Users with the *itil* and *asset* roles can access ESXii and ESXi configuration item (CI) records. To run discovery on vCenter servers, users must have the *discovery_admin* role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credentials

To run a complete Discovery of vCenter/ESXi servers, you need vCenter credentials. If the vmapp port probe is turned off, you must use Windows credentials to access the Windows host on which the vCenter server runs.

Note: Make sure to select a credential *Type* of *VMware*.

ESXi server Discovery Components

Discovery identifies ESXi servers based on the correlation ID (BIOS UUID), when the hardware manufacturer is on a certified inclusion list. So if the manufacturer is in the compatible manufacturers list, the correlation-id must be unique. If the manufacturer isn't on the certified inclusion list, we check for the Managed Object Reference ID (MORID) and Serial Number as well. After running the vCenter classifier, Discovery launches the VMware - vCenter datacenters probe, which launches the probes that explore the ESXi server. For the complete list of vCenter probes, see List of Discovery probes.

Component and Name	Description
Classifier [vCenter]	Classifies standalone vCenter servers.
Probe [VMWare - vCenter ESX Hosts]	Creates records for ESXi servers and host mounts. Creates relationships between ESXi servers and vCenter components. Triggers probes for storage Discovery.
Probe [VMWare - vCenter ESX Hosts Storage]	Creates records for ESXi host hardware: network adapters, disks, HBAs, FC ports, iSCSI, and FC disks. Creates relationships between DAS/iSCSI/FC disks and datastore disks.

Data collected

Basic server data from ESXi hosts is collected by the VMware - vCenter ESX Hosts probe.

Data collected by Discovery on ESXi hosts

Field label and Name	Table Name	Description
Operating System [os]	cmdb_ci_esx_server	The OS version running on the ESXi host.
OS Version [os_version]	cmdb_ci_computer	Detailed OS build or release information.
Name [name]	cmdb_ci_esx_server	Hostname of the ESXi server.
DNS domain [dns_domain]	cmdb_ci_esx_server	Domain associated with the ESXi host.
Manufacturer [manufacturer]	cmdb_ci_computer	Hardware vendor of the ESXi host.
Serial number [serial_number]	cmdb_ci_computer	Physical server hardware serial number.
CPU type [cpu_type]	cmdb_ci_esx_server	The CPU model used by the ESXi host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery on ESXi hosts (continued)

Field label and Name	Table Name	Description
CPU speed (MHz) [cpu_speed]	cmdb_ci_esx_server	Clock speed of the processors.
CPU count [cpu_count]	cmdb_ci_esx_server	Number of physical CPUs.
CPU core count [cpu_core_count]	cmdb_ci_computer	Total number of cores across all CPUs.
CPU manufacturer [cpu_manufacturer]	cmdb_ci_esx_server	Vendor of the CPU hardware.
Model ID [model_id]	cmdb_ci_computer	Hardware model identifier.
RAM (MB) [ram]	cmdb_ci_esx_server	Total physical memory installed.
Disk space (GB) [disk_space]	cmdb_ci_esx_server	Total disk capacity available.
Type [type]	cmdb_ci_disk	Type of disk.
Model ID [model_id]	cmdb_ci_disk	Hardware identifier for the disk.
Disk space (GB) [disk_space]	cmdb_ci_disk	Disk capacity of the individual disk.
Name [name]	cmdb_ci_disk	Disk name as detected by ESXi.
Name [name]	cmdb_ci_network_adapter	Network adapter name.
IP address [ip_address]	cmdb_ci_network_adapter	IP assigned to the NIC.
MAC address [mac_address]	cmdb_ci_network_adapter	Hardware address of the NIC.
Netmask [netmask]	cmdb_ci_network_adapter	Subnet mask of the NIC.
Managed object reference ID [morid]	Visualization Server [cmdb_ci_virtualization_server]	Managed Object Reference ID.
Serial Number [serial_number]	Serial Number [cmdb_serial_number]	Serial number record for CMDB correlation.

Relationships

Discovery collects the following relationship data for ESXi servers.

Base Class	Relationship	Dependent Class
ESXi Resource Pool [cmdb_ci_esx_resource_pool]	Defines resources for:Gets resources from	ESXi Server [cmdb_ci_esx_server]
VM Instance [cmdb_ci_vmware_instance]	Registered on:Has registered	ESXi Server [cmdb_ci_esx_server]
ESXi Server [cmdb_ci_esx_server]	Provides storage for:Stored on	VM Template [cmdb_ci_vmware_template]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Base Class	Relationship	Dependent Class
ESXi Server [cmdb_ci_esx_server]	Provided by:Provides	Networks [cmdb_ci_vcenter_network]
ESXi Server [cmdb_ci_esx_server]	Members of:Members	Cluster [cmdb_ci_vcenter_cluster]
ESXi Server [cmdb_ci_esx_server]	Contains:Contained by	Datacenter [cmdb_ci_vcenter_datacenter]

ESXi resource pools

Resource pools are configured in vCenter and define the maximum amount of resources that virtual machines using that pool can consume. If the ESXi server has additional resources to spare, an ESXi server property enables resource pools to expand. The **Name** and **Owner** fields of each resource pool on the ESXi server must be configured within the ServiceNow AI Platform

in the ESXi Resource Pool [cmdb_ci_esx_resource_pool] table. When Orchestration for VMware executes its manual provisioning tasks, the provisioner must select the proper resource pool for the virtual server requested. Discovery finds resource pools on ESXi machines and populates the fields on the ESXi Resource Pool form automatically. For more information, see Configure ESXi resource pools.

ESXi resource pools require the Orchestration - VMware Support plugin.

Note: Confirm that vCenter and the ESXi server have been fully configured, including the creation of the templates and resource pools.

Data collected

Field label and Name	Table	Source	Description
CPU expandable [cpu_expendable]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	Whether the resource pool can exceed CPU allocation limits.
CPU limit (MHz) [cpu_limit_mhz]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	Maximum CPU allocation.
CPU reserved (MHz) [cpu_reserved_mhz]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	CPU guaranteed to the pool.
CPU shares [cpu_shares]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	CPU priority weight.
Full path [fullpath]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	vCenter folder path.
Memory expandable [mem_expandable]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter	Whether memory can exceed the configured maximum.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected (continued)

Field label and Name	Table	Source	Description
		Clusters probe	
Memory limit (MB) [mem_limit_mb]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	Maximum memory.
Memory reserved (MB) [mem_reserved_mb]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	Memory guaranteed to the pool.
Memory shares [mem_shares]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	Memory scheduling weight.
Owner [owner]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	Host or VM owning the pool.
Owner Managed Object Reference ID [owner_morid]	ESXi Resource Pool [cmdb_ci_esx_resource_pool]	VMWare - vCenter Clusters probe	VMware reference ID of the owner.

Standalone ESXi discovery

Standalone ESXi discovery supports discovery of individual ESXi servers that host virtual machines (VMs) and related components without a vCenter. Various CIs and relationships are discovered as part of a discovery schedule.

Required roles:

Users with the *itil* and *asset* roles can access ESXi configuration item (CI) records. To run a standalone ESXi discovery, users must have the *discovery_admin* role.

VMware credentials

To run a standalone ESXi discovery, you need VMware credentials. Create the credentials by navigating to **Discovery** > **Credentials** > **VMware Credentials.**

If you use a domain account to access the ESXi host, specify the domain with the user name in the credential record. Use one of the supported formats, such as **Domain\UserName**.

Note: The VMware credentials must have a read-only role in the ESXi host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Requirements

- Make sure the Discovery (com.snc.discovery) plugin is installed and activated and that you have upgraded to Brazil or later.
- Activate the ESXi trigger probe. Navigate to the Trigger Probe [trigger_probe_m2m] table. The esxi record is inactive by default. Mark Active as true to enable Standalone ESXi discovery.
- Create a Discovery schedule for the host with the appropriate IP address of the ESXi host.

Note: If both SSH and ESXi are triggered, SSH is launched first and may cause Discovery to complete with the message "ESX Discovery is only supported through the vCenter." In this case, open the UNIX - Classify probe and set ESX - OS inactive.

ESXi server Discovery Components

Discovery identifies ESXi servers based on the correlation ID (BIOS UUID), when the hardware manufacturer is on a certified inclusion list. If the manufacturer is on the list, the correlation ID must be unique. If the manufacturer isn't on the certified inclusion list, the Managed Object Reference ID (MORID) and Serial Number are checked as well.

After Shazzam runs, it checks for the port probe esxi. Discovery then launches the VMWare - Standalone ESXi Server probe, which then launches the probes that explore the ESXi server. Other existing Discovery probes are also launched. For the complete list of probes, see List of Discovery probes.

ESXi server Discovery Components

Component and Name	Description
IP Service [ESXi]	IP Service ESXi - VMWare VM console is defined for the Port 902.
Port Probe [esxi]	ESXi Server Appliance web user interface. It's triggered by the IP Service ESXi and it triggers the probe VMWare - Standalone ESXi Server.
Probe [VMWare - Standalone ESXi Server]	Probe to get information about an ESXi server.
Probe [VMWare - vCenter ESX Hosts]	Creates records for ESXi servers and host mounts. Triggers other probes. • VMware - vCenter ESX Hosts Storage • VMware - vCenter Datastores • VMware - vCenter Networks • VMware -vCenter VMs
Probe [VMWare - vCenter ESX Hosts Storage]	Creates records for ESXi host hardware: network adapters, disks, HBAs, FC ports, iSCSI, and FC disks. Creates relationships between DAS/iSCSI/FC disks and datastore disks.

Basic server data from ESXi hosts is collected by the VMware - vCenter ESX Hosts probe.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ESXi Standalone server data

Discovery uses multiple existing probes to collect this data from ESXi. The data is saved in various tables. Some of the CIs, which have the "server" field have a reference to ESXi Host (for example, cmdb_ci_esx_server).

VMware Virtual Machine Instance [cmdb_ci_vmware_instance]

Field label and Name	Description
Name [name]	VM name.
Memory (MB) [memory]	RAM assigned.
CPUs [cpus]	Number of vCPUs.
Disks [disks]	Number of VM disks.
Network adapters [nics]	Count of NICs.
Object ID [object_id]	Unique object ID.
Server [server]	ESXi host running the VM.
State [state]	Power state.
Correlation ID [correlation_id]	Unique VM identifier for correlation.
VM Instance UUID [vm_instance_uuid]	ESXi VM ID.
Status [install_status]	CMDB install status.

VMware vCenter Datastore [cmdb_ci_vcenter_datastore]

Field label and Name	Description
Name [name]	Datastore name.
Capacity (GB) [capacity]	Total storage size.
Free space (GB) [freespace]	Available storage.
Accessible [accessible]	Whether the datastore is reachable.
Type [type]	Filesystem type.
Object ID [object_id]	Unique datastore ID.
Server [server]	ESXi host providing the datastore.
URL [url]	Datastore path.
Status [install_status]	CMDB install status.

VMware vCenter Network [cmdb_ci_vcenter_network]

Field label and Name	Description
Name [name]	Network name.
Object ID [object_id]	Unique network ID.
Server [server]	ESXi host providing network connectivity.
Status [install_status]	CMDB install status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware Network Adapter [cmdb_ci_vmware_nic]

Field label and Name	Description
Name [name]	NIC name.
MAC Address [mac_address]	NIC hardware address.
IP Address [ip_address]	Assigned IP.
Netmask [netmask]	Subnet mask.
Configuration Item [cmdb_ci]	CI that owns this NIC.
Object ID [object_id]	Unique NIC ID.
Mac manufacturer [mac_manufacturer]	Vendor of the NIC.
DHCP Enabled [dhcp_enabled]	Whether DHCP is active.
Status [install_status]	CMDB install status.

VMware Datastore HostMount [vcenter_datastore_hostmount]

Field label and Name	Description
VMware vCenter Datastore [datastore]	Mounted datastore.
ESX Server [esx_server]	Host using the datastore.
Accessible [accessible]	Whether the datastore is accessible.
Access Mode [access_mode]	Read/write permissions.

Datastore Disk [cmdb_ci_vcenter_datastore_disk]

Field label and Name	Description
Name [name]	Disk name.
Manufacturer [manufacturer]	Disk vendor.
Location [location]	Disk physical/virtual location.
Description [short_description]	Short description.
Class [sys_class_name]	CMDB class.
Updated [sys_updated_on]	Timestamp of last update.
Maintenance schedule [maintenance_schedule]	Maintenance timing.
Correlation ID [correlation_id]	Identifier linking disk to datastore.
Datastore [datastore]	Associated datastore.
Status [install_status]	CMDB status.

ESX Resource Pool [cmdb_ci_esx_resource_pool]

Field label and Name	Description
Name [name]	Resource pool name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ESX Resource Pool [cmdb_ci_esx_resource_pool] (continued)

Field label and Name	Description
CPU reserved (MHz) [cpu_reserved_mhz]	Guaranteed CPU.
CPU limit (MHz) [cpu_limit_mhz]	Maximum CPU.
CPU shares [cpu_shares]	CPU priority.
Memory reserved (MB) [mem_reserved_mb]	Guaranteed memory.
Memory limit (MB) [mem_limit_mb]	Maximum memory.
Memory shares [mem_shares]	Memory weight.
Object ID [object_id]	Unique pool ID.
Server [server]	ESXi host.
Managed object reference ID [morid]	Managed Object Reference ID.
Status [install_status]	CMDB install status.

ESX Server [cmdb_ci_esx_server]

Field label and Name	Description
Name [name]	ESXi hostname.
Manufacturer [manufacturer]	Vendor name.
Model ID [model_id]	Server hardware model.
Operating System [os]	OS version.
OS Version [os_version]	Detailed OS version.
Description [short_description]	Summary information.
Class [sys_class_name]	CMDB class.
Status [install_status]	CMDB status.

Network Adapter [cmdb_ci_network_adapter]

Field label and Name	Description
Name [name]	NIC name.
Mac Address [mac_address]	NIC hardware address.
Netmask [netmask]	Subnet mask.
Configuration Item [cmdb_ci]	The CI owning this adapter.
Mac manufacturer [mac_manufacturer]	NIC vendor.
DHCP Enabled [dhcp_enabled]	Whether DHCP is used.
Status [install_status]	CMDB install status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk [cmdb_ci_disk]

Field label and Name	Description
Name [name]	Disk name.
Computer [computer]	System owning the disk.
Size [size]	Disk capacity.
Manufacturer [manufacturer]	Disk vendor.
Model ID [model_id]	Model identifier.
Status [install_status]	CMDB install status.

Storage HBA [cmdb_ci_storage_hba]

Field label and Name	Description
Name [name]	HBA name.
Model ID [model_id]	HBA model number.
Computer [computer]	ESXi host.
WWNN [wwnn]	Worldwide Node Name.
Status [install_status]	CMDB install status.

Fibre Channel Port [cmdb_ci_fc_port]

Field label and Name	Description
Name [name]	Port name.
WWNN [wwnn]	Node Name.
WWPN [wwpn]	Port Name.
Speed [speed]	Link speed.
Controller [controller]	HBA controller.
Computer [computer]	Host system.
Status [install_status]	CMDB status.

iSCSI Disk [cmdb_ci_iscsi_disk]

Field label and Name	Description
Name [name]	Disk name.
Computer [computer]	Host accessing the disk.
Size [size]	Disk size.
Provided by [provided_by]	Storage system provider.
IQN [iqn]	iSCSI Qualified Name.
Device LUN [device_lun]	Logical Unit Number.
Storage type [storage_type]	Storage type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

iSCSI Disk [cmdb_ci_iscsi_disk] (continued)

Field label and Name	Description
Status [install_status]	CMDB status.

Fibre Channel Disk [cmdb_ci_fc_disk]

Field label and Name	Description
Name [name]	Disk name.
Computer [computer]	Host using this disk.
Size [size]	Disk capacity.
Provided by [provided_by]	Storage array.
Device LUN [device_lun]	LUN identifier.
WWN [wwn]	Worldwide Name.
Status [install_status]	CMDB install status.

IP Address [cmdb_ci_ip_address]

Field label and Name	Description
IP Address [ip_address]	Assigned network address.
IP version [ip_version]	IPv4 or IPv6.
Netmask [netmask]	Subnet mask.
Nic [nic]	NIC owning the IP.
Status [install_status]	CMDB status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships

Standalone ESXi discovery relationships

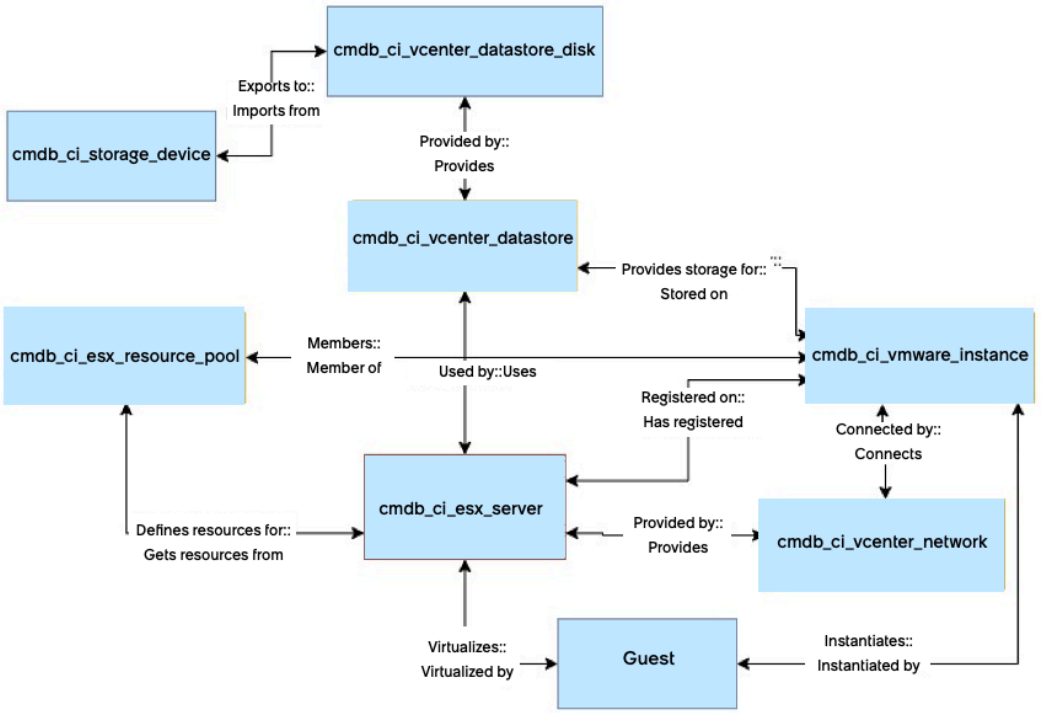

Resource pools

Standalone ESXi discovery also fetches the resource pools on the host including the root resource pool. This root resource pool is hidden for every ESXi host. The root resource pool may not be visible in the VSphere web client for the ESXi host, but you can view it using the mob browser.

Navigate to this URL: **<domain name/or ip_address>/mob/?moid=ha-root-pool**

The root resource pool groups the resources of that host. Other child resource pools can also be created from the root resource pool. The root is identified in the ESXi host with the Managed Object ID: ha-root-pool.

Forward migration

If you were using standalone ESXi discovery and now the same ESXi is part of vCenter, you can use vCenter discovery instead. Create a vCenter discovery schedule and trigger it. Triggering a vCenter discovery creates duplicate CIs in the following tables as the identifiers for the CIs are different when ESXi is standalone or part of vCenter:

- VMware VCenter Network [cmdb_ci_vcenter_network]
- ESX Resource Pool [cmdb_ci_esx_resource_pool]
- VMware VCenter Datastore [cmdb_ci_vcenter_datastore]
- Datastore Disk [cmdb_ci_vcenter_datastore_disk]

To avoid duplicates, you must mark the CIs created by standalone ESXi discovery in the above four tables as retired. When vCenter discovery is triggered, the vCenterESXHostsSensor script

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

include checks for all the ESXi servers whether they were previously discovered as standalone ESXi server. If yes, it automatically triggers the ESXMigrationUtil script to mark all the previously discovered duplicate CIs as retired.

Note: If you want to trigger migration manually, you can do so by executing the following script from the background script: // @params esx_sys_ids – array of sys ids of all the ESXi servers that must be migrated.

```
ESXMigrationUtil.
 retireCIsForESXForwardMigration(esx_sys_ids)
```

Once an ESXi server is migrated to vCenter, triggering a standalone ESXi discovery schedule on the same ESXi host results in an error. Discovery is aborted with an error message that "This ESXi is part of vCenter <IP_address of Vcenter> discovery schedule. Aborting discovery".

Operating systems discovery

Discovery identifies the following computers, clusters, and virtual machines.

AIX server discovery

Discovery identifies and classifies information about AIX servers.

Pattern-based discovery and mapping requirements

Configure SSH credentials

For more information, see SSH credentials.

Verify that you're using at least version 7.5.102 of OpenSSH

Use OpenSSH starting with version 7.5.102 for AIX machines since lower versions of OpenSSH can impact the discovered devices that are encountered with Discovery.

Verify privileged commands for AIX discovery

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Reference the main CI on discovered IP addresses

Starting with Discovery and Service Mapping Patterns version 1.35.0, you can populate the **Owned By Configuration Item** field on discovered IP address records with a reference to the main CI by setting the *add.owned.by.attribute* property to true. For more information, see Reference the main CI on discovered IP addresses.

Discover only the latest OS patch versions

Starting with Visibility Content version 6.35.0, you can collect only the current operating system patch level by setting the *sn_itom_pattern.discover_latest_os_patches* MID Server property to true. For more information, see Discover only the latest operating system patch versions.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier, probes, and pattern

Classifier	Probes	Pattern
AIX (UNIX Classification)	• Horizontal discovery probe: launches patterns • AIX - Installed Software^ • AIX - ADM^ • AIX - Identity* • AIX - CPU* • AIX - Hardware Model* • AIX - Memory* • AIX - OS Patches* • AIX - OS Release* • AIX - OS Revision* • UNIX - Find FQDN • UNIX - OS Filesystems • UNIX - OS Uptime	AIX server

*These probes aren't active on the classifier, as Discovery uses patterns by default for these discoveries.

^These probes remain active by default, even when Discovery uses pattern discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Label	Table Name	Field Name	Source
Operating System	cmdb_ci_computer	os	uname
OS version	cmdb_ci_computer	os_version	oslevel
OS service pack	cmdb_ci_computer	os_service_pack	oslevel
Short description	cmdb_ci_aix_server	short_description	uname
Name	cmdb_ci_aix_server	name	DNS, NBT
Hostname	cmdb_ci_aix_server	host_name	DNS, NBT
DNS domain	cmdb_ci_aix_server	dns_domain	DNS
Start date	cmdb_ci_aix_server	start_date	uptime
CPU type	cmdb_ci_computer	cpu_type	lsdev, lsattr

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Table Name	Field Name	Source
CPU speed (MHz)	cmdb_ci_computer	cpu_speed	lsdev, lsattr
CPU count	cmdb_ci_computer	cpu_count	lsdev, lsattr
Manufacturer	cmdb_ci_aix_server	manufacturer	lsattr
Model ID	cmdb_ci_aix_server	model_id	lsattr
RAM (MB)	cmdb_ci_computer	ram	lsdev, lsattr
Name	cmdb_ci_file_system	name	df
Capacity (MB)	cmdb_ci_file_system	capacity	df
Available Space (MB)	cmdb_ci_file_system	available_space	df
Mount point	cmdb_ci_file_system	mount_point	df
Name	cmdb_ci_patches	name	instfix
Name	cmdb_running_process	name	ps
Command	cmdb_running_process	command	ps
Connects to	cmdb_running_process	connects_to	lsof
Listening on	cmdb_running_process	listening_on	lsof
Type	cmdb_running_process	type	ps
PID	cmdb_running_process	pid	ps
Parameters	cmdb_running_process	parameters	ps
Name	cmdb_ci_network_adapter	name	ifconfig, netstat
IP address	cmdb_ci_network_adapter	ip_address	ifconfig, netstat
MAC address	cmdb_ci_network_adapter	mac_address	ifconfig, netstat
Netmask	cmdb_ci_network_adapter	netmask	ifconfig, netstat

Apache HBase instance discovery

Discovery creates or updates a CMDB record when it detects a running instance of HBase on a UNIX server.

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

In the following table, the `HBase Instance@`*`hostname`* source may contain the information on whether the instance is a primary or a secondary instance. The name contains *`Hmaster`* if the instance is a primary instance. The name contains *`HRegionServer`* if the name is a secondary instance.

Data collected

Label	Table Name	Field Name	Source
Name	cmdb_ci_db_hbase_instance	name	HBase Instance@*`hostname`*
Root Directory	cmdb_ci_db_hbase_instance	root_dir	hbase-site.xml
TCP port	cmdb_ci_db_hbase_instance	tcp_port	running process

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected (continued)

Label	Table Name	Field Name	Source
Site XML	cmdb_ci_db_hbase_instance	site_xml	hbase-site.xml
Version	cmdb_ci_db_hbase_instance	version	HBase shell
HBase Home	cmdb_ci_db_hbase_instance	hbase_home	running process
ZooKeeper Quorum	cmdb_ci_db_hbase_instance	zookeeper	hbase-site.xml

HP-UX discovery

Discovery identifies and classifies information about HP-UX computers.

Credentials

Configure SSH credentials.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Classifier, probes, and pattern

Classifier	Probes	Pattern
HP-UX (Unix classifier)	• Horizontal discovery probe: launches patterns • HP-UX - Installed Software^ • HP-UX - ADM^ • HP-UX - OS Patches* • HP-UX - CPU Speed* • HP-UX - Total Memory* • HP-UX - Hardware Model* • HP-UX - CPU Info* • HP-UX - Identity* • UNIX - Find FQDN* • UNIX - OS Filesystems* • UNIX - OS Uptime*	HP-UX Server

*These probes aren't active on the classifier, as Discovery uses patterns by default for these discoveries.

^These probes remain active by default, even when Discovery uses pattern discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Label	Table Name	Field Name	Source
Operating System	cmdb_ci_hpux_server	os	uname
Short description	cmdb_ci_hpux_server	short_description	uname
Name	cmdb_ci_hpux_server	name	DNS, NBT
Hostname	cmdb_ci_hpux_server	host_name	DNS, NBT
DNS domain	cmdb_ci_hpux_server	dns_domain	DNS
Start date	cmdb_ci_hpux_server	start_date	uptime
Manufacturer	cmdb_ci_computer	manufacturer	dmidecode
Serial number	cmdb_ci_hpux_server	serial_number	uname
CPU type	cmdb_ci_hpux_server	cpu_type	cpuinfo
CPU speed (MHz)	cmdb_ci_hpux_server	cpu_speed	adb
CPU count	cmdb_ci_hpux_server	cpu_count	cpuinfo
Model ID	cmdb_ci_hpux_server	model_id	model
RAM (MB)	cmdb_ci_hpux_server	ram	adb
Name	cmdb_ci_file_system	name	df
Capacity (MB)	cmdb_ci_file_system	capacity	df
Available Space (MB)	cmdb_ci_file_system	available_space	df
Mount point	cmdb_ci_file_system	mount_point	df
Name	cmdb_ci_patches	name	swlist
Name	cmdb_running_process	name	ps
Command	cmdb_running_process	command	ps
Type	cmdb_running_process	type	ps
PID	cmdb_running_process	pid	ps
Parameters	cmdb_running_process	parameters	ps
Name	cmdb_ci_network_adapter	name	hifconfig
IP address	cmdb_ci_network_adapter	ip_address	hifconfig
MAC address	cmdb_ci_network_adapter	mac_address	hifconfig
Netmask	cmdb_ci_network_adapter	netmask	hifconfig

Hyper-V discovery

The ServiceNow Discovery application finds Microsoft Hyper-V hypervisors in your environment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hyper-V is a hypervisor developed by Microsoft that enables you to create and manage virtual machines (VMs) on a physical server or computer. It enables hardware virtualization, enabling multiple virtual machines to run on a single physical machine while sharing the underlying hardware resources, such as CPU, memory, storage, and networking.

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Hyper-V supports the following functionality:

- Failover clustering: Failover is managed with Failover Cluster Manager.
- Live migration: Virtual machines can be moved between failover cluster nodes without bringing down the virtual machine.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Supported versions

Discovery is supported for Hyper-V running on the following Windows Server versions:

- 2008
- 2012
- 2012 R2
- 2016
- 2019
- 2022
- 2025

Important: Discovery of Hyper-V running on Windows starting from version 2016 supports pattern discovery only. If you upgraded and are still using probes and sensors for discovery, see KB0694773 for instructions on configuring Hyper-V probes.

Hyper-V data models

There are 2 types of Hyper-V:

- Hyper-V on Windows Server: This model, installed as a server role on Windows Servers, is designed for enterprises needing both virtualization and other server roles (for example: file servers, Active Directory, web servers) on the same machine.
- Hyper-V Server (Standalone): A free, lightweight version of Windows Server specifically designed for virtualization. It includes only the Hyper-V hypervisor and no graphical user interface (GUI) or additional server roles.

Discovery supports identifying and mapping both types of Hyper-V environments. However, the CMDB data models differ between the two deployment types.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example of Windows Server with Hyper-V role vs. Hyper-V Server (VM shown with Windows OS)

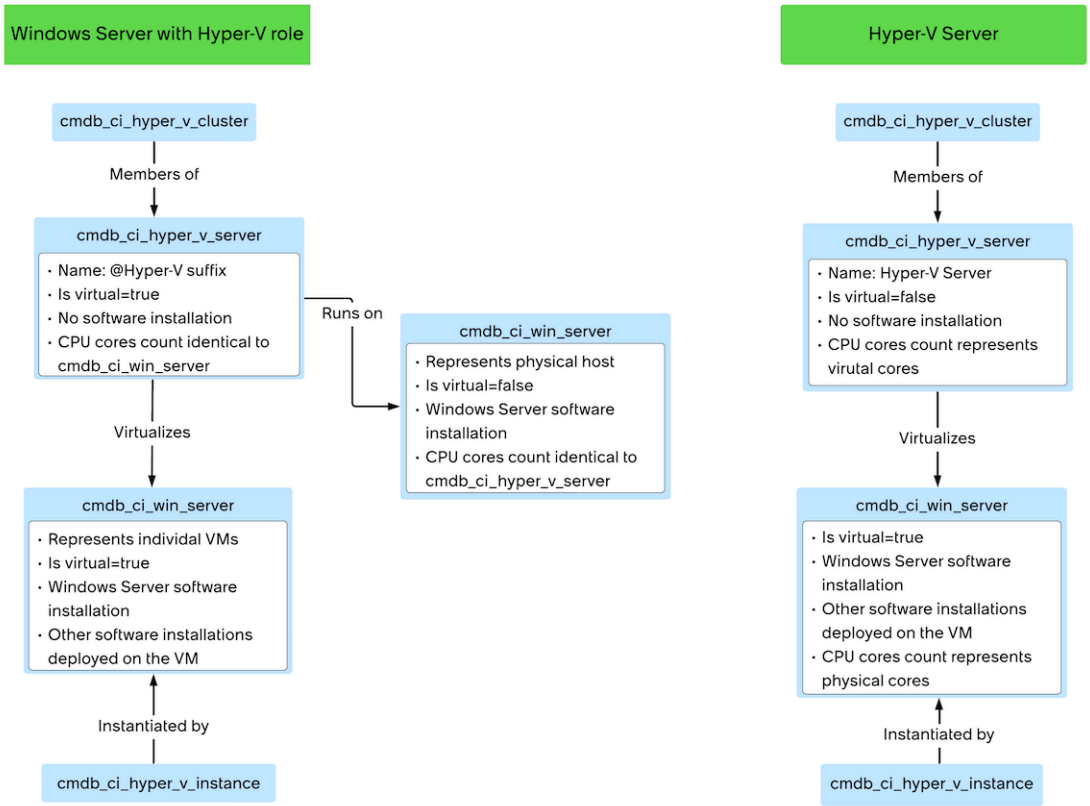

Comparison between Windows Server with Hyper-V role and Hyper-V Server

Key aspects	Windows Server with Hyper-V role	Hyper-V Server
Hyper-V Server CI	• Name of the CI includes the **@Hyper-V Server** suffix. For example: hyperv1@Hyper-V Server. • The **Windows Host** field references the parent Windows Servers table, which represents the physical host of the Hyper-V Server.	Name of the CI is the Hyper-V Server.
Hyper-V Cluster has **Members::Member of** relationship with Hyper-V Server	Yes The Hyper-V Server table: • Has **Is virtual** field set to true • Doesn't have software installation	Yes The Hyper-V Server table: • Has **Is virtual** field set to false • Doesn't have software installation
Hyper-V Server has **Runs on::Runs** relationship with Windows Server	Yes The Windows Server table:	No

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Comparison between Windows Server with Hyper-V role and Hyper-V Server (continued)

Key aspects	Windows Server with Hyper-V role	Hyper-V Server
	• Represents the physical host • Has **Is virtual** field set to false • Has Windows Server software installation	
Operating system (OS) installed on the VM has **Virtualized by::Virtualizes** relationship with Hyper-V Server	Yes The OS table (for example: Windows Server/Linux Server): • Represents the individual virtual machines (VMs) • Has **Is virtual** field set to true • Has OS software installation, such as Windows or Linux Server • Has other software installations deployed on the VM	Yes The OS table (for example: Windows Server/Linux Server): • Has **Is virtual** field set to true • Has OS software installation, such as Windows or Linux Server • Has other software installations deployed on the VM
OS installed on the VM has **Instantiates::Instantiated by** relationship with Hyper-V Virtual Machine Instance	Yes	Yes
CPU cores count	Is identical between the physical Windows Server host and the Hyper-V Server	The parent Windows Servers table reflects the physical cores, while the Hyper-V Server table displays the virtual cores

Credentials

Configure Windows credentials with Domain administrator rights. You should also Enable PowerShell for the MID Server used to discover Hyper-V servers and instances.

Classifier, probes, and pattern

Classifier	Probes	Pattern
Hyper-V Server The classification applies to standalone Hyper-V servers only (servers NOT running on Windows 2008).	• Horizontal discovery probe: Launches patterns • Windows - ADM^ • Windows - Installed Software^ • Hyper-V - Cluster* • Hyper-V - Resource Pools* • Hyper-V - Virtual Machines*	Hyper-V Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier	Probes	Pattern
	• Hyper-V - Virtual Networks* • Several other Windows probes also run, such as Windows - Identity, Windows - software, and so on.*	

*These probes aren't active on the classifier, as Discovery uses patterns by default for these discoveries.

^These probes remain active by default, even when Discovery uses pattern discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected by Discovery during horizontal discovery

Virtual instances

Label	Field Name	Data Description	Stored in Table
Object ID ^	object_id	ID	cmdb_ci_hyper_v_instance
Name^	name	name	cmdb_ci_hyper_v_instance
State	state	• On • Off	cmdb_ci_hyper_v_instance*
CPUs	cpus	Count	cmdb_ci_hyper_v_instance*
Memory	memory	Quantity in MB	cmdb_ci_hyper_v_instance*
Network adapters	nics	Count	cmdb_ci_hyper_v_instance*
Disks	disks	Count+	cmdb_ci_hyper_v_instance*
Disks size	disks_size	Capacity in GB+	cmdb_ci_hyper_v_instance*
Virtual Base Board Serial Number	baseboard_serial	Virtual serial number	cmdb_ci_hyper_v_instance
BIOS Serial Number	bios_serial	Virtual serial number	cmdb_ci_hyper_v_instance
Chassis Serial Number	chassis_serial	Virtual serial number	cmdb_ci_hyper_v_instance
BIOS GUID	bios_guid	Globally unique identifier (GUID)	cmdb_ci_hyper_v_instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

*Attributes that are put in the `cmdb_ci_hyper_v_instance` table when using the Hyper-V Server pattern. For upgrades, probes and sensors continue to put these attributes in the `cmdb_ci_vm_instance` table.

^Newly discovered attribute when using the Hyper-V Server pattern.

+Discovery can only return this information if the virtual machine is running.

Virtual servers

Label	Field	Store in table
Name	name	`cmdb_ci_hyper_v_server`
Chassis type	chassis_type	`cmdb_ci_hyper_v_server`
Description	short_description	`cmdb_ci_hyper_v_server`
OS Address Width (bits)	os_address_width	`cmdb_ci_hyper_v_server`
Operating System	OS	`cmdb_ci_hyper_v_server`
Host name	hostname	`cmdb_ci_hyper_v_server`
IP Address	ip_address	`cmdb_ci_hyper_v_server`
OS Version	os_version	`cmdb_ci_hyper_v_server`
OS Service Pack	os_service_pack	`cmdb_ci_hyper_v_server`
Is virtual	virtual	`cmdb_ci_hyper_v_server`

Virtual networks

Label	Field Name	Data Description	Stored in Table
Name	name	Name of the virtual network	`cmdb_ci_hyper_v_network` *
ID	object_id	Globally unique identifier (GUID)	`cmdb_ci_hyper_v_network` *

*Attributes that are put in the `cmdb_ci_hyper_v_network` table when using the Hyper-V Server pattern. For upgrades, probes and sensors continue to put these attributes in the `cmdb_ci` table (name) and the `cmdb_ci_vm_object` table (object_id).

Resource pools

Label	Field Name	Data Description	Stored in Table
Name^	name	Name of the resource pool.	`cmdb_ci_hyper_v_resource_pool`
ID	object_id	Globally unique identifier (GUID)	`cmdb_ci_hyper_v_rpool_comp`
Capacity	capacity	Maximum amount of the appropriate allocation unit	`cmdb_ci_hyper_v_rpool_comp`
Allocation units	allocation_units	Units of measurement used (for example, MB or GB)	`cmdb_ci_hyper_v_rpool_comp`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource pools (continued)

Label	Field Name	Data Description	Stored in Table
Resource type	resource_type	Type of resource discovered (for example, memory, or hard drive space)	cmdb_ci_hyper_v_rpool_comp

Clusters

Label	Field Name	Data Description	Stored in Table	Reference Table
Windows Cluster	name	Reference field displaying the Windows Cluster on which the Hyper-V cluster resides.	cmdb_ci_hyper_v_cluster	cmdb_ci_win_cluster

Relationships

Discovery creates records in the CMDB for the host machine, the parent partition running Windows Server 2008, and all the virtual machines running on the Hyper-V server. Discovery finds the stored templates and maps all the relationships between components.

The instance modifies these tables for use with multiple virtualization products:

Table	Purpose
Virtualization Server [cmdb_ci_virtualization_server]	Contains data on all discovered virtualization servers.
Virtual Machine Instance [cmdb_ci_vm_instance]	Contains data on all discovered virtual machine instances.
Virtual Machine Object [cmdb_ci_vm_object]	Contains data about various objects associated with a Hyper-V server, such as partitions, networks, resource pools, and clusters.

Parent table	Relationship	Child table
Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance]	Registered on::Has registered	Hyper-V Server [cmdb_ci_hyper_v_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parent table	Relationship	Child table
Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance]	Connected by ::Connects	Hyper-V Network [cmdb_ci_hyper_v_network]
Hyper-V Network [cmdb_ci_hyper_v_network]	Provided::Provides	Hyper-V Server [cmdb_ci_hyper_v_server]
Hyper-V Resource Pool [cmdb_ci_hyper_v_resource_pool]	Defines resources for:Gets resources from	Hyper-V Server [cmdb_ci_hyper_v_server]
Hyper-V Resource Pool Components [cmdb_ci_hyper_v_rpool_comp]	Contained by::Contains	Hyper-V Resource Pool [cmdb_ci_hyper_v_resource_pool]
Hyper-V Cluster [cmdb_ci_hyper_v_cluster]	Members::Member of	Hyper-V Server [cmdb_ci_hyper_v_server]
Hyper-V Server [cmdb_ci_hyper_v_server]	Runs on::Runs	Windows Servers [cmdb_ci_win_server]
Windows Servers [cmdb_ci_win_server]	Instantiates::Instantiated by	Hyper-V Virtual Machine Instance [cmdb_ci_hyper_v_instance]
Windows Servers [cmdb_ci_win_server]	Virtualized by::Virtualizes	Hyper-V Server [cmdb_ci_hyper_v_server]

Clone Hyper-V virtual machines

When importing (cloning) Hyper-V virtual machines, make sure each virtual machine has a unique ID. Discovery identifies virtual machines with duplicate IDs as the same machine.

1. Export the virtual machine from the Hyper-V server.
2. Import the virtual machine into a different server, making one of the these choices:
   - Move or restore the virtual machine: This selection clones the image, using the same ID, and moves it to another server. The original image remains on the Hyper-V server.
   - Copy the virtual machine: This selection copies the virtual machine, using a different ID, and moves it from the Hyper-V server. This is the best selection and does not require any further action before you run Discovery.
3. If you import the Hyper-V clone using the **Move or restore** selection, be sure to delete the original image from the Hyper-V server.

   When Discovery encounters two virtual machines with the same equivalent serial numbers, it creates only one configuration item (CI).

Red Hat JBoss server discovery

Discovery can detect JBoss application servers running on Linux and Windows systems.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Requirements

For JBoss application servers running on Linux systems.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Set probe permissions to use these Bourne shell commands: *find*, *cat*, and *dirname*.
- Enable SSH on the JBoss application server. The SSH credential must also have read permissions on the `web.xml` and `jboss-service.xml` files.

For JBoss application servers running on Windows systems, enable PowerShell on the MID Server.

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Credentials

- Applicative credentials
- Windows credentials
- SSH credentials

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Classifiers, patterns, and probes

Classifier	Trigger probes	Patterns
JBoss Server	• Horizontal Pattern: launches patterns • Windows - JBoss Find web.xml List* • Windows - JBoss Get jboss-service.xml* • JBoss - Find web.xml List* • JBoss - Get jboss-service.xml*	JBoss pattern

*For new instances, these probes are inactive on the classifier. Discovery uses patterns for discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

Discovery populates the following fields in the JBoss [cmdb_ci_app_server_jboss] table:

Label	Field name
Name	name
Category	category
Version	version

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field name
TCP port(s)	tcp_port

Discovery populates the following fields on the Web Service [cmdb_ci_web_service] and Web Application [cmdb_ci_web_application] servers tables:

Label	Table name	Field name	Source
Name	cmdb_ci_web_service	name	jboss-service.xml
App server	cmdb_ci_web_service	app_server	Internal reference
Description	cmdb_ci_web_application	short_description	web.xml
Name	cmdb_ci_web_application	name	web.xml
Document base	cmdb_ci_web_application	document_base	web.xml
App server	cmdb_ci_web_application	app_server	web.xml
Servlet Name*	cmdb_ci_web_application	servlet_name	web.xml
Servlet Class*	cmdb_ci_web_application	servlet_class	web.xml

*Windows servers only

Relationships

Relationships created

Parent class	Relationship	Child class
cmdb_ci_web_service	Runs on::Runs	cmdb_ci_windows_server
cmdb_ci_app_server_jboss	Contains::Contained by	cmdb_ci_web_application
cmdb_ci_app_server_jboss	Runs:Runs on	cmdb_ci_web_application
cmdb_ci_app_server_jboss	Runs:Runs on	cmdb_ci_web_service

Legacy: JBoss probes and sensors

Discovery identifies a Linux JBoss application server using probes and sensors.

Linux

This process describes the use of probes, which are active on your instance if you upgrade. For new instances, Discovery uses the JBoss pattern.

1. The UNIX - Active Processes probe detects a running process that matches an `org.jboss.main` entry point parameter.
2. If there's a match, a record is created in the JBoss Application Server [cmdb_ci_app_server_jboss] table. The following probes are also triggered:
   - JBoss - Find web.xml list: The sensor of this probe populates information in the Web Application [cmdb_ci_web_application] table if applicable.
   - JBoss - Get jboss-service.xml: The sensor of this probe populates information in the Web Service [cmdb_ci_web_service] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. The JBoss - Find web.xml list probe searches for the `web.xml` files of JBoss application server. The probe uses the *`classpath`* parameter in the running process, and then searches in the related `server\default\deploy` directory for the JBoss installation.
4. If associated web applications reside in the `server\default\deploy` directory, the JBoss - Get web.xml probe triggers for each application. This probe reads the `web.xml` file for each web application and the sensor populates additional information to the Web Service [cmdb_ci_web_service] table.
5. The Boss - Get jboss-service.xml probe uses the *`classpath`* parameter in the running process to search for the `jboss-service.xml` file in the related `server\default\conf\` directory for the JBoss installation.
6. If the probe successfully finds the `jboss-service.xml` file in the `server\default\conf\` directory, the sensor reads the contents of the XML file. It then creates additional records in the Web Service [cmdb_ci_web_service] table as necessary.

Windows

This process describes the use of probes, which are active on your instance if you upgrade. For new instances, Discovery uses the JBoss pattern.

1. The Windows - Active Processes probe detects a running process that matches an *`org.jboss.main`* entry point parameter.
2. If there's a match, a record is created in the JBoss Application Server [cmdb_ci_app_server_jboss] table. The following probes are also triggered:
   - JBoss - Find web.xml list: The sensor of this probe populates information in the Web Application [cmdb_ci_web_application] table if applicable.
   - JBoss - Get jboss-service.xml: The sensor of this probe populates information in the Web Service [cmdb_ci_web_service] table.
3. The JBoss - Find web.xml list probe searches for the `web.xml` files of JBoss application server. The probe uses the *`classpath`* parameter in the running process, and then searches in the related `server\default\deploy` directory for the JBoss installation.
4. If associated web applications reside in the `server\default\deploy` directory, the JBoss - Get web.xml probe triggers for each application. This probe reads the `web.xml` file for each web application and the sensor populates additional information to the Web Service [cmdb_ci_web_service] table.
5. The Boss - Get jboss-service.xml probe uses the *`classpath`* parameter in the running process to search for the `jboss-service.xml` file in the related `server\default\conf\` directory for the JBoss installation.
6. If the probe successfully finds the `jboss-service.xml` file in the `server\default\conf\` directory, the sensor reads the contents of the XML file. It then creates additional records in the Web Service [cmdb_ci_web_service] table as necessary.

Linux discovery

Discovery and Service Mapping applications use probes and patterns to discover and map information about Linux computers and servers. The information is populated in the CMDB. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Supportability details

The Linux server pattern can discover only English-based Linux.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Supported Linux operating systems

For IPv4 discovery:

- Red Hat
- Oracle
- Fedora
- Debian
- SUSE
- CentOS
- Ubuntu
- Rocky Linux
- AlmaLinux OS
- Debian GNU/Linux
- VMware Photon OS
- Amazon Linux
- Alpine Linux (starting from Visibility Content version 6.30.0)*

* Operating System and OS Version only

For IPv6 discovery:

- CentOS 7
- Ubuntu 20

The following operating systems or devices haven't been verified for IPv6 discovery:

- Red Hat cluster
- Oracle Clusterware
- Linux Pacemaker cluster
- Network storage

Classifier, probes, and pattern

Classifier	Probes	Pattern
Linux (Unix classification)	• Horizontal discovery probe: Launches patterns • Linux - Installed Software^ • Unix - ADM^ • Linux - Identity* • Linux - CPU* • Linux - Distribution* • Linux - Find FQDN* • Linux - Memory* • Linux - Memory Modules*	Linux Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier	Probes	Pattern
	• Linux - Network ARP Tables* • Linux - Storage* • Linux - Amazon EC2* • Unix - OS Uptime* • UNIX - OS Filesystems* • UNIX - Find FQDN* • Unix - ADM Enhanced*	

*These probes aren't active on the classifier, as Discovery uses patterns by default for these discoveries.

^These probes remain active by default, even when Discovery uses pattern discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Requirements for Linux discovery

Before running a discovery, you must verify that all the required configurations are complete.

Verify the plugins installation and activation

- Plugins or applications installed with ITOM Visibility
- Install Discovery and Service Mapping Patterns
- Install Visibility Content

Verify the required credentials configuration

Configure the credentials by the required SSH permissions. Discovering sensitive Linux data may require a user with appropriate sudo privileges to run the sudo commands. To configure the required sudo privileges, modify the `/etc/sudoers` file to include the commands that Discovery needs. For `/etc/sudoers` line examples for each command, see Privileged SSH commands for probe-based discovery.

For more information about the commands that require a user with elevated rights, refer to the following documents.

- Commands that require root privileges for Discovery, Orchestration, and Integration Hub.
- For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- MID Server privileged commands
- View Patterns commands through the Discovery Patterns module
- Validate commands used in pattern-based discovery

SSH private key credentials are preferable over SSH password credentials for security reasons.

Configure PermitTTY for SSH sessions

Set `PermitTTY yes` in `/etc/ssh/sshd_config` on the target Linux server. This setting is required for privileged commands such as `dmidecode` to run successfully.

Verify the MID Server configuration

For detailed information, see:

- Test MID Server connectivity
- MID Server properties

Verify the configuration of the discovery schedule

For information on creating a discovery schedule, see Schedule a horizontal discovery.

Define IP ranges so that only the specified network adapters and their primary IP addresses are used during Linux discovery. You can configure **Quick Ranges** by entering comma-separated IPv4 address ranges or single IPv6 address.

Alternatively, you can use the *`network_adapter_exclusion_list`* property to limit Linux discovery to specific network adapters and their primary IP addresses. For more information, see Omit network adapter secondary IP addresses in Linux discovery.

(Optional) Populate Virtual Machine Object field in Hardware [cmdb_ci_hardware] table

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating the Virtual Machine Object field in the Hardware [cmdb_ci_hardware] table. For more information, see Improved query performance with direct field population in CI tables.

Set the preferred IP address version for network adapter discovery

Starting with Visibility Content version 6.32.0, if your network adapters support both IPv4 and IPv6, the IPv4 address is populated by default in the IP address [ip_address] field on the Network Adapter [cmdb_ci_network_adapter] table. To control which IP version is populated, see Set the preferred IP version for network adapter discovery.

Configure the serial number type

Starting with Visibility Content version 6.35.0, you can populate a specific serial number type instead of the first available one. For more information, see Configure the serial number type for Linux Server discovery.

Data collected

The Linux classifier triggers probes that perform the discovery. Several probes are launched during the discovery. See the classifier for a list of the trigger probes.

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Table Name	Field Name	Source
Operating System	cmdb_ci_linux_server	os	uname -a
OS Version	cmdb_ci_computer	os_version	uname -a or cat /etc/*release
Short description	cmdb_ci_linux_server	short_description	uname -a
Name	cmdb_ci_linux_server	name	DNS, NBT
Hostname	cmdb_ci_linux_server	host_name	DNS, NBT
DNS domain	cmdb_ci_linux_server	dns_domain	DNS
Start date	cmdb_ci_linux_server	start_date	uptime
Manufacturer	cmdb_ci_computer	manufacturer	dmidecode
Serial number	cmdb_ci_computer	serial_number	dmidecode
CPU type	cmdb_ci_linux_server	cpu_type	/proc/cpuinfo
CPU speed (MHz)	cmdb_ci_linux_server	cpu_speed	/proc/cpuinfo
CPU count	cmdb_ci_linux_server	cpu_count	/proc/cpuinfo
CPU core count	cmdb_ci_computer	cpu_core_count	/proc/cpuinfo
CPU core thread	cmdb_ci_computer	cpu_core_thread	/proc/cpuinfo
CPU manufacturer	cmdb_ci_linux_server	cpu_manufacturer	/proc/cpuinfo
Model number	cmdb_ci_computer	model_number	dmidecode
Model ID	cmdb_ci_computer	model_id	dmidecode
RAM (MB)	cmdb_ci_linux_server	ram	meminfo
Disk space (GB)*	cmdb_ci_linux_server	disk_space	/proc/ide, /proc/scsi, /var/log/dmesg
Type	cmdb_ci_disk	type	/proc/ide, /proc/scsi, /var/log/dmesg
Model ID	cmdb_ci_disk	model_id	/proc/ide, /proc/scsi, /var/log/dmesg
Disk space (GB)	cmdb_ci_disk	disk_space	/proc/ide, /proc/scsi, /var/log/dmesg
Name	cmdb_ci_disk	name	/proc/ide, /proc/scsi, /var/log/dmesg
Name	cmdb_ci_file_system	name	df
Capacity (MB)	cmdb_ci_file_system	capacity	df
Free Space Bytes (MB)	cmdb_ci_file_system	free_space_bytes	df
Mount point	cmdb_ci_file_system	mount_point	df
Name	cmdb_running_process	name	ps

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Table Name	Field Name	Source
Command	cmdb_running_process	command	ps
Type	cmdb_running_process	type	ps
PID	cmdb_running_process	pid	ps
Parameters	cmdb_running_process	parameters	ps
Name	cmdb_ci_network_adapter	name	ifconfig or ip address show
IP address**	cmdb_ci_network_adapter	ip_address	ifconfig or ip address show
MAC address	cmdb_ci_network_adapter	mac_address	ifconfig or ip address show
Netmask	cmdb_ci_network_adapter	netmask	ifconfig or ip address show
Default gateway	cmdb_ci_hardware	default_gateway	route

* The value in the **disk_space** field is an aggregation of the total capacity (to include used space) for all non-removable disks, including both directly attached and SAN storage.

** Starting with Visibility Content version 6.32.0, for network adapters that support both IPv4 and IPv6, the IPv4 address is populated by default. Before this release, the populated value was selected randomly. To control which IP version is populated, see Set the preferred IP version for network adapter discovery.

Note: Make sure to define unique serial numbers for the RAMs on the Linux server.

Discovery also identifies and classifies information about Linux KVM. Discovery identifies Linux kernel-based virtual machines (KVM) when the process classifier detects `libvirtd` running on a Linux server. The classification triggers the creation of a cmdb_ci_kvm record and launches the SSH Command probes to explore the Linux server with `virsh`, `lbvert` utility, and virtual machine configuration data.

Discovery creates a [cmdb_ci_kvm_instance] record for each virtual machine on the server. It then matches the [cmdb_ci_kvm_instance] record to a corresponding [cmdb_ci_computer] record using the MAC addresses of installed network adapters.

Tables used by Discovery on Linux KVM

Table name	Extends	Description	Source
cmdb_ci_kvm	cmdb_ci_vm	A hypervisor that manages kernel-based virtual machines (KVMs)	Process classifier detects libvirtd running on Linux servers
cmdb_ci_kvm_vm_instance	cmdb_ci_vm_instance	A virtual machine instance on this hypervisor	virsh list --all and dumpxml command
cmdb_ci_kvm_object	cmdb_ci_vm_object	An object connected to a virtual machine instance	<network>, <storage pool>, and <storage volume> elements from the dumpxml command

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables used by Discovery on Linux KVM (continued)

Table name	Extends	Description	Source
cmdb_kvm_device	Not applicable	A device connected to a virtual machine instance	<devices> element from the dumpxml command

Data collected by Discovery on Linux KVM

Label	Table name	Field name	Source
Linux Host	cmdb_ci_kvm	linux_host	Reference to the cmdb_ci_linux_server that is running this virtual machine
Details	cmdb_ci_kvm	details_xml	dumpxml
Object ID	cmdb_ci_kvm_vm_instance	object_id	virsh dumpxml
State	cmdb_ci_kvm_vm_instance	state	virsh list --all
CPUs	cmdb_ci_kvm_vm_instance	cpus	virsh dumpxml
Memory	cmdb_ci_kvm_vm_instance	memory	virsh dumpxml
Disks	cmdb_ci_kvm_vm_instance	disks	virsh dumpxml
Disks size	cmdb_ci_kvm_vm_instance	disks_size	virsh domblkinfo
Network adapters	cmdb_ci_kvm_vm_instance	nics	virsh dumpxml
Name	cmdb_ci_kvm_vm_instance	name	virsh dumpxml
Short description	cmdb_ci_kvm_vm_instance	short_description	virsh desc
Details	cmdb_ci_kvm_object	details_xml	XML element from dumpxml
KVM instance	cmdb_kvm_device	kvm_instance	Reference to cmdb_ci_kvm_instance
Device	cmdb_kvm_device	device	disk, controller, interface, and so on
Type	cmdb_kvm_device	type	depends on the device
Details	cmdb_kvm_device	details_xml	XML element from dumpxml

Data collected by Discovery on KVM Relationship

Relationship	Parent table	Child table
Registered On::Has Registered	KVM [cmdb_ci_kvm]	KVM Virtual Machine instance [cmdb_ci_kvm_vm_instance]
Provided By::Provides	KVM [cmdb_ci_kvm]	Network [cmdb_ci_kvm_network]
Defines resource for::Gets resources from	KVM [cmdb_ci_kvm]	Storage Pool [cmdb_ci_kvm_storage_pool]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery on KVM Relationship (continued)

Relationship	Parent table	Child table
Connected By::Connects	KVM Virtual Machine instance [cmdb_ci_kvm_vm_instance]	Network [cmdb_ci_kvm_network]
Instantiated By::Instantiates	KVM Virtual Machine instance [cmdb_ci_kvm_vm_instance]	Computer [cmdb_ci_computer]
Virtualized By::Virtualizes	Computer [cmdb_ci_computer]	KVM [cmdb_ci_kvm]
Provides storage for::Stored on	Storage Pool [cmdb_ci_kvm_storage_pool]	KVM Virtual Machine instance [cmdb_ci_kvm_vm_instance]

Related topics

- Linux discovery fails with "Cannot connect, status is SSH_CONNECTION_FAILURE. Could not agree on signature algorithm Client (KB1425502)
- Linux discovery fails in Unix Classification with the error - Name of unclassified CI is 'servername' (KB0860486)
- Linux discovery error "rbash: PATH: readonly variable Exit status: 1" (KB0866860)

Linux Red Hat cluster discovery

Discovery can find Linux Red Hat clusters that offer high availability and load balancing.

Required credentials

Configure SSH credentials.

Additional requirements

You also need permissions to read the `/etc/hosts/` and `/etc/cluster/cluster.conf` file, which associates host names with IP addresses.

The following commands are run for Oracle clusterware:

Command
`ps -ef | grep corosync | grep -v grep`
`clustat -x`
`ifconfig | grep 'inet addr:' |awk '{print $1, $2 }'`
`hostname -s`
`hostname -s`

Sudo permission to run: `sudo /u01/app/12.1.0.2/grid/bin/ocrcheck | egrep -v 'error|return code'`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier, probes, and pattern

Classifier	Probes	Pattern
Linux (Unix classification)	Horizontal discovery probe: launches patterns	These patterns are an extension of the Linux Server pattern: • UNIX Cluster - Linux Red hat cluster • Linux - Red Hat Cluster The pattern uses the `clustat -x` command and looks at this cluster configuration file: `/etc/cluster/cluster.conf`

Data collected

Table	Labels and field names
Unix Cluster [cmdb_ci_unix_cluster]	Name [name] Cluster status [cluster_status] Cluster type [cluster_type] IP address [ip_address] Cluster ID [cluster_id] Cluster version [cluster_version]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Name [name] Node state [node_state] IP address [ip_address] Node status [node_status]
Unix Cluster Resource [cmdb_ci_unix_cluster_resource]	Name [name] Resource type [resource_type] Resource status [resource_status] Properties [properties]
Cluster Virtual IP [cmdb_ci_cluster_vip]	Name [name] IP address [ip_address]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Labels and field names
	Cluster ID [cluster_id]

Relationships

CI	Relationship	CI
Unix Cluster [cmdb_ci_unix_cluster]	Virtualizes::Virtualized by	Cluster Virtual IP [cmdb_ci_cluster_vip]
Unix Cluster Resource [cmdb_ci_unix_cluster_resource]	Defines resources for::Gets resources from	Unix Cluster Node [cmdb_ci_unix_cluster_node]
Cluster Virtual IP [cmdb_ci_cluster_vip]	Uses::Used by	Unix Cluster Node [cmdb_ci_unix_cluster_node]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Cluster of::Cluster	Unix Cluster [cmdb_ci_unix_cluster]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Hosted on:Hosts	Unix Cluster [cmdb_ci_unix_cluster]
Unix Cluster Resource [cmdb_ci_unix_cluster_resource]	Defines resources for::Gets resources from	Unix Cluster [cmdb_ci_unix_cluster]
Unix Cluster [cmdb_ci_unix_cluster]	Hosts::Hosted on	Linux Server [cmdb_ci_linux_server]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Hosted on::Hosts	Linux Server [cmdb_ci_linux_server]

Omit network adapter secondary IP addresses in Linux discovery

Limit Linux discovery to specific network adapters and their primary IP addresses to improve performance by ignoring secondary IP addresses.

About this task

Configuring the *network_adapter_exclusion_list* system property ensures that only specified network adapters and their primary IP addresses are used during Linux discovery. Secondary IP addresses that are not named are automatically omitted.

Before you begin

Role required: admin

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** field, search for `network_adapter_exclusion_list`.
3. Select the *network_adapter_exclusion_list* property.
4. In the **Value** field, enter the network adapter associated with the NIC IP address to exclude it from deletion.

   The network adapter must be entered in one of the formats listed in the following table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Format	Description
Single name	Single adapter name. For example: `bond0`.
Regular expression matching multiple adapter names	Regular expression that matches multiple adapter names ending with one or more digits. For example: `bond\d+`.
Multiple adapter names	Multiple adapter names, separated by commas, including regular expression. For example: `nic1,bond\d+`.

5. Select **Update**.

Result
When the predefined schedule triggers the Linux server discovery or a scheduled horizontal discovery runs, only the specific network adapters and their primary IP addresses are included in the discovery process.

For more information, see Schedule a horizontal discovery.

Related topics

Linux discovery

Configure the serial number type for Linux Server discovery

Control which serial number type populates the **Serial Number** field on discovered Linux server CIs by configuring the *sn_itom_pattern.linux_server_preferred_serial_number_type* system property.

Before you begin
Verify that you have at least version 6.35.0 of Visibility Content.

Role required: admin

About this task
By default, Discovery populates the **Serial Number** field on the Linux Server [cmdb_ci_linux_server] table with the first available serial number found on each discovery run. The *sn_itom_pattern.linux_server_preferred_serial_number_type* system property enables you to specify which serial number type to use. If the preferred type isn't found on a given host, Discovery falls back to the first available serial number. The value is empty by default.

Procedure
1. In the navigation filter, enter `sys_properties.list`.
2. In the **Name** column, search for `sn_itom_pattern.linux_server_preferred_serial_number_type`.
3. Select the *sn_itom_pattern.linux_server_preferred_serial_number_type* system property.
4. In the **Value** field, enter one of the following serial number types:
   - `system`
   - `chassis`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `board`
- `bios`

To revert to the default behavior and use the first available serial number, leave the value empty.

5. Select **Update**.

What to do next

Run Discovery again to apply the change.

Related topics

Linux discovery

Netware discovery

Discovery identifies and classifies information about Netware.

Data collected

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Label	Table name	Field name	Source
Name	cmdb_ci_netware_server	name	snmp
Serial number	cmdb_ci_netware_server	serial_number	snmp
OS Version	cmdb_ci_netware_server	os_version	snmp
RAM	cmdb_ci_netware_server	ram	snmp
CPU count	cmdb_ci_netware_server	cpu_count	snmp

Mac (OS/X) discovery

Discovery identifies and classifies information about Mac (OS/X) computers.

Credentials

Configure SSH credentials.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Label	Table Name	Field Name	Source
Operating System	cmdb_ci_computer	os	uname

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Table Name	Field Name	Source
OS Version	cmdb_ci_computer	os_version	system_profiler
OS Service pack	cmdb_ci_computer	os_service_pack	system_profiler
Short description	cmdb_ci_computer	short_description	uname
Name	cmdb_ci_computer	name	DNS, NBT
Hostname	cmdb_ci_computer	host_name	DNS, NBT
DNS domain	cmdb_ci_computer	dns_domain	DNS
Start date	cmdb_ci_computer	start_date	uptime
Manufacturer	cmdb_ci_computer	manufacturer	Assumed to be Apple
Serial number	cmdb_ci_computer	serial_number	system_profiler
CPU type	cmdb_ci_computer	cpu_type	system_profiler
CPU speed (MHz)	cmdb_ci_computer	cpu_speed	system_profiler
CPU count	cmdb_ci_computer	cpu_count	system_profiler
Model ID	cmdb_ci_computer	model_id	system_profiler
RAM (MB)	cmdb_ci_computer	ram	system_profiler
Disk space (GB)	cmdb_ci_computer	disk_space	system_profiler
Volume name	cmdb_ci_disk	volume_name	system_profiler
Volume serial number	cmdb_ci_disk	volume_serial_number	system_profiler
Disk space (GB)	cmdb_ci_disk	disk_space	system_profiler
Name	cmdb_ci_disk	name	system_profiler
Device ID	cmdb_ci_disk	device_id	system_profiler
Free space (GB)	cmdb_ci_disk	free_space	system_profiler
File system	cmdb_ci_disk	file_system	system_profiler
Name	cmdb_running_process	name	ps
Command	cmdb_running_process	command	ps
Connects to	cmdb_running_process	connects_to	lsof
Listening on	cmdb_running_process	listening_on	lsof
PID	cmdb_running_process	pid	ps
Parameters	cmdb_running_process	parameters	ps
Name	cmdb_ci_network_adapter	name	system_profiler
IP address	cmdb_ci_network_adapter	ip_address	system_profiler
MAC address	cmdb_ci_network_adapter	mac_address	system_profiler
Netmask	cmdb_ci_network_adapter	netmask	system_profiler
MAC manufacturer	cmdb_ci_network_adapter	mac_mfr	Assumed to be Apple

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Table Name	Field Name	Source
DHCP enabled	cmdb_ci_network_adapter	dhcp_enabled	system_profiler

Solaris discovery

Discovery identifies and classifies information about Solaris computers.

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Supported operating systems

- Oracle Solaris 10
- Oracle Solaris 11

Prerequisites

Handle truncated command line parameters

Discovery stores data in the [cmdb_running_process] table with command line parameters truncated to 80 characters. This can cause multiple applications to be merged into one CI. To get the full command line and prevent this issue, run `pargs -a` and parse the result.

Install SNEEP

To use patterns to find Solaris computers, install SNEEP. Without it, Discovery cannot find the serial number.

Configure unique serial numbers for Solaris zones

Configure each zone to return a unique serial number. If zones return the same number, CI identification fails.

Note: Solaris discovery only supports local zones installed on the same host as the global zone. If the local zone is installed on a different host from the global zone, the `zoneadm list -cip` command can't retrieve the UUID for the local zone.

Install the LSOF command

The Solaris - ADM probe requires the LSOF command to be installed.

Configure SSH credentials

Configure SSH credentials for the systems you are discovering.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Set the preferred IP address version for network adapter discovery

Starting with Visibility Content version 6.32.0, if your network adapters support both IPv4 and IPv6, the IPv4 address is populated by default in the IP address [ip_address] field on the Network Adapter [cmdb_ci_network_adapter] table. To control which IP version is populated, see Set the preferred IP version for network adapter discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier, probes, and pattern

Classifier, probes, and pattern

Classifier	Probes	Pattern
Solaris (UNIX Classification)	• Horizontal discovery probe: launches patterns • Solaris - Installed Software^ • Solaris - Identity* • Solaris - CPU* • Solaris - Hardware Model* • Solaris - Memory* • Solaris - Network ARP Tables* • Solaris - OS Patches* • Solaris - Storage* • Solaris - Zones* • Solaris - Disks* • DNS*	Solaris Server

*For new instances, these probes are inactive on the classifier. Discovery uses patterns for discovery.

^This probe remains active by default, even when Discovery uses pattern discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

Data collected on Solaris computers

Label	Table name	Field name	Source
Operating System	cmdb_ci_computer	os	uname
Short description	cmdb_ci_solaris_server	short_description	uname
Name	cmdb_ci_solaris_server	name	DNS, NBT
Hostname	cmdb_ci_solaris_server	host_name	DNS, NBT
DNS domain	cmdb_ci_solaris_server	dns_domain	DNS
Start date	cmdb_ci_solaris_server	start_date	uptime
CPU type	cmdb_ci_computer	cpu_type	kstat
CPU speed (MHz)	cmdb_ci_computer	cpu_speed	kstat
CPU count	cmdb_ci_computer	cpu_count	kstat
CPU core count	cmdb_ci_computer	cpu_core_count	kstat

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected on Solaris computers (continued)

Label	Table name	Field name	Source
CPU core thread	cmdb_ci_computer	cpu_core_thread	kstat
Model number	cmdb_ci_solaris_server	model_number	suntype
Model ID	cmdb_ci_solaris_server	model_id	suntype
RAM (MB)	cmdb_ci_computer	ram	prtconf
Disk space (GB)	cmdb_ci_solaris_server	disk_space	iostat
Serial Number	cmdb_ci_solaris_server	serial_number	sneep *
Manufacturer	cmdb_ci_disk	manufacturer	iostat
Model ID	cmdb_ci_disk	model_id	iostat
Volume serial number	cmdb_ci_disk	volume_serial_number	iostat
Disk space (GB)	cmdb_ci_disk	disk_space	iostat
Name	cmdb_ci_disk	name	iostat
Name	cmdb_ci_file_system	name	df
Capacity (MB)	cmdb_ci_file_system	capacity	df
Available Space (MB)	cmdb_ci_file_system	available_space	df
Mount point	cmdb_ci_file_system	mount_point	df
Name	cmdb_ci_patches	name	showrev
Name	cmdb_running_process	name	ps
Command	cmdb_running_process	command	ps
Connects to	cmdb_running_process	connects_to	lsof
Listening on	cmdb_running_process	listening_on	lsof
Type	cmdb_running_process	type	ps
PID	cmdb_running_process	pid	ps
Parameters	cmdb_running_process	parameters	ps
Name	cmdb_ci_network_adapter	name	ifconfig
IP address**	cmdb_ci_network_adapter	ip_address	ifconfig
MAC address	cmdb_ci_network_adapter	mac_address	ifconfig
Netmask	cmdb_ci_network_adapter	netmask	ifconfig
Default gateway	cmdb_ci_hardware	default_gateway	netstat

* To discover Fujitsu PRIMEPOWER devices, you must install Oracle SNEEP and run Solaris discovery with root credentials.

** Starting with Visibility Content version 6.32.0, for network adapters that support both IPv4 and IPv6, the IPv4 address is populated by default. Before this release, the populated value was selected randomly. To control which IP version is populated, see Set the preferred IP version for network adapter discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solaris zone discovery

Discovery maps the relationships between global and local Solaris zones upon detection.

In the following example, a Solaris global zone contains two local zones: zone01 and zone02. Each local zone is represented by a physical Solaris CI record and a Virtual Machine Instance record. Each of the local zones is tied to a Zone Server, demonstrating how virtualization relates to the global zone (mmp1).

Solaris zone relationship map

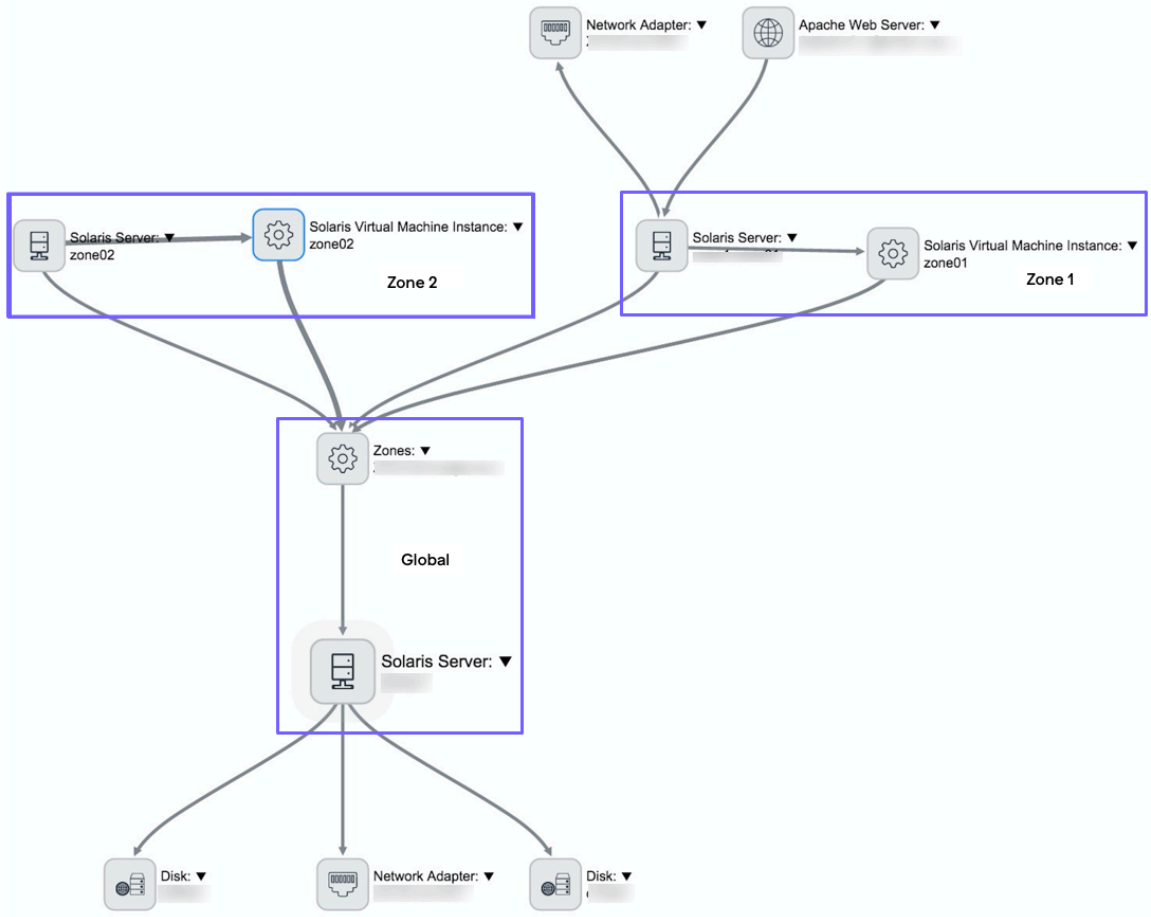

The TCP connection and process information for local zone servers must be collected by running commands on their parent global zone. The relationship path between the local and global zone physical machines must be established before TCP connection and process information for local zone servers can be collected.

Case 1: Global zone is discovered first.

- The system creates the Solaris server CI for the global zone.
- Discovery detects the local zones, creates a hypervisor zone server record, and creates a virtual machine instance record for each Solaris device in the local zone.
- Discovery creates the relationship between the hypervisor record and the VM instance record.

Case 2: Local zone is discovered first.

- The system creates the Solaris server CI for the local zone.
- Discovery sets the correlation ID, so that it can be reconciled during later global zone discoveries.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Case 3: Global zone is discovered after local zone Solaris server CIs are created.

- Global zone Discovery detects local zones.
- Discovery creates a hypervisor zone server record, and creates a virtual machine instance record for each Solaris device in the local zone.
- Discovery creates the relationship between the hypervisor record and the VM instance record. In addition, it creates the relationship between the physical local zone VM and its virtual machine instance record.
- The global zone runs the Solaris - ADM probe on itself, filtering by the local zone, and updates the physical local zone VMs with this data.

Case 4: The relationship path between physical local and global zone machines is established.

Subsequent discoveries of the global zone refresh the TCP connection and process information for the contained local zones.

When the system discovers a global zone, the Solaris - Zones & ADM Launcher probe triggers the Solaris - ADM probe to explore the global zone and each local zone found. Because the Solaris - ADM probe must run on the global zone to detect TCP connection and process information from its local zones, you might see multiple ECC queue records that appear identical.

ECC queue entries for a zone Discovery

2016-01-13 02:02:29	mid.server.demouk	MultiProbe	Solaris - Application Dependency Mapping	10.196.39.206	input	processed	2016-01-13 02:02:31
2016-01-13 02:02:29	mid.server.demouk	MultiProbe	Solaris - Application Dependency Mapping	10.196.39.206	input	processed	2016-01-13 02:02:31
2016-01-13 02:02:29	mid.server.demouk	MultiProbe	Solaris - Application Dependency Mapping	10.196.39.206	input	processed	2016-01-13 02:02:31

Examining the payload shows that each probe targets a different zone CI to filter on and update.

Local zone payload

```
<parameter name="sequence" value="1523a01c9470000001
<parameter name="from_host" value=""/>
<parameter name="zone" value="zone01"/>
<parameter name="sys_created_on" value="2016-01-13 (
<parameter name="used_by_discovery" value="true"/>
```

Data collected on Solaris zones

Label	Table name	Field name	Source
Version	cmdb_ci_vm_zones	version	zoneadm, zonename
Correlation ID	cmdb_ci_vm_zones	correlation_id	zoneadm, zonename
Name	cmdb_ci_solaris_instance	name	zoneadm, zonename
Parent	cmdb_ci_solaris_instance	parent	Internal
CMDB CI	cmdb_ci_solaris_instance	cmdb_ci	Internal
Correlation ID	cmdb_ci_solaris_instance	correlation_id	zoneadm, zonename

Note: cmdb_ci_solaris_instance extends cmdb_ci_vm_instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery for VMware vCenter

Discovery can explore the VMware vCenter process running on a Windows or Linux host. IPv6 is supported for discovery in VMware vCenter.

Note:

If a Windows or Linux server is running vCenter application, after creating the vCenter CI, a "Runs on" relationship to the Windows or Linux server host is created.

When a vCenter application is running on a dedicated vCenter appliance, a "Runs on" relationship to the vCenter CI isn't created.

Tested vCenter versions

The ServiceNow® platform supports vCenter API versions 4.0 and higher. The following versions were tested with the Brazil release:

- vCenter versions 9.0 and earlier
- vCenter appliance version 6.7 and earlier

For installations based on the vCenter appliance, a server CI is created for the vCenter appliance and the following fields are populated:

- NameIP
- AddressMAC
- AddressDiscovery
- Source

Note: If you're discovering SUSE Linux hosts for vCenter appliances, version 6.0 and earlier, observe these SSH restrictions.

See Data collected for VMware Cloud Discovery for a description of the VMware architecture and component relationships.

vCenter discovery process

After classifying vCenter, Discovery launches the VMware - vCenter Datacenters probe, which in turn launches specific probes that return information about ESX machines, virtual machines, and other vCenter objects. The vmapp port probe is also configured to launch the VMware - vCenter Datacenters probe.

Note: For accurate relationship mapping between server CIs and vCenter resources (such as ESX hosts and datastores), you should initiate discovery of the vCenter environment before discovering Server CIs. This approach helps Discovery associate servers with their corresponding vCenter-managed infrastructure more effectively.

VMWare credentials

To access vCenter with a domain account, specify the domain with the user name in the credential record using a supported format, such as **Domain\UserName** or **username@domain**. The VMware credentials must have a read-only role in vCenter. For Software Asset Management (SAM) tracking, the credentials also require the "Assign license" privilege in vCenter, also known as the License Admin privilege.

Important: Windows credentials aren't necessary for vCenter Discovery when valid VMware credentials are used.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Updating the CMDB with vCenter event collector

In addition to finding vCenter data through the standard discovery process, Discovery can also update the CMDB by detecting vCenter events through a MID Server extension called the vCenter event collector.

The event collector allows the CMDB to be updated with changes to virtual machines (VMs), in addition to the updates detected by Discovery. A change to a VM is sent as an event from the vCenter server to the vCenter event collector. When an event is received, the CMDB is updated accordingly. Full vCenter Discovery does not need to rerun. For some events, such as powered on and powered off events, Discovery does not need to run again at all. For most events, Discovery runs only on the necessary vCenter resource.

For instructions on configuring vCenter events, see Configure and run the vCenter event collector extension.

VM deleted from vCenter

If the VM is deleted from vCenter, the cmdb_ci_vm_instance state changes to terminated and the Status field changes to retired.

CIs removed from vCenter

When a vCenter CI, such as a virtual machine, is removed, the ServiceNow instance marks it as "stale" in the CMDB, using either of these procedures:

- When Discovery runs, it creates an audit record in the Components installed with CMDB Health table for the missing CI and marks the CI "stale".
- If the instance is configured to collect vCenter events, the system can also create a "stale" audit record for the CI in the CMDB Health Result [cmdb_health_result] table from the VmRemovedEvent event, without having to run Discovery.

Note: When the **Staleness** setting is configured, the dependency view (BSM map) grays out stale CIs in its relationship diagram to indicate that they were removed from vCenter.

To avoid stale CI health indicators from being generated during VMware discovery, set the system property *glide.cmdb.health.src.cmdb_health_audit_only* to **true**. This disables stale CI reporting from the VMware discovery source, enabling you to manage the CI life cycle through other means. For more information, see CMDB Health system properties.

You have the option of creating a CMDB remediation rule to automatically execute a remediation workflow that can, for example, delete stale CIs. For more information on stale CIs, see CMDB Health Metrics.

vCenter Discovery on Windows host

Windows credentials aren't necessary for vCenter Discovery, when valid VMware credentials are used.

Discovery for VMware virtual machines

Discovery gathers information about virtual machines managed by VMware vCenter.

The Discovery product can extract information from VMware configuration items (CI), including their relationships. Discovery is configured to work with the Orchestration VMware Support Plugin, which enables the automatic cloning of virtual machines by an ESX server managed by vCenter. For information about data collected by Discovery on vCenter, see Data collected for VMware Cloud Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows discovery

Discovery identifies and classifies information about Windows computers that use IPv4 addresses, IPv6 addresses, or both.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Note: admin$ share access is required for Windows discovery from Madrid P3 and later. A Remote Access role must be configured in the target Windows Server to enable connectivity with ServiceNow Discovery.

Supported Windows versions

For IPv4 discovery:

- Windows Workstation 7, 8, 10
- Windows Server versions:
  - 2008
  - 2008R2
  - 2012R2
  - 2016
  - 2019
  - 2022
  - 2025

For IPv6 discovery:

- Windows 2019 (10.0.17763)
- Windows Server 2016
- Windows Desktop 10
- Windows 2019 Datacenter

Note: For fiber channel discovery on a Windows 2008 host, the Microsoft Fibre Channel Information Tool (`fcinfo.exe`) must be installed on that machine. The `fcinfo` executable should be available on the environment path. The Microsoft Fibre Channel Information Tool tool is available for download at http://www.microsoft.com.

IPv6 supportability limitations

The following device types haven't been verified for IPv6 discovery:

- Printers
- Network Storage (NFS, CIFS, NAS, FC, ISCSI)
- Azure virtual machines (IPv6 native mode isn't supported by Microsoft)

Requirements

Verify PowerShell support

ServiceNow now supports PowerShell 3.0 up to 5.1.

Configure Windows credentials

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Windows credentials.

Verify user access

Verify that the user configured for the credential has the following permissions:

- Local admin access to the Windows machine.
- Access to the WMI service to the current namespace and subnamespaces.
- Access to the PowerShell service.
- Membership in the Distributed COM Users local security group.

(Optional) Populate Virtual Machine Object field in Hardware [cmdb_ci_hardware] table

Starting with Discovery and Service Mapping Patterns version 1.30.2, you can improve query performance by populating the Virtual Machine Object field in the Hardware [cmdb_ci_hardware] table. For more information, see Improved query performance with direct field population in CI tables.

Set the preferred IP address version for network adapter discovery

Starting with Visibility Content version 6.32.0, if your network adapters support both IPv4 and IPv6, the IPv4 address is populated by default in the IP address [ip_address] field on the Network Adapter [cmdb_ci_network_adapter] table. To control which IP version is populated, see Set the preferred IP version for network adapter discovery.

Classifiers, probes, and patterns

Classifiers	Probes	Patterns
• Windows • Windows 2008 Server • Windows 2012 Server • Windows 2016 Server • Windows 2025 Server	• Horizontal discovery probe: Launches patterns • WMIRunner-Windows - Installed Software^ • MultiProbe-Windows - ADM^ • Windows - Identity* • Windows - Network ARP Table* • Windows - Network NDP Table* • Windows - OS Information* • Windows - Cluster* • Windows - CPU/Memory* • Windows - Installed Software* • Windows - Printers* • Windows - Storage 2008* • Windows - Storage 2012* • Windows - Amazon EC2* • Windows - Azure*	• Windows OS - Server • Windows OS - Desktops

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifiers	Probes	Patterns
	• DNS* • SNMP - Routing*	
Windows 2019 Server	N/A	Windows OS - Server

*These probes aren't active on the classifier, as Discovery uses patterns by default for these discoveries.

^These probes remain active by default, even when Discovery uses pattern discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Label	Table name	Field name	Source
Assigned to	cmdb_ci_win_server	assigned_to	wmi
Chassis type	cmdb_ci_win_server	chassis_type	wmi
Command	cmdb_running_process	command	wmi
Connects to	cmdb_running_process	connects_to	wmi
CPU core count*	cmdb_ci_computer	cpu_core_count	wmi
CPU core thread*	cmdb_ci_computer	cpu_core_thread	wmi
CPU count*	cmdb_ci_computer	cpu_count	wmi
CPU manufacturer	cmdb_ci_computer	cpu_manufacturer	wmi
CPU name	cmdb_ci_computer	cpu_name	wmi
CPU speed (MHz)	cmdb_ci_computer	cpu_speed	wmi
Default gateway	cmdb_ci_win_server	default_gateway	wmi
Description	cmdb_ci_disk	short_description	wmi
Disk space (GB)**	cmdb_ci_computer	disk_space	wmi
Disk space (GB)**	cmdb_ci_disk	disk_space	wmi
DHCP enabled	cmdb_ci_network_adapter	dhcp_enabled	wmi
DNS domain	cmdb_ci_win_server	dns_domain	DNS
Free space (GB)	cmdb_ci_file_system	free_space	wmi
Hostname	cmdb_ci_win_server	host_name	DNS, NBT
IP address***	cmdb_ci_network_adapter	ip_address	wmi
Listening on	cmdb_running_process	listening_on	wmi
MAC address	cmdb_ci_network_adapter	mac_address	wmi

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Table name	Field name	Source
Manufacturer	cmdb_ci_win_server	manufacturer	wmi
Model ID	cmdb_ci	model_id	wmi
Name	cmdb_ci_win_server	name	DNS, NBT
Name	cmdb_ci_disk	name	wmi
Name	cmdb_running_process	name	wmi
Name	cmdb_ci_network_adapter	name	wmi
Netmask	cmdb_ci_network_adapter	netmask	wmi
Operating System	cmdb_ci_computer	os	wmi
OS domain	cmdb_ci_computer	os_domain	NBT
OS service pack	cmdb_ci_computer	os_service_pack	wmi
OS version	cmdb_ci_computer	os_version	wmi
Parameters	cmdb_running_process	parameters	wmi
PID	cmdb_running_process	pid	wmi
RAM (MB)	cmdb_ci_computer	ram	wmi
Serial number	cmdb_ci_win_server	serial_number	wmi
Short description	cmdb_ci_win_server	short_description	wmi
Type	cmdb_ci_disk	type	wmi

* Core counts and threads per core might not be accurate, due to issues with Microsoft reporting.

** The value in the **disk_space** field is an aggregation of the total capacity (to include used space) for all non-removable disks, including both directly attached and SAN storage.

*** Starting with Visibility Content version 6.32.0, for network adapters that support both IPv4 and IPv6, the IPv4 address is populated by default. Before this release, the populated value was selected randomly. To control which IP version is populated, see Set the preferred IP version for network adapter discovery.

The Windows registry

Discovery can find software that has been installed on a Windows machine by looking at the Windows Registry. Discovery can find the following attributes of discovered software:

- **Product Name**: Combination of name and version, such as Windows Imaging Component 3.0.
- **Name**: Name of the product only without the version.
- **Version**: Version of the product.
- **Uninstall String**: Path to the uninstaller, such as `C:\Program Files\Notepad++\uninstall.exe`.
- **Part of**: Update for which this registry is a part, such as Windows Internet Explorer 8 - Software U.
- **Install Date**: The date the software was installed. The Windows - Installed Software sensor appends a timestamp of 00:00:00 to the install_date retrieved from the registry. The installation time of all Windows software is independent of the time zone and is set to midnight

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

of the day it was installed. For example, an install date of **2.19.2017** in the Windows registry appears as **2.19.2017 00:00:00** in the CMDB.

- **Installed on**: The name of the asset on which the software is installed.

Related topics

Windows probes and permissions

Windows server cluster discovery

Discovery establishes the relationships between a Windows server cluster and its nodes.

When a Windows server is found, and Discovery detects that it is part of a cluster, the system launches the Windows - Cluster probe. Navigate to **Discovery Definition** > **Probes** and open the record for the Windows - Cluster probe to see the specific data Discovery collects.

Note: Discovery collects cluster resources based on the MSCluster_ResourceType WMI class. To see a list of default cluster resources, see Resource Types. To understand how cluster resources are related to process classifiers, see Relate the process classifier to Windows cluster resources.

Cluster data tables

Discovery creates CIs for clusters in these tables:

- Windows Cluster [cmdb_ci_win_cluster]
- Windows Cluster Resource(s) [cmdb_ci_win_cluster_resource]
- Windows Cluster Node(s) [cmdb_ci_win_cluster_node]
- Cluster Node Resource(s) [cmdb_ci_cluster_node_resource]
- Cluster Virtual IP(s) [cmdb_ci_cluster_vip]

Cluster relationships

Discovery creates a CI Relationship [cmdb_rel_ci] record for each node, using these relationships:

- **Cluster of::Cluster**: Relationship between the cluster nodes and the cluster. Service Mapping uses this relationship to map the Windows cluster and it's nodes.
- **Members::Member of**: Relationship between the cluster and the cluster node.
- **Hosted on::Hosts**: Relationship between the cluster nodes and the Windows server.
- **Uses::Used by**: Cluster VirtualIP and Windows Cluster Node.
- **Defines resources for::Gets resources from**: Windows Cluster Resources and Windows Cluster.
- **Defines resources for::Gets resources from**: Windows Cluster Resources and Windows Cluster Node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows cluster map

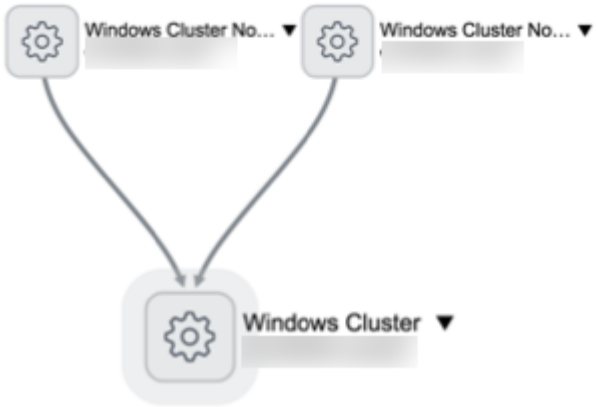

Windows cluster node relationships

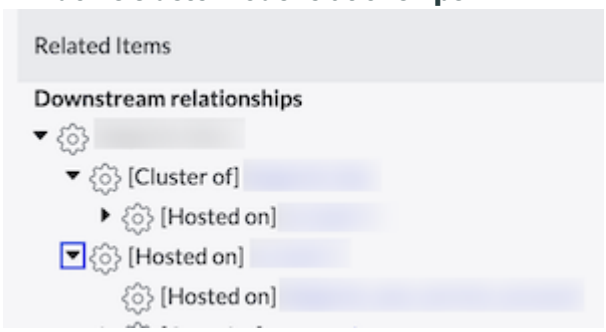

Cluster references

Discovery creates the following cluster references:

- From Windows Cluster Nodes to the associated Windows Cluster.
- From Windows Cluster Resources to the associated Windows Cluster.
- From Cluster Virtual IPs to the associated Windows Cluster.
- From Cluster Node Resources to the associated Windows Cluster Node.
- From Cluster Virtual IPs to the associated Windows Cluster Node.

Note: References are maintained in Discovery to provide backward compatibility for systems that use them.

z/OS discovery

Discovery and Service Mapping Patterns finds computers running the z/OS operating system using the IBM zOS Server pattern. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The IBM zOS Server pattern performs infrastructure-level discovery for z/OS servers. To discover applications on z/OS, use application patterns, such as DB2 on Linux, WMQ On Unix, or Websphere On Unix. For more information, see IBM z/OS Discovery requirements and expectations [KB0623031] article in the Now Support Knowledge Base.

Prerequisites

Verify that the UNIX System Services (USS) is running on port 22

Verify that the USS service is running on port 22 on z/OS.

Verify the required commands on the host machine

Verify that you can run the following commands on the host machine:

- `uname`
- `hostname`
- `sysvar SYSNAME`
- `netstat`
- `df`
- `ls`
- `cat`
- `grep`
- `ps`
- `cut`

Verify required permissions

Verify that you have the following permissions:

- Write permission to `/tmp`
- Read access to all user processes
- Ability to run REXX scripts

Create SSH credentials

For more information, see SSH credentials.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the IBM zOS Server pattern. The pattern discovers USS-mounted file systems.

IBM zOS server [cmdb_ci_ibm_zos_server]

Field	Description
Name [name]	Hostname of the z/OS system.
OS Version [os_version]	z/OS operating system version. For example: V2R5 or 3.1.
Operating System [os]	Operating system type. The value is set to **IBM_ZOS**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM zOS server [cmdb_ci_ibm_zos_server] (continued)

Field	Description
Manufacturer [manufacturer]	Hardware manufacturer. The value is set to **IBM**.
Serial number [serial_number]	Hardware chassis serial number.
Model ID [model_id]	Hardware model display name based on model code lookup. For example: z14, z15, or zEnterprise 196. If the model code isn't in the lookup table, this field remains empty.
IP Address [ip_address]	Management IP address used for discovery connection.

IP Address [cmdb_ci_ip_address]

Field	Description
IP Address [ip_address]	IPv4 or IPv6 address.
Netmask [netmask]	Network mask for the IP address. The value is set to **0.0.0.0**.
MAC Address [mac_address]	Hardware MAC address. If unavailable, the value is set to **NOT_FOUND**.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	Network interface name. If unavailable, the value is set to **VIP**.
IP Address [ip_address]	IPv4 or IPv6 address assigned to this network adapter.
Netmask [netmask]	Network mask for the adapter. The value is set to **0.0.0.0**.
MAC Address [mac_address]	Hardware MAC address for the network adapter. If unavailable, the value is set to **NOT_FOUND**.
Configuration Item [cmdb_ci]	References the IBM zOS server [cmdb_ci_ibm_zos_server] table.

File System [cmdb_ci_file_system]

Field	Description
Name [name]	Name of the USS file system set to the mount point path. For example: `/`, `/tmp`, or `/u`.
Volume ID [volume_id]	Volume identifier, set to the mount point path.
Mount point [mount_point]	Directory path where the file system is mounted in the USS hierarchy.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File System [cmdb_ci_file_system] (continued)

Field	Description
File system [file_system]	USS dataset name of the zFS file system. For example: OMVS . ROOT . ZFS or OMVS . TMP . ZFS.
Size bytes [size_bytes]	Total capacity of the file system in kilobytes (KB), as retrieved by the *df -Pkv* command.
Free space bytes [free_space_bytes]	Available free space in the file system in KB, as retrieved by the *df -Pkv* command.
Computer [computer]	References the IBM zOS server [cmdb_ci_ibm_zos_server] table.

CI relationships and references

The IBM zOS Server pattern creates the following relationships and references to support z/OS discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
IBM zOS server [cmdb_ci_ibm_zos_server]	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
IBM zOS server [cmdb_ci_ibm_zos_server]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
IBM zOS server [cmdb_ci_ibm_zos_server]	Contains::Contained by	File System [cmdb_ci_file_system]

CI references

CI	Field	Referenced CI
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	IBM zOS server [cmdb_ci_ibm_zos_server]
Serial Number [cmdb_serial_number]	Configuration item [cmdb_ci]	IBM zOS server [cmdb_ci_ibm_zos_server]
File System [cmdb_ci_file_system]	Computer [computer]	IBM zOS server [cmdb_ci_ibm_zos_server]

Related topics

- IBM Db2 on Linux database discovery
- IBM MQ discovery
- IBM WebSphere Application Server discovery using patterns

Set the preferred IP version for network adapter discovery

Control which IP version Discovery populates on a Network Adapter CI during Linux, Windows, and Solaris discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Verify that you have at least version 6.32.0 of Visibility Content.

Role required: discovery_admin

About this task
When a network adapter supports both IPv4 and IPv6, Discovery populates the IP Address [ip_address] field on the Network Adapter [cmdb_ci_network_adapter] CI with one IP address. Before Visibility Content version 6.32.0, Discovery selected the value randomly. Starting from this release, IPv4 is populated by default. Use the *network_adaptor_preferred_ip_version* MID Server property to specify which IP version to populate.

Procedure
1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** column, search for `network_adaptor_preferred_ip_version`.
3. Select the *network_adaptor_preferred_ip_version* property.
4. In the **Value** field, enter the accepted value for the IP version to populate.

   Values are case-sensitive. Use only the accepted values listed in the table.

   Accepted IP version values

IP version	Accepted values
IPv4	ipv4, IPv4, or v4.
IPv6	ipv6, IPv6, or v6.

5. Select **Update**.

What to do next
Run discovery again or wait for the scheduled run to apply the changes.

Related topics

Linux discovery

Solaris discovery

Windows discovery

Discover only the latest operating system patch versions

Collect only the current operating system patch level by configuring the *sn_itom_pattern.discover_latest_os_patches* MID Server property. Collecting only the current patch is supported only for AIX Server discovery.

Before you begin
Verify that you have at least version 6.35.0 of Visibility Content.

Role required: discovery_admin

About this task
By default, Discovery adds a record to the Patches [cmdb_ci_patches] table for every patch reported by the server, including older patch versions. The *sn_itom_pattern.discover_latest_os_patches* MID Server property controls this behavior; the default value is false. When set to true, Discovery evaluates all patches reported

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

by the server and identifies the latest patch level. Only that patch is populated as active in the CMDB.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. In the **Name** column, search for `sn_itom_pattern.discover_latest_os_patches`.
3. Select the *`sn_itom_pattern.discover_latest_os_patches`* MID Server property.
4. In the **Value** field, enter `true`.
   To revert to the default behavior and collect all patch versions on each run, set the value to `false`.
5. Select **Submit**.

What to do next

Run Discovery again to apply the change.

Related topics

AIX server discovery

Operating system-level virtualization discovery

Discovery can collect image and container information from Operating system-level virtualization (OS-level virtualization) engines.

Discovery can determine container status and size and identify images and their tags provided by operating-system-level virtualization.

Note: Currently, the ServiceNow® platform supports the discovery of Docker virtualization containers only.

Table schema

Table	Description
cmdb_ci_oslv_engine	Base table for all OSLV engines.
cmdb_ci_oslv_image	Base table for globally unique OSLV images.
cmdb_ci_oslv_local_image	Base table for local instances of OSLV images.
cmdb_ci_oslv_image_tag	Base table for tags on OSLV images.
cmdb_ci_oslv_container	Base table to store containers found on the host.

Data collected

Table	Fields
Operating System Level Virtualization Engine [cmdb_ci_oslv_engine]	Server [server]
Operating System Level Virtualization Image [cmdb_ci_oslv_image]	• Image ID [image_id] • Image digest [image_digest]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Fields
	• Image created [image_created_at] • Size (bytes) [size_bytes]
Operating System Level Virtualization Local Image [cmdb_ci_oslv_local_image]	Image ID [image_id]
Operating System Level Virtualization Image Tag [cmdb_ci_oslv_image_tag]	• Image ID [image_id] • Repository [repository] • Tag [tag]
Operating System Level Virtualization Container [cmdb_ci_oslv_container]	• Container ID [container_id] • Status [status] • Size (bytes) [size_bytes] • Image ID [image_id]

Docker virtualization

Discovery uses the *Docker Pattern* to collect data about specific objects in a Docker engine, running on a Linux host.

The ServiceNow® platform supports the discovery of Docker release 1.11.0 or later.

Discovery runs the Docker Engine process classifier in the network. If the classifier identifies the **dockerd** or **docker daemon** process, the classifier triggers the Horizontal Pattern (HorizontalDiscoveryProbe) probe, which launches the Docker Pattern and begins collecting data from Docker components.

Starting with Discovery and Service Mapping Patterns version 1.31.0, you can choose to discover Docker image CIs only, without discovering Docker container CIs. Check your entitlements to determine whether you have access to 2026 Packaging SKU. For more information, see Disable Docker container CI discovery.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Prerequisites

User privileges

The user whose credentials are used to perform Docker Discovery must have privileges defined by one of these methods:

- Provide a user with elevated rights for running commands, since the Docker daemon runs as the root user. The Docker pattern supports the use of privileged commands, such as *sudo* or *pbrun*, to run as the root user.
- Assigned to a group named **docker**, which has special privileges for running Docker commands. For instructions on setting up a group, see documentation on Docker.

Latest patterns

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Deploy the latest Discovery and Service Mapping Patterns application from ServiceNow Store.

Docker restrictions and considerations

When using Docker virtualization, consider the following:

- The initial Discovery process scan can identify an application in a container and classify it correctly. However, subsequent probes launched to explore that application cannot see inside the container and cannot return details about the application.
- Discovery scans all containers it finds, including inactive containers, which can slow down Discovery. You should delete containers that are not running.
- Before setting *sn_itom_pattern.manifest_digest_image_id* property to **true** and running discovery: prevent duplicate Docker records from being created by deleting all Docker image records.
- Only one Docker engine is permitted per computer (on either a physical or virtual machine).

Data collected by Discovery during horizontal discovery

These attributes are discovered, in addition to the attributes derived from the parent OSLV tables.

Table and fields	Description
Docker Engine [cmdb_ci_docker_engine]	
Name [name]	Stores information about instances of the Docker engine.
OS Arch [os_arch]	
GIT Commit [git_commit]	
Build Date [build_date]	
Version [version]	
API Version [api_version]	
Go Version [go_version]	
Is Clustered [is_clustered]	
Running Process [running_process]	
Running Process Command [running_process_command]	
Running Process Key Parameters [running_process_key_parameters]	
Docker Image [cmdb_ci_docker_image]	
Name [name]	Stores information on the globally unique representation of Docker images.
Image id [image_id]	
Image digest [image_digest]	
Size (byte) [size_byte]	
Image created [Image_created_at]	
Docker Local Image [cmdb_ci_docker_local_image]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and fields	Description
Name [name]	Stores local instances of Docker images.
Image id [image_id]	
Docker Image Tag [cmdb_ci_docker_image_tag]	
Name [name]	Stores tags from local Docker images.
Image id [image_id]	
Repository [repository]	
Tag [tag]	
Docker Container [cmdb_ci_docker_container]	
Name [name]	Stores Docker containers found on the host. In cases where duplicate records are created, deduplication tasks appear once discovery runs. For information on how to resolve these tasks, see the Making docker container identifier independent [KB1443042] article in the ServiceNow® Knowledge Base.
Image id [image_id]	
Container id [container_id]	
Size (bytes) [size_bytes]	
Command [command]	
Container created [container_created]	
Status [status]	

The *Collect Container Repository* extension section discovers this information.

Container Repository [cmdb_ci_container_repository]

Field	Description
Name [name]	The name of the container repository.

Container Repository Entry [cmdb_ci_container_repository_entry]

Field	Description
Name [name]	The name of the container repository entry.
Category [category]	The category of the container repository entry.

CI relationships

The Docker Pattern identifies the following relationships between discovered CIs.

CI	Relation	CI
cmdb_ci_server	Runs::Runs On	cmdb_ci_docker_engine
cmdb_ci_docker_engine	Manages::Managed By	cmdb_ci_docker_container
cmdb_ci_docker_engine	Manages::Managed By	cmdb_ci_docker_local_image
cmdb_ci_docker_image	Instantiates::Instantiated By	cmdb_ci_docker_container

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relation	CI
cmdb_ci_docker_local_image **Note:** The direction of this relation is the opposite for all store versions earlier than 1.0.99, February 2023.	Instantiates::Instantiated By	cmdb_ci_docker_container
cmdb_ci_docker_image_tag	Registered On::Has Registered	cmdb_ci_docker_local_image
cmdb_ci_docker_local_image	Instantiates::Instantiated By	cmdb_ci_docker_image

The *Collect Container Repository* extension section of the Docker Pattern identifies these relationships.

CI	Relationship	CI
Docker Image [cmdb_ci_docker_image]	Provisioned From::Provisioned	Container Repository Entry [cmdb_ci_container_repository_entry]
Container Repository Entry [cmdb_ci_container_repository_entry]	Hosted on::Hosts	Container Repository [cmdb_ci_container_repository]

Identification, containment, and hosting rules

Discovery uses an application rule identifier to find the Docker engine and then applies other rules to identify specific Docker objects.

Application rule identifier

The system creates the cmdb_ci_docker_engine configuration item (CI) during process classification. Based on this, Discovery uses the Application Rule identifier on the Application [cmdb_ci_appl] table to identify the particular Docker engine encountered. After establishing this identity, Discovery applies the relationships defined in the containment and hosting rules. It then creates and updates the individual Docker component CIs related to that engine.

Identifiers

Name	Table	Attributes
Docker Container	Docker Container [cmdb_ci_docker_container]	container_id
Docker Global Image	Docker Image [cmdb_ci_docker_image]	image_id
Docker Local Image	Docker Local Image [cmdb_ci_docker_local_image]	image_id
Docker Image Tag	Docker Image Tag [cmdb_ci_docker_image_tag]	repository, tag

Containment and hosting rules

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Docker Discovery uses these Create or edit a collection of containment rules and Create or edit a collection of hosting rules rules to create configuration items (CI) from the data returned by the Docker Pattern. After Discovery identifies the Docker engine by its relationship to the Application [cmdb_ci_appl] table, it uses these rules to identify the specific CIs connected to that engine from their relationships to one another. By connecting the components to one another in this fashion, from the application down, starting with the engine, Discovery avoids creating duplicate CIs for components from other Docker engines that use the same name or image_id.

Containment rule

Parent	Child	Relationship
Docker Local Image	Docker Image Tag	Has registered

Hosting rules

Parent	Child	Relationship
Docker Container	Docker Engine	Managed by
Docker Local Image	Docker Engine	Managed by

Software discovery

Discovery identifies several types of software.

Active Directory Domain Controller discovery

The Discovery and Service Mapping Patterns application uses the Active Directory Domain Controller On Windows pattern to find Active Directory domain controllers running on a Windows server. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

When discovery runs on a domain controller, the pattern also creates LDAP endpoint connections to child domains in the same forest. This triggers discovery of those domain controllers as well.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Prerequisites

Verify that the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

Verify an LDAP or TCP entry point is configured

Verify that the relevant port is open and listening on the domain controller.

Create Windows credentials with WMI access

Verify that the credentials have permission to run WMI queries on the `\\root\CIMV2` namespace. For more information, see Windows credentials.

Schedule a horizontal discovery

For more information, see Schedule a horizontal discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Active Directory Domain Controller On Windows pattern.

Active Directory Domain Controller [cmdb_ci_ad_controller]

Field	Description
Name [name]	Name of the domain controller, which is the domain name.
Domain name [domain_name]	Name of the Active Directory domain.
Domain Controller name [domain_controller_name]	Name of the domain controller server.
Forest name [forest_name]	DNS name of the Active Directory forest.

CI relationships

The Active Directory Domain Controller On Windows pattern creates the following relationships to support Active Directory Domain Controller discovery.

CI relationships

CI	Relationship	CI
Active Directory Domain Controller [cmdb_ci_ad_controller]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

Adobe JRun discovery

Discovery creates or updates a CMDB record when it detects a running instance of Adobe JRun.

By default, Discovery uses the *Jrun* pattern to perform the Discovery.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Privileged commands

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Data collected for Adobe JRun

The following data is collected in the Jrun `[cmdb_ci_app_server_jrun]` table:

Label	Field Name
Name	name
Version	version
Installation directory	install_directory
Configuration directory	config_directory

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field Name
Configuration file	config_file

Apache web server discovery

Discovery identifies and classifies information about Apache web servers on both Windows and Linux computers.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Credentials

- Applicative credentials
- Windows credentials
- SSH credentials

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Classifier, patterns, and probes

Classifier	Trigger probes	Patterns
Apache Server	• Horizontal discovery probe: Launches patterns • Apache - Version* • Apache - Get Configuration* • Apache - Get JK ModuleApache - Get Proxy Module*	• Apache On Windows Pattern • Apache On Unix Pattern

*For new instances, these probes are inactive on the classifier. Discovery uses patterns for discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected using probes

Discovery collects the following data on the Apache [cmdb_ci_apache_web_server] table for both Unix and Windows:

Label	Field name
Name	name
Version	version
Description	description

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field name
TCP port(s)	tcp_port

Connections found by Service Mapping during top-down discovery

Service Mapping uses discovery patterns to discover the following connections:

- Any application using java
- Any application using the LDAP port
- Oracle WebLogic Module Oracle, and WebLogic Server (version 10.3)
- (On Windows only) Jiva applications
- (On Windows only) Tomcat server
- (On UNIX only) Applications discovered using the mod_jk process
- (On UNIX only) Hosts to which the Apache web server is connecting in the load balancing mode
- (On UNIX only) Epic application

Apache mod_jk and mod_proxy discovery

Discovery uses SSH to find Apache connectors mod_jk and mod_proxy.

Discovery uses the **Unix - Active Processes** probe to identify an Apache web server process by matching one of the following criteria: the **name** of the process is:

- `httpd`, or
- `apache`, or
- `httpd2`, or
- `httpd2-prefork`
- `httpd.worker`

If there is a match on one of these criteria, a record is created in the Web Server [cmdb_ci_web_server] table if one does not already exist for that running process.

The following probes are triggered after classification:

Apache web server probes

Probe	Description	Commands
Apache – Version	The sensor of this probe populates the Apache version information in the Web Server record.	httpd
Apache – Get Configuration	This probe contains a Bourne shell script and an argument that determines the path of the Apache configuration file. The sensor of this probe populates some additional information in the Web Server record.	echo, sed, httpd, cut, grep, egrep (within the Borne shell script)

The sensor processing of the **Apache – Get Configuration** probe identifies whether either the mod_jk or mod_proxy modules are present and triggers the appropriate probe.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache module probes

Probe	Description	Commands
Apache – Get JK Module	If the mod_jk module is running as a load balancer on the server, the sensor of this probe populates the information in the Load Balancer Service [cmdb_ci_lb_service], Load Balancer Pool [cmdb_ci_lb_pool] and Load Balancer Pool Member [cmdb_ci_lb_pool_member] tables.	echo, sed, httpd, cut, grep, egrep (within the Borne shell script)
Apache – Get Proxy Module	If the mod_proxy module is running as a load balancer on the server, the sensor of this probe populates the information in the Load Balancer Service [cmdb_ci_lb_service], Load Balancer Pool [cmdb_ci_lb_pool] and Load Balancer Pool Member [cmdb_ci_lb_pool_member] tables.	grep, egrep (within the Borne shell script)

In addition to data population, the following relationships are created in the CI Relationship [cmdb_rel_ci] table:

- The records in the cmdb_ci_lb_appl table run on the cmdb_ci_web_server table records.
- The records in the cmdb_ci_lb_service table use the cmdb_ci_lb_pool table records.
- The records in the cmdb_ci_lb_pool table are used by the cmdb_ci_service table record.
- The records in the cmdb_ci_lb_pool table are members of the cmdb_ci_lb_pool_member table.
- The records in the cmdb_ci_lb_pool_member table are members of the cmdb_ci_lb_pool table.

Legacy: Add sudo access for the Unix - Active Processes probe

The Unix - Active Processes probe requires sudo privileges on Apache web servers.

Before you begin

Role required: admin

About this task

You can configure the Unix - Active Processes probe to elevate privileges.

Procedure

1. Navigate to **All** > **Discovery** > **Probes**.
2. In the **Search** field, search for `Unix - Active Processes`.
3. Select the **Unix - Active Processes** link from the table.
4. In the **Probe Parameters** related list, select **New**.
5. Use the following information to fill out the form:

Probe	Commands
Name	must_sudo
Value	true

6. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Mod_jk module

The Apache mod_jk module forwards requests from the Apache web server to a Servlet container, such as Tomcat.

Additional mod_jk directives can also manage load balancing. Discovery populates the CMDB when it detects an Apache Server. If the mod_jk module is running on the web server as a load balancer, Discovery populates the related information to the CMDB.

- The MID Server user account to explore the target server must have access to the httpd.conf configuration file in the /etc/httpd/conf/ folder.
- Discovery uses secure shell (SSH) commands to identify the following associated elements:
  - Apache Get Configuration
  - Apache Version
  - Apache Get JK Module
- The following probes require execute privileges to run commands:

Probes that require execute privileges

Probe	Command
Apache – Get Configuration	echo, sed, httpd, cut, grep, egrep (within the Borne shell script)
Apache – Version	httpd
Apache – Get JK Module	echo, sed, httpd, cut, grep, egrep (within the Borne shell script)

Probes and sensors

Discovery uses the Unix - Active Processes probe to identify an Apache server that contains the mod_jk module:

1. The Unix - Active Processes probe detects a running process that matches one of the following criteria:
   - The name of the process is httpd.
   - The name of the process is apache.
2. If there is a match on one of these criteria, a record is created in the Web Server table [cmdb_ci_web_server] if one does not already exist for that running process. The following probes are also triggered:
   - Apache – Version: the sensor of this probe populates the Apache version information in the Web Server record.
   - Apache – Get Configuration: this probe contains a Bourne shell script and an argument that determines the path of the Apache configuration file. The sensor of this probe populates some additional information in the Web Server record.
3. The sensor processing of Apache – Get configuration probe results triggers the following probes if the mod_jk module is running on the web server:

   Apache – JK Module:

   If the mod_jk module is running as a load balancer on the server, the sensor of this probe populates information in the following tables. The Load Balancer Service [cmdb_ci_lb_service], Load Balancer Pool [cmdb_ci_lb_pool], and Load Balancer Pool Member [cmdb_ci_lb_pool_member] tables are populated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data Collected

For the mod_jk module with no load balancer, the following data is collected by default:

Default Data Collected by Discovery

Table Name	Field and Label Name	Probe
cmdb_ci_web_server	Name [name]	Apache – Get Configuration
cmdb_ci_web_server	Version [version]	Apache – Version
cmdb_ci_web_server	Description [short_description]	Apache – Version

If the mod_jk module is enabled for load balancing, Discovery connects the following data:

Data Collected by Discovery When the mod_jk Module is Identified as a Load Balancer

Table Name	Field and Label Name	Probe
cmdb_ci_lb_appl	Name [name]	Apache – Apache - Get JK Module
cmdb_ci_lb_appl	IP Address [ip_address]	Apache – Apache - Get JK Module
cmdb_ci_lb_appl	Last Discovered [last_discovered]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool_member	Name [name]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool_member	Last Discovered [last_discovered]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool_member	IP Address [ip_address]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool_member	Load Balancer [load_balancer]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool_member	Port [port]	Apache – Apache - Get JK Module
cmdb_ci_lb_service	Input URL [Input_url]	Apache – Apache - Get JK Module
cmdb_ci_lb_service	Last Discovered [last_discovered]	Apache – Apache - Get JK Module
cmdb_ci_lb_service	IP Address [ip_address]	Apache – Apache - Get JK Module
cmdb_ci_lb_service	Name [name]	Apache – Apache - Get JK Module
cmdb_ci_lb_service	Load Balancer [load_balancer]	Apache – Apache - Get JK Module
cmdb_ci_lb_service	Port [port]	Apache – Apache - Get JK Module

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data Collected by Discovery When the mod_jk Module is Identified as a Load Balancer (continued)

Table Name	Field and Label Name	Probe
cmdb_ci_lb_pool	Last Discovered [last_discovered]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool	Load balancing Method [load_balancing_method]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool	Load Balancer [load_balancer]	Apache – Apache - Get JK Module
cmdb_ci_lb_pool	Name [name]	Apache – Apache - Get JK Module

Relationships

In addition to data population, the following relationships are created in the CI Relationship [cmdb_rel_ci] table:

- The records in the cmdb_ci_lb_appl table run on the cmdb_ci_web_server table records.
- The records in the cmdb_ci_lb_service table use the cmdb_ci_lb_pool table records.
- The records in the cmdb_ci_lb_pool table are used by the cmdb_ci_service table record.
- The records in the cmdb_ci_lb_pool table are members of the cmdb_ci_lb_pool_member table.
- The records in the cmdb_ci_lb_pool_member table are members of the cmdb_ci_lb_ pool table.

Data collected by Mod_proxy module

The Apache mod_proxy module implements a proxy, gateway, or cache for the Apache web server.

An additional mod_proxy_balancer can also manage load balancing. Discovery populates the CMDB when it detects an Apache server. When the Apache Web Server process is detected, if the mod_proxy module is running on the web server as a load balancer the related information populates to the CMDB.

Consider the following requirements for discovering an Apache server that contains the mod_proxy module:

- The MID Server user account to explore the target server must have access to the [httpd.conf] configuration file in the /etc/httpd/conf/ folder.
- Discovery uses secure shell (SSH) commands to identify the following associated elements:
  - Apache Get Configuration
  - Apache Version
  - Apache Get Proxy Module
- The following probes require execute privileges to run commands:

Data collected from the mod_jk module

Probe	Commands
Apache – Get Configuration	echo, sed, httpd, cut, grep, egrep (within the Borne shell script)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	Commands
Apache – Get Proxy Module	grep, egrep (within the Borne shell script)
Apache – Version	httpd

Probes and sensors

Discovery uses the Unix - Active Processes probe to identify an Apache server that contains the mod_proxy module. The probes and sensors operate in the following manner:

1. The Unix - Active Processes probe detects a running process that matches one of the following criteria:
   - The name of the process is httpd.
   - The name of the process is apache2.
2. If there is a match on one of these criteria, a record is created in the Web Server table [cmdb_ci_web_server] if one does not already exist for that running process. The following probes are also triggered:
   - Apache – Version: the sensor of this probe populates the Apache version information in the Web server [cmdb_ci_web_server] record.
   - Apache – Get Configuration: this probe contains a Bourne shell script and an argument that determines the path of the Apache configuration file. The sensor of this probe populates some additional information in the Web server [cmdb_ci_web_server] record.
3. The sensor processing of the Apache – Get configuration probe results triggers the following probes if the mod_proxy module is running on the web server:

   Apache - Get Proxy Module:

   If the mod_proxy module is running as a load balancer on the server, the sensor of this probe populates information in the following tables. Load Balancer Service [cmdb_ci_lb_service], Load Balancer Pool [cmdb_ci_lb_pool], and Load Balancer Pool Member [cmdb_ci_lb_pool_member].

Data Collected

For the mod_proxy module with no load balancer, the following data is collected by default:

Data collected by Discovery by default

Table Name	Field and Label Name	Probe
cmdb_ci_web_server	Name [name]	Apache – Get Configuration
cmdb_ci_web_server	Version [version]	Apache – Version
cmdb_ci_web_server	Description [short_description]	Apache – Version

If the mod_proxy module is enabled for load balancing, Discovery connects the following data:

Data collected by Discovery if the mod_proxy module is identified as a load balancer

Table Name	Field and Label Name	Probe
cmdb_ci_lb_appl	IP Address [ip_address]	Apache - Get Proxy Module

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery if the mod_proxy module is identified as a load balancer (continued)

Table Name	Field and Label Name	Probe
cmdb_ci_lb_appl	Last Discovered [last_discovered]	Apache - Get Proxy Module
cmdb_ci_lb_pool_member	Name [name]	Apache - Get Proxy Module
cmdb_ci_lb_pool_member	Last Discovered [last_discovered]	Apache - Get Proxy Module
cmdb_ci_lb_pool_member	IP Address [ip_address]	Apache - Get Proxy Module
cmdb_ci_lb_pool_member	Load Balancer [load_balancer]	Apache - Get Proxy Module
cmdb_ci_lb_pool_member	Port [port]	Apache - Get Proxy Module
cmdb_ci_lb_service	Input URL [Input_url]	Apache - Get Proxy Module
cmdb_ci_lb_service	Last Discovered [last_discovered]	Apache - Get Proxy Module
cmdb_ci_lb_service	IP Address [ip_address]	Apache - Get Proxy Module
cmdb_ci_lb_service	Name [name]	Apache - Get Proxy Module
cmdb_ci_lb_service	Load Balancer [load_balancer]	Apache - Get Proxy Module
cmdb_ci_lb_service	Port [port]	Apache - Get Proxy Module
cmdb_ci_lb_pool	Last Discovered [last_discovered]	Apache - Get Proxy Module
cmdb_ci_lb_pool	Load balancing Method [load_balancing_method]	Apache - Get Proxy Module
cmdb_ci_lb_pool	Load Balancer [load_balancer]	Apache - Get Proxy Module
cmdb_ci_lb_pool	Name [name]	Apache - Get Proxy Module

Relationships

In addition to data population, the following relationships are created in the CI Relationship [cmdb_rel_ci] table:

- The records in the cmdb_ci_lb_appl table run on the cmdb_ci_web_server table records.
- The records in the cmdb_ci_lb_service table use the cmdb_ci_lb_pool table records.
- The records in the cmdb_ci_lb_pool table are used by the cmdb_ci_service table records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The records in the cmdb_ci_lb_pool are members of the cmdb_ci_lb_pool_member table records.
- The records in the cmdb_ci_lb_pool_member is a member of the cmdb_ci_lb_pool table records.

Application Mapping for UNIX discovery

To perform the mapping that establishes application relationships, Discovery must be able to detect TCP connections.

The best way to accomplish this is to run the lsof command on target machines. For AIX machines, Discovery is configured to use the active connection probe for AIX targets. However, the user that executes the shell script on AIX machines might require additional privileges to execute commands, such as kdb, employed in the script.

On Linux machines, Discovery uses the `lsof` command (installed by default on Linux) to detect TCP connections. The Linux classification probe triggers the Linux - Active Connections probe, which uses `lsof` to discover application relationships and does not produce any pauses. The `lsof` command is the recommended method of returning active TCP connections and can be installed on Solaris and AIX target machines to eliminate any issues produced by the shell script.

Note: Discovery of active connections on HP-UX is not supported.

Using the lsof Command

Use the ls of command to return active TCP connections and can be installed on Solaris and AIX target machines.

Before you begin

Role required: admin

Procedure

1. Install `lsof` on the Solaris and AIX machines on which you want to gather application relationship data.
2. In the instance, navigate to **Discovery Definition** > **CI Classification** > **UNIX**.
3. Select **Solaris** from the list.
4. In the UNIX Classification form for Solaris, select the **Triggers probes** related list, and then select **Edit**.
5. Move **Linux - Active Connections** from the Collections list to the Triggers probe List.
   This probe is configured to use `lsof`.
6. Select **Save**.
7. Repeat this procedure for the AIX classification probe.

Citrix License Server and Delivery Controller discovery

Discovery can find Citrix License Server instances and Delivery Controller components that run on Windows machines.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Requirements

- Configure Windows administrative credentials to access the host machine.
- Make sure that PowerShell for Discovery is installed on your MID Servers. The patterns used for Citrix discovery require the ability to run PowerShell commands.
- For OS users, make sure you can execute “Get-XDSite.” The Citrix Studio Admin should provide permissions.
- (Optional) For top-down discovery with Service Mapping, on Citrix Delivery Controller, install the Citrix.Broker module.
- (Optional) For top-down discovery with Service Mapping, a user with permission to run PowerShell commands.
- (Optional) (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.

Patterns

Discovery uses these patterns to perform horizontal discovery:

Pattern	Process detected
Citrix License Server	lmadmin.exe
Delivery Controller	BrokerService

Data collected and relationships

Table	Fields
Citrix License Server [cmdb_ci_appl_license_server]	Name [name] TCP ports [tcp_port] Operational status [operational_status]
Citrix Delivery Controller [cmdb_ci_appl_delivery_controler]	Name [name] Version [version] Operational status [operational_status]

CI	Relationship	CI
Citrix License Server [cmdb_ci_appl_license_server]	Runs on::Runs	Windows Server [cmdb_ci_win_server]
Citrix Delivery Controller [cmdb_ci_appl_delivery_controler]	Runs on::Runs	Windows Server [cmdb_ci_win_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Clustered application discovery on Windows

A process and its corresponding resource information can be used to determine whether the process is a clustered process.

The **Microsoft SQL Server** process classifier that comes with your instance identifies clustered SQL Servers. You can leverage the same technique to detect other clustered applications on Windows using Discovery.

To detect Windows clustered applications with Discovery, you create a process classifier that detects whether a process is a clustered process. Perform the following steps:

Create a classifier for clustered processes

To detect Windows clustered applications, create a process classifier for that CI type.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **Processes**.
2. Click **New**.
3. Fill out the form using the fields from the table.
4. Click **Submit**.

Process Classification form

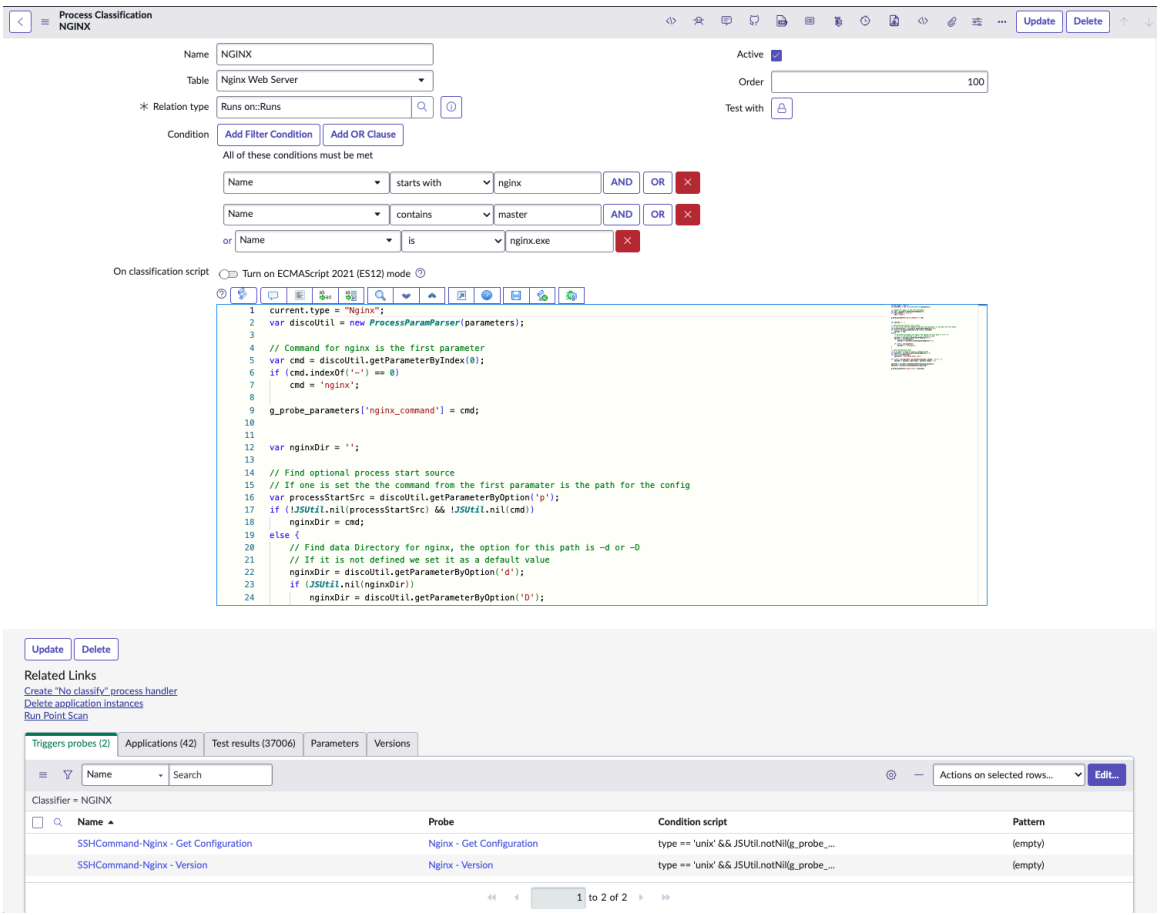

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Process classification form

Field	Description
Name	A unique name for the process classifier.
Table	Select **Application**.
Relation type	Select the CI relationship type for this classification. The relationship field is only available for Process and Scan Application classifications. Discovery process classifications typically use one of these relationship types: ○ **Runs on::Runs**: Defines the relationship of an application to the host on which it runs. This relationship is expressed from the perspective of the host and the application. For example: My database application runs on server001::server001 runs my database application. ○ **Depends on::Used by**: Defines the relationship of an application that communicates with another application. This relationship is expressed from the perspective of each application. For example: The Tomcat application depends on the MySQL database:: The MySQL database is used by Tomcat. ○ **Virtualized by::Virtualizes**: Defines the relationship of a virtual machine to its host. This relationship is expressed from the perspective of the virtual machine and of the host. For example: server001 is virtualized by ServerESX::Server ESX virtualizes server001. ○ **Cluster of::Cluster**: Defines the relationship of a cluster node to the cluster to which it belongs. This relationship is expressed from the perspective of the cluster node and of the cluster. ○ **Hosted on::Hosts**: Defines the relationship of a cluster node and the Windows host. This relationship is expressed from the perspective of the cluster node and of the Windows host.
Active	Select this box to activate the process classification. Deselect to deactivate it.
Order	Enter the order in which Discovery should run this process classification when there are multiple classifications available for a table. Discovery runs process classifications from the lowest to the highest order.
Test with	Lists the host CI where an automatically generated process classification conditions were met. Use this field to test changes to the process classification to ensure that the updated classification behaves as expected.
Condition	Configure the appropriate condition. The example screenshot triggers this process classification when the process command contains the string `MSFT SQL Server 2000`.

Relate the process classifier to Windows cluster resources

In the second step, relate the process classifier to the information that is in the Windows Cluster resources table.

Before you begin

Role required: admin

About this task

The Windows Cluster resources table is cmdb_ci_win_cluster_resource. This relationship is built through JavaScript.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate back to the process classification you just made.
2. Select the **Parameters** tab.
3. Select **New**.
4. Fill out the fields on the form.

Classifier parameter form

Field	Description
Name	Select a unique name.
Application	Select **Global**.
Type	Enter **Cluster**.
Value	This field defines two JavaScript objects, called `resourceType` and `isMatch`. For more information, see Populate the Value field on the Classifier Parameter form.

5. Select **Submit**.

Populate the Value field on the Classifier Parameter form

The value field defines two JavaScript objects, called `resourceType` and `isMatch`.

Before you begin
Role required: admin

About this task
The `resourceType` value refers to the Resource Type column in the Windows Cluster Resources table [cmdb_ci_win_cluster_resource]. The `resourceType` value can't be empty or null.

Procedure

1. Navigate to a Windows cluster page to access the Windows Cluster Resources table.
2. Select the **Windows Cluster Resources** tab.

Example:
The **isMatch** value evaluates whether the process is a clustered process.

- An evaluation that returns **true** indicates the process is a clustered application, and the process is classified as a clustered application in the cluster.
- An evaluation that returns **false** classifies the process as a regular application running on the cluster node.

The **isMatch** value is a function that contains two input variables, process and resource:

- process: Process is the GlideRecord of the process application. It is determined by the **Table** field in the classifier. In this example, it is the GlideRecord entry of the Application table (cmdb_ci_app) for the process that is being classified. You have access to any field values for the CI type such as name or version.
- resource: Resource is the GlideRecord entry in the Windows Cluster Resource table after the resourceType condition has been applied. In the example, it is the GlideRecord entry of the sixth row.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Examples of JavaScript for the Value field.

In the following example code: resourceType: "SQL Server" the query filters the Windows Cluster Resources table entry for the Resource Type field that has a value that is equal to SQL Server. In the following example table, the sixth record would be returned.

The following script indicates that if there is a resource of type SQL Server, the process is classified as a clustered application.

```
resourceType: "SQL Server"
isMatch: function(process, resource) {
    return true;
}
```

The following script indicates that if there is a resource of type SQL Server, and the application name is equal to the resource name, then the process is classified as a clustered application.

```
resourceType: "SQL Server"
isMatch: function(process, resource) {
    if (process.name == resource.name)
        return true;
}
```

If there are multiple matches to the resourceType condition, the matching function is called multiple times. For example, if there are two entries with the same resource type (such as Physical Disk), the matching function runs twice—once for each matching entry (for example, Cluster Disk 1 and Cluster Disk 2).

Discovery with Software Asset Management

The table structure for managing software installations behaves differently when the Software Asset Management (SAM) application is activated.

The presence of this application affects the way Discovery stores the software installation data it returns from a scan and require additional configuration within configuration item (CI) records.

The following tables are created when the Software Asset Management plugin is activated:

- Software Discovery Model [cmdb_sam_sw_discovery_model]
- Software Model [cmdb_software_product_model]
- Software Installation [cmdb_sam_sw_install]

> **Note:** The original schema used by previous versions is not affected when the Software Asset Management plugin is not activated.

When it runs, Discovery populates the Software Installation [cmdb_sam_sw_install] table. The appropriate configuration item (CI) references the software data. A business rule on this table runs, searching for a matching record on the Software Discovery Model [cmdb_sam_sw_discovery_model] table. If a record exists in the Software Discovery Model table, then the Software Installation table updates the record. If no record exists in the Software Discovery Model table, then one is created.

These tables are not used by Discovery with Software Asset Management:

- Software [cmdb_ci_spkg]
- Software Model [cmdb_software_product_model]

> **Important:** When SAM is installed, the Software Installation [cmdb_sam_sw_install] table is the appropriate source for all current software data. This means you need to update any related lists or customized reference fields you added to CI records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Starting with the Kingston release, the Discovery populates the VMWare License Key [samp_vmware_license_key] and VMWare License Key Usage [samp_vmware_license_key_usage] tables. See Data collected for VMware Cloud Discovery for more information on where Discovery saves vCenter data.

Configure a CI to display Software Asset Management data

Configure computer configuration item (CI) forms to display data from the Software Installations [cmdb_sam_sw_install] table.

Before you begin

Confirm you have installed and activated the Software Asset Management (com.sn_samp_master_ws) plugin.

Role required: admin

Procedure

1. Navigate to **All** > **Configuration** > **Servers** > **All**.
   Alternatively, enter `cmdb_ci_computer.list` in the filter navigator to open the Computer [cmdb_ci_computer] table list view.
2. Select a **Name** from the list to open any computer CI record.
3. Right-click in the header of the record and select **Configure** > **Related Lists** from the context menu.
   The Configuring related lists on Computer form page displays.
4. Scroll through the Available list and select **Software Installation->Installed on**.
5. Select the add icon.
   Software Installation->Installed on now displays in the Selected list.
6. Select **Save**.
7. Select the back icon to return to the computer record.
   The Software Installations tab now displays in the record's related list.

Microsoft Exchange Client Access Server discovery

Discovery creates or updates a CMDB record when it detects a running instance of Microsoft Exchange Client Access Server (Microsoft Exchange CAS).

By default, Discovery uses the *CAS On Windows* pattern to discover the Exchange Client Access Server.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery on the Exchange Client Access Server

The following data is collected on the Exchange Client Access Server `[cmdb_ci_exchange_cas]` table:

Label	Field Name
Name	name
Version	version
Class	sys_class_name
Fully qualified domain name	fqdn

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field Name
IP Address	ip_address
Installation directory	install_directory
Type	type

Exchange Hub Transport Servers discovery

Discovery creates or updates a CMDB record when it detects a running instance of Exchange Hub.

By default, Discovery uses the Hub Transport Server On Windows pattern to discover Exchange Hub Transport Servers.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the Hub Transport Server On Windows pattern.

ExchangeHub [cmdb_ci_exchange_hub]

Label	Field Name
Name	name
Version	version
Class	sys_class_name
Fully qualified domain name	fqdn
IP Address	ip_address
Installation directory	install_directory
Type	type

CI relationships

The Hub Transport Server On Windows pattern doesn't discover any relationships.

Exchange MailBox discovery

Discovery creates or updates a CMDB record when it detects a running instance of Exchange Mailbox.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Prerequisites

The `Microsoft.Exchange.Management.PowerShell` module must be installed on the server that you want to discover.

The service `winRM` must be running on the host machine. Run the following command in the Powershell window to verify: `New-PSSession -ConfigurationName`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Microsoft.Exchange -ConnectionUri "http://$hostname/PowerShell/"
-Authentication Kerberos
```

Pattern

By default, Discovery uses the *MailBox On Windows* pattern to discover Exchange MailBox.

Credentials

Configure these credentials:

Credential type	Additional notes
Applicative credentials	The user must be able to run Powershell commands against the Exchange hosts. Use **Exchange Mailbox** as the CI type for the applicative credential. **Important:** Do not use the **Exchange Mailbox server** CI type.
Windows credentials	The OS user must be a domain user, not a local user for the server.

Data collected

The following data is collected on the Exchange Mailboxes [cmdb_ci_exchange_mailbox] table:

Label	Field Name
Name	name
Version	version
Class	sys_class_name
Fully qualified domain name	fqdn
IP Address	ip_address
Installation directory	install_directory
Type	type

General software package discovery

General software packages are collections of programs that work together, often with similar user interfaces. Examples include Microsoft Office 365 (Word, Excel, PowerPoint), and Apple's iWork (Pages, Numbers, Keynote). Discovery identifies and classifies information about general software packages. Data it collects include, Name, Version, Install Count, License Count, Installation Date, Software, and more.

Data collected by Discovery on General Software Packages

Label	Table name	Field name	Source
Name	cmdb_ci_spkg	name	Various
Version	cmdb_ci_spkg	version	Various

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery on General Software Packages (continued)

Label	Table name	Field name	Source
Install count	cmdb_ci_spkg	install_count	Various
License count	cmdb_ci_spkg	license_count	Various
Microsoft SMS ID	cmdb_ci_spkg	msft_sms_id	Various
Installed on	cmdb_software_instance	installed_on	Various
Software	cmdb_software_instance	software	Various

GlassFish Server discovery

Discovery creates or updates a CMDB record when it detects a running instance of GlassFish Server.

By default, Discovery uses the following patterns to perform the discovery: *GlassFish Server* and *GlassFish WAR*.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery for GlassFish

The following data is collected in the GlassFish [cmdb_ci_appl_glassfish] table:

Label	Field name
Inst name	inst_name
Port	port
Class	sys_class_name
Version	version
Configuration file	config_file
Domain directory	domain_directory
Install Status	install_status
Installation directory	install_directory
PID	pid
Running process	running_process
Running process command	running_process_command
Running process key parameters	running_process_key_parameters
TCP port(s)	tcp_port
Operational status	operational_status

The following data is collected in the GlassFish War [cmdb_ci_appl_glassfish_war] table:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field name
Name	name
Class	class

HP Operations Manager discovery

Discovery creates or updates a CMDB record when it detects a running instance of HP Operations Manager.

By default, Discovery uses the *HP Operations Manager for Windows* pattern to discover the HP Operations Manager.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery for HP Operations Manager

The following data is collected in the HP Operations Manager `[cmdb_ci_appl_hp_operations]` table:

Label	Field Name
Name	name

HP Service Manager discovery

Discovery creates or updates a CMDB record when it detects a running instance of HP Service Manager.

By default, Discovery uses the *HP Service Manager Application Server* pattern to discover HP Service Manager.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery for HP Service Manager

The following data is collected on the HP Service Manager `[cmdb_ci_appl_hp_service]` table.

Label	Field Name
Name	name
Vendor	vendor

IBM App Connect Enterprise and HTTP listener discovery

Discovery and Service Mapping Patterns application uses the WMB patterns to discover Integration Bus (formerly WebSphere Message Broker and IBM Integration Bus) and HTTP listeners running on both Linux and Windows. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

By default, Discovery uses the following patterns to perform the discovery:

- WMB On UNIX
- WMB On Windows
- WMB HTTP Listener On UNIX
- WMB HTTP Listener On Windows

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

To learn about IBM WebSphere Message Broker (WMB) and WMB HTTP Listener components and versions that you can discover, see Detailed information on products discovered by ITOM Visibility.

Prerequisites

Verify the applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

Data collected by Discovery during horizontal discovery

Resources discovered by the WMB On UNIX pattern

IBM WebSphere Message Broker [cmdb_ci_appl_ibm_wmb]

Field	Description
Name [name]	Name of the message broker
Installation directory [install_directory]	Install directory of the message broker
Version [version]	Version of the message broker

Resources discovered by the WMB On Windows pattern

IBM WebSphere Message Broker [cmdb_ci_appl_ibm_wmb]

Field	Description
Name [name]	Name of the message broker
Installation directory [install_directory]	Install directory of the message broker
Version [version]	Version of the message broker

Resources discovered by the WMB HTTP Listener On UNIX pattern

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM WMB Http Listener [cmdb_ci_appl_ibm_wmb_listener]

Field	Description
Name [name]	Name of the HTTP listener
Installation directory [install_directory]	Install directory of the HTTP listener

Resources discovered by the WMB HTTP Listener On Windows pattern

IBM WMB Http Listener [cmdb_ci_appl_ibm_wmb_listener]

Field	Description
Name [name]	Name of the HTTP listener
Installation directory [install_directory]	Install directory of the HTTP listener

IBM MQ discovery

The ServiceNow Discovery application uses the WMQ On Unix and WMQ On Windows patterns to find IBM MQ (formerly IBM WebSphere MQ) resources. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Discovery uses the following patterns to perform horizontal discovery:

- *WMQ On Unix*
- *WMQ On Windows*

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery for IBM MQ

Discovery populates the data in the CMDB when running the WMQ patterns.

IBM MQ Manager (formerly IBM WebSphere MQ) [cmdb_ci_appl_ibm_wmq]

Label	Field Name
Name	name
Version	version
Installation directory	install_directory
Class	sys_class_name

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft IIS server discovery

Discovery identifies and classifies information about Microsoft IIS servers.

Prerequisites

Install IIS Management Scripts and Tools on a Microsoft IIS Server.

Credentials

Configure these credentials:

- Applicative credentials
- Windows credentials

Classifier, pattern, and probe

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Discovery uses these classifiers, probes, and patterns for IIS discovery:

Classifier	Trigger probes	Pattern
Microsoft IIS Server	• Horizontal discovery probe: launches patterns • Windows - Get IIS Information*	IIS

*Probes inactive on new instances.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected by Discovery for IIS Web Server

The following data is collected in the IIS Web Server [cmdb_ci_microsoft_iis_web_server] table.

Label	Field name
Name	name
Version	version
Type*	type
Running process command	running_process_command
Running process key parameters	running_process_key_parameters

*The Type is always IIS.

The following data is collected in the Web site [cmdb_ci_web_site] table.

Label	Field name	Source
Name	name	wmi

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field name	Source
TCP Port	tcp_port	wmi

Connections found by Service Mapping during top-down discovery

When top-down discovery reaches an IIS server, Service Mapping identifies the specific website and application or virtual directory serving the entry point, not just the IIS server itself. Service Mapping creates a connection from the entry point to the matching IIS Virtual Directory [cmdb_ci_iisdirectory] CI. The resulting service map shows the application running on IIS rather than a generic host-level node.

Microsoft SharePoint discovery

Discovery creates or updates a CMDB record when it detects a running instance of Microsoft SharePoint.

By default, Discovery uses the *SharePoint* pattern to discover Microsoft SharePoint. For information on supported Microsoft SharePoint versions, see Detailed information on products discovered by ITOM Visibility.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery for Microsoft SharePoint

The following data is collected in the SharePoints `[cmdb_ci_appl_sharepoint]` table:

Label	Field Name
Name	name
Class	sys_class_name
Fully qualified domain name	fqdn
IP Address	ip_address
Version	version

NGINX web server discovery

NGINX is an open source web server with a load balancer. Discovery identifies the web server and information related to the load balancer.

Discovery identifies NGINX server software using the following process:

1. The NGINX Process Classifier detects a running process that matches the following criteria during the exploration of a Linux server:
   - Name starts with nginx
   - Name contains master
2. If there is a match:
   - A record is created in the NGINX Web Server [cmdb_ci_nginx_web_server] table.
   - A Runs on relationship is created in the CI Relationship [cmdb_rel_ci] table for a Linux server (Linux Server [cmdb_ci_linux_server]) and for an NGINX web server (NGINX Web Server [cmdb_ci_nginx_web_server]).

     The following two probes are triggered:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- NGINX – Version: This probe contains a Bourne shell script. It determines the version of NGINX and populates the NGINX Web Server [cmdb_ci_nginx_web_server] table.
- NGINX – Get Configuration: This probe contains a Bourne shell script and an argument that determines the path of the NGINX configuration file. The probe identifies configuration parameters based on keywords within the configuration file and returns them as a single payload result.

The sensor on the ServiceNow instance parses the payload result and populates the CMDB.

You must have the following requirements to discover an NGINX webserver:

- NGINX is installed and running on the server.
- The MID Server has access to the NGINX configuration file, which is `/etc/nginx/nginx.conf` by default.

Requirements

For NGINIX servers:

- Ensure that the NGINX software is installed and running on the server.
- Grant the MID Server has access to the NGINX configuration file, which is `/etc/nginx/nginx.conf` by default.
- Enable secure shell (SSH) commands to identify the following associated elements:
  - NGINX Version
  - NGINX Get Configuration

The following probes require execute privileges to run commands:

Commands

Probe	Commands
Nginx – Version	nginx
Nginx – Get Configuration	echo, sed, httpd, cut, grep, egrep (within the Bourne shell script)

Note: The Nginx pattern performs top-down discovery on both Windows and Linux servers. For information on prerequisites for top-down discovery for NGINX web servers, see Prerequisites for performing top-down discovery using Service Mapping.

Probes and Sensors

Discovery identifies NGINX server software using the following process:

1. The Nginx Process Classifier detects a running process that matches the following criteria during the exploration of a UNIX server:
   - The name of the process starts with `nginx`.
   - The name of the process contains `master`.
2. If there is a match:
   - A record is created in the Web Server [cmdb_ci_web_server] table.
   - A Runs on relationship is created in the CI Relationship [cmdb_rel_ci] table for the Linux Server [cmdb_ci_linux_server] table and the Web Server [cmdb_ci_web_server] table.
3. The following two probes are triggered:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Nginx – Version**: this probe contains a Bourne shell script. It determines the version of NGINX and populates the Web Server [cmdb_ci_web_server] table.
- **Nginx – Get Configuration**: this probe contains a Bourne shell script and an argument that determines the path of the NGINX configuration file. The probe identifies configuration parameters based on keywords within the configuration file and returns them as a single payload result.

4. The sensor on the ServiceNow instance processes the payload and populates the CMDB.

Data Collected

Discovery creates or updates CMDB records when it detects a running NGINX process. The following data is collected.

Data Collected by Discovery for NGINX

Table Name	Field and Label Name	Probe
Name	cmdb_ci_lb_appl [name]	Nginx– Get Configuration
IP Address	cmdb_ci_lb_appl [ip_address]	Nginx– Get Configuration
Last Discovered	cmdb_ci_lb_appl [last_discovered]	Nginx– Get Configuration
Version	cmdb_ci_lb_appl [version]	Nginx– Version
Name	cmdb_ci_lb_pool_member [name]	Nginx– Get Configuration
Last Discovered	cmdb_ci_lb_pool_member [last_discovered]	Nginx– Get Configuration
IP Address	cmdb_ci_lb_pool_member [ip_address]	Nginx– Get Configuration
Load Balancer	cmdb_ci_lb_pool_member [load_balancer]	Nginx– Get Configuration
Fully Qualified Domain Name	cmdb_ci_lb_pool_member [fqdn]	Nginx– Get Configuration
Port	cmdb_ci_lb_pool_member [port]	Nginx– Get Configuration
Last Discovered	cmdb_ci_lb_service [last_discovered]	Nginx– Get Configuration
IP Address	cmdb_ci_lb_service [ip_address]	Nginx– Get Configuration
Name	cmdb_ci_lb_service [name]	Nginx– Get Configuration
Load Balancer	cmdb_ci_lb_service [load_balancer]	Nginx– Get Configuration
Input URL	cmdb_ci_lb_service [input_url]	NGINX – Get Configuration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Relationships

In the table [cmdb_rel_ci] the following relationships are populated:

- The records in the cmdb_ci_lb_appl table run on the cmdb_ci_web_server table records.
- The records in the cmdb_ci_lb_service table use the cmdb_ci_lb_pool table records.
- The records in the cmdb_ci_pool table are used by the cmdb_ci_service table records.
- The records in the cmdb_ci_pool table are members of the cmdb_ci_pool_member table.
- The records in the cmdb_ci_pool_member table are members of cmdb_ci_pool table.

Oracle Tuxedo discovery

The ServiceNow Discovery application uses the Tuxedo pattern to find Oracle Tuxedo resources. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Discovery uses the **Tuxedo** pattern to run horizontal discovery.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery for Oracle Tuxedo

Discovery populates the data in the CMDB when running the Tuxedo pattern.

Tuxedo [cmdb_ci_appl_tuxedo]

Name	Description
Port [port]	The network communication port used by Oracle Tuxedo services and processes for network communication.
Domain [domain]	The logical grouping of Oracle Tuxedo components and services operating within a defined environment.
Installation directory [install_directory]	The location that contains all binaries and configuration files necessary for running Oracle Tuxedo.
Log dir [log_dir]	The directory that stores all log files generated by Oracle Tuxedo for monitoring and troubleshooting.
Version [version]	The specific release of Oracle Tuxedo that is installed and currently in use.
Mode [mode]	The definition of the interaction protocol between Oracle Tuxedo and various clients or other domains, such as Jolt, Workstation, or Domain Gateway, based on the operational configuration.

Puppet automation software discovery

Discovery identifies Puppet Masters running on UNIX servers.

Discovery uses secure shell (SSH) commands to collect information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To access Puppet Master records, navigate to **All** > **Configuration** > **Automation Servers** > **Puppet Masters**. You can view the resources in the related list of the Puppet Master record.

How the Puppet Master probe works

Discovery identifies the Puppet Master using the following process:

1. The **UNIX - Active Processes** probe detects a running process that matches one of the following criteria:
   - The name of the process is `pe-httpd`.
   - The name of the process is `ruby` and the parameters of the process contain *`puppet master`*.
2. If there is a match on one of these criteria:
   - A record is created in the Puppet Master table [cmdb_ci_puppet_master] if one does not already exist for that running process.
   - The **Puppet - Master Info** probe is triggered. The sensor of this probe populates additional information in the Puppet Master [cmdb_ci_puppet_master] record.

Requirements to discover a Puppet master

By default, Discovery identifies Puppet Masters running on UNIX servers. Discovery uses secure shell (SSH) commands to collect information.

The credentials used to discover the UNIX server must have privileges to execute the following commands. The use of sudo is supported, but you must add the *`must_sudo`* parameter to the probe.

Probe	Commands
Puppet – Master Info	puppet, echo, hostname (within the Borne shell script)
Puppet – Certificate Requests	puppet
Puppet – Manifests	echo, sed, find (within the Bourne shell script)
Puppet – Modules	puppet

Data collected

Data collected by Discovery for Puppet automation software, by default

Name	Puppet Master [cmdb_ci_puppet_master]	name	Puppet - Master Info
Configuration directory	Puppet Master [cmdb_ci_puppet_master]	config_directory	Puppet - Master Info
Manifest directory	Puppet Master [cmdb_ci_puppet_master]	manifestdir	Puppet - Master Info
Module path	Puppet Master [cmdb_ci_puppet_master]	modulepath	Puppet - Master Info
Fully qualified domain name	Puppet Master [cmdb_ci_puppet_master]	fqdn	Puppet - Master Info

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery for Puppet automation software, by default (continued)

IP Address	Puppet Master [cmdb_ci_puppet_master]	ip_address	Puppet - Master Info
Version	Puppet Master [cmdb_ci_puppet_master]	version	Puppet - Master Info

Add the must_sudo parameter to the Puppet probe

To use the Puppet probe with sudo, you must add the *must_sudo* parameter to each probe that requires it.

Before you begin
Role required: admin

About this task
You add the must_sudo parameter for each Puppet probe.

Procedure

1. Navigate to **All** > **Discovery** > **Probes**.
2. In the **Search** field, search on the name **Puppet**.
3. Click the name of the probe.
   For example, **Puppet Master - Info**.
4. In the **Related Links** pane, select the **Probe parameters** tab.
5. Click **New**.
6. In the **Name** field, enter `must_sudo`
7. In the **Value** field, enter `true`.
8. Click **Submit**.

SAP products discovery

ITOM Visibility can discover SAP applications, SAP HANA Database, and SAP HANA DB Catalog. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

To learn about the SAP components and their versions that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

Prerequisites

Verify that the applications are up to date

Discovery and Service Mapping Patterns

Visibility Content

CMDB CI Class Models

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify the configuration of the credentials

- (If hosted on a UNIX server) On the ServiceNow AI Platform, configure applicative credentials and SSH credentials. The SSH credentials are used to create connections from SAP Central Instance (CI) to other applications and devices. In addition, provide the user elevated rights to be able to run commands requiring a privileged user: Service Mapping commands requiring a privileged user.
- (If hosted on a Windows server) On the ServiceNow AI Platform, configure Windows credentials to create connections from SAP CI to other applications and devices.
- (Optional) For discovering using the WinRM protocol, see Configure WinRM trusted hosts on MID Server.
- (Optional for top-down discovery by Service Mapping) Provide a user with permissions to run SOAP on RFC read table function.
- (Optional for top-down discovery by Service Mapping) If your organization deploys SAP Digital Business Services using web services and you want to map all connected services, perform the following actions:
  - Configure applicative credentials for a user with permissions to access the following tables: CVERS, DBCONS, and RFCDES.
  - Ensure that you have the web service URL. For information on how to create web services, see https://blogs.sap.com/2012/10/24/how-to-create-web-services-abap-based/.

SAP HANA 2.0 DB Catalog on UNIX

Verify that the Applicative Credentials for **SAP HANA DB App Creds** are configured with **SAP Hana Db [cmdb_ci_appl_sap_hana_db]** as CI type.

Verify the permissions to run Service Mapping commands

Service Mapping commands not requiring a privileged user

Verify the user has the permissions to run the Patterns queries

The following commands are executed when running the SAP HANA DB and SAP HANA DB Catalog Patterns.

- `select * from schemas;`
- `select * from M_HOST_INFORMATION where upper(KEY) like '%TIMEZONE%';`
- `select SCHEMA_NAME, round(sum(TABLE_SIZE) /1024/1024) from M_TABLES GROUP BY SCHEMA_NAME`

The query type	Fallback commands executed in order
Schemas Info	1. `$hdbsql_path + " -n " + $entry_point.ip_address + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select * from schemas;\""`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The query type	Fallback commands executed in order
	2. `$hdbsql_path + " -n " + $ip_address + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select * from schemas;\""` 3. `$hdbsql_path + " -n " + $server_name + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select * from schemas;\""` 4. `$hdbsql_path + " -n localhost -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select * from schemas;\""`
Time Zone:	1. `$hdbsql_path + " -n " + $entry_point.ip_address + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select * from M_HOST_INFORMATION where upper(KEY) like '%TIMEZONE%';\""` 2. `$hdbsql_path + " -n " + $ip_address + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select * from M_HOST_INFORMATION where upper(KEY) like '%TIMEZONE%';\""` 3. `$hdbsql_path + " -n " + $server_name + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select * from M_HOST_INFORMATION where upper(KEY) like '%TIMEZONE%';\""` 4. `$hdbsql_path + " -n localhost -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The query type	Fallback commands executed in order
	* from M_HOST_INFORMATION where upper(KEY) like '%TIMEZONE%';\""
Disk Size	1. $hdbsql_path + " -n " + $entry_point.ip_address + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select SCHEMA_NAME, round(sum(TABLE_SIZE) /1024/1024) from M_TABLES GROUP BY SCHEMA_NAME;\"" 2. $hdbsql_path + " -n " + $ip_address + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select SCHEMA_NAME, round(sum(TABLE_SIZE) /1024/1024) from M_TABLES GROUP BY SCHEMA_NAME;\"" 3. $hdbsql_path + " -n " + $server_name + ":" + $query_port + " -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select SCHEMA_NAME, round(sum(TABLE_SIZE) /1024/1024) from M_TABLES GROUP BY SCHEMA_NAME;\"" 4. $hdbsql_path + " -n localhost -i " + $digits + " -u '$$username$$' -p '$$password$$' \"select SCHEMA_NAME, round(sum(TABLE_SIZE) /1024/1024) from M_TABLES GROUP BY SCHEMA_NAME;\""

Verify the configuration of the Discovery schedule

For more information, see Running discoveries in your network

Data collected by Discovery during horizontal discovery

Table and field	Description
SAP ASCS [cmdb_ci_appl_sap_ascs]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Name [name]	Data collected for SAP ABAP Central Services (ASCS)
Version [version]	
Installation directory [install_directory]	
Configuration directory [config_directory]	
Configuration file [confi_file]	
Instance name [instance_name]	
Instance number [instance_number]	
System ID [sid]	
System directory [system_directory]	
System type [system_type]	
Transport domain [transport_domain]	
SAP Business Objects CMS server [cmdb_ci_appl_sap_bo]	
Name [name]	Data collected for SAP Business Objects CMS Server
Server Name [server_name]	
Version [version]	
Installation directory [install_directory]	
SAP SCS Application [cmdb_ci_appl_sap_scs]	
Name [name]	Data collected for SAP Central Services (SCS)
Version [version]	
Installation directory [install_directory]	
Configuration directory [config_directory]	
Configuration file [config_file]	
Instance name [instance_name]	
Instance number [instance_number]	
System ID [sid]	
System directory [system_directory]	
System type [system_type]	
Transport domain [transport_domain]	
SAP CI Application [cmdb_ci_appl_sap_ci]	
Name [name]	Data collected for SAP CI
Version [version]	
Installation directory [install_directory]	
Configuration directory [config_directory]	
Configuration file [config_file]	
Instance name [instance_name]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Instance number [instance_number]	
System ID [sid]	
System directory [system_directory]	
System type [system_type]	
Transport domain [transport_domain]	
SAP DI Applications [cmdb_ci_appl_sap_di]	
Name [name]	Data collected for SAP Dialog Instance (DI)
Version [version]	
Installation directory [install_directory]	
Configuration directory [config_directory]	
Configuration file [config_file]	
Instance name [instance_name]	
Instance number [instance_number]	
System ID [sid]	
System directory [system_directory]	
System type [system_type]	
Transport domain [transport_domain]	
SAP ERS Application [cmdb_ci_appl_sap_ers]	
Name [name]	Data collected for SAP Enqueue Replication Server (ERS)
Version [version]	
Installation directory [install_directory]	
Configuration directory [config_directory]	
Configuration file [config_file]	
Instance name [instance_name]	
Instance number [instance_number]	
System ID [sid]	
System directory [system_directory]	
System type [system_type]	
Transport domain [transport_domain]	
HANA Database [cmdb_ci_appl_sap_hana_db]	
Name [name]	Data collected for SAP HANA Database
IP Address [id_address]	
Class [sys_class_name]	
Fully qualified domain name [fqdn]	
Version [version]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Installation directory [install_directory]	
SAP SID [cmdb_ci_sap_sid]	
Name [name]	Data collected for the SAP System Identification (SID) code is a unique identification code for an SAP R/3 installation. An SID is a three-character code like D10, PRD, or E56. An R/3 installation consists of a database server plus several application servers. ITOM Visibility uses the SAP SID shared library pattern to discover SIDs for SAP applications.
Version [version]	
Patch Level Info [sap_patch_info]	

Data collected by the SAP HANA DB Catalog on UNIX pattern during horizontal discovery

SAP HANA Catalog [cmdb_ci_db_hana_catalog]

Field	Description
Catalog Owner [catalog_owner]	The user ID of the schema owner.
Catalog Create Time [catalog_create_time]	The catalog creation time.
Disk Allocated Size (MB) [disk_allocated_size_mb]	The allocated Disk size in MB.
Database Reference [database_instance]	The reference to the SAP HANA DB.

CI Relationships

The SAP SID shared library pattern creates CI relationships between all major CI applications and SAP System ID.

CI	Relationship	CI
SAP ERS Application [cmdb_ci_appl_sap_ers]	Uses::Used by	SAP SID [cmdb_ci_sap_sid]
SAP DI Applications [cmdb_ci_appl_sap_di]	Uses::Used by	SAP SID [cmdb_ci_sap_sid]
SAP ASCS [cmdb_ci_appl_sap_ascs]	Uses::Used by	SAP SID [cmdb_ci_sap_sid]
SAP CI Application [cmdb_ci_appl_sap_ci]	Uses::Used by	SAP SID [cmdb_ci_sap_sid]
SAP SCS Application [cmdb_ci_appl_sap_scs]	Uses::Used by	SAP SID [cmdb_ci_sap_sid]
SAP JC Application [cmdb_ci_appl_sap_jc]	Uses::Used by	SAP SID [cmdb_ci_sap_sid]

SAP HANA DB Catalog on UNIX CI relationships

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
SAP HANA DB [cmdb_ci_appl_sap_hana_db]	Contains::Contained by	SAP HANA DB Catalog [cmdb_ci_db_hana_catalog]

Limitations

Service Mapping can map only one instance of ASCS per deployment.

The SAP HANA DB Catalog pattern can only run on UNIX.

Map application services containing SAP applications

Map a service instance containing an SAP application and verify that Service Mapping discovered all SAP components correctly.

Before you begin

Role required: service_mapping_admin

Procedure

1. Create a service instance.
   a. Navigate to **All** > **Service Mapping** > **Service Instances**.
   b. Select **New**.
   The **Discoverable by Service Mapping** option is selected by default.
   c. In the **Name** field, enter a descriptive name for the service instance.
2. Define the entry point for this application service:
   a. Select the **SAP application** tile.
   The **Discoverable by Service Mapping** option is selected by default.
   b. Define the entry point as follows:
      - A URL to any SAP module
      - A TCP connection to any SAP module
      - To discover environments with connected services, provide the URL to the web service in the following format:

```
 http://<hostname>:<port>/sap/bc/srt/scs/sap/z_servicenow_read?sap-client=003
```

   c. Select **Add**.
   d. Select **Save**.
3. Select **Additional Info** on the left and clear the **Traffic based discovery** check box.
4. Select **Update**.
5. Verify that the mapping result is satisfactory:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select **View Map** and wait until the mapping is complete.
It may take a few minutes.

b. Review the application service map and make sure that the service layout is as in the following examples.

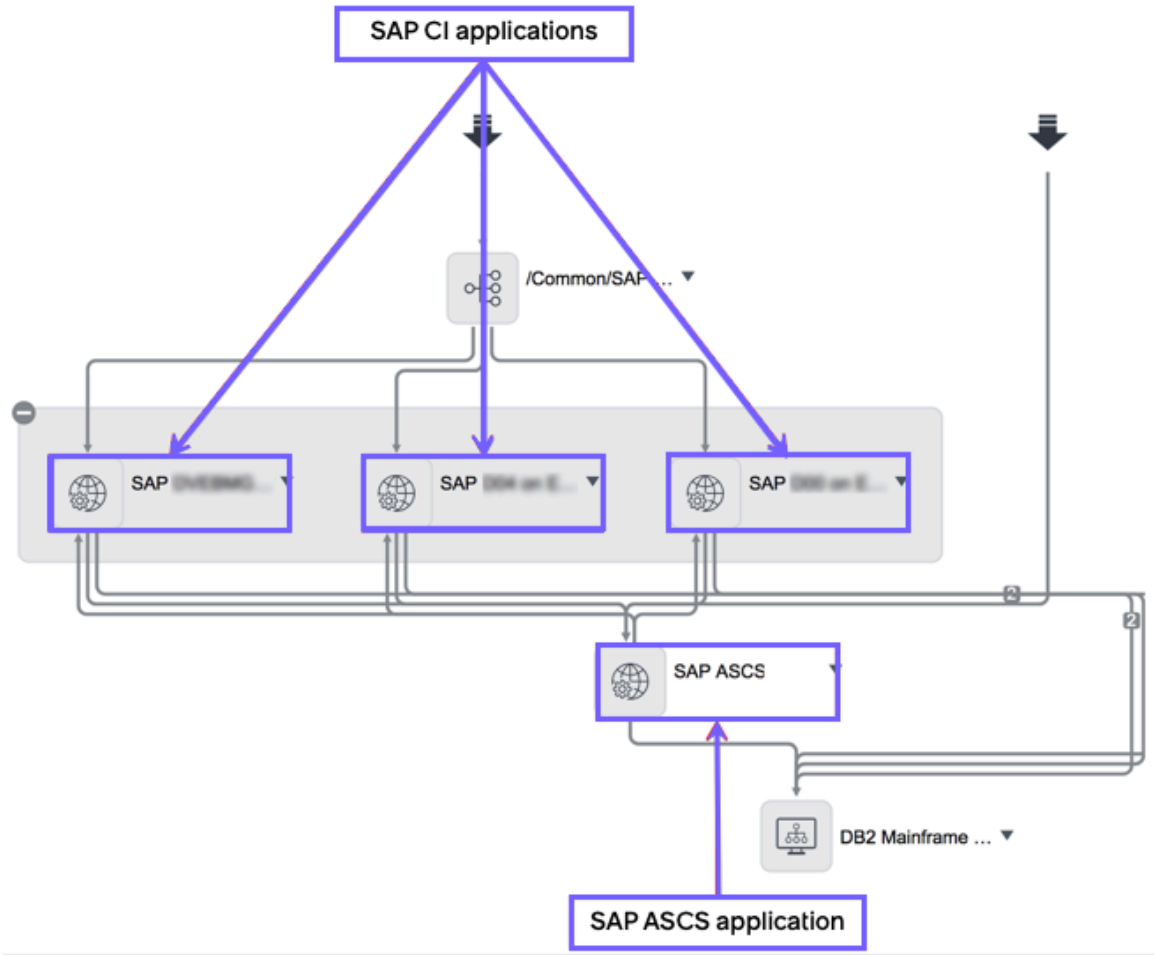

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the SAP environment integrates SAP components with Tibco services, you can expect this mapping result:

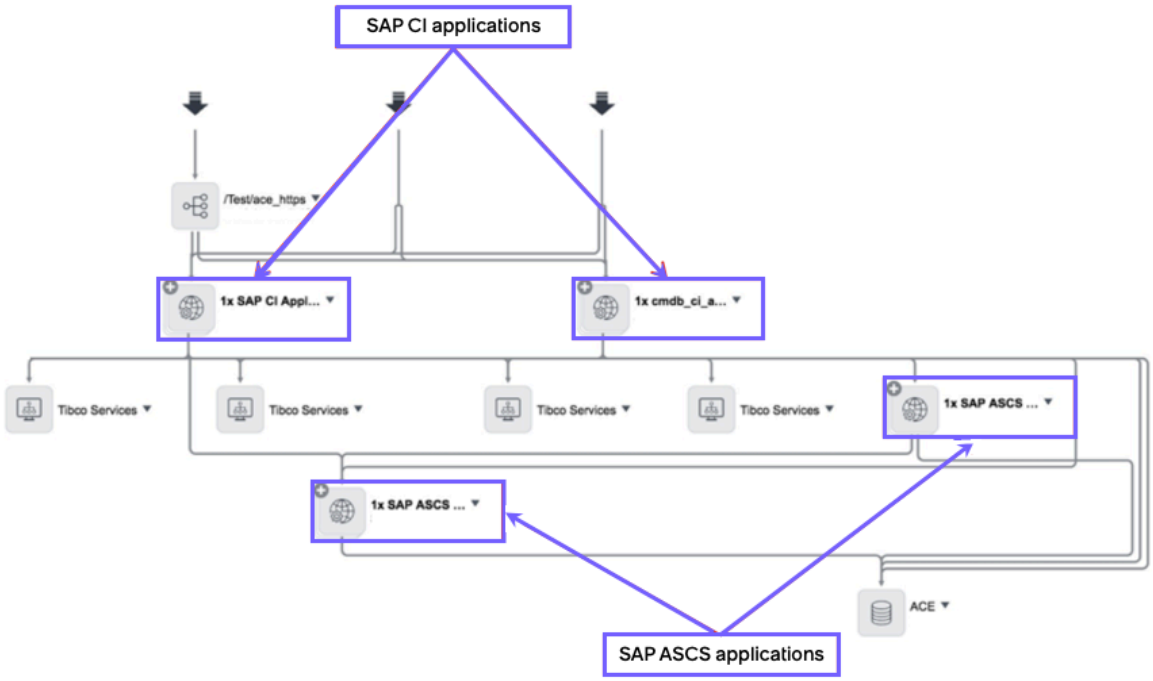

Troubleshoot application services containing SAP applications

Fix basic discovery issues based on the symptoms and discovery messages.

Symptom

The application service map in Service Mapping displays the following message `Failed to discover <SAP component name>.`

Cause

The CI failed to identify, and process detection retrieves the process attributes.

Solution

Perform the following steps:

1. Open the failed pattern step in Debug mode:
   a. On the business service or service instance map, right-click the error message and select **Show discovery log**.

      The Discovery Log window opens.

   b. In the right pane, locate the failed identification section and the failed step within this section.
   c. Select this failed identification section and click **Debug**.

      The Debug Identification Section window opens.

   d. If necessary, fill in the required details and click **Connect**.
2. Click the **set variables** step under **Steps**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Make sure that the regular expression includes SYS as shown in the following figure.

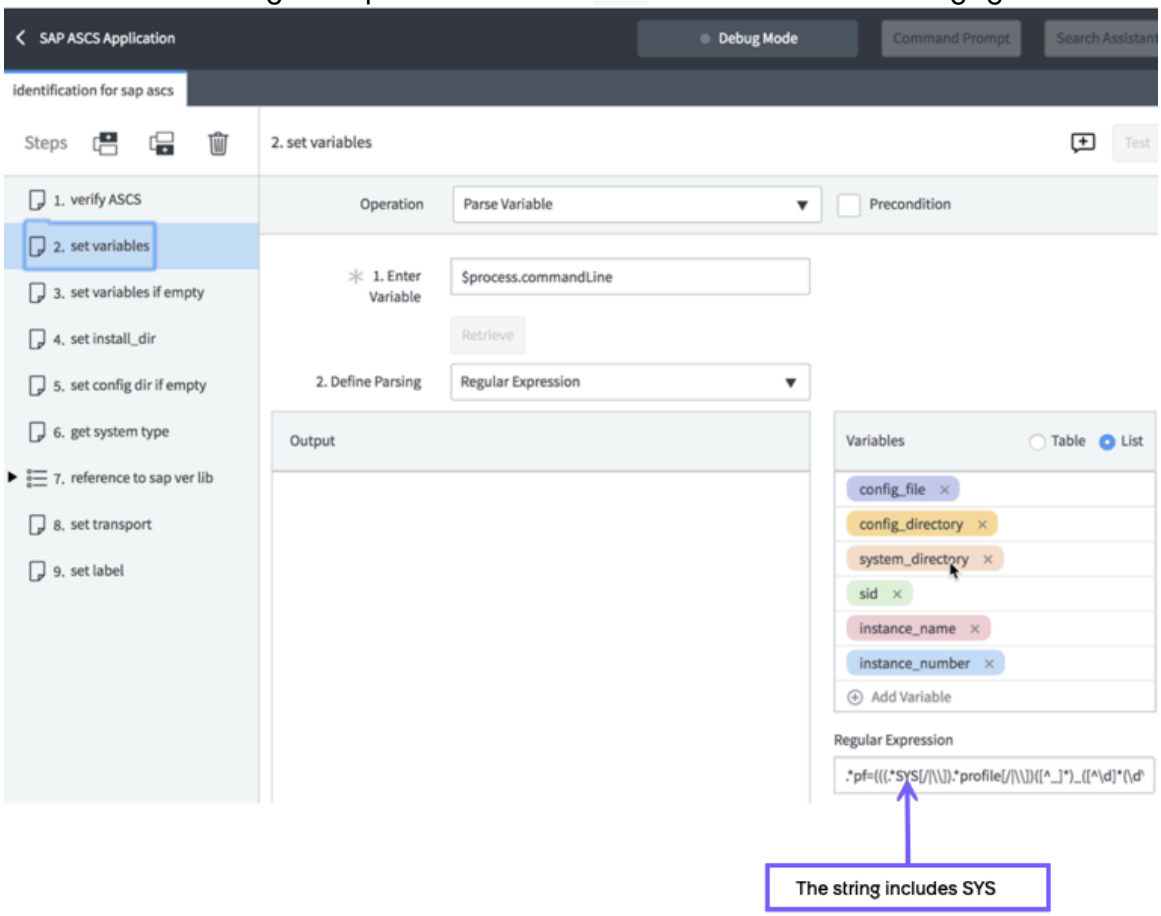

4. Click the **set install_dir** step under **Steps**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Make sure that the regular expression includes `sap` as shown in the following graphic.

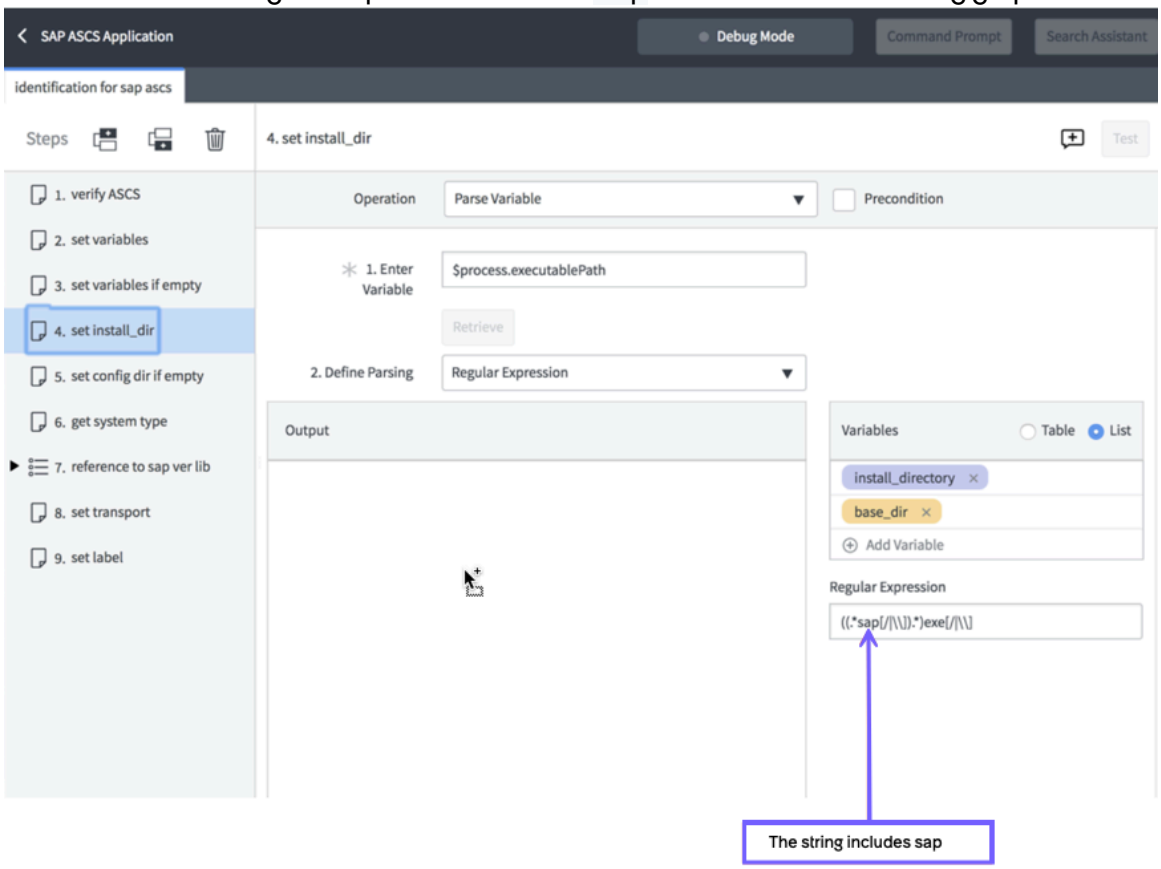

Symptom

Service Mapping fails to discover the web service your organization uses to deploy SAP digital business services.

Cause 1

The pattern fails at the EVAL function for identifying the web service.

Solution 1

1. Verify that the format of the entry point is the URL to the web service as described in the mapping procedure.
2. Verify that you configured the necessary applicative credentials as described in the mapping procedure.
3. Check that the discovery log contains the following message: GROOVE_CODE_FAILURE: No such property: com for class: Script1.
4. Open the failed pattern step in Debug mode:
   a. On the business service or service instance map, right-click the error message and select **Show discovery log**.

      The Discovery Log window opens.

   b. In the right pane, locate the failed identification section and the failed step within this section.
   c. Select this failed identification section and click **Debug**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Debug Identification Section window opens.

d. If necessary, fill in the required details and click **Connect**.

5. Pattern Designer opens showing the selected identification section.
6. Select the failed step on the right.
7. Verify that the function running on EVAL is `com.snc.sw.util.HttpInvokerUtil.post`.

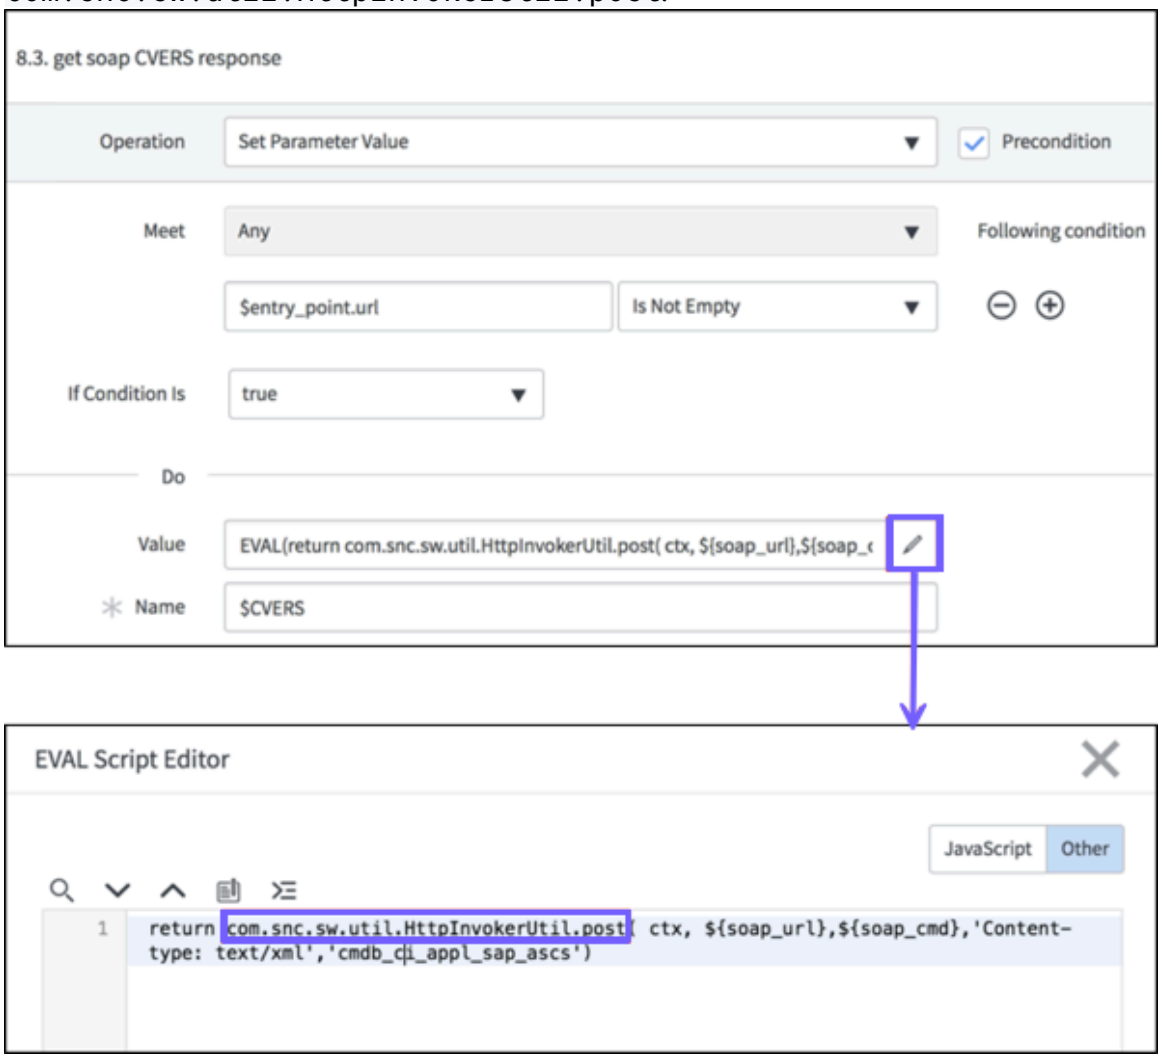

8. Try to map the application service again and check the result.

Cause 2

The pattern fails at the EVAL function for identifying the web service, because the web service was configured incorrectly.

Solution 2

1. Verify that the format of the entry point is the URL to the web service as described in the mapping procedure.
2. Verify that you configured the necessary applicative credentials as described in the mapping procedure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Check that the discovery log contains the following message: "Debug run railed: Test step failed. GROOVE_CODE_FAILURE: All command implementations (HTTP) failed on host 10.x.x.x. Host must be down or unreachable. ()"
4. Recreate the web service as described in https://blogs.sap.com/2012/10/24/how-to-create-web-services-abap-based/ .
5. Try to map the application service again and check the result.

Cause 3

The pattern fails at the EVAL function for identifying the web service. This occurs because the URL configured in the entry point is not correct or because the MID Server cannot access this URL.

Solution 3

1. Verify that the format of the entry point is the URL to the web service as described in the mapping procedure.
2. Verify that you configured the necessary applicative credentials as described in the mapping procedure.
3. Check that the discovery log contains the following message: "Debug run railed: Test step failed. GROOVE_CODE_FAILURE: All command implementations (HTTP) failed on host 10.x.x.x. Host must be down or unreachable. ()"
4. Verify that you used the correct URL for the service entry point:
   a. In the TCode SOAMANAGER, navigate to **Simplified Web Service Configuration**.
   b. On the **Configuration** tab of the Simplified Service Configuration window, search for the web service that is part of the service instance.
   c. Click the check box under **User Name/Password (Basic)** for this web service.
   d. Click **Show Details**.
   e. On the Details of the Service Definition window, verify the actual URL of the web service under WSDL URLs.
   f. If necessary, fix the URL you configure for the entry point and try to map the service instance again.

Cause 4

MID Server cannot access this URL configured in the entry point of the service instance.

Solution 4

1. Verify that the format of the entry point is the URL to the web service as described in the mapping procedure.
2. Verify that you configured the necessary applicative credentials as described in the mapping procedure.
3. Check that the discovery log contains the following message: "Debug run failed: Test step failed. GROOVE_CODE_FAILURE: All command implementations (HTTP) failed on host 10.x.x.x. Host must be down or unreachable. ()"
4. Verify that you used the MID Server that can access the URL you configured in the service entry point:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Log into the MID Server host remotely.

b. Open a browser.

c. Paste the URL into the browser address bar and press Enter.

d. If the page does not open, it means that there is a network issue, for example, there is a firewall that prevents the MID Server from accessing this URL.

5. Resolve the network issue and verify that the MID Server can access the URL as described in step 4.

Cause 5

MID Server cannot access this URL configured in the entry point of the service instance.

Solution 5

1. Verify that the format of the entry point is the URL to the web service as described in the mapping procedure.
2. Verify that you configured the necessary applicative credentials as described in the mapping procedure.
3. Check that the discovery log contains the following message: "Debug run railed: Test step failed. GROOVE_CODE_FAILURE: All command implementations (HTTP) failed on host 10.x.x.x. Host must be down or unreachable. ()"
4. Verify that the applicative credentials you configured on the ServiceNow AI Platform for the web service are valid:

   a. Log into the MID Server host remotely.

   b. Open a browser.

   c. Paste the URL into the browser address bar and press Enter.

   d. The browser page displays the credentials pop-up window.

   e. Enter the credentials and check that the page opens.

   The page does not open, it means that the configured applicative credentials are incorrect.

5. Find out the right credentials.
6. Verify the new credentials as described in step 4.
7. If the page opens after you use the new credentials, configure these credentials as applicative credentials on the ServiceNow AI Platform.

Cause 6

MID Server cannot access this URL configured in the entry point of the service instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solution 6

1. Verify that the format of the entry point is the URL to the web service as described in the mapping procedure.
2. Check that the discovery log contains the following message: "Debug run railed: Test step failed. GROOVE_CODE_FAILURE: Missing applicative credentials for CI type <name of the CI type>".
3. Verify that you defined applicative credentials for the relevant CI according as described in Map application services containing SAP applications.

SQL Server Analysis Services (SSAS) discovery

The Discovery and Service Mapping application uses the SSAS pattern to find SQL Server Analysis Services (SSAS) on your infrastructure. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store..

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the applications are up to date.

Discovery and Service Mapping Patterns

CMDB CI Class Models

Visibility Content

Verify the configuration of credentials

On the ServiceNow AI Platform®, verify the configuration of the Windows credentials. For more information, see Windows credentials

Data collected during horizontal discovery

Discovery populates the following data when running the SSAS pattern.

SQL Server Analysis Services [cmdb_ci_db_mssql_analysis]		
Label	**Field**	**Description**
Name	name	The SSAS on the Windows sever name
IP Address	ip_address	The SSAS IP address
Version	version	The SSAS version
Edition	edition	The SSAS edition
Installation directory	install_directory	The SSAS install directory
Instance Name	instance	The SSAS instance name
TCP port(s)	tcp_port	The SSAS TCP port

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Microsoft SQL Server Integration Services (SSIS) discovery

The Discovery and Service Mapping application uses the SSIS pattern to find SQL Server Integration Services (SSIS) on your infrastructure. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the plugins setup

Discovery and Service Mapping Patterns

CMDB CI Class Models

Visibility Content

Verify the configuration of credentials

On the ServiceNow AI Platform®, verify the configuration of the Windows credentials. For more information, see Windows credentials

Data collected during horizontal discovery

Discovery populates the following data when running the SSIS pattern.

SQL Server Integration Services [cmdb_ci_db_mssql_integration_list]		
Label	**Field**	**Description**
Name	name	The SSIS on the Windows sever name
IP Address	ip_address	The SSIS IP address
Version	version	The SSIS version
Edition	edition	The SSIS edition
Installation directory	install_directory	The SSIS install directory
Instance Name	instance	The SSIS instance name

Sun Java Enterprise System (JES) discovery

Discovery creates or updates a CMDB record when it detects a running instance of Sun Java Enterprise System (JES).

By default, Discovery uses the *Sun JES pattern* pattern to discover Sun Java Enterprise System.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Data collected by Discovery for Sun Java Enterprise System

The following data is collected in the JES [cmdb_ci_email_server_jes] table:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field Name
Name	name
IP Address	ip_address
Class	sys_class_name
Fully qualified domain name	fqdn

Tomcat server discovery

Discovery identifies and classifies information about Tomcat server Web applications running on UNIX machines.

Discovery can identify and classify Web applications present in either the `CATALINA_BASE` directory or the `CATALINA_HOME` directory.

The classifier that finds Tomcat server processes uses the condition: [Parameters] [contains] [org.apache.catalina.startup.Bootstrap].

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Credentials

Configure these credentials:

- Applicative credentials
- SSH credentials

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Classifier, pattern, and probes

Classifier	Trigger probes	Pattern
Tomcat	• Horizontal pattern: launches patterns • SSHCommand-Tomcat - Find web.xml List* • SSHCommand-Tomcat - Get server.xml* • SSHCommand-Tomcat - Get Version*	Tomcat

*Probes inactive on new instances.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected

Discovery collects the following data on the Tomcat [cmdb_ci_app_server_tomcat] table for both Unix and Windows:

Label	Field name	Source
Name	name	server.info
Install Directory	install_directory	server.xml
Version	version	server.info
Server Port	server_port	server.xml
Attributes	attributes	server.xml

The following data is collected on the Tomcat Connector [cmdb_ci_tomcat_connector] table for both Unix and Windows:

Label	Field name	Source
Tomcat	tomcat	server.xml
Port	port	server.xml

The following data is collected on the Web Application [cmdb_ci_web_application] table for both Unix and Windows:

Label	Field name	Source
Description	description	web.xml
Document base	document_base	web.xml
App server	app_server	web.xml

Relationships

Parent class	Relationship	Child class
cmdb_ci_app_server_tomcat	Contains::Contained by	cmdb_ci_web_application

WebLogic application server discovery

Discovery creates or updates a CMDB record when it detects an instance of an Oracle or BEA Weblogic application server running on a Windows or Linux system.

Requirements

For Windows

- PowerShell must be enabled on the MID Server.
- The WebLogic Administration Server instances that start via WebLogic NodeManager must have the `-Dweblogic.RootDirectory=<path>` parameter defined upon server startup. The Windows credential must also have read permissions on the `config.xml` file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For Linux

- Discovery requires the use of these Bourne shell commands:
  - `find`
  - `cat`
  - `dirname`
- WebLogic administration server instances started via NodeManager must have the *`-Dweblogic.RootDirectory=<path>`* parameter defined and visible through the Linux *`ps process stat`* command (for each AdminServer) for the rest of the Linux WebLogic application server and web application information to be populated in the CMDB.
- The following configuration must be set on the sudorers file:

  `User ALL=(root) NOPASSWD: /bin/sh *`

  `User` is the user name on the SSH credential that you create for this discovery.
- This command must be able to be run:

  `/bin/sh /tmp/snc-*-findcat.sh * *"`

  Do not make any changes to `findcat.sh`.
- If you are using probes, not patterns, for Weblogic discovery, the *`Linux - Weblogic - Find config.xml`* probe must specify this parameter and value: *`must_sudo`* = *`true`*

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Credentials

- Applicative credentials
- Windows credentials
- SSH credentials

Classifier, pattern, and probes

Classifier	Trigger probes	Pattern
Weblogic Server	• Horizontal Pattern: launches patterns • Windows - Weblogic - Get web.xml* • Windows - Weblogic - Find config.xml* • Linux - Weblogic - Find config.xml*	WebLogic

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

*For new instances, these probes are inactive on the classifier. Discovery uses patterns for discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected on Windows

Label	Table Name	Field Name	Source
Name	cmdb_ci_app_server_weblogic	name	running process
Version	cmdb_ci_app_server_weblogic	version	config.xml
Weblogic domain	cmdb_ci_app_server_weblogic	weblogic_domain	config.xml
TCP port(s)	cmdb_ci_app_server_weblogic	tcp_port	web.xml and running process
Name	cmdb_ci_web_application	name	config.xml
Document base	cmdb_ci_web_application	document_base	config.xml
Description	cmdb_ci_web_application	description	web.xml
Servlet class	cmdb_ci_web_application	servlet_class	web.xml
Servlet name	cmdb_ci_web_application	servlet_name	web.xml
App server	cmdb_ci_web_application	app_server	config.xml
TCP port(s)	cmdb_ci_web_application	tcp_port	web.xml

Data collected on Linux

Table	Field	Source
cmdb_ci_app_server_weblogic	Name (name)	running process
cmdb_ci_app_server_weblogic	TCP port (tcp_port)	running process
cmdb_ci_app_server_weblogic	Version (version)	config.xml
cmdb_ci_app_server_weblogic	Weblogic domain (weblogic_domain)	config.xml
cmdb_ci_web_application	Name (name)	config.xml
cmdb_ci_web_application	Context path (context_path)	config.xml
cmdb_ci_web_application	App server (app_server)	config.xml

Relationship

Parent class	Relationship type	Child class
cmdb_ci_app_server_weblogic	Runs on::Runs	cmdb_ci_windows_server or cmdb_ci_linux_server
cmdb_ci_app_server_weblogic	Contains::Contained by	cmdb_ci_web_application

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parent class	Relationship type	Child class
cmdb_ci_app_server_weblogic	Depends on::Used by	cmdb_ci_app_server_weblogic This relationship is made between an AdminServer and any managed servers it encapsulates.

Example Weblogic relationships

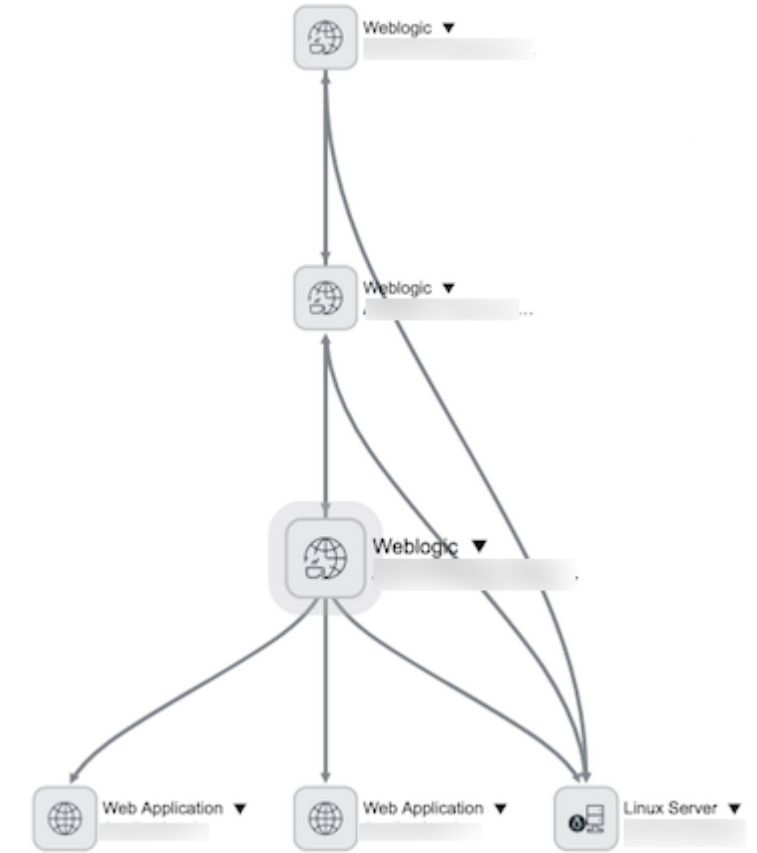

Add sudo access for the Weblogic - Find config.xml probe

If you are not using the Weblogic pattern, you can add sudo privileges to the Linux, the Weblogic - Find config.xml probe.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Discovery Definition** > **Probes**.
2. Search for and select the **Linux - Weblogic - Find config.xml** probe.
3. In the **Probe Parameters** related list, click**New**
4. Use the following information to complete the form

Probe	Commands
Name	must_sudo

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	Commands
Value	true

5. Click **Submit**

Identifying Windows WebLogic application servers

If you are not using the Weblogic pattern, Discovery follows this process when it uses the Windows - Active Processes or Linux - Active Processes probes.

This process describes the use of probes, which are active on your instance if you upgrade. For new instances, Discovery uses the Weblogic pattern.

1. The Windows - Active Processes probe or Linux - Active Processes detects a running process that matches one of the follow criteria:
   - The parameters of the process contain *`weblogic.Server`*.
   - The parameters of the process contain *`-Dweblogic.name`*.
2. If there is a match on one of the criteria:
   - A record is created in the BEA Weblogic [cmdb_ci_app_server_weblogic] table. The record is populated with the server name and TCP port, which is gathered from the running process.
   - The Windows - Weblogic - Find config.xml probe or the Linux - Weblogic - Find config.xml probe triggers. The sensor of this probe populates additional information in the BEA Weblogic [cmdb_ci_app_server_weblogic] record and the Web Application [cmdb_ci_web_application] record if applicable.
3. The Windows - Weblogic - Find config.xml probe or the Linux - Weblogic - Find config.xml probe attempts to find the related `config.xml` file for the server by:
   - Using the *`-Dweblogic.RootDirectory=<path>`* parameter defined in the running process.
   - [Windows] Searching for `config.xml` files under the *`-Dplatform.home=<path>`* parameter defined in the running process (not as efficient using the parameters of the process).
   - [Linux] Searching for the parent process that started the WebLogic server. This is viable only if the weblogic jvm was started via the *`startWeblogic.sh`* or related custom script and not the init process.
4. If there are associated web applications found in the WebLogic `config.xml` file, the Windows – Weblogic find web.xml probe triggers for each application. This probe reads the WebLogic `web.xml` file for each web application and the sensor, and then populates additional information.

IBM WebSphere server discovery using probes

The IBM WebSphere application server is a software framework with middleware that hosts Java-based web applications. Discovery creates or updates a CMDB record when it detects an instance of a WebSphere application server running on a Windows or Linux system.

Note: Starting with Visibility Content version 6.1.0, pattern-based discovery is the default method for discovering IBM WebSphere Application Servers. The WebSphere probes are disabled by default. For more information, see IBM WebSphere Application Server discovery using patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Websphere discovery on Windows

Requirements for discovery:

For WebSphere application servers running on Windows systems, enable PowerShell on the MID Server.

Websphere discovery on Linux

Requirements for discovery:

- Enable SSH on the WebSphere Application Server.
- Set execute privileges to enable the following probes to run commands:
  - WebSphere - Get serverindex.xml: cat/read permissions on the `serverindex.xml` file
  - WebSphere - Get cell.xml: cat/read permissions on the `cell.xml` file
  - WebSphere - Get server.xml: cat/read permissions on the `server.xml` file

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Probes and sensors

Discovery uses this process to identify Linux WebSphere application servers:

1. The Unix - Active Processes probe detects a running process that matches the *`com.ibm.ws.runtime.WsServer`*. For example, the *`...com.ibm.ws.bootstrap.WSLauncher com.ibm.ws.runtime.WsServer`* process output has a parameter of */opt/IBM/WPS/profiles/ccmdb/config localhostNode01Cell ccmdb01 server1*. The values are:
   - last parameter = server_name = server1
   - last parameter - 1 = node_name = ccmdb01
   - last parameter - 2 = cell_name = localhostNode01Cell
   - last parameter - 3 = config_path = '/opt/IBM/WPS/profiles/ccmdb/config'
2. If there is a match, a record is created in the IBM Websphere [cmdb_ci_app_server_websphere] table. The following probes are triggered:
   - WebSphere - Cell
   - WebSphere - Web Applications
   - WebSphere –Web Services
3. The WebSphere - Cell probe searches for the `cell.xml` file for the instance by using the parameters in the running process, and then searching in the related `<config_path>\cells\<cell_name>\` directory.
4. If the probe successfully finds the `cell.xml` file, the sensor reads its contents and populates additional Websphere Cell [cmdb_ci_websphere_cell] table records as necessary.
5. If the probe successfully finds the `serverindex.xml` file, the sensor reads its contents and populates additional Web Application [cmdb_ci_web_application] table records as necessary.
6. If the probe successfully finds the `server.xml` file, the sensor reads its contents and populates additional Web Service [cmdb_ci_web_service] table records as necessary.

Discovery uses this process to identify Windows WebSphere application servers:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. The Windows - Active Processes probe detects a running process that matches the *com.ibm.ws.runtime.WsServer*. For example, the *...com.ibm.ws.bootstrap.WSLauncher com.ibm.ws.runtime.WsServer* process output has a parameter of */opt/IBM/WPS/profiles/ccmdb/config localhostNode01Cell ccmdb01 server1*. The values are:
   - last parameter = server_name = server1
   - last parameter - 1 = node_name = ccmdb01
   - last parameter - 2 = cell_name = localhostNode01Cell
   - last parameter - 3 = config_path = '/opt/IBM/WPS/profiles/ccmdb/config'
2. If there is a match, a record is created in the IBM Websphere [cmdb_ci_app_server_websphere] table. The following probes are triggered:
   - Windows - WebSphere - Cell
   - Windows - WebSphere - Web Applications
   - Windows - WebSphere –Web Services
3. The Windows - WebSphere - Cell probe searches for the `cell.xml` file for the instance by using the parameters in the running process, and then searching in the related `<config_path>\cells\<cell_name>\` directory.
4. If the probe successfully finds the `cell.xml` file, the sensor reads its contents and populates additional Websphere Cell [cmdb_ci_websphere_cell] table records as necessary.
5. The Windows - WebSphere - Web Applications probe searches the `serverindex.xml` file for the instance by using the parameters in the running process, and then searching in the related `<config_path>\cells\<cell_name>\nodes\<node_name>` directory.
6. If the probe successfully finds the `serverindex.xml` file, the sensor reads its contents and populates additional Web Application [cmdb_ci_web_application] table records as necessary.
7. The Windows WebSphere - Web Services probe searches for the `server.xml` file for the instance by using the parameters in the running process, and then searching in the related `<config_path>\cells\<cell_name>\nodes\<node_name>\servers\<server_name>` directory.
8. If the probe successfully finds the `server.xml` file, the sensor reads its contents and populates additional Web Service [cmdb_ci_web_service] table records as necessary.

Data collected

Data collected on Windows Websphere servers

Table	Field	Source
cmdb_ci_app_server_websphere	Name (name)	Running process
cmdb_ci_websphere_cell	Name (name)	cell.xml
cmdb_ci_websphere_cell	Cell ID (cell_id)	cell.xml
cmdb_ci_websphere_cell	Cell type (cell_type)	cell.xml
cmdb_ci_websphere_cell	Cell discovery protocol (cell_discovery_protocol)	cell.xml
cmdb_ci_web_service	Name (name)	server.xml
cmdb_ci_web_service	Service ID (service_id)	server.xml

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected on Windows Websphere servers (continued)

Table	Field	Source
cmdb_ci_web_application	Name (name)	serverindex.xml

Data collected on Linux Websphere servers

Label	Table Name	Field Name	Source
Name	cmdb_ci_websphere_cell	name	server.xml
App server	cmdb_ci_web_service	app_server	Internal reference
Name	cmdb_ci_web_service	name	server.xml
Name	cmdb_ci_web_application	name	serverindex.xml
App server	cmdb_ci_web_application	app_server	serverindex.xml

Relationships

These relationships are created in the CI Relationship [cmdb_rel_ci] table.

Parent class	Relationship type	Child class
cmdb_ci_app_server_websphere	Runs on::Runs	cmdb_ci_linux_server
cmdb_ci_win_server	Runs on::Runs	cmdb_ci_win_server
cmdb_ci_app_server_websphere	Contains::Contained by	cmdb_ci_web_application
cmdb_ci_web_service	Runs on::Runs	cmdb_ci_app_server_websphere
cmdb_ci_websphere_cell	Contains::Contained by	cmdb_ci_app_server_websphere

Database discovery

Discovery can find database applications, such as MySQL, Oracle, and MongoDB. Discovery can also find database management system software, such as MSSQL Server.

Database catalog discovery

The database catalog lists all the catalog objects, or databases, discovered for an instance of a database.

Databases vs. catalogs

Most database manufacturers use the term catalog to mean database, while Microsoft SQL uses the term database. Database catalogs can be imported into your instance from a third-party discovery tool, entered into the platform manually, or found by Discovery.

Important: Discovery does not support MySQL database catalog discovery in the base system. The Oracle DB on Unix and Oracle DB on Windows patterns discover Oracle catalog objects through the pattern extension Get Catalog info. The MSSql DB On Windows and MSSql 2000 DB patterns discover both database instances and catalogs and store them in the database instance records.

For more information about Oracle catalog discovery, see Oracle Catalog discovery.

For more information about Microsoft SQL database instances, see Microsoft SQL Server and Cluster discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you discovered database catalogs using probes before migrating to patterns, you can view the probe-based database catalogs by navigating to **All** > **Configuration** > **Database Catalogs** and selecting a database.

Database Administrator (DBA) report discovery

Discovery and Service Mapping uses the Patterns extension sections to provide a Database Administrator report (DBA report) for the Apache Cassandra, Microsoft SQL, MySQL, MongoDB, and Oracle databases. The extension sections for each DB pattern populate the related entries tables. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify the plugins setup

- CMDB CI Class Models 1.49.0 (September 2023)
- Discovery and Service Mapping Patterns 1.8.0 (September 2023)

Verify the system properties are enabled

By default, the system properties for DBA reports are set to **false**. To enable the report collection, verify that the system properties value for each pattern is set to **true**.

Pattern	System property
Cassandra	cassandra_dba_report_enabled
MongoDB	mongo_dba_report_enabled
MSSQL DB On Windows	mssql_dba_report_enabled
My SQL server On Windows and Linux	mysql_dba_report_enabled
Oracle DB	oracle_dba_report_enabled

Verify the configuration of the credentials

For more information, see Applicative credentials.

MongoDB execute command	Applicative credentials should have read access to run `"getSiblingDB (databaseName)"`
Apache Cassandra execute command	Discovery credentials should have read access to `"nodetool"`
MySQL tables and execute commands	Discovery credentials should have read access to:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

	• `"information_schema.tables"` table • run `"show databases"` command
Microsoft SQL tables	Discovery credentials should have read access to the following tables: • *master.sys.databases* • *sys.sysprocesses* • *sys.dm_os_sys_info* • *sys.dm_os_ring_buffers* • *msdb.dbo.backupset* • *master.dbo.sysdatabases* • *master..sysprocesses* • *msdb..sysjobactivity* • *sys.master_files* • *sys.dm_tran_session_transactions* • *sys.dm_os_sys_memory* • *sys.database_files*
Oracle DB tables and execute commands	Discovery credentials should have read access to: • Run `"lsnrctl services"` and `"sqlplus"` commands • *dba_data_files* table • *dba_segments* table • *v$database* table • *gv$instance* table • *dba_hist_osstat* table • *dba_hist_snapshot* table • *v$asm_disk* table • *v$asm_diskgroup* table • *dba_audit_session* table • *gv$session_longops* table • *gv$session* table • *gv$sqlarea* table • *gv$process* table

Verify that a Discovery schedule is configured

For more information, see create a discovery schedule

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by the extension section of the Apache Cassandra pattern

Discovery populates the following data when running the Apache Cassandra pattern.

Note: See the Prerequisites section for important information about enabling the system property.

Table name	Table label	Fields
cassandra_health_report	Cassandra Health Report	• database_host • local_write_latency • read_latency • snapshot_used_space • space_used • table • keyspace • ci

Data collected by the extension section of the MongoDB pattern

Discovery populates the following data when running the MongoDB pattern.

Note: See the Prerequisites section for important information about enabling the system property.

Table name	Table label	Fields
mongo_db_statistics	Mongo DB Statistics	• database_host • database_name • fs_total_size • fs_used_size • scale_factor • index_size • index_count • extents_count • storage_size • data_size • average_object_size • objects_count • collection_count • view_count • ci

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by the extension section of the MSSQL pattern

Discovery populates the following data when running the MSSQL pattern.

Note: See the Prerequisites section for important information about enabling the system property.

Table name	Table label	Fields
ms_sql_db_blocking_info	MSSQL DB Blocking Info	• database_host • database_instance • definition • object_name • blocking_spid • spid • ci
ms_sql_db_cpu_util	MSSQL DB CPU Util	• database_host • database_instance • load_date • other_process_cpu_util • other_process • system_idle_process • sql_server_cpu_util • event_time • server_name • ci
ms_sql_db_not_backup	MSSQL DB Not BackUp	• database_host • database_instance • last_backup_date • database_name • ci
ms_sql_db_running_job	MSSQL DB Running Job	• database_host • database_instance • login_name • nt_domain • cmd • hostname • name

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Table label	Fields
		• status • open_tran • Spid • last_wait_type • login_time • ci
ms_sql_disk_info	MSSQL DB Disks Info	• database_host • database_instance • percentage_free_space • volume_free_space • volume_capacity • volume_label • volume_name • ci
ms_sql_files_space	MSSQL DB Files Space	• database_host • database_instance • database_name • free_space_mb • file_size_mb • file_type • file_logic_name • ci
ms_sql_fregment_info	MSSQL DB Fragmentation Info	• database_host • database_instance • index_status • page_count • avg_frag • index_name • object_name • schema_name • ci
ms_sql_last_restart	MSSQL DB Last Restart	• database_host • database_instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Table label	Fields
		• uptime_in_days • current_date • last_recycle • ci
ms_sql_long_trans	MSSQL DB Long Transactions	• database_host • database_instance • blocked • sql • cmd • program_name • hostname • login_name • status • last_batch • duration • login_time • database_name • user_tran • trans_id • spid • ci
ms_sql_memory_info	MSSQL DB Memory Info	• database_host • database_instance • system_memory_state_desc • percentage_used • available_page_file_mb • total_page_file_mb • available_physical_memory_mb • total_physical_memory_mb • ci
ms_sql_temp_files_space	MSSQL DB Temp Files Space	• database_host • database_instance • percent_full • available_space_mb

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Table label	Fields
		• file_physical_name • file_logic_name • database_name • ci

Data collected by the extension section of the MySQL pattern

Discovery populates the following data when running the MySQL pattern.

Note: See the Prerequisites section for important information about enabling the system property.

Table name	Table label	Fields
mysql_table_schemas	MySql Table Schemas	• database_host • database_instance • data_index_lengh • table_schema • ci

Data collected by the extension section of the Oracle DB pattern

Discovery populates the following data when running the Oracle DB pattern.

Note: See the Prerequisites section for important information about enabling the system property.

Table name	Table label	Fields
oracle_active_session	Oracle Active Session	• database_host • database_instance • Key • sql_id • program • module • action • status • last_call_et_hrs • sid • spid

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Table label	Fields
		• inst • ci
oracle_asm_info	Oracle, ASM Disk Info	• database_host • database_instance • used_gb • free_gb • total_gb • state • block_size • asm • ci
oracle_asm_nodes_info	Oracle ASM Nodes Disk Info	• database_host • database_instance • redundancy • free_gb • total_gb • state • asm_node • ci
oracle_cpu_usage	Oracle CPU Usage	• database_host • database_instance • idle • io • sys • nice • user • snapid • Inst • end_time • begin_time • ci
oracle_inactive_sessions	Oracle Inactive Sessions	• database_host • database_instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Table label	Fields
		• key • status • spid • sid • program • Module • last_call_et_hrs • inst • ci
oracle_instance_info	Oracle Instance Info	• database_host • database_instance • platform_name • current_scn • dbid • db_unique_name • log_mode • db_name • open_mode • ci
oracle_nodes_info	Oracle Nodes Info	• database_host • database_instance • instance_mode • instance_state • instance_role • instance_status • instance_version • hostname • instance_name • ci
oracle_table_space	Oracle Table Space	• database_host • database_instance • used_mb • free_percent • free_mb

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Table label	Fields
		• tablespace • total_mb • ci
oracle_top_10_queries	Oracle Top 10 Queries	• database_host • database_instance • unique • status • sql_id • spid • sid • program • Module • last_call_et_hrs • ci

Microsoft SQL Server and Cluster discovery

The Discovery and Service Mapping Patterns application uses the MSSql DB On Windows pattern to find Microsoft SQL DB servers and clusters on your infrastructure. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

For detailed information on SQL Server supported versions see: Detailed information on products discovered by ITOM Visibility

Prerequisites

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns starting with version 1.7.0 (August 2023)
- Visibility Content starting with version 6.12.1 (August 2023)

Create Windows credentials

After verifying that the MID Server is connected to the target Windows server, create Windows credentials on the ServiceNow AI Platform. For more information, see Windows credentials.

Create applicative credentials when using SQL authentication

For more information, see Applicative credentials.

Verify the configuration of the MID Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Verify that the MID Server has the necessary permissions to read the Windows Registry on a remote mode. The MSSql DB On Windows pattern uses registry values to discover SQL cluster details.

Verify SQL Server discovery report is active

Verify that the SQL Server discovery report is active on your server. Discovery uses the SQL Server discovery report to discover Microsoft SQL Server components.

Verify the permissions to run WMI queries

Verify remote read-only access to the following WMI namespaces:

- `Root\CIMv2`
- `Root\Microsoft\SqlServer\ComputerManagement*`

Verify the SQL Server table permissions

The MSSql DB On Windows pattern requires read-only permissions to read the following tables:

- `@@version`
- `sys.configurations`
- `sys.dm_exec_sessions`
- `sys.dm_os_schedulers`
- SERVERPROPERTY

Verify availability group table permissions

Verify the following table permissions to discover Microsoft SQL Server Always On availability groups:

- **VIEW SERVER STATE** access to the `sys.dm_hadr_availability_group_states` table
- Read access permission to the following tables:
  - `sys.availability_groups`
  - `sys.availability_group_listeners`
  - `sys.availability_group_listener_ip_addresses`

SQL failover cluster instance (FCI)

Starting with version 1.6.1 of the Patterns available from the ServiceNow® Store (August 2023), discovery of Always On availability group details and SQL Failover Clusters is supported. When a failover scenario occurs, the MSSql DB On Windows pattern discovers the failover cluster (Node 2). A new MSFT SQL instance and databases are then created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

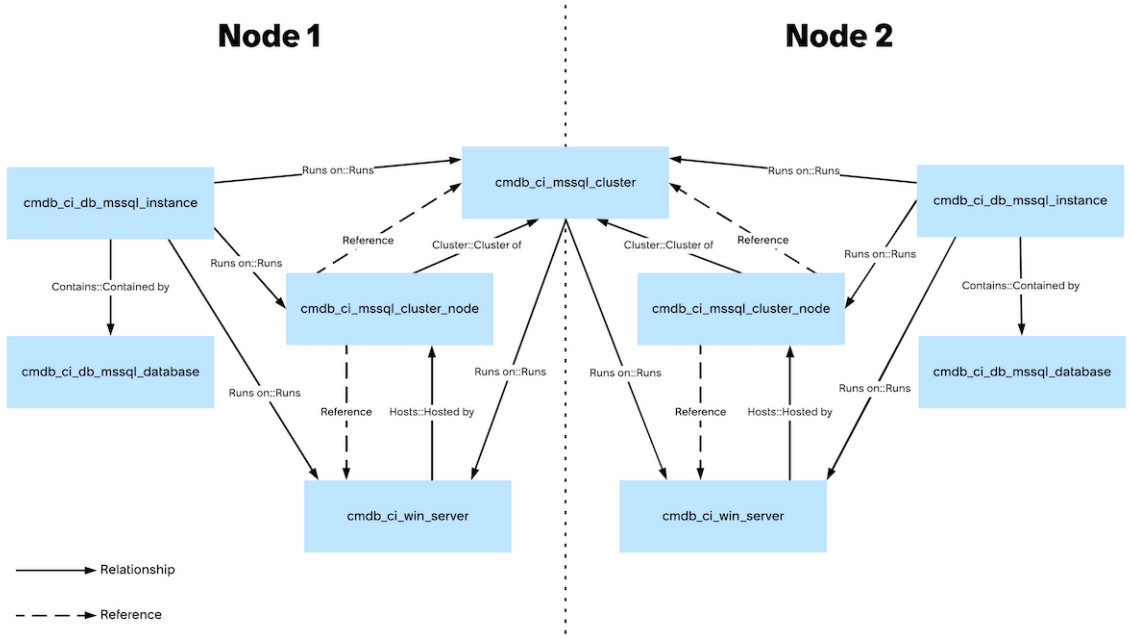

The following procedure enables appropriate visibility of the SQL clusters.

1. Set the *sn_itom_pattern.discover_mssql_cluster.load_cluster_name_from_registry=true* MID Server property to **true** to enable the **Collect MSSQL Cluster info- MSSQL Cluster** extension section.
2. Ensure that an `sqlservr.exe` process to trigger the MSSql DB On Windows pattern is running on the server.
3. Run a CI discovery.

   For more information, see Running discoveries in your network.
4. Either keep duplicate instances and databases or delete them.
   - To keep the duplicate instances and databases, keep the default configuration.
     - The install status of the nodes of the passive cluster changes to **Absent**.
     - The install status of the nodes of the active cluster changes to **Installed**.
   - To delete the passive instances and databases, set the *discovery.mssql.cluster.instance.postsensor.delete* system property to **True**.

Microsoft SQL Server Always On availability groups support

Starting from Discovery and Service Mapping Patterns version 1.27.0, the MSSql DB On Windows pattern extension **Collect MSSQL HADR - Availability Group Info** supports the discovery of Microsoft SQL Server Always On availability groups.

The Collect MSSQL HADR - Availability Group Info pattern extension introduces the following CI classes that extend existing CMDB classes:

CI classes introduced by this pattern

CI class	Extends from
MSSQL Availability Group [cmdb_ci_mssql_ag]	Cluster [cmdb_ci_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI classes introduced by this pattern (continued)

CI class	Extends from
MSSQL Availability Group Listener [cmdb_ci_mssql_ag_listener]	Endpoint [cmdb_ci_endpoint]
MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica]	MSSQL Cluster Node [cmdb_ci_mssql_cluster_node]

Data model of Collect MSSQL HADR - Availability Group Info

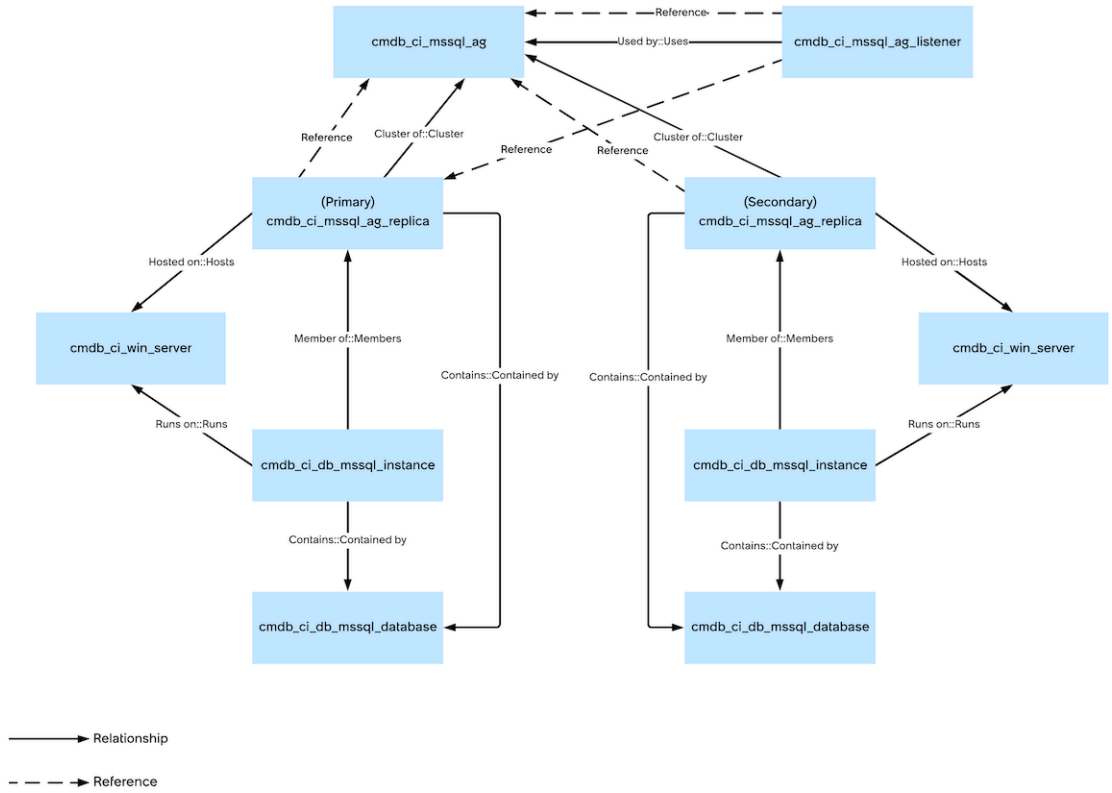

Prerequisites

Verify availability group table permissions

Discovered resources

Resources discovered by the pattern extension: Collect MSSQL HADR - Availability Group Info

Discovered relationships

Relationships discovered by the pattern extension: Collect MSSQL HADR - Availability Group Info

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the MSSql DB On Windows pattern.

MSFT SQL instance [cmdb_ci_db_mssql_instance]

The data collected for the running processes of the database (the actual SQL server) is referred to as the database instance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name [name]	The field name. For example: MSSQLSERVER@windows_server_name
Instance Name [instance_name]	The SQL server instance name. For example: MSSQLSERVER
Is clustered [is_clustered]	Indicates the type of installation: • If this value is set to **true**,the SQL server is a part of a cluster, and the failover mechanism is enabled. • If this value is set to **false**, the SQL server is a standalone deployment and the failover mechanism is disabled.
Comments [comments]	The type of instance: • standalone_instance • failover_cluster • always_on
Version [version]	The SQL server version number. For example: 16.0.1000.6
Version Name [version_name]	The SQL server version by release dates. For example: The value **2022** represents version **16.0.1000.6**
Service Pack [service_pack]	The service pack version. For example: SP2.
Edition [edition]	The type of the edition. For example: Enterprise Edition.
TCP Port [tcp_port]	The TCP port that the MSSQL instance uses to accept connections.
Install Status [install_status]	The install status of the instance: **Installed**- The instance is currently running. **Absent**- The instance isn't running and it's the backup instance for a failover scenario.

MSSQL database [cmdb_ci_db_mssql_database]

Field	Description
Name [name]	The name of the database. For example: msdb.
Install Status [install_status]	The install status of the database.
Instance name [instance_name]	References the MSFT SQL instance [cmdb_ci_db_mssql_instance] table.

MSSQL Cluster Node [cmdb_ci_mssql_cluster_node]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name [name]	The server host name
IP Address [ip_address]	The hosting Windows server's IP address
Cluster [cluster]	Reference to the SQL cluster record.
Server [server]	Reference to the hosting Windows server record.

MSSQL Cluster [cmdb_ci_mssql_cluster]

Field	Description
Name [name]	The SQL cluster network name
IP Address [ip_address]	The SQL cluster's virtual IP address.

MSSQL Service Info [mssql_sqlservice_info]

Field	Description
Description [description]	Describes the MSSQL service. For example: "Provides storage, processing, and controlled access of data and rapid transaction processing."
Service Name [service_name]	The MSSQL service name as presented on the server. For example: MSSQL$NAMEDMSSQL2016
Binary Path [binary_path]	The file path of the service. For example: `C:\Program Files\Microsoft SQL Server\MSSQL13.NAMEDMSSQL2016\MSSQL\Binn\sqlservr.exe`
CI [ci]	Reference the MSFT SQL instance [cmdb_ci_db_mssql_instance] table.
Installed On [installed_on]	Reference to the Windows server record [cmdb_ci_win_server] **Note:** This field isn't populated for standalone deployments.

MSSQL Components Info [mssql_components_info]

Field	Description
Install Path [install_path]	The install path of the component service. Foe example: `C:\Program Files\Microsoft SQL Server\MSSQL16.SQLEXPRESS\MSSQL`
Service Name [service_name]	The components service type of the server. For example: **Database Engine Services**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
SKU Name [sku_name]	The MSSQL version of the component service. For example: **Express Edition (64-bit)**
Version [version]	The MSSQL version number of the component service
CI [ci]	Reference the MSFT SQL instance [cmdb_ci_db_mssql_instance] table.
Installed On [installed_on]	Reference to the Windows server record [cmdb_ci_win_server]. **Note:** This field is populated only for SQL Clusters and not for standalone deployments.

MSSql DB On Windows pattern extension: Collect MSSQL HADR - Availability Group Info

MSSQL Availability Group [cmdb_ci_mssql_ag]

Field	Description
Cluster ID [cluster_id]	Globally unique identifier (GUID) of the availability group.
Name [name]	Cluster name.
IP Address [ip_address]	IP address of the cluster.

MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica]

Field	Description
Replica ID [replica_id]	Unique ID of the replica.
Role [role]	Role of the replica. Possible values are: • Primary • Secondary • Resolving
Availability Mode [availability_mode]	Availability mode of the replica. Possible values are: • Asynchronous Commit • Synchronous Commit • Configuration Only
Failover Mode [failover_mode]	Failover mode of the availability replica. Possible values are: • Automatic • Manual

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica] (continued)

Field	Description
Readable Secondary [readable_secondary]	Whether an availability replica that's performing the secondary role can accept connections from clients. Possible values are: • true • false
Availability Group [availability_group]	References the MSSQL Availability Group [cmdb_ci_mssql_ag] table.
Name [name]	Name of the replica server.
Serial number [serial_number]	Unique ID of the replica.
State [state]	Failover state. Possible values: • Pending • Online • Offline • Failed • Failed No Quorum

MSSQL Availability Group Listener [cmdb_ci_mssql_ag_listener]

Field	Description
Name [name]	Name of the availability group listener.
Listener ID [listener_id]	Resource globally unique identifier (GUID) from Windows Server Failover Clustering (WSFC) cluster.
IP Subnet Mask [ip_subnet_mask]	IP subnet mark for the IPv4 address, if any, that's configured for the availability group listener.
Availability Group [availability_group]	References the MSSQL Availability Group [cmdb_ci_mssql_ag] table.

CI relationships and references

The MSSql DB On Windows pattern creates the following relationships and references to support Microsoft SQL Server and Cluster discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Windows Server [cmdb_ci_win_server]	Runs::Runs on	MSSQL Cluster [cmdb_ci_mssql_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
MSFT SQL instance [cmdb_ci_db_mssql_instance]	Runs::Runs on	MSSQL Cluster [cmdb_ci_mssql_cluster]
MSSQL Cluster Node [cmdb_ci_mssql_cluster_node]	Cluster::Cluster of	MSSQL Cluster [cmdb_ci_mssql_cluster]
MSSQL Cluster Node [cmdb_ci_mssql_cluster_node]	Hosted on::Hosts	Windows Server [cmdb_ci_win_server]
MSFT SQL instance [cmdb_ci_db_mssql_instance]	Runs on::Runs	Windows Server [cmdb_ci_win_server]
MSFT SQL instance [cmdb_ci_db_mssql_instance]	Runs on::Runs	MSSQL Cluster Node [cmdb_ci_mssql_cluster_node]
MSFT SQL instance [cmdb_ci_db_mssql_instance]	Contains::Contained by	MS SQL Database [cmdb_ci_db_mssql_database]

CI references

CI	Field	Referenced CI
MSSQL Cluster Node [cmdb_ci_mssql_cluster_node]	Cluster [cluster]	MSSQL Cluster [cmdb_ci_mssql_cluster]
MSSQL Cluster Node [cmdb_ci_mssql_cluster_node]	Server [server]	Windows Server [cmdb_ci_win_server]
MSSQL Service Info [mssql_sqlservice_info]*	CI [ci]	MSFT SQL instance [cmdb_ci_db_mssql_instance]
MSSQL Components Info [mssql_components_info]	CI [ci]	MSFT SQL instance [cmdb_ci_db_mssql_instance]
MS SQL Database [cmdb_ci_db_mssql_database]	Instance name [instance_name]	MSFT SQL instance [cmdb_ci_db_mssql_instance]

*This reference is created only when the pattern uses the WMI queries.

The MSSql DB On Windows pattern creates the following relationships and references using the **Collect MSSQL HADR - Availability Group Info** pattern extension. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
MSFT SQL instance [cmdb_ci_db_mssql_instance]	Member of::Members	MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica]
MSSQL Availability Group Listener [cmdb_ci_mssql_ag_listener]	Used by::Uses	MSSQL Availability Group [cmdb_ci_mssql_ag]
MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica]	Contains::Contained by	MS SQL Database [cmdb_ci_db_mssql_database]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica]	Cluster of::Cluster	MSSQL Availability Group [cmdb_ci_mssql_ag]
MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica]	Hosted on::Hosts	Windows Server [cmdb_ci_win_server]

CI references

CI	Field	Referenced CI
MSSQL Availability Group Listener [cmdb_ci_mssql_ag_listener]	Availability Group [availability_group]	MSSQL Availability Group [cmdb_ci_mssql_ag]
MSSQL Availability Group Replica [cmdb_ci_mssql_ag_replica]	Availability Group [availability_group]	MSSQL Availability Group [cmdb_ci_mssql_ag]

Related topics

Add the Horizontal Pattern probe to a classifier

Horizontal Pattern probe

Database Administrator (DBA) report discovery

MySQL discovery

Discovery can identify an instance of MySQL that is running on UNIX or Windows operating systems.

Requirements

Discovery searches for the MySQL configuration file location from the following areas:

- UNIX: Discovery searches for the MySQL configuration file location from the mysqld process, or port 3306.
- Windows: Discovery searches for the MySQL configuration file location from the mysqld.exe process, or port 3306.

For each process, the following process parameters are explored in the following order:

--defaults-extra-file

--defaults-file

If the MYSQL configuration file location is not found from that search, then the following occurs:

- UNIX: The configuration file location defaults to `/etc/my.cnf`.
- Windows: No default configuration file location exists, and the probe to read the configuration file location is skipped.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Credentials

Configure these credentials:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Windows credentials
- SSH credentials
- Applicative credentials (When applications require credentials in addition to the credentials the that host machine requires.)

Note: There's no need to configure applicative credentials unless your Discovery prerequisites mention them.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Classifiers, patterns, and probes

Discovery uses these classifiers, probes, and patterns for MySQL discovery:

Classifier	Trigger probes	Pattern
MySQL Server	• Horizontal Pattern: launches patterns • Windows - MySQL Configuration* • Windows - MySQL Version* • MySQL - Configuration* • MySQL - Version*	My SQL server On Windows and Linux Pattern

*For new instances, these probes are inactive on the classifier. Discovery uses patterns for discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

Discovery populates these fields in the MySQL Instance [cmdb_ci_db_mysql_instance] table:

Label	Field name	Source
MySQL configuration	myconf	my.cnf
TCP port(s)	tcp_port	running_process
Version	version	mysqld

Related topics

Database Administrator (DBA) report discovery

MongoDB discovery

Discovery creates or updates a CMDB record when it detects a running instance of MongoDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credentials

Configure SSH credentials.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

The following data is collected:

Data Collected by Discovery on MongoDB instances

Label	Table Name	Field Name	Source
Version	cmdb_ci_db_mongodb_instance	version	mongod (UNIX) or mongod.exe (Windows)
Mongo configuration	cmdb_ci_db_mongodb_instance	mongodb_conf	mongod.conf
TCP port(s)	cmdb_ci_db_mongodb_instance	tcp_port	Process Classification or mongod.conf

PostgreSQL discovery

Discovery can find running instances of PostgreSQL on Windows and Linux systems.

Credentials and other prerequisites

Create SSH credentials for PostgreSQL discovery

For more information, see SSH credentials.

Verify root-level access to the database

The user must have root-level access to the database to access the `postgresql.conf` file.

Verify privileged commands for for PostgreSQL discovery

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Differentiate multiple PostgreSQL instances on the same host by port number

Starting with Discovery and Service Mapping Patterns version 1.35.0, you can make multiple PostgreSQL instances on the same host distinguishable by port number by creating the *`mid.discovery.postgresql.include_port_in_name`* MID Server property. For more information, see Include the port number in PostgreSQL instance names.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifiers, patterns, and probes

Classifier	Trigger probes	Patterns
PostgreSQL Instance	• Horizontal Pattern: launches patterns • PostgreSQL - Configuration* (add the *must_sudo* parameter to this probe) • PostgreSQL - Version*	PostgreSQL DB

*For new instances, these probes are inactive on the classifier. Discovery uses patterns for discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

Discovery populates the data in the CMDB when running the PostgreSQL DB pattern.

The following fields gather specified information from the target. If the source is not configured, it brings back default information. For instance, for PostgreSQL Instance@hostname (default name), the source needs to be modified. If not, all the "PostgreSQL Instance@hostname" will be added in the source for the PostgreSQL Instance [cmdb_ci_db_postgresql_instance] table.

PostgreSQL Instance [cmdb_ci_db_postgresql_instance]

Field	Description
Name [name]	The display name of the PostgreSQL instance.*
Data Directory [data_dir]	The data directory of the PostgreSQL instance, parsed from the process command line.
TCP port(s) [tcp_port]	The port on which the PostgreSQL instance is listening, determined from the running process.
Config File [postgres_conf]	The path to the `postgresql.conf` configuration file.
Version [version]	The PostgreSQL version, retrieved from the process executable.

* To populate the **Name** field in the format **instance-port-tcp_port@hostname** instead of the default **instance@hostname**, create the *mid.discovery.postgresql.include_port_in_name* MID Server property. For more information, see Include the port number in PostgreSQL instance names.

Relationships

CI	Relationship	CI
PostgreSQL Instance [cmdb_ci_db_postgresql_instance]	Runs on::Runs	Windows Server [cmdb_ci_windows_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
		or Linux Server [cmdb_ci_linux_server]

Include the port number in PostgreSQL instance names

You can make multiple PostgreSQL instances on the same host distinguishable by port number by creating the *mid.discovery.postgresql.include_port_in_name* MID Server property. This property is supported for UNIX hosts only.

Before you begin

Verify that you have at least version 1.35.0 of Discovery and Service Mapping Patterns.

Role required: mid_server, agent_admin, or agent_security_admin

About this task

By default, Discovery assigns every PostgreSQL instance the display name **instance@hostname**, regardless of port, making multiple instances on the same host indistinguishable in list views. The *mid.discovery.postgresql.include_port_in_name* MID Server property controls this behavior. When created and set to true, Discovery uses the format **instance-port-tcp_port@hostname** instead (for example, **postgres-port-5432@host**). Existing CI records aren't affected.

Procedure

1. Navigate to **All** > **MID Server** > **Properties**.
2. Select **New**.
3. On the form, fill in the fields.

MID Server Property New record fields

Field	Description
Name	Enter `mid.discovery.postgresql.include_port_in_name`.
Value	Enter `true`.
MID server	MID Servers to use for the property. Leave this field empty to apply the property to all MID Servers.

4. Select **Submit**.

What to do next

Run Discovery again to apply the change.

Related topics

PostgreSQL discovery

Oracle database discovery

Discovery can identify an Oracle database instance that is running on UNIX or Windows operating systems. It can also find Oracle clusterware, which runs Real Application Clusters (RAC).

Patterns perform the discovery of Oracle installations. These patterns replace several probes for new instances. If you want to continue to use probes, you can deactivate the Horizontal discovery probe, which launches the pattern, and then activate the necessary probes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Oracle DB on Unix and Oracle DB on Windows do not support RAC single-node discovery.

The Oracle DB on Unix and Oracle DB on Windows patterns discover Oracle catalog objects through the pattern extension Get Catalog info. For more information, see Oracle Catalog discovery. For general information about database catalog discovery, see database catalogs.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Credentials

Configure these credentials:

Credential type	Additional information
Applicative credentials	Applicative credentials are mandatory for two use cases: • Executing SQL queries against the remote Oracle DB Instance • Creating connections to database links • Software Asset Management activated You must add the database user you configured during Oracle database installation. Specify `Oracle instance [cmdb_ci_db_ora_instance]` as the CI type. The user must also have access to the *`tnsnames.ora`* parameters file. Make sure that the tnsnames.ora is configured correctly and exists in the Oracle path on the target Windows host. Also, verify that the utilized Discovery Credentials can read tnsnames.ora file. The user needs permissions for the SQL*Plus command line tool and permissions to execute SELECT queries against the following tables: • PRODUCT_COMPONENT_VERSION • V$VERSION • V$DATABASE • V$PARAMETER
Oracle Wallet authentication	You can use Oracle Wallet instead of applicative credentials.
Windows credentials	Required for Windows targets. The user must have read permissions to the `tnsnames.ora` file. For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.
SSH credentials	Required for UNIX targets. The user must have read permission to the `oratab` file and read permission to the System Parameter file. For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credential type	Additional information
	list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Additional requirements

Discovery identifies a running instance of an Oracle database on UNIX from the process that starts with `ora_pmon_`. Ensure this process is running in the IP range you designate for Discovery. Discovery identifies a running instance of an Oracle database on Windows from the `oracle.exe` process. Ensure this process is running in the IP range you designate.

For Oracle clusterware, you also need permissions to read the `/etc/hosts/` file, which associates host names with IP addresses.

The following commands are run for Oracle clusterware:

Command	Description
`ps -ef | grep 'crsd.bin' |grep -v grep`	Verifies if Oracle CRS process is running.
`ps -ef |grep LISTENER | grep -v $$ | head -1 | awk '{print $8}'`	Gets Oracle home from the listener process path. Used when the ORACLE_HOME environment variable is unavailable.
`ps -ef | grep 'd.bin' | grep -v grep | grep -v bash | awk '{print $8}'`	Gets Oracle Grid home from the CRS daemon process path.
`/u01/app/12.1.0.2/grid/bin/olsnodes -c | egrep -v 'error|return code'`	Gets Oracle CRS cluster name.
`sudo /u01/app/12.1.0.2/grid/bin/ocrcheck | egrep -v 'error|return code'`	Gets Oracle CRS cluster ID.
`/u01/app/12.1.0.2/grid/bin/crsctl query crs softwareversion| egrep -v 'error|return code'`	Gets Oracle CRS version.
`/u01/app/12.1.0.2/grid/bin/olsnodes | egrep -v 'error|return code'`	Gets Oracle CRS nodes.
`/u01/app/12.1.0.2/grid/bin/crsctl status server | egrep -v 'error| return code'`	Gets Oracle CRS node status.
`/u01/app/12.1.0.2/grid/bin/crsctl stat res -f | tr '\n' '#'`	Gets Oracle CRS resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Command	Description
`/u01/app/12.1.0.2/grid/bin/crsctl stat res -p | grep -v CHECK_R |grep -v error |grep -v failed`	Gets Oracle CRS VIP resources.
`/u01/app/12.1.0.2/grid/bin/srvctl config scan | tr '\n' ' '`	Gets Oracle RAC SCAN VIP configurations.
`/u01/app/12.1.0.2/grid/bin/srvctl config vip -node ol7-121-rac1| egrep ' Name' | awk '{print $3}'`	Gets Oracle CRS VIPs name.
`/u01/app/12.1.0.2/grid/bin/srvctl config vip -node ol7-121-rac1| grep ' IPv4 Address:' | awk '{print $4}'`	Get Oracle CRS VIPs ip_address.

Important: Sudo permission are required to run `sudo /u01/app/12.1.0.2/grid/bin/ocrcheck | egrep -v 'error|return code'`.

Classifiers, patterns, and probes

Discovery uses these classifiers, probes, and patterns for Oracle database discovery:

Classifier	Trigger probes	Pattern
Oracle Instance	• Horizontal discovery probe: launches patterns • Powershell-Oracle - Instance PFile (for Windows)* • Powershell-Oracle - Instance Version (for Windows)* • SSHCommand-Oracle - Instance Details (for Linux)*	Application patterns: • Oracle DB On Windows Pattern • Oracle DB on UNIX Pattern This shared library pattern used by the Linux Server pattern: UNIX Cluster - ORACLE Clusterware **Note:** This shared library pattern is triggered if the `crsd.bin` process is running on the Oracle database server.
Oracle Listener	SSHCommand-Oracle - Listener Details	none

*Probes inactive for new instances.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected

The following data is collected on the Oracle Instances [cmdb_ci_db_ora_instance] table for both UNIX and Windows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected on UNIX machines

Label	Field name	Additional information
SID	sid	Obtained from the name of the process that starts with `ora_pmon_`
Install directory	install_directory	Obtained from the path of ORA_HOME
Version	version	Obtained in this order: • From the output of the `ORA_HOME/bin/sqlplus /NOLOG` command • From the output of the `ORA_HOME/bin/lsnrctl` status command • From the path of ORA_HOME
Parameter file	pfile	The following locations are explored for the location of the System Parameter File. If this file does not exist in one of the explored locations, Discovery does not find the file and reports an error. • *oracle_home*/dbs/spfile*SID*.ora • *oracle_home*/dbs/spfile.ora • *oracle_home*/dbs/init*SID*.ora
Oracle Home	oracle_home	Obtained from the *ORATAB* file.
Edition	edition	Obtained from the `context.xml` file.
Name	name	Combination of SID name Host name in the format *SID@HOST_NAME*.
TCP port(s)	tcp_port	Obtained from the `tnsnames.ora` file, or *lsnrctl* status command. If no valid information is retrieved, the default port 1521 is assigned.

Note: See Oracle options for pattern extensions and additional information collected.

Data collected on Windows machines

Label	Field name	Additional information
SID	sid	Obtained from the process parameter that is passed to the *oracle.exe* process.
Install directory	install_directory	
Version	version	Obtained from the output of the *ORA_HOME/bin/sqlplus.exe -V* command.
Parameter file	pfile	The following locations are explored for the location of the System Parameter File. If this file does not exist in one of the explored locations, Discovery does not find the file and reports an error.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected on Windows machines (continued)

Label	Field name	Additional information
		• *oracle_home*\database\spfile*SID*.ora • *oracle_home*\database\spfile.ora • *oracle_home*\database\init*SID*.ora
Oracle Home	oracle_home	Parsed from the path of *oracle.exe*
Edition	edition	

Note: See Oracle options for pattern extensions and additional information collected.

Oracle options

For general information on Oracle options, see https://docs.oracle.com/en/.

Discovery can also find the Oracle options that you enable on the instance using the legacy pattern extension, Oracle option extension for Windows/Unix. Note that to discover Oracle options, the Software Asset Management Professional for Oracle (com.snc.samp.oracle) plugin is required.

Oracle Options [samp_oracle_options]

Label	Field name
Instance	instance
Option	option
Licensable	licensable
Currently used	currently_used
First usage date	first_usage
Last usage date	last_usage
Last sample date	last_sample
Detected usage	detected_usage
Sample interval	sample_interval

Table privileges needed	Additional information
The following tables are accessed as part of the Oracle option extension for UNIX and Windows: • DBA_FEATURE_USAGE_STATISTICS • CDB_FEATURE_USAGE_STATISTICS • V$DATABASE • V$VERSION • GV$PARAMETER	The following pattern extensions use the Oracle script: options_packs_usage_statistics.sql. This sql script is for DB Server version 11.2 and later. You need the **Select any dictionary** privilege to execute the script. See the support document ID13172651.1 on https://support.oracle.com/portal/ for more information about usage scripts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table privileges needed	Additional information
• V$CONTAINERS • V$INSTANCE	• Oracle option extensions for UNIX • Oracle option extensions for Windows

Data collected on clusterware

Table	Labels and field names
Unix Cluster [cmdb_ci_unix_cluster]	Name [name] Cluster status [cluster_status] Cluster type [cluster_type] IP address [ip_address] Short description [short_description] Cluster ID [cluster_id] Cluster version [cluster_version]
cmdb_ci_unix_cluster_node	Name [name] Node state [node_state] IP address [ip_address] Node status [node_status]
cmdb_ci_unix_cluster_resource	Name [name] Resource type [resource_type] Resource status [resource_status] Properties [properties]
cmdb_ci_cluster_vip	Name [name] IP address [ip_address] Cluster name [cluster_name] Cluster ID [cluster_id]

On the dependency view, a **+1** on an Oracle database server signifies that it is part of a cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Extend Oracle instance extension

Note: Discovery can also find additional attributes for the instance using the Extend Oracle instance extension. See the Oracle website for more information on options https://docs.oracle.com/en/. Discovery saves these attributes to the Oracle Instances [cmdb_ci_db_ora_instance] table.

Label	Field name
Instance	instance
Host Name	host_name
IS_CDB	cdb_yn
DB Unique Name	db_unique_name
DB Name	database_name
Replication Enabled	replication_enabled

Table privileges needed	Additional information
The following tables are accessed as part of the Extend Oracle Instance extension for UNIX and Windows: • V$DATABASE • V$INSTANCE	The Extend Oracle Instance extension needs to be installed. This extension is included in the Discovery and Service Mapping Patterns application, version 1.0.72 or higher, which can be downloaded from the ServiceNow Store.

Oracle instances on virtual machines

When Oracle instances are hosted on virtual machines, the system creates relationships between CIs.

CI	Relationship	CI
Oracle instance	Hosted on	Computer [cmdb_ci_computer] or Server [cmdb_ci_server] **Note:** This is a virtual machine. The **Is virtual** field is **true**.
Computer [cmdb_ci_computer] or Server [cmdb_ci_server] **Note:** This is a virtual machine. The **Is virtual** field is **true**.	Virtualized by	Computer [cmdb_ci_computer] or Server [cmdb_ci_server] **Note:** This is not a virtual machine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle instances on virtual machines that are hosted on AWS or Azure clouds

Base class	Relationship	Dependent class
Oracle instance	Hosted on	Computer [cmdb_ci_computer] **Note:** This is a virtual machine. The **Is virtual** field is **true**.
Computer [cmdb_ci_computer] **Note:** This is a virtual machine. The **Is virtual** field is **true**.	Virtualized by	Computer [cmdb_ci_computer] **Note:** This is not a virtual machine.
Computer [cmdb_ci_computer] **Note:** This is a virtual machine. The **Is virtual** field is **true**.	Virtualized by	Virtual Machine Instance [cmdb_ci_vm_instance] **Note:** The vCPU count is populated in this table.
Virtual Machine Instance [cmdb_ci_vm_instance]	Hosted on	AWS Datacenter [cmdb_ci_aws_datacenter] or Azure Datacenter [cmdb_ci_azure_datacenter]

Relationships for clusterware

CI	Relationship	CI
Unix Cluster [cmdb_ci_unix_cluster]	Virtualizes::Virtualized by	Cluster Virtual IP [cmdb_ci_cluster_vip]
Unix Cluster Resource [cmdb_ci_unix_cluster_resource]	Defines resources for::Gets resources from	Unix Cluster Node [cmdb_ci_unix_cluster_node]
Cluster Virtual IP [cmdb_ci_cluster_vip]	Uses::Used by	Unix Cluster Node [cmdb_ci_unix_cluster_node]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Cluster of::Cluster	Unix Cluster [cmdb_ci_unix_cluster]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Hosted on:Hosts	Unix Cluster [cmdb_ci_unix_cluster]
Unix Cluster Resource [cmdb_ci_unix_cluster_resource]	Defines resources for::Gets resources from	Unix Cluster [cmdb_ci_unix_cluster]
Unix Cluster [cmdb_ci_unix_cluster]	Hosts::Hosted on	Linux Server [cmdb_ci_linux_server]
Unix Cluster Node [cmdb_ci_unix_cluster_node]	Hosted on::Hosts	Linux Server [cmdb_ci_linux_server]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Add the Horizontal Pattern probe to a classifier
- Horizontal Pattern probe
- Database Administrator (DBA) report discovery
- Oracle Catalog discovery
- Oracle pluggable database and container database discovery
- Oracle GLAS data collection using patterns

SAP Sybase ASE discovery

Discovery and Service Mapping uses the Sybase pattern to find SAP Sybase Adaptive Server Enterprise (ASE) databases and catalogs in your infrastructure. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Verify that the applications are up to date.

Discovery and Service Mapping Patterns

CMDB CI Class Models

Visibility Content

For details on the product discovered see Detailed information on products discovered by ITOM Visibility

Note: Discovery finds Sybase Catalog that run on Linux only.

Verify the configuration of OS server credentials.

Verify the configuration of the Applicative Credentials

When configuring the Applicative Credentials, make sure you specify `Sybase Instance [cmdb_ci_db_syb_instance]` as the CI type.

For more information on Applicative Credentials, see Applicative credentials

Verify that the user have permission to run the following commands

Execute the *`isql`* command.

Have the Sybase environment run one of the commands:

- `"echo -ne 'set nocount on select name from sysdatabases \ngo\n' | " + $isql_path + " -U $$username$$ -P '$ $password$$' -S " + $instance + " -w 999 -b -y " + $process.environmentVariables.SYBASE.value`
- `"export SYBASE=" + $process.environmentVariables.SYBASE.value + " && . " + $process.environmentVariables.SYBASE.value`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
+ "/SYBASE.sh && echo -ne 'set nocount on
select name from sysdatabases\ngo\n' | " +
$isql_path + " -U $$username$$ -P '$$password
$$' -S " + $instance + " -w 999 -b -y " +
$process.environmentVariables.SYBASE.value
```

Verify the configuration of a Discovery schedule

For detailed information, see Running discoveries in your network

Data collected during Horizontal discovery

Discovery populates the following data in the CMDB when running the Sybase pattern.

Sybase instance [cmdb_ci_db_syb_instance]	
Label	**Field Name**
Name	name
Version	version
Installation directory	install_directory
Instance name	Instance
Configuration file	config_file

Sybase Catalog [cmdb_ci_db_syb_catalog]	
Label	**Field Name**
Name	name

CI relationships

These relationships are created to support the Sybase ASE discovery.

CI	Relationship	CI
[cmdb_ci_db_syb_instance]	Contains::Contained by	[cmdb_ci_db_syb_catalog]

Network device discovery

Discovery identifies several network devices, including load balancers, routers, TCP connections, IP networks, and so on. It can find devices through traditional port probes, and then through identifications and exploration probes, or through patterns. It also supports discovery by identifying HTTP and TCP connections.

Note: See the knowledge article KB0687582 for information on model_id and manufacturer.

Load balancer discovery

Discovery can collect data about network routers, switches, and applications.

Discovery supports data collection from the following applications and hardware:

- A10
- Apache mod_jk and Apache mod_proxy

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Big-IP F5 Traffic Manager
- Cisco
- Citrix NetScaler
- HAProxy
- NGINX
- Alteon
- ACE
- Radware

Important: Discovery treats load balancers as licensable entities and attempts to discover them primarily using SNMP. If a load balancer in your system, running on a Linux host, has SNMP and SSH ports open, Discovery might classify it based on the SSH port. This classification has priority over SNMP. To ensure that Discovery properly classifies your hardware load balancers, create a Discovery behavior for load balancers that includes SNMP but not SSH. Software load balancers are treated as applications.

Load balancer patterns

Starting with the Jakarta release, Discovery uses out-of-box patterns to find some types load balancers for new instances (see table).

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Load balancer	Pattern names and description	Main CI table populated
A10	A10 LoadBalancer	A10 Load Balancer [cmdb_ci_lb_a10]
ACE	ACE Load Balancer by SSH (formerly ACE Load Balancer)	ACE [cmdb_ci_lb_ace]
Alteon	Alteon Load Balancer	Alteon [cmdb_ci_lb_alteon]
F5 BIG-IP	F5 - HD, F5 Load Balancer	F5 BIG-IPs [cmdb_ci_lb_bigip]
NetScaler	Netscaler Load Balancer	Citrix NetScalers [cmdb_ci_lb_netscaler]
Radware - AppDirector	AppDirector Load Balancer	Radware Load Balancer [cmdb_ci_lb_radware]
HAProxy load balancer discovery with patterns	HA Proxy	HAProxy Load Balancer [cmdb_ci_lb_haproxy]

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Load balancer discovery, Service Mapping, and customized CI fields

Discovery collects information about the install status and operational status of load balancers. If you are using Discovery with Service Mapping, do not customize the [operational_status] or [install_status] fields. By default, Service Mapping ignores all host CIs for which the value of the *[operational_status]* field is not *1* (Operational) or the value of the status

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

[install_status] field is *100* (absent). For additional information, see Preparing customized ServiceNow deployments to work with Service Mapping [KB0647574] in the HI Knowledge Base.

A10 load balancer discovery

Discovery and Service Mapping uses patterns to collect information about A10 load balancers.

Note: Probe-to-pattern migration is supported for the New York release and later. See the knowledge article KB0694477 for more information.

The following patterns support the A10 load balancer discovery:

- A10 Load Balancer
- A10 Load Balancer SSH
- A10

The A10 IP Source NAT (Network Address Translator) is an extension of the A10 Load Balancer Pattern. The IP Source NAT Pattern extension discovers IP Source NAT pools. This extension is available on Store from the February 2023 release, Store version 1.0.98.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Credentials

Configure SNMP credentials or (optionally) SSH credentials.

Discovery and Service Mapping Patterns application

To discover A10 load balancer partitions, deploy Discovery and Service Mapping Patterns application 1.0.98 release from the ServiceNow Store.

Data collected by Discovery during horizontal discovery

Table and field	Description
A10 Load Balancer [cmdb_ci_lb_a10]	
Serial number [serial_number]	Serial number of the load balancer.
Load Balancer Interface [cmdb_ci_lb_interface]	
Name [name]	Name of the load balancer interface.
MAC Address [mac_address]	MAC address of the load balancer interface, in the following format: 55:0d:f9:98:3i:ae.
Operational status [operational_status]	Operational status of the load balancer interface.
Load Balancer Pool [cmdb_ci_lb_pool]	
Name [name]	Name of the load balancer pool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Load balancing method [load_balancing_method]	The method used by the load balancer.
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	
Name [name]	Name of the pool member.
IP Address [ip_address]	IP address of the pool member host.
Service port [service_port]	Port used by the pool member.
Operational status [operational_status]	Operational status of the load balancer pool members.
Pool [pool]	Pool member.
Load Balancer Services [cmdb_ci_lb_service]	
Name [name]	Name of the load balancer, in the following format: <LB name>:<LB port>.
IP Address [ip_address]	IP address of the FQDN.
Port [port]	Port of the load balancer.
Operational status [operational_status]	Operational status of the load balancer service.
Install Status [install_status]	Install status of the load balancer service.
Pool [pool]	Pool member of the load balancer service.
DNS [cmdb_ci_dns_name]	
Name [name]	DNS name.
IP Address [ip_address]	DNS IP address.
A10 Load Balancer Partition [cmdb_ci_lb_a10_partition]	
Name [name]	Name of the partition.
Index [index]	Partition ID. An admin assigns this ID in the process of creating a partition.
Max Aflex [max_aflex]	Maximum number of aFleX scripts that can belong to the partition.
Admin Count [admin_count]	Number of admins assigned to the partition.
L3V [l3v]	Value indicating whether Layer 2/3 virtualization is enabled for the partition.
Network Adapter [cmdb_ci_network_adapter]	
Operational status [operational_status]	Operational status of the network adapter.
Mac address [mac_address]	MAC address of the network adapter.
Alias [alias]	Alias of the network adapter.
IP Address [ip_address]	IP address of the network adapter.
Name [name]	Name of the network adapter.
Netmask [netmask]	Netmask of the network adapter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
PrimaryIP [primaryIP]	PrimaryIP address of the network adapter.
IP Address [cmdb_ci_ip_address]	
IP Address [ip_address]	IP address.
Netmask [netmask]	Netmask to which the IP address belongs.
Version [version]	IP version.
Operational Status [operational_status]	Operational status of the IP address.
SNAT IP Pool [cmdb_ci_snat_ip_pool]	
NAT Pool Failed [nat_pool_failed]	The total misses in the pool.
HA Group [ha_group]	High Availability group number associated with the NAT pool.
NET Mask [net_mask]	Netmask of the NAT pool address range.
NAT Pool Total Free [nat_pool_total_free]	The total of free addresses in the pool.
NAT Pool Total Used [nat_pool_total_usage]	The total of used addresses in the pool.
NAT Pool Port Usage [nat_pool_total_used]	The total port usage in the pool.
NAT Pool End Address [nat_pool_end_address]	The IP NAT pool end address (IPv4 or IPv6).
NAT Pool Start Address [nat_pool_start_address]	The IP NAT pool start address (IPv4 or IPv6).
Object ID [object_id]	A unique identifier for the NAT pool, assigned by the A10 device.
Configuration Item [cmdb_ci]	References the A10 Load Balancer [cmdb_ci_lb_a10] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dependency Views showing A10 load balancer and connected components and applications

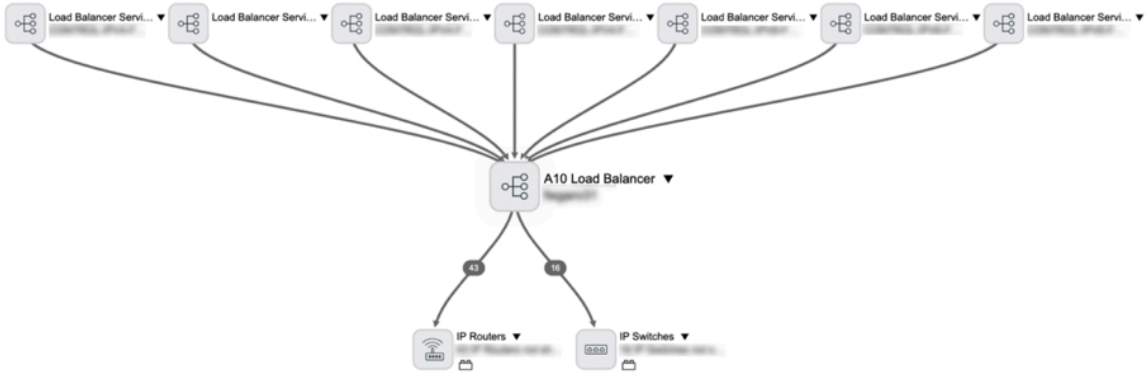

CI relationships

These relationships are created to support A10 load balancer discovery:

CI	Relationship	CI
Router Interface [dscy_router_interface]	Uses::Used by	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Interface [cmdb_ci_lb_interface]	Owns::Owned by	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	Owns::Owned by	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Services [cmdb_ci_lb_service]	Used by::Uses	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Services [cmdb_ci_lb_service]	Runs on::Runs	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Pool [cmdb_ci_lb_pool]	Owns::Owned by	Load Balancer Pool [cmdb_ci_lb_pool]
DNS Name[cmdb_ci_dns_name]	Used by::Uses	IP Address [cmdb_ci_ip_address]
Network Adapter [cmdb_ci_network_adapter]	Owns::Owned by	Load Balancer Pool [cmdb_ci_lb_pool]
Load Balancer Pool [cmdb_ci_lb_pool]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
A10 Load Balancer [cmdb_ci_lb_a10]	Owns::Owned by	Source NAT IP Pool [cmdb_ci_snat_ip_pool]

Data collected by Service Mapping during top-down discovery

Service Mapping discovers the connections between the A10 load balancers and applications in the deployment.

Cisco ACE load balancer discovery

Discovery identifies and classifies information about ACE load balancers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Credentials

Configure SNMP credentials and SSH credentials.

Classifier, pattern, and probe

Classifier	Trigger probes	Pattern
ACE Load Balancer	Horizontal discovery probe: launches patterns	ACE Load Balancer by SSH (formerly ACE Load Balancer)

Data collected

The following data is collected on the Ace Load Balancers [`cmdb_ci_lb_ace`] table.

Label	Field Name
Serial Number	serial_number
Operational Status	operational_status

The following data is collected on the Load Balancer Services [cmdb_ci_lb_service] table.

Label	Field Name
Name	name
IP address	ip_address
Port	port
Operational status	operational_status
Install status	install_status
Short description	short_description

The following data is collected on the DNS [cmdb_ci_dns_name] table.

Label	Field name
Name	name
IP address	ip_address

The following data is collected on the Load Balancer Pool [cmdb_ci_lb_pool] table.

Label	Field name
Name	name
Context	context
Load balancing method	load_balancing_method

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following data is collected on the Load Balancer Pool Members [cmdb_ci_lb_pool_member] table.

Label	Field name
Name	name
IP address	ip_address
Service port	service_port
Pool	pool
Context	context

The following data is collected on the Serial Number [cmdb_serial_number] table.

Label	Field name
Serial number	serial_number
Serial number type	serial_number_type
Valid	valid

The following data is collected on the Load Balancer Interface table [cmdb_ci_lb_interface] table:

Label	Field name
Name	name
MAC address	mac_address
Operational status	operational_status

Cisco GSS load balancer discovery

Discovery of Cisco GSS load balancers is performed by both SNMP and SSH.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Classifiers

Discovery uses the Cisco GSS Load Balancer classifier, which contains the condition: `sysdescr contains Global Site Selector` to classify the device.

Probes

Discovery uses several probes, including an SNMP probe to discover the specific GSS device and a **Serial Number** probe to identify the serial number of the device. See the table below for the list of related probes.

Warning: You must have SNMP and SSH credentials configured correctly for the GSS load balancers on your network.

The following probes are triggered after classification:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	Description	Type
Cisco GSS - Identity	A multiprobe that identifies the GSS device and serial number.	MultiProbe
Cisco GSS - Network	This is included in the Cisco GSS - Identify probe.	Java-SNMP
Cisco GSS - Serial Number	This is included in the Cisco GSS - Identify probe.	JavaScript
Cisco GSS - Get Domains	A probe that discovers global domain names corresponding to the device.	JavaScript

Tables

Discovery creates a record for each GSS device in the Cisco GSS [cmdb_ci_lb_cisco_gss] table. Domains are populated in the DNS Names [cmdb_ci_dns_name] table. Host names and IP addresses are stored in the IP Address to DNS Name [cmdb_ip_address_dns_name] table. Service Mapping uses this information. You can see the DNS name information on the **DNS Names for CIs** related list of the Load Balancer form.

Note: The Cisco GSS pattern is available by default to use with the Pattern Designer.

Cisco CSS load balancer

Discovery of Cisco CSS load balancers is performed by SNMP.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Classifiers

Discovery uses the Cisco CSS Load Balancer classifier, which contains the OID Classification: 1.3.6.141.9.9.368.4.5.

Probes

The following probes are triggered:

Probe	Description	Type
SNMP - Load Balancer - Identity	A multiprobe that identifies load balancers.	Java-SNMP
Cisco CSS - Get Services	A Java probe that includes the Cisco CSS sensor to identify services defined on the load balancer in the Load Balancer Services [cmdb_ci_lb_service] table. For every service, Discovery populates **Name**, **ip_address**, and **port**.	Java-SNMP

Tables

Discovery creates a record for each CSS device in the Cisco CSS [cmdb_ci_lb_cisco_css] table, which includes the device's **Name**, **Model**, **Serial Number**, **Manufacture**, and **NIC** (ip_address). It also creates a record for each service in the Load Balancer Services [cmdb_ci_lb_service] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

.

Citrix NetScaler load balancer discovery

Discovery and Service Mapping find Citrix NetScaler load balancers including Server Load Balancing (GSLB).

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Requirements and support

Important: Discovery treats load balancers as licensable entities and attempts to discover them primarily using SNMP. If a load balancer in your system, running on a Linux host, has SNMP and SSH ports open, Discovery might classify it based on the SSH port. This classification has priority over SNMP. To ensure that Discovery properly classifies your hardware load balancers, create a Discovery behavior for load balancers that includes SNMP but not SSH. Software load balancers are treated as applications.

Virtual NetScaler load balancers are not fully supported. Discovery supports local server load balancers and GSLB.

Credentials

- Configure SNMP credentials. The SNMP community string requires read-only access to the Citrix NetScaler MIB (`1.3.6.1.4.1.5951`).

  Note: You also have the option of using SSH credentials. The SSH user requires permission to run NetScaler `show` commands only. However, if the payload size of the data that's returned to the instance is large, the discovery process can fail. In this case, use SNMP credentials instead.

- (Optional) When using the NS.CONF pattern, provide a user with read-only access to the *NS.conf* configuration file if you don't want to provide SNMP or SSH credentials.
- Use basic authentication credentials for discovery via REST. Configure the credential as a credential alias. The NITRO API user requires read-only access to the following resources:
  - `lbvserver`
  - `csvserver`
  - `service`
  - `servicegroup`
  - `interface`
  - `lbvserver_binding`
  - `csvserver_binding`
  - `nspartition`
  - `nshardware`
  - `nshostname`
  - `nsconfig`
  - `hanode`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Classifier, pattern, and probes

Classifier	Trigger probes	Patterns

NetScaler Load Balancer	• Horizontal discovery probe: launches patterns (see the Patterns column)
	• NetScaler - Network Path Information (multiprobe)
	  ○ NetScaler - ARP Information
	  ○ NetScaler - Route Information
	• SNMP - NetScaler - Identity
	• SNMP - NetScaler - Identity - Serial
	• SNMP - NetScaler - Identity Info
	• SNMP - NetScaler - System
		• NetScaler Load Balancer
		• NetScaler Load Balancer NS.CONF
		• NetScaler Load Balancer SSH
		• Netscaler Loadbalancer NITRO

To use probes for the horizontal discovery, make sure that the *NetScaler - Network Path Information* multiprobe is active on the **Trigger probes** related list for the *NetScaler Load Balancer* classifier. Then deactivate or remove the *Horizontal discovery* probe from the **Trigger probes** related list.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Note: When using probes and sensors and the system properties *glide.discovery.hostname.snmp_trusted* and/or *glide.discovery.hostname.dns_nbt_trusted* are set to false, the system uses *dns_name* as the load balancer name. When using the Netscaler Load Balancer pattern, the system currently doesn't recognize these system properties when they are set to false.

Connections discovered by Service Mapping during the top-down discovery

Service Mapping performs the top-down discovery of the NetScaler Global Server Load Balancer in the context of application services. It discovers configured URLs of the NetScaler Global Server Load Balancer.

Data collected by Discovery during horizontal discovery

Table and field	Description
DNS name [cmdb_ci_dns_name]	
Name [name]	The name or the DNS name alias of the NetScaler Global Server Load Balancing (GSLB).
IP address [ip_addresses]	Management IP address of the NetScaler Global Server Load Balancing (GSLB).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Port [port]	The DNS port that Global Server Load Balancer listens to.
Citrix NetScaler [cmdb_ci_lb_netscaler]	
Serial Number [serial_number]	The serial number of the NetScaler device and its current operational status.
Operational Status [operational_status]	
Load Balancer Pool [cmdb_ci_lb_pool]	
Name [name]	The name of the pool and the method of load balancing it enables.
Load Balancing Method [load_balancing_method]	
Load Balancer Pool Member [cmdb_ci_lb_pool_member]	
Name [name]	The name, IP address and service port that each pool member uses.
IP Address [ip_address]	
Service Port [service_port]	
Load Balancer VLAN [cmdb_ci_lb_vlan]	
Name [name]	The name of the VLAN and the tag that is associated with this VLAN to identify it.
Tag [tag]	
Load Balancer Services [cmdb_ci_lb_service]	
Name [name]	Details about the service running on the load balancer, including
IP Address [ip_address]	
Port [port]	
Operational Status [operational_status]	
Install Status [install_status]	
Hit count [hit_count]	
Exit Interface Routing Rules [dscy_route_interface]	
Router Interface [router_interface]	Details about the routing rules for the exit interface, including the interface on the router that ...
Configuration item [cmdb_ci]	
Destination IP network [dest_ip_network]	
Next Hop Routing Rules [dscy_route_next_hop]	
Next hop IP address [router_interface]	Details about the routing rules on the load balancer, including the IP addresses of the next hop router on the network. The Configuration item field is the router which is the next hop device.
Router Interface [router_interface]	
Configuration item [cmdb_ci]	
Destination IP network [dest_ip_network]	

cmdb_ci_lb_netscaler

Column name	Type	Reference
ip_address	IP Address (Validated IPV4, IPV6)	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

cmdb_ci_lb_netscaler (continued)

name	String
operational_status	Integer
serial_number	String

Data collected during Netscaler Loadbalancer NITRO.

cmdb_serial_number

Column label	Column name	Type	Reference
Configuration Item	cmdb_ci	Reference	Configuration Item (Load Balancer)
Serial Number	serial_number	String	
Serial Number Type	serial_number_type	String	
Valid	valid	True/False	

cmdb_ci_lb_vlan

Column label	Column name	Type	Reference
Load balancer	load_balancer	Reference	Configuration Item (Load Balancer)
Name	name	String	
Tag	tag	Integer	

cmdb_ci_lb_interface

Column label	Column name	Type	Reference
Load balancer	load_balancer	Reference	Configuration Item (Load Balancer)
Name	name	String	
Operational status	operational_status	Integer	
MAC Address	mac_address	String	

cmdb_ci_lb_service

Column label	Column name	Type	Reference
IP Address	ip_address	IP Address (Validated IPV4, IPV6)	
Load balancer	load_balancer	Reference	Configuration Item (Load Balancer)
Name	name	String	
Port	port	Integer	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

cmdb_ci_lb_pool

Column label	Column name	Type	Reference
Load balancer	load_balancer	Reference	Configuration Item (Load Balancer)
Name	name	String	

cmdb_ci_lb_pool_member

Column label	Column name	Type	Reference
IP Address	ip_address	IP Address (Validated IPV4, IPV6)	
Load balancer	load_balancer	Reference	Configuration Item (Load Balancer)
Name	name	String	
Pool	pool	Reference	Load Balancer Pool
Service port	service_port	Integer	

To see the DNS names and alases for a NetScaler load balancer, you must configure the form and add the **DNS Names for CIs** related list. DNS names also include aliases.

Note: Discovery populates load balancer interface information that is necessary to map the network path in Service Mapping.

Configuration item (CI) Relationships

These relationships are created to support NetScaler Global Server Load Balancer discovery:

CI	Relationship	CI
DNS name [cmdb_ci_dns_name]	Uses:Used by	IP address [cmdb_ci_ip_addresses]

These relationships are created to support NetScaler Loadbalancer NITRO:

CI	Relationship	CI
Netscaler [cmdb_ci_lb_netscaler]	Owns::Owned by	interfaces [cmdb_ci_lb_interface]
Netscaler [cmdb_ci_lb_netscaler]	Owns::Owned by	vlan [cmdb_ci_lb_vlan]
Netscaler [cmdb_ci_lb_netscaler]	Owns::Owned by	service groups [cmdb_ci_lb_pool]
service groups	Owns::Owned by	lbvservers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI	Relationship	CI
[cmdb_ci_lb_pool]		[cmdb_ci_lb_pool_member]
services [cmdb_ci_lb_service]	Runs on::Runs	Netscaler [cmdb_ci_lb_netscaler]

NetScaler device sysName

To verify the NetScaler device sysName, check the snmp mib and verify the sysName value by running the command: *`show snmp mib`*.

To change the value of the sysName, run the command: *`set mib -name <DNS_NAME or NAME_OF_THE_NETSCALER_LB>`* to set the name of the SNMP MIB.

F5 BIG-IP load balancer discovery

Discovery and Service Mapping can find F5 BIG-IP load balancers via SNMP, SSH, and through the REST API.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Important: Discovery treats load balancers as licensable entities and attempts to discover them primarily using SNMP. If a load balancer in your system, running on a Linux host, has SNMP and SSH ports open, Discovery might classify it based on the SSH port. This classification has priority over SNMP. To ensure that Discovery properly classifies your hardware load balancers, create a Discovery behavior for load balancers that includes SNMP but not SSH. Software load balancers are treated as applications.

You can download VMware images of BIG-IP with a free 90-day trial from https://www.f5.com/trial.

For F5 Global Traffic Manager (GTM) BIG-IP load balancers, Discovery resolves DNS names for the F5 GTM hardware and all associated servers that receive distributed traffic. To view this data, navigate to **All** > **Configuration** > **Load Balancers** > **LB Hardware**, open the F5 load balancer record, and select the **DNS Names for CIs** related list. Service Mapping uses this information to map F5 relationships.

Note: If your F5 BIG-IP device is part of a failover cluster, it can be associated with two cluster nodes. In this case, one of the nodes appears as Operational in the **Operational Status** field, and the other appears as Non-Operational.

Prerequisites

- Verify that the following applications are up to date:
  - Discovery and Service Mapping Patterns
  - Visibility Content
  - CMDB CI Class Models
- To successfully discover the load balancer pool members, set the **display service names** option (*`bigpipe.displayservicenames`*) of the load balancer to **false** on the F5 load balancer.
- For F5 Load Balancer pattern:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Configure SNMP credentials on the ServiceNow AI Platform.
- (Optional) If there are iRules or SNMP community credentials aren't enough for discovering outgoing connections, configure SSH credentials.

  Note: If you don't want to use SSH credentials, you can use the REST API to create a connection to F5 BIG-IP devices.

  Service Mapping uses the SSH credentials to retrieve connections that aren't from CMDB. Discovering connections using the SSH protocol is a failover mechanism for the SNMP-based discovery.

- For F5 Load Balancer SSH pattern:
  - Configure SSH credentials.
  - Verify permissions for the following SSH commands. In newer versions of F5 BIG-IP, only A-type GTM pools are supported.
    - `list auth partition all`
    - `list sys global-settings`
    - `show sys hardware`
    - `list sys db failover.state`
    - `list net self one-line`
    - `show sys mac-address | grep interface`
    - `show net interface field-fmt`
    - `show net self | grep -e 'Mac Address' 'IP:'`
    - `list ltm pool | grep -e 'pool' 'mode'`
    - `list ltm pool /<partition_name>/* | grep -e 'pool' 'mode'`
    - `show ltm pool | grep -e 'Ltm::' 'Total Connections'`
    - `show ltm pool /<partition_name>/* |grep -e 'Ltm::' 'Total Connections'`
    - `list sys db bigpipe.displayservicenames | grep value`
    - `show /net vlan | grep -e 'Interface Name' 'Mac' 'Tag'`
    - `show ltm virtual detail | grep -e 'Server:' 'State' 'Pool:' 'Destination' 'Total Connections'`
    - `show ltm virtual detail /<partition_name>/* | grep -e 'Server:' 'State' 'Pool:' 'Destination' 'Total Connections'`
    - `show gtm wideip | grep -e 'WideIp' 'State'`
    - `list wideip a /<partition_name>/* | grep -e 'WideIp' 'State'`
    - `show gtm wideip all | grep -e 'WideIp' 'State'`
    - `list gtm wideip`
    - `list gtm wideip all`
    - `list gtm pool one-line`
    - `list gtm pool a one-line`
    - `list gtm pool /<partition_name>/* one-line`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `list gtm pool a /<partition_name>/* one-line`
- `list gtm server one-line`

- For F5REST pattern:
  - For horizontal discovery using Discovery, provide read-only permissions to run the following APIs:
    - `"https://" + $ipAddress + "/mgmt/tm/ltm/pool"`
    - `"https://" + $ipAddress + "/mgmt/tm/sys/global-settings"`
    - `"https://" + $ipAddress + "/mgmt/tm/sys/hardware"`
    - `"https://" + $ipAddress + "/mgmt/tm/sys/failover"`
    - `"https://" + $ipAddress + "/mgmt/tm/net/self"`
    - `"https://" + $ipAddress + "/mgmt/tm/ltm/virtual?$select=destination,mask,vlans"`
    - `"https://" + $ipAddress + "/mgmt/tm/net/interface"`
    - `"https://" + $ipAddress + "/mgmt/tm/net/interface/stats"`
    - `"https://" + $ipAddress + "/mgmt/tm/net/trunk"`
    - `"https://" + $ipAddress + "/mgmt/tm/net/vlan"`
    - `"https://" + $ipAddress + "/mgmt/tm/net/vlan/"+<vlan id>+"/interfaces"`
    - `"https://" + $ipAddress + "/mgmt/tm/gtm/wideip"`
    - `"https://" + $ipAddress + "/mgmt/tm/cm/traffic-group/"`
    - `"https://" + $ipAddress + "/mgmt/tm/cm/device"`
    - `"https://" + $ipAddress + "/mgmt/tm/gtm/server"`

    - For BIG-IP starting with version 12:
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/pool/a"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/pool/aaaa"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/pool/cname"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/pool/mx"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/pool/naptr"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/pool/srv"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/wideip/a"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/wideip/aaaa"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/wideip/cname"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/wideip/mx"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/wideip/naptr"`
      - `"https://" + $ipAddress + "/mgmt/tm/gtm/wideip/srv"`
  - Create basic authentication credentials.
- For F5 REST OAuth 2.0 authentication pattern:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Verify read-only permission for the following APIs:
  - "https://" + $ipAddress + "/mgmt/shared/authn/login"
  - "https://" + $ipAddress + "/mgmt/tm/ltm/pool"
  - "https://" + $ipAddress + "/mgmt/tm/sys/global-settings"
  - "https://" + $ipAddress + "/mgmt/tm/sys/hardware"
  - "https://" + $ipAddress + "/mgmt/tm/sys/failover"
  - "https://" + $ipAddress + "/mgmt/tm/net/self"
  - "https://" + $ipAddress + "/mgmt/tm/ltm/virtual?$select=destination,mask,vlans"
  - "https://" + $ipAddress + "/mgmt/tm/net/interface"
  - "https://" + $ipAddress + "/mgmt/tm/net/interface/stats"
  - "https://" + $ipAddress + "/mgmt/tm/net/trunk"
  - "https://" + $ip_address + "/mgmt/tm/net/vlan"
  - "https://" + $ipAddress + "/mgmt/tm/net/vlan/"+<vlan id>+"/interfaces"
  - "https://" + $ipAddress + "/mgmt/tm/gtm/wideip"
  - "https://" + $ipAddress + "/mgmt/tm/cm/traffic-group/"
  - "https://" + $ipAddress + "/mgmt/tm/cm/device"
  - "https://" + $ipAddress + "/mgmt/tm/gtm/server"

  - For BIG-IP starting with version 12:
    - "https://" + $ipAddress + "/mgmt/tm/gtm/pool/a"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/pool/aaaa"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/pool/cname"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/pool/mx"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/pool/naptr"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/pool/srv"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/a"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/aaaa"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/cname"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/mx"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/naptr"
    - "https://" + $ipAddress + "/mgmt/tm/gtm/wideip/srv"
- Create basic authentication credentials.
- Create a credential alias for the basic authentication credential.
- Create a serverless discovery schedule.

- For top-down discovery using Service Mapping, provide permissions to run the following APIs:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `https://" get_attr {"managementIP"}"/mgmt/tm/ltm/"get_attr {"irule"}`
- `https://" get_attr {"managementIP"}"/mgmt/tm/ltm/virtual/~Common~"get_attr {"vip_name"}`
- `https://" get_attr {"computer_system.managementIP"}"/mgmt/tm/ltm/pool/"get_attr {"pool_name"}"/members?`

- (Optional) For top-down discovery with Service Mapping, configure applicative credentials.
- (Optional) For top-down discovery with Service Mapping, provide a user with permissions to run the following commands:
  - bigpipe commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 9)
  - bigpipe and Traffic Management Shell (TMSH) commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 10)
  - Traffic Management Shell (TMSH) commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 11)
  - Traffic Management Shell (TMSH) advanced commands (for BIG-IP LTM F5 or BIG-IP GTM F5 version 10, 11, and 12)
  - The *`show cm traffic-group - get command`* for discovering F5 BIG-IP Device Service Clustering
- (Optional) For top-down discovery with Service Mapping, give the user elevated rights to be able to run Service Mapping commands requiring a privileged user.

F5 cluster discovery

The horizontal discovery process can find F5 clusters. To discover clusters, the credentials must have read permission to traffic groups. For SNMP, this is the regular read permission. For SSH, you need to have **Traffic Management Shell (TMSH shell)** enabled and accessible.

Note: Because F5 clusters share the same virtual IPs and addresses for all devices, the horizontal discovery process cannot find specific IP addresses for each device.

Classifier, patterns, and probes

Discovery uses these classifiers, probes, and pattern for F5 discovery:

Classifiers	Trigger probes	Patterns
F5 BIG-IP Load Balancer	Horizontal discovery probe: launches patterns	These patterns are used for standard horizontal discovery including identifying clusters: • F5 Load Balancer • F5 Load Balancer SSH • F5REST • F5 REST OAuth 2.0 authentication

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifiers	Trigger probes	Patterns
		Note: The F5 Load Balancer and F5 Load Balancer SSH patterns include the extension section that discovers F5 certificates.
F5 HTTP	The F5 HTTP classifier does not use a trigger probe. It triggers the F5 Load Balancer REST pattern to perform data collection via REST.	These patterns use iRules to get information via REST. Service Mapping uses these patterns for top-down discovery: • F5 Load Balancer REST • F5 BigIP GTM • F5 BigIP LTM

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the F5 patterns.

F5 BIG-IP [cmdb_ci_lb_bigip]

Field	Description
Name [name]	The name of the load balancer CI.
Serial number [serial_number]	The serial number of the load balancer hardware.
Operational status [operational_status]	The operational status of the load balancer.

Load Balancer Interface [cmdb_ci_lb_interface]

Field	Description
Name [name]	The name of the load balancer interface.
MAC Address [mac_address]	The MAC address of the interface.
Operational status [operational_status]	The operational status of the interface.

Load Balancer Pool [cmdb_ci_lb_pool]

Field	Description
Name [name]	The name of the load balancer pool. GTM pool names end with the suffix **- GTM** to differentiate them from LTM pools with the same name.
Load balancing method [load_balancing_method]	The method used to distribute traffic across pool members.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer Pool Member [cmdb_ci_lb_pool_member]

Field	Description
Name [name]	The name of the pool member.
IP Address [ip_address]	The IP address of the pool member.
Service port [service_port]	The port on which the pool member receives traffic.
Pool [pool]	The load balancer pool that this member belongs to.

Load Balancer VLAN [cmdb_ci_lb_vlan]

Field	Description
Name [name]	The name of the VLAN.
Tag [tag]	The tag that identifies the VLAN.
MAC Address [mac_address]	The MAC address associated with the VLAN.

Load Balancer Service [cmdb_ci_lb_service]

Field	Description
Name [name]	The name of the load balancer service.
IP Address [ip_address]	The IP address of the load balancer service.
Port [port]	The port on which the load balancer service receives traffic.
Operational status [operational_status]	The operational status of the load balancer service.
Hit count [hit_count]	The number of connections the load balancer service has received.

DNS Name [cmdb_ci_dns_name]

Field	Description
Name [name]	The DNS name associated with the load balancer or its GTM WideIPs.
IP Address [ip_address]	The IP address associated with the DNS name.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	The name of the network adapter.
IP Address [ip_address]	The IP address of the network adapter.
Netmask [netmask]	The netmask of the network adapter.
MAC Address [mac_address]	The MAC address of the network adapter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cluster [cmdb_ci_cluster]

Field	Description
Name [name]	The name of the cluster.
Cluster status [cluster_status]	The status of the cluster.
Cluster ID [cluster_id]	The unique identifier of the cluster.
IP Address [ip_address]	The IP address of the cluster.

Cluster Node [cmdb_ci_cluster_node]

Field	Description
Name [name]	The name of the cluster node.
Cluster ID [cluster_id]	The unique identifier of the cluster this node belongs to.
Cluster status [cluster_status]	The status of the cluster node.

CI relationships

The following relationships are created to support F5 BIG-IP load balancer discovery.

CI	Relationship	CI
Cluster [cmdb_ci_cluster]	Hosts::Hosted on	F5 BIG-IP [cmdb_ci_lb_bigip]
Cluster Node [cmdb_ci_cluster_node]	Hosted on::Hosts	F5 BIG-IP [cmdb_ci_lb_bigip]
Cluster Node [cmdb_ci_cluster_node]	Cluster of::Cluster	Cluster [cmdb_ci_cluster]
F5 BIG-IP [cmdb_ci_lb_bigip]	Owns::Owned by	DNS Name [cmdb_ci_dns_name]

Connections discovered by Service Mapping during the top-down discovery

Service Mapping performs the top-down discovery of the GTM BIG-IP load balancer in the context of application services. It discovers configured URLs of the GTM BIG-IP load balancers.

Create a serverless schedule for F5 BIG-IP REST OAuth 2.0 authentication-based discovery

Create a serverless discovery schedule to discover F5 BIG-IP load balancers through REST OAuth 2.0 authentication (tokens).

Before you begin

- Verify you have the IP address for the F5 BIG-IP load balancer.
- Create a basic authentication credential. For more information, see Basic authentication credentials.
- Create a alias for the basic authentication credential. For more information, see Credential aliases for Discovery.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Discovery** > **Discovery Schedules**.
2. Create the discovery schedule record.
   a. Select **New**.
   b. On the form, fill in the fields.

Discovery Schedule New record form

Field	Description
Name	Unique name for this discovery schedule.
Discover	Scan type, which should be Serverless.
MID server	Name of the MID Server to use for this schedule.
Active	Enables this schedule for discovery.

   c. Select **Submit**.
3. Create the execution pattern.
   a. In the Discovery Schedules page, select the record you created.
   b. In the **Serverless Execution Patterns** tab, select **New**.
   c. On the form, fill in the fields.

Serverless Execution Pattern New record form

Field	Description
Name	Descriptive name for this record.
Pattern	Pattern to be used for this schedule, which should be the F5 REST OAuth 2.0 authentication pattern.
Proxy Host	Fully qualified domain name of the machine on which you're installing the proxy server, which should be Global.
Active	Enables this schedule for discovery.

   d. Select **Submit**.
4. Set the pattern launcher parameters.
   a. In the **Discovery Pattern Launcher Parameters** tab, select the record you created.
   b. On the form, fill in the fields.

Parameter	Value
ip_address	IP address of the F5 BIG-IP load balancer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Value
credential_alias	Name of the credential alias you created.

c. Select **Submit**.

What to do next

Either execute discovery immediately by selecting **Discover now** or wait until the predefined schedule triggers the discovery.

HAProxy load balancer discovery using probes

Discovery of HAProxy Community edition load balancers is performed by SSH.

Note: Starting with Discovery and Service Mapping Patterns 1.15.0 version, pattern-based discovery is the default method for discovering HAProxy load balancers.

- For information about enabling probe-based discovery, see Enable HAProxy load balancer discovery with probes.
- For information about the HA Proxy pattern, see HAProxy load balancer discovery with patterns.
- For information about Probe to Pattern migration see the knowledge article KB0694477.

HAProxy is an open-source load balancer that can manage any TCP service. It is particularly suited for HTTP load balancing because it supports session persistence and Layer 7 processing. Discovery supports HAProxy for HTTP load balancing. TCP load-balancing is not supported.

Consider the following requirements for discovering the HAProxy:

- The HAProxy software is installed and running on a Linux server.
- The MID Server is deployed to explore the server and the MID Server has access to the server HAProxy configuration file.
- The configuration probe checks for the `haproxy.cfg` file using one of the following methods:
  - Using the *`f`* parameter for the HAProxy process output.
  - Using the default `/etc/haproxy/haproxy.cfg` path.
- The HAProxy probes require credentials and execute privileges to run commands.

Discovery uses the **Unix - Active Processes** probe to identify an HAProxy load balancer when the name of the process is *`haproxy`*. If this criterion matches, a record is created in the HAProxy Load Balancers [cmdb_ci_lb_haproxy] table if one does not already exist for that running process.

The following probes are triggered after classification:

HAProxy probes

Probe	Description	Commands
HAProxy – Version	The sensor of this probe populates the HAProxy version in the HAProxy Load Balancers [cmdb_ci_lb_haproxy] table.	haproxy
HAProxy – Get Configuration	The sensor of this probe populates additional information in the HAProxy Load Balancers [cmdb_ci_lb_haproxy] table. The probe also populates information in the Load Balancer Service [cmdb_ci_lb_service], Load Balancer Pool [cmdb_ci_lb_pool], Load	echo, sed, cut, grep, egrep (within

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HAProxy probes (continued)

Probe	Description	Commands
	Balancer Pool Member [cmdb_ci_lb_pool_member], Load Balancer Interface [cmdb_ci_lb_interface], and Load Balancer Application [cmdb_ci_lb_appl] tables.	the Borne shell script)

In addition to populating the data, the following relationships records are created in CI Relationships [cmdb_rel_ci] table:

- The records in the cmdb_ci_lb_appl table run on the cmdb_ci_web_server table records.
- The records in the cmdb_ci_lb_service table use the cmdb_ci_lb_pool table records.
- The records in the cmdb_ci_pool table are used by the cmdb_ci_service table records.
- The records in the cmdb_ci_pool table are members of the cmdb_ci_pool_member table records.
- The records in the cmdb_ci_pool_member table are members of the cmdb_ci_pool table records.

Enable HAProxy load balancer discovery with probes

If you prefer not to start using the default HAProxy pattern-based discovery, enable probe-based HAProxy discovery to continue using it instead.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **Processes**.
2. In the **Name** field, search for `HAProxy`.
3. Select the HAProxy classifier.
4. In the **Triggers probe** related list, select **Edit**.
5. Remove **Horizontal Pattern** from the Triggers probes List.
6. From the Collection list, add **HAProxy - Get Configuration** and **HAProxy - Version** to the Triggers probe list.
7. Select **Save**.

What to do next

Schedule a horizontal discovery to discover HAProxy load balancer with probes. For more information, see Schedule a horizontal discovery.

Related topics

HAProxy load balancer discovery using probes

NGINX load balancer discovery

Discovery of NGINX load balancers is performed by SSH.

NGINX is an open source web server with a load balancer. Discovery identifies the web server and information related to the load balancer.

Consider the following requirements for discovering NGINIX servers:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Ensure that the NGINX software is installed and running on the server.
- Grant the MID Server has access to the NGINX configuration file, which is `/etc/nginx/nginx.conf` by default.
- Enable secure shell (SSH) commands.
- The NGINX probes require credentials and execute privileges to run commands.

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

The Nginx Process Classifier detects a running process that matches the following criteria during the exploration of a UNIX server:

- The name of the process starts with *`nginx`*.
- The name of the process contains *`master`*.

Both of these conditions must match:

- A record is created in the Web Server [cmdb_ci_web_server] table.
- A **Runs on** relationship is created in the CI Relationship [cmdb_rel_ci] table for the Linux Server [cmdb_ci_linux_server] table and the Web Server [cmdb_ci_web_server] table.

The following probes are triggered after classification:

NGINX probes

Probe	Description	Commands
Nginx – Version	This probe contains a Bourne shell script. It determines the version of NGINX and populates the Web Server [cmdb_ci_web_server] table.	nginx
Nginx – Get Configuration	This probe contains a Bourne shell script and an argument that determines the path of the NGINX configuration file. The probe identifies configuration parameters based on keywords within the configuration file and returns them as a single payload result.	echo, sed, cut, grep, egrep (within the Bourne shell script)

In addition to populating the data, the following relationships records are created in CI Relationships [cmdb_rel_ci] table:

- The records in the cmdb_ci_web_server table run on the cmdb_ci_linux_server table records.
- The records in the cmdb_ci_lb_service table use the cmdb_ci_lb_pool table records.
- The records in the cmdb_ci_pool table are used by the cmdb_ci_service table records.
- The records in the cmdb_ci_pool table are members of the cmdb_ci_pool_member table records.
- The records in the cmdb_ci_pool_member table are members of the cmdb_ci_pool table records.

Radware Alteon load balancer discovery

Discovery identifies and classifies information about Alteon load balancers.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credentials

Configure SNMP credentials.

Classifier, pattern, and probe

Classifier	Trigger probe	Pattern
Alteon Load Balancer	Horizontal pattern	Alteon Load Balancer by SNMP (formerly Alteon Load Balancer)

Data collected

The following data is collected on the Alteon `[cmdb_ci_lb_alteon]` table.

Label	Field Name
Serial Number	serial_number
Operational Status	operational_status

The following data is collected on the Load Balancer Interface table [cmdb_ci_lb_interface] table:

Label	Field name
Name	nam
MAC address	mac_address
Operational status	operational_status
Index	index

The following data is collected on the Load Balancer Pool [cmdb_ci_lb_pool] table.

Label	Field name
Name	name
Index	index
Load balancing method	load_balancing_method

The following data is collected on the Load Balancer Pool Member [cmdb_ci_pool_member] table:

Label	Field name
IP address	ip_address
Service port	service_port

The following data is collected on the Load Balancer Services [cmdb_ci_lb_service] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field Name
Name	name
IP address	ip_address
Port	port
Operational status	operational_status
Install status	install_status
Hit count	hit_count

The following data is collected on the DNS [cmdb_ci_dns_name] table.

Label	Field name
Name	name
IP address	ip_address

Radware-appDirector load balancer discovery

Discovery of Radware load balancers is performed by SNMP.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Requirements

Important: Discovery treats load balancers as licensable entities and attempts to discover them primarily using SNMP. If a load balancer in your system, running on a Linux host, has SNMP and SSH ports open, Discovery might classify it based on the SSH port. This classification has priority over SNMP. To ensure that Discovery properly classifies your hardware load balancers, create a Discovery behavior for load balancers that includes SNMP but not SSH. Software load balancers are treated as applications.

Credentials

Configure SNMP credentials.

Classifier, pattern, and probes

Classifier	Probe	Pattern
Radware - AppDirector - Load Balancer	Horizontal discovery probe: launches patterns	AppDirector Load Balancer

Data discovery

The following data is collected on the Radware Load Balancers `[cmdb_ci_lb_radware]` table.

Label	Field name
Serial number	serial_number

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following data is collected on the DNS Name [cmdb_ci_dns_name] table.

Label	Field name
Name	name
IP address	ip_address

The following data is collected on the Load Balancer Services [cmdb_ci_lb_service] table.

Label	Field Name
Name	name
IP Address	ip_address
Port	port
Operational Status	operational_status
Install Status	install_status

Load balancer fields and probes

Discovery stores load balancer information in several tables.

The information in this topic pertains to discovery with probes and sensors only. If you are using patterns, such as to discover F5 Big-IP load balancers, see the relevant topic for that type of load balancer.

Note: By default, the system uses the discovered IP address of a load balancer for the CI record. This can be the management IP created for the device that is used in the Discovery schedule. For instructions on how to force Discovery to use the IP address of the load balancer's NIC rather than that of a management IP, see IP address selection properties.

Fields on the Citrix Netscalers [cmdb_ci_lb_netscaler] table

Field label and name	Probes that gather data
Serial Number [serial_number]	SNMP - Netscaler - Identity - Serial
Asset tag [asset_tag]	SNMP - Netscaler - Identity
Model ID [model_id]	SNMP - Netscaler - Identity

Fields on the Cisco GSS [cmdb_ci_lb_cisco_gss] table

Field label and name	Probes that gather data
Serial Number [serial_number]	Cisco GSS - Serial Number
Model ID [model_id]	Cisco GSS - Identity
IP Address [ip_address]	Cisco GSS - Network
Manufacturer [manufacturer]	Cisco GSS - Identity

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields on the Cisco CSS [cmdb_ci_lb_cisco_css] table

Field label and name	Probes that gather data
Serial Number [serial_number]	SNMP - Load Balancer – Identity
Model ID [model_id]	
IP Address [ip_address]	
Manufacturer [manufacturer]	

Fields on the F5 BIG-IP [cmdb_ci_lb_bigip] table

Field label and name	Probes that gather data
Active modules [active_modules]*	SNMP - F5 BIG-IP - System
Asset tag [asset_tag]	
Model ID [model_id]	
Failover mode [failover_mode]*	
Failover peer [failover_peer]*	
Serial Number [serial_number]	Cisco GSS - Serial Number

*not visible on the form by default. Customize the form to add this field.

Fields on the Load Balancer [cmdb_ci_lb] table

Field label and name	Probes that gather data
Operational status [operational_status]	• Apache - Get JK Module • Apache - Get Proxy Module • HAProxy - Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System • SNMP - Load Balancer - Identity
Name [name]	

Fields on the Load Balancer Application [cmdb_ci_lb_appl] table

Field label and name	Probes that gather data
Configuration file [config_file]*	Nginx - Get Configuration
Name [name]	• Apache - Get JK Module • Apache - Get Proxy Module • HAProxy - Get Configuration • Nginx - Get Configuration
IP Address [ip_address]*	
Most recent discovery [last_discovered]*	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields on the Load Balancer Application [cmdb_ci_lb_appl] table (continued)

Field label and name	Probes that gather data
Version [version]	• Apache - Version • HAProxy - Version • Nginx - Version

Fields on the Load Balancer Interface [cmdb_ci_lb_interface] table

Field label and name	Probes that gather data
Status [install_status]*	• Apache - Get JK Module • Apache - Get Proxy Module • HAProxy - Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System
Load Balancer [load_balancer]	
Name [name]	

*not visible on the form by default. Customize the form to add this field.

Fields on the Load Balancer Pool [cmdb_ci_lb_pool] table

Field label and name	Probes that gather data
Load Balancing Method [load_balance_method]	• Apache - Get JK Module • Apache - Get Proxy Module • HAProxy - Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System
Load Balancer [load_balancer]	
Most recent discovery [last_discovered]*	
Name [name]	

*not visible on the form by default. Customize the form to add this field.

Fields for Load Balancer Pool Member [cmdb_ci_lb_pool_member] data

Field label and name	Probes that gather data
Fully qualified domain name [fqdn]*	• HAProxy - Get Configuration • Nginx - Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System
IP Address [ip_address]	• Apache - Get JK Module • Apache - Get Proxy Module
Load balancer [load_balancer]	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields for Load Balancer Pool Member [cmdb_ci_lb_pool_member] data (continued)

Field label and name	Probes that gather data
Most recent discovery [last_discovery]*	• HAProxy - Get Configuration • Nginx- Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System
Name [name]	
Pool [pool]*	
Service Port [service_port]	

*not visible on the form by default. Customize the form to add this field.

Fields for Load Balancer Service [cmdb_ci_lb_service] data

Field label and name	Probes that gather data
Input URL [input_url]	• Apache - Get JK Module • Apache - Get Proxy Module • Nginx - Get Configuration
IP Address [ip_address]	• Apache - Get JK Module • Apache - Get Proxy Module • Cisco GSS - Identity • HAProxy - Get Configuration • Nginx - Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System
Last port [last_port]*	HAProxy - Get Configuration
Load Balancer [load_balancer]	• Apache - Get JK Module • Apache - Get Proxy Module • Cisco GSS - Identity • HAProxy - Get Configuration • Nginx - Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System
Most recent discovery [last_discovery]*	
Service Port [service_port]*	

*not visible on the form by default. Customize the form to add this field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields for Load Balancer VLAN [cmdb_ci_lb_vlan] data

Field label and name	Probes that gather data
Load Balancer [load_balancer] Name [name]	• Apache - Get JK Module • Apache - Get Proxy Module • HAProxy - Get Configuration • Nginx - Get Configuration • SNMP - F5 BIG-IP - System • SNMP - Netscaler - System
Tag [tag]	• SNMP - F5 BIG-IP - System • SNMP - Netscaler - System

Apache web server data

Discovery also collects data on Apache web serer load balancing modules using SSH. See Apache mod_jk and mod_proxy discovery for more information on probes for Apache web server data. For information on the tables, fields, and data sources that discovery populates for Apache web servers, see Apache web server discovery.

Layer 2 discovery

Discovery can detect the physical connections, known as layer 2, between network devices.

Layer 2 discovery process

Discovery uses multiple probes to gather information about network adapters and their Layer 2 connections. The following diagram shows the probes that Discovery launches for network Discovery, including those used for Layer 2 Discovery.

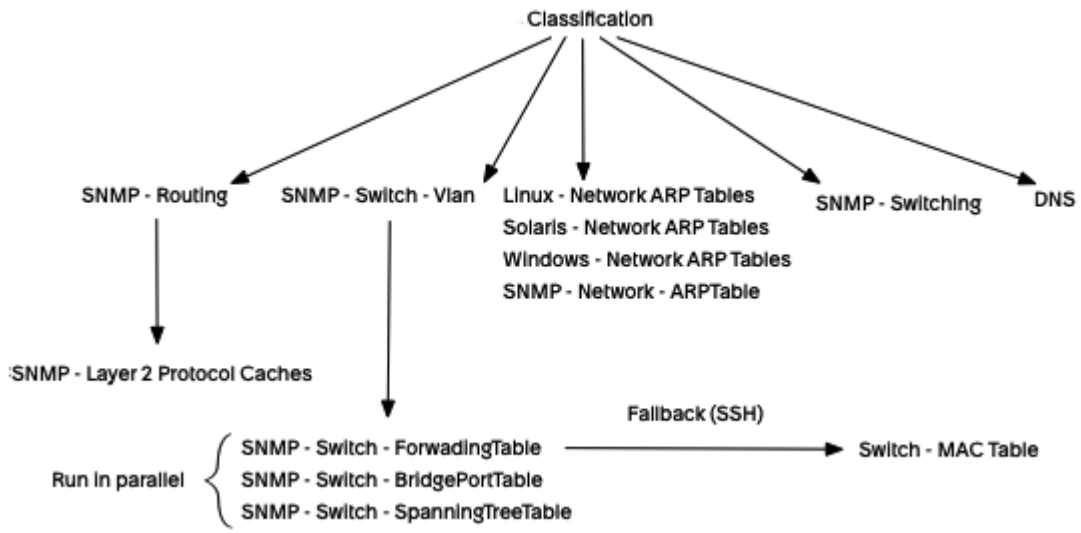

For example, if Discovery finds a switch in a network, it triggers the *SNMP - Switch - Vlan* probe and the *SNMP - Network - ARPTable* probe. For every Vlan that Discovery finds, it triggers various switch probes. If a switch has routing capabilities, Discovery triggers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

the *SNMP - Routing* probe to collect network adapter information in the Network Adapter [cmdb_ci_network_adapter] table. If Discovery finds a server, it triggers the appropriate Address Resolution Protocol (ARP) probe for that operating system. Discovery also supports the use of patterns, such as the *Network Switch* and *Network Router* patterns, which are available by default in Discovery. See Network switch and router discovery for more information.

During the discovery of a network device, Discovery creates records in the Router Interface [dscy_router_interface] table and the Switchport [dscy_switchport] table. This information contains network adapter information for that device. For SNMP-enabled devices, Discovery gathers the information from a routing probe during the exploration phase. The Layer 2 protocol cache probe runs next to collect neighbor data from the device.

Retrieving neighbor data

As Discovery gathers network information from the probes on a device, the SNMP - Layer 2 Protocol Caches identifies device neighbors via CDP and LLDP. This probe updates the Device Neighbors [discovery_device_neighbors] table, which contains the Layer 2 connections between switches. Sometimes, the neighbors of this device might not yet be known to the instance. The neighbor's interface cannot be resolved to a record until Discovery eventually finds the neighbor's side of the relationship. When Discovery runs on the neighboring device, Discovery completes the information for the neighbor's interface for the original reporting device.

Discovery can retrieve neighbor data from these caches on a network device:

- Cisco Discovery Protocol (CDP): Cache on Cisco devices that contains device neighbor information in the form of a protocol specific neighbor ID. Cisco Discovery Protocol (CDP).
- Link Layer Discovery Protocol (LLDP): Generic cache that contains device neighbor information in the form of a protocol specific neighbor ID.
- Address Resolution Protocol (ARP): Cache that contains the IP and MAC addresses of all connecting devices and servers.

How Layer 2 Discovery works

When device discovery is complete, the system runs the Service Discovery - Device Complete business rule on the Device [discovery_device_history] table. This business rule uses several strategies, together with information from various tables to create *Connects to:Connected by* relationships in the CMDB. Relationships can be created between:

- A server and a switch
- Two switches
- A device and the port of another device
- The ports of two devices

In the following example, Discovery found a server running AIX, and was also able to find two IP switches on the network. These relationships were created:

- A *Connects to* relationship between the AIX Server and the two IP switches A and B.
- A reference between the AIX server and its own network adapter.
- A *Connects to* relationship between the adapters on the two IP switches (not shown in the following image).
- A *Connects to* relationship between the network adapter of the AIX server and the switch port of IP switch A (highlighted in red). This kind of relationship is created by default.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

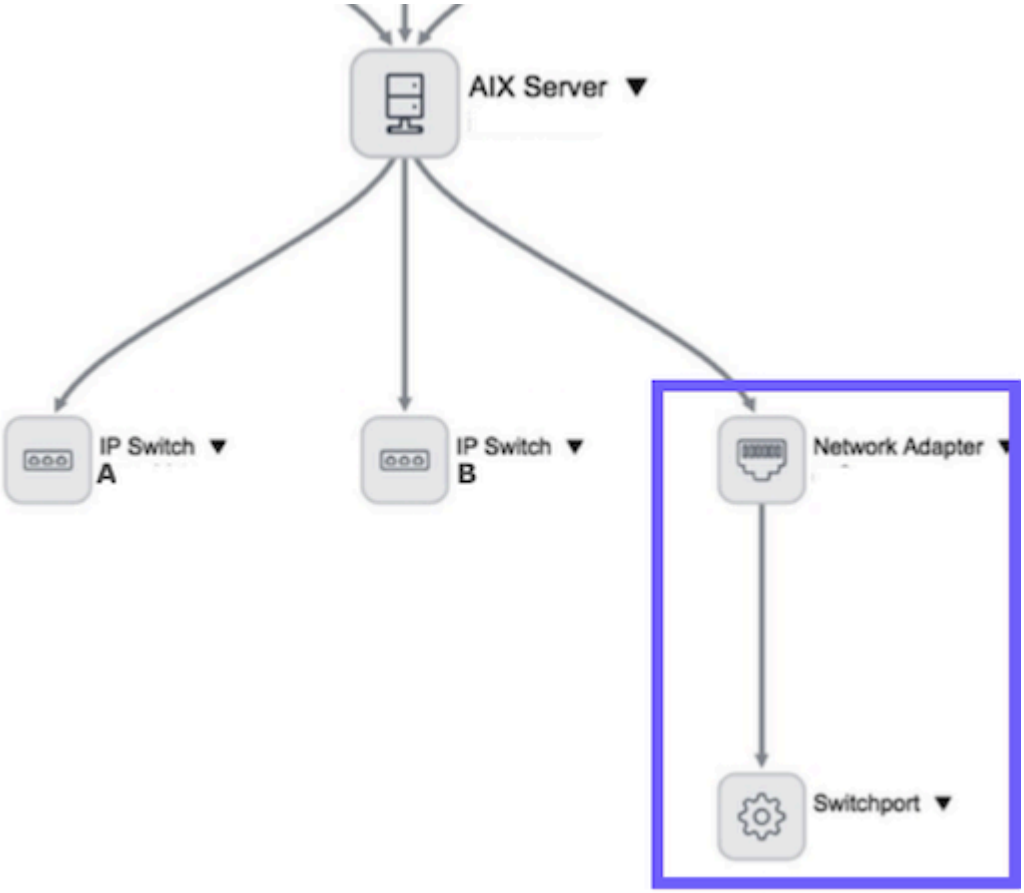

To view these relationships, open the dependency view for the server. To view the relationship between the two IP switches, open the dependency view from one of the switches and select the **Physical Network Connections** option for the **Dependency Type** in the map settings.

Layer 2 connection strategies

Layer 2 Discovery uses the following strategies to create relationships between devices in a network:

PhysicalHostConnectionStrategy

This strategy creates a connection between a discovered server, which is not a network device, and a network device. This strategy can only create a Layer 2 connection if there is a single MAC address on the port in the switch's forwarding table. If the port has multiple MAC addresses, Discovery does not create a connection between the server and the switch.

If the switch's port is identified, Discovery creates a **Connects to::Connected by** relationship between the server adapter and the switch port. Otherwise, Discovery creates a **Connects to::Connected by** relationship between the server adapter and the switch itself.

VMLayer2ConnectionStrategy

This strategy creates a connection between a VM and a network device. If the number of exceptional MAC addresses is very low (less than 3 and less than 15% of all MAC address of the ESX server and its VMs), Discovery creates a Layer 2 connection.

NetworkDeviceLayer2ConnnectionStrategy

This strategy creates a connection between a network device and its neighbors. Discovery creates a Layer 2 connection if there is a single neighbor with a neighbor

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

address, a neighbor interface, or both. If the neighbor's interface exists, Discovery creates a **Connects to::Connected by** relationship between the origin interface and the neighbor's interface. Otherwise, Discovery creates a **Connects to::Connected by** relationship between the origin interface and the switch with the neighbor's address.

SpanningTreeLayer2ConnectionStrategy

This strategy creates a connection between a network device and the parent of the network device in the spanning tree.

JavaScriptLayer2ConnectionStrategy

This strategy calls a Javascript function with empty implementation. This allows for the creation of an additional strategy.

Address Resolution Protocol (ARP) in Layer 2 discovery

The probes for Address Resolution Protocol (ARP) map the IP address of a computer or network device to a MAC address. These probes retrieve the IP address and MAC address for a CI from the Network Infrastructure Item [dscy_net_base] table. Devices that support SNMP, such as Linux computers and network devices, cache two types of address information:

- **Static**: Manually added address resolutions.
- **Dynamic**: Hardware name and IP address pairs added to the cache by previous, successful ARP resolutions.

When the ARP table Discovery completes, the system collects all static and dynamic table entries from devices via SNMP. If a new ARP entry is available, it is added to the Network ARP Table [discovery_net_arp_table]. If any previously discovered ARP entries are no longer cached in the device ARP table, the system removes the corresponding records from the CMDB using the reconciliation process.

Note: If new ARP entries are created after Discovery runs, they are not discovered until the next Discovery schedule. If ARP entries are removed from the device after Discovery runs, the CMDB ARP table is not updated until Discovery runs again.

Discovery provides these probes for extracting IP and MAC address resolution information:

Probe	ECC queue topic	Command	Description
Linux - Network ARP Tables	SSHCommand	sh ${file:linux_arp.sh} This runs a script file to try the net-tools command (arp -n) first. If that fails, it will then try the iproute2 command (ip neighbor).	SSH command probe that retrieves the network information from the ARP table on a Linux server.
Solaris - Network ARP Tables	SSHCommand	arp -an	SSH command probe that retrieves the network information from the ARP table on a Solaris server.
Windows - Network ARP Table	Powershell	arp -a	Powershell probe that retrieves the network information

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Probe	ECC queue topic	Command	Description
			from the ARP table on a Windows server.

SNMP - Network - ArpTable	SNMP	Table
		
		The SNMP probe uses this OID first:
		`iso.org.dod.internet.mgmt.mib-2.ip.ipNetToM`
		`ipNetToMediaPhysAddress,ipNetToMediaNetAddr`
		
		If the probe fails to return results, it uses this OID:
		`iso.org.dod.internet.mgmt.mib-2.ip.ipNetToP`
		`ipNetToPhysicalNetAddress,ipNetToPhysicalPh`
			SNMP probe that collects information from the ARP table on a switch or router.

SNMP switch probes

These probes return bridging information from VLANs connected across network switches, including port selection, forwarding tables, and the use of the spanning tree protocol.

SNMP - Switch - BridgePortTable

This probe returns all the ports from a switch that are used to create a bridge between network segments.

Bridging data returned

Table	Switch Bridge Port Table [discovery_switch_bridge_port_table] This table is used to map between a port number in the discovery_switch_fwd_table and an interface index.
OID	iso.org.dod.internet.mgmt.mib-2.dot1dBridge.dot1dBase.dot1dBasePortTable dot1dBasePort,dot1dBasePortIfIndex
Fields populated	• cmdb_ci • port • interface_index

SNMP - Switch - SpanningTreeTable

This probe returns the active path between any two network nodes bridged by a switch.

Spanning tree data returned

Table	Switch Spanning Tree Table [discovery_switch_spanning_tree_table]. This table is used to find Layer 2 connections between switches.
OID	iso.org.dod.internet.mgmt.mib-2.dot1dBridge.dot1dStp.dot1dStpPortTable dot1dStpPort,dot1dStpPortState,dot1dStpPortEnable,dot1dStpPortDesignatedRoot,dc

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Spanning tree data returned (continued)

Fields populated	• cmdb_ci • port • port_state • port_enable • designated_root • designated_bridge_mac

SNMP - Switch - ForwardingTable

This probe returns information from a switch's forwarding table.

Forwarding table data returned

Table	Switch Forwarding Table [discovery_switch_fwd_table] This is the forwarding table for the switches.
OIDs	These OIDs are built as needed by the *`DiscoveryVlanSwitchProcessor`* script include. • **Non-Cisco**: ◦ **Q-BRIDGE MIB**: `oid_spec_list = 'table iso.org.dod.internet.mgmt.mib-2.dot1dBridge.qBridgeMIB dot1qTpFdbAddress.' + vlanIndex + ',dot1qTpFdbPort.' + vlanIndex + ',dot1qTpFdbStatus.' + vlanIndex;` ◦ **BRIDGE MIB**: oid_spec_list = 'table iso.org.dod.internet.mgmt.mib-2.dot1dBridge.dot1dTp.dot1dTpFdbTable dot1dTpFdbAddress,dot1dTpFdbPort,dot1dTpFdbStatus'; • **Cisco BRIDGE MIB**: `oid_spec_list = 'table iso.org.dod.internet.mgmt.mib-2.dot1dBridge.dot1dTp.dot1 dot1dTpFdbAddress,dot1dTpFdbPort,dot1dTpFdbStatus';`
Additional probe called	The SSH probe `Switch - MAC Table` can be triggered by the *`DiscoveryVlanSwitchProcessor`* script include to populate additional Layer 2 data. If you are missing Layer 2 data after you run discovery on multiple VLANs, you should enable the triggering of the `Switch - MAC Table` probe and run discovery again. To do this, add this property to the System Property [sys_properties] table: `glide.discovery.sensors.forwardingtable.ssh`. Set the value of the property to **true**. By default this property is false.
Command	show mac address-table

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Forwarding table data returned (continued)

Fields populated	• cmdb_ci • vlan_id • port • status • mac_address (from the **cmdb_ci** field in the Network Infrastructure Item [discovery_net_base] table)

SNMP - Switch - Vlan

This probe returns VLAN IDs from a network switch and these OIDs:

```
iso.org.dod.internet.private.enterprises.cisco.ciscoMgmt.ciscoVtpMIB.vtpMIBObjects.vlanInfo.vtpVlanTable vtpVlanState

iso.org.dod.internet.private.enterprises.cisco.ciscoMgmt.ciscoVlanMembershipMIB.ciscoVlanMembershipMIBObjects.vmMembership.vmMembershipSummaryTable vmMembershipSummaryVlanIndex,vmMembershipSummaryMemberPorts,vmMembershipSummaryMember2kPorts

iso.org.dod.internet.private.enterprises.juniperMIB.jnxMibs.jnxExMibRoot.jnxExSwitching.jnxExVlan.jnxVlanMIBObjects.jnxExVlanTable jnxExVlanTag

iso.org.dod.internet.mgmt.mib-2.system.sysDescr
```

Other switch types are not supported.

Port tables

These port tables are populated during horizontal Discovery of a device:

- Network Adapter [cmdb_ci_network_adapter]
- Switchport [dscy_switchport]
- Router Interface [dscy_router_interface]
- Load Balancer Interface [cmdb_ci_lb_interface]

TCP connection discovery

Discovery identifies and classifies information about TCP connections using the ADM and ADME probes.

Discovery populates this data in the TCP Connections [cmdb_tcp] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Table Name	Field Name	Description
Absent	cmdb_tcp	absent	If the process is no longer running on this device. The first time the process is detected, the value is not selected (meaning that it is not absent). If Discovery runs again and does not detect the process, it is selected.
Computer	cmdb_tcp	computer	The device that is connected to or is listening on the TCP port.
IP	cmdb_tcp	ip	The IP address that this device or process is listening on or is connected to. If the IP address is 127.0.0.1, it is the host computer itself.
PID	cmdb_tcp	pid	The ID of the process that is connecting to or listening on this device (IP address).
Port	cmdb_tcp	port	The port number through which TCP traffic was detected.
Process	cmdb_tcp	process	The process that is connecting to or listening on the port.
Type	cmdb_tcp	type	• **Listening on**: the process is listening to the device via TCP. • **Connecting to**: if the process is connected to this device via TCP.
count	cmdb_tcp	count	The number of times that Discovery detects a specific TCP connection. This field is used by ADM enhanced probes.

Discovery populates this data in the Running Process [cmdb_running_process] table.

Label	Field Name	Description
Name	name	The name of the running process.
Computer	computer	The device that the process is running on.
Parent	parent	The parent process of this process, if any.
Listening on	listening_on	The ports that the processes is listening on.
Connected to	connected_to	The ports that the processes is connected to.
Command	command	The command for this process, including the path.
PID	pid	The ID for the process.
PPID	ppid	The ID for the parent process.
Absent	absent	If the process is no longer running on this device. The first time the process is detected, the value is not selected (meaning that it is not absent). If Discovery runs again and does not detect the process, it is selected.
Classify	classify	If the running process should be classified.
Key parameters	key_parameters	The key parameters for the command.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field Name	Description
Parameters	parameters	The parameters for the command.
Count	count	The number of times that Discovery detects a specific TCP connection. This field is used by ADM enhanced probes.

Standard ADM probes

Discovery launches these probes with standard ADM discovery:

Name	Description	Includes these sub-probes
AIX - ADM	Queries for information about active (running) AIX processes and active connections - the information required to perform application dependency mapping.	• AIX - Active Connections • AIX - Active Processes
HP-UX - ADM	Queries for information about active (running) HP-UX processes and active connections - the information required to perform application dependency mapping. This probe requires that lsof be installed and the UNIX - Active Connections probe be activated, which is inactive by default.	• UNIX - Active Connections • HP-UX - Active Processes
Solaris - ADM	Queries for information about active (running) Solaris processes and active connections - the information required to perform application dependency mapping. This probe is triggered by the Solaris Zones probe. The system triggers an ADM probe for each local zone contained in a global zone. This probe requires that lsof be installed on the global zone.	• Solaris - Active Processes • Solaris - Active Connections
UNIX - ADM	Queries for information about active (running) UNIX processes and active connections - the information required to perform application dependency mapping.	• UNIX - Active Connections • UNIX - Active Processes
Windows - ADM	Queries for information about active (running) Windows processes and active connections - the information required to perform application dependency mapping.	• Windows - Active Connections • Windows - Active Processes
UNIX - ADM netstat	Performs the same functions as the UNIX ADM probe relying on netstat to support continuous TCP-based discovery. Service Mapping uses this probe for traffic-based discovery.	• UNIX - Active Connections with Netstat. This probe retrieves TCP connections to populate the cmdb_tcp and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Includes these sub-probes
		cmdb_tcp_connection tables. • UNIX - Active Processes. This probe retrieves TCP processes and populates the cmdb_running_process table.

Enhanced ADM probes

Discovery launches these probes when enhanced ADM is activated:

Probe	Description
Windows - ADM Enhanced	Performs the same functions as the Windows ADM probe with additional parameters to control continuous TCP-based discovery. This probe is compatible with WinRM and non-WinRM supported MID Servers. Configure your MID Server using the `mid.windows.management_protocol` parameter. See MID Server parameters for the values you can use for this parameter.
UNIX - ADM Enhanced	Performs the same functions as the UNIX ADM probe with additional parameters to control continuous TCP-based discovery.

Use these probe parameters with the enhanced ADM probes:

Parameter	Description
mid.adme.max_result_size	Control the size of the TCP tracking file.
max.adme_result_payload_size	Controls the size of the result, including the payload. **Note:** If this is greater than the MID Server properties mid.discovery.max_payload_size, a configuration error can occur.

Related topics

Application Dependency Mapping (ADM) for Discovery

IP address discovery

Discovery identifies and classifies information about IP addresses.

Data Collected by Discovery on IP Addresses

Label	Table Name	Field Name	Source
IP Address	cmdb_ci_ip_address	ip_address	Various internal
IP version	cmdb_ci_ip_address	ip_version	Various internal

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data Collected by Discovery on IP Addresses (continued)

Label	Table Name	Field Name	Source
Netmask	cmdb_ci_ip_address	netmask	Various internal
Nic	cmdb_ci_ip_address	nic	Various internal

Related topics

Discovery IP address configuration

IP network discovery

Discovery identifies and classifies information about IP networks.

Data collected by Discovery on IP networks

Label	Table Name	Field Name	Source
Discover	cmdb_ci_ip_network	discover	Various internal
Subnet	cmdb_ci_ip_network	subnet	Various internal
Network discovery	cmdb_ci_ip_network	network_discovery	Various internal
Last discovered	cmdb_ci_ip_network	last_discovered	Various internal
MID server	cmdb_ci_ip_network	mid_server	Various internal
Router	cmdb_ci_ip_network	router	Various internal
State	cmdb_ci_ip_network	state	Various internal

Related topics

Network discovery

Network printer discovery

Discovery identifies and classifies information about network printers.

Discovery can identify and classify Network Printers. For more information, see Network discovery.

Data collected by Discovery on network printers

Label	Table name	Field name	Source
Serial number	cmdb_ci	serial_number	SNMP, various MIBs
Start date	cmdb_ci	start_date	SNMP, RFC1213 MIB
Manufacturer	cmdb_ci	manufacturer	SNMP, RFC1213 MIB
Model ID	cmdb_ci	model_id	SNMP, RFC1213 MIB
IP address	cmdb_ci_network_adapter	ip_address	SNMP, IP MIB
MAC address	cmdb_ci_network_adapter	mac_address	SNMP, IF MIB
Printer type	cmdb_ci_printer	print_type	SNMP, PRINT MIB
Use count	cmdb_ci_printer	use_count	SNMP, PRINT MIB
Use count units	cmdb_ci_printer	use_units	SNMP, PRINT MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery on network printers (continued)

Label	Table name	Field name	Source
Colors	cmdb_ci_printer	colors	SNMP, PRINT MIB
Horizontal resolution	cmdb_ci_printer	horizontal_resolution	SNMP, PRINT MIB
Vertical resolution	cmdb_ci_printer	vertical_resolution	SNMP, PRINT MIB
Resolution units	cmdb_ci_printer	resolution_units	SNMP, PRINT MIB
Description	discovery_printer_supplies	description	SNMP, PRINT MIB
Supply type	discovery_printer_supplies	supply_type	SNMP, PRINT MIB
Supply class	discovery_printer_supplies	supply_class	SNMP, PRINT MIB
Current level	discovery_printer_supplies	current_level	SNMP, PRINT MIB
Max capacity	discovery_printer_supplies	max_capacity	SNMP, PRINT MIB
Name	cmdb_ci	name	Shazzam, SNMP, various MIBs

Network switch and router discovery

Discovery identifies and classifies information about network switches and routers.

Network switches and routers often have very similar capabilities. It is very common for some switches known as Layer 3 switches to have IP routing. Larger routers with optional modules might accept switching modules. Because of these overlaps and the resulting ambiguity of a particular device's classification, Discovery collects the same data for both routers and switches, if it is available in any given device. For details on how Discovery collects data about connections between network devices and other components, see Layer 2 discovery.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Prerequisites

- The router or switch needs to have IP Address information.
- The system property glide.discovery.L3_mapping is set to true to discover routers and switches.
- The router or Layer 3 switch that provides the IP Address needs to have been successfully discovered with populated Exit Interface Routing Rules.
- If you must populate the [discovery_device_neighbors] table during horizontal layer 2 discovery of the bgp-enabled devices, set the **BGP router exploration disable** property to **No**. Notice that enabling this property can cause performance issues including out-of-memory issues on the MID Server.

System properties

You can configure many aspects of Discovery to work according to your personal preferences. The following properties apply to network routers and switches. To view the properties, navigate to **Discovery Definition** > **Properties** unless indicated otherwise.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
glide.discovery.L3_mapping	This creates IP Connecton::IP Connection relation between servers and network devices. • **Type:** string • **Default value:** true • **Location:** System Property [sys_properties] table
glide.discovery.max_concurrent_invocations_per_schedule	Sets a maximum number of scheduled invocations of the same Discovery schedule. It prevents a backlog of scheduled runs if Discovery does not finish before the next invocation is scheduled to run. The value is an integer that represents the max number of automated invocations of the same schedule that may proceed at one time. If the limit has been reached, subsequent scheduled invocations are canceled. A value of 0 or any negative number disables this property. • **Type:** integer • **Default value:** 3 • **Location: Discovery Definition > Properties** **Note:** This property does not apply to schedules that have a 'Run after' configuration set to 'Even if canceled'.

For a list of properties used in discovery in general, see Discovery properties.

Credentials

Discovery explores many kinds of devices, such as switches, routers, and printers, using the SNMP protocol. SNMP Community Credentials do not include a user name, just a password, which is the community string. The default read-only community string for many SNMP devices is public, and Discovery will try that automatically. SNMPv3 credentials require a user name and a key, depending on the protocol selected. Enter the appropriate SNMP credentials if they differ from the public community string.

Note: SSH is not supported for routers and switches, but is supported for load balancers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier, pattern, and probes

Classifiers	Trigger probes	Pattern
Standard Network Router Standard Network Switch	• Horizontal Pattern: launches patterns • SNMP-Routing^ • DNS* • SNMP-Identity* • SNMP-Switch-Vlan* • SNMP-Switching* • SNMP - Network - ArpTable* (Standard Network Switch only)	Network Router Network Switch

*For new instances, these probes are inactive on the classifier. Discovery uses patterns for discovery.

^This probe remains active by default for the Standard Network Router classifier, even when Discovery uses pattern discovery.

To use patterns, verify that the correct pattern is specified in the horizontal pattern probe on the classifier. See Add the Horizontal Pattern probe to a classifier for instructions.

Discovering stacked switches

The Network Switch shared library supports Cisco stacked switches. The CI table for this is IP Switch [cmdb_ci_ip_switch]. Multiple switches in the stack are identified by the primary IP address. Each member of the stack is identified by a combination of elements, including the primary IP address, a name associated with the switch, and the stack mode. The relationship is created between a primary and all of the secondaries. The primary switch has all the fields listed in the table below. The secondary switch has name, serial_number, model_id, manufacturer, stack, and stack_mode.

Note: If all the ports on a switch are open for port 161, that switch will have multiple IPs (one for each port) but the serial number for all the ports is the same. The CI identification is performed based on the serial number, therefore the same record is updated every time instead of creating a new record.

Router and switch data

Table name	Label/Field name	Source
cmdb_ci	Serial Number serial_number	SNMP, various MIBs
cmdb_ci	Manufacturer manufacturer	SNMP, RFC1213 MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Label/Field name	Source
cmdb_ci	Model ID model id	SNMP, RFC1213 MIB
cmdb_ci_ip_router	Name name	
cmdb_ci_ip_router	Model ID model ID	SNMP, RFC1213 MIB
cmdb_ci_ip_router	Manufacturer manufacturer	SNMP, RFC1213 MIB
cmdb_ci_ip_router	Serial Number serial_number	SNMP, various MIBs
cmdb_ci_ip_router	IP Address ip_address	SNMP, IP MIB
cmdb_ci_ip_router	Description short_description	
cmdb_ci_ip_router	Can route IP can_route	SNMP, IP MIB, BGP MIB
cmdb_ci_ip_router	Can switch IP can_switch	SNMP, dot1dBridge MIB
cmdb_ci_ip_router	Port port	
cmdb_ci_ip_switch	Name name	
cmdb_ci_ip_switch	Model ID model ID	SNMP, RFC1213 MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Label/Field name	Source
cmdb_ci_ip_switch	Manufacturer manufacturer	SNMP, RFC1213 MIB
cmdb_ci_ip_switch	Serial Number serial_number	SNMP, various MIBs
cmdb_ci_ip_switch	IP Address ip_address	SNMP, IP MIB
cmdb_ci_ip_switch	Description short_description	
cmdb_ci_ip_switch	Can route IP can_route	SNMP, IP MIB, BGP MIB
cmdb_ci_ip_switch	Ports ports	SNMP, dot1dBridge MIB
cmdb_ci_ip_switch	Can switch IP can_switch	SNMP, dot1dBridge MIB
cmdb_ci_ip_switch	Can partition VLANs can_partitionvlans	SNMP, dot1dBridge MIB
cmdb_ci_ip_switch	Stack stack	SNMP, entPhysicalTable
cmdb_ci_ip_switch	Stack Mode stack_mode	SNMP, entPhysicalTable
cmdb_ci_netgear	Can route IP can_route	SNMP, IP MIB, BGP MIB
cmdb_ci_netgear	Can switch IP can_switch	SNMP, dot1dBridge MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Label/Field name	Source
cmdb_ci_netgear	Can partition VLANs cdscy_router_interface	SNMP, dot1dBridge MIB
cmdb_ci_netgear	Can Hub can_hub	SNMP, IP MIB
cmdb_ci_network_adapter	IP Address ip_address	SNMP, IP MIB
cmdb_ci_network_adapter	MAC Address mac_address	SNMP, IF MIB
discovery_device_neighbors	Neighbor Address neighbor_address	ciscoCdpMIB
discovery_device_neighbors	Neighbor ID neighbor_ID	ciscoCdpMIB, lldpMIB
discovery_device_neighbors	Neighbor Interface neighbor_interface	ciscoCdpMIB, lldpMIB
discovery_device_neighbors	Neighbor Source neighbor_source	ciscoCdpMIB, lldpMIB
discovery_device_neighbors	Origin Interface origin_interface	ciscoCdpMIB, lldpMIB
discovery_net_arp_table	IP Address IP_address	SNMP, IP MIB
discovery_net_arp_table	MAC Address mac_address	SNMP, IP MIB
discovery_switch_bridge_port_table	Interface Index interface_index	SNMP, dot1dBridge MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Label/Field name	Source
discovery_switch_bridge_port_table	Port port	SNMP, dot1dBridge MIB
discovery_switch_fwd_table	MAC Address mac_address	SNMP, dot1dBridge MIB
discovery_switch_fwd_table	Port port	SNMP, dot1dBridge MIB
discovery_switch_fwd_table	Status status	SNMP, dot1dBridge MIB
discovery_switch_fwd_table	VLAN ID vlan_id	SNMP, dot1dBridge MIB
discovery_switch_spanning_tree_table	Designated Bridge MAC designated_bridge_mac	SNMP, dot1dBridge MIB
discovery_switch_spanning_tree_table	Designated Root designated_root	SNMP, dot1dBridge MIB
discovery_switch_spanning_tree_table	Port port	SNMP, dot1dBridge MIB
discovery_switch_spanning_tree_table	Port Enable port_enable	SNMP, dot1dBridge MIB
discovery_switch_spanning_tree_table	Port State port_state	SNMP, dot1dBridge MIB
dscy_router_interface	Name name	SNMP, IP MIB
dscy_router_interface	Type type	SNMP, IP MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Label/Field name	Source
dscy_router_interface	Number number	SNMP, IP MIB
dscy_router_interface	IP Address ip_address	SNMP, IP MIB
dscy_router_interface	MAC Address mac_address	SNMP, IP MIB
dscy_route_interface	Destination Network dest_ip_network	SNMP, IP MIB
dscy_route_interface	Type type	SNMP, IP MIB
dscy_route_next_hop	Destination Network dest_ip_network	SNMP, IP MIB
dscy_route_next_hop	Type type	SNMP, IP MIB
dscy_route_next_hop	Next Hop next_hop_ip_address	SNMP, IP MIB
dscy_switchport	Type type	SNMP, dot1dBridge MIB
dscy_switchport	Status status	SNMP, dot1dBridge MIB
dscy_switchport	MAC Address mac_address	SNMP, dot1dBridge MIB
dscy_switchport	Port Number port_number	SNMP, dot1dBridge MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Label/Field name	Source
dscy_switchport	Interface Name interface_name	SNMP, dot1dBridge MIB
dscy_switchport	Interface Number interface_number	SNMP, dot1dBridge MIB
dscy_swtch_fwd_rule	MAC Address mac_address	SNMP, dot1dBridge MIB
dscy_swtch_fwd_rule	MAC Manufacturer mac_mfr	SNMP, dot1dBridge MIB
dscy_swtch_fwd_rule	Status status	SNMP, dot1dBridge MIB
dscy_swtch_fwd_rule	IP address ip_address	SNMP, dot1dBridge MIB
dscy_swtch_fwd_rule	Netmask netmask	SNMP, dot1dBridge MIB
dscy_swtch_partition	Base IP Address base_ip_address	SNMP, dot1dBridge MIB
dscy_swtch_partition	Base MAC Address base_mac_address	SNMP, dot1dBridge MIB
dscy_swtch_partition	Base Netmask base_netmask	SNMP, dot1dBridge MIB
dscy_swtch_partition	Type type	SNMP, dot1dBridge MIB
dscy_swtch_partition	Transparent transparent	SNMP, dot1dBridge MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table name	Label/Field name	Source
dscy_swtch_partition	Source Route sourceroute	SNMP, dot1dBridge MIB
dscy_swtch_partition	Name name	SNMP, dot1dBridge MIB
dscy_swtch_partition	Status status	SNMP, dot1dBridge MIB
dscy_swtch_partition	Interface Number interface_number	SNMP, dot1dBridge MIB

Turn off collection of port data for SNMP network devices

In some environments you may not want to collect the port data for SNMP network devices.

Before you begin

Role required: admin

About this task

Network devices can contain an extremely large number of ports. If each individual port's information is not critical to the overall management of the devices, you may not want to collect the port data for SNMP network devices. This would affect records from being created in the Switchports table (dscy_switchport). This process only impacts probe discovery.

Procedure

1. Navigate to **Discovery Definition** > **Probes**.
2. Go to the **SNMP - Switching** probe.
3. In the **SNMP Fields** tab, find the field named **ports**, double-click the value in the **Active** column, and change the value to **false**.
   The collection and processing of ports information for switches is then disabled. The **ports** field for the switch CI should have the correct ports count even though the related list **switchports** is empty.
4. Click the check mark to save the change.

IP service and daemon discovery

Discovery identifies and classifies information about services and daemons.

Data Collected by Discovery on Services and Daemons

Label	Table Name	Field Name	Source
Name	`cmdb_ip_service`	name	Various
Port	`cmdb_ip_service`	port	Various
Description	`cmdb_ip_service_ci`	description	Various

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data Collected by Discovery on Services and Daemons (continued)

Label	Table Name	Field Name	Source
CI	cmdb_ip_service_ci	ci	Internal reference
Service	cmdb_ip_service_ci	service	Internal reference

Related topics

Network discovery

Uninterruptible power supply (UPS) discovery

Discovery identifies and classifies information about an uninterruptible power supply (UPS).

Discovery can identify and classify uninterruptible power supplies.

Data collected by Discovery on uninterruptible power supplies

Label	Table name	Field name	Source
Serial number	cmdb_ci	serial_number	SNMP, various MIBs
Start date	cmdb_ci	start_date	SNMP, RFC1213 MIB
Manufacturer	cmdb_ci	manufacturer	SNMP, RFC1213 MIB
Model ID	cmdb_ci	model_id	SNMP, RFC1213 MIB
IP address	cmdb_ci_network_adapter	ip_address	SNMP, IP MIB
MAC address	cmdb_ci_network_adapter	mac_address	SNMP, IF MIB
UPS software version	cmdb_ci_ups	ups_software_version	SNMP, UPS MIB
Agent software version	cmdb_ci_ups	agent_software_version	SNMP, UPS MIB
Attached devices	cmdb_ci_ups	attached_devices	SNMP, UPS MIB
Battery status	cmdb_ci_ups	battery_status	SNMP, UPS MIB
Seconds remaining on battery	cmdb_ci_ups	seconds_on_battery	SNMP, UPS MIB
Estimated minutes remaining on battery	cmdb_ci_ups	est_mins_remaining	SNMP, UPS MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery on uninterruptible power supplies (continued)

Label	Table name	Field name	Source
Estimated charge % remaining	cmdb_ci_ups	est_charge_remaining	SNMP, UPS MIB
Battery voltage	cmdb_ci_ups	battery_voltage	SNMP, UPS MIB
Battery current	cmdb_ci_ups	battery_current	SNMP, UPS MIB
Battery temperature (C)	cmdb_ci_ups	battery_temperature	SNMP, UPS MIB
Input line bads	cmdb_ci_ups	input_line_bads	SNMP, UPS MIB
Output source	cmdb_ci_ups	output_source	SNMP, UPS MIB
Output frequency	cmdb_ci_ups	output_freq	SNMP, UPS MIB
Bypass frequency	cmdb_ci_ups	bypass_freq	SNMP, UPS MIB
Nominal input voltage	cmdb_ci_ups	nom_input_volt	SNMP, UPS MIB
Nominal input frequency	cmdb_ci_ups	nom_input_freq	SNMP, UPS MIB
Nominal output voltage	cmdb_ci_ups	nom_output_volt	SNMP, UPS MIB
Nominal output frequency	cmdb_ci_ups	nom_output_freq	SNMP, UPS MIB
Rated output VA	cmdb_ci_ups	rated_output_va	SNMP, UPS MIB
Rated output power	cmdb_ci_ups	rated_output_power	SNMP, UPS MIB
Low battery threshold minutes	cmdb_ci_ups	low_battery_threshold_mins	SNMP, UPS MIB
Audible alarm status	cmdb_ci_ups	audible_alarm_status	SNMP, UPS MIB
Low voltage transfer point	cmdb_ci_ups	low_voltage_transfer_point	SNMP, UPS MIB
High voltage transfer point	cmdb_ci_ups	high_voltage_transfer_point	SNMP, UPS MIB
Input index	cmdb_ci_ups_input	input_index	SNMP, UPS MIB
Input frequency (Hz)	cmdb_ci_ups_input	input_freq	SNMP, UPS MIB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collected by Discovery on uninterruptible power supplies (continued)

Label	Table name	Field name	Source
Input voltage (RMS VAC)	cmdb_ci_ups_input	input_volt	SNMP, UPS MIB
Input current (RMS AAC)	cmdb_ci_ups_input	input_current	SNMP, UPS MIB
Input power (Watts)	cmdb_ci_ups_input	input_power	SNMP, UPS MIB
Output index	cmdb_ci_ups_output	output_index	SNMP, UPS MIB
Output load (%)	cmdb_ci_ups_output	output_load	SNMP, UPS MIB
Output voltage (RMS VAC)	cmdb_ci_ups_output	output_volt	SNMP, UPS MIB
Output current (RMS AAC)	cmdb_ci_ups_output	output_current	SNMP, UPS MIB
Output power (Watts)	cmdb_ci_ups_output	output_power	SNMP, UPS MIB
Bypass index	cmdb_ci_ups_bypass	bypass_index	SNMP, UPS MIB
Bypass voltage (RMS VAC)	cmdb_ci_ups_bypass	bypass_volt	SNMP, UPS MIB
Bypass current (RMS AAC)	cmdb_ci_ups_bypass	bypass_current	SNMP, UPS MIB
Bypass power (Watts)	cmdb_ci_ups_bypass	bypass_power	SNMP, UPS MIB
Alarm index	cmdb_ci_ups_alarm	alarm_index	SNMP, UPS MIB
Alarm type	cmdb_ci_ups_alarm	alarm_type	SNMP, UPS MIB
Alarm time	cmdb_ci_ups_alarm	alarm_time	SNMP, UPS MIB

IBM WebSphere DataPower discovery

ITOM Visibility discovers IBM WebSphere DataPower. The Discovery feature uses the DataPower Server pattern for horizontal discovery. The Service Mapping feature uses the DataPower pattern for top-down discovery.

Prerequisites

Configure the following credentials before discovering the IBM WebSphere DataPower devices:

SNMP credentials

SNMP credentials are necessary for using the following SNMP queries:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For horizontal discovery of the Data Power Hosting Server [cmdb_ci_datapower_server]:
  - 1.3.6.1.2.1.1.1
  - 1.3.6.1.2.1.1.5
  - 1.3.6.1.2.1.4.20.1
  - 1.3.6.1.2.1.2.2.1
  - 1.3.6.1.4.1.14685.3.1.196.1
  - 1.3.6.1.4.1.14685.3.1.11.1
  - 1.3.6.1.4.1.14685.3.1.112.1
- For top-down discovery of the Data Power [cmdb_ci_datapower_server]:
  - 1.3.6.1.2.1.1.1
  - 1.3.6.1.2.1.1.5
  - 1.3.6.1.4.1.14685.3.1.11.1

Applicative credentials

Together with SNMP credentials, applicative credentials are required to discover DataPower application domains, otherwise known as logical CIs.

In addition, SNMP and applicative credentials are used for top-down discovery of applicative relationships (corresponding application DataPower domain).

User with permissions to run SOAP

Permissions to run SOAP queries are necessary for top-down discovery of the DataPower Web Service and MQ connection.

Data collected by Discovery during horizontal discovery

By default, Discovery uses the DataPower Server pattern to perform the discovery.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Table and field	Description
Data Power Hosting Server [cmdb_ci_datapower_server]	
Model number [model_number]	The model of the server hosting the WebSphere DataPower.
Host name [host_name]	The host name of the server hosting the WebSphere DataPower.
Serial number [serial_number]	The serial number of the server hosting the WebSphere DataPower.
IP Address [cmdb_ci_ip_address]	
IP Address [ip_address]	The IP address of the server hosting the WebSphere DataPower.
Netmask [netmask]	The netmask of the server hosting the WebSphere DataPower.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and field	Description
Network Adapters [cmdb_ci_network_adapter]	
MAC address [mac_address]	The MAC address of the network adapter.
Name [name]	The name of the network adapter.

CI Relationships

CI	Relationship/Reference	CI
Data Power Hosting Server [cmdb_ci_datapower_server]	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Network Adapters [cmdb_ci_network_adapter]	Owns::Owned by	Data Power Hosting Server [cmdb_ci_datapower_server]

Data collected by Service Mapping during top-down discovery

Table and field	Description
Data Power [cmdb_ci_app_server_datapower]	
Host name [host_name]	The attributes of the DataPower application.
SOAP URL [soap_url]	
Version [version]	
IP address [ip_address]	
Name [name]	
Domain [domain]	

Cisco Unified Computing System (UCS)-HD device discovery

Discovery and Service Mapping Patterns application uses the to find Cisco UCS equipment, including chassis and blades. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Prerequisites

Create SNMP credentials

For more information, see SNMP credentials.

Create applicative credentials

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create applicative credentials for the **Cisco UCS Equipment [cmdb_ci_ucs_equipment]** CI type. For more information, see Applicative credentials.

Verify network connectivity

Verify that the MID Server has HTTPS access to the UCS Manager on port 443 and is able to reach the management IP address of the UCS equipment.

Verify UCS Manager account requirements

Verify that the account has at least the **read-only** role in UCS Manager and has permissions to execute XML API queries for the following object classes:

- `computeRackUnit`
- `equipmentChassis`
- `computeBlade`
- `mgmtIf`
- `topSystem`
- `versionEp`
- `equipmentFan`

Classifier, pattern, and OIDs

The UCS classifier uses HorzontalDiscoveryProbe to launch the UCS - HD pattern. The classifier also uses these SNMP OIDs, which specify the Cisco UCS Equipment `[cmdb_ci_ucs_equipment]` table:

OID	Manufacturer	Model	Table
1.3.6.1.4.1.9.12.3.1.3.847	Unified Computing Cisco UCS 6120XP 20-Port Fabric Interconnect	1.4(3u)(NX-OS) 5.0 Aggregation Switch	Cisco UCS Equipment
1.3.6.1.4.1.9.12.3.1.3.899	Unified Computing Cisco UCS 6140XP 40-Port Fabric Interconnect	1.4(3u)(NX-OS) 5.0 Aggregation Switch	Cisco UCS Equipment
1.3.6.1.4.1.9.12.3.1.3.1062	Unified Computing Cisco UCS 6248UP 48-Port Fabric Interconnect	2.0(1t)(NX-OS) 5.0.2 Aggregation Switch	Cisco UCS Equipment
1.3.6.1.4.1.9.12.3.1.3.1063	Unified Computing Cisco UCS 6296UP 96-Port Fabric Interconnect	2.0(1t)(NX-OS) 5.0.2 Aggregation Switch	Cisco UCS Equipment

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the UCS - HD pattern pattern. Only B-series devices are discovered. The UCS - HD pattern does not find C, S, or E series devices.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cisco UCS Equipment [cmdb_ci_ucs_equipment]

Label	Field Name
Name	name
IP address	ip_address
Model	model_id
Distinguished name	dn
UCS Manager Version	version

Cisco UCS Chassis [cmdb_ci_ucs_chassis]

Label	Field Name
Name	name
Model number	model_number
Serial number	serial_number
Chassis ID	chassis_id
Operational status	operational_status
Distinguished name	dn

Cisco UCS Blade [cmdb_ci_ucs_blade]

Label	Field Name
Name	name
Model number	model_number
Serial number	serial_number
CPU count	cpu_count
CPU core count	cpu_core_count
Ram	ram
Number of Adapters	number_of_adapters
Chassis ID	chassis_id
Slot ID	slot_id
Description	short_description
Distinguished name	dn
Blade Serial Number	blade_serial_number
Operational status	operational_status

Cisco UCS Rack Mount Units [cmdb_ci_ucs_rack_unit]

Label	Field Name
Name	name

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cisco UCS Rack Mount Units [cmdb_ci_ucs_rack_unit] (continued)

Label	Field Name
Vendor	vendor
Server ID	server_id
Availability	availability
UUID	uuid
Rack serial	rack_serial
Number of CPUs	num_of_cpus
Number of cores	num_of_cores
Model ID	model_id
Total Memory	total_memory
Distinguished name	dn
State	state

Relationships

These relationships are created to support UCS - HD discovery.

CI	Relationship	CI
cmdb_ci_ucs_chassis	Contains::Contained by	cmdb_ci_ucs_blade
cmdb_ci_ucs_equipment	Contains::Contained by	cmdb_ci_ucs_rack_unit
cmdb_ci_ucs_equipment	Contains::Contained by	cmdb_ci_ucs_chassis
cmdb_ci_server	Runs on::Runs	cmdb_ci_ucs_blade

Dell Remote Assistant Card discovery

The Dell™ Remote Assistant Card (DRAC) and Integrated Remote Assistance Card (iDRAC) provide users with tools and functionality to monitor, troubleshoot, and repair servers. You can generate DNS URLs to access out-of-band devices when security prevents access from IP addresses.

To identify the DRAC, Discovery uses the **SNMP – DRAC** probe. This probe uses SNMPv1 and SMNPv2c. Out-of-band devices have a dedicated management channel for device maintenance. System administrators can remotely monitor and manage devices with the dedicated management channel, including hardware sensor monitoring, access to local media drives, or adjusting BIOS settings. DNS URLs can be generated for out-of-band devices to access the out-of-band device page when the SSL certificate forbids navigating through the IP address.

Prerequisites

Create SNMP credentials

SNMP credentials are required to discover DRAC devices. For more information, see SNMP credentials.

Run network discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information on configuring and running network discovery to discover DRAC devices, see Run network Discovery.

Data collected

Data collected by Discovery for Dell DRAC

Label	Table	Name	Command
Firmware version	cmdb_ci_outofband_device	firmware_version	SNMP walk: drsFirmwareVersion (racFirmwareVersion for iDRAC7)
Host *	cmdb_ci_outofband_device	host	SNMP walk: drsSystemServiceTag (systemServiceTag for iDRAC7)
IP Address	cmdb_ci_outofband_device	ip_address	DNS probe
Name	cmdb_ci_outofband_device	name	SNMP Identity probe
Product version	cmdb_ci_outofband_device	product_version	SNMP walk: drsProductVersion (racVersion for iDRAC7)
Type	cmdb_ci_outofband_device	type	SNMP walk: drsProductType (racType for iDRAC7)
URL	cmdb_ci_outofband_device	url	SNMP walk: drsProductURL (racURL for iDRAC7)

* **Host** is a reference to the cmdb_ci_computer table via the serial number. For the cmdb_ci_outofband_device.host field to be populated correctly, the host machine must be discoverable or exist within the CMDB with the appropriate serial number.

Out-of-band devices

A **DNS URL** is generated for configuration items in the **Out-Of-Band Device** list if the configuration item has a **Name** and **DNS Domain**. The **DNS URL** enables you to access the out-of-band device page when the SSL certificate forbids navigating through the **IP URL**. The **DNS URL** has the format `https://Name.DNS_Domain`. If either the **Name** or **DNS Domain** field is empty, the **DNS URL** contains the value in the **URL** field.

HTTP device discovery

Discovery can find devices that use the HTTP protocol.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Classifier, probes, and patterns

HTTP discovery relies on a the *HTTP* port probe, which Discovery runs after trying to run SSH and SNMP port probes. This port probe launches the *HTTP - Classify* probe which runs an HTTP GET request for every entry in the HTTP classification table. You should configure your own classifications that the *HTTP - Classify* probe can access.

The HTTP port probe uses mappings between a protocol and port number in the IP Service [cmdb_ip_service] table. You can add additional port-to-protocol definitions in this table, or you can override the protocol used for a specific HTTP classifier.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classifier	Probes	Pattern
	HTTP (port probe): this probe launches the HTTP - Classify probe, which runs an HTTP GET request for every entry in the HTTP classification table.	

Run discovery through an HTTP or HTTPS REST call

Discovery can classify devices using the HTTP(S) protocol. You can create your own HTTP classifier to find devices and access them with Basic Auth credentials, rather than using SNMP or SSH credentials.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Credentials**.
2. Select **New**.
3. Select **Basic Auth Credentials**.
   In most cases, you can use Basic Auth Credentials to access a device through an HTTP or HTTPS REST call.

   Important: The HTTP Classify probe no longer attempts credentials over the HTTP protocol by default. To override this behavior, you can enable *mid.http_classy.allow_credentials_over_http*. However, enabling this setting can expose credentials to man-in-the-middle (MitM) attacks. Therefore, it's strongly recommended to keep this property set to **false** and use HTTPS whenever possible.

4. Add the user name and password that you need to access the device.
5. Select **Submit**.
6. Set up an HTTP classifier by navigating to **Discovery Definition** > **CI Classification** > **HTTP**.
7. Select **New**.
8. Fill out the form fields (see table):

The F5 HTTP classification

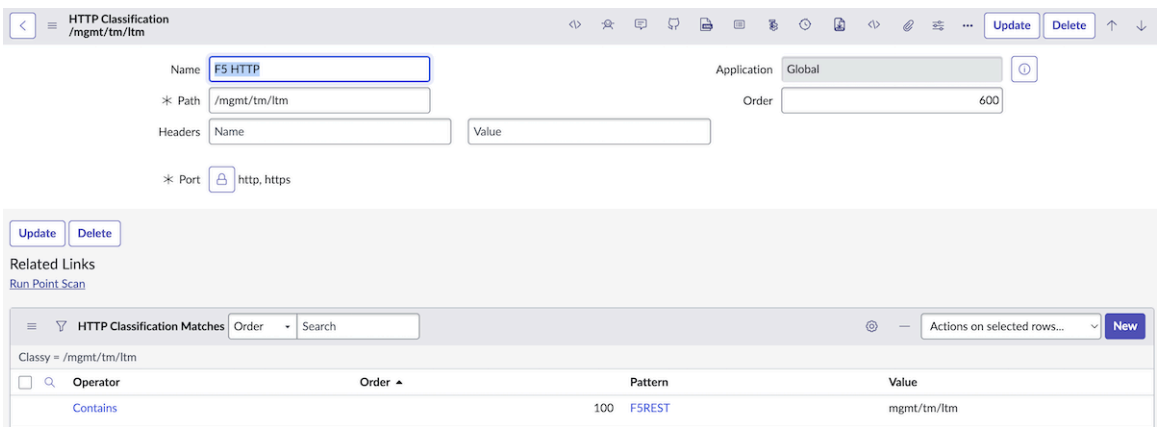

Field	Description
Name	Enter a name for the classifier.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Path	Enter the path for the HTTP GET request.
Headers	Add an option header with a value. You can use a variable in the format `${variable}` instead of a static value. Supported headers depend on the REST web service provider that you want to connect to. See the documentation for your web service provider to identify which headers are valid or required.
Port (IP Service)	If you want this classifier to use a custom port that is not already a default value in the IP Service [cmdb_ip_service] table, do the following: **a.** Select the lookup icon to see the list of IP Services. **b.** Select **New** and create a port-to-protocol mapping. **c.** Select that port-to-protocol mapping for this **Port** field. Many commonly used protocols are already mapped to ports by default, such as `http` to port 80 and the `https` to port 443. If you want to use one of these default port-to-protocol mappings, don't enter any values in the **Port** field. The HTTP port probe uses the correct one by default.
Order	Enter a number that represents the priority of this classifier. If Discovery finds more than one classifier that applies, it uses the classifier with lower **Order** number.
Protocol [optional] This field isn't visible by default. You must configure the form to add it.	If you configure a unique port-to-protocol mapping for the **Port** field, enter an optional protocol override that Discovery uses with this classifier. For example, if you want to use port 5000 and the HTTP protocol, add the 5000/HTTP port-to-protocol mapping to the **Port** field, and then choose **http** in the **Protocol** field.

9. Right-click the header and select **Save**.
   The HTTP Classification Match related list appears. This related list shows criteria that you can configure to match against the body of the response to the GET Request that Discovery sent.
10. In the HTTP Classification Match related list, define the match criteria and specify the pattern to launch for this classifier:

    a. Select **New**.

    b. On the HTTP Classification Match form, fill in the form fields (see table).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

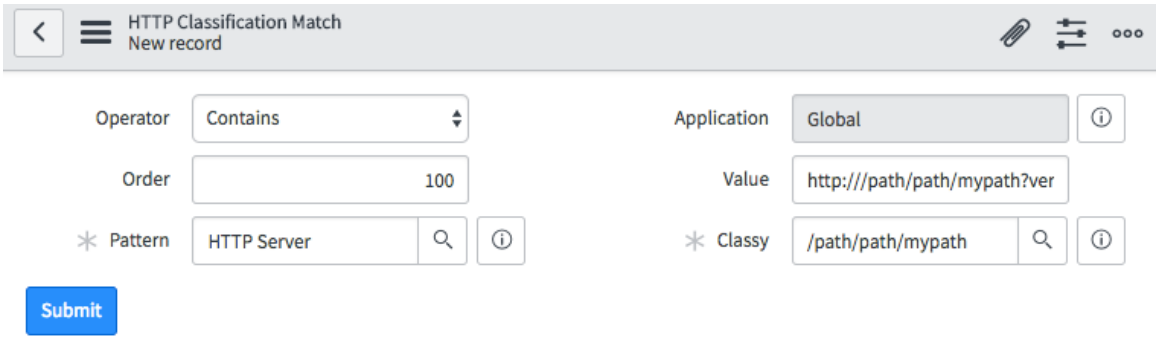

Field	Description
Operator	Choose the operator that the system uses to evaluate the condition: ▪ **Contains** ▪ **Equals** ▪ **Regex**
Order	Enter a number that represents the priority of this condition. If Discovery finds more than one condition that applies, it uses the condition with lower Order number.
Pattern	Specify the pattern to launch upon a correct match. This pattern must use HTTP to discover the device.
Value	Enter the value that the body of the returned REST method must have.

c. Select **Submit**.

d. Define as many HTTP Classification Matches as necessary.
You might want to do this if there are different patterns for different versions of the device you're trying to discover.

Example:
In this example, HTTP classification is configured to use the F5 load balancer API:

What to do next
If you don't have other types of credentials that can find devices, such as Linux credentials or SNMP credentials, you can see how the Discovery application processes HTTP discovery:

1. Navigate to **Discovery** > **Discovery Schedules** and run a discovery.
2. After the discovery process completes, check the ECC Queue to see the results of the HTTP discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You should see that after the Shazzam probe was launched, and that UNIX and SNMP classification probes failed. Both the HTTP - Classify probe and the pattern that you specify in the HTTP classification record should launch.

Topic	Name	Source	Queue
Multiprobe	Unix - Classify	10.196.39.128	input
Multiprobe	Unix - Classify	10.196.39.128	output
SNMP	SNMP - Classify	10.196.39.128	output
SNMP	SNMP - Classify 108 OIDs	10.196.39.128	input
HTTPClassyProbe	HTTPClassyProbe	10.196.39.128	input
HTTPClassyProbe	HTTPClassyProbe	10.196.39.128	output
Shazzam	wmi,snmp,ssh,wins,dns,slp,wbem,vmapp,win..	See Payload	input
Shazzam	wmi,snmp,ssh,wins,dns,slp,wbem,vmapp,win..	See Payload	output

Storage discovery

Discovery collects information on Direct Attached Storage (DAS), Storage Area Networks (SAN), and Network Attached Storage (NAS).

Storage can be located on specialized devices, such as Storage Arrays, Fibre Channel Switches, iSCSI disks, or on host operating systems, including Windows, Linux, and Solaris.

Discovery finds and maps dependencies for the following types of storage:

- Direct-attached storage (DAS), network-attached storage (NAS), or storage area network (SAN).
- NAS or SAN storage that is discovered via a Storage Management Initiative Specification (SMI-S) and Common Information Model (CIM).
- Virtual storage for VMware ESX servers and Linux Kernel-based Virtual Machines (KVM). Discovery maps this storage to the underlying physical storage.

Discovery of storage via a host reconciles data and creates relationships between the host's file systems and associated local storage devices. The local storage devices represent the storage available to the host, whether it's directly attached or provided by Fibre Channel or iSCSI. This reconciliation assumes that the storage server has been discovered first.

Note: Tape storage drivers are not discovered.

Discovery collects and creates CIs in the CMDB for the following information:

- File systems (local and NAS).
- Disks (both SAN disks and DAS drives).
- Fibre Channel (FC) HBAs and ports.
- Linux Volume Manager (LVM) volumes. LVM volume data resides in the Storage Pool [cmdb_ci_storage_pool] table.
- Veritas Volume Manager disks, subdisks, disk groups, plexes, and volumes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: For details about the discovery of direct attached or multipath block storage provisioned on a Linux host, see KB0622583.

Probes, sensors, and patterns

- KVM - Storage Pools: identifies storage attached to KVM virtual machines.
- Linux - Storage: identifies storage attached to systems running the Linux operating system.
- Solaris - Storage: identifies storage attached to systems running the Solaris operating system.
- Windows - Storage 2008: identifies storage attached to systems running Windows 2008.
  - Windows - Storage 2008 - PS: identifies storage attached to systems running Windows 2008, using the PowerShell.
  - Windows - Storage 2008 - WMI: identifies storage attached to systems running Windows 2008, using WMI Runner.
- Windows - Storage 2012: identifies storage attached to systems running Windows 2012 and later.
  - Windows - Storage 2012 - PS: identifies storage attached to Windows systems, using PowerShell.
  - Windows - Storage 2012 - WMI: identifies storage attached to Windows systems, using WMI Runner.
- VMWare - vCenter ESX Hosts Storage: collects information about ESX servers and creates relationships from datastores to underlying disks.
- Patterns for NetApp storage discovery:
  - NetApp 7-mode: finds NetApp servers via REST with two nodes.
  - NetApp cluster mode: finds NetApp servers via REST when more than two nodes are connected through a cluster interconnect switch.

Note: For information on Probe to Pattern migration see the knowledge article KB0694477.

Requirements

Windows

- Supports the following configurations:
  - Windows Server 2012, DAS or NAS with Fibre Channel (FC) or Internet Small Computer System Interface (iSCSI).
  - Windows Server 2008, DAS or NAS with FC or iSCSI.
- Install the `fcinfo.exe` tool on Windows 2008 and 2012 servers that attach to storage via FC.
- Install Windows Management Instrumentation (WMI).
- Enable Powershell on the MID Server host server.
- Provide the instance with the necessary credentials to the host server.
- Put the MID Server and the target machine in the same domain or add the target machine to the trusted host list on the MID Server machine.
- Optionally, install Windows Remote Management (WinRM) on the host server to discover Fibre Channel information. WinRM is on by default for Windows 2012 and Windows 2016 machines, but not for Windows 2008.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux

- Supports the following configurations:
  - Solaris, DAS, NAS, or SAN with iSCSI
  - CentOS, DAS, NAS, or SAN with FC or iSCSI
  - Ubuntu Server, DAS, NAS, or SAN with iSCSI
- Provide the device with SSH credentials that have root or sudo access.
- Provide the MID Server with the necessary credentials to the host server.

Relationships

Discovery creates the following relationships for storage CIs:

Parent Component	Relationship	Child Component
Storage Export [cmdb_ci_storage_export]	Exports to::Imports from	Storage Device [cmdb_ci_storage_device]
Fibre Channel Disk [cmdb_ci_fc_disk]	Provides::Provided by	File System [cmdb_ci_file_system]
iSCSI Disk [cmdb_ci_iscsi_disk]	Provides::Provided by	File System [cmdb_ci_file_system]

Configure discovery of hosts with attached storage

Discovery can find information about UNIX, Linux, and Windows hosts with attached storage.

Configure Discovery for UNIX or Linux hosts with attached storage

Discovery can find information about a UNIX or Linux host and attached storage.

Before you begin

Role required: discovery_admin

Procedure

1. **Optional:** On the host, assign elevated privileges to the account that Discovery will use.
2. On the ServiceNow instance, add credentials to the Discovery Credentials table.
3. Create a Discovery Schedule for each host IP address.
4. **Optional:** Create a behavior that uses a functionality definition with a wbem port probe to make the initial port-scanning phase (Shazzam) more efficient.
5. Run network discovery.

> **Note:** If the host also connects to a NAS or SAN storage array, set up the SMI-S Provider and CIM credentials.

Configure Discovery for Windows hosts with attached storage

Discovery can find information about a Windows host and attached storage.

Before you begin

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Follow the steps if you're using the Windows 2008 servers that is attached to storage:
   a. Download and install the Microsoft Fibre Channel Information Tool from http://www.microsoft.com and search the website for Microsoft Fibre Channel Information Tool.
   b. Set the environment path for the fcinfo executable and run the `fcinfo.exe` file.

   Note: For servers that are using newer operating systems, similar functionality is present with the operating system; this step is not required.

2. For all Windows hosts including Windows 2008, add Windows credentials to the Discovery Credentials table.
3. On the ServiceNow instance, create a Discovery Schedule for each host IP address.
4. **Optional:** Create a behavior that uses a functionality definition with a wbem port probe to make the initial port-scanning phase (Shazzam) more efficient.
5. Run network discovery.

   Note: If the host also connects to a NAS or SAN storage array, set up the SMI-S Provider and CIM credentials.

Discovery data collected for storage via a host

Discovery gathers information about storage units that connect to Linux, Solaris, and Windows hosts via a local I/O port or Host Bus Adapter (HBA).

Probes used to discover storage from a host

Table	Probes
Disk [cmdb_ci_disk]	• KVM • Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
Disk Partition [cmdb_ci_disk_partition]	• KVM • Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
Fibre Channel Disk [cmdb_ci_fc_disk]	• Linux - Storage • Solaris - Storage • Windows - Storage 2008

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Probes
Fibre Channel Port [cmdb_ci_fc_port]	• Linux - Storage • Solaris - Storage • Windows - Storage 2008
File System [cmdb_ci_file_system]	• KVM • Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
iSCSI Disk [cmdb_ci_iscsi_disk]	• KVM • Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
NAS File System [cmdb_ci_nas_file_system]	• KVM • Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
Storage HBA [cmdb_ci_storage_hba]	• Linux - Storage • Solaris - Storage • Windows - Storage 2008
Storage Pool [cmdb_ci_storage_pool]	Linux - Storage
Storage Pool Member [cmdb_ci_storage_pool_member]	Linux - Storage
VMware vCenter Datastore [cmdb_ci_vcenter_datastore]	VMware - vCenter Datastores

Storage references discovered

Discovery maps these references for storage units that connect to a host via a local I/O port or Host Bus Adapter (HBA).

Table and column	Reference	Target data element
cmdb_ci_storage_device.computer	refers to the	cmdb_ci_computer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table and column	Reference	Target data element
cmdb_ci_storage_device.provided_by	refers to the	cmdb_ci_fc_port (for FC only)
cmdb_ci_disk_partition.disk	is a partition of	cmdb_ci_disk
cmdb_ci_storage_hba.computer	is the	cmdb_ci_computer
cmdb_ci_fc_port.controller	contains the	cmdb_ci_storage_hba
cmdb_ci_fc_port.computer	is the same	cmdb_ci_computer as cmdb_ci_storage_hba.computer
cmdb_ci_storage_volume.computer	is the	cmdb_ci_computer
cmdb_ci_storage_volume.provided_by	is the	cmdb_ci_storage_pool or cmdb_ci_storage_pool or cmdb_ci_storage_device (providing storage)
cmdb_ci_storage_pool.hosted_by	is the	cmdb_ci_computer
cmdb_ci_computer	the pool is on	cmdb_ci_storage_pool.container
cmdb_ci_storage_pool.container	is the	cmdb_ci_storage_pool or cmdb_ci_storage_pool_member containing the pool (if the pool is present)
cmdb_ci_storage_pool_member.pool	is the	cmdb_ci_storage_pool
cmdb_ci_storge_pool_member.storage	is the	cmdb_ci_storage_pool, cmdb_ci_disk_partition or cmdb_ci_storage_device providing storage

Data collected for HBAs on Linux and ESX servers

Discovery collects information on host bus adapters (HBA) for fiber channel enabled devices connected to Linux and ESX servers.

HBA tables and fields

Label	Table	Field name	Source
Name	Storage HBA [cmdb_ci_storage_hba]	name	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage
WWNN	Storage HBA [cmdb_ci_storage_hba]	wwnn	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HBA tables and fields (continued)

Label	Table	Field name	Source
Model ID	Storage HBA [cmdb_ci_storage_hba]	model_id	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage
Computer	Storage HBA [cmdb_ci_storage_hba]	computer	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage

Fibre Channel ports

Label	Table	Field name	Source
WWPN	Fibre Channel Port [cmdb_ci_fc_port]	wwpn	VMWare - vCenter ESX Hosts Storage • probe • Linux - Storage
WWNN	Fibre Channel Port [cmdb_ci_fc_port]	wwnn	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage
Port type	Fibre Channel Port [cmdb_ci_fc_port]	port_type	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage
Speed	Fibre Channel Port [cmdb_ci_fc_port]	speed	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage
Controller	Fibre Channel Port [cmdb_ci_fc_port]	controller	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage
Computer	Fibre Channel Port [cmdb_ci_fc_port]	computer	• VMWare - vCenter ESX Hosts Storage probe • Linux - Storage

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fibre Channel targets

Label	Table	Field name	Source
FC Disk	Fibre Channel Targets [cmdb_fc_target]	fc_disk	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
Created by	Fibre Channel Targets [cmdb_fc_target]	sys_created_by	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
Created	Fibre Channel Targets [cmdb_fc_target]	sys_created_on	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
SysID	Fibre Channel Targets [cmdb_fc_target]	sys_id	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
Updates	Fibre Channel Targets [cmdb_fc_target]	sys_mod_count	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
Updated by	Fibre Channel Targets [cmdb_fc_target]	sys_updated_by	• Linux - Storage • Solaris - Storage

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fibre Channel targets (continued)

Label	Table	Field name	Source
			• Windows - Storage 2008 • Windows - Storage 2012
Updated	Fibre Channel Targets [cmdb_fc_target]	sys_updated_on	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
WWNN	Fibre Channel Targets [cmdb_fc_target]	wwnn	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012
WWPN	Fibre Channel Targets [cmdb_fc_target]	wwpn	• Linux - Storage • Solaris - Storage • Windows - Storage 2008 • Windows - Storage 2012

Data collected for HBAs on Solaris servers

Discovery on Solaris supports HBAs manufactured by:

- Emulex
- QLogic

Discovery populates these fields in the Storage HBA [cmdb_ci_storage_hba] table, using the Solaris - Storage probe and sensor.

Label	Field name
Name	name
Computer	computer
Device ID	device_id

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Field name
Manufacturer	manufacturer
Model ID	model_id
Serial number	serial_number

Discovery populates these fields in the Fibre Channel Port [cmdb_ci_fc_port] table.

Label	Field name
Name	name
Computer	computer
Controller	controller
WWPN	wwpn
WWNN	wwnn
Port type	port_type
Speed	speed
Operational status	operational_status

Discovery populates these fields in the Fibre Channel Disk [cmdb_ci_fc_disk] table.

Label	Name
Name	name
Storage type	storage_type
Device interface	device_interface
Device LUN	device_lun

Discovery populates these fields in the Fibre Channel Targets [cmdb_fc_target] table.

Label	Name
FC Disk	fc_disk
WWNN	wwnn
WWPN	wwpn

Data collected for HBAs on Windows servers

Discovery on Windows supports HBAs on any operating system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HBA tables and fields

Label	Table	Field Name	Source
Name	Storage HBA [cmdb_ci_storage_hba]	name	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012
WWNN	Storage HBA [cmdb_ci_storage_hba]	wwnn	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012
Model ID	Storage HBA [cmdb_ci_storage_hba]	model_id	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012
Computer	Storage HBA [cmdb_ci_storage_hba]	computer	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012

Fibre Channel ports

Label	Table	Field Name	Source
WWPN	Fibre Channel Port [cmdb_ci_fc_port]	wwpn	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012
WWNN	Fibre Channel Port [cmdb_ci_fc_port]	wwnn	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fibre Channel ports (continued)

Label	Table	Field Name	Source
Port type	Fibre Channel Port [cmdb_ci_fc_port]	port_type	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012
Speed	Fibre Channel Port [cmdb_ci_fc_port]	speed	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012
Controller	Fibre Channel Port [cmdb_ci_fc_port]	controller	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012
Computer	Fibre Channel Port [cmdb_ci_fc_port]	computer	• VMWare - vCenter ESX Hosts Storage probe • Windows - Storage 2008 • Windows - Storage 2012

Storage Discovery via SMI-S and CIM

Discovery can explore storage devices that contain a Storage Management Initiative Specification (SMI-S) provider that is a specialized Common Information Model (CIM) server.

To see the current list of vendors and products conforming to SMI-S as tested by SNIA, see the SNIA website.

Other types of storage, such as storage that is attached via a host, VM storage that is available on VMware ESX servers, and storage for Kernel VM (KVM), use a separate set of probes and sensors.

Discovery collects information about storage area networks (SAN) and network-attached storage (NAS) from specialized devices, such as storage arrays and Fibre Channel (FC) switches. Discovery collects and creates CIs in the CMDB for the following information items:

- Array disks, pools, and volumes
- Fibre Channel HBAs, ports, and controllers
- FC exports
- Fabrics, including endpoints, zoning, and switches
- Dependencies between storage sub-components

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: For more information on storage discovery and using the CIM Query tool for troubleshooting, see the knowledge article KB0782305.

Requirements

- A CIM server using SMI-S 1.4 or later.
- One or more CIM server per storage vendor.
- NAS and SAN systems from major vendors such as EMC, Hitachi, HP, and NetApp. SAN storage devices must use FC.
- FC switches from major vendors such as Brocade and Cisco.
- The CIM credentials must be available for SMI-S configuration. The CIM credentials can be different than the credentials for the system hosting the CIM server.
- Ensure that SMI-S is enabled for your storage product.
- Ensure the connectivity between the MID Server and CIM server.
- New CIM credentials need to be configured in the instance. CIM server credentials could be different from the system hosting the prerequisite CIM server.

Note: Because the SMI-S Provider caches storage device information, the Discovery query to the provider does not affect storage device performance.

SMI-S Discovery architecture

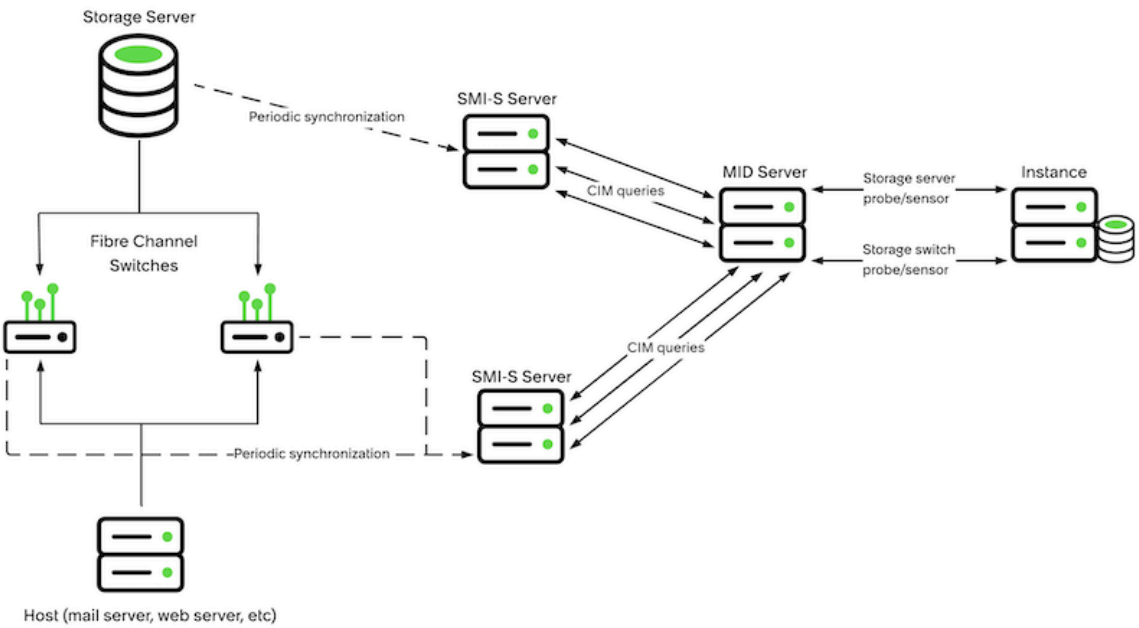

CIM architecture

CIM probes can explore any device based on the Common Information Model (CIM) by querying a CIM server, also referred to as a CIMOM - Common Information Model Object Manager. By default, Discovery uses CIM probes to explore storage systems as well as to get the serial numbers of ESX servers.

Discovery queries SMI-S compliant storage devices using CIMIQL queries.

The following components are part of CIM:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Common Information Model (CIM): CIM allows multiple parties to exchange information about managed elements. CIM represents these managed elements and the management information, while providing the mechanism to actively control and manage the elements.
- Storage Management Initiative Specification (SMI-S): SMI-S is a standard of use that describes methods for storage discovery on the vendor's side. ServiceNow uses SMI-S to determine how to discover CIM. SMI-S is based on the Common Information Model (CIM) and the Web-Based Enterprise Management (WBEM) standards, which define management functionality via HTTP. The main objective of SMI-S is to enable management of dissimilar storage products. ServiceNow supports SMI-S version 1.4 or higher.

CIM SMI-S Standard Diagram

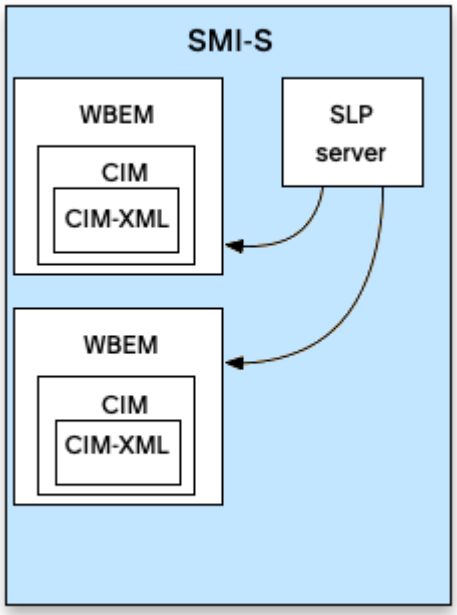

- Web-Based Enterprise Management (WBEM): WBEM defines a particular implementation of CIM, including protocols for discovering and accessing each CIM implementation.
- Service Location Protocol (SLP): SLP is an ad hoc protocol for retrieving and associating configuration information about CIM servers, such as default paths, capabilities, and the exact interop namespace. Discovery retrieves the interop namespace of a CIM server via SLP and passes that information to the CIM Classify probe. SLP, referred to here as the SLP server, uses service agents (SA) to gather and disseminate information about a CIM server on a subnet. A subnet can have multiple service agents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

The `mid.cim.interop.namespace` system property defines four default storage namespaces:

- interop
- root/interop
- root/pg_interop
- pg_interop

If you're using multiple storage vendors with custom namespaces not specified as one of the defaults, add the new namespaces to the comma-separated list in this property. If you intend to continue using any of the default namespaces, make sure to include them in the property.

CIM Agents Diagram

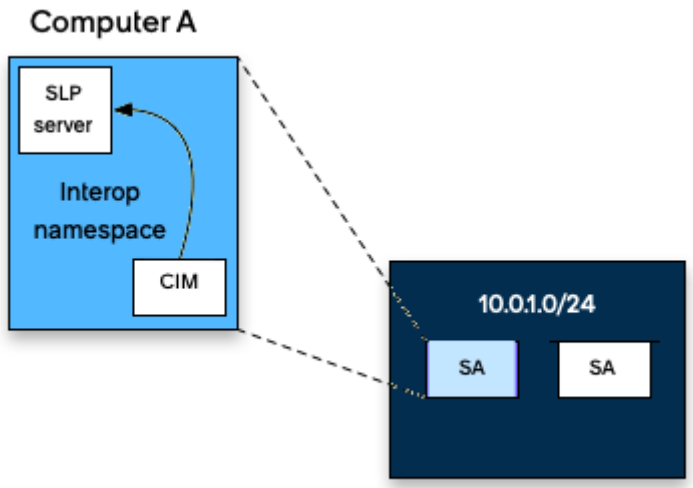

SLP and WBEM support

SLP is required for CIM Discovery as it is part of the Storage Management Initiative Specification (SMI-S) stack. Some storage devices may support the WBEM protocol, but may not support SLP.

You can manually register the WBEM services on SLP using a common Linux tool like `slptool`. This tool has a command line interface that you can use to make SLPv2 User Agent (UA) requests, which usually come with the SLP daemon package. To register a service, provide a URL and list of attributes. An example can be extracted from a working SLP server by using the same tool.

Storage Discovery table schema

This diagram displays the disk hierarchical schema for storage Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk hierarchical schema

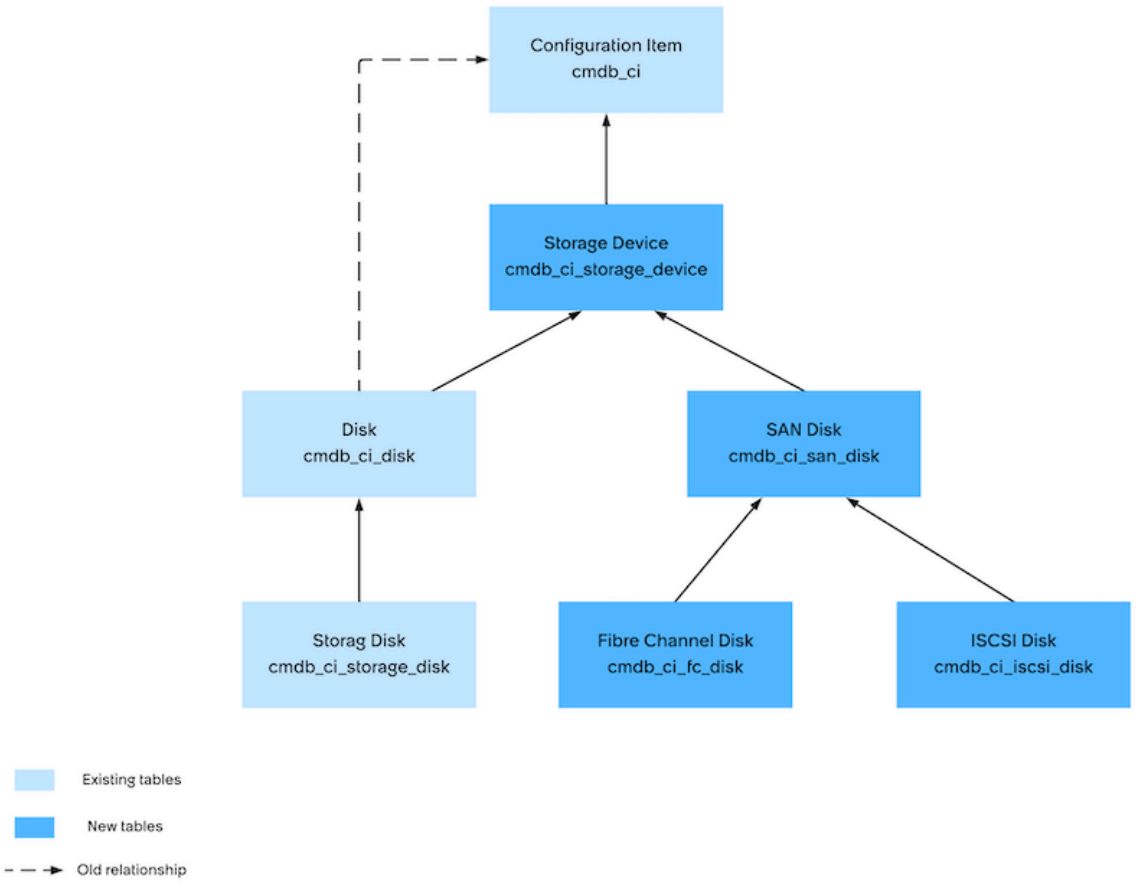

Storage Discovery tables and probes

Discovery uses the following tables and probes to gather information about storage devices that are managed by a SMI-S provider.

Data Collected

Table	Probe
Disk [cmdb_ci_disk]	SMI 1.5 - Storage Server
Fibre Channel Export [cmdb_ci_fc_export]	SMI 1.5 - Storage Server
Fibre Channel Port [cmdb_ci_fc_port]	SMI 1.5 - Storage Server
SAN Export [cmdb_ci_san_export]	SMI 1.5 - Storage Server
Storage Controller [cmdb_ci_storage_controller]	SMI 1.5 - Storage Server
Storage Export [cmdb_ci_storage_export]	SMI 1.5 - Storage Server
Storage File Share [cmdb_ci_storage_fileshare]	SMI - Array - File Shares
Storage Pool [cmdb_ci_storage_pool]	SMI 1.5 - Storage Server
Storage Server [cmdb_ci_storage_server]	CIM - Identity
Storage Switch [cmdb_ci_storage_switch]	CIM - Identity
Storage Volume [cmdb_ci_storage_volume]	SMI 1.5 - Storage Server
Storage Device [cmdb_ci_storage_device]	SMI 1.5 - Storage Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data Collected (continued)

Table	Probe
Storage Area Network [cmdb_ci_san]	SMI - Fabric
SAN Connection [cmdb_ci_san_connection]	SMI - Fabric
SAN Endpoint [cmdb_ci_san_endpoint]	SMI - Fabric
SAN Fabric [cmdb_ci_san_fabric]	SMI - Fabric
SAN Zone [cmdb_ci_san_zone]	SMI - Fabric
SAN Zone Alias [cmdb_ci_san_zone_alias]	SMI - Fabric
SAN Zone Alias Member [cmdb_ci_san_zone_alias_member]	SMI - Fabric
SAN Zone Member [cmdb_ci_san_zone_member]	SMI - Fabric
SAN Zone Set [cmdb_ci_san_zone_set]	SMI - Fabric

These probes are available but not currently used:

- SMI - Storage Server
- SMI - Array - Controllers
- SMI - Array - Disks
- SMI - Array - Pools
- SMI - Array - Ports
- SMI - Array - Volumes

Discovery of storage area networks (SAN)

Discovery collects information about storage area networks from specialized devices, such as storage arrays and Fibre Channel (FC) switches. Discovery creates specific references between the tables in the SAN schema.

ServiceNow SAN schema

This diagram shows the tables in the SAN schema and the default references defined between them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

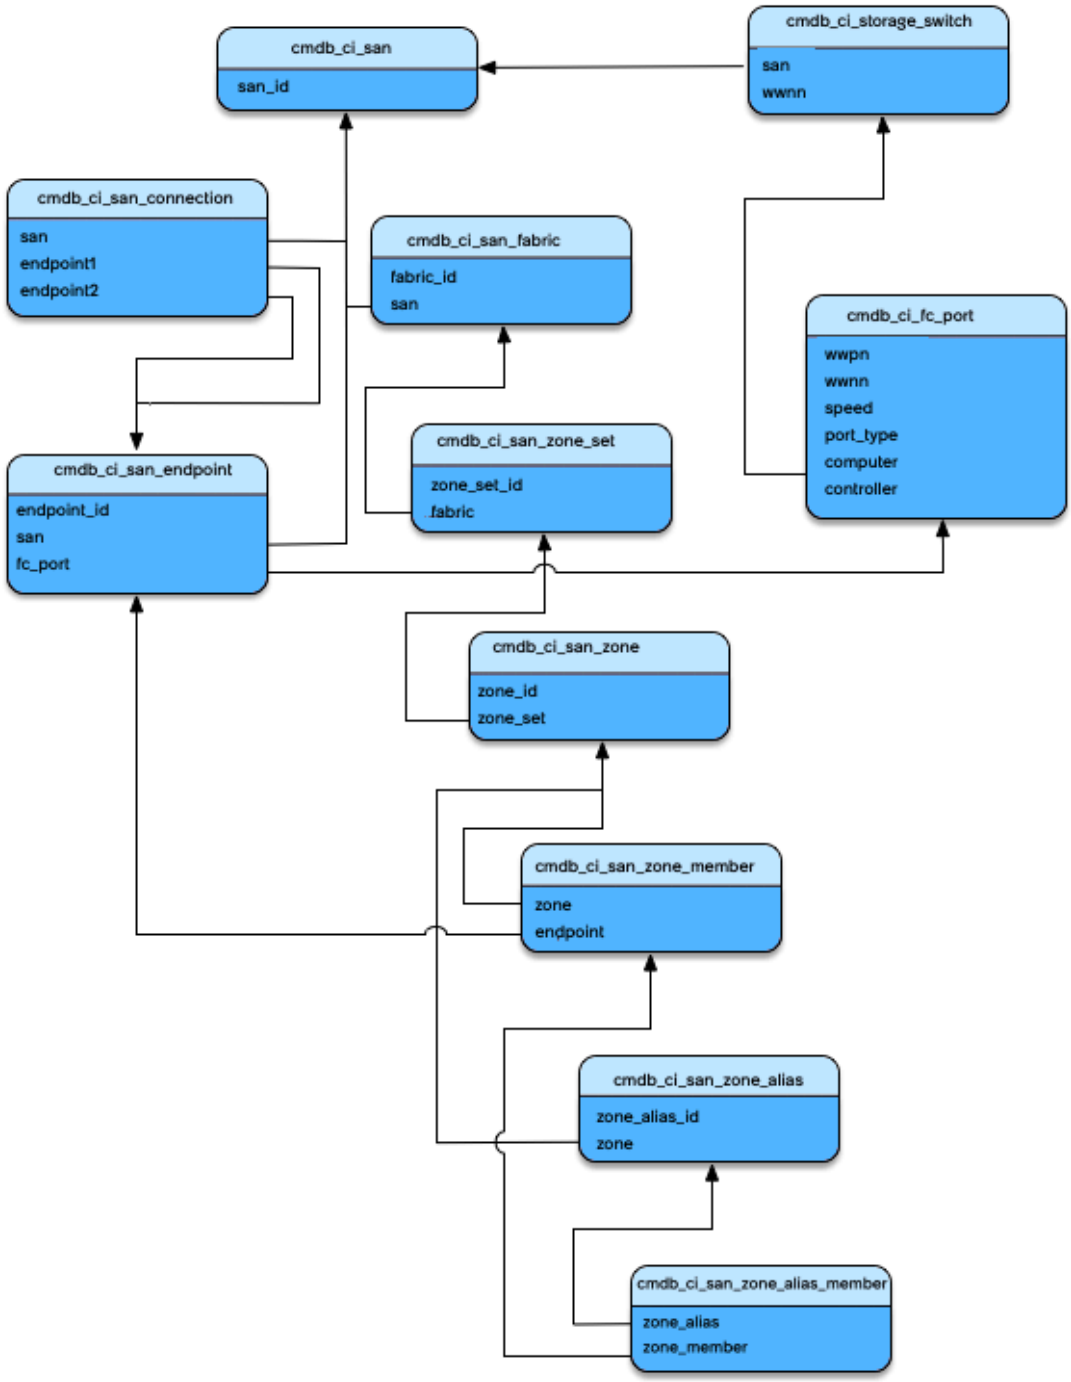

Multipath SAN storage example

This diagram shows the relationship of SAN storage volumes to a host computer. Storage accessible to the Linux host consists of two physical fibre channel (FC) disks, mpatha and mpathb, connected to the host via a multipath FC SAN. The two HBAs on the host interface are connected with two fibre cables each to separate FC switches for failover capability. The FC fabric switches are connected to two FC storage processors on the storage server. In this example, each storage processor is connected to the storage volume's, **LUN 1** and **LUN 2**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

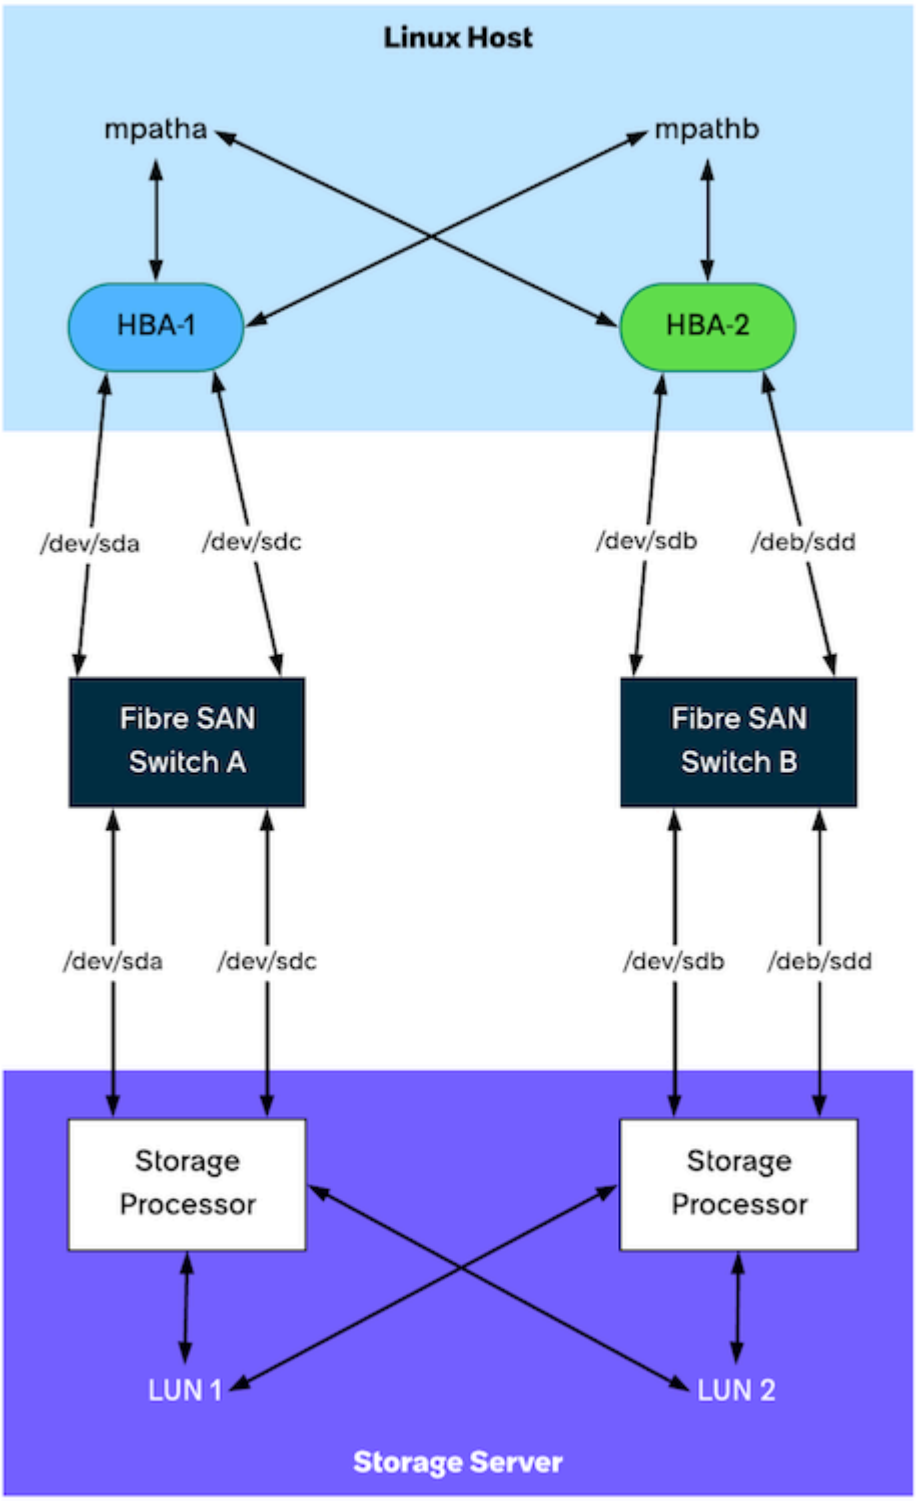

How CIM Discovery works

This is the processing flow for classifying Common Information Model (CIM) storage systems.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Processing flow

1. The Shazzam probe launches the wbem port probe as part of network discovery.
2. The wbem port probe detects activity on target ports SLP 427, CIM 5989 and 5988, and then examines the **Service Registry Queries** related list, at **Discovery Definition** > **Port Probes**, for the SLP query. The base system provides this query to detect the service:wbem service type, which indicates the presence of an SLP server.
3. The Shazzam probe launches a scanner for the WBEM service type. The scanner retrieves:
   - The attributes of the service from the SLP server.
   - The interop namespaces of CIM servers in the network.
4. The scanner appends the namespace values it finds to the port probe results.
5. The wbem port probe appends the SLP data it carries to the CIM Classify probes.
6. The CIM Classify probe uses that information to explore the CIM servers.

CIM Shazzam Processing Diagram

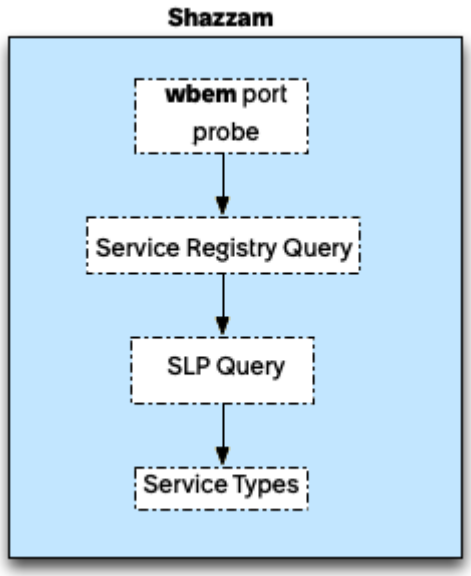

The wbem port Probe

The **wbem** probe stores the data it retrieves in the CIM Classification `[discovery_classy_cim]` table. To view the wbem port probe, navigate to **Discovery Definition** > **Port Probes**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

wbem port probe record

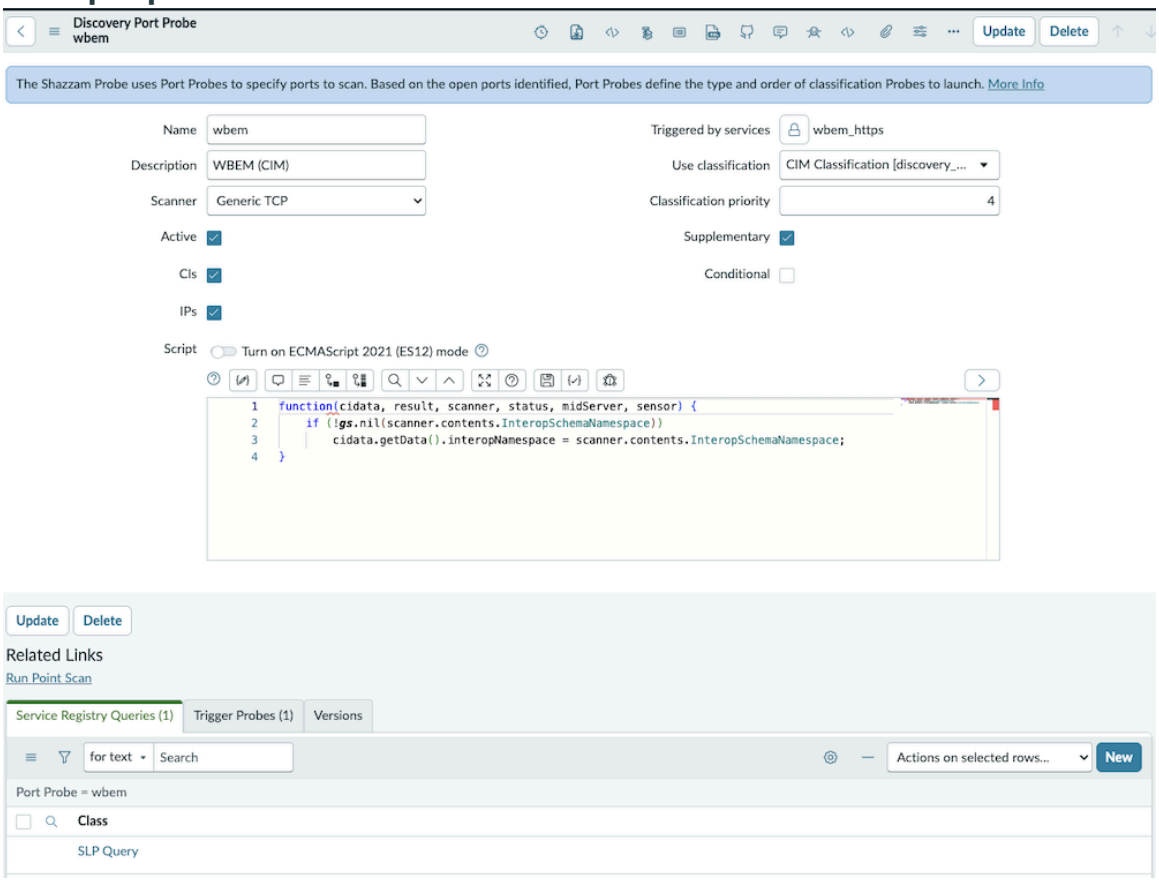

SLP query

The SLP query detects the wbem service (service:wbem) on an SLP server and gathers the attributes of the service. To view the SLP Query record, open the wbem port probe record and select **SLP Query** from the **Service Registry Queries** related list.

SLP Query record

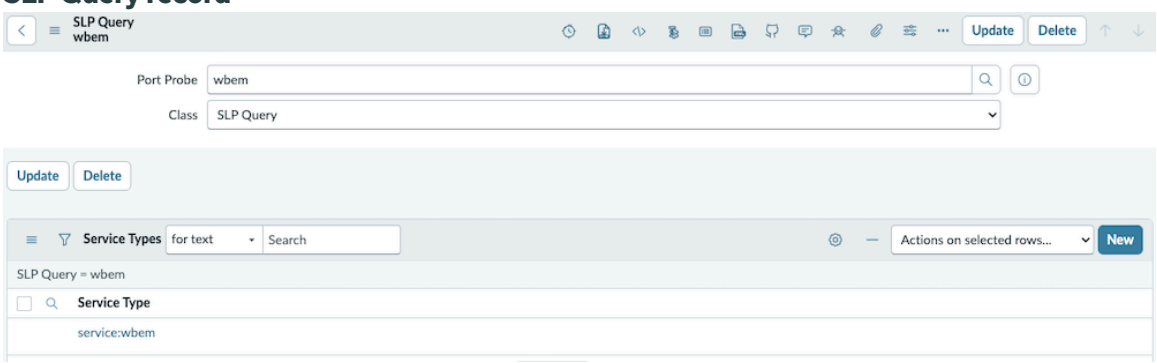

CIM - Classify probe

The wbem port probe appends the SLP data it carries, including namespaces, to the CIM - Classify probe before launching it. The CIM classification probe extracts VMware ESX serial numbers and connector relationships between the SAN and NAS components from CIM Servers in the network.

To access the CIM classification probe, navigate to **Discovery Definition** > **Probes** and select **CIM - Classify** from the list of probes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

The `mid.cim.interop.namespace` system property defines four default storage namespaces:

- interop
- root/interop
- root/pg_interop
- pg_interop

If you're using multiple storage vendors with custom namespaces not specified as one of the defaults, add the new namespaces to the comma-separated list in this property. If you intend to continue using any of the default namespaces, make sure to include them in the property.

CIM - Classify probe record

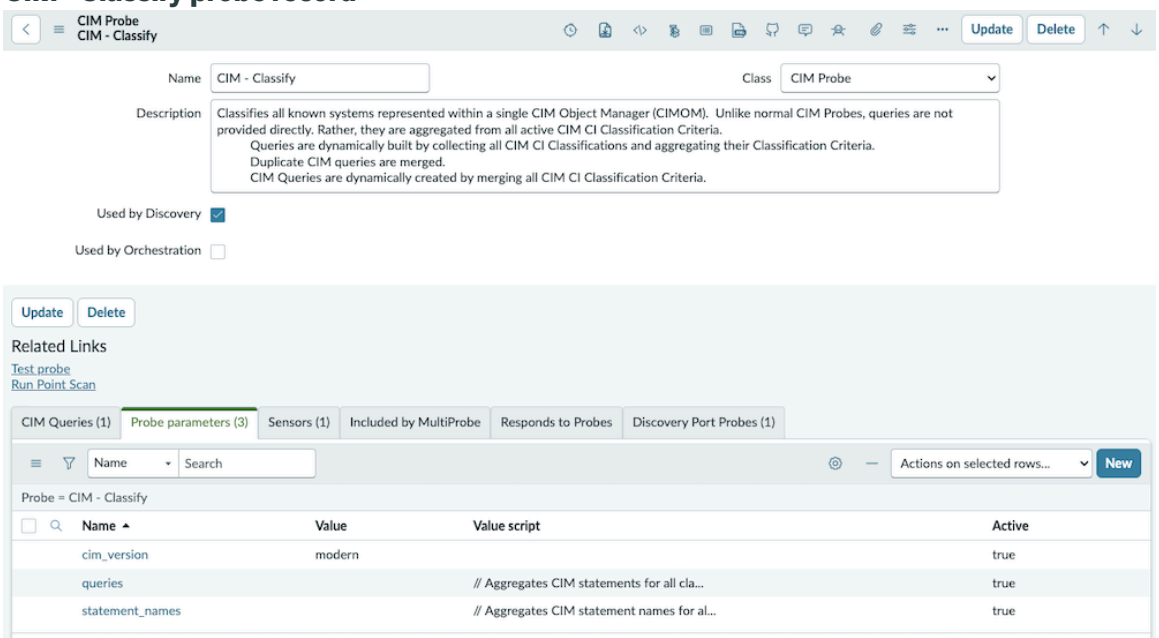

SMI-S and CIM probes and sensors

SMI-S and CIM probes and sensors

Probe/Sensor	Description
CIM - Identity	Identifies a system via CIM per SMI-S.
SMI - Array - Controllers	Retrieves controller information.
SMI - Array - Disks	Retrieves storage disk information.
SMI - Array - File Shares	Enumerates NAS file shares from a storage server.
SMI - Array - Pools	Retrieves storage pools.
SMI - Array - Ports	Retrieves storage ports.
SMI - Array - Volumes	Retrieves storage volumes.
SMI - Fabric	Retrieves SANs, fabrics, zone sets, zones, zone aliases, endpoints, and connections.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SMI-S and CIM probes and sensors (continued)

Probe/Sensor	Description
SMI - Fibre Channel Switch	Retrieves FC switches.
SMI - NAS Head - Component Systems	Retrieves all virtual file servers in a NAS head profile.
SMI - NAS Head - File Server IPs	Retrieves IP addresses for each NAS file server.
SMI - NAS Head - File Servers	Retrieves NAS file servers such as Common Internet File System (CIFS) and Network File System (NFS).
SMI - NAS Head - File Shares	Retrieves file shares for each NAS file server.
SMI - Storage Server	Retrieves SAN and NAS arrays and servers.
SMI - WBEM Service	Retrieves WBEM Service information such as profiles and SMI-S version.

Configure SMI-S Provider for storage Discovery

Use this procedure for configuring a standalone storage device with the required SMI-S Provider for Discovery.

Before you begin

Role required: admin

Procedure

1. On the storage device, if the SMI-S Provider is not present, install the SMI-S Provider software. The SMI-S Provider software is often part of the device management software. For more information, download the SMI-S Provider instructions from the storage provider manufacture.
2. For NetApp storage devices, install the SMI-S agent on the storage device host.

   Note: Discovery can also perform native discovery of NetApp servers without accessing the SMI-S server. See NetApp Server and Cluster discovery for more information.

3. Start the SMI-S Provider service.
4. In the SMI-S Provider or agent, configure the **Discovery Interval** with a synchronization rate that allows the wbem probe to receives the most current information during discovery.
5. On the ServiceNow instance, set up CIM credentials.
6. Create a Discovery Schedule with the IP addresses of each SMI-S Provider.
7. **Optional:** Create a Discovery behavior that uses a functionality definition with a *wbem* port probe to make the initial port-scanning phase (Shazzam) more efficient.
8. Run a basic IP address Discovery.

Data collected by Discovery on storage devices

Discovery identifies and classifies information about storage devices.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resources discovered on the Storage Disks [cmdb_ci_storage_disk] table

Field	Description
Name	Name of the storage disk.
Computer	Reference to the host CI.
Interface	Interface type with which the storage disk is connected to the host.
Size	Size of the storage disk.
Free disk space (GB)	Free space in GV available on the disk.
Manufacturer	Manufacturer of the storage disk.
Model ID	Model name of the disk.

Resources discovered on the Storage File Shares [cmdb_ci_storage_fileshare] table

Field	Description
Name	Name of the CI.
Path	Path to the file share.

Resources discovered on the Storage Pools [cmdb_ci_storage_pool] table

Field	Description
Name	Name of the storage pool.
Size	Size of the storage pool.
Free space	Free space available in the pool.
Hosted by	Reference to the CMDB CI.

Resources discovered on the Storage Ports [cmdb_ci_storage_port] table

Field	Description
Name	Name of the storage port.
WWPN	Unique Worldwide Port Name (WWPN) for the storage port.
Speed	Speed of the storage port.

Resources discovered on the Storage HBAs [cmdb_ci_storage_hba] table

Field	Description
Name	Name of the Storage Host Bus Adapter (HBA).
Model ID	Model name of the HBA.
Computer	Reference to the computer CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resources discovered on the Storage Servers [cmdb_ci_storage_server] table

Field	Description
Name	Name of the CI.
Manufacturer	Manufacturer of the storage server.
Model ID	Model name of the storage server.
Operating System	Operating system (OS) of the storage server.
OS Version	OS version of the storage server.
Description	Short description of the storage server.

Resources discovered on the Storage Volumes [cmdb_ci_storage_volume] table

Field	Description
Name	Name of the storage volume.
Object ID	Unique ID for the storage volume.
State	State of the storage volume.
Size	Size of the storage volume.
Storage Type	Type of storage volume.

Storage relationships

Discovery creates relationships for storage CIs running on Storage Area Networks (SAN) and Network Attached Storage (NAS).

SAN relationships

Discovery creates the following relationships for storage CIs running on Storage Area Networks.

SAN CI relationships

Parent Component	Relationship	Child Component
Storage Volume [cmdb_ci_storage_volume]	Exports to::Imports from	Fibre Channel Disk [cmdb_ci_fc_disk]
Storage Volume [cmdb_ci_storage_volume]	Exports to::Imports from	iSCSI [cmdb_ci_iscsi_disk]

NAS relationships

Discovery maps NAS file shares by resolving the hostname of a remotely mounted disk on the client computer to an IP address. It then matches the IP address to the storage server that provides the exported file system.

Discovery extracts the hostname or IP address from the NAS path to identify the storage server. If the NAS path contains a hostname, Discovery resolves it to an IP address and stores the value in the *nas_ip_address* field of the NAS File System [cmdb_ci_nas_file_system] table.

Discovery creates the following relationships for storage CIs running on Network Attached Storage (NAS). These relationships are the same for Linux and Windows hosts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NAS CI relationships

Parent Component	Relationship	Child Component
NAS File System [cmdb_ci_nas_file_system]	Allocated from::Allocated to	Storage File Share [cmdb_ci_storage_fileshare]

Data collected for Veritas Volume Manager on Linux

Discovery collects disk and volume information for Veritas Volume Manager (VxVM) on Linux hosts and maps file systems mounted on Veritas volumes to the upstream storage provider.

Note: Ensure that VxVM is correctly installed and configured. If you are using 3rd party drivers, you must configure sudo permission for *vxdmpadm*.

Table schema

VxVM tables

Table	Extends table
Veritas Disk [cmdb_ci_veritas_disk]	Storage Pool [cmdb_ci_storage_pool]
Veritas Disk Group [cmdb_ci_veritas_disk_group]	Configuration Item [cmdb_ci]
Veritas Plex [cmdb_ci_veritas_plex]	Storage Pool [cmdb_ci_storage_pool]
Veritas Subdisk [cmdb_ci_veritas_subdisk]	Storage Pool Member [cmdb_ci_storage_pool_member]
Veritas Volume [cmdb_ci_veritas_volume]	Storage Pool [cmdb_ci_storage_pool]

Data collected

VxVM data collected

Table	Fields	Relationship
Veritas Disk [cmdb_ci_veritas_disk]	• name • size_bytes • disk_group	Provided by::Provides relationship to the physical disk.
Veritas Disk Group [cmdb_ci_veritas_disk_group]	name	Contains::Contained by relationship to Veritas Disks in the group.
Veritas Plex [cmdb_ci_veritas_plex]	• name • size_bytes • layout (CONCAT, STRIPE, RAID)	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VxVM data collected (continued)

Table	Fields	Relationship
Veritas Subdisk [cmdb_ci_veritas_subdisk]	• partition_number • source_start_offset: Offset in the disk from which this subdisk is partitioned. • source_end_offset: Offset in the disk from which this subdisk is partitioned. • start_offset: Offset of this subdisk in the plex that is using it. • end_offset: Offset of this subdisk in the plex that is using it. • size_bytes	
Veritas Volume [cmdb_ci_veritas_volume]	name	Stored on::Provides storage for relationship to Veritas Disks on which the volume is stored.

Displaying data

VxVM Discovery maps file systems to the disks that supply storage. By default this is the only information displayed in the UI. The file system to disk relationship is shown as an upstream **Stored on** relationship on the File System form. To see the remainder of the information that Discovery captures for VxVM, you must add the appropriate related list to the Linux Server form. The available lists are:

- Veritas Disk Group
- Veritas Disk
- Veritas Plex
- Veritas Subdisk (Computer)
- Veritas Subdisk (Storage)
- Veritas Volume

NetApp Server and Cluster discovery

Discovery and Service Mapping find NetApp servers and clusters using patterns. Discovering some of these resources might require updating to the latest version of the Discovery and Service Mapping Patterns application from the ServiceNow Store.

Request new or enhanced Patterns on the ServiceNow® Store

Visit the ServiceNow Store to view all the available updates and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NetApp discovery isn't supported for NetApp servers in the AWS cloud. To learn about the components and versions of NetApp servers and clusters that you can discover, refer to Detailed information on products discovered by ITOM Visibility.

NetApp Server and Cluster discovery methods

Method	Supported modes	Upgrades
Discovery with patterns using ONTAP API. These patterns support native discovery of NetApp servers without requiring the discovery of the SMI-S server.	• NetApp 7-Mode • NetApp Cluster-Mode	Active by default.
Discovery with patterns using HTTP REST API.	NetApp Cluster-Mode	Available on Discovery and Service Mapping Patterns store application starting August 2023.
Discovery with probes and sensors using SMI-S.	• NetApp 7-Mode • NetApp Cluster-Mode	Legacy

NetApp discovery properties

To use discovery with patterns, navigate to **Discovery Definition** > **Properties** and enable these properties:

Property label	Name
For 7-Mode NetApp storage servers, use native discovery instead of SMI-S.	glide.discovery.sensors.netapp_native_7mode
For Cluster-Mode NetApp storage servers, use native discovery instead of SMI-S.	glide.discovery.sensors.netapp_native_cluster_mode

The following patterns support the NetApp storage discovery:

- NetApp Storage 7-Mode
- NetApp Storage Cluster-Mode
- NetApp Storage Cluster-Mode - HTTP

The Cluster-Mode patterns support chassis and disk shelf chassis discovery through an extension section.

Note: If you have already performed NetApp discovery with Probes and then enable these properties, a migration script runs the next time you perform NetApp discovery. This script could delete configuration items (CIs) for items like storage pools, which probes discover but the patterns don't. It also creates relationships.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SNMP-based discovery prerequisites

Configure SNMP community credentials

Configure SNMP community credentials to classify the NetApp server. For more information, see SNMP community credentials.

Enable read-only access to the ONTAP API

Enable **read-only** access to the ONTAP API on the NetApp server.

- NetApp Storage 7-Mode: The user must have a custom role with the following capabilities:
  - `login-snmp, login-ndmp, login-sp, and login-http-admin`
  - `Api-*`
- NetApp Storage Cluster-Mode: The user can have a read-only role with default access.

Verify read-only permission for ONTAP SDK APIs

Verify that the user has **read-only** permissions for the following ONTAP SDK APIs:

- NetApp Storage 7-Mode:
  - `aggr-list-info`
  - `cifs-share-list-iter-start`
  - `disk-list-info`
  - `fcp-adapter-initiators-list-info`
  - `fcp-adapter-list-info`
  - `igroup-list-info`
  - `iscsi-node-get-name`
  - `lun-list-info`
  - `lun-map-list-info`
  - `net-config-get-active`
  - `nfs-exportfs-list-rules`
  - `system-get-info`
  - `system-get-version`
  - `volume-list-info`
- NetApp Storage Cluster-Mode:
  - `aggr-get-iter`
  - `cifs-share-get-iter`
  - `cluster-identity-get`
  - `fcp-interface-get-iter`
  - `fcp-initiator-get-iter`
  - `igroup-get-iter`
  - `iscsi-initiator-get-iter`
  - `lun-get-iter`
  - `lun-map-get-iter`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `net-interface-get-iter`
- `net-port-get-iter`
- `storage-disk-get-iter`
- `storage-pool-get-iter`
- `storage-shelf-info-get-iter`
- `system-get-node-info-iter`
- `system-get-version`
- `system-node-get-iter`
- `volume-get-iter`
- `vserver-get-iter`

Configure basic authentication credentials

Configure basic authentication credentials for the read-only user. For more information, see Basic authentication credentials.

REST-based discovery prerequisites

Disable SNMP-based discovery

For more information, see Disable SNMP-based NetApp storage cluster discovery.

Verify read-only role and permissions

Verify that the user has **read-only** role and permissions for the following ONTAP REST APIs:

- `/api/cluster`
- `/api/cluster/chassis?fields=nodes.name,nodes.uuid,id`
- `/api/cluster/nodes?fields=name,serial_number,model,system_id,controller,version`
- `/api/cluster/nodes?fields=name,state`
- `/api/cluster/nodes?fields=name,uuid,model`
- `/api/network/ethernet/ports?fields=mac_address,name,node.name`
- `/api/network/fc/interfaces?fields=svm.name,wwnn,wwpn,port_address,name,location.node.name`
- `/api/network/fc/logins?fields=initiator,interface.name`
- `/api/network/ip/interfaces?fields=ip.netmask,svm.name,ip.address,location.home_node.name,location.`
- `/api/network/ip/interfaces?fields=location,svm,state`
- `/api/network/ip/interfaces?fields=name,ip,state,enabled,svm,location`
- `/api/protocols/cifs/shares?fields=path,name,svm.name`
- `/api/protocols/san/igroups?fields=initiators.name,name,os_type,svm.name`
- `/api/protocols/san/iscsi/sessions?fields=svm.name,initiator.alias`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `/api/protocols/san/lun-maps?fields=svm.name,lun.uuid,lun.name,lun.node.name,logical_unit_number,igr`
- `/api/storage/aggregates?fields=space.block_storage.size,space.block_storage.available,node.name`
- `/api/storage/disks?fields=sector_count,bytes_per_sector,node.name,serial_number,model,type`
- `/api/storage/luns?fields=space.used,space.size,svm.name,location.volume.name,location.vol`
- `/api/storage/pools?fields=name,capacity.total,capacity.remaining,nodes.name`
- `/api/storage/shelves?fields=name,model,serial_number,paths.node.name,id`
- `/api/storage/volumes?fields=aggregates.name,space.size,space.available,svm.name&return_timed`
- `/api/svm/svms?fields=name,subtype`
- `/api/svm/svms?fields=subtype,name`

Configure basic authentication credentials

Configure basic authentication credentials for the read-only user. For more information, see Basic authentication credentials.

General prerequisites

Verify that the following applications are up to date

- Discovery and Service Mapping Patterns
- Visibility Content

Disable SSLv3 and enable TLS for the target NetApp server

Disable SSLv3 and enable TLS for the target NetApp server, as the Java Runtime used by the MID Server doesn't enable SSLv3 by default. The key to encrypt the certificate must be at least 1024 bits.

Schedule a horizontal discovery

To run discovery using Patterns, create a Discovery schedule that targets the management IP address of the NetApp server. For more information, see create a discovery schedule.

Note: NetApp creates different storage volume identifiers when running with different protocols. When running discovery with HTTP after running discovery with SNMP, duplicate CIs are created for the same storage volume. One CI is created for the unique identifier (UUID) NetApp ONTAP API provides, and one CI is created for the unique identifier (UUID) NetApp HTTP REST API provides.

Data collected by Discovery during horizontal discovery

Discovery populates the data in the CMDB when running the NetApp patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Server [cmdb_ci_storage_server]

Field	Description
Name [name]	Name of the NetApp storage server or Storage Virtual Machine (SVM)
Manufacturer [manufacturer]	Manufacturer of the storage server. The value is set to NetApp.
Model ID [model_id]	Model identifier of the storage server. For SVMs in Cluster-Mode, the value is set to SVM.
Firmware version [firmware_version]	ONTAP firmware version running on the storage server.
Serial number [serial_number]	Unique serial number of the storage server.
Description [short_description]	Full version information of the ONTAP software.
IP Address [ip_address]	Management IP address of the storage server.
Cluster [cluster]	References the Storage Cluster [cmdb_ci_storage_cluster] table.
Physical Capacity (GB) [physical_capacity]	Total amount of physical storage space available within the NetApp system, measured in gigabytes (GB). This capacity includes both allocated and unallocated storage resources.
Free Physical Capacity (GB) [free_physical_capacity]	Available physical storage space within the NetApp system, measured in GB.
Used Physical Capacity (GB) [used_physical_capacity]	Used physical storage space within the NetApp system, measured in GB.

Storage Cluster [cmdb_ci_storage_cluster]

Field	Description
Description [short_description]	Full version information of the ONTAP software running on the cluster.
Name [name]	Name of the NetApp storage cluster.
Correlation ID [correlation_id]	Unique identifier (UUID) of the cluster.
IP Address [ip_address]	Cluster management IP address.
Manufacturer [manufacturer]	Manufacturer of the storage server. The value is set to NetApp.
Serial number [serial_number]	Unique serial number of the cluster, constructed from cluster UUID and name. This field is applicable to HTTP pattern only.
Cluster Version [cluster_version]	ONTAP version running on the cluster, for example: 9.14.1P4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Cluster [cmdb_ci_storage_cluster] (continued)

Field	Description
	This field is applicable to HTTP pattern only.

Storage Cluster Node [cmdb_ci_storage_cluster_node]

Field	Description
Name [name]	Name of the cluster node.
Operational status [operational_status]	Operational status of the cluster node. Default status is Operational.
Cluster [cluster]	References the Storage Cluster [cmdb_ci_storage_cluster] table.
Server [server]	References the Storage Node Element [cmdb_ci_storage_node_element] table.

Storage Node Element [cmdb_ci_storage_node_element]

Field	Description
Name [name]	Name of the storage node element.
Manufacturer [manufacturer]	The hardware manufacturer of the node element, which is: NetApp. This field is applicable to HTTP pattern only.
Model ID [model_id]	Model identifier of the node hardware.
Serial number [serial_number]	Serial number of the physical node.
Description [short_description]	Full version information of the ONTAP software running on the node.
RAM (MB) [ram]	Memory size of the node in megabytes (MB).
CPU core count [cpu_core_count]	The total number of CPU cores in the node element.
CPU type [cpu_type]	The processor type used in the node element.
CPU manufacturer [cpu_manufacturer]	References the Company [core_company] table.
Chassis [chassis]	References the Chassis [cmdb_ci_chassis] table. This field is applicable to Cluster-Mode patterns only.

Network Adapter [cmdb_ci_network_adapter]

Field	Description
Name [name]	Name of the network adapter port.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network Adapter [cmdb_ci_network_adapter] (continued)

Field	Description
MAC Address [mac_address]	MAC address of the network adapter.
IP Address [ip_address]	IP address.
Netmask [netmask]	Network subnet mask for the IP address.
Configuration Item [cmdb_ci]	References the Storage Server [cmdb_ci_storage_server] or Storage Node Element [cmdb_ci_storage_node_element] table. Physical resources are owned by Storage Node Elements, while logical resources are owned by Storage Servers.

IP Address [cmdb_ci_ip_address]

Field	Description
Name [name]	IP address.
IP Address [ip_address]	IP address.
Netmask [netmask]	Network subnet mask for the IP address.
MAC Address [mac_address]	MAC address of the network adapter.
Nic [nic]	References the Network Adapter [cmdb_ci_network_adapter] table.
IP version [ip_version]	Indicates whether the IP address is IPv4 or IPv6. This field is applicable for storage 7-Mode pattern only.

Disk [cmdb_ci_disk]

Field	Description
Name [name]	Name of the disk in ONTAP format.
Device ID [device_id]	Device identifier of the disk.
Manufacturer [manufacturer]	Manufacturer name of the disk.
Model ID [model_id]	The model identifier for the disk.
Serial number [serial_number]	Serial number of the disk.
Size bytes [size_bytes]	The total capacity of the disk in bytes.
Device interface [device_interface]	The interface type used to connect the disk.
Interface [interface]	Interface type of the disk.
Drive type [drive_type]	Drive type. The value is set to disk.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk [cmdb_ci_disk] (continued)

Field	Description
Storage type [storage_type]	The type of storage used by the disk. The value is set to disk.
Disk space (GB) [disk_space]	The total capacity of the disk in GB.
Size [size]	Display value of the total disk capacity in GB or terabytes (TB).
Computer [computer]	• Storage 7-Mode pattern: References the Storage Server [cmdb_ci_storage_server] table • Cluster-Mode patterns: References the Storage Node Element [cmdb_ci_storage_node_element] table.
Diskshelf chassis [diskshelf_chassis]	References the Disk Shelf Chassis [cmdb_ci_diskshelf_chassis] table. This field is applicable to Cluster-Mode patterns only.

Storage Pool [cmdb_ci_storage_pool]

Field	Description
Name [name]	The unique name of the storage pool.
Description [short_description]	A brief summary describing the storage pool.
Pool ID [pool_id]	The unique identifier assigned to the storage pool.
Size bytes [size_bytes]	The total capacity of the storage pool, in bytes.
Free space bytes [free_space_bytes]	The amount of available space in the pool, in bytes.
Size [size]	Display value of the total storage pool capacity in GB or TB.
Free space [free_space]	Display value of the available storage pool capacity in GB or TB.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] or Storage Node Element [cmdb_ci_storage_node_element] table. Physical resources are owned by Storage Node Elements, while logical resources are owned by Storage Servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage Volume [cmdb_ci_storage_volume]

Field	Description
Name [name]	Name of the storage volume.
Object ID [object_id]	Unique object identifier (UUID) for the storage volume. This field is applicable to Cluster-Mode patterns only.
Volume ID [volume_id]	The UUID of the storage volume.
LUN [lun]	The logical unit number (LUN) associated with the volume.
Size bytes [size_bytes]	Total capacity of the volume in bytes.
Free space bytes [free_space_bytes]	Available free space in the volume in bytes.
Size [size]	Display value of the total storage volume capacity in GB or TB.
Consumable Size [free_space]	Display value of the available storage volume capacity in GB or TB.
Provided by [provided_by]	References the Storage Pool [cmdb_ci_storage_pool] table.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

Storage File Share [cmdb_ci_storage_fileshare]

Field	Description
Path [path]	The file system path or mount location of the shared storage.
Name [name]	The unique name identifying the file share.

Storage Controller [cmdb_ci_storage_controller]

Field	Description
Name [name]	Name of the initiator group (igroup) that acts as the storage controller.
Device ID [device_id]	UUID of the LUN associated with this storage controller.
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

Fibre Channel Port [cmdb_ci_fc_port]

Field	Description
Name [name]	The name identifying the Fibre Channel (FC) port.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fibre Channel Port [cmdb_ci_fc_port] (continued)

Field	Description
WWNN [wwnn]	The World Wide Node Name (WWNN) of the FC port.
WWPN [wwpn]	The World Wide Port Name (WWPN) of the FC port.
Speed [speed]	Speed of the FC port in Gigabit Fibre Channel (GFC).
Computer [computer]	References the Storage Server [cmdb_ci_storage_server] table.

iSCSI Export [cmdb_ci_iscsi_export]

Field	Description
Name [name]	The name of the iSCSI export target.
LUN [lun]	The LUN mapped to the iSCSI export.
Export ID [export_id]	Unique identifier combining LUN path and controller name.
Initiator IQN [initiator_iqn]	The iSCSI Qualified Name of the connecting initiator.
IQN [iqn]	The iSCSI Qualified Name of the target being exported.
Exported by [exported_by]	References the Storage Controller [cmdb_ci_storage_controller] table.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] table.
Storage [storage]	References the Storage Volume [cmdb_ci_storage_volume] table.

Fibre Channel Export [cmdb_ci_fc_export]

Field	Description
Name [name]	Name of the FC export combining LUN path and initiator WWPN.
LUN [lun]	LUN ID number for the FC export.
Export ID [export_id]	Unique identifier combining LUN path and controller name.
Initiator WWPN [initiator_wwpn]	The WWPN of the Fibre Channel initiator.
Exported by [exported_by]	References the Storage Controller [cmdb_ci_storage_controller] table.
Hosted by [hosted_by]	References the Storage Server [cmdb_ci_storage_server] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fibre Channel Export [cmdb_ci_fc_export] (continued)

Field	Description
Storage [storage]	References the Storage Volume [cmdb_ci_storage_volume] table.

Discovery populates the data in the CMDB when running the NetApp Cluster-Mode Patterns Only.

Chassis [cmdb_ci_chassis]

Field	Description
Name [name]	Name of the cluster that the chassis belongs to.
Manufacturer [manufacturer]	The value is set to NetApp.
Model ID [model_id]	Model identifier of the chassis hardware.
Serial number [serial_number]	Serial number of the chassis.
Description [short_description]	The value is set to NetApp Storage Chassis.

Disk Shelf Chassis [cmdb_ci_diskshelf_chassis]

Field	Description
Name [name]	Identifier of the disk shelf.
Manufacturer [manufacturer]	The value is set to NetApp.
Model ID [model_id]	Model identifier of the disk shelf.
Serial number [serial_number]	Serial number of the disk shelf.
Description [short_description]	The value is set to NetApp Storage Disk Shelf Chassis.
Chassis [chassis]	References the Chassis [cmdb_ci_chassis] table.

Note: For more information on NetApp discovery with CIM, see Storage Discovery via SMI-S and CIM.

CI relationships and references

The NetApp Storage patterns create the following relationships and references to support NetApp Server and Cluster discovery. References link to records in other tables and don't appear in the CI Relationship [cmdb_rel_ci] table.

CI relationships

CI	Relationship	CI
Storage Cluster Node [cmdb_ci_storage_cluster_node]	Hosted on::Hosts	Storage Node Element [cmdb_ci_storage_node_element]
Storage Cluster Node [cmdb_ci_storage_cluster_node]	Cluster of::Clusters	Storage Cluster [cmdb_ci_storage_cluster]
Storage Node Element [cmdb_ci_storage_node_element]	Managed by::Manages	Storage Cluster [cmdb_ci_storage_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI relationships (continued)

CI	Relationship	CI
Storage Node Element [cmdb_ci_storage_node_element] or Storage Server [cmdb_ci_storage_server]*	Owns::Owned by	Network Adapter [cmdb_ci_network_adapter]
Storage Node Element [cmdb_ci_storage_node_element] or Storage Server [cmdb_ci_storage_server]*	Owns::Owned by	IP Address [cmdb_ci_ip_address]
Storage Node Element [cmdb_ci_storage_node_element] or Storage Server [cmdb_ci_storage_server]**	Contains::Contained by	Disk [cmdb_ci_disk]
Storage Node Element [cmdb_ci_storage_node_element] or Storage Server [cmdb_ci_storage_server]*	Contains::Contained by	Storage Pool [cmdb_ci_storage_pool]
Storage Server [cmdb_ci_storage_server]	Allocated to::Allocated from	Storage Node Element [cmdb_ci_storage_node_element]
Storage Server [cmdb_ci_storage_server]	Runs on::Runs	Storage Cluster [cmdb_ci_storage_cluster]
Storage Server [cmdb_ci_storage_server]	Owns::Owned by	Fibre Channel Port [cmdb_ci_fc_port]
Storage File Share [cmdb_ci_storage_fileshare]	Provided By::Provides	Storage Server [cmdb_ci_storage_server]
Storage Server [cmdb_ci_storage_server]	Contains::Contained by	Storage Volume [cmdb_ci_storage_volume]
Storage Controller [cmdb_ci_storage_controller]	Controller for::Controlled by	Fibre Channel Port [cmdb_ci_fc_port]
Storage Controller [cmdb_ci_storage_controller]	Controller for::Controlled by	Storage Server [cmdb_ci_storage_server]
iSCSI Export [cmdb_ci_iscsi_export]	Hosted on::Hosts	Storage Server [cmdb_ci_storage_server]
Fibre Channel Export [cmdb_ci_fc_export]	Hosted on::Hosts	Storage Server [cmdb_ci_storage_server]

CI references

CI	Field	Referenced CI
Storage Cluster Node [cmdb_ci_storage_cluster_node]	Server [server]	Storage Node Element [cmdb_ci_storage_node_element]
Storage Cluster Node [cmdb_ci_storage_cluster_node]	Cluster [cluster]	Storage Cluster [cmdb_ci_storage_cluster]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CI references (continued)

CI	Field	Referenced CI
Storage Node Element [cmdb_ci_storage_node_element]	CPU manufacturer [cpu_manufacturer]	Company [core_company]
Network Adapter [cmdb_ci_network_adapter]	Configuration Item [cmdb_ci]	Storage Server [cmdb_ci_storage_server] or Storage Node Element [cmdb_ci_storage_node_element]*
IP Address [cmdb_ci_ip_address]	Nic [nic]	Network Adapter [cmdb_ci_network_adapter]
Disk [cmdb_ci_disk]	Computer [computer]	Storage Server [cmdb_ci_storage_server] or Storage Node Element [cmdb_ci_storage_node_element]**
Storage Pool [cmdb_ci_storage_pool]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server] or Storage Node Element [cmdb_ci_storage_node_element]*
Storage Server [cmdb_ci_storage_server]	Cluster [cluster]	Storage Cluster [cmdb_ci_storage_cluster]
Fibre Channel Port [cmdb_ci_fc_port]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Storage Volume [cmdb_ci_storage_volume]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
Storage Volume [cmdb_ci_storage_volume]	Provided by [provided_by]	Storage Pool [cmdb_ci_storage_pool]
Storage Controller [cmdb_ci_storage_controller]	Computer [computer]	Storage Server [cmdb_ci_storage_server]
iSCSI Export [cmdb_ci_iscsi_export]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server]
iSCSI Export [cmdb_ci_iscsi_export]	Exported by [exported_by]	Storage Controller [cmdb_ci_storage_controller]
iSCSI Export [cmdb_ci_iscsi_export]	Storage [storage]	Storage Volume [cmdb_ci_storage_volume]
Fibre Channel Export [cmdb_ci_fc_export]	Hosted by [hosted_by]	Storage Server [cmdb_ci_storage_server]
Fibre Channel Export [cmdb_ci_fc_export]	Exported by [exported_by]	Storage Controller [cmdb_ci_storage_controller]
Fibre Channel Export [cmdb_ci_fc_export]	Storage [storage]	Storage Volume [cmdb_ci_storage_volume]
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	Storage Server [cmdb_ci_storage_server] or Storage Node Element [cmdb_ci_storage_node_element]*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

* The CI table (Storage Server or Storage Node Element) is determined by the resource type and owner information returned by the NetApp ONTAP API. Physical resources are owned by Storage Node Elements, while logical resources are owned by Storage Servers.

** The CI table (Storage Server or Storage Node Element) is determined by the pattern type. The Storage 7-Mode pattern uses the Storage Server table, while the Cluster-Mode patterns use the Storage Node Element table.

Discovery creates these relationships and references when running the NetApp Cluster-Mode patterns only.

CI relationships

CI	Relationship	CI
Storage Node Element [cmdb_ci_storage_node_element]	Hosted on::Hosts	Chassis [cmdb_ci_chassis]
Disk Shelf Chassis [cmdb_ci_diskshelf_chassis]	Hosted on::Hosts	Chassis [cmdb_ci_chassis]
Disk [cmdb_ci_disk]	Hosted on::Hosts	Disk Shelf Chassis [cmdb_ci_diskshelf_chassis]

CI references

CI	Field	Referenced CI
Serial Number [cmdb_serial_number]	Configuration Item [cmdb_ci]	Chassis [cmdb_ci_chassis] or Disk Shelf Chassis [cmdb_ci_diskshelf_chassis]*
Storage Node Element [cmdb_ci_storage_node_element]	Chassis [chassis]	Chassis [cmdb_ci_chassis]
Disk Shelf Chassis [cmdb_ci_diskshelf_chassis]	Chassis [chassis]	Chassis [cmdb_ci_chassis]
Disk [cmdb_ci_disk]	Diskshelf chassis [diskshelf_chassis]	Disk Shelf Chassis [cmdb_ci_diskshelf_chassis]

* The CI table (Chassis or Disk Shelf Chassis) is determined by the serial number type field.

Disable SNMP-based NetApp storage cluster discovery

Disable the default SNMP-based NetApp storage cluster discovery and use REST-based discovery instead for enhanced security.

Before you begin

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery Definition** > **CI Classification** > **SNMP**.
2. In the **Name** field, search for `NetApp Storage Server Cluster-Mode`.
3. Select the NetApp Storage Server Cluster-Mode record.
4. Clear the **Active** check box.
5. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Configure basic authentication credentials for a user with a read-only role. For more information, see Basic authentication credentials.

Related topics

NetApp Server and Cluster discovery

Manage large storage payloads in Linux and Solaris

Large payloads for Linux and Solaris direct attached storage can cause out of memory errors if not configured to serialize the processing of the payload.

Before you begin

Role required: discovery_admin or admin

About this task

When Linux or Solaris storage sensors process payloads from large disk arrays, worker threads run simultaneously and can run an instance node out of memory, requiring a restart. To prevent this, the [discovery_sensor_yield] table controls the behavior of worker threads used to process large payloads. When Discovery processes a sensor, it checks the [discovery_sensor_yield] table for that sensor. If the sensor is listed in the table, and if the payload size exceeds the configured limit in the table, then the worker attempts to secure a lock before processing the payload. If the lock is not available because another worker is already processing a large payload, the waiting worker re-schedules the next job for 30 seconds later and moves on to other tasks. The job then waits in the Schedule Item [sys_trigger] table for an idle worker to pick it up. This serialization occurs per node, enabling multiple nodes to process large payloads simultaneously. The processing interval of 30 seconds is not configurable.

Important: These threshold values were created for Linux and Solaris storage sensors only and, in most cases, should not be changed. If your instance is running out of memory while simultaneously processing multiple, large storage payloads, consider scheduling Discovery of these systems to times of lower activity. If this is not possible, examine the ECC queue to determine which storage sensor might be causing the issue and note the size of the payload. If the payload size is below that of the threshold in the [discovery_sensor_yield] table, lower the threshold value enough to serialize that payload. If a sensor other than the Linux or Solaris storage sensor is responsible, you can add that sensor to the [discovery_sensor_yield] table. Make sure to set the threshold value correctly for the payload size. Serializing sensor processing can slow Discovery.

Procedure

1. In the application navigator filter, enter **discovery_sensor_yield.list**.
   The base Discovery system serializes these sensors when their payloads exceed 512 KB:
   - **Solaris - Storage**
   - **Linux - Storage**
   - **Unix - ADM Enhanced**
   - **Windows - ADM Enhanced**
2. To change the threshold, double-click in the **Value** column in the list and edit the size.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Editing the sensor yield threshold

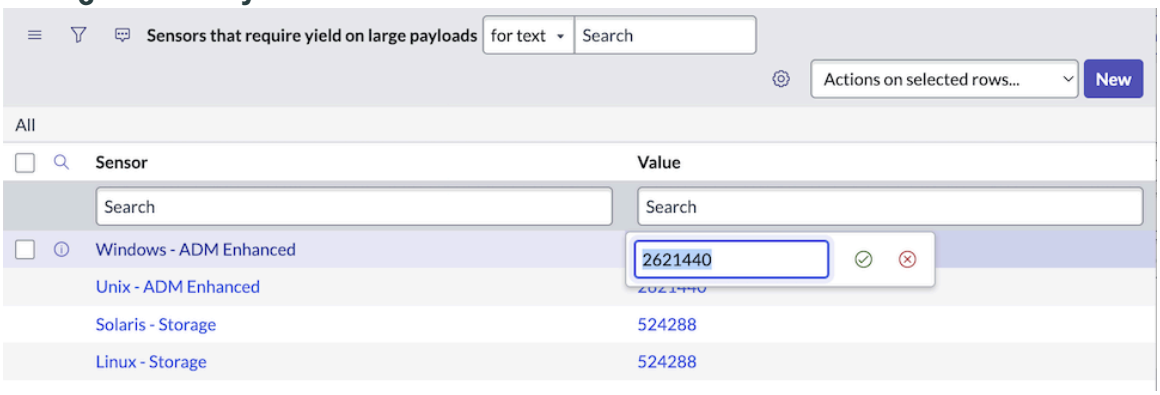

3. To add another sensor and configure it for payload serialization, click **New** and complete the fields in the form.
   - **Sensor**: Storage sensor to configure.
   - **Value**: Threshold size for the sensor.
4. Click **Submit**.

Storage discovery examples

Discovery creates configuration items (CIs) and CI relationships for physical and logical storage components. Storage can be attached directly to servers or through fibre channel switched fabric.

Direct attached storage

In this example of direct attached storage (DAS), a SCSI drive with two partitions, `/dev/sda1` and `/dev/sda2`, is attached to a Linux host. The `/dev/sda1` partition is bootable and supports the system software. The `/dev/sda2` partition contains a logical volume configured as a storage pool and mounted to the Linux root file system by Logical Volume Management (LVM). The storage pool uses only 1.5GB of the partition, leaving 13GB of storage available for additional logical volumes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Direct attached storage example

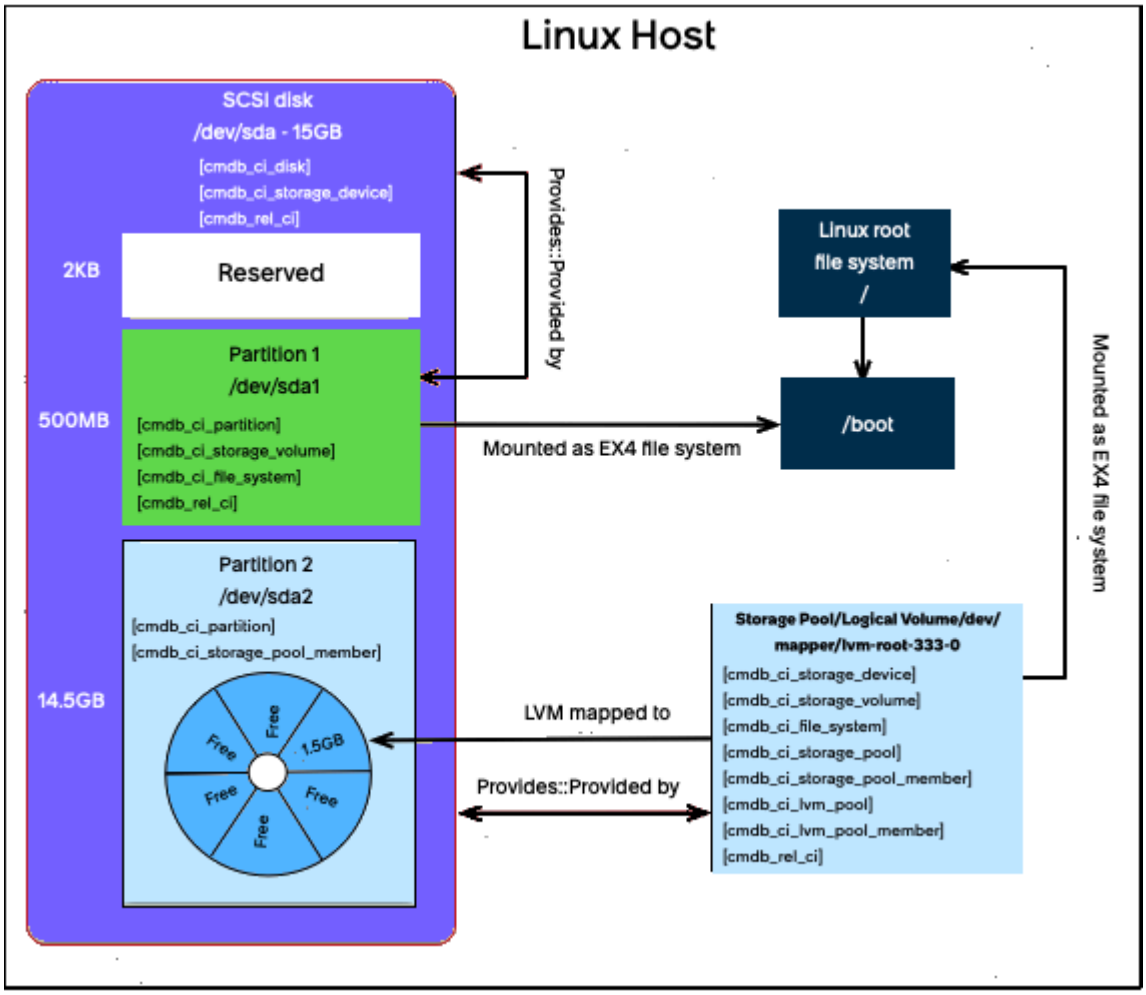

CIs and relationships created for direct attached storage (DAS)

Configuration item	Description	Tables	Key reference and Relationships
/dev/sda	SCSI physical storage device	• [cmdb_ci_disk] • [cmdb_ci_storage_device] • [cmdb_rel_ci]	[cmdb_rel_ci] • **Provides**: /dev/sda • **Provided by:**/dev/sda1
			[cmdb_rel_ci] • **Provides**: /dev/sda • **Provided by:**/dev/mapper/lvm-root-333-0

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CIs and relationships created for direct attached storage (DAS) (continued)

Configuration item	Description	Tables	Key reference and Relationships
/dev/sda1	Partition 1 on the SCSI storage device	• [cmdb_ci_partition] • [cmdb_ci_storage_volume] • [cmdb_ci_file_system] • [cmdb_rel_ci]	[cmdb_ci_file_system] • **Mount point**: /boot • **File system**: Ext4 [cmdb_rel_ci] • **Provides**: /dev/sda • **Provided by:**/dev/sda1
/dev/sda2	Partition 2 on the SCSI storage device	• [cmdb_ci_partition] • [cmdb_ci_storage_pool_member]	[cmdb_ci_storage_pool_member] • **Pool**: /dev/mapper/lvm-root-333-0 • **Storage**: /dev/sda2
/dev/mapper/lvm-root-333-0	Linux logical volume, mapped with LVM to a physical disk storage partition.	• [cmdb_ci_storage_device] • [cmdb_ci_storage_volume] • [cmdb_ci_file_system] • [cmdb_ci_storage_pool] • [cmdb_ci_storage_pool_member] • [cmdb_ci_lvm_pool] • [cmdb_ci_lvm_pool_member] [cmdb_rel_ci]	[cmdb_ci_file_system] • **Mount point**: / • **File system**: Ext4 [cmdb_rel_ci] • **Provides**: /dev/sda • **Provided by:**/dev/mapper/lvm-root-333-0 [cmdb_ci_storage_pool_member] • **Pool**: /dev/mapper/lvm-root-333-0 • **Storage**: /dev/sda2

Multipath fibre channel storage

In this example of a fibre channel storage area network (SAN), two physical storage devices, `mpatha` and `mpathb`, are attached to a Linux host through fibre switches. The fibre switches provide failover capabilities. The `mpatha` drive contains two partitions, `mpatha1` and `mpatha2`. The first partition is mounted directly to `/boot` on the Linux host. Three logical volumes are mapped to the `mpatha2` partition and to the physical device `mpathb`. The logical volumes are mounted as Ext4 file systems in folders on the Linux root structure. This example shows the CIs that Discovery manages for each component and the mounting points for the logical volumes on the Linux host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Multipath fibre channel storage example

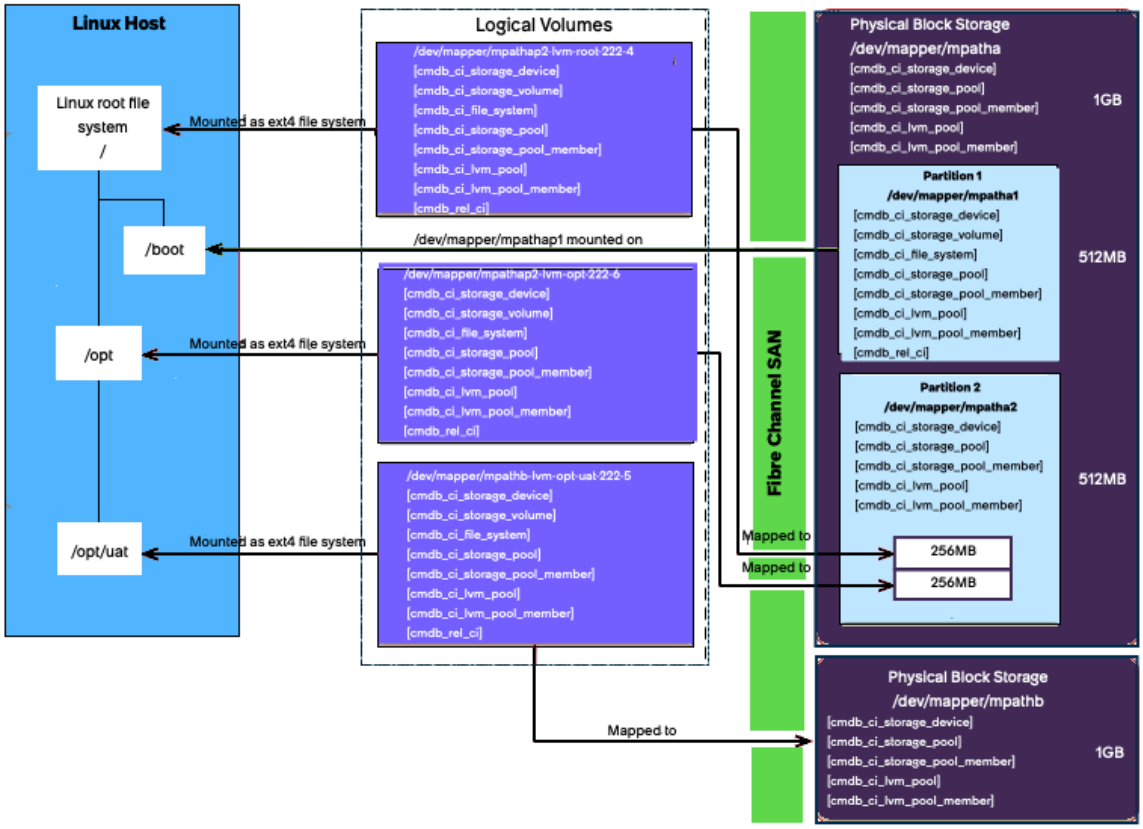

Switched fibre fabric details

Discovery creates CIs for the logical sub-components in NAS and SAN environments, such as fibre channel disks and pool components. Discovery also creates CIs for host bus adapters (HBA) and physical block storage. In multipath environments, Discovery creates CI relationships within the switched fibre fabrics that connects the Linux host to the physical storage devices. In this diagram, the fibre fabrics have redundant paths that the SAN environment can use for failover if connections fail.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fibre fabric redundant paths

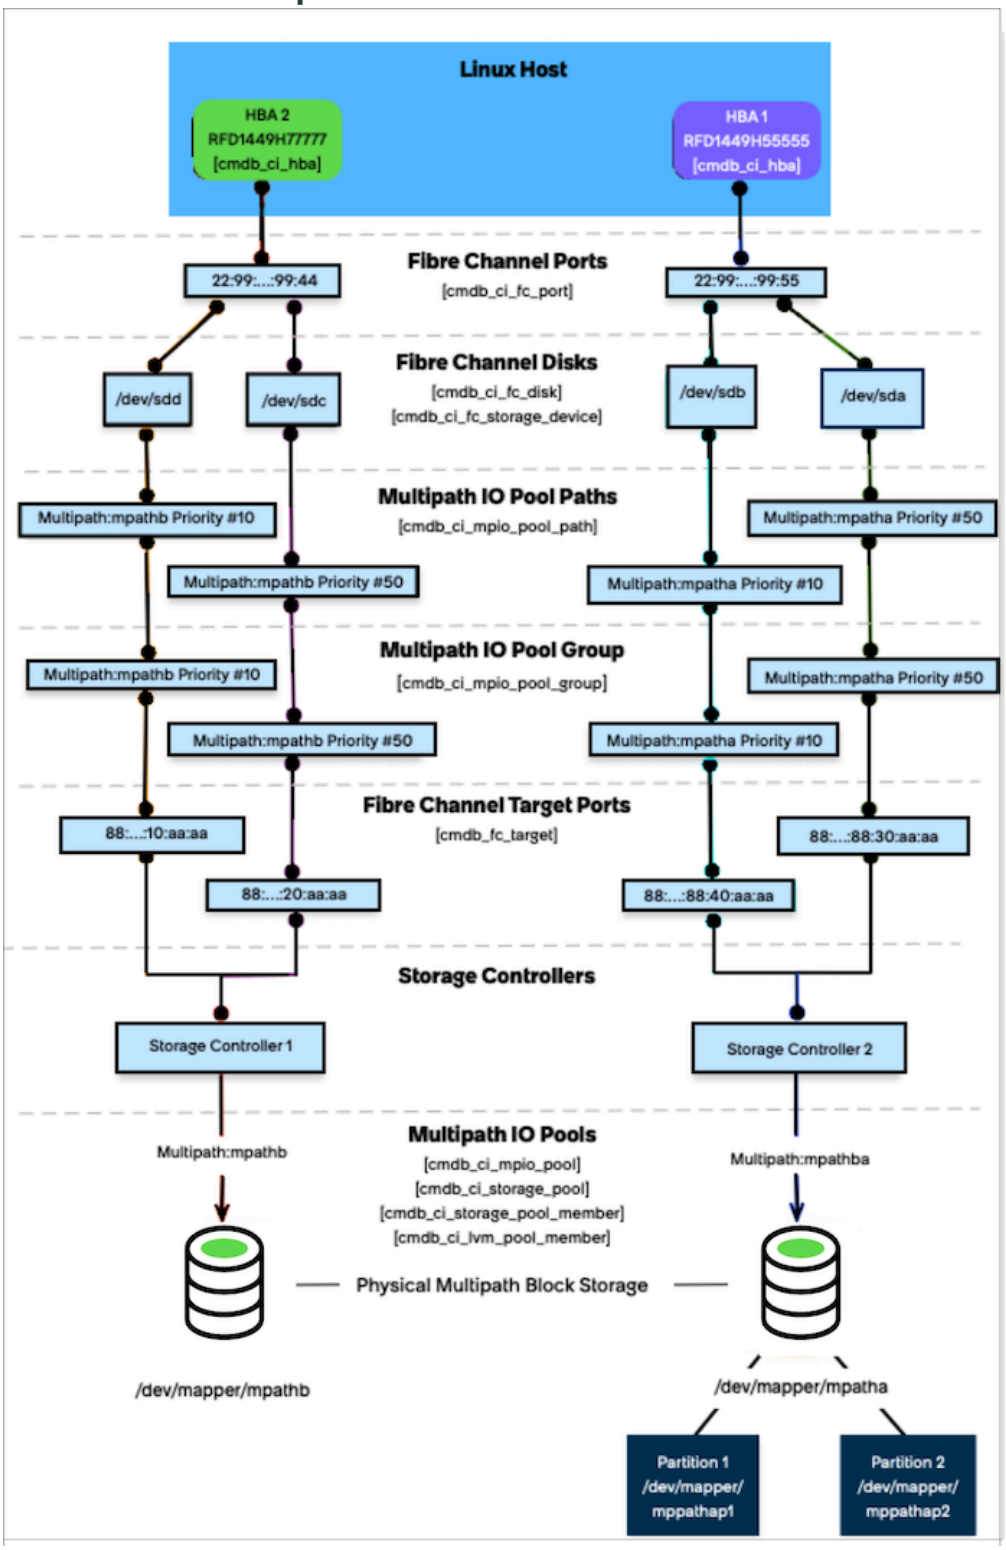

ITOM Visibility and CSDM reference

ITOM Visibility consists of two ServiceNow products: Discovery and Service Mapping. These products are responsible for creating Configuration Items (CIs) in the CMDB and relating them. The goal of this product view is to help you to understand how ITOM Visibility works with the core CSDM framework.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Aspects of ITOM Visibility

Discovery

Discovery finds computers, servers, printers, a variety of IP-enabled devices, and the applications that run on them. It can then update the CIs in your CMDB with the collected data.

Service Mapping

Service Mapping discovers all application services in your organization and builds a comprehensive map of all devices, applications, and configuration profiles that support the associated business services.

Service Mapping maps dependencies based on the connections between devices and applications. This method, called top-down mapping, helps you immediately see the impact of a problematic object on the rest of the application services.

Implementing CMDB data

For details on implementing CMDB data to conform with the CSDM framework, see Implementing the CSDM framework in stages.

For more information

Exploring the CSDM model.

Implementing the CSDM framework in stages.

Discovery basics.

Service Mapping.

ITOM Visibility and CMDB tables

ITOM Visibility manages and uses CMDB tables. Several ServiceNow products benefit from and add value to ITOM Visibility.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables that ITOM Visibility manages

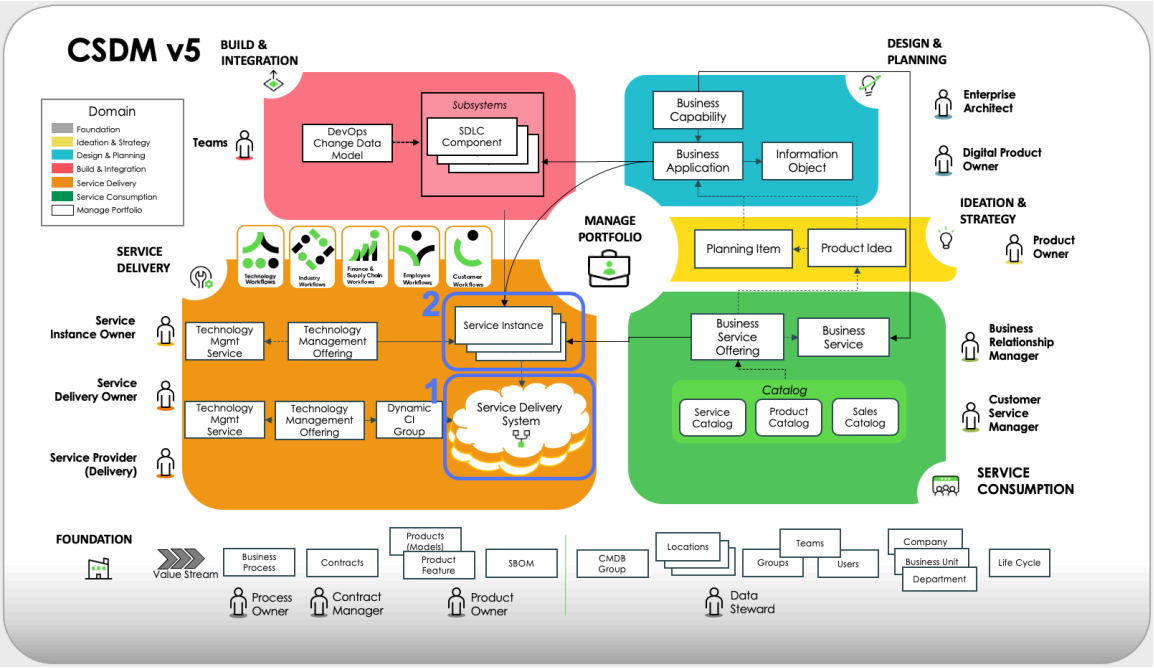

Mapped application service tables

The tables that ITOM Visibility manages depend on the method used to map the service instance (the label for the table was formerly application service) . You can use the Application Service wizard to populate an application service.

- Map the Service instance using top-down discovery of CI connections. This method requires Service Mapping, the sm_admin role, and the current domain must be a leaf domain. This is the best option for accuracy and complete mapping. ITOM Visibility manages the Mapped Application Service table [cmdb_ci_service_discovered]
- Map the service instance using a tag list or tag family. This method requires Service Mapping and is preferred for public or private cloud where tags are typically used. ITOM Visibility manages the tag-based maps table [cmdb_ci_service_by_tags].
- Manually map the service instance. This method requires that you select specific CIs — no automation is used. ITOM Visibility manages the service maps table [cmdb_ci_service_manual].
- Use Dynamic Service to synchronize manually-created CI relationships to a Service Map view.

Configuration items (CIs) tables [cmdb_ci_*]

- Application table [cmdb_ci_appl]
- Server table [cmdb_ci_server]
- Virtual machines table [cmdb_ci_vm_instance]
- Load balancer table [cmdb_ci_lb]
- Network gear table [cmdb_ci_netgear]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

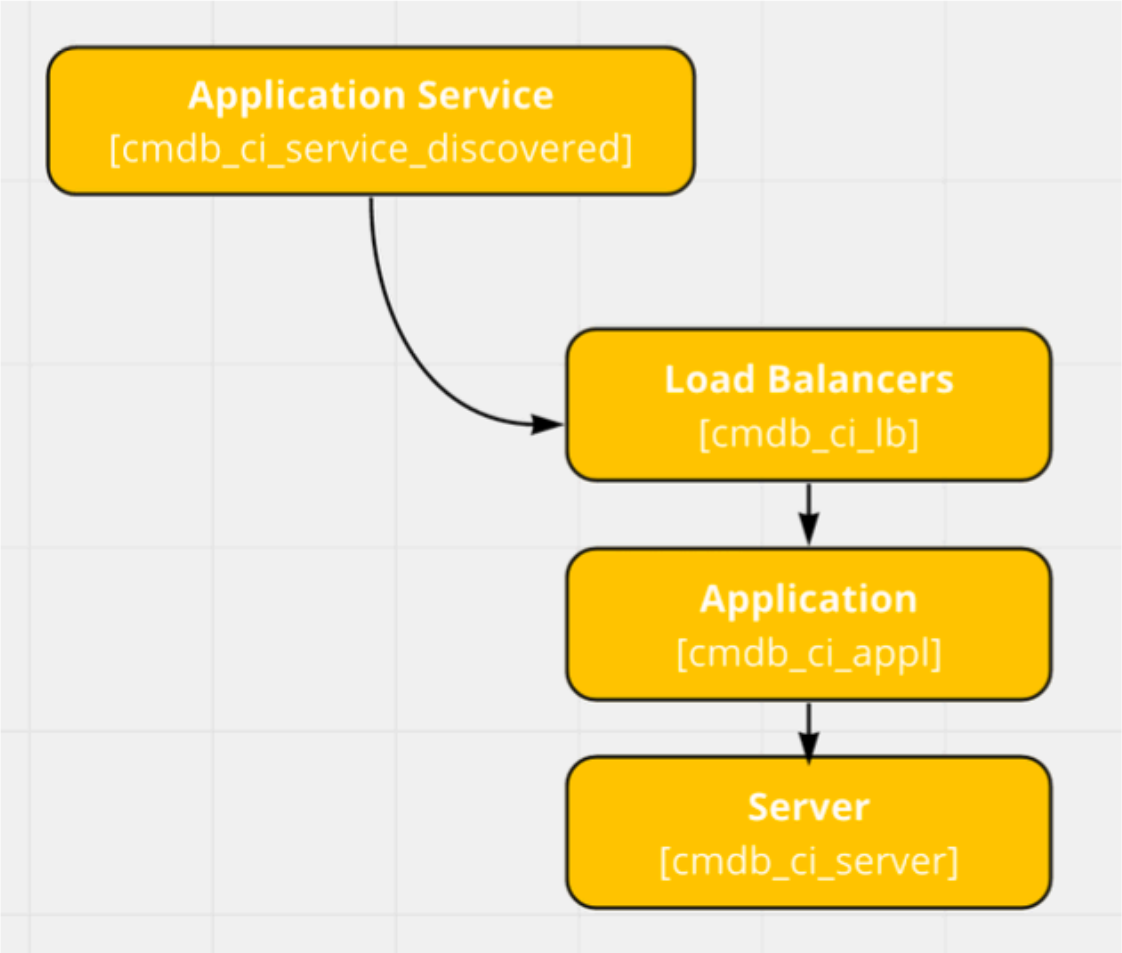

Tables that ITOM Visibility uses

1. Dynamic CI Group table [cmdb_ci_query_based_service]
2. Product model tables (for Technology Portfolio Management [TPM]) software and hardware models

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

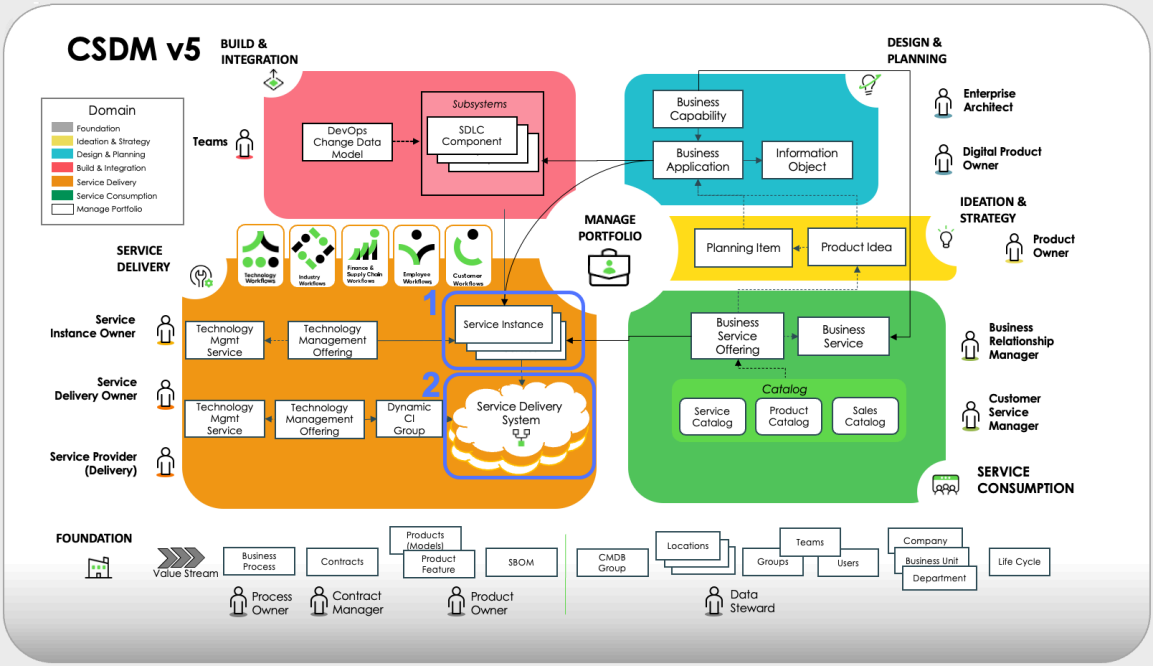

Products that benefit from ITOM Visibility

- IT Service Management
- Customer Service Management (CSM)
- Event Management
- Cloud Management
- Asset Management (Hardware Asset Management and Software Asset Management)
- Strategic Portfolio Management (Financial Management — Showback statements, Application Portfolio Management [APM])
- Security Operations (Incident Response, software vulnerability management)

Implementing CMDB data

For details on implementing CMDB data to conform with the CSDM framework, see Implementing the CSDM framework in stages.

ITOM Visibility use case

The ITOM Visibility use cases are described in this section.

Discovery use case

- Gain increased visibility into your on-premises and cloud resources.
- Keep track of changes in your on-premises, cloud, and serverless infrastructure in the Configuration Management Database (CMDB).
- Set a strong foundation with accurate data and relationship views for ServiceNow products, including Change Management, Software Asset Management, Customer Service Management (CSM), and Security Operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key features of the Discovery use case

- Set up and manage Discovery jobs.
- See IT resources and dependencies at-a-glance.
- Build queries to validate discovered IT resources.
- Manage your public key infrastructure (PKI) certificates in one dashboard.

Results of the Discovery use case

The CSDM framework gives Discovery a prescribed model for how infrastructure and software CIs relate to other areas, such as application services, business applications, and assets.

Service Mapping use case

You can create application services manually, using an API, or by having Service Mapping discover them. Regardless of creation method, all application services are stored in the Mapped Application Service table [cmdb_ci_service_discovered].

Results of the Service Mapping use case

Service Mapping automates a critical aspect of CSDM. It provides a consistent, automated approach to connect the logical layer of the CSDM model to the physical model CIs in the CMDB. This approach lets you more effectively manage your business applications. You can automate modeling of your business applications for impact assessment and analysis. For impact assessments and analysis, you can use a number of ServiceNow products including Change Management, Incident Management, or the CMDB Query Builder.

Application Services flow

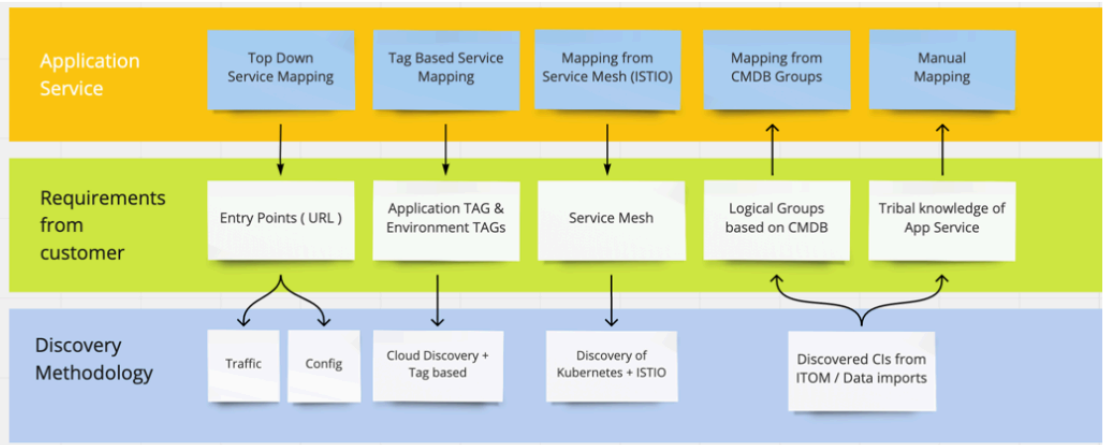

Application services view in Operator Workspace

You can specify a life-cycle status for an application service. Application services with an Operational life-cycle status can be used on service maps for the relevant workflows, such as Artificial Intelligence for IT Operations (AIOPs) workflows.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

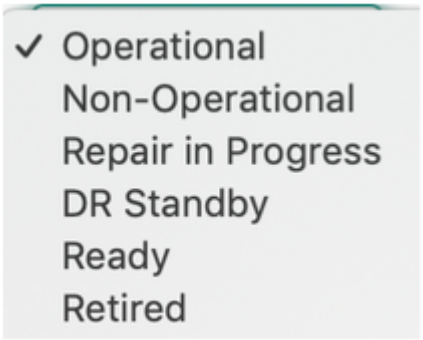

You can view information about the services, such as the criticality, in the Operator Workspace.

Operator Workspace

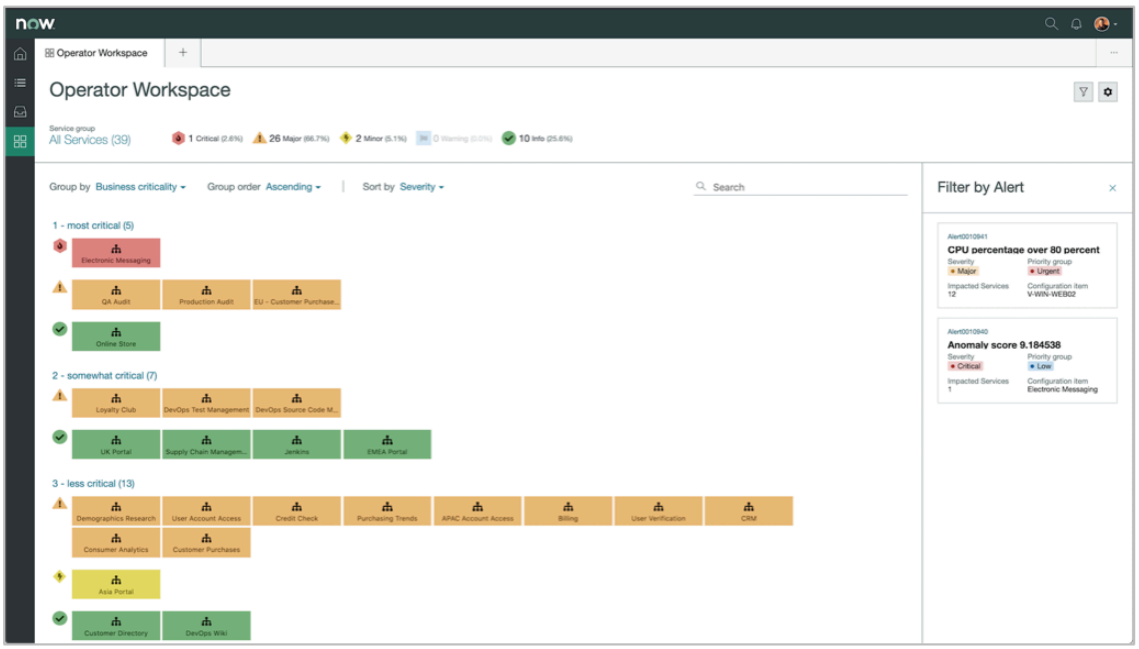

The Service Map tab displays a visual representation of the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service map diagram

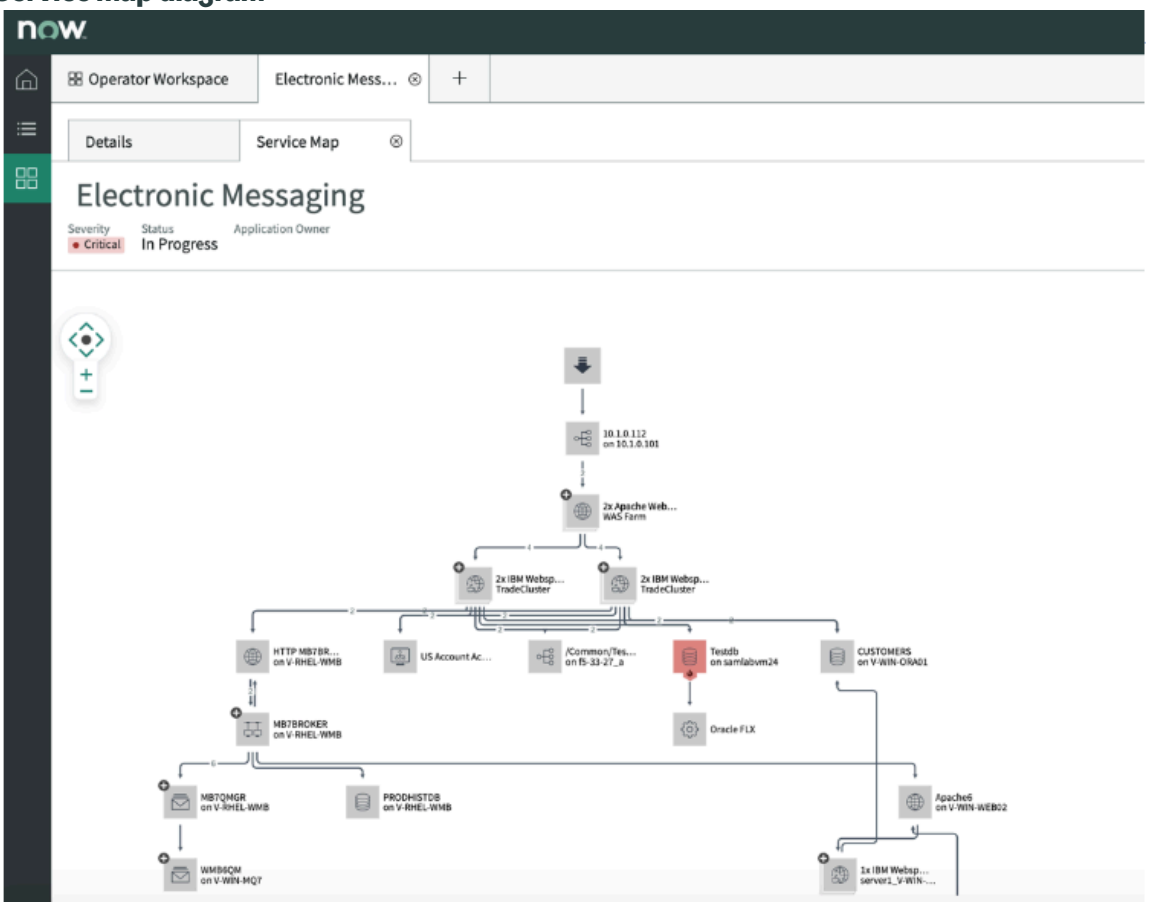

Related topics

- Life cycle of document and contract entities
- Life cycle of tangible/physical CIs
- Life cycle of location entities
- Life cycle of intangible/logical entities
- Life cycle of product entities

ITOM Visibility considerations for CSDM framework

Consider these points while implementing the CSDM framework.

You can use Service Mapping to define application services.

Here are alternative methods for defining application services:

- Define manually.
- Use a Dynamic CI group, based on a query.
- Use a tag-based approach (commonly used for Cloud integrations).

For more information about creating application services, see Create an application service.

Run options for discovery schedules

Both horizontal and top-down discovery use these run options.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery schedule run options

Run option	Description
On Demand	Doesn't run on a schedule. Select **Discover now** to run Discovery. This is the default
Daily	Runs every day. Use the **Time** field to specify the time of day.
Weekly	Runs on one designated day of each week. Use the **Time** field to specify the time of
Monthly	Runs on one designated day of each month. Use the **Day** field to select the day of th specify the time of day. If the designated day doesn't occur in the month, the schedu example, if you designate day 30, the schedule doesn't run in February.
Periodically	Runs every designated time. Use the **Repeat Interval** field to define the time in days, The first Discovery runs at the point in time defined in the **Starting** field. The subsequ **Repeat Interval** period passes. **Note:** For schedules configured as "Periodically", the "Time" column won't disp time. To determine the next execution time for schedule that is configured to run Schedule [sys_trigger] table and locate the corresponding schedule record. The the next scheduled run.
Once	Run one time as designated by the date and time defined in the **Starting** field.
Weekdays	Runs every Monday, Tuesday, Wednesday, Thursday, and Friday. Use the **Time** field t
Weekends	Runs every Saturday and Sunday. Use the **Time** field to select the time of day.
Month Last Day	Run the last day of every month. Use the **Time** field to select the time of day.
Calendar Quarter End	Runs on March 31, June 30, September 30, and December 31. Use the **Time** field to s the dates, modify the *DiscoveryScheduleRunType* script include.
After Discovery	Allows you to sequentially stagger the schedule. Use this option to run this schedule in the **Run after** field finishes. Select the **Even if canceled** check box to designate tha even if the **Run after** Discovery is canceled before it finishes. • This option isn't valid when Discovery is started via DiscoverNow, or when using th • You can't designate an inactive Discovery schedule. • You can't create a loop by designating the run after Discovery to be the same Disc • This Discovery doesn't run if the **Run after** Discovery doesn't finish, with the excep check box is selected and the Discovery is canceled.

Related topics

Schedule a top-down discovery by Service Mapping

Schedule a horizontal discovery

Plugins or applications installed with ITOM Visibility

Tables that list the plugins or applications that are installed with ITOM Visibility applications. When you update your application, any newly required application dependencies are installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery

Plugins or applications installed with Discovery (com.snc.discovery)

Plugin or application name	Description
Cloud Provisioning and Governance Core (com.snc.cloud.core)	Contains resource blocks and cloud Application Programming Interface (cloud API) mappings related to Cloud API-based discovery and events processed by either response mapping or calling patterns.
Service Watch Suite commons (com.snc.service-watch.commons)	Contains common UI components for Discovery and Service Mapping.
Discovery - IP Based (com.snc.discovery.ip_based)	Contains shared components of horizontal discovery used by Discovery and Service Mapping.

Discovery Admin Workspace

Plugin activated with Discovery Admin Workspace (sn_disco_workspace)

Plugin or application name	Description
Discovery (com.snc.discovery)	Finds applications and devices on your network, and then updates the CMDB with the information it finds.

Service Mapping

Plugins or applications installed with Service Mapping (com.snc.service-mapping)

Plugin or application name	Description
Service Mapping – Map (com.itom-map-app)	A dependent internal plugin necessary to support the core components of the Service Mapping, Event Management, and Service Operations Workspace ITOM applications.
Discovery (com.snc.discovery)	Finds applications and devices on your network, and then updates the CMDB with the information it finds.
Event Management and Service Mapping Core (com.snc.service-watch)	Comprises the common components of Event Management and Service Mapping.
Cloud Provisioning and Governance Core (com.snc.cloud.core)	Contains resource blocks and cloud Application Programming Interface (cloud API) mappings related to Cloud API-based discovery and events processed by either response mapping or calling patterns.
Pattern Designer (NG version) (com.snc.ng.pattern.designer)	The next generation (NG) version of the Pattern Designer that enables you to create, edit, and debug patterns used by horizontal and top-down discovery.
Service Mapping Plus (sn_sm_scoped_app)	Offers extended features for Service Mapping, and if related ServiceNow® Store applications and plugins aren't already installed, it installs them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag Governance

Plugins or applications installed with Tag Governance (sn_itom_tag)

Plugin or application name	Description
Cloud Access Interface (sn_cai)	Provides easy access to the cloud through scripting. You can use this ServiceNow Store application to perform actions on virtual machines provisioned on Amazon Web Services (AWS) and Microsoft Azure.
Certification Core (com.snc.certification_core)	Contains core certification structures such as filters and templates, and provides compliance functionality.
ITOM Visibility License (com.snc.itom.vis.license)	Enables the activation of any ServiceNow Store application that requires ITOM Visibility licensing.
ITOM/OT SU Licensing (sn_itom_licensing)	Calculates and displays how many subscriptions for ITOM applications your organization purchases and consumes.
Discovery (com.snc.discovery)	Finds applications and devices on your network, and then updates the CMDB with the information it finds.
Flow Designer Action Step - Script on MID (com.glide.hub.action_step.script_mid)	Provides an action step to execute custom JavaScript within a reusable action on a MID Server.

Cloud Discovery Workspace

Plugins or applications installed with Cloud Discovery Workspace (sn_cloud_ops_ws)

Plugin or application name	Description
Discovery and Service Mapping Patterns (sn_itom_pattern)	Provides the latest versions of the Discovery and Service Mapping Patterns released on the ServiceNow Store. ServiceNow releases new patterns on the ServiceNow® Store on a monthly basis, enabling your organization to discover the latest industry-standard devices and applications.
Cloud Provisioning and Governance Core (com.snc.cloud.core)	Contains resource blocks and cloud Application Programming Interface (cloud API) mappings related to Cloud API-based discovery and events processed by either response mapping or calling patterns.
Performance Analytics Content Pack for Cloud Resources (sn_disco_cd)	Displays the aggregated view of all your Cloud Resources. It also provides analytics tables and views that associate tags with resources, enabling you to build your own custom dashboards and reports.
Pattern Designer (NG version) (com.snc.ng.pattern.designer)	The next generation (NG) version of the Pattern Designer that enables you to create, edit, and debug patterns used by horizontal and top-down discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector for Visibility Content

Plugins or applications installed with Agent Client Collector for Visibility Content (sn_acc_vis_content)

Plugin or application name	Description
Agent Client Collector for Visibility Content Content (sn_acc_vis_content)	Provides agent-based infrastructure discovery and collects Software Asset Management (SAM) metrics to support Discovery and Digital End-User Experience (DEX) applications.
Discovery (com.snc.discovery)	Finds applications and devices on your network, and then updates the CMDB with the information it finds.
Agent Client Collector Framework (sn_agent)	Manages Agent Client Collector on the ServiceNow instance and the MID Server.
CMDB CI Class Models (sn_cmdb_ci_class)	Adds class models that extend the CMDB class hierarchy, including class descriptions, identification rules, identifier entries, and dependent relationships if applicable. Applications such as Discovery and Service Mapping can use these class extensions to populate configuration items (CIs) and discover various technologies and software.

Agent Client Collector for Visibility Content Content

Plugins or applications installed with Agent Client Collector for Visibility Content Content (sn_acc_vis_content)

Plugin or application name	Description
Expanded Model and Asset Classes (sn_ent)	Provides new, expanded model classes, asset classes, and model categories that come with the base system, defined by ServiceNow. Install the Asset Management (com.snc.asset_management) plugin as a prerequisite for using this application.
Discovery – IP Based (com.snc.discovery.ip_based)	Contains shared components of horizontal discovery used by Discovery and Service Mapping.
Agent Client Collector Framework (sn_agent)	Manages Agent Client Collector on the ServiceNow instance and the MID Server.
CMDB CI Class Models (sn_cmdb_ci_class)	Adds class models that extend the CMDB class hierarchy, including class descriptions, identification rules, identifier entries, and dependent relationships if applicable. Applications such as Discovery and Service Mapping can use these class extensions to populate configuration items (CIs) and discover various technologies and software.

Related topics

- Plugins or applications installed with ITOM Cloud Accelerate
- Plugins or applications installed with ITOM AIOps

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins activated with ITOM Optimization

Applications installed with ServiceNow Otto for ITOM

Agent Client Collector (ACC) admin workspace agent actions

When working in the Agent Client Collector (ACC) admin workspace, you can select actions to perform on the agent by selecting the **Agent actions** button.

Agent actions

Agent action	Description
Restart	Restart the agent.
Test Check	On the resulting pop-up window, select a check definition that you want to check on the agents.
Validate ACC plugins	Validate plugins on the selected agent.
Pause data collection	Pause data collection run on the agent when CPU consumption is getting too high.
Upgrade	Upgrade the agent to a more recent version.
Clear ACC plugins	Clear all plugins on the agent to remove unwanted plugins.
Collect host data	Manually refresh host data.
Grab agent log	Retrieves the log for the agent.
Reset failed upgrade attempts	Agents that failed to upgrade are reset to their original state.

Agent Client Collector (ACC) admin workspace information

When selecting an agent in the ACC admin workspace, select from the displayed options to view information relating to the agent.

Agent information modules

Option	Description
Agent details	Details of the selected agent.
Diagnostic tests	Diagnostic tests that have run on the agent. Select **Run full diagnostics** to run a diagnostics test on demand.
Configured checks	The checks configured to run on the agent.
Disabled checks	Checks associated with the agent that are disabled.
Error messages	Error messages associated with the agent. Select the name of the error to open the error's details page. The page contains the following tabs:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent information modules (continued)

Option	Description
	• Details: Details about the error. • Error Message Suggestions: Suggestions on how to fix the selected error. • Agents for error: The agents on which the selected error appears.

ITOM AIOps

ITOM AIOps enables IT operations teams, site reliability engineers, and DevOps professionals to proactively monitor, analyze, and optimize the health and performance of their IT infrastructure and services.

Get started

Explore Learn about ITOM AIOps	Configure Configure the Service Operations Workspace to work with ITOM AIOps
Reference Find reference information about ITOM AIOps	

Install ITOM AIOps apps using Setup Hub

Install two plugins to enable ITOM AIOps functionality: ITOM AIOps and ServiceNow Otto for ITOM.

Before you begin
Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
   The Admin Home page opens.

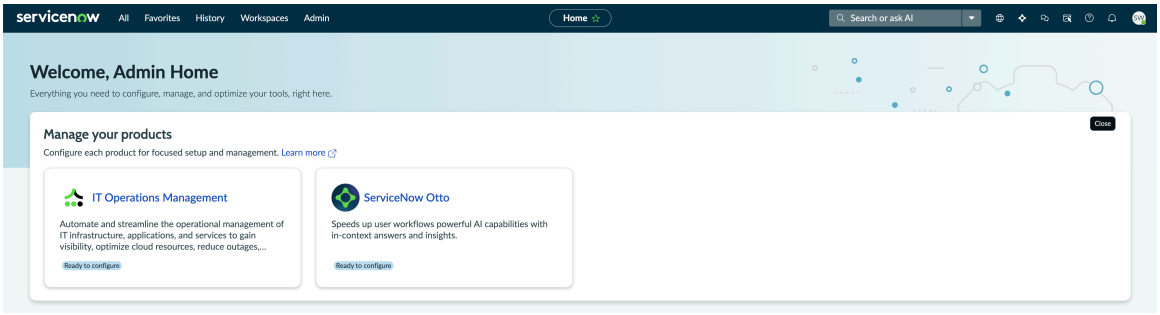

2. In the IT Operations Management tile, select **Ready to configure**.
   The Welcome! Let's get started dialog box appears.

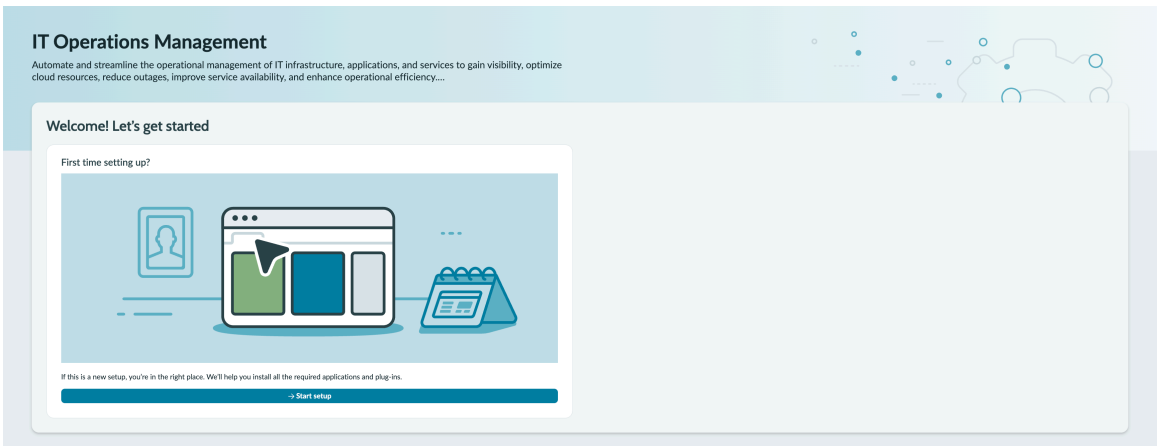

3. Select **Start setup**.

   View all available plugins in the **Not Installed** tab.

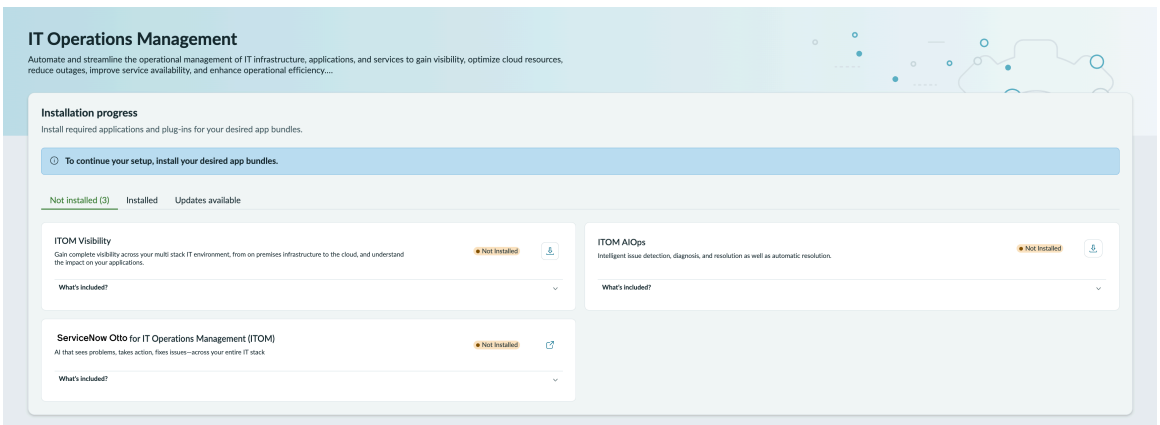

4. Install the ITOM AIOps plugin by selecting the install icon ( ).
5. To select the ServiceNow Otto for ITOM plugin, perform the following steps:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select the Application manager icon ( ).

b. On the Application manager page, select **Install**.

Both plugins are installed and listed in the **Installed** tab in the IT Operations Management page.

Exploring ITOM AIOps

Overview of ITOM AIOps applications and capabilities that enable proactive IT operations management through AI-powered monitoring, analytics, and automation.

ITOM AIOps overview

ITOM AIOps enables IT operations teams, site reliability engineers, and DevOps professionals to proactively monitor, analyze, and optimize the health and performance of their IT infrastructure and services.

ServiceNow® ITOM AIOps provides AIOps capabilities that transform how organizations manage IT operations. By combining advanced analytics, machine learning, and automation, ITOM AIOps helps teams avoid service outages, reduce mean time to resolution, and optimize infrastructure performance.

Service Operations Workspace

Network Operations Center (NOC) operators and IT operations teams use Service Operations Workspace as their primary interface for managing IT operations powered by ITOM AIOps. The workspace consolidates alerts, incidents, and operational data from multiple AIOps applications into unified dashboards and workflows.

Operators use the workspace to monitor service health across the entire IT infrastructure, investigate alerts with full context from multiple data sources, and coordinate response activities. The workspace presents correlated events, suggested remediation actions, and service impact information in a single interface. This enables operators to quickly understand and respond to issues.

The workspace integrates data from all AIOps products to provide operators with comprehensive situational awareness and streamlined incident management capabilities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM AIOps workflow

Each AIOPs application focuses on specific aspects of IT operations while contributing to a unified AIOps

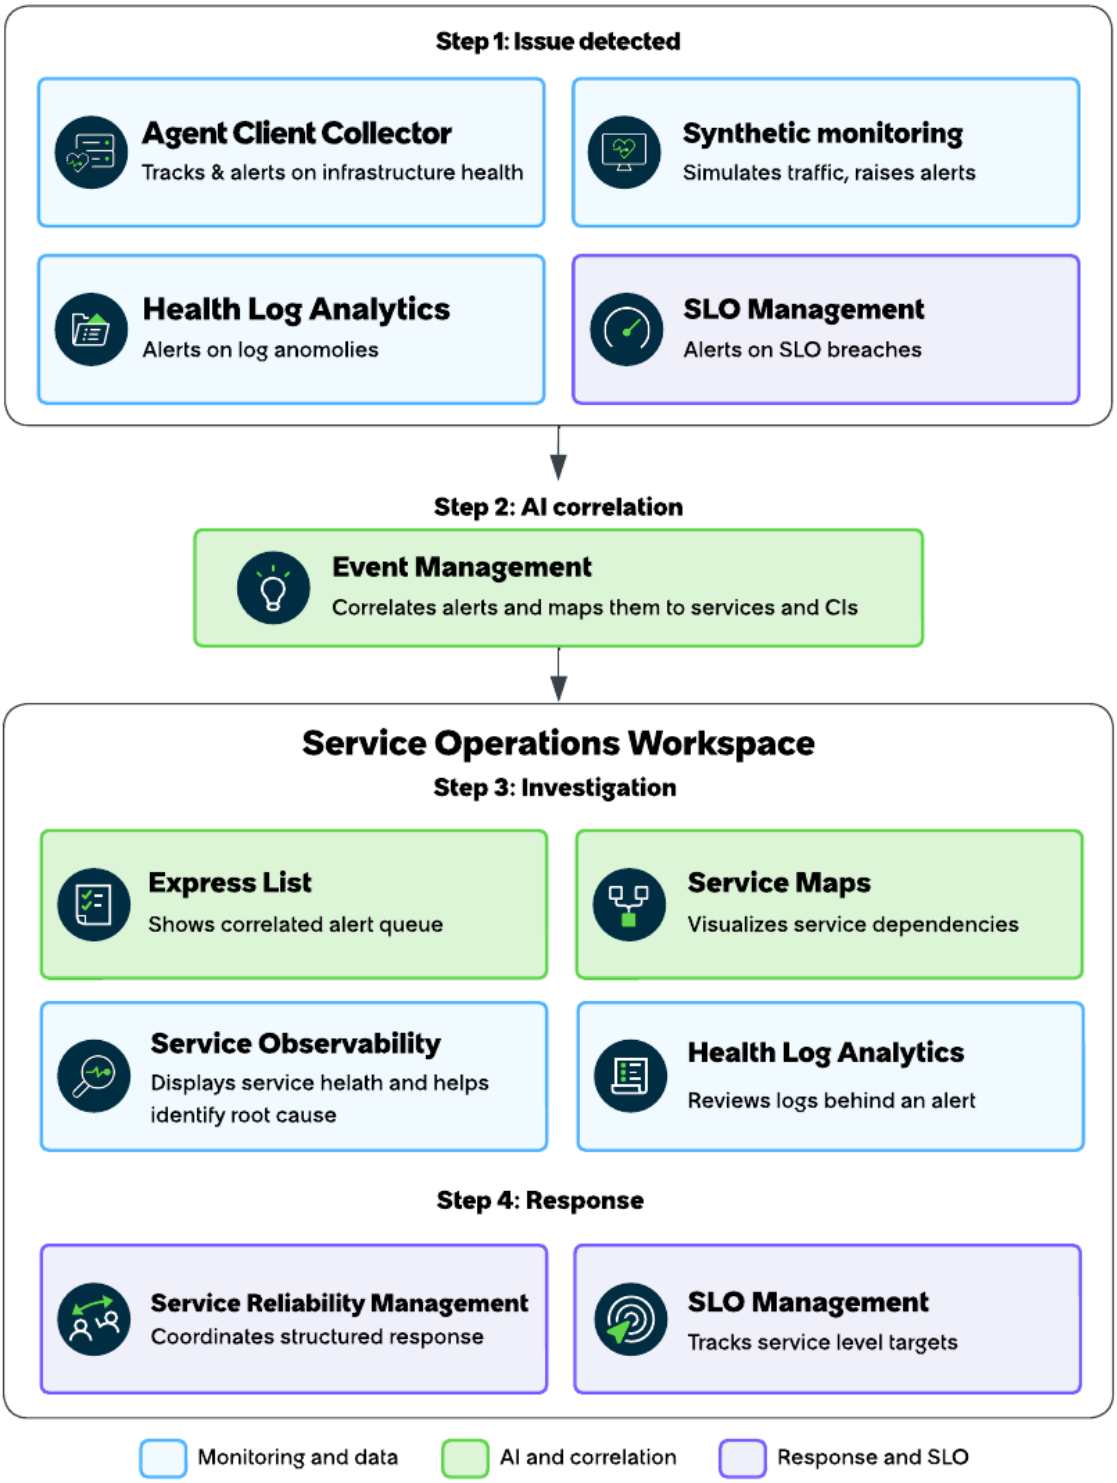

platform.

Agent Client Collector

Monitors service availability and infrastructure performance in real-time. Combined with Metric Intelligence, it establishes dynamic thresholds and detects anomalies that may indicate potential service outages before they occur.

Synthetic monitoring

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Monitors critical services by simulating user transactions on API endpoints, identifying performance bottlenecks and helping maintain optimal user experiences. It provides early warning of service degradation before real users are affected.

Health Log Analytics

Collects log data in real time and uses machine learning to identify patterns and detect anomalies. It typically ingests logs through the MID Server, identifies normal operating patterns, and using Event Management, raises actionable alerts when it detects significant deviations that might indicate emerging issues.

Event Management

Serves as the central nervous system for your IT operations. It receives alerts from monitoring tools and correlates related events to reduce noise. The system maps alerts to configuration items in the Configuration Management Database (CMDB) and calculates business impact using service dependencies from Service Mapping. This correlation transforms hundreds of individual alerts into prioritized, context-rich incidents that focus teams on the most critical issues.

Express List

Provides the Service Operations Workspace view into Event Management, enabling operators to efficiently monitor systems and services, resolve alerts, evaluate the alert impact, track issues, and report incidents.

Service Observability

Combines external observability telemetry with Configuration Management Database (CMDB) data, to reveal dependencies that may not be obvious from individual alerts. Charts provide critical insights by showing how infrastructure components may be affecting a single service.

Service Reliability Management

Enables teams to respond to incidents with structured workflows, automated escalation, and collaborative tools that streamline incident response.

SLO Management

Helps organizations define, monitor, and report on service level objectives and service level indicators. It provides automated SLO tracking and alerting when services approach or breach defined thresholds, enabling proactive service quality management.

ITOM AIOps users

Role	Description
Admin [evt_mgmt_admin]	Configures and sets up Event Management properties and rules.
Operator [evt_mgmt_operator]	Manages alerts, including closing and acknowledging them.
User [evt_mgmt_user]	Manages the lifecycle of alerts, including performing basic operations such as viewing and acknowledging them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ITOM AIOps benefits

Benefit	Feature	Users
Predictive issue detection	• Health Log Analytics • Synthetic monitoring	Admin, operator
Automated alert response	Alert automation in Service Operations Workspace for ITOM	Admin
Performance optimization	Exploring Agent Client Collector	Admin, operator
Service health and investigation	Exploring Service Observability	Admin, operator
Service level management	• Exploring Service Reliability Management • Exploring SLO Management	Admin, operator

Products that add value to ITOM AIOps

In the context of ITOM AIOps, which focuses on maintaining the health and performance of IT systems and infrastructure, there are several ServiceNow products that can add significant value:

Discovery

Discovery automatically identify and collect information about IT assets, configurations, and relationships, providing organizations with a comprehensive inventory to effectively manage and monitor their IT infrastructure.

Service Mapping

Service Mapping provides information about application instance services within the [cmdb_ci_service_discovered] table. This data helps establish connections between infrastructure and application configuration items (CIs) stored in the [cmdb_ci_appl] table, enhancing visibility into IT environments and facilitating efficient management and monitoring processes.

Service Portfolio Management

Service Portfolio Management (SPM) offers the associated product model, while Software Asset Management (SAM) and Hardware Asset Management (HAM) provide lifecycle data for Technology Lifecycle Management (TLM). Together, they enable comprehensive management of IT assets, ensuring effective utilization, compliance, and optimization throughout their life cycles.

Products that benefit from ITOM AIOps

Incident Management

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Incident Management tools use downstream information from Event Management to create incidents swiftly, ensuring timely resolution of issues.

Customer Service Management (CSM)

Customer Service Management (CSM) systems benefit from ITOM AIOps by utilizing application service impact data to identify affected users promptly, enhancing customer support efficiency and satisfaction.

What to know before you begin

Before implementing ITOM AIOps, verify that your instance has the necessary prerequisites and that you understand the configuration requirements for each application.

A well-populated Configuration Management Database (CMDB) is crucial to get the most out of AIOps. ITOM AIOps relies on accurate configuration item data to map events to infrastructure components, calculate service impact, and provide context for alert correlation. Use Discovery to populate your CMDB with current infrastructure data before activating AIOps applications.

Configure the MID Web Server extension to enable ITOM AIOps features. The MID Web Server is an extension that enables external clients to push metric data and events to the MID Server. This extension is required for Event Management, and many instances of Agent Client Collector, and Health Log Analytics.

ServiceNow Otto for Setup provides a sequence of tasks that help you configure Event Management on your ServiceNow instance. For more information about using the ServiceNow Otto for Setup, see Configure Event Management using ServiceNow Otto for Setup.

Service Operations Workspace for ITOM

ServiceNow® Service Operations Workspace (SOW) is a configurable workspace designed to streamline IT Operations Management (ITOM) workflows, offering a unified experience for multiple IT operational processes. Configurable and user-friendly, it consolidates various tasks into a single interface, allowing seamless navigation and easy access to tools and information.

By centralizing workflows, it enhances collaboration, facilitates efficient task management, and provides real-time insights through different metrics. Highly customizable, it enables you to tailor the workspace to your specific needs, optimizing productivity and workflow efficiency. With its comprehensive set of features, including generative AI assistance, automation and collaboration tools, Service Operations Workspace empowers users to work more effectively across different workflows, boosting productivity organization-wide. Scalable to accommodate evolving business needs, it ensures continued support for IT environments of all sizes, delivering superior services and driving operational excellence.

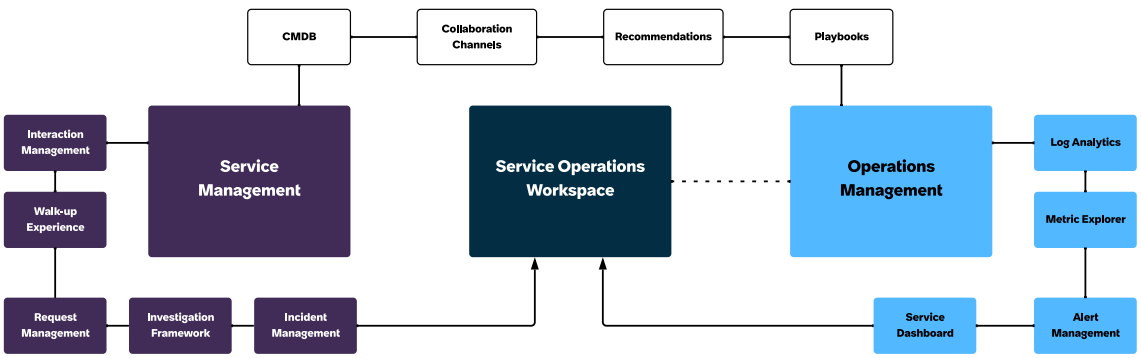

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information regarding subscriptions for ITOM Service Operations Workspace, check the ITOM AIOps subscription at Subscription ITOM AIOps.

When setting up the Service Operations Workspace, it's crucial to configure both IT Operations Management (ITOM) for managing alerts and IT Service Management (ITSM) for resolving incidents and delivering services. This integrated setup ensures access to the full range of features, facilitating comprehensive management of IT operations and service quality.

To know about the ITSM features, see Service Operations Workspace for ITSM.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Troubleshoot and get help

Contact Customer Service and Support

Service Operations Workspace for ITOM user interface

Explore the interface to understand how an operator can prioritize tasks and offer solutions.

Access Service Operations Workspace

Open the workspace by navigating to **Workspaces** > **Service Operations Workspace**.

Home page

The **Overview** tab is your team's command center. It gives you a real-time snapshot of alert activity, workload distribution, and the impact AI is having on your team's performance — all in one place. Spanning four areas—AI efficiency, team workload, severity, and assignments—this reveals AI's impact on resolution speed and workflow bottlenecks.

Use it to spot imbalances, track resolution speed, and see at a glance whether your team is keeping up with incoming alerts.

Use the assignment group filter (for example, em_admin — 19 items selected) to filter your view. If you're an evt_admin, you can filter by groups you manage and groups you're a member of. If you're an aiops_manager, you can filter only by groups you manage. Use the **Teammates** list to filter by teammate.

You can also reach the homepage by navigating to **All** > **Event Management** > **Operational Home Pages** > **AIOps Managers**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

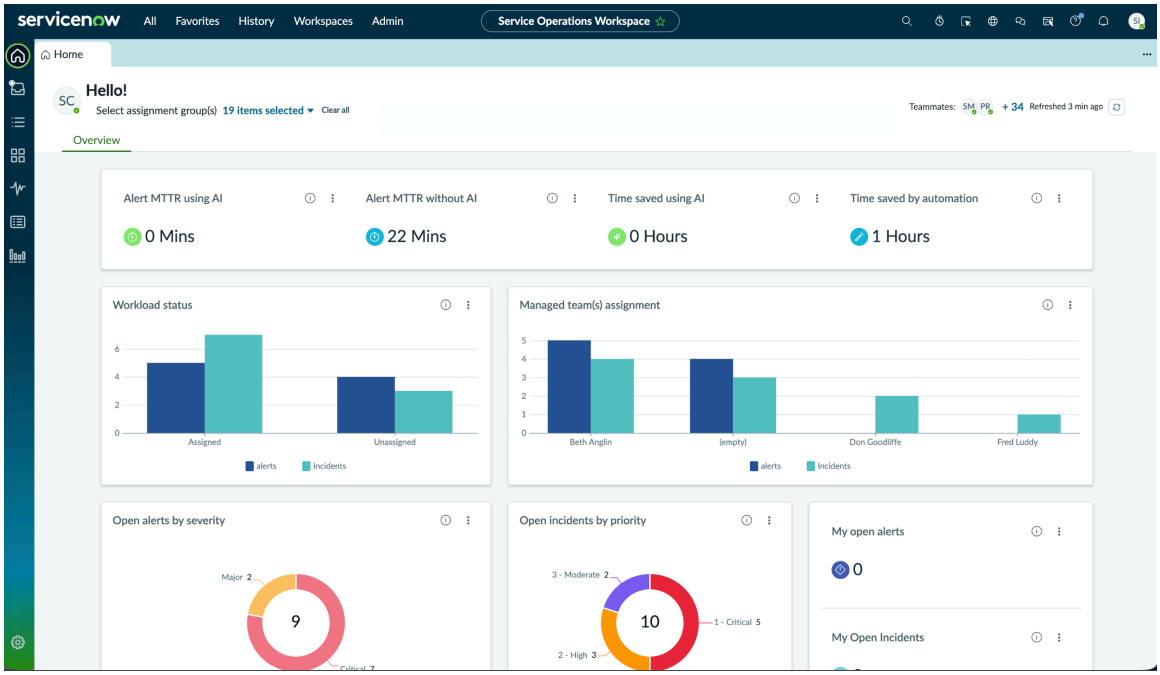

Note: If you don't install the ServiceNow Otto for IT Operations Management (ITOM) plugin, AI-related KPIs aren't visible. If you have the plugin enabled, you can see the **AI supervising** tab.

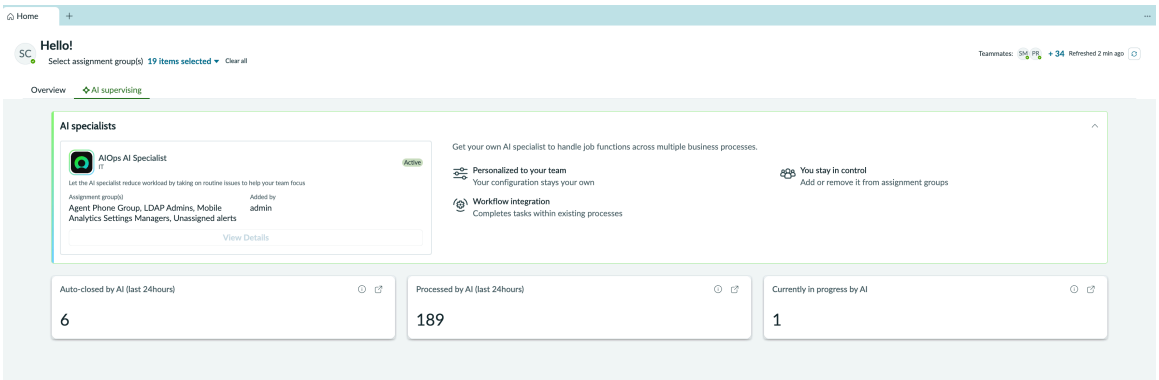

For each tile, select the icon to view the list of alerts in the Express List.

List page

To access information on alerts from List, select the List icon ( ).

The **Lists** tab displays predefined lists, while the **My Lists** tab allows users to create and customize their own lists for personalized task management. For detailed instructions on creating a list, see Create a list in Service Operations Workspace for ITOM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Lists page

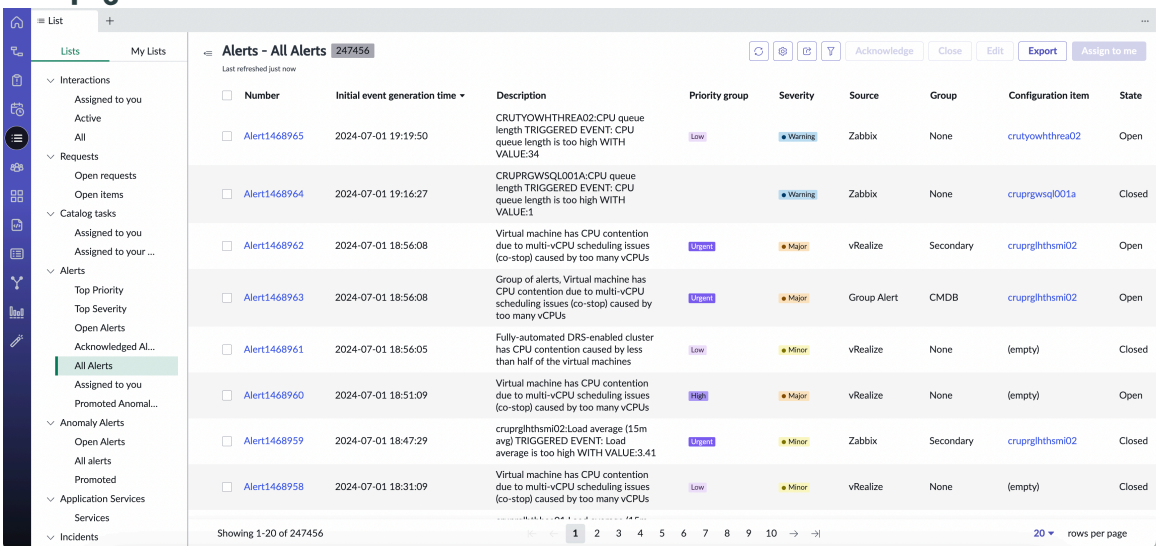

Record view

Access detailed information for any alert by selecting its alert number to open the record.

Alert details page

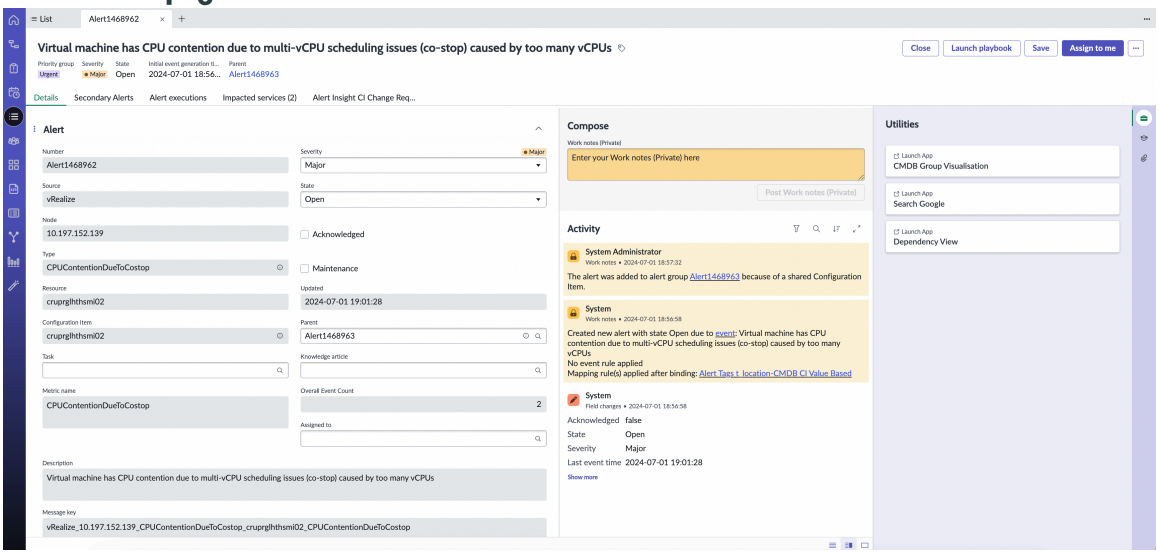

Express List

Express List offers quick alert information for efficient monitoring, resolution, impact evaluation, tracking, and incident reporting.

To access information on Express List, select the Express List icon ( ).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Express List page

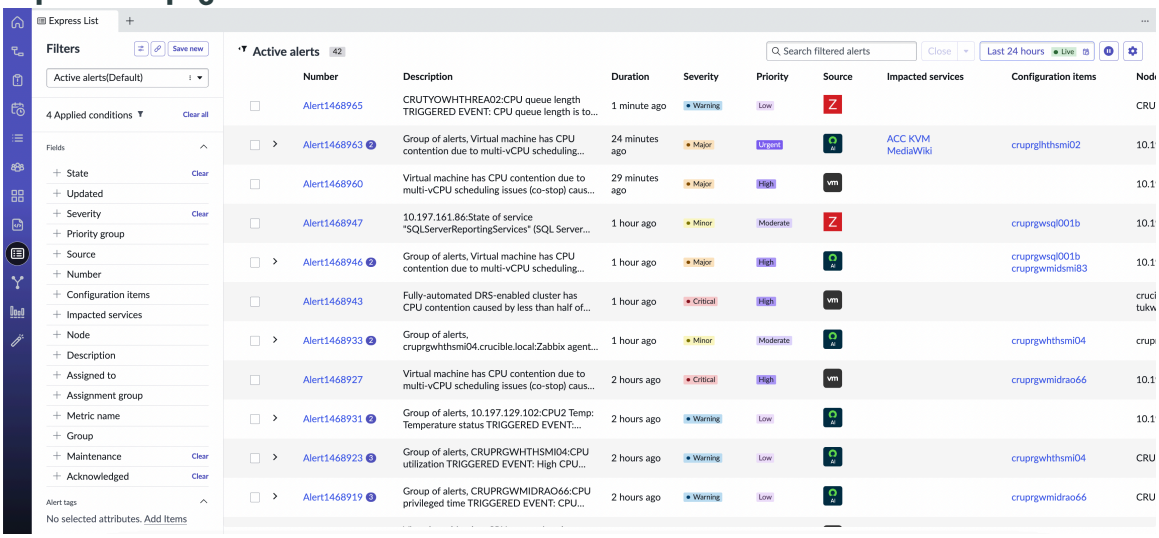

Integrations Launchpad

The Integrations Launchpad consolidates events, metrics, and logs from your infrastructure, providing insights into a unified system.

To access information on Integrations Launchpad, select the Integrations Launchpad icon ( ).

Integrations Launchpad page

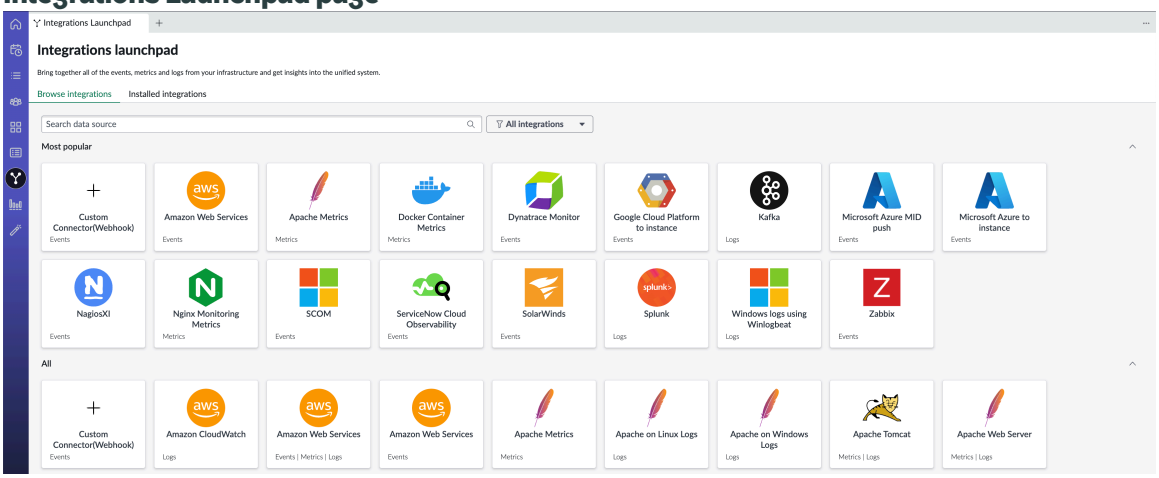

Alert automation

Automations streamline processes, enabling timely responses to potential issues by ignoring, enriching, grouping, and responding to alerts. They reduce manual intervention, lower the risk of errors, and free up IT teams to focus on strategic tasks, boosting system reliability and customer satisfaction.

To access information on automations, select the alert automations icon ( ).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Automations page

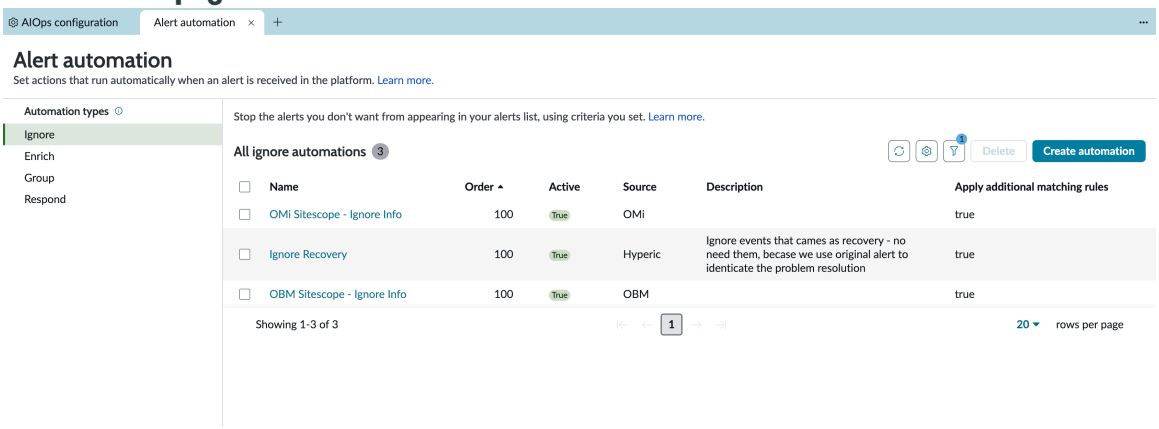

AIOps Dashboards

The AIOps dashboards offer comprehensive visualizations of critical operational data. You can monitor KPIs, track service health, and gain valuable context on the IT environment, enabling proactive responses to issues.

To access information on dashboards, select the AIOps Dashboards icon ( )

AIOps Dashboards page

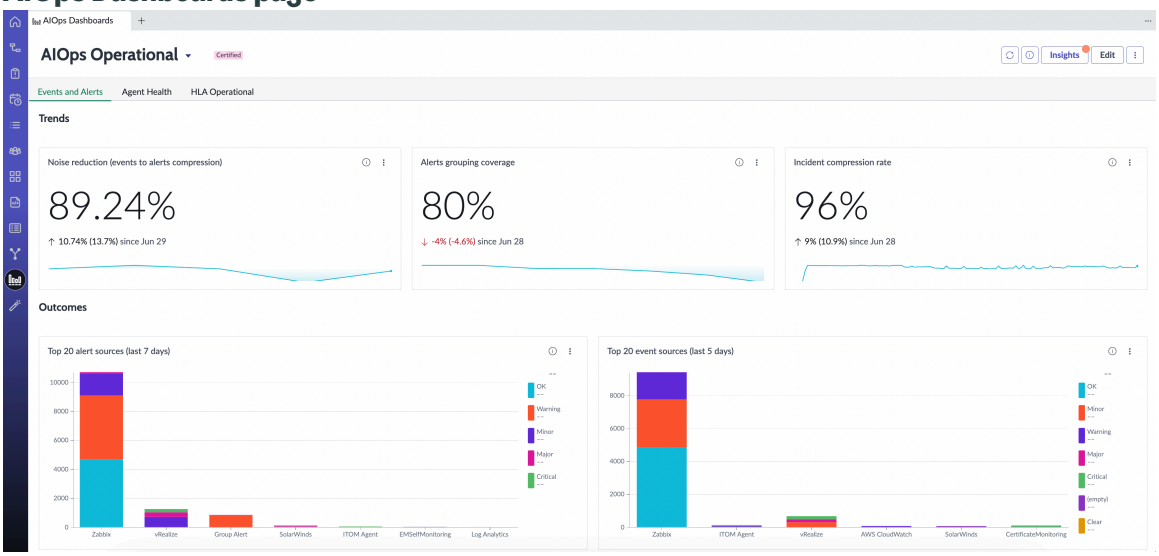

Admin configuration use case for AIOps

Learn how IT administrators configure and optimize ITOM AIOps for enterprise environments. This use case demonstrates the setup journey for a multi-cloud financial services organization.

This use case demonstrates how IT administrators configure ITOM AIOps to transform operations from reactive to proactive. The scenario follows an IT administrator implementing AIOps capabilities for a global financial services company with hybrid cloud infrastructure.

Configuration scenario

The IT administrator is responsible for implementing ITOM AIOps across a complex environment including on-premises data centers and cloud services. The organization runs critical trading applications, customer-facing banking services, and internal business systems that require 99.9% availability.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The configuration goals include reducing alert noise, enabling predictive issue detection, and providing unified visibility across all environments. The goals also include empowering NOC operators with intelligent workflows and AI-powered insights.

Phase 1: Event sources and intelligent correlation

The administrator configures Event Management to aggregate and intelligently correlate alerts from multiple monitoring tools and data sources.

Event source integration

The administrator integrates existing monitoring tools including network monitoring systems, application performance monitoring platforms, and infrastructure monitoring solutions. The administrator also integrates cloud-native monitoring services. Each integration is configured with appropriate event parsing rules and severity mappings.

Correlation rules configuration

Intelligent correlation rules are established based on CMDB relationships, time-based patterns, and business service dependencies. The administrator configures advanced correlation algorithms that consolidate related infrastructure alerts into a single root cause alert.

Alert prioritization

Trading system alerts are automatically designated as Priority 1 with immediate escalation during market hours. Customer-facing banking services receive Priority 2 classification with 15-minute response requirements. Internal tools receive lower priority during business hours with standard response times. Business service criticality scoring uses revenue impact, customer count, and regulatory requirements as weighting factors.

Phase 2: Predictive analytics and machine learning

The administrator configures advanced analytics capabilities that enable proactive issue detection and intelligent insights for operational teams.

Health Log Analytics setup

HLA is configured to ingest log data from critical applications and infrastructure components. The administrator sets up log parsing rules, anomaly detection thresholds, and predictive models. These models learn normal operational patterns and identify deviations that indicate potential issues.

Metric Intelligence configuration

Machine learning algorithms are trained on historical performance data to establish baseline behaviors for key metrics. The administrator configures anomaly detection sensitivity.

Synthetic monitoring

Synthetic transactions are configured to continuously test critical user journeys through API endpoints. The administrator sets up monitoring with 5-minute frequency intervals.

Service Reliability Management configuration

Service Reliability Management is configured to establish incident response workflows, collaboration channels, and automated remediation capabilities. The administrator sets up on-call schedules, escalation policies, and integration with communication platforms. This enables rapid response when alerts require immediate attention.

Service Level Objective Management setup

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SLO targets are defined for critical business services based on customer expectations and business requirements. The administrator configures SLO monitoring for availability, performance, and error rate metrics. The administrator establishes error budgets and burn rate alerts that provide early warning when services risk missing reliability targets.

Service Observability integration

Service Observability is configured to combine Observability telemetry from external monitoring tools with CMDB relationship data. This data integration enables rapid identification of issues in specific microservices and their related entities without having to use an external tool.

Phase 3: Unified workspace and user experience

The administrator configures the Service Operations Workspace to provide role-based dashboards and workflows that enable efficient operations management.

Dashboard customization

Role-specific dashboards are created for different user types including NOC operators, incident managers, and executive stakeholders. Each dashboard displays relevant metrics, alerts, and service health information appropriate for the user's responsibilities and decision-making needs.

Workflow integration

Automated workflows are configured for common operational scenarios including alert escalation, and notification routing. The administrator establishes integration with IT Service Management (ITSM) processes. This creates seamless handoffs between monitoring and service management.

AI-Powered assistance

ServiceNow Otto for IT Operations Management (ITOM) is configured to provide intelligent alert analysis, recommended actions, and contextual guidance. AI models use organization-specific knowledge including historical incident patterns, common error code translations, and proven remediation procedures to provide insights. Approval workflows are established for high-risk AI-recommended actions while allowing automatic execution of low-risk remediation steps.

Phase 4: Performance tuning and optimization

After initial deployment, the administrator continuously optimizes the system based on operational feedback and performance data to maximize effectiveness and user adoption.

Alert tuning

Based on operator feedback and false positive analysis, the administrator refines correlation rules and adjusts thresholds. The administrator optimizes alert prioritization to reduce noise while maintaining comprehensive coverage of critical issues.

Machine learning model refinement

Anomaly detection models are tuned based on operational patterns and seasonal variations. The administrator adjusts sensitivity settings, incorporates feedback on false positives, and expands training data sets. This improves prediction accuracy.

Configuration outcomes

The comprehensive ITOM AIOps configuration delivers measurable improvements in operational efficiency and service reliability:

Alert reduction

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert volume reduced through intelligent correlation and noise suppression

Proactive detection

Critical issues identified and resolved before customer impact during typical operational shifts

Faster resolution

Mean time to resolution reduced through AI-powered insights and automated workflows

Unified visibility

Single pane of glass for hybrid cloud and on-premises infrastructure

Predictive capabilities

Capacity issues and performance degradation predicted in advance

Business alignment

Alert prioritization and impact analysis directly tied to business service criticality

This configuration enables the operational workflows demonstrated in the NOC operator use case for AIOps. The configuration transforms reactive operations into intelligent, proactive service management.

Related topics

NOC operator use case for AIOps

ITOM AIOps

Service Operations Workspace for ITOM

NOC operator use case for AIOps

Follow a typical NOC operator through their daily workflow using ITOM AIOps capabilities to transform reactive operations into proactive, intelligent IT management.

This use case demonstrates how ITOM AIOps transforms a Network Operations Center (NOC) operator's daily workflow from reactive firefighting to proactive, intelligent operations management. The scenario follows a senior NOC operator at a global financial services company through a typical shift managing critical trading applications and customer-facing banking services.

Starting the shift with unified visibility

At 9:00 a.m., the operator begins the shift by logging into the Service Operations Workspace for ITOM. This unified dashboard consolidates all AIOps capabilities, providing immediate visibility into:

- Alert prioritization showing active alerts with intelligent ranking
- Service health overview with color-coded status of critical business services
- Trending anomalies detected by Metric Intelligence
- Predicted issues identified by Health Log Analytics

Instead of checking multiple monitoring tools, the operator has one consolidated view with intelligent prioritization. This enables faster situational awareness and decision-making.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Proactive issue detection with predictive analytics

At 9:15 a.m., the operator notices a warning on the Health Log Analytics dashboard indicating unusual patterns in trading application logs. The system detects:

- Increasing error rates in authentication services
- Memory utilization patterns that historically precede outages
- Database connection pool exhaustion warnings

The operator drills down into the anomaly. Health Log Analytics provides root cause analysis pointing to a recent configuration change, predicts high probability of service degradation within 45-60 minutes, and recommends specific remediation actions. This enables the operator to address the issue before it impacts users, preventing a potential trading system outage during market hours.

Intelligent alert correlation and impact analysis

At 10 a.m., multiple alerts flood in from various monitoring tools. These alerts concern network latency, application response time degradation, and infrastructure CPU spikes. Event Management uses CMDB topology data to:

- Correlate individual alerts to a single root cause: failed network switch
- Map the impact to affected business services using service topology
- Prioritize based on business criticality with trading applications marked as Priority 1

The operator sees one consolidated alert with full impact analysis. The operator uses the Service Operations Workspace to acknowledge the alert, run automated diagnostics, and initiate network team escalation workflows. This reduces multiple alerts to one actionable item with clear business context.

Performance monitoring with synthetic transactions

At 11:30 a.m., synthetic monitoring alerts the operator that the customer mobile banking application is experiencing issues. These include transaction failures and API response times exceeding SLA thresholds.

The operator uses the Service Operations Workspace to view the realted alerts and starts an investigation. The operator creates an incident with pre-populated impact assessment. Issues are detected before customers contact the help desk.

AI-powered alert analysis and resolution

At 1:15 p.m., a complex alert from the core banking system arrives with cryptic error codes. The operator uses ServiceNow Otto for IT Operations Management (ITOM) to receive:

- Human-readable explanation of database connection timeout caused by connection pool exhaustion
- Historical context showing similar patterns occurred three times in past six months
- Recommended step-by-step remediation actions with risk assessment

Following the AI-recommended steps, the operator resolves the issue. Complex technical issues are resolved quickly without escalation using AI-powered insights from historical data patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Proactive capacity management with machine learning

At 3:00 p.m., Agent Client Collector and Metric Intelligence alert the operator to unusual storage growth patterns, CPU utilization trending toward critical thresholds, and memory consumption anomalies.

The operator reviews anomaly details, creates a proactive capacity planning ticket, and schedules automated cleanup jobs for non-critical data. Proactive capacity management helps prevent weekend outages and emergency procurement needs.

Complex incident resolution with unified telemetry

At 3:30 p.m., a critical incident affects the loan processing application. The operator uses Service Observability to combine observability telemetry from external monitoring tools with CMDB relationship data. The operator correlates performance metrics with infrastructure health data.

The operator quickly identifies the issue in a specific microservice and its database dependency, enabling faster resolution by the development team. Complex distributed system issues are resolved with unified telemetry and topology views.

Key business outcomes achieved

This use case demonstrates how ITOM AIOps transforms NOC operations by delivering:

- Proactive issue prevention with three issues resolved before user impact
- Alert noise reduction with 70+ individual alerts consolidated into eight actionable items
- Faster resolution times reduced from 45 minutes average to 12 minutes
- Enhanced context with every alert including business impact and recommended actions
- Unified operational experience replacing multiple disparate tools
- AI-powered insights using historical patterns and machine learning predictions
- Automated workflows freeing operator time for strategic activities

The transformation from reactive firefighting to proactive, intelligent operations management enables NOC operators to maintain higher service levels. This reduces operational overhead and improves job satisfaction.

Applying the CSDM guidelines to ITOM AIOps

ITOM AIOps helps you track and maintain the health of services in your organization. ITOM AIOps CSDM reference serves as a foundational blueprint for maintaining optimal IT operations by standardizing data models and processes. The main goal is to help you understand the two main capabilities that work together within the CSDM framework.

ITOM AIOps consists of two main capabilities that use the CSDM framework.

Event Management

Event Management gathers alerts from infrastructure events captured by third-party monitoring tools. Event Management uses IT-related information gathered by Discovery to map alerts to CIs. Based on the collected information, Event Management then provides dashboards showing a consolidated view of all service-impact events.

Metric Anomaly Detection

Proactively analyzes your IT infrastructure to spot issues and prevent service outages. Using advanced machine learning to analyze information about your IT

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

infrastructure, the application automatically determines dynamic thresholds and identifies anomalies that may indicate potential service outages.

For CSDM set-up information, see Implementing the CSDM framework in stages.

For more information on CSDM framework, see Exploring the CSDM model.

ITOM AIOps considerations for CSDM framework

ITOM AIOps aligns with CSDM framework to ensure standardized data management and optimize IT operations.

Key considerations for implementing ITOM AIOps for CSDM framework include the following.

Identifying necessary Configuration Items (CIs) for optimal Event Management benefits

Event Management utilizes Configuration Items (CIs) monitored by the operations team to correlate alerts with the appropriate CI and assess their impact on the monitored service.

Defining Technical Services within Event Management

Dynamic CI Group services, formerly known as Technical Services, leverage the CMDB CI Group capability to query the CMDB. This enables efficient grouping of Configuration Items (CIs) for various operational purposes.

Recognizing common use cases for the Dynamic CI Group service

In Event Management (EM), the Dynamic CI Group service functions as a logical grouping mechanism for Configuration Items (CIs). It offers technology or service owners insight into the health status of the grouped entities.

CMDB tables used by ITOM AIOps

ITOM AIOps relies on accurate data in CMDB tables to function as expected. Following CSDM guidelines when populating these tables improves the accuracy of service mapping and helps route alerts and incidents correctly.

ITOM AIOps uses the following CMDB tables:

- Mapped Application Service table `[cmdb_ci_service_discovered]`

  Note: With CSDM v5, the label for the table has changed from Application Service to Service Instance.

- Configuration Item table `[ci_*]`
- Dynamic CI Group table `[cmdb_ci_query_based_service]`

The diagram shows a conceptual map of CMDB tables and their relationships, structured according to CSDM guidelines. ITOM AIOps uses only the highlighted elements: the Service Instance table `[cmdb_ci_service_discovered]` and the Dynamic CI Group table `[cmdb_ci_query_based_service]`. The Configuration Item table `[ci_*]` is not highlighted separately, as it represents the broader CI class hierarchy that underlies most elements in the model. The remaining elements in the model support other ServiceNow products and processes, and are shown here for reference only.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CMDB tables used by ITOM AIOps within the CSDM v5 model

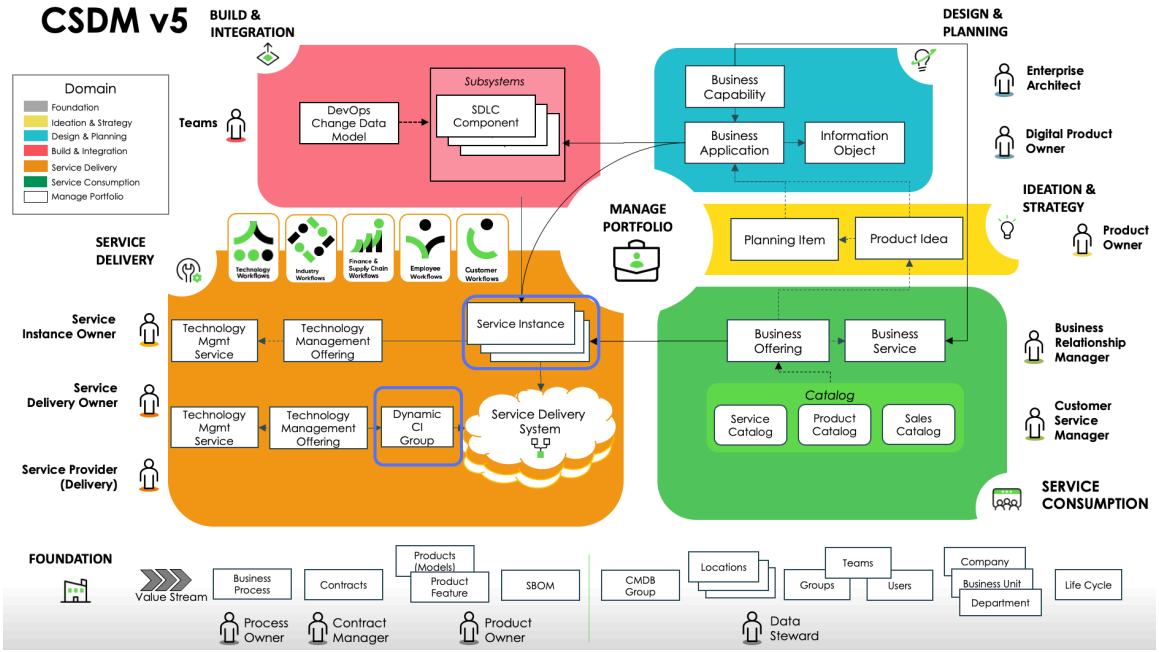

For more information on the CSDM framework, see Exploring the CSDM model.

Plugins or applications installed with ITOM AIOps

Tables that list the plugins or applications that are installed with ITOM AIOps applications. When you update your application, any newly required application dependencies are installed.

Event Management

Plugins or applications installed with Event Management (com.glideapp.itom.snac)

Plugin or application name	Description
Alert Management Content (com.em-alert-mgmt-content)	Contains all the necessary base system subflows that can be used for alert remediation.
ITOM Licensing (com.snc.itom.license)	Central components for the ITOM Licensing Model, which calculates and displays how many subscriptions for ITOM applications your organization purchases and consumes.
Mid Web Server (com.snc.sa.mid.webserver)	A MID Server extension that enables developing REST APIs to send events and metrics to the MID Server. The plugin also enables the communication between the MID Server and Agent Client Collector.
Service Analytics (com.snc.sa.analytics)	Provides alert analytics based on machine learning (ML) and topology learning.
Event Management and Service Mapping Core (com.snc.service-watch)	Comprises the common components of Event Management and Service Mapping.
Tag Based Alert Clustering Engine (sn_em_tbac)	Provides a zero-code method that helps simplify alert grouping and performs alert correlation without the need for CMDB, model training, or fine-tuning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Event Management (com.glideapp.itom.snac) (continued)

Plugin or application name	Description
Alert Rules Management (sn_em_arm)	Enables Event Management users to respond automatically to alerts. An alert management rule determines the required alert response.

Health Log Analytics

Plugins or applications installed with Health Log Analytics (sn_occ)

Plugin or application name	Description
Integration Commons for CMDB (sn_cmdb_int_util)	Contains a collection of utility functions that provide a common methodology for the processing of key CMDB attributes and is available to all CMDB certified integrations.
ITOM Cloud Services Core (sn_itom_cloud_svc)	Generates relevant topics for tenants, enabling them to transmit data from their customer infrastructure to ServiceNow automatically. These topics are predominantly intended for use with the Hermes Messaging Service.
Service Operations Workspace ITOM Apps (sn_sow_itom_cont)	A configurable workspace that provides a unified experience for multiple IT Operations Management workflows. You can manage the life cycle of alerts in IT Operations Management workflows through simple, intuitive, and clear user interface, automated recommendations, and targeted experience.
Health Log Analytics Core (com.glideapp.itom.occultus)	Essential components for the Health Log Analytics scoped application.
Microsoft Teams Spoke for ServiceNow IntegrationHub (com.sn.ms_teams.ah)	Provides actions that a process analyst can use when designing flows to automate sending messages about ServiceNow incident and problem records to a Microsoft Teams channel.
Slack Webhooks Spoke for ServiceNow IntegrationHub (com.sn.slack.ah)	Use Slack webhooks to post messages and record details for ServiceNow incidents, problems, and changes to Slack channels.
Normalization Data Services Client (com.glide.data_services_canonicalization.client)	Helps maintain consistency by checking that records for a given company match the list of standard company names. If the plugin finds a company-name field with a variant of that name, it substitutes the standard name in its place.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Reliability Management

Plugins or applications installed with Service Reliability Management (sn_sow_srm)

Plugin or application name	Description
Service Operations Workspace Service Reliability Management (SRM) Common (sn_sow_srm_common)	Delivers common utilities and configuration settings used by Service Reliability Management such as core functionality related to managing teams and services, and metric calculation for services.
Service Reliability Management (Global dependencies) (com.snc.srm)	Supports global scripts and utilities for Service Reliability Management for Service Operations Workspace.
Event Management (com.glideapp.itom.snac)	Helps you to identify health issues across business services and infrastructure on a single management console. Event Management provides intelligent event and alert processing, including automatic alert remediation, aggregation, and root cause analysis (RCA)
Service Operations Workspace ITOM Apps (sn_sow_itom_cont)	A configurable workspace that provides a unified experience for multiple IT Operations Management workflows. You can manage the life cycle of alerts in IT Operations Management workflows through simple, intuitive, and clear user interface, automated recommendations, and targeted experience.
On-Call Scheduling (com.snc.on_call_rotation)	Provides the ability to create on-call schedules and escalation trees. When an incident is created, you can dynamically route the escalation to an on-call resource.
SOW - Alert Automation (sn_sow_alert_rules)	Provides operators in Service Operations Workspace with the ability to create and manage alert automation for their team. This automation can be used to enrich, process, and group alerts.
Service Level Objective Management for Service Operations Workspace (sn_sow_slo)	Includes features to monitor and manage service health using Service Level Indicators (SLIs), Service Level Objectives (SLOs), and error budgets.

Service Operations Workspace for ITOM

Plugins or applications installed with Service Operations Workspace ITOM Apps (sn_sow_itom_cont)

Plugin or application name	Description
Service Operations Workspace Core (sn_sow)	Core capabilities and workspace configurations for Service Operations Workspace.
Event Management Core (sn_em_ai)	Enables you to track and monitor alerts in the ServiceNow Agent Workspace.
Service Operations Workspace Integrations	Provides a new and enhanced experience for viewing, configuring, and managing all your available ITOM AIOps integrations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Service Operations Workspace ITOM Apps (sn_sow_itom_cont) (continued)

Plugin or application name	Description
launchpad (sn_itom_integ_app)	
Service Operations Workspace Alert Mngmt (sn_sow_em)	A detailed alert record page with supplementary data such as Quick Overview, Playbook, and Metrics tabs.
Service Operations Workspace Express List (sn_itom_aiops_list)	A centralized and consolidated platform that enables operators to triage, analyze, and remediate alerts, all in one place. Express List provides a live view of alerts, dynamic filtering, a preview panel with additional context and essential details about each alert, and automated technical analysis using AI.
Service Operations Workspace Log Analytics (sn_sow_hla)	A log viewer in Service Operations Workspace. The application also includes the custom Elasticsearch data source for Performance Analytics (PA) dashboards.
AIOps Dashboards (sn_aiops_dashboard)	Offers comprehensive visualizations of critical operational data and actionable insights that facilitate proactive responses to potential issues. The application targets multiple ITOM personas with different needs. You can monitor key performance indicators (KPIs), track service health, and gain valuable context on the overall IT environment.
Event Management Connectors (sn_em_connector)	A collection of predefined interfaces (connectors) for ingesting events from third-party monitoring tools or Agent Client Collector into a ServiceNow instance.
Service Operations Workspace UI Components (sn_itom_uib_comp)	A joined repository for holding common internal UI components.

Metric Intelligence

Plugins activated with Metric Intelligence (com.snc.sa.metric)

Plugin name	Description
Operational Intelligence - Extension Health Monitoring (com.snc.sa.metric.health)	Contains some features that enable self-monitoring capabilities for the Metric Intelligence components in the MID Server.
Highcharts Reporting (com.glide.highcharts)	Provides some UI functionality to Metric Intelligence.
Event Management (com.glideapp.itom.snac)	Helps you to identify health issues across business services and infrastructure on a single management console. Event Management provides intelligent event and alert processing, including automatic alert remediation, aggregation, and root cause analysis (RCA)
ITOM UI Components (com.snc.itom.ui)	Shared UI components for ITOM products that provide some UI functionality to Metric Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins activated with Metric Intelligence (com.snc.sa.metric) (continued)

Plugin name	Description
MID Server Distributed Cluster (com.snc.agent.distributed.cluster)	Enables management of the mid cache component of Metric Intelligence from the ServiceNow instance.
List v3 Components (com.glide.ui.list_v3_components)	Provides some UI functionality to Metric Intelligence and components required for List v3, including REST endpoints, directives, and templates.
MetricBase (com.snc.clotho)	The database where the raw metric data such as CPU usage and Disk Free is saved and where model fitting for anomaly detection occurs.
Mid Web Server (com.snc.sa.mid.webserver)	A MID Server extension that enables developing REST APIs to send events and metrics to the MID Server. The plugin also enables the communication between the MID Server and Agent Client Collector.

Agent Client Collector

Plugins or applications installed with Agent Client Collector Framework (sn_agent)

Plugin or application name	Description
Agent Client Collector Global Utilities (com.snc.sa.accutils)	A set of globally scoped utilities that facilitates cross-scope transactions between Agent Client Collector Framework and a ServiceNow instance.
Configuration Management For Scoped Apps (CMDB) (com.snc.cmdb.scoped)	Provides scoped applications with access to Identification Engine APIs.
Mid Web Server (com.snc.sa.mid.webserver)	A MID Server extension that enables developing REST APIs to send events and metrics to the MID Server. The plugin also enables the communication between the MID Server and Agent Client Collector.

Plugins or applications installed with Agent Client Collector Monitoring (sn_itmon)

Plugin or application name	Description
Agent Client Collector Framework (sn_agent)	Manages Agent Client Collector on the ServiceNow instance and the MID Server.
Metric Intelligence (sn_metrics_ng)	A companion application to the family Metric Intelligence plugin that provides additional functionality and enhancements.
Service Operations Workspace Metric Explorer (sn_itom_metric_exp)	Enables operators to investigate metric behavior and compare and correlate different metrics across hosts, applications, and services so they can troubleshoot issues and find the root cause.
Event Management (com.glideapp.itom.snac)	Helps you to identify health issues across business services and infrastructure on a single management console. Event Management

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Agent Client Collector Monitoring (sn_itmon) (continued)

Plugin or application name	Description
	provides intelligent event and alert processing, including automatic alert remediation, aggregation, and root cause analysis (RCA)
Metric Intelligence (com.snc.sa.metric)	Provides the ability to capture, explore and analyze application metrics data, identifying and indicating anomalies. Metric Intelligence generates anomaly alerts, which can be promoted to IT alerts and appear on the Service Operations Workspace and Service Health dashboard.

Plugins or applications installed with Agent Client Collector for Visibility Content (sn_acc_vis_content)

Plugin or application name	Description
Agent Client Collector for Visibility Content Content (sn_acc_vis_content)	Provides agent-based infrastructure discovery and collects Software Asset Management (SAM) metrics to support Discovery and Digital End-User Experience (DEX) applications.
Discovery (com.snc.discovery)	Finds applications and devices on your network, and then updates the CMDB with the information it finds.
Agent Client Collector Framework (sn_agent)	Manages Agent Client Collector on the ServiceNow instance and the MID Server.
CMDB CI Class Models (sn_cmdb_ci_class)	Adds class models that extend the CMDB class hierarchy, including class descriptions, identification rules, identifier entries, and dependent relationships if applicable. Applications such as Discovery and Service Mapping can use these class extensions to populate configuration items (CIs) and discover various technologies and software.

Applications installed with Agent Client Collector Log Analytics (sn_accl)

Application name	Description
Agent Client Collector Monitoring (sn_itmon)	Provides the ability to monitor your configuration items (CIs). This application comes with predefined monitoring checks and policies that are applied on the agent.
Health Log Analytics (sn_occ)	Helps you solve issues faster by ingesting, analyzing, and correlating machine-generated log data in real-time to predict IT issues before they affect you.

Plugins or applications installed with Agent Client Collector Security Incident Response (sn_secops_acc)

Plugin or application name	Description
Security Incident Response (sn_si)	Enables you to manage and automate the life cycle of your security incidents from initial prioritization to containment and resolution. Use the automated workflows to respond quickly and consistently, and understand trends and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugins or applications installed with Agent Client Collector Security Incident Response (sn_secops_acc) (continued)

Plugin or application name	Description
	bottlenecks through the data provided by analytics-driven dashboards and comprehensive reporting systems.
Security Support Orchestration (sn_sec_cmn_orch)	Dependency for the Vulnerability Response and Security Incident Response applications.
Threat Intelligence Support Common (sn_ti)	Enables you to find Indicators of Compromise (IoCs) and enrich security incidents with threat intelligence data. Threat Intelligence enables you to access and provide a reference point for your company's Structured Threat Information Expression (STIX) data.
Agent Client Collector Spoke (sn_acc_spoke)	Provides a framework to write customized workflows or automations to execute on agents or CIs.
Security Support Core (com.snc.security_support.core)	Provides a robust core data model and UI framework specifically for security operations integrations. This comprehensive model serves as the backbone for seamlessly integrating various security tools and systems, promoting smooth data flow and efficient operations.
Security Support Common (sn_sec_cmn)	Enables core functionality used by the Security Operations applications.
Security Integration Framework (sn_sec_int)	Manages the common functionality for Security Operations integrations including both Security Incident Response and Vulnerability Response. Security Integration Framework enables you to use existing features and functionality without having to reimplement them.
Agent Client Collector Framework (sn_agent)	Manages Agent Client Collector on the ServiceNow instance and the MID Server.
Vulnerability Response Dependencies (com.snc.si_dep)	Activates all the dependent plugins required to support the Vulnerability Response application.

Service Observability

Plugins or applications installed with Service Observability (sn_service_obs)

Plugin or application name	Description
Service Observability UI (sn_service_obs_ui)	Enables you to use Service Observability in the Service Operations Workspace (SOW) and the SOW Admin Center.
Service Reliability Management (sn_sow_srm)	Service Reliability Management and Service Observability display service performance metrics together to give a holistic view into service health.
Platform Analytics (sn_app_analytics_w)	Enables customizing the dashboards.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Plugins or applications installed with ITOM Cloud Accelerate
- Plugins or applications installed with ITOM Visibility
- Plugins activated with ITOM Optimization
- Applications installed with ServiceNow Otto for ITOM

AI in ITOM AIOps

AI features in ITOM AIOps help operators triage alerts, investigate incidents, analyze service health, and improve service reliability using generative AI and agentic workflows.

AI features in ITOM AIOps span several products. The following tables summarize all available AI features by type and product.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

Agentic workflows

ServiceNow Otto for ITOM includes agentic workflows that use AI agents to complete alert management and investigation tasks autonomously or with minimal operator input.

Agentic workflows in ITOM AIOps

Product	AI feature	Description	Use case	Resources
Event Management, Express List	Analyze alert impact agentic workflow	Uses observability AI agents to investigate an alert and surface information about its severity, impact, ownership, and potential root causes. Integrates with observability tools such as Dynatrace, Kentik, and New Relic.	An operator selects an alert from an observability-connected source and wants to understand its blast radius, identify responsible teams, and find potential root causes before escalating.	• Analyze alert impact agentic workflow • Analyze alert impact in the ServiceNow Otto panel
Event Management, Express List	Manage alerts autonomously agentic workflow	Provides a unified AI-driven process that automates alert triage, impact analysis, and root cause investigation. Generates reports, summarizes key insights and possible next steps, and stores structured findings in the alert record. Also analyzes Health Log Analytics alerts in	An operations team wants alerts to be triaged, analyzed, and summarized automatically as they	• Manage alerts autonomously agentic workflow • Configure the manage alerts autonomously

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agentic workflows in ITOM AIOps (continued)

Product	AI feature	Description	Use case	Resources
		an alert group and can classify alerts as proactive or reactive.	arrive, with AI insights available in Express List without manual intervention.	agentic workflow • Review AI-generated alert insights in Express List
Event Management, Express List	Analyze potential impact agentic workflow	Analyzes how a change request might affect relevant servers and services. The agent verifies prerequisites, retrieves the change request, selects up to 10 affected servers, identifies matches with suggested services, and displays an impact analysis in the ServiceNow Otto panel. The analysis is also saved to the change request work notes. Includes Service Mapping Candidate and Service Mapping Candidates Impact skills.	An operator or change manager wants to assess the risk of a change request before approving it, without manually reviewing all affected CIs.	• Analyze potential impact agentic workflow • Assess a change request with the Analyze potential impact workflow

AI agents

Standalone AI agents in ITOM AIOps run autonomously on a schedule or in response to a trigger, without requiring user input.

AI agents in ITOM AIOps

Product	AI feature	Description	Use case	Resources
Service Reliability Management	SLO creator agent	Runs every 14 days and processes up to 25 services or CIs that do not already have auto-generated SLOs. Analyzes historical alerts, incidents, and outages to generate one SLO per service or CI with up to 10 service level indicators (SLIs). Generated SLOs are active by default. Email notifications are sent to team managers when SLOs are generated and when the error budget falls to 25% or lower. Requires ServiceNow Otto for ITOM and Service Reliability Management.	A team wants to adopt SLOs without manually defining them, using historical incident and alert data to generate a starting point for each service.	Generating service level objectives

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generative AI skills

Generative AI skills in ITOM AIOps are discrete, on-demand capabilities that operators invoke to generate analyses, summaries, or descriptions for a specific alert, group, or service.

Generative AI skills in ITOM AIOps

Product	AI feature	Description	Use case	Resources
Event Management, Express List	Alert analysis skill	Generates a human-readable brief of an alert and additional technical information to support investigation. Available in the ServiceNow Otto panel, on the alert Overview tab in Service Operations Workspace, and in the Express List preview panel. Analyses are provided in English regardless of the language used in the alert description.	An operator opens an alert and wants a quick summary of what it means and what to investigate next.	• View an alert analysis by ServiceNow Otto in Service Operations Workspace • View an alert analysis by ServiceNow Otto in Express List
Express List	Alert group analysis skill	Generates a simplified, human-readable description of an alert group and technical information to support investigation. Available in the Express List preview panel.	An operator selects an alert group and wants to understand its overall significance before deciding how to respond.	View an alert group analysis by ServiceNow Otto in Express List
Express List	Alert group description generation skill	Generates a comprehensive description of an alert group that encompasses all alerts within the group. The generated description replaces the original alert group description.	An operator wants to replace a system-generated alert group description with a meaningful, AI-generated summary that reflects the full scope of the group.	Generate an alert group description in Express List using ServiceNow Otto
Event Management, Express List	Past incident analysis skill	Queries historical records to find related past incidents and analyzes their frequency, criticality, work notes, and resolution. Presents a summary of the most relevant incidents in the ServiceNow Otto panel, including resolution strategies	An operator is investigating an alert and wants to know whether similar issues have occurred before and	• Accelerate alert resolution with past incident analysis

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generative AI skills in ITOM AIOps (continued)

Product	AI feature	Description	Use case	Resources
		and contact details for teams who resolved similar incidents.	how they were resolved.	• Generate a ServiceNow Otto summary of past related incidents
Express List	AI-generated incident creation skill	Creates an incident with a human-readable, AI-generated description directly from an alert in Express List.	An operator determines that an alert requires an incident and wants the incident description generated automatically rather than written manually.	Create an incident from an alert with ServiceNow Otto in Express List
Service Observability	Analyze dashboard skill	Summarizes a single Service Observability dashboard using generative AI, including calling out insights found in the charts. The skill analyzes data within the selected time period and reports any insights it finds. Results remain available for one hour. Requires manual activation.	An operator wants a quick orientation to an unfamiliar dashboard or needs to identify anomalies in chart data without manually reviewing each chart.	• Activate the analyze Service Observability dashboard skill • Analyze a dashboard in Service Observability
Service Observability	Analyze service health skill	Analyzes all available Service Observability dashboards for a selected service. Detects insights, reports on general service health, and helps operators assess service impact, identify root causes, and streamline incident response. When activated, the skill also runs automatically during incident investigation from the **Investigate** tab. Requires manual activation.	An operator is triaging an incident and needs to quickly assess service impact, identify root causes, and determine next steps across all dashboards for the affected service.	• Activate the analyze service health skill • Analyze service health in Service Observability

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuring Service Operations Workspace for ITOM

Configure Service Operations Workspace for ITOM so that you can manage IT alert volumes efficiently and customize the workspace to enhance service delivery and operational visibility.

The Service Operations Workspace for ITOM is a centralized hub designed to optimize IT Operations Management by focusing on the handling and resolution of alerts. By leveraging this workspace, administrators can streamline the process of monitoring, triaging, and resolving alerts, thus enhancing service delivery and operational efficiency.

Note: Installing the AIOps Experience application [sn_sow_aiops] provides you with Service Operations Workspace for ITOM.

If you want information regarding subscriptions for ITOM Service Operations Workspace, check the ITOM AIOps subscription at Subscription for ITOM Health.

When setting up the Service Operations Workspace, it's crucial to configure both IT Operations Management (ITOM) for managing alerts and IT Service Management (ITSM) for resolving incidents and delivering services. This integrated setup ensures access to the full range of features, facilitating comprehensive management of IT operations and service quality.

To know about the ITSM configurations, see Configuring Service Operations Workspace for ITSM.

Install Service Operations Workspace for ITOM application

You can install the AIOps Experience [sn_sow_aiops] application to get Service Operations Workspace for ITOM if you have the admin role.The application includes demo data and installs related ServiceNow® Store applications if they are not already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the AIOps Experience application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- For a list of plugins or applications installed with Service Operations Workspace for ITOM, see Plugins or applications installed with ITOM AIOps.
- Service Operations Workspace supports domain separation. For details, see Domain separation and Event Management.

Role required: admin

About this task

The following items are installed with AIOps Experience [sn_sow_aiops] plugin.

- Roles
- Plugins

For more information, see Components installed with AIOps Experience.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the AIOps Experience [sn_sow_aiops] using the filter criteria and search bar.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

In the list next to the **Install** button, the versions that are available to you are displayed.

3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.

4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. **Optional:** If demo data is available and you want to install it, select the **Load demo data** check box.
   Demo data are the sample records that describe application features for common use cases. Load the demo data when you first install the application on a development or test instance.
6. Select **Install**.

ITOM AIOps configuration center

The ITOM AIOps configuration center is a centralized hub to manage AIOps settings. It helps you integrate monitoring tools, optimize alerts, metrics, and logs, manage services, dashboards, teams, and authorization—streamlining event-to-response workflows for better visibility, noise reduction, and operational efficiency.

Navigate to **Workspaces** > **Service Operations Workspace**. Select the ITOM AIOps configuration center icon ( ) at the bottom of the left navigation pane.

The ITOM AIOps configuration center page appears. On the ITOM AIOps configuration center page, under the **Optimize** > **Dashboards** section, select the dashboard.

Note: The ITOM AIops configuration center is accessible only to evt_mgmt_admin.

The AIOps configuration center provides a single interface to configure and administer AIOps capabilities in Service Operations Workspace (SOW). You can:

- Integrate monitoring tools to bring in events, logs, and metrics.
- Optimize how alerts are enriched, grouped, ignored, and responded to.
- Fine-tune metrics and log rules for accurate anomaly detection and alerting.
- Set up synthetic monitoring for app uptime.
- Manage services (service monitoring dashboards, SLIs/SLOs, error budgets) for Service Observability and Service Reliability Management.
- Access AIOps dashboards to track performance, operations, and business value.
- Configure authorization for admins, integrations, APIs, and notifications.
- Govern teams and approvals for service creation and on-call schedules.

The AIOps configuration center consolidates all administrative tasks needed to set up and operate ITOM AIOps effectively. It streamlines monitoring, improves incident response, and maintains service reliability.

Configure observability agents for ServiceNow Otto

Configure observability agents for third-party application performance monitoring (APM) or network performance monitoring (NPM) vendors. These agents are invoked by the analyze alert

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

impact agentic workflow. You must configure connections to those vendors before they can be invoked.

Important: This AI agent is turned on by default. For more information, see AI agents, skills, and agentic workflows on by default.

After you configure the agent(s), they can surface information from alerts generated by third-party systems to help you investigate alerts and incidents in the Service Operations Workspace.

Connections to vendors use MCP when possible, otherwise they use an API connection. You need connection and credential information to complete the connection process as shown in the following tables.

Note: These agents are different from the data sources used in the Service Observability dashboards.

Before you begin

Before configuring the integration agents, you must do the following:

- Install ServiceNow Otto for IT Operations Management (ITOM).
- Integrate third-party alerts with Event Management.
- Copy the connection requirements for the vendor as noted in the following tables. You will need these to create credentials and connections to the third-party vendor.

Role required: connection_admin and credential_admin

AWS CloudWatch MCP

Connection information	Value
Agent name	AWS CloudWatch MCP Server Agent
Overview of data returned	Alarm details, metric trend analysis, CloudWatch logs (anomalies, error patterns, log insights queries), CMDB resource context, correlated service metrics, and root cause analysis with recommended next steps
Credential & Connection Alias name	AWS CloudWatch MCP server
Connection type	MCP
Returned data type	APM
Connection URL	`https://your-mcp-server-host/mcp` This assumes you have deployed the CloudWatch MCP server using a MID Server instead of a publicly exposed EC2 instance. For more information about deploying the MCP server, see the AWS CloudWatch MCP Server — MID Server Deployment Guide [KB3030674] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
Required credentials	• AWS access key ID • AWS secret access key
Required scope	AWS IAM permissions: • `cloudwatch:Describe*` • `cloudwatch:Get*` • `cloudwatch:List*` • `logs:Describe*` • `logs:Get*` • `logs:StartQuery` • `logs:StopQuery` • `logs:GetQueryResults`

AWS CloudWatch API

Connection information	Value
Agent name	AWS CloudWatch API Agent
Overview of data returned	Alarm details, metric trend analysis, CloudWatch logs (anomalies, error patterns, log insights queries), CMDB resource context, correlated service metrics, and root cause analysis with recommended next steps
Credential & Connection Alias name	AWS CloudWatch API Credentials
Connection type	API (MCP fallback mechanism)
Returned data type	APM
Connection URL	N/A
Credential type	AWS Credentials
Authentication algorithm	AWS CloudWatch Algorithm
Required credentials	• AWS access key ID • AWS secret access key

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
Required scope	AWS IAM permissions: • `cloudwatch:Describe*` • `cloudwatch:Get*` • `cloudwatch:List*` • `logs:Describe*` • `logs:Get*` • `logs:StartQuery` • `logs:StopQuery` • `logs:GetQueryResults`

Datadog

Connection information	Value
Agent name	Datadog APM MCP Server Agent
Overview of data returned	Service health, distributed traces, triggered monitors, log analysis, incidents, SLO compliance, deployment events, and service dependencies
Credential & Connection Alias name	Datadog APM MCP Connection
Connection type	MCP
Returned data type	APM
Connection URL	`https://mcp.datadoghq.com/api/unstable/mcp-server/mcp?toolsets=core,alerting,apm,error-tracking`
Required credentials	• Datadog API key • Datadog application key
Required scope	• `mcp_read` • `apm_read` • `timeseries_query` • `dashboards_read` • `logs_read_data` • `logs_read_index_data` • `user_access_read` • `apm_service_catalog_read`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
	• `teams_read` • `slos_read` • `metrics_read` • `cloud_cost_management_read` • `incident_read` • `monitors_read`

Dynatrace

Connection information	Value
Agent name	Dynatrace MCP Server Agent
Overview of data returned	Insights about logs, topology, recent changes, root causes, impacted entities, and environments.
Credential & Connection Alias name	Dynatrace MCP server
Connection type	MCP
Returned data type	APM
Connection URL	URL of your Dynatrace instance. Dynatrace URLs follow this format: `https://<your-resource-name>.apps.dynatrace.com/platform-reserved/mcp-gateway/v0.1/servers/dynatrace-mcp/mcp`
Required credentials	Platform token (must be prefixed with `Bearer`). For example, `Bearer dt0s01.STABCDEF12345.G3HIJKLMNOP`.
Required scope	IAM policy and group assignment that allows the following scopes: • `davis-copilot:nl2dql:execute` • `davis-copilot:dql2nl:execute` • `davis-copilot:conversations:execute` • `davis:analyzers:read` • `davis:analyzers:execute` • `mcp-gateway:servers:invoke` • `mcp-gateway:servers:read` • `storage:buckets:read` • `storage:logs:read` • `storage:events:read`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
	• `storage:security.events:read` • `storage:metrics:read` • `storage:bizevents:read` • `storage:spans:read` • `storage:entities:read` • `storage:smartscape:read` • `storage:system:read`
Required Dynatrace Intelligence settings	• Enable generative AI • Enable document suggestions • Enable environment-aware queries

Google Cloud

Note: Before configuring this connection, create a keystore file by following the steps in Create a Java KeyStore certificate.

Connection information	Value
Agent name	Gemini Cloud Assist A2A Investigation Agent
Overview of data returned	Summary, findings, root-cause hypotheses, next steps, impact on affected Google Cloud resources, ownership, recommended actions, supporting evidence, and sources
Credential & Connection Alias name	Google Gemini Cloud Assist Agent
Connection type	Agent-to-agent
Returned data type	APM
Required credentials	• Google Project ID • Service Account Email • Private Key ID • Keystore Password • Keystore file (`.jks`)
Required scope	Google Cloud IAM permissions: • Assign the Investigation Creator role to the service account • Give the service account read-access to the resources that you want it to investigate

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
Required Google Cloud settings	• Enable the required APIs for your Google Cloud project: `logging.googleapis.com`, `monitoring.googleapis.com`, `cloudresourcemanager.googleapis.com`, and `geminicloudassist` • Gemini Cloud Assist must be set up in the Google Cloud project

Kentik

Connection information	Value
Agent name	Kentik analysis AI agent
Overview of data returned	Service network performance, connectivity, DDOS attacks, and anomalies
Credential & Connection Alias name	Kentik analysis AI agent
Connection type	API
Returned data type	APM
Connection URL	URL of your Kentik instance. Kentik URLs follow this format: `https://<your-resource-name>.api.kentik.com`
Required credentials	• User email • API token
Required scope	Can view devices

LogicMonitor

Connection information	Value
Agent name	LogicMonitor API Agent
Overview of data returned	Alert details, metric trend analysis, device properties and classification, collector health, SDT/maintenance windows, related entity topology, CMDB service associations, recent ServiceNow Change records, top processes, and root cause analysis with verdict (degraded / healthy / cleared-and-sustained / likely-false-positive / inconclusive) and recommended next steps.
Credential &	LogicMonitor API connection

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
Connection Alias name	
Connection type	API
Returned data type	APM
Connection URL	`https://<account>.logicmonitor.com/santaba/rest` Replace `<account>` with your LogicMonitor portal subdomain.
Required credentials	API bearer token: An LMv2 personal access token
Required scope	LogicMonitor API v3 read-only access: Devices (view), Alerts (view), Dashboards/ Metrics (view), Collectors (view), SDTs (view), Device Groups (view). The API token inherits the creating user's role. Any user with the default read only or administrator role has sufficient access.

Microsoft Azure

Connection information	Value
Agent name	Azure Monitor MCP Agent
Overview of data returned	Alert details, resource health status, activity logs, Log Analytics query results, metric trend analysis, CMDB resource context, correlated service metrics, and root cause analysis with recommended next steps
Credential & Connection Alias name	Azure Monitor MCP Connection
Connection type	MCP
Returned data type	APM
Connection URL	`http://<your-azure-mcp-server>`
Required credentials	• Azure connection URL • Azure username • Azure password
Required scope	• Reader role on target Azure subscriptions • Log Analytics Reader role on target Log Analytics workspaces

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
Required Microsoft Azure settings	• Microsoft Azure MCP Server must be deployed and reachable • Log Analytics workspaces must exist in the target subscription • Azure Monitor must be enabled for target resources

New Relic

Connection information	Value
Agent name	New Relic MCP Server Agent
Overview of data returned	Service and user impact, root cause theories, and responsible teams.
Credential & Connection Alias name	New Relic MCP Connection
Connection type	MCP
Returned data type	APM
Connection URL	`https://mcp.newrelic.com/mcp/`
Required credentials	API key (also known as "User Key") Header name: `api-key`
Required scope	N/A

Prometheus

Connection information	Value
Agent name	Prometheus API Agent
Overview of data returned	PromQL metric queries (CPU, memory, disk, network), active alerts, alert rule definitions, and scrape target health
Credential & Connection Alias name	Prometheus connection
Connection type	API
Returned data type	APM
Connection URL	`https://<your-onprem-prometheus-server>`
Required credentials	• MID Server Prometheus is installed on • User name (for outbound connection) • Password (for outbound connection)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
Required scope	N/A

SolarWinds

Connection information	Value
Agent name	SolarWinds analysis AI agent
Overview of data returned	On-premises data from SolarWinds Orion: node health/status, CPU/memory, packet loss/latency, interface utilization/errors/discards, active alerts/history, affected entities/services, and trend/baseline metrics used for root-cause investigation
Credential & Connection Alias name	SolarWinds AI Agent
Connection type	API
Returned data type	NPM
Connection URL	`https://<your-onprem-solarwinds-server>`
Required credentials	• MID Server SolarWinds is installed on • User name • Password
Required scope	SolarWinds Orion API/SWQL read access (NPM/APM) via MID Server

Splunk

Connection information	Value
Agent name	Splunk MCP Server Agent
Overview of data returned	SPL query results from Splunk indexes, index/sourcetype metadata, and structured investigation findings including affected entities, root cause analysis, and recommended actions
Credential & Connection Alias name	Splunk MCP Connection
Connection type	MCP

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connection information	Value
Returned data type	APM
Connection URL	`https://<your-splunk-instance>.splunkcloud.com`
Required credentials	Splunk MCP token
Required scope	For Splunk token generation: • `mcp_user` role (or any role with `mcp_tool_execute`) required for all MCP tool calls • Search capability and read access to the relevant indexes for running queries • Knowledge object read permissions for Get Knowledge Objects • Token must be generated using the Splunk MCP Server app with `audience = mcp` (not a standard Splunk API token)

ThousandEyes

Connection information	Value
Agent name	ThousandEyes MCP Server Agent
Overview of data returned	Test configuration and status, aggregated metrics (response time, packet loss, latency, jitter, throughput, availability), metric anomalies with deviation analysis, network events and routing changes, ISP/network outages with provider and ASN details, hop-by-hop path visualization, and root cause analysis with ranked probable causes and recommended next steps
Credential & Connection Alias name	ThousandEyes MCP Connection
Connection type	MCP
Returned data type	NPM
Connection URL	`https://api.thousandeyes.com/mcp`
Required credentials	API key prefixed with `Bearer`. For example, `Bearer <api-key>`.
Required scope	ThousandEyes API access with permissions to read tests, metrics, anomalies, events, outages, and path visualization data

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **sys_alias.LIST**.
2. Search for and select the vendor's connection name as shown in the preceding tables.
3. Select **Create New Connection & Credential**.
4. Fill in the form, using the information in the preceding tables.
5. Select **Create**.

   Your connection appears in the **Connections** tab.

The agents are now ready to be used by the analyze alert impact agentic workflow

Customize the Service Operations Workspace for ITOM home page

You can customize the data displayed in sub pages of the home page that is available in the base system. For any visual styling changes, you need to create a copy of that landing page and edit it.

Modify the greeting text on the landing page

Customize the header message that is a part of the greeting text on the home page. This creates a personalized user experience and sets the right tone for the information presented.

Before you begin

Role required: evt_mgmt_admin or admin

About this task

The Get service desk agent shift data broker server script calls the `SOWServiceDeskAgentShift` script includes to display the shift information in the greeting text.

You can customize the display text for the default landing page. To change the visual styling of the text, create a copy of the landing page.

Suppose you are on call until Monday, March 21. You can update your greetings to reflect your on-call status.

Header

Procedure

1. Navigate to **All** > **Service Operations Workspace Admin Center** > **Overview**.
2. On the **Overview** tab, from the **Initial Setup** section, select the **Configure the landing page** option.
3. On the Landing Page form, select **Configure** next to **Greetings**.
   The Script Includes form is displayed for SOWServiceDeskAgentShift, which internally calls the SOWServiceDeskAgentShiftSNC
4. Copy the **Script** value of SOWServiceDeskAgentShiftSNC.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **All** > **System Definition** > **Script Includes**.

b. From the list of Script Includes, open SOWServiceDeskAgentShiftSNC.

Script Includes list

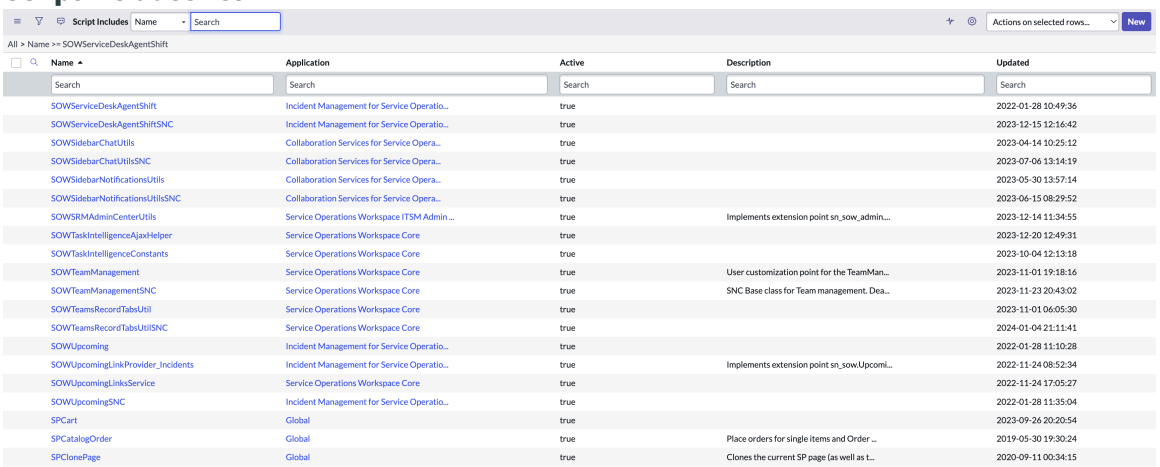

Script Includes | Name | Search | Actions on selected rows... | New

All > Name >= SOWServiceDeskAgentShift

Name	Application	Active	Description	Updated
SOWServiceDeskAgentShift	Incident Management for Service Operatio...	true		2022-01-28 10:49:36
SOWServiceDeskAgentShiftSNC	Incident Management for Service Operatio...	true		2023-12-15 12:16:42
SOWSidebarChatUtils	Collaboration Services for Service Opera...	true		2023-04-14 10:25:12
SOWSidebarChatUtilsSNC	Collaboration Services for Service Opera...	true		2023-07-06 13:14:19
SOWSidebarNotificationsUtils	Collaboration Services for Service Opera...	true		2023-05-30 13:57:14
SOWSidebarNotificationsUtilsSNC	Collaboration Services for Service Opera...	true		2023-06-15 08:29:52
SOWSRMAdminCenterUtils	Service Operations Workspace ITSM Admin ...	true	Implements extension point sn_sow_admin....	2023-12-14 11:34:55
SOWTaskIntelligenceAjaxHelper	Service Operations Workspace Core	true		2023-12-20 12:49:31
SOWTaskIntelligenceConstants	Service Operations Workspace Core	true		2023-10-04 12:13:18
SOWTeamManagement	Service Operations Workspace Core	true	User customization point for the TeamMan...	2023-11-01 19:18:16
SOWTeamManagementSNC	Service Operations Workspace Core	true	SNC Base class for Team management. Dea...	2023-11-23 20:43:02
SOWTeamsRecordTabsUtil	Service Operations Workspace Core	true		2023-11-01 06:05:30
SOWTeamsRecordTabsUtilSNC	Service Operations Workspace Core	true		2024-01-04 21:11:41
SOWUpcoming	Incident Management for Service Operatio...	true		2022-01-28 11:10:28
SOWUpcomingLinkProvider_Incidents	Incident Management for Service Operatio...	true	Implements extension point sn_sow.Upcomi...	2022-11-24 08:52:34
SOWUpcomingLinksService	Service Operations Workspace Core	true		2022-11-24 17:05:27
SOWUpcomingSNC	Incident Management for Service Operatio...	true		2022-01-28 11:35:04
SPCart	Global	true		2023-09-26 20:20:54
SPCatalogOrder	Global	true	Place orders for single items and Order ...	2019-05-30 19:30:24
SPClonePage	Global	true	Clones the current SP page (as well as t...	2020-09-11 00:34:15

c. Copy the **Script** field value.

Script includes details page

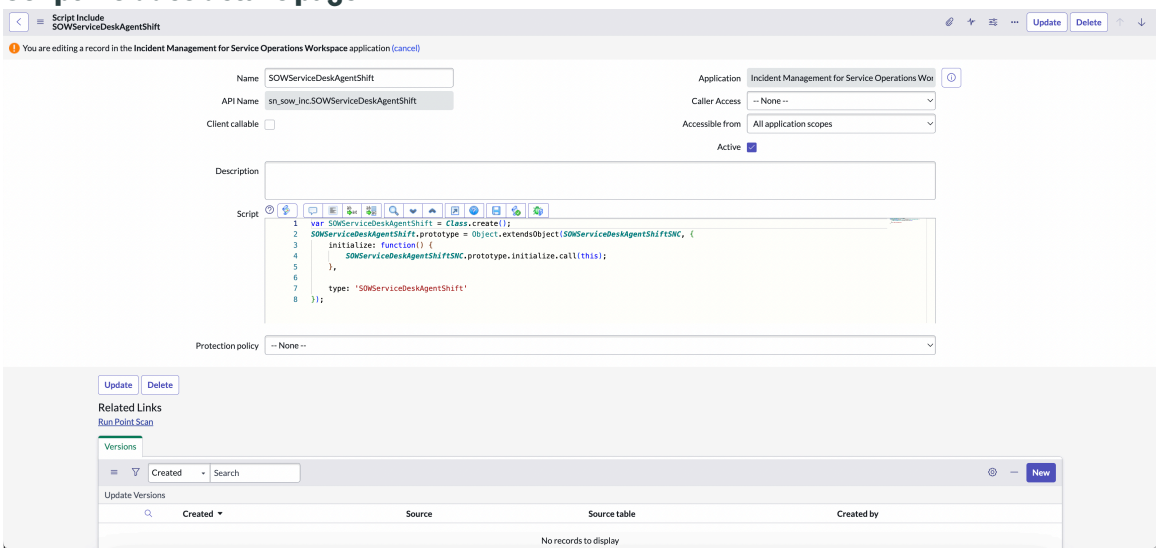

Script Include SOWServiceDeskAgentShift | Update | Delete

You are editing a record in the **Incident Management for Service Operations Workspace** application (cancel)

Name: SOWServiceDeskAgentShift
API Name: sn_sow_inc.SOWServiceDeskAgentShift
Client callable
Application: Incident Management for Service Operations Wor
Caller Access: -- None --
Accessible from: All application scopes
Active

Description

Script

```
var SOWServiceDeskAgentShift = Class.create();
SOWServiceDeskAgentShift.prototype = Object.extendsObject(SOWServiceDeskAgentShiftSNC, {
    initialize: function() {
        SOWServiceDeskAgentShiftSNC.prototype.initialize.call(this);
    },

    type: 'SOWServiceDeskAgentShift'
});
```

Protection policy: -- None --

Update | Delete

Related Links
Run Point Scan

Versions

Created | Search | New

Update Versions

Created	Source	Source table	Created by
No records to display			

5. Edit the SOWServiceDeskAgentShift script.

a. Navigate to **All** > **System Definition** > **Script Includes**.

b. From the list of Script Includes, open SOWServiceDeskAgentShift.

c. Paste the **Script** field value that you copied from SOWServiceDeskAgentShiftSNC.

```
initialize: function() {

    this.PLUGINS = {

      ON_CALL: "com.snc.on_call_rotation",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      SHIFT_PLANNING: "com.sn_shift_planning"

    };

    this.DAYS_AHEAD =
 gs.getProperty('com.snc.on_call_rotation.upcoming_rota_day
s', 30);

    this.DATE_FORMAT = gs.getDateFormat();

    this.TIME_FORMAT = gs.getTimeFormat();

    this.WFO_SCHEMA_IDENTIFIER = 'readInTimezone';

    this.WFO_ONCALL_TYPE_IDENTIFIER =
 'oncallAPIReadInTimezone';

    this.WEEK_DAY_FORMAT = 'EEEE';

    this.FORMAT = {

      SERVICE_DESK_HEADER: 'SERVICE_DESK_HEADER'

    };

    this.DATE_TIME_FORMAT_FULL = "full";

  },



  getUpcomingShift: function(userId, outputFormat, wfoShift)
 {

    var upcomingShifts, gd = new GlideDate();

    var from = gd.getDisplayValueInternal();

    gd.addDaysLocalTime(this.DAYS_AHEAD);

    var to = gd.getDisplayValueInternal();

    var mayBeWFOUser = userId == gs.getUserID() ?
 gs.hasRole("sn_shift_planning.agent") : true;

    if
 (GlidePluginManager.isActive(this.PLUGINS.SHIFT_PLANNING) &&
 mayBeWFOUser && wfoShift)

      upcomingShifts = this.getWFOShifts(from, to, userId);

    var isWFOResponseValid = Array.isArray(upcomingShifts) &&
 upcomingShifts.length > 0;

    if (!isWFOResponseValid &&
 GlidePluginManager.isActive(this.PLUGINS.ON_CALL))
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      upcomingShifts = this.getOnCallShifts(from, to,
userId);

    return this.getFormattedShift(upcomingShifts,
outputFormat);

  },



  getWFOShifts: function(from, to, userId) {

    var smUtil = new sn_shift_planning.ScheduleManagerUtil();

    var defaultEventCategories = [

      smUtil.WORK_CATEGORY_SYS_ID,

      smUtil.ON_CALL_CATEGORY_SYS_ID

    ];

    var params = {

      "assignmentGroups": "",

      "individualAgents": userId,

      "startDate": from,

      "endDate": to,

      "published": true,

      "eventCategories": defaultEventCategories,

      "filter": {

        "sn_shift_planning_agent_schedule": {

          "event": ["8d03fd9b0f310010717cc562ff767e4a",
"936319993bd02010e81a4ee1f3efc419"]

        }

      }

    };

    var upcomingShifts = [];

    try {

      var agentScheduleRESTAPIUtils = new
sn_shift_planning.AgentScheduleRESTAPIUtils();
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      var spans =
 agentScheduleRESTAPIUtils.retrieveAgentSchedulesFromParams(p
arams);

      if (spans && spans.shiftData &&
 spans.shiftData.user && spans.shiftData.user.data
 && spans.shiftData.user.data.length > 0 &&
 spans.shiftData.user.data[0].schedules) {

        var schedules =
 spans.shiftData.user.data[0].schedules,

          arrayUtil = new global.ArrayUtil();

        for (var schedule in schedules)

          upcomingShifts = arrayUtil.union(upcomingShifts,
 schedules[schedule]);

      }

    } catch (e) {

      gs.error("SOWServiceDeskAgentShiftSNC: Error fetching
 WFO shifts");

    }



    return upcomingShifts;

  },



  getOnCallShifts: function(from, to, userId) {

    var onCallCommon = new global.OnCallCommon();

    var onCallGroups = onCallCommon.getMyGroups();

    var spans = new global.OCRotationV2(null,
 global.OCRotationV2.DEFAULT_FORMAT_OPTION).setStartDate(from
).setEndDate(to).setUserIds(userId).setGroupIds(onCallGroups)
.getSpans();

    var upcomingShifts = [];

    for (var i = 0; i < spans.length; i++)

      if ((spans[i].type == "roster" || spans[i].type ==
 "override") && spans[i].user_id == userId)

        upcomingShifts.push(spans[i]);

    return upcomingShifts;
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 },


 getFormattedShift: function(upcomingShifts, outputFormat) {
   if (gs.nil(upcomingShifts))
     return;
   switch (outputFormat) {
     case this.FORMAT.SERVICE_DESK_HEADER:
       return
this.getServiceDeskFormattedShift(upcomingShifts);
     default:
       return upcomingShifts;
   }
 },


 getServiceDeskFormattedShift: function(upcomingShifts) {
   var response = {
     shiftText: "",
     nearestShift: {}
   };
   var startAttr = "start",
     endAttr = "end",
     isOnCall = true;
   var dateFormat = 'EEEE, MMM dd',
     timeFormat = 'hh:mm a';
   var isWFOSchema = Array.isArray(upcomingShifts) &&
upcomingShifts.length > 0 && this.WFO_SCHEMA_IDENTIFIER in
upcomingShifts[0];
   if (isWFOSchema)
     startAttr = "scheduledStartTime", endAttr =
"scheduledEndTime";
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    response.nearestShift =
 this.getNearestShift(upcomingShifts, startAttr, endAttr);

    if (response.nearestShift) {

      if (isWFOSchema) {

        if (response.nearestShift.schedule
 && !(this.WFO_ONCALL_TYPE_IDENTIFIER in
 response.nearestShift.schedule))

          isOnCall = false;

      }

      if (isOnCall) {

        if (response.nearestShift.isCurrent) {

          if (response.nearestShift.roster_name)

            response.shiftText = gs.getMessage("You are
 on-call until {0} as {1}",
 [this.getFormattedDateTime(response.nearestShift[endAttr]),
 response.nearestShift.roster_name]);

          else

            response.shiftText = gs.getMessage("You are
 on-call until {0}",
 [this.getFormattedDateTime(response.nearestShift[endAttr])
]);

        } else {

          if (response.nearestShift.roster_name)

            response.shiftText = gs.getMessage("Your
 next on-call shift starts on {0} as {1}",
 [this.getFormattedDateTime(response.nearestShift[startAttr
]), response.nearestShift.roster_name]);

          else

            response.shiftText = gs.getMessage("Your
 next on-call shift starts on {0}",
 [this.getFormattedDateTime(response.nearestShift[startAttr])
]);

        }

      } else {

        if (response.nearestShift.isCurrent) {

          response.shiftText = gs.getMessage("Your shift ends
 on {0}",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 [this.getFormattedDateTime(response.nearestShift[endAttr])
]);

        } else

          response.shiftText = gs.getMessage("Your next shift
 starts on {0}",
 [this.getFormattedDateTime(response.nearestShift[startAttr])
]);

      }

    }



    return response;

  },



  getNearestShift: function(upcomingShifts, startAttr,
 endAttr) {

    if (gs.nil(upcomingShifts) || gs.nil(startAttr) ||
 gs.nil(endAttr))

      return;

    var nearestShift;

    upcomingShifts.sort(function(shift_1, shift_2) {

      var gdt1 = new GlideDateTime();

      gdt1.setDisplayValueInternal(shift_1[startAttr]);

      var gdt2 = new GlideDateTime();

      gdt2.setDisplayValueInternal(shift_2[startAttr]);

      return GlideDateTime.subtract(gdt2,
 gdt1).getNumericValue();

    });

    var currentDateTime = new GlideDateTime(),

      startDateTime = new GlideDateTime(),

      endDateTime = new GlideDateTime();

    for (var i = 0; i < upcomingShifts.length; i++) {
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 startDateTime.setDisplayValueInternal(upcomingShifts[i][star
tAttr] + "");

 endDateTime.setDisplayValueInternal(upcomingShifts[i][endAt
tr] + "");
      if (currentDateTime.onOrBefore(endDateTime)) {
        nearestShift = upcomingShifts[i];
        if (currentDateTime.onOrAfter(startDateTime))
          nearestShift.isCurrent = true;
        else
          nearestShift.isCurrent = false;
        break;
      }
    }
    return nearestShift;
  },


  getFormattedDateSpan: function(date, dateFormat,
timeFormat) {
    var formattedDate = this.getFormattedDate(date,
dateFormat);
    var dateAmPmTime = this.getFormattedTime(date,
timeFormat);
    return formattedDate + " " + dateAmPmTime;
  },


  getFormattedDate: function(inputDate, format) {
    var date = new GlideDateTime();
    date.setDisplayValueInternal(inputDate);
    format = !format ? this.DATE_FORMAT : format;
    var gd = date.getLocalDate();
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    return gd.getByFormat(format);
  },


  getFormattedTime: function(inputDate, format) {
    var date = new GlideDateTime();
    date.setDisplayValueInternal(inputDate);
    format = !format ? this.TIME_FORMAT : format;
    var gt = date.getLocalTime();
    return gt.getByFormat(format);
  },


  getFormattedDateTime: function(date) {
    var dateTime = new GlideDateTime();
    dateTime.setDisplayValueInternal(date);
    return
dateTime.getDisplayValueLang(this.DATE_TIME_FORMAT_FULL);
  },


  endsToday: function(inputDate) {
    var inputGdt = new GlideDateTime();
    inputGdt.setDisplayValueInternal(inputDate);
    return new GlideDateTime().getLocalDate().getValue() ==
inputGdt.getLocalDate().getValue();
  },


  endingTomorrow: function(inputDate) {
    var inputGdt = new GlideDateTime();
    inputGdt.setDisplayValueInternal(inputDate);
    var tomorrowDateTime = new GlideDateTime();
    tomorrowDateTime.addDaysLocalTime(1);
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
   return tomorrowDateTime.getLocalDate().getValue() ==
inputGdt.getLocalDate().getValue();

 },
```

d. Edit the following part of the script as mentioned.

```
     if (isOnCall) {

       if (response.nearestShift.isCurrent) {

         if (response.nearestShift.roster_name)

           response.shiftText = gs.getMessage("You are
on-call until {0} as {1}",
[this.getFormattedDateTime(response.nearestShift[endAttr]),
response.nearestShift.roster_name]);

         else

           response.shiftText = gs.getMessage("You are
on-call until {0}",
[this.getFormattedDateTime(response.nearestShift[endAttr])
]);

       } else {

         if (response.nearestShift.roster_name)

           response.shiftText = gs.getMessage("Your
next on-call shift starts on {0} as {1}",
[this.getFormattedDateTime(response.nearestShift[startAttr
]), response.nearestShift.roster_name]);

         else

           response.shiftText = gs.getMessage("Your
next on-call shift starts on {0}",
[this.getFormattedDateTime(response.nearestShift[startAttr])
]);

       }

     } else {

       if (response.nearestShift.isCurrent) {

         response.shiftText = gs.getMessage("Your shift ends
on {0}",
[this.getFormattedDateTime(response.nearestShift[endAttr])
]);

       } else

         response.shiftText = gs.getMessage("Your next shift
starts on {0}",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 [this.getFormattedDateTime(response.nearestShift[startAttr])
]);

      }
```

6. Select **Update**.

Modify the first-level cards in the Service Operations Workspace for ITOM Overview section

Customize the data displayed in the first-level cards in the Overview section. Configure parameters such as the header label, data source, metric, group by field, and viewAllQuery for a tailored display.

Before you begin
Role required: evt_mgmt_admin

Procedure

1. Navigate to **sys_ux_client_script_include.list**.
2. From the UX Client Script Includes list, select SowEMLandingPageUtilsSNC client script include definition.
3. On the UX Client Script Include form, edit the **Script** field.
   For example:

```
  }
    };
   const labelMap = {
       'em_alert': {
           'alerts_assigned_to_you': {
               '1': 'Alerts assigned to you - Urgent',
               '2': 'Alerts assigned to you - High',
               '3': 'Alerts assigned to you - Moderate',
               '4': 'Alerts assigned to you - Low'
           },
           'open_alerts': {
               '1': 'Open alerts - Urgent',
               '2': 'Open alerts - High',
               '3': 'Open alerts - Moderate',
               '4': 'Open alerts - Low'
           },
           'team_alerts_assigned': 'Assigned team alerts - to
',
           'team_alerts_unassigned': {
               '1': 'Unassigned team alerts - Urgent',
               '2': 'Unassigned team alerts - High',
               '3': 'Unassigned team alerts - Moderate',
               '4': 'Unassigned team alerts - Low'
           }
       }
```

4. Click **Update**.

Create a copy of the Service Operations Workspace for ITOM landing page

Make visual styling changes to the sub-pages in the home page. This will enhance the overall user experience by providing a consistent and visually appealing interface, improving navigation, and ensuring that key information is easily accessible and engaging.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: evt_mgmt_admin or ui_builder_admin

About this task

A page is content that lives at a specific URL. When a user opens the URL, they will see a different variant of this page based on the audience, conditions, and order that you set.

Procedure

1. Navigate to **All** > **Now Experience Framework** > **UI Builder**.
2. From the **Display name** field, filter and search for **Service Operations Workspace**.

   The Experiences page opens.

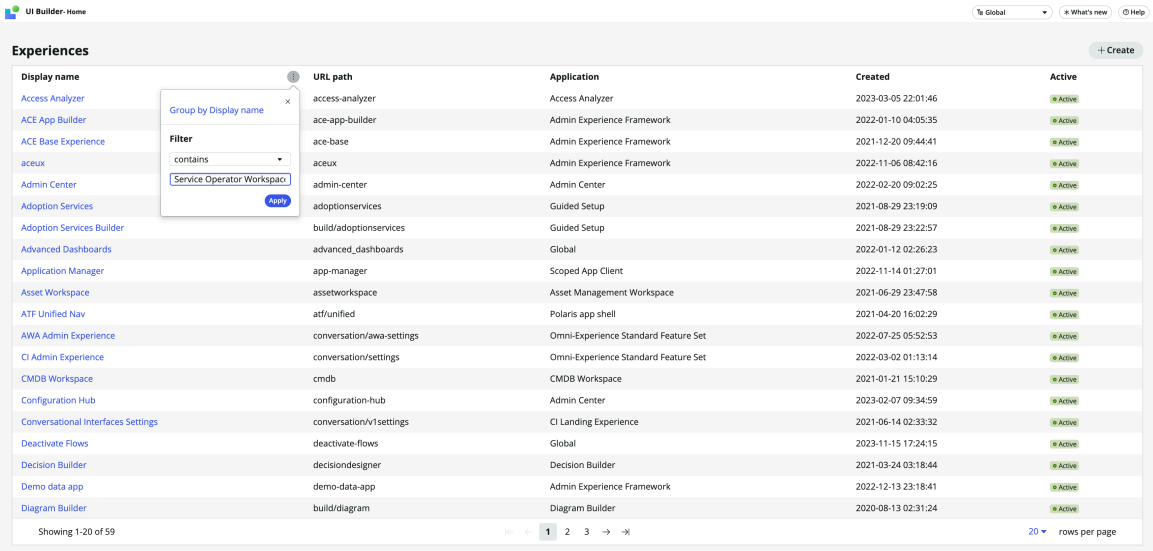

3. Open **Service Operations Workspace** by selecting it.
4. Next to **Pages and variants**, select the plus icon (⊕).

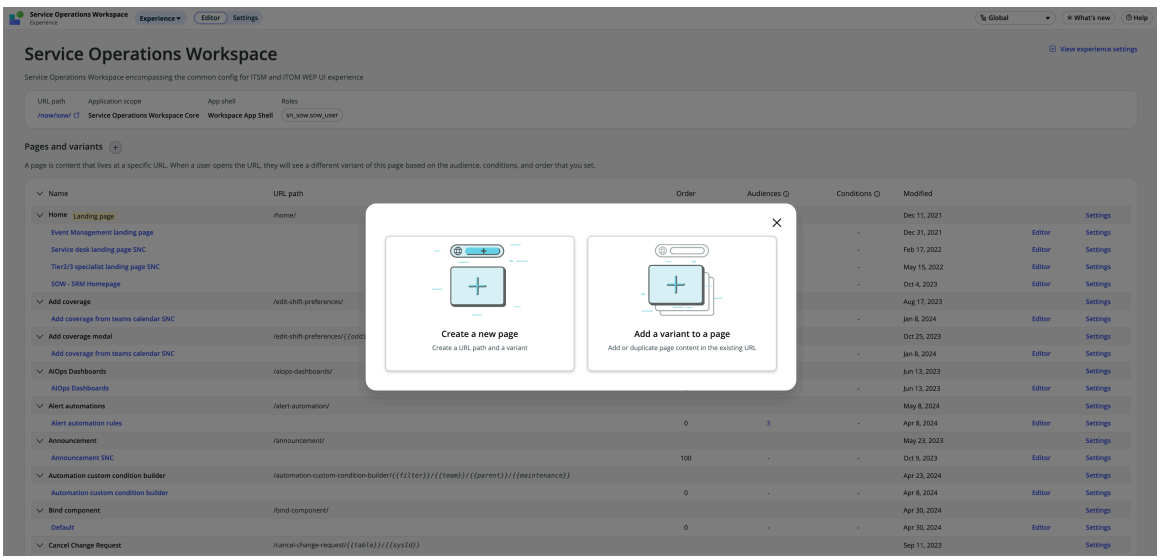

5. Select **Add a variant to a page**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Add variant to an existing page window opens.

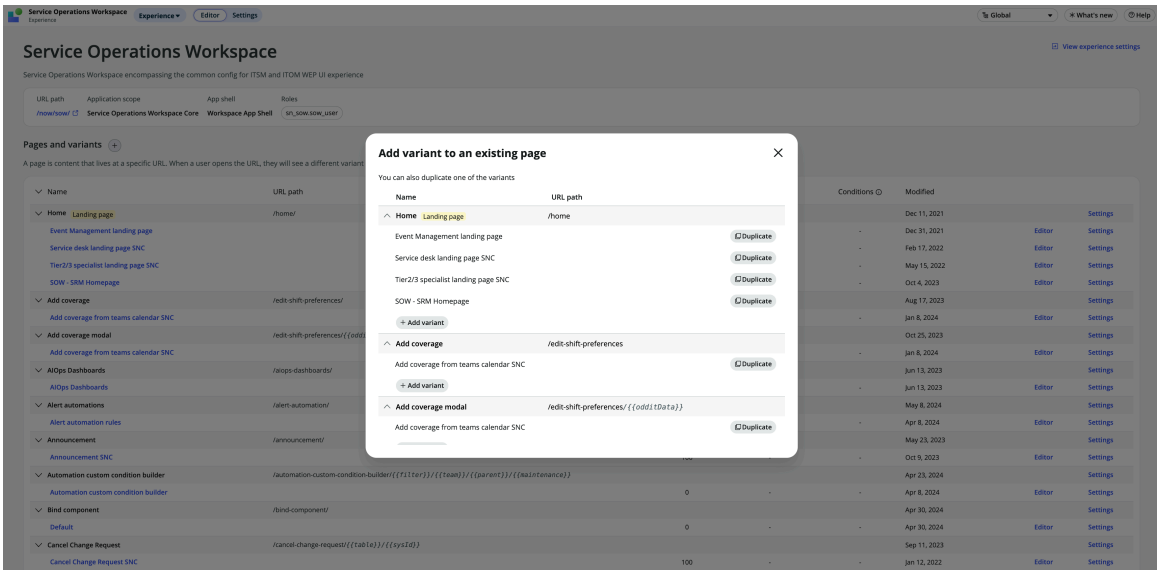

6. To create a duplicate page content in the existing URL path, select **Duplicate** and follow the steps.

Open your configurable workspace experience in UI Builder

Access your Configurable Workspace experience in UI Builder for editing.

Before you begin

Role required: ui_builder_admin or admin

Procedure

1. In the navigation filter, navigate to **All** > **Now Experience Framework** > **UI Builder**.
2. Select **Experiences**.
3. Select your experience.

Result

Your configurable experience opens in UI Builder.

Redirect to the Service Operations Workspace home page

Enable quick task prioritization by directing agents to the homepage immediately upon logging into the ServiceNow instance. This improves workflow efficiency by granting agents immediate access to essential tools and information, facilitating the timely resolution of critical tasks and issues.

Before you begin

Ensure that the value of the *`glide.login.home property`* is set to `/now/nav/ui/home`. This property specifies the default homepage that is displayed upon logging in.

Role required: evt_mgmt_admin

Procedure

1. Select **All** and then search for `sys_properties.list`.
2. Add the *`sn_sow.service_desk_groups`* property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the **Value** field, enter the sys_id of the user group to which the operator belongs.

   Note: If you want to specify multiple service desk groups, add a comma-separated list of group sys_ids in the value field.

4. Select **Update**.

Open your Configurable Workspace experience in UI Builder

Access your Configurable Workspace experience in UI Builder for editing.

Before you begin

Role required: ui_builder_admin or admin

Procedure

1. In the navigation filter, navigate to **All** > **Now Experience Framework** > **UI Builder**.
2. Select **Experiences**.
3. Select your experience.

Result

Your configurable experience opens in UI Builder.

Service Operations Workspace

The Service Operations Workspace monitors services within your Event Management environment, enabling you to review their status and identify those at risk of suboptimal performance. This empowers you to promptly address any issues.

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

You can configure the Service Operations Workspace to remain logged in for up to 24 hours. For more information, see Manage user sessions.

View application service health information in Service Operations Workspace

Monitor the health of services and filter the information that you want to view, allowing you to focus on the most relevant data and make informed decisions quickly.

Before you begin

Role required: evt_mgmt_user or evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the left navigation bar, select the Service Dashboard icon: .
   The Service Dashboard page appears.

   Note: The application services that appear in the Service Dashboard are those added to the Impact Filter Services list. For more information, see Add application services for impact calculation.

3. Perform the following actions as required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Features and actions

Features	Actions
Make the tile size larger or smaller.	**a.** Select the Display settings icon ( ). **b.** From the **Display settings** pop-up window, select the size. 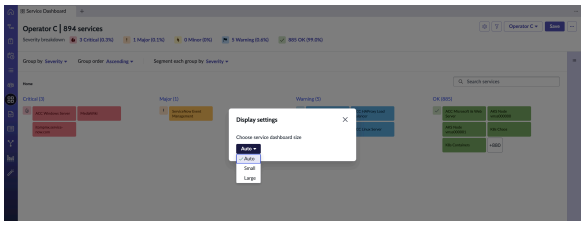
Group services by selected criteria	Select the arrow next to **Group By** and then select the criteria by which you want to group services. 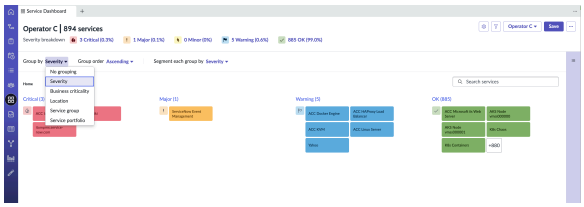
Group order	Select the arrow next to **Group order** and then select the criteria by which you want to group services. 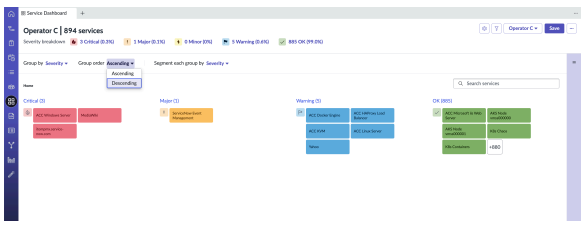
Segment each group by	Segment each group by severity, business criticality, and cost to categorize and organize them into distinct groups based on these attributes. 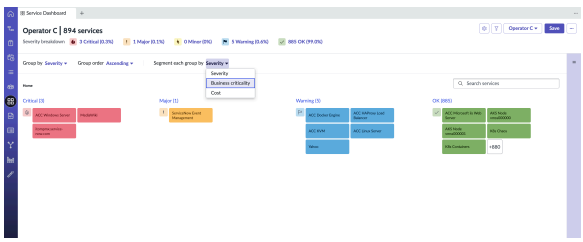
Filter information	Filter the information displayed on the dashboard ensuring you focus on relevant data for effective monitoring.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Features	Actions
	○ To set a global filter, select the Filter icon ( ) and set the global filter condition, add filter conditions, and then select **Apply**. 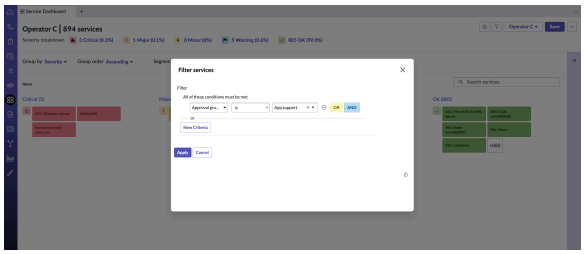 ○ To exclude a severity from the display, select it in the **Severity Breakdown** ribbon. Filtered-out severities will appear dimmed. To restore a dimmed severity to the display, simply select it again. 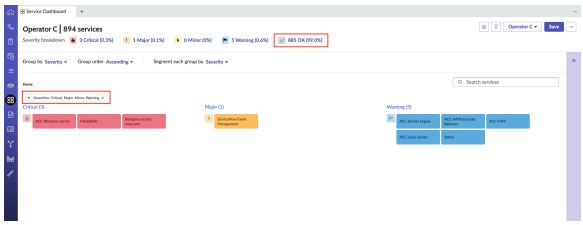

Create a user view in Service Operations Workspace

Create and customize a user view in the Service Operations Workspace to enable viewing information in the format you want.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the left navigation bar, select the Service Dashboard icon: .
   The Service Dashboard page appears.

   Note: The application services that appear in the Service Dashboard are those added to the Impact Filter Services list. For more information, see Add application services for impact calculation.

3. Customize the dashboard appearance by selecting output criteria from the **Group by**, **Group order**, or **Segment each group by** filters, or by selecting the relevant **Severity breakdown**. Selecting **Service group** > **All** in the **Group by** filter displays all service groups.
4. Configure the information for the view to display by clicking the filter services icon ( ) and setting a global filter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, you can configure a filter to display only those services assigned to a specific user.

Note: To discard your changes and revert to the view's most recently saved state, select the More actions icon (•••) and select **Discard changes**.

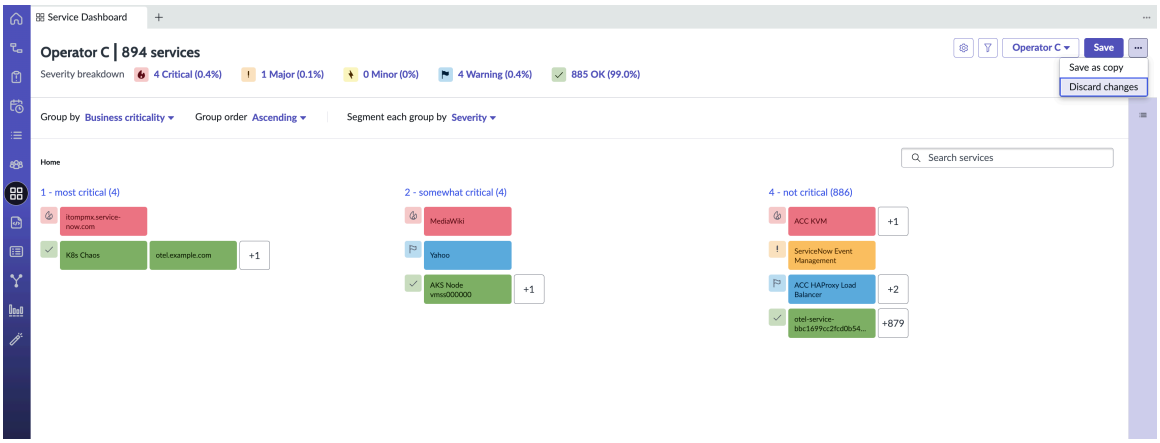

5. Save your new view by clicking the More actions icon (•••) next to the **Save** button and select **Save as copy**.

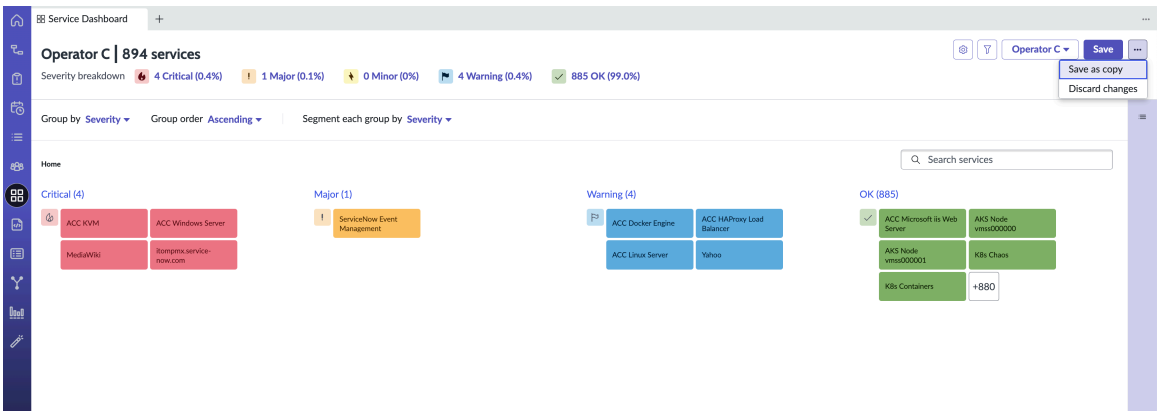

6. In the Save as copy dialog box, enter the name of your view in the **View name** field.

   The view is saved as a **Private view** (visible only to the logged-in user) and is available for use.

   Views are grouped in the following categories:

   - **Global views**: Visible to all users. Only an Event Management administrator can create a global view.
   - **Private views**: Visible only to the logged-in user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can switch between different views by selecting a view from the list next to the **Save** button.

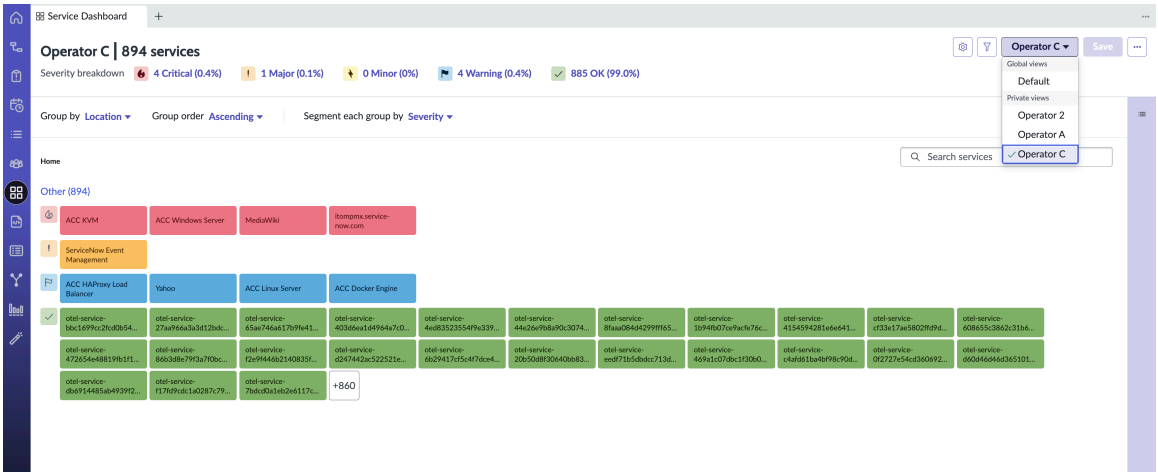

Delete a user view in Service Operations Workspace

Delete a Service Operations Workspace user view you no longer need. You perform this action from the Event Management menu, outside of Service Operations Workspace.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. In your Event Management instance, in the navigation pane, search for *`noc_user_view.list`*.
   The Operator Workspace Views page opens.

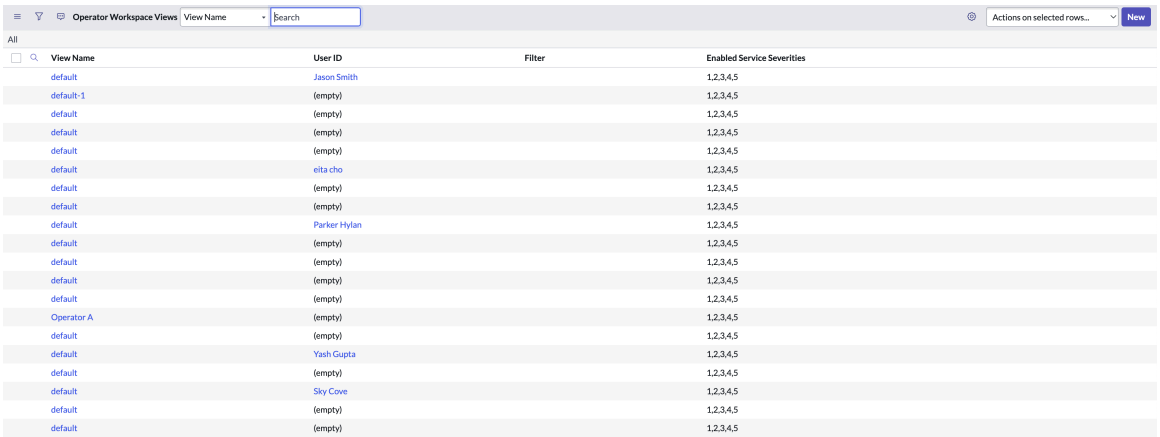

View Name	User ID	Filter	Enabled Service Severities
default	Jason Smith		1,2,3,4,5
default-1	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	eita cho		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	Parker Hylan		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
Operator A	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	Yash Gupta		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	Sky Cove		1,2,3,4,5
default	(empty)		1,2,3,4,5
default	(empty)		1,2,3,4,5

2. Select the check box next to the view you want to delete.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the **Actions on selected rows** list, select **Delete**.

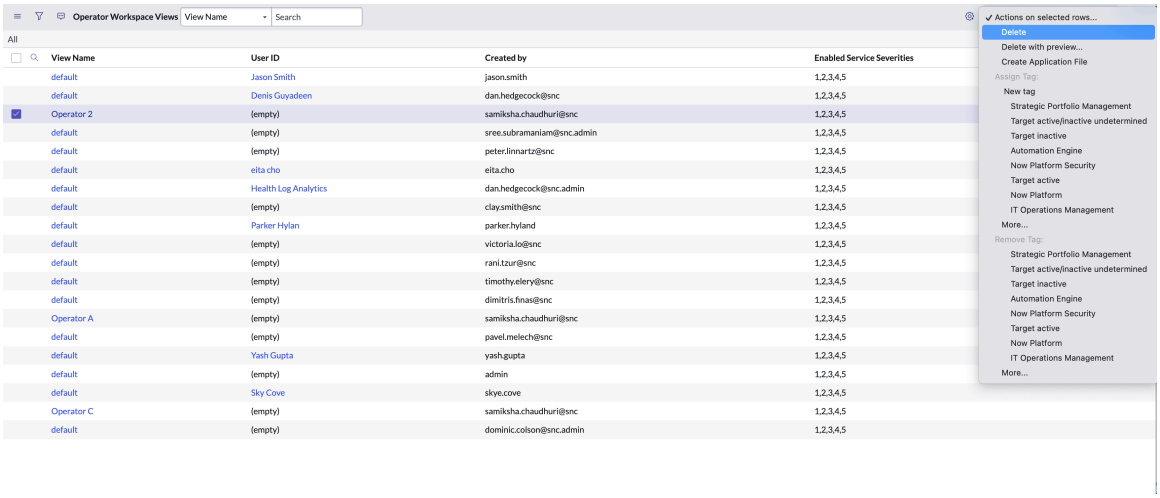

View application service details in Service Operations Workspace

View comprehensive information about an application service, including its name, business criticality, ownership, process and operational status, traffic-based discovery details, discovery status, and relevant comments. This enables thorough analysis of the application service's current state and usage, provides clarity on ownership and criticality, and facilitates effective prioritization of resources and actions.

Before you begin

Role required: evt_mgmt_user, evt_mgmt_operator, or evt_mgmt_admin

About this task

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the left navigation bar, select the Service Dashboard icon: .
   The Service Dashboard page appears.

   Note: The application services that appear in the Service Dashboard are those added to the Impact Filter Services list. For more information, see Add application services for impact calculation.

3. Select a service tile.
   A pop-up window displays the business criticality and severity of a service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

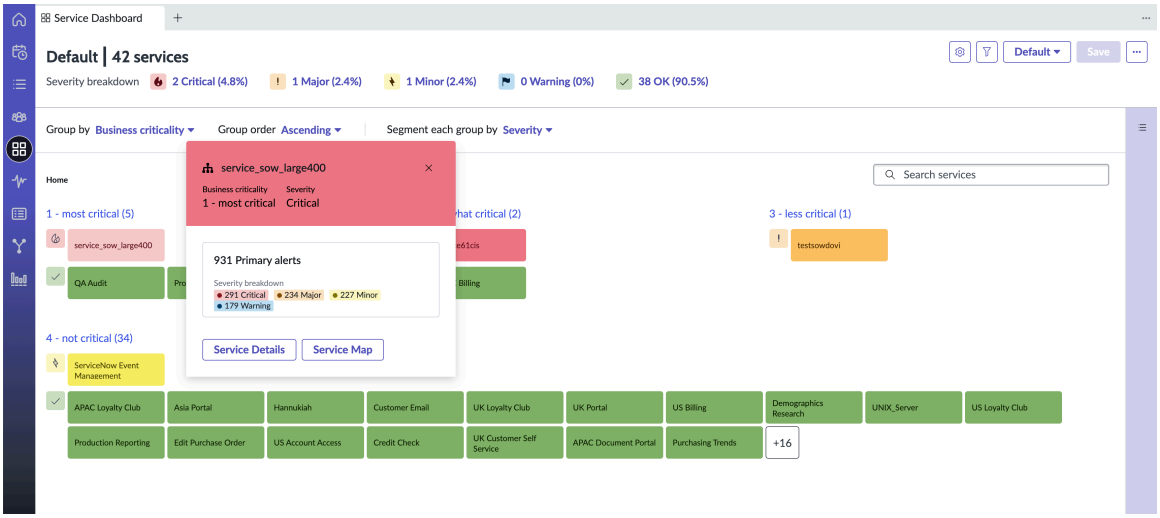

4. Select **Service Details**.
   By default, the **Details** tab is selected.

   The details page provides comprehensive information about an application service.

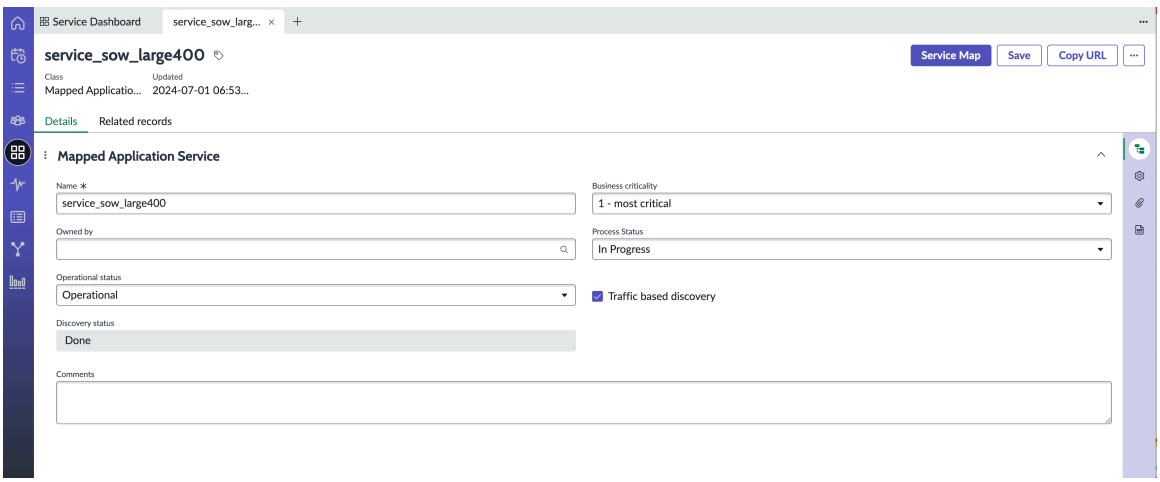

   For information on the fields, see Application service details page.

View primary alerts associated with a service

View the primary alerts associated with a service to help troubleshoot issues with the service.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the left navigation bar, select the Service Dashboard icon: .
   The Service Dashboard page appears.

> **Note:** The application services that appear in the Service Dashboard are those added to the Impact Filter Services list. For more information, see Add application services for impact calculation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select a service tile.
   A pop-up window displays the business criticality and severity of a service.

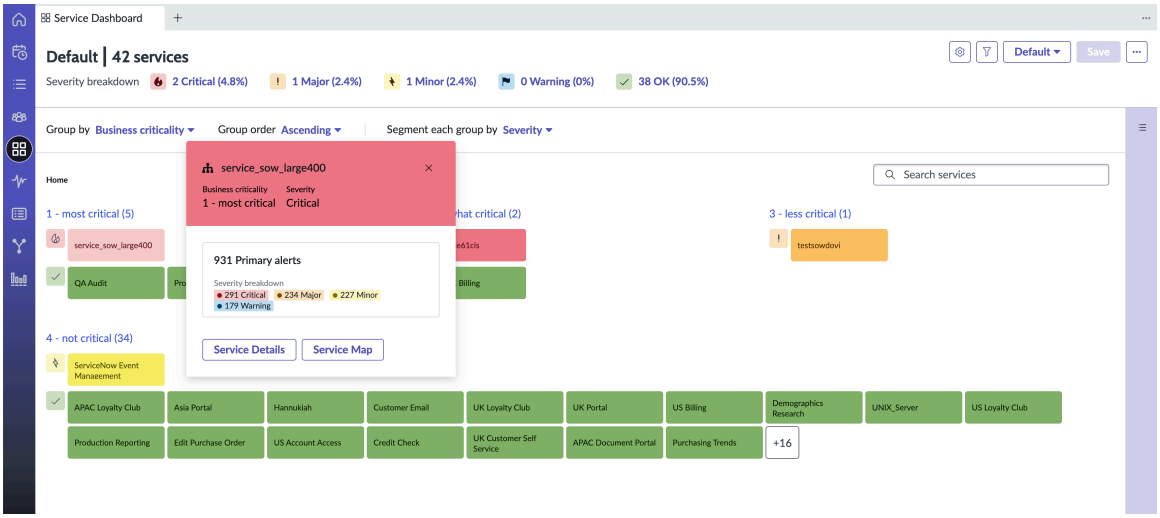

4. Select **Service Details**.
   By default, the **Details** tab is selected.
5. Select the **Related records** tab.
6. From the left pane, select **Related Alerts**.

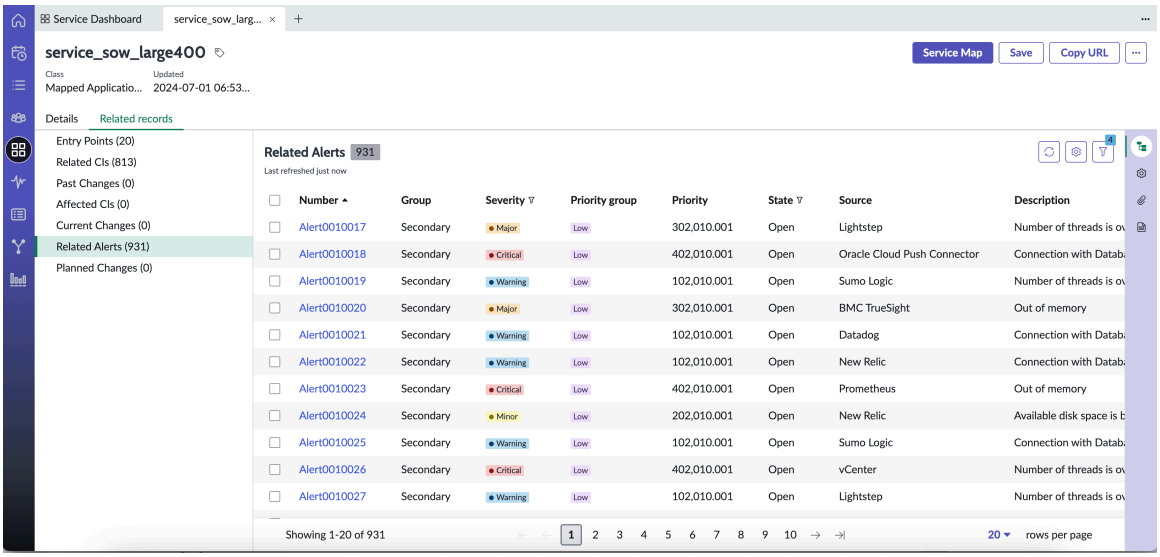

View details about an alert by selecting its number in the **Number** column.

Create a list in Service Operations Workspace for ITOM

Customize lists in Service Operations Workspace by creating new lists, changing their order, modifying their content, or deleting lists as needed. This flexibility allows operators to organize information according to their specific requirements, facilitating easier navigation, quicker access to relevant data, and streamlined workflow management.

Before you begin

Role required: itil

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Create custom groupings of records listed under the **Lists** tab, such as incidents tied to a specific company. Customize existing lists or create new ones exclusively visible under the **My Lists** tab for personalized organization.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the List icon: .
3. Select the **My Lists** tab.
4. Select **Add new list**.

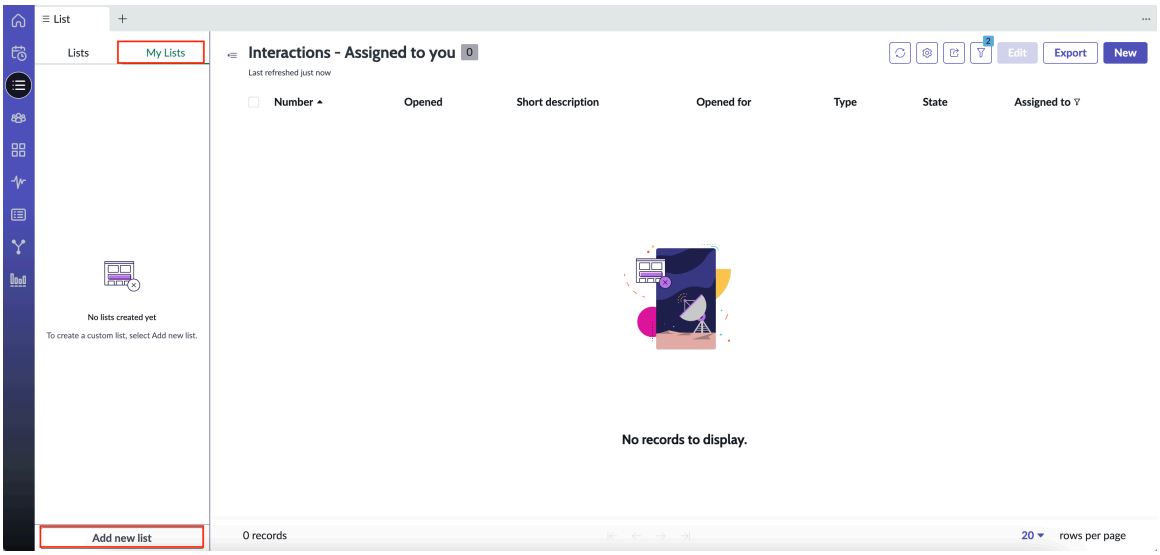

5. Create a list by either basing it on an existing list or creating a new one.
   - To base your list on an existing list, select **Start from existing**.
   - To begin a new list, select **Create your own.**
6. On the form, fill in the fields.
   For more information on the fields, see Create list fields.
7. Select **Create**.
   The list appears in the **My Lists** tab.

Rename a list in Service Operations Workspace for ITOM

Rename a list to better reflect its content or purpose. This customization allows operators to align list names with their workflow terminology, improving clarity and usability.

Before you begin

Role required: itil

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the List icon: .
3. Select the **My Lists** tab.
4. Select the list you wish to rename.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Select the settings icon ( ) and then select **Rename**.

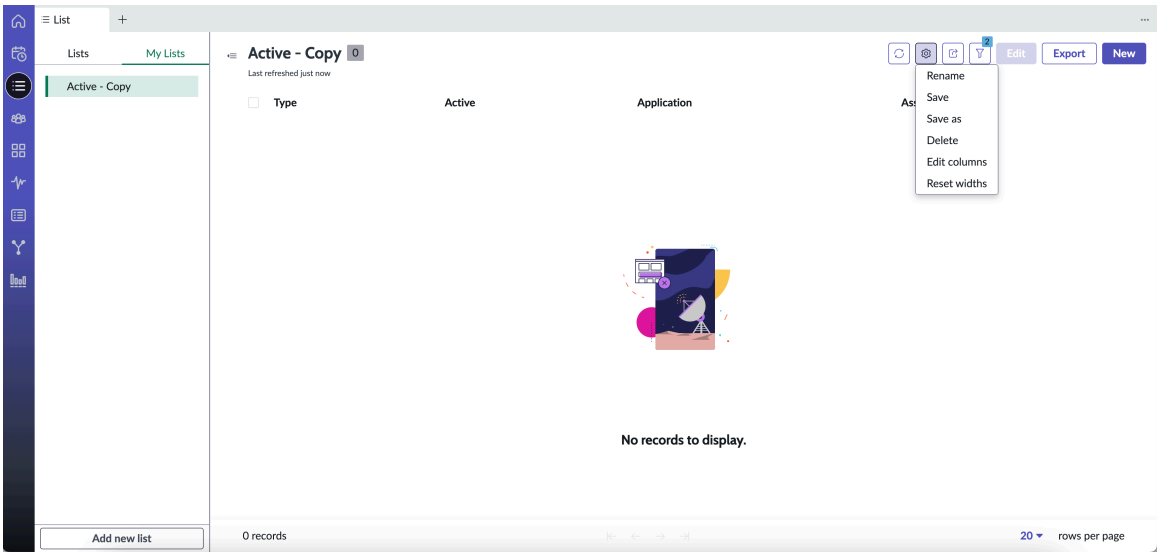

6. In the **Rename** pop-up window, in the **Title** field, rename the name of the list and select **Save**. The list is renamed.

Reorder list in Service Operations Workspace for ITOM

Reorder lists to prioritize information and enhance workflow organization.

Before you begin

Role required: itil

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the List icon: .
3. Select the **My Lists** tab.

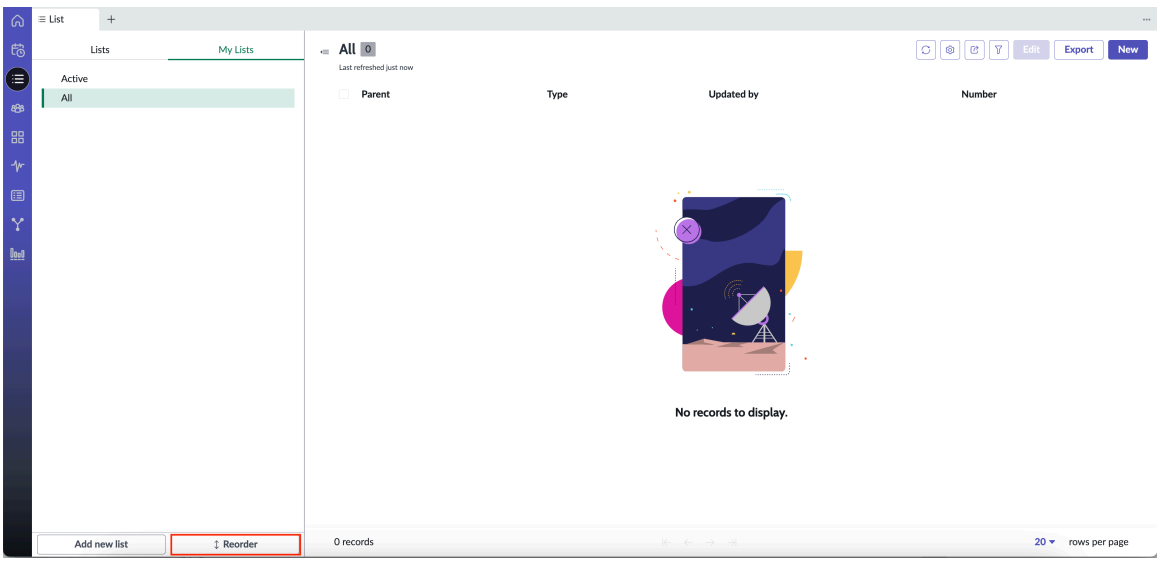

4. Select the list you want to reorder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Select **Reorder** to adjust the position of the list.

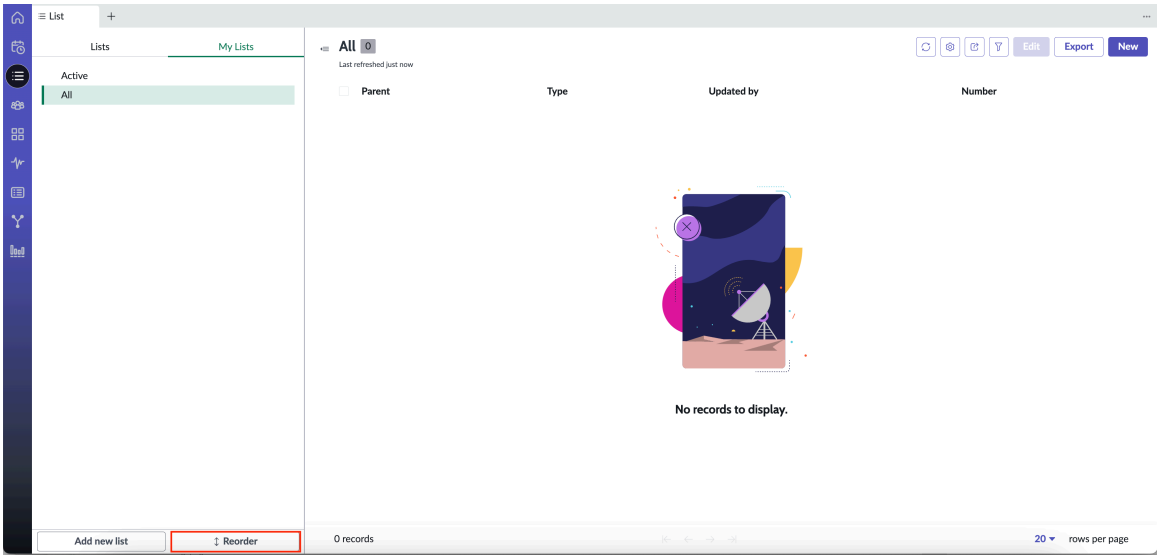

Delete a list in Service Operations Workspace for ITOM

Delete a list to efficiently manage data sets, declutter your workspace, and maintain organized information.

Before you begin

Role required: itil

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the List icon: .
3. Select the **My Lists** tab.
4. Select the list you wish to delete.
5. Select the settings icon ( ) and then select **Delete**.

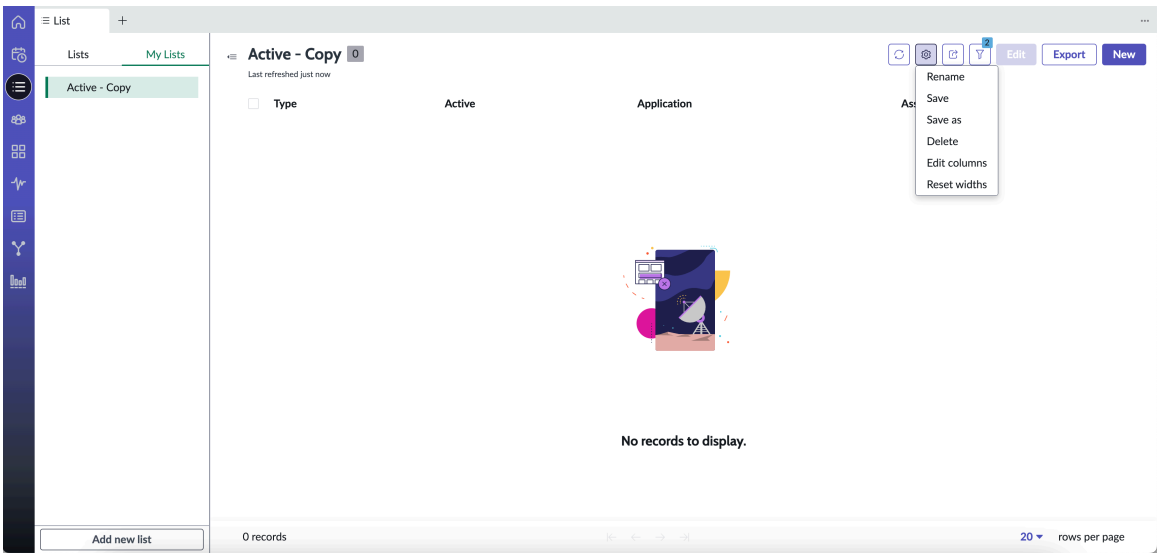

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. In the confirmation pop-up window, select **Delete**.
   The list is delete.

View unified service map and the impact paths in Service Operations Workspace

Visualize relationships between Configuration Items (CIs) and alerts with real-time updates and detailed impact paths. Enhance troubleshooting and proactive management by quickly identifying root causes and dependencies for both discovered services and application services.

Before you begin

Role required: evt_mgmt_user, evt_mgmt_operator, or evt_mgmt_admin

About this task

The enhanced service map in Service Operations Workspace displays the impacted path of an alert to enable you to quickly assess its impact on a service. You can access the enhanced service map directly from the Service Dashboard, or from the preview panel or from an impacted service in the Express List.The service map has several modes of display:

- For a service with up to 60 CIs, regardless of alert status, the map shows all CIs.
- For a service with more than 60 CIs without alerts, the map shows up to 60 CIs.
- For a service with more than 60 CIs with alerts, the map shows the impacted path of the most recent alert with the highest severity.

You can also investigate a wider view of the service topology on the service map for up to 250 CIs regardless of alerts. For a broader view, select **Clear** and choose a mode of either up to 60 or up to 250 nodes.

Note: Admins can continue using the legacy service map to manage the propagation of severity. Access the map by navigating to **Services** > **Application Services**.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the left navigation bar, select the Service Dashboard icon: .
   The Service Dashboard page appears.

   Note: The application services that appear in the Service Dashboard are those added to the Impact Filter Services list. For more information, see Add application services for impact calculation.

3. Select a service tile.
   A pop-up window displays the business criticality and severity of a service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

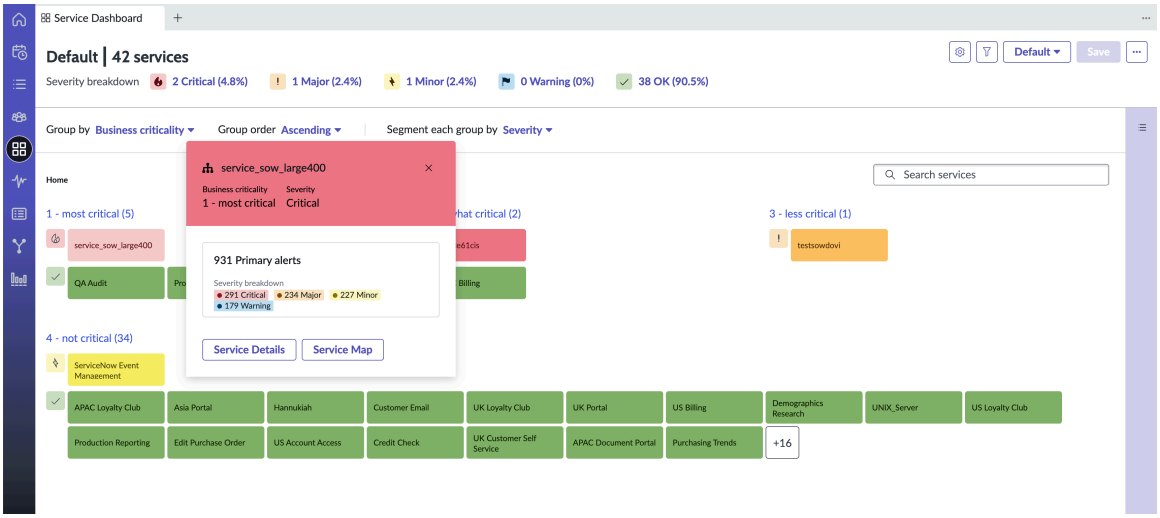

4. Select **Service Map**.

The application map along with the CI relationships appears.

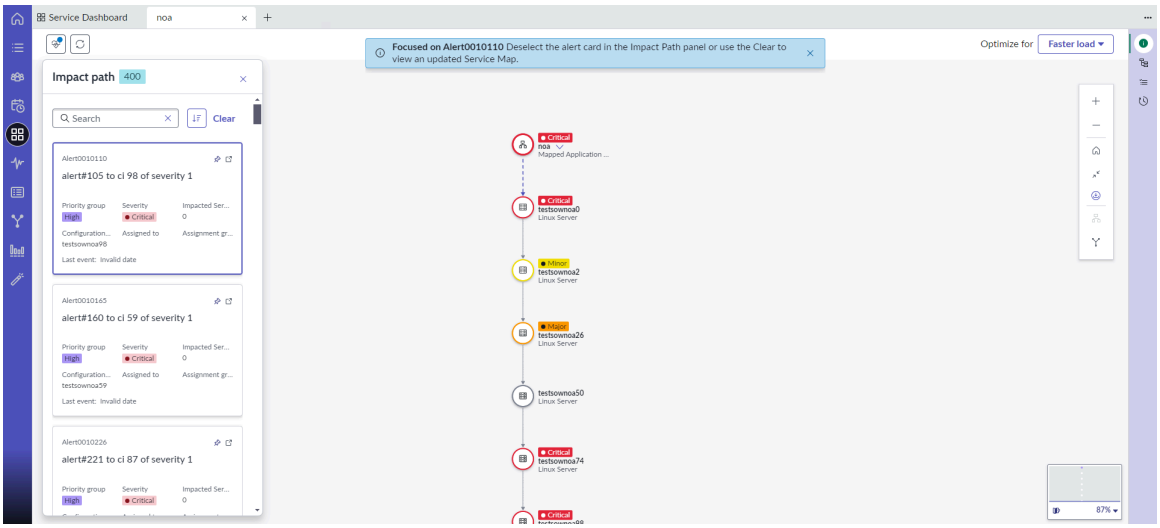

Note: Depending on the size of the service, the unified map either shows the full map (if it contains up to 60 CIs) or an Impact Path leading to the most critical alert (for maps with over 60 CIs). For large maps, you can reveal more CIs by clearing the existing selecting and navigating to the **Showing more nodes** optimization option.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

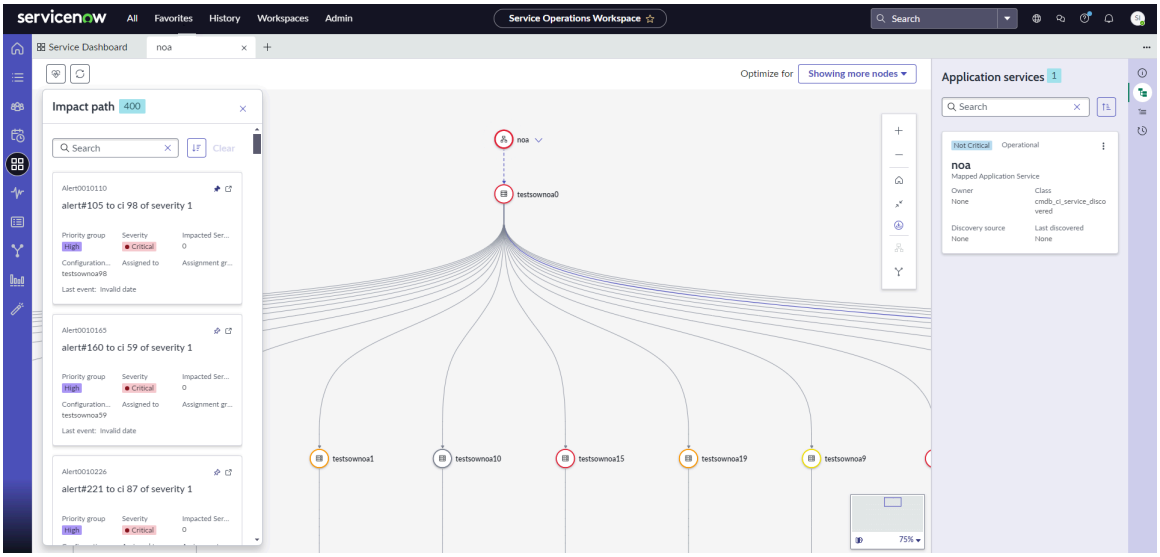

If the impact path panel is closed, you can open it by selecting the Impact path icon ( ).

5. Navigate between alerts by selecting an alert card.

   Alternatively, use the search function to find an alert by description, CI, or alert number and focus on its impact path.

   Selecting any alert tile displays the unified map for that alert. To return to the full map for the service, select **Clear**.

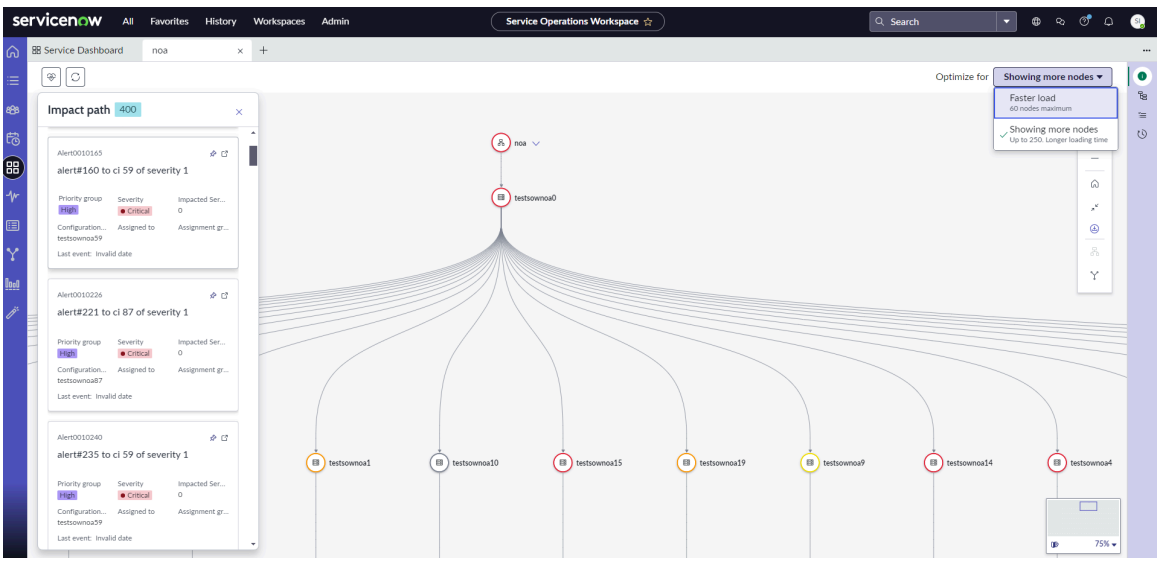

6. To view alert details, select the Open Record Form icon ( ) on the alert card.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The alert details page opens.

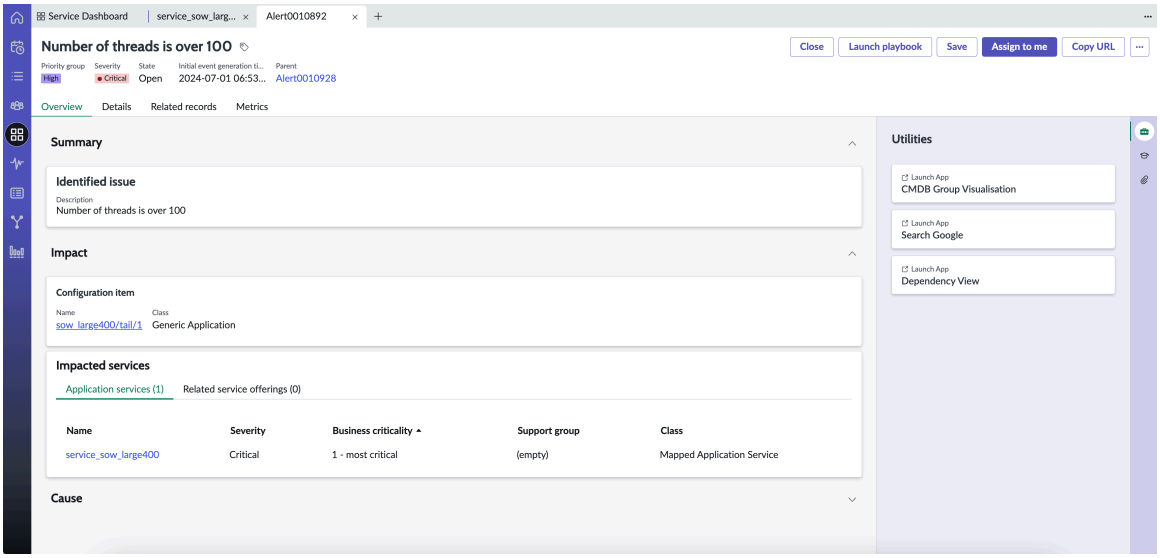

Configure live updates for alert lists

Choose live updates for alert lists to receive real-time updates, or disable live updates and refresh the alert list manually.

Before you begin

Role required: evt_mgmt_user, evt_mgmt_operator, or evt_mgmt_admin

About this task

By default, the toggle for live updates is off, and you must enable it to receive live alert updates. Once enabled, the alert list automatically updates whenever an action is performed on an alert. Live updates enhance the accuracy and timeliness of the information displayed, enabling quicker response times, and improving overall system reliability and performance.

To update the alert list manually, you can switch the live updates toggle to off. Manual updates are useful when many alerts are being processed continuously, to avoid excessive refreshing.

Note: Set Read Role permissions to the Event Management user (evt_mgmt_user). Otherwise, the property won't be available.

If an alert is updated and, as a result, no longer matches the active filter, Live List grays it out.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the List icon: .
3. Select the **Default lists** tab.
4. Under **Alerts**, select the list where you want to enable or disable live updates.

   Note: Live updates aren't available for the **All Alert** list.

5. At the top of the alert list, select the toggle switch to turn live updates on or off.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- When live updates are turned on, the setting is saved for future logins by the same user.
- When live updates are off, the refresh button displays a counter to show the number of changes since the last refresh. You can select the refresh button anytime to update the list manually.

Configure alert metrics

Configure the metrics you want to be visible for Service Operations Workspace alerts or use the metrics configured with the base system to focus on relevant data for faster troubleshooting.

Before you begin

Ensure you have the Metric Intelligence (sn_metrics_ng) plugin installed.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Metrics View Configuration**.

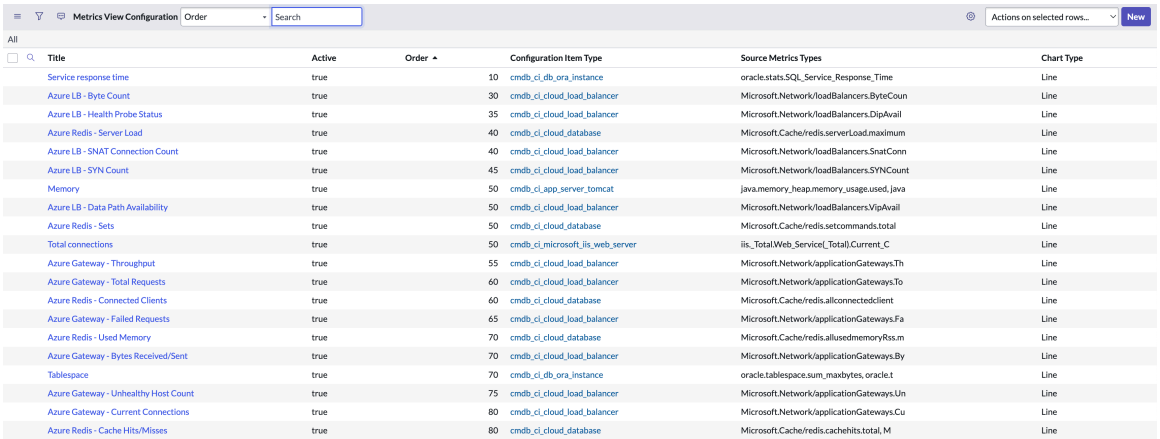

Title	Active	Order	Configuration Item Type	Source Metrics Types	Chart Type
Service response time	true	10	cmdb_ci_db_ora_instance	oracle.stats.SQL_Service_Response_Time	Line
Azure LB - Byte Count	true	30	cmdb_ci_cloud_load_balancer	Microsoft.Network/loadBalancers.ByteCoun	Line
Azure LB - Health Probe Status	true	35	cmdb_ci_cloud_load_balancer	Microsoft.Network/loadBalancers.DipAvail	Line
Azure Redis - Server Load	true	40	cmdb_ci_cloud_database	Microsoft.Cache/redis.serverLoad.maximum	Line
Azure LB - SNAT Connection Count	true	40	cmdb_ci_cloud_load_balancer	Microsoft.Network/loadBalancers.SnatConn	Line
Azure LB - SYN Count	true	45	cmdb_ci_cloud_load_balancer	Microsoft.Network/loadBalancers.SYNCount	Line
Memory	true	50	cmdb_ci_app_server_tomcat	java.memory_heap.memory_usage.used, java	Line
Azure LB - Data Path Availability	true	50	cmdb_ci_cloud_load_balancer	Microsoft.Network/loadBalancers.VipAvail	Line
Azure Redis - Sets	true	50	cmdb_ci_cloud_database	Microsoft.Cache/redis.setcommands.total	Line
Total connections	true	50	cmdb_ci_microsoft_iis_web_server	iis_Total:Web_Service(_Total).Current_C	Line
Azure Gateway - Throughput	true	55	cmdb_ci_cloud_load_balancer	Microsoft.Network/applicationGateways.Th	Line
Azure Gateway - Total Requests	true	60	cmdb_ci_cloud_load_balancer	Microsoft.Network/applicationGateways.To	Line
Azure Redis - Connected Clients	true	60	cmdb_ci_cloud_database	Microsoft.Cache/redis.allconnectedclient	Line
Azure Gateway - Failed Requests	true	65	cmdb_ci_cloud_load_balancer	Microsoft.Network/applicationGateways.Fa	Line
Azure Redis - Used Memory	true	70	cmdb_ci_cloud_database	Microsoft.Cache/redis.allusedmemoryRss.m	Line
Azure Gateway - Bytes Received/Sent	true	70	cmdb_ci_cloud_load_balancer	Microsoft.Network/applicationGateways.By	Line
Tablespace	true	70	cmdb_ci_db_ora_instance	oracle.tablespace.sum_maxbytes, oracle.t	Line
Azure Gateway - Unhealthy Host Count	true	75	cmdb_ci_cloud_load_balancer	Microsoft.Network/applicationGateways.Un	Line
Azure Gateway - Current Connections	true	80	cmdb_ci_cloud_load_balancer	Microsoft.Network/applicationGateways.Cu	Line
Azure Redis - Cache Hits/Misses	true	80	cmdb_ci_cloud_database	Microsoft.Cache/redis.cachehits.total, M	Line

2. On the Metrics View Configuration page, select **New**.
3. On the form, fill in the fields.
   For information on the fields, see Metrics View Configuration form fields.
4. Select **Submit**.

View CI health in Metric Explorer

View the health and performance of an alert's CI using Metric Explorer, to help troubleshoot problematic areas in the CI.

Before you begin

All plugins related to Event Management, Metric Intelligence, and Metric Explorer, are required.

Roles required:

- To access Metric Intelligence metrics in Metric Explorer, the workspace_agent and itil roles are required.
- To configure the set of metrics that appears in Metric Explorer for a class, the admin or evt_mgmt_admin role is required.

About this task

Metric Explorer displays the following Metric Intelligence data:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Charts, control bounds, and aggregations for key metrics for Configuration Items (CIs), in a single location.
- Aggregation of metrics for CIs with resources, in a single consolidated time series.
- Ability for administrators to configure which metric types appear for each CI class.

Metrics are displayed only for alerts containing a CI.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the List icon: .
3. Select **Lists** > **Alerts** > **All Alerts**.
   The Alerts - All Alerts page opens.
4. Select the alert number to open it.
   The alert opens in a separate tab, displaying the **Overview** tab.
5. Select the **Metrics** tab.
6. Select the **Metric explorer** button.
   The **Metric explorer** interface opens.

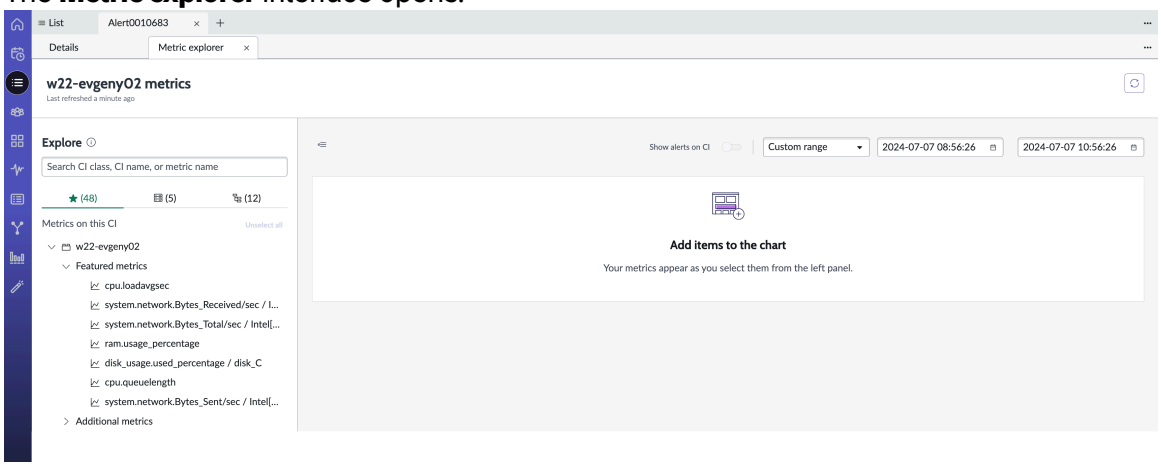

> **Note:** You can also access the **Metrics** subtab and **Metric Explorer** button when selecting a CI.

7. Select metrics in the Explore section to display the metrics in graph format in the **Custom chart** section.
   Metrics are divided according to the following subtabs:

Explore section subtabs

Subtab	Description
★	Displays metrics on the current CI.
	Displays metrics on other CIs on the same host.
	Displays metrics on entry points that relate to the service that the CI is impacting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subtab	Description
	Note: When working with large quantities of entry points (100K or more), collecting entry point metrics may cause system degradation. To prevent this, create the *sn_disable_selection_entry_point* property (**All** > **System Properties** > **All Properties**) and set it to **true** to turn off metric collection for entry points.

8. In the **Custom chart** section, configure a time range for which you want to view the displayed metrics.
9. Label alerts on the Custom chart by turning on the **Show alerts on CI** option.
   An alert icon ( ) appears at the top of the graph at the point where the alert occurred.

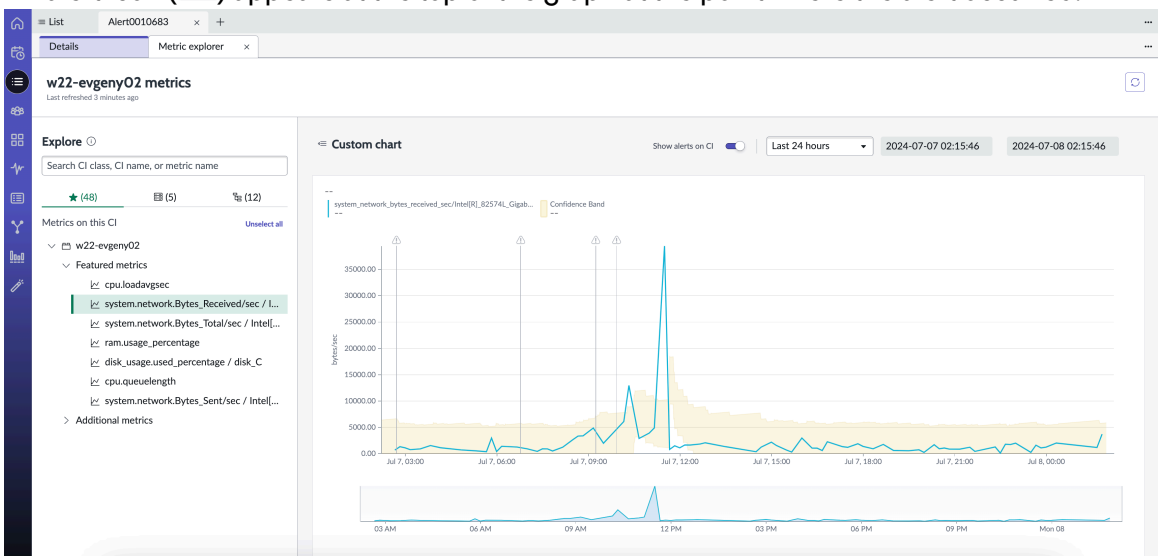

View alert metric trends in Service Operations Workspace

View the metrics for an alert and adjust the time range to see the trends for the alert or its associated Configuration Item (CI).

Before you begin

Ensure you have the Metric Intelligence (sn_metrics_ng) plugin installed.

Configure the metrics to be displayed for the alert, as described in Configure alert metrics.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the List icon: .
3. Select **Lists** > **Alerts** > **All Alerts**.
   The Alerts - All Alerts page opens.
4. Select the alert number to open it.
   The alert opens in a separate tab, displaying the **Overview** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Select the **Metrics** tab.

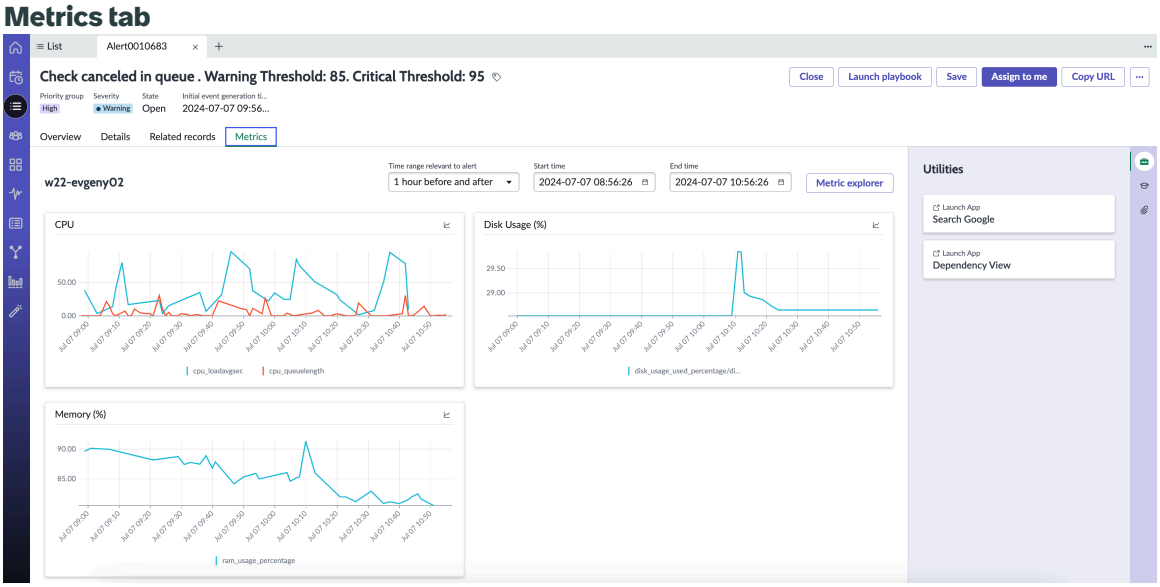

6. In the fields above the metric charts, configure the time range for the alert's metrics.

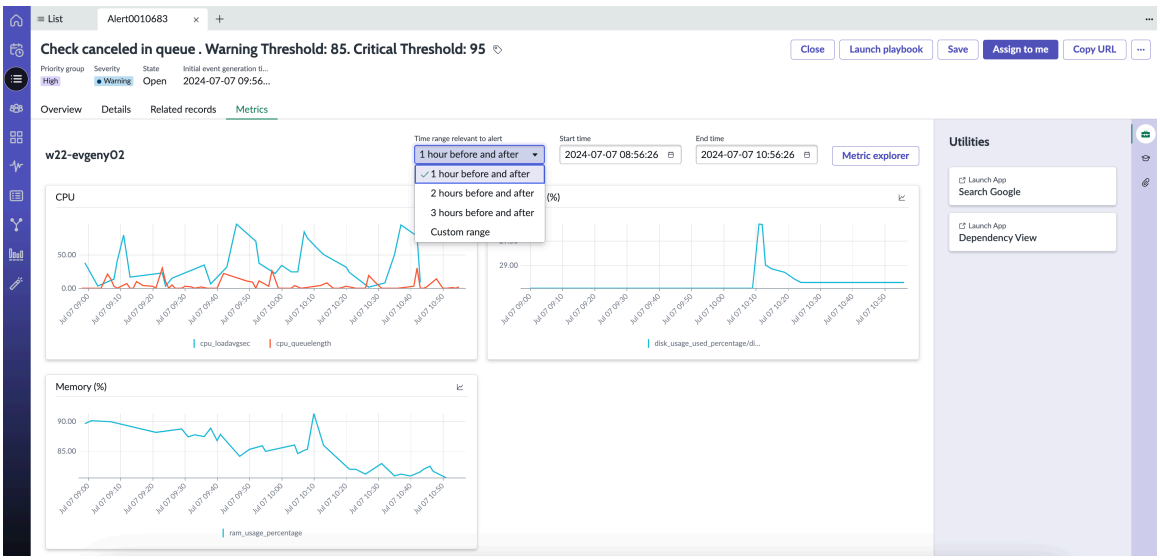

To customize the metrics display, configure the following properties on the **System Properties** > **All Properties** page:

- *ws_metric.widgets_number_limit*: Specify the number of widgets/records that can be displayed on the metrics tab.
- *ws_metric.datasets_per_widget_limit*: Specify the number of charts/datasets that can be displayed on a chart.

Field name	Description
Time range relevant to alert	Duration for calculating metrics around the alert.
Start time	Start date and time for metric calculations.
End time	End date and time for metric calculations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Select a parameter below a chart to hide the specified parameter's data on the chart.
8. **Optional:** Select the **Insights Explorer** button to view metrics by dragging and dropping them onto a timeline.
9. Select **Save**.
10. **Optional:** View a time range for an alert's CI.

    a. Select the **Overview** tab.

    b. Under the Impact section, on the **Configuration item** card, select the link to the CI that appears under the **Name** field.

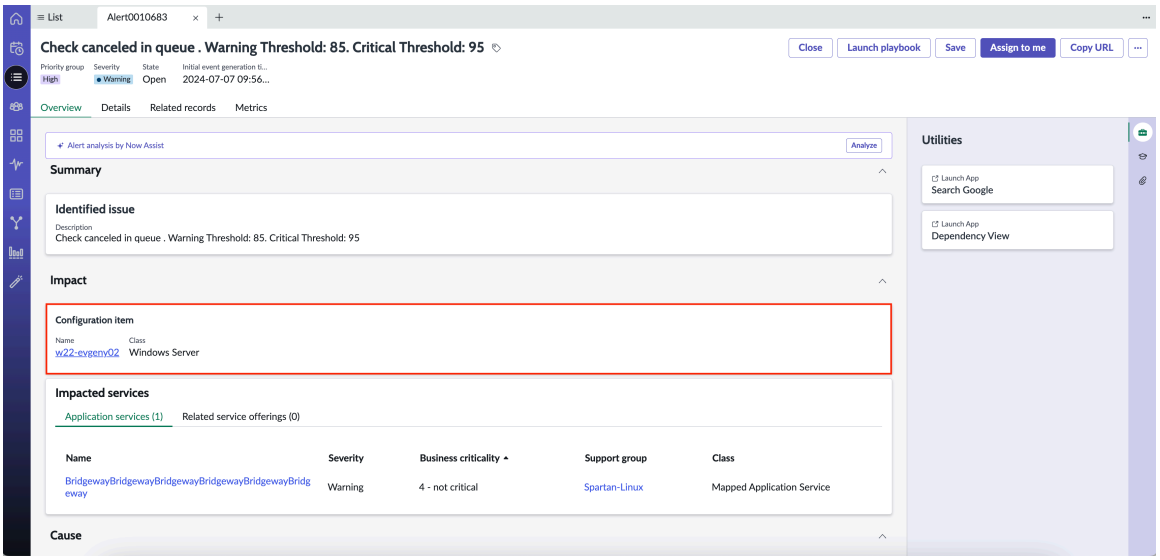

    c. Select the **Metrics** tab.

    d. Select **Save**.

On-call scheduling in Service Operations Workspace for ITOM

ITOM Service Operations Workspace supports creating and managing your on-call shifts and schedules. On-call schedule determines who is on-call at a given time. Learn how to configure on-call calendars and escalation policies to ensure timely responses to critical issues, reducing downtime and maintaining service continuity. Streamlining these processes enhances team readiness, minimizes disruptions, and improves service reliability.

Getting started with on-call scheduling

To enable on-call scheduling for alerts, first incorporate the Alert table [em_alert] into the Trigger Rule Table Config list. For more information, see Add alert table to enable on-call scheduling for Service Operations Workspace for ITOM.

Teams in Service Operations Workspace

Use the Teams menu ( ) in Service Operations Workspace to view all teams. Select a team card to view and manage a team's schedule, define escalation policies, and team preferences for the selected shift. This feature verifies that your team knows who is on call, is accountable, and is confident that critical alerts or incidents are acknowledged. For information on creating a shift, see Create on-call shift for your team .

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add alert table to enable on-call scheduling for Service Operations Workspace for ITOM

Enable on-call scheduling by incorporating the Alert table [em_alert] into the Trigger Rule Table Config list. This allows for creating shifts and implementing escalation triggers and policies, ensuring teams can promptly respond to alerts and maintain operational continuity.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** and in the search field, type `trigger_rule_table_cfg.list`.
   The Trigger Rule Table Configs page opens.
2. On the form, fill in the fields with the following values.

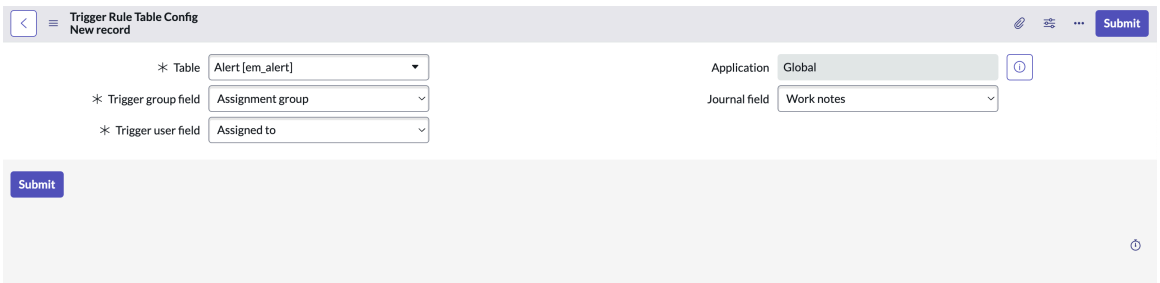

   - In the **Table** field, select **Alert [em_alert]**.
   - In the **Trigger group** field, select **Assignment group**.
   - In the **Trigger user** field, select **Assigned to**.
   - In the **Journal** field, select **Work notes**.
3. Select **Submit**.

What to do next

Create on-call shift for your team.

AIOps Dashboards in Service Operations Workspace for ITOM

The AIOps dashboards offer comprehensive visualizations of critical operational data. You can monitor key performance indicators (KPIs), track service health, and gain valuable context on the overall IT environment. These dashboards provide actionable insights that facilitate proactive responses to potential issues.

Key benefits of using the AIOps dashboards:

- Configure the dashboards to focus on the metrics and data most relevant to your role.
- Enjoy interactive visual representations of critical operational data.
- Monitor KPIs, service health metrics, and other important performance data at a glance.
- Leverage actionable insights derived from data analysis for better decision-making.

Note: Opening a set of dashboards shows the most recently accessed dashboard in that category.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Access the AIOps Operational dashboard

To open the dashboard, navigate to **Workspaces** > **Service Operations Workspace**. From the bottom of the navigation pane, select the AIOps configuration center icon .The AIOps configuration center page appears. On the ITOM AIOps configuration center page, under the **Optimize** > **Dashboards** section, select **View AIOps Operational dashboard**.

AIOps Dashboard home page

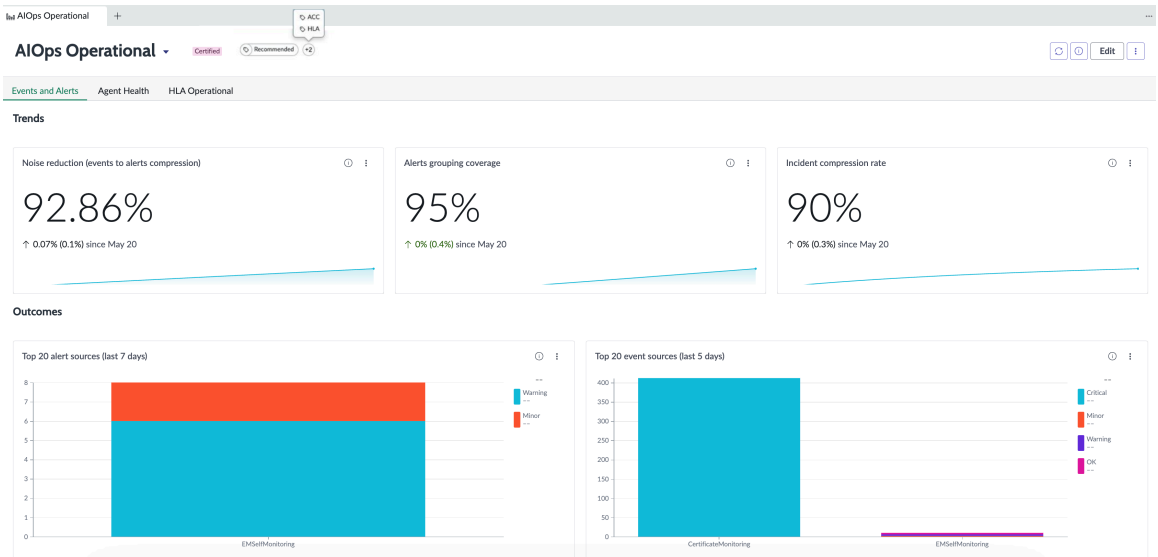

To view AIOps Operational dashboards, such as AIOps Value Realization, select the AIOps Operational list and then select the dashboard name.

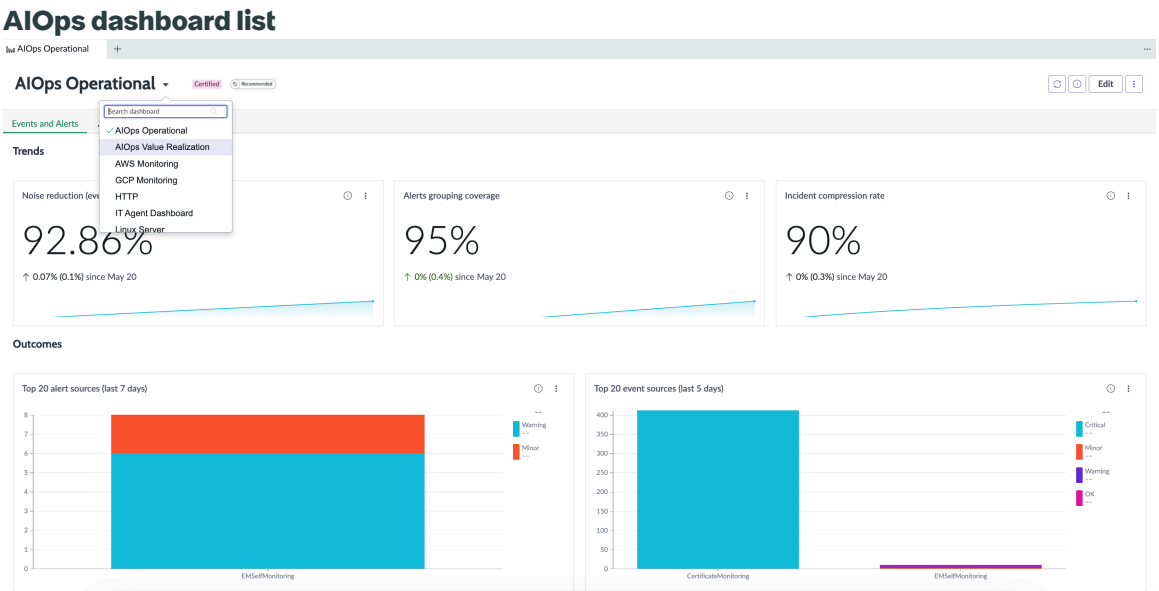

AIOps 360 overview dashboard

The AIOps 360 Overview dashboard offers a unified view of value and performance across IT operations, helping teams track efficiency, monitor alert handling, and assess automation outcomes. It enhances visibility into AIOps impact, supports informed decisions, and drives faster, smarter operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prerequisites

Ensure that the Event Management application is installed.

Required ServiceNow AI Platform roles

- evt_mgmt_admin
- evt_mgmt_operator

Access the AIOps 360 overview dashboard

To open the dashboard, navigate to **Workspaces** > **Service Operations Workspace**. From the left navigation pane, select the AIOps Dashboard icon ( ). The AIOps 360 overview dashboard page opens.

Breakdowns

Breakdowns available in the AIOps 360 overview dashboard are:

- Value overview
- Performance overview

Value overview section

Use filters such as **Date** and **Assignment Group** to customize the dashboard view and focus on specific time periods and teams. This helps you pinpoint trends, identify areas for improvement, and make data-driven decisions tailored to your operational needs.

The Value overview section highlights the overall impact of AIOps on operational efficiency. It showcases how much manual effort has been reduced through automation and intelligent alert grouping. It also reflects noise reduction levels, outage history, and overall workload savings—giving a clear picture of time and cost benefits realized.

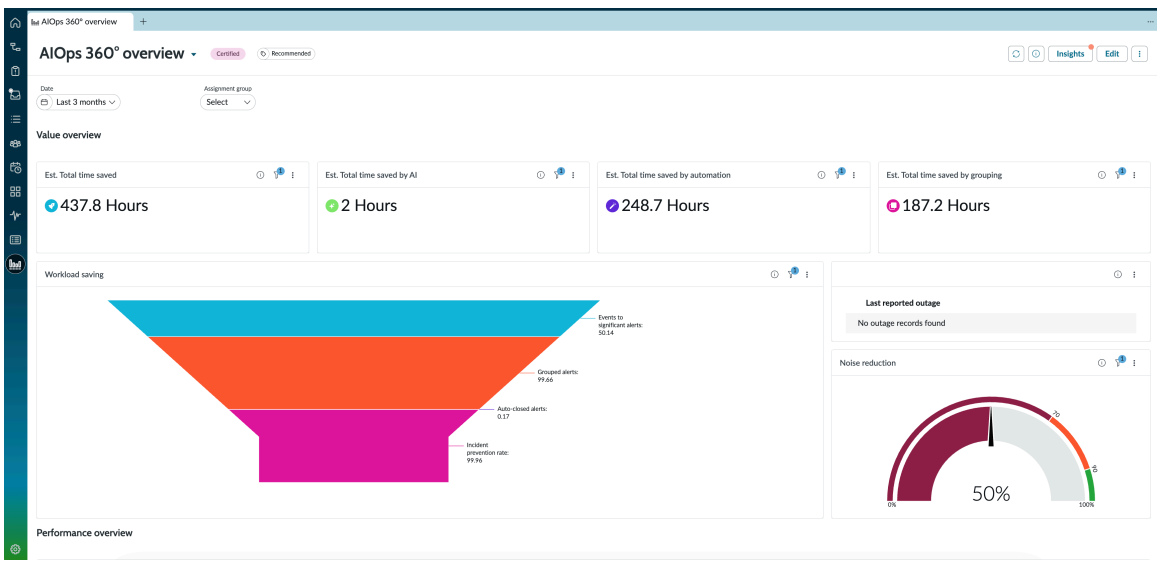

Metrics	Description
Est. Total time saved	Sum of time saved through alert automation and alert grouping, calculated in hours.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metrics	Description
Est. Total time saved by AI	Sum of time saved through AI-driven actions such as automatic noise classification and AI-generated analysis, calculated in hours. This is distinct from time saved by rule-based alert automation and grouping.
Est. Total time saved by automation	Time saved based on the number of alert automation rules executed multiplied by the mean resolution time per alert, calculated in hours. The mean resolution time can be adjusted by administrators in **All** > **System Properties** > **All Properties**. Use the *sn_aiops_dashboard.ITOM.Dashboard.Avg_time_* property to get the actual time saved by automation, in minutes, as used by the Automation Time Saved indicator. By default, it's set to 5 minutes.
Est. Total time saved by grouping	Time saved based on the number of alerts grouped versus the time required to triage each alert individually, calculated in hours. The mean resolution time can be adjusted by administrators in **All** > **System Properties** > **All Properties**. Use the *sn_aiops_dashboard.ITOM.Dashboard.MTTR_for_* property to define the MTTR (in minutes) used in the alert grouping time saved calculation. By default, it's set to 5 minutes.
Workload saving	Measures the reduction in manual effort by filtering, grouping, automating alert closures, and preventing incident escalations. • Events to significant alerts: Percentage of events that were filtered out and did not convert into alerts, reduced through AIOps settings like ignore automations. • Grouped alerts: Percentage reduction in the number of stand-alone alerts due to grouping rules defined in alert grouping automations. • Auto-closed alerts: Percentage of alerts automatically closed by Respond

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metrics	Description
	automation, Clear event actions, or closed as part of alert group resolution. • Incident prevention rate: Percentage of alerts resolved successfully without needing escalation to incidents.
Last reported outage	Most recent date of a major outage (e.g., CMDB) affecting a service or server, based on reported outages on all classes inherited from cmdb_ci_service or cmdb_ci_server.
Noise reduction	Percentage effectiveness in ignoring duplicate and meaningless events. Typical customer value exceeds 90%.

Performance Overview section

This section focuses on how effectively alerts and incidents are being handled. It includes metrics like resolution times across groups, trends in detection and response, automation usage, and the distribution of resolution ownership. It helps identify performance gaps, track improvements, and optimize response strategies across the organization.

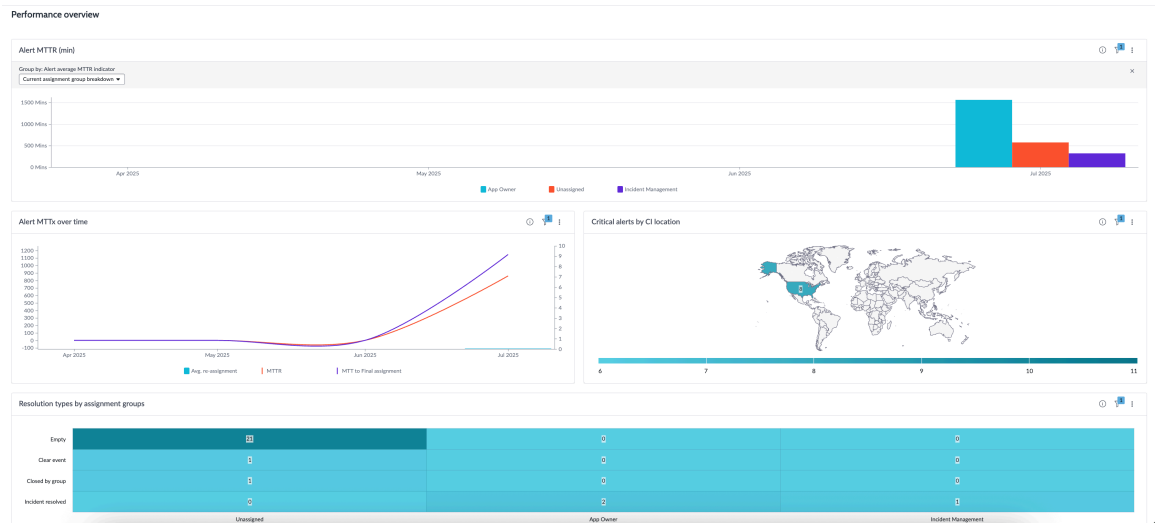

Metrics	Description
Alert MTTR	Alert resolution time by assignment group to identify slower-responding teams and enforce SLAs. Can also be grouped by severity using the grouping drop-down menu.
Alert MTTx over time	MTTR (mean time to resolve), MTTA (mean time to final group assignment), and average group reassignment count tracked over time. Indicates assignment efficiency and the ability to resolve alerts within a single group.
Critical alert by CI location	Critical alerts on a geographic map based on the physical or logical location of Configuration Items (CIs). Helps identify

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metrics	Description
	recurring hotspots for retrospective analysis and to spot regional trends.
Resolution types by assignment groups	Closure reasons by top assignment groups, highlighting automation closures (within 10 minutes of execution) to encourage quick resolution, and time threshold closures (after the auto-close interval), which indicate unattended alerts. You should consider the number of alerts that are automatically closed when they reach their time limit. This figure can reveal if your team is unable to resolve or close all alerts, leaving some unaddressed.

Note: If you're starting with Brazil, some charts, graphs, or metrics in the dashboard may appear blank or show limited data. This is expected if there's no historical data or only a few weeks or months of data available. As your system continues to collect and process alerts, the dashboard will populate with more complete insights over time.

If you want to know about funnel visualization, see Create a funnel visualization.

AIOps Value Realization dashboard

The AIOps Value Realization dashboard uses Performance Analytics to offer comprehensive visibility into the business outcomes of alerts, events, and incidents. It features metrics such as noise reduction (events to alerts compression), average MTTR for incidents created by Event Management[var.event-mgmt], and the most critical services impacted by hours.

The AIOps Value Realization dashboard includes a calendar feature that enables users to select various date ranges such as custom dates, the last 3 months, the last 30 days, the last 6 months, the last 7 days, the last year, today, and year-to-date (YTD). You can specify a custom start and end date and apply these selections to view reports specifically for the chosen date range. This feature facilitates tailored analysis and reporting based on specific time periods of interest.

AIOps Value Realization dashboard date range selection

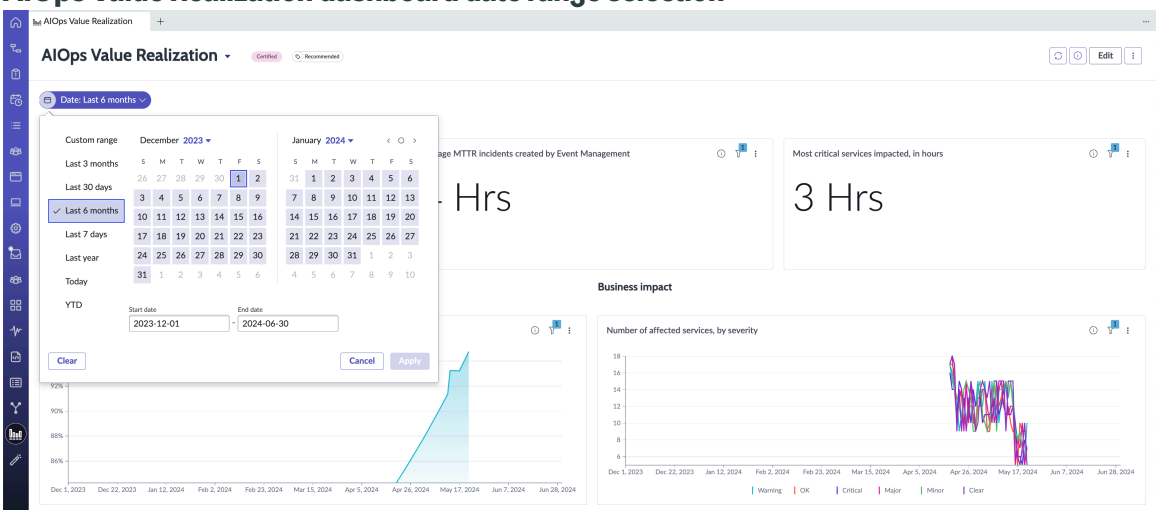

The AIOps Value Realization dashboard offers key insights and metrics crucial for operational efficiency and business impact assessment. It includes features such as noise reduction to enhance focus on critical issues by reducing event-to-alert conversions. Tracking average

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MTTR for incidents and measuring downtime on critical services aids in improving operational efficiency and prioritizing resolutions. The dashboard is categorized into Operational Excellence and Business Impact sections. It provides detailed coverage on alerts grouping, swift resolution of critical incidents reported by users, and quantifying service impacts by severity to optimize operational stability.

Run the [PA EM] Historic Data Collection job once to enable the partial collection of historical Event Management data:

1. Navigate to **All** > **Performance Analytics** > **Data Collector** > **Jobs**.
2. Select [PA EM] Historic Data Collection.
3. Select **Execute Now**.

AIOps Value Realization dashboard key features overview

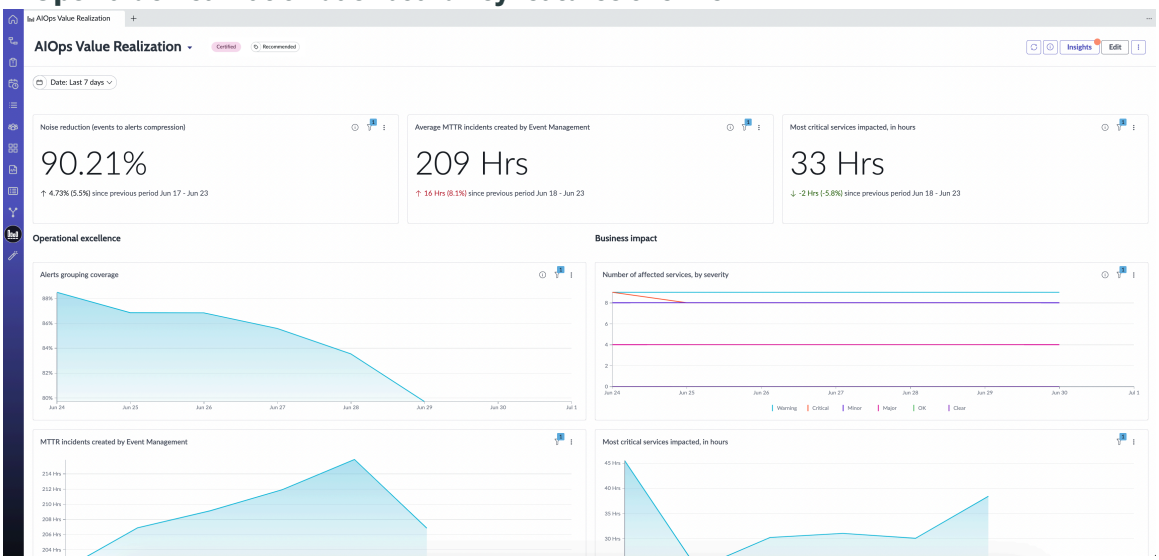

Prerequisites

Ensure that the Event Management application is installed.

Required ServiceNow AI Platform roles

- evt_mgmt_admin
- evt_mgmt_operator

Access the AIOps Value Realization dashboard

To open the dashboard, navigate to **Workspaces** > **Service Operations Workspace**. From the bottom of the navigation pane, select the AIOps configuration center icon .The AIOps configuration center page appears. On the ITOM AIOps configuration center page, under the **Optimize** > **Dashboards** section, select **View AIOps Value Realization dashboard**.

Use cases

For examples of how different people in your organization would use this dashboard, refer to these use cases.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

User	Dashboard use
evt_mgmt_admin or admin	Edit the dashboard and grant view and share permissions.
evt_mgmt_operator or admin	View the dashboard and review record details to visualize and track the business impact of alerts, events, and incidents.

Breakdowns

Breakdowns available in the AIOps Value Realization dashboard are:

- Operational Excellence
- Business Impact

Reports

Title	Type	Description
Noise reduction (events to alerts compression)	Line graph	The compression rate from events to alert creation: a higher number indicates fewer alerts are being created.
Average MTTR incidents created by Event Management	Line graph	The average Mean Time to Repair (MTTR) of an incident created by Event Management: a lower number indicates greater efficiency of your IT team.
Most critical services impacted, in hours	Line graph	The total time your organization's most critical business services are impacted: a higher number indicates more time critical services are affected by open alerts.
Alerts grouping coverage	Line graph	The percentage of total alerts that are grouped: a higher number indicates more alerts are being correlated and grouped within an alert group.
Number of affected services, by severity	Line graph	The number of services impacted by alerts, grouped by severity, such as the count of services experiencing critical issues per day.
Number of new P1 incidents on monitored services created by users	Line graph	The number of incidents created on services monitored by Event Management: a lower number indicates fewer users are being impacted

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Description
		by a service outage or performance degradation.

Event and Alert dashboard

The Event and Alert dashboard uses Performance Analytics to provide real-time visibility into events and alerts in Event Management, showcasing key trends, outcomes, and the most impacted configuration items. It highlights metrics such as noise reduction, alert grouping coverage, and top alert sources.

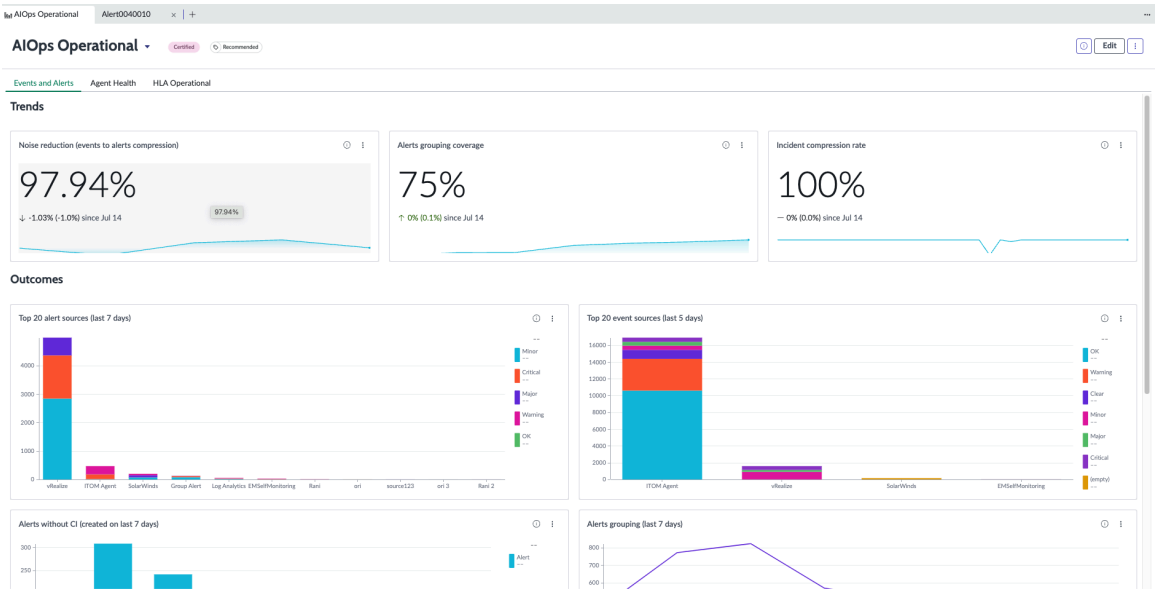

Run the [PA EM] Historic Data Collection job once to enable the partial collection of historical Event Management data:

1. Navigate to **All** > **Performance Analytics Administration** > **Data Collector** > **Jobs**.
2. Select [PA EM] Historic Data Collection.
3. Select **Execute Now**.

Prerequisites

Ensure that the AIOps Experience [`sn_sow_aiops`] application is installed. The application installs the Event Management and Performance Analytics components required for the dashboard. For installation instructions, see Install Service Operations Workspace for ITOM application.

Required ServiceNow AI Platform roles

- evt_mgmt_admin
- evt_mgmt_operator

Access the Events and alerts dashboard

To open the dashboard, use one of the following methods:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Navigate to **All** > **AIOps Dashboards** > **AIOps Operational** > **Events and Alerts**.
- Navigate to **Workspaces** > **Service Operations Workspace** and select the AIOps Dashboards icon ( ).

  By default, the **Events and Alerts** tab is selected.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
evt_mgmt_admin or admin	Edit the dashboard and grant view and share permissions.
evt_mgmt_operator or admin	View the dashboard and details of the records contained in it to visualize and track events, alerts, trends, outcomes, and the most impacted Configuration Items in your organization.

Breakdowns

Breakdowns available in the Event and Alert dashboard are:

- Trends
- Outcomes

Reports

Title	Type	Description
Noise reduction (events to alerts compression)	Line graph	The compression rate from events to alert creation. The higher the number, the fewer alerts are being created.
Alerts grouping coverage	Line graph	The percentage of alerts aggregated into grouped alerts over time.
Incident compression rate	Line graph	The percentage of alerts that did not result in incident creation. A higher percentage means more alerts were resolved without generating incidents.
Top 20 alert sources (last 7 days)	Bar chart	The number of alerts per source categorized by severity over the last 7 days.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Description
Top 20 event sources (last 5 days)	Bar chart	The number of events per source categorized by severity over the last 5 days.
Alerts without CI (created on last 7 days)	Line graph	The number of alerts without CI binding created over the last 7 days.
Alerts grouping (last 7 days)	Bar chart	The distribution of grouped alerts over the last 7 days.

Most impacted Configuration Items

The Most Impacted Configuration Items section of the Event and Alert dashboard provides a comprehensive overview of the configuration items (CIs) that are most impacted by the issue. This section lists key details such as the name of each CI, the number of associated alerts, their classification (such as application service), and location. Additionally, it identifies the owner and support group responsible for each CI, facilitating targeted and efficient incident resolution. By highlighting the most impacted CIs, this section helps prioritize critical assets and resources. This approach confirms that the most significant issues are addressed promptly to maintain system stability and performance.

Agent Health dashboard

The Agent Health dashboard provides visibility into the status and distribution of agents installed in your system. It allows you to quickly assess the number of deployed agents and their distribution, helping you prioritize your work more efficiently.

The following example shows the agents in your system that are currently up, which agents are configured to collect data, and which agents have stopped data collection due to having reached their CPU threshold (Silent/Off). It also displays the agents according to their version, OS, and connected MID Server.

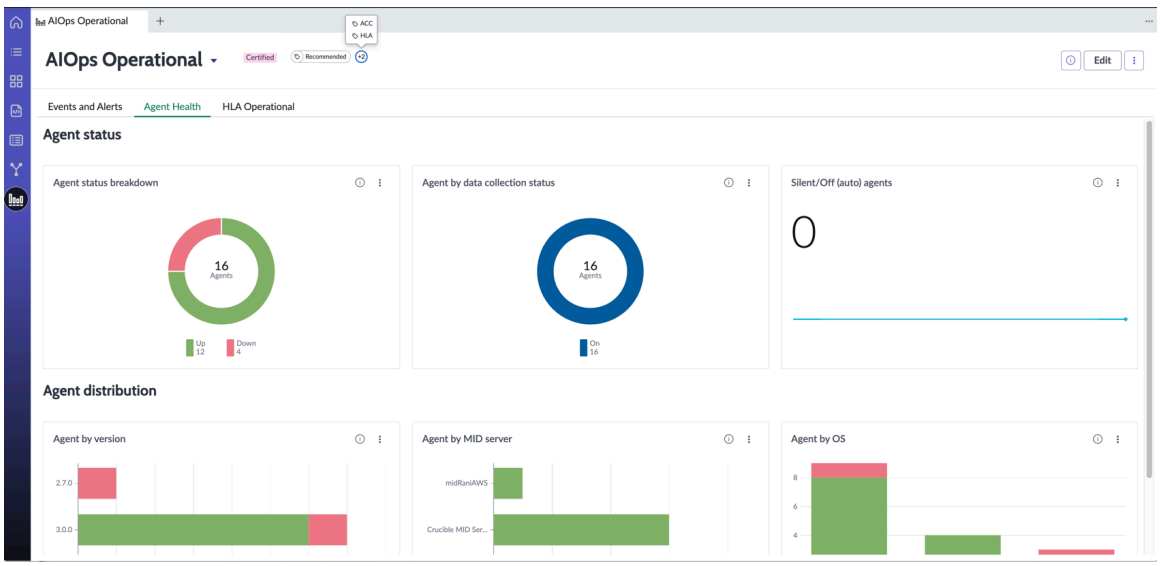

Prerequisites

The Agent Client Collector application must be installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Required ServiceNow AI Platform roles

A user must have the agent_client_collector_admin role to see all the widgets in the dashboard.

Access the Events and alerts dashboard

To open the dashboard, use one of the following methods:

- Navigate to **All** > **AIOps Dashboards** > **AIOps Operational** > **Agent Health**.
- Navigate to **Workspaces** > **Service Operations Workspace** and select the AIOps Dashboards icon ( ) and then select the **Agent Health** tab.
- After upgrading to Discovery Admin Workspace version 1.3.1 (August 2024 Store), you can navigate to **Workspaces** > **Discovery Admin Workspace** > **Insights** and use the enhanced dashboard.

Use case

For examples of how different people in your organization would use this dashboard, see the use case in the following table.

Dashboard use case

User	Dashboard use
admin	Visibility into the status of the agents.

Breakdowns

Breakdowns available in the Agent Health are:

- Agent status: The status of the agents currently deployed in your system.
- Agent distribution: The distribution of the agents by the agent version number, the MID Server that the agent is assigned to, and the OS that the agent is working on.

Reports

Title	Type	Description
Agent status breakdown	Pie chart	Agents deployed in your system that are up and down. Select the relevant color in the widget to display a list of up or down agents.
Agent by data collection status	Pie chart	Agents deployed in your system with data collection turned on and off. Select the relevant color in the widget to display a list of agents with data collection on or off.
Silent/Off (auto) agents	Line graph	Agents that automatically stopped collecting data after

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Description
		crossing the CPU usage threshold. The CPU usage threshold is configurable in the `acc.yml` configuration file. For more information on the `acc.yml` configuration file, see Configuration file options.
Agent by version	Bar chart	Breakdown of your deployed agents by their version numbers.
Agent by *`MID Server`*	Bar chart	Breakdown of the deployed agents according to the *`MID Server`* they're assigned to.
Agent by OS	Bar chart	Breakdown of your deployed agents according to the operating systems they are working on.

HLA Operational dashboard

The Health Log Analytics Operational dashboard uses Performance Analytics to monitor log data, alerts, and error rates in Service Operations Workspace.

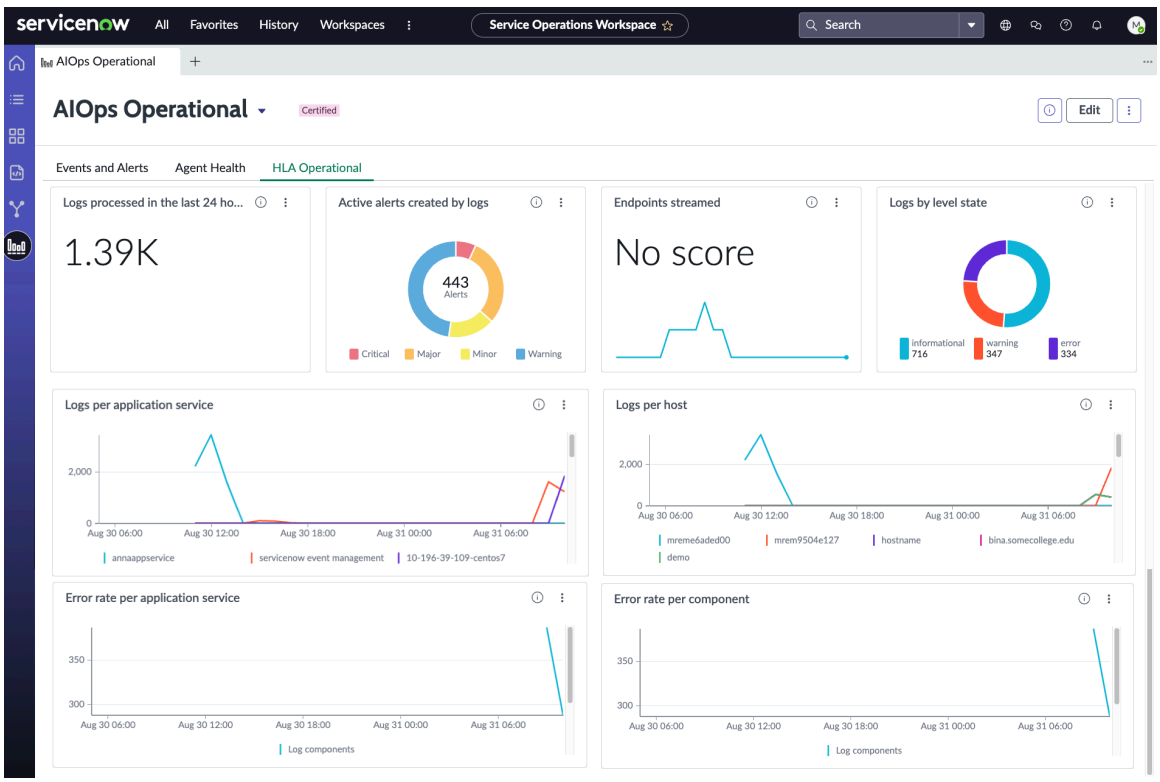

Prerequisites

Ensure that the Health Log Analytics application is installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Required ServiceNow AI Platform roles

- evt_mgmt_admin
- evt_mgmt_operator

Access the HLA Operational dashboard

To open the dashboard, use one of the following methods:

- Navigate to **All** > **AIOps Dashboards** > **AIOps Operational** > **HLA Operational**.
- Navigate to **Workspaces** > **Service Operations Workspace** and select the AIOps Dashboards icon ( ), then display the HLA Operational dashboard.

Use cases

These use cases provide examples of how different people in your organization would use this dashboard.

User	Dashboard use
evt_mgmt_admin	Edits the dashboard and grants view and share permissions.
evt_mgmt_operator or evt_mgmt_admin	Monitors the dashboard to visualize and track HLA log and alerts data, error rates, and other relevant information.

Breakdowns

Breakdowns available in the HLA Operational dashboard are:

- HLA Overview
- Error rate

Reports

Title	Type	Description
Logs processed	Bar chart	The number of logs that were processed and analyzed in the selected period. By default, the dashboard visualizes log data that was ingested in the last 24 hours.
Active alerts created by logs	Table	Detailed data on all currently active alerts.
Endpoints streamed	Line graph	The number of endpoints (IPs, networks, and devices) that are streaming into Health

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Title	Type	Description
		Log Analytics, by application service.
Logs by level state	Bar chart	The severity level of the logs for the streaming sources.
Logs per application service	Line graph	The number of logs streamed for each application service in the selected period. Monitor this data to ensure that log data is streaming as expected for each application service.
Logs per host	Line graph	The number of processed logs for each host in the selected period. Monitor this data to detect hosts that are logging more or less data than expected.
Error rate per application service	Line graph	The number of logs with errors for each application service analyzed in the selected period. Monitor this data to further investigate the errors. **Note:** Health Log Analytics automatically creates alerts for Error anomalies.
Error rate per component	Line graph	The number of logs with errors for each component analyzed in the selected period. Monitor this data to ensure that the error rate for each component is normal.

Data visualization in ITOM

Gain insights faster with data visualization in IT Operations Management (ITOM). Visualize metrics, logs, events, and alerts through intuitive dashboards and charts to identify trends, monitor performance, and resolve issues proactively.

Anyone with access to data can create a visualization of that data on a dashboard that they own.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a funnel visualization

A funnel visualization shows how data progresses through sequential stages, helping you identify drop-offs or bottlenecks at each step. Use it to track conversions, workflows, or process completion rates. You can add a funnel visualization to an existing dashboard or to a new dashboard.

Before you begin

Role required: report_user, admin, or viz_creator

Ensure you have an admin role or permission from the dashboard owner to edit the dashboard.

About this task

If you have the appropriate role, you can add and edit funnel reports on a dashboard. The **Edit** button is not available on the ITOM dashboards unless the user has rights to edit it. For more details on creating dashboards, see Common dashboard tasks in the in-line editor.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and select the AIOps Dashboards icon ( ).
2. From the Dashboard picker, select **AIOps 360 overview**.

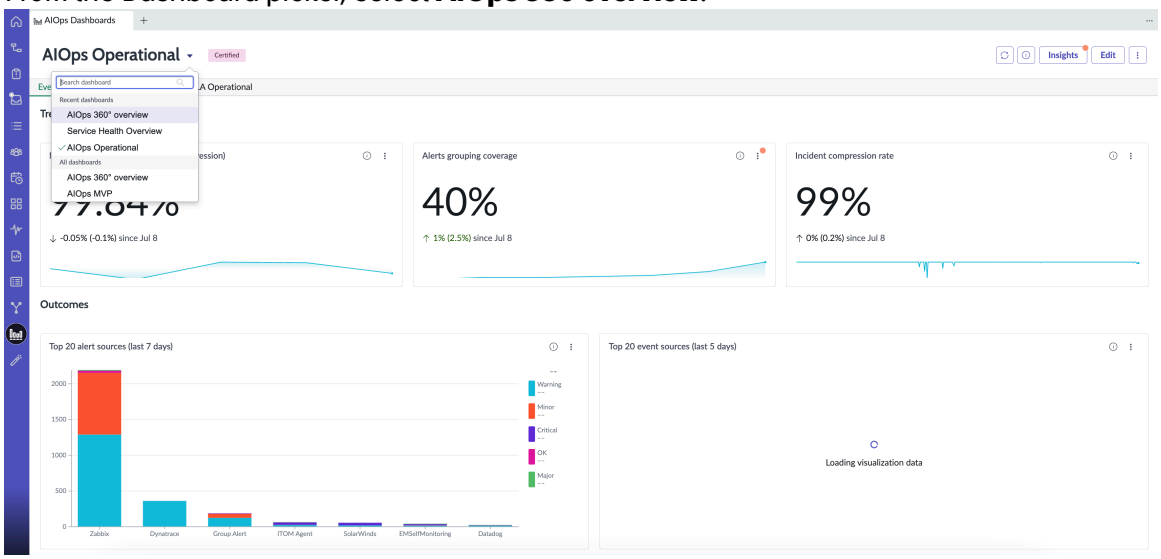

3. Select **Edit** to put dashboard in the edit mode.

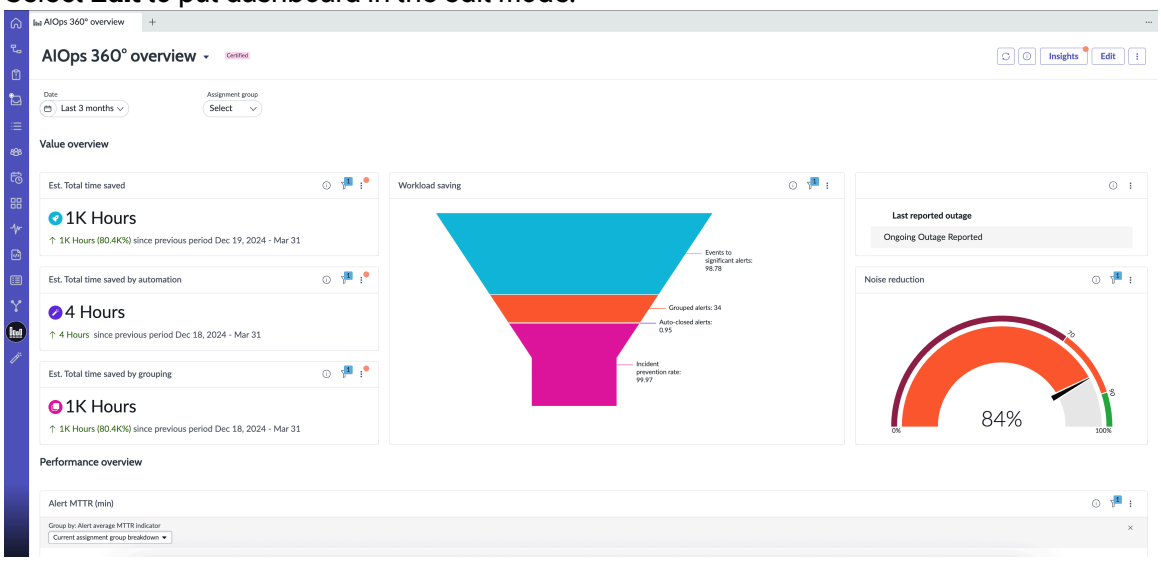

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. From the **Add new element** drop-down menu, select **Data visualization**.

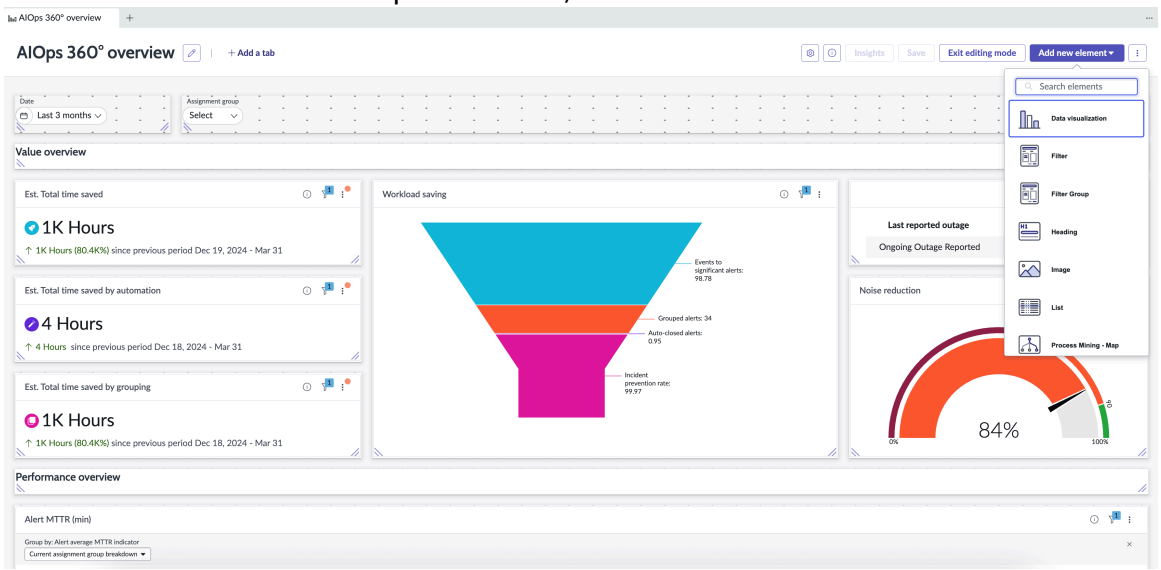

5. In the Add Visualization window, select **New visualization**.

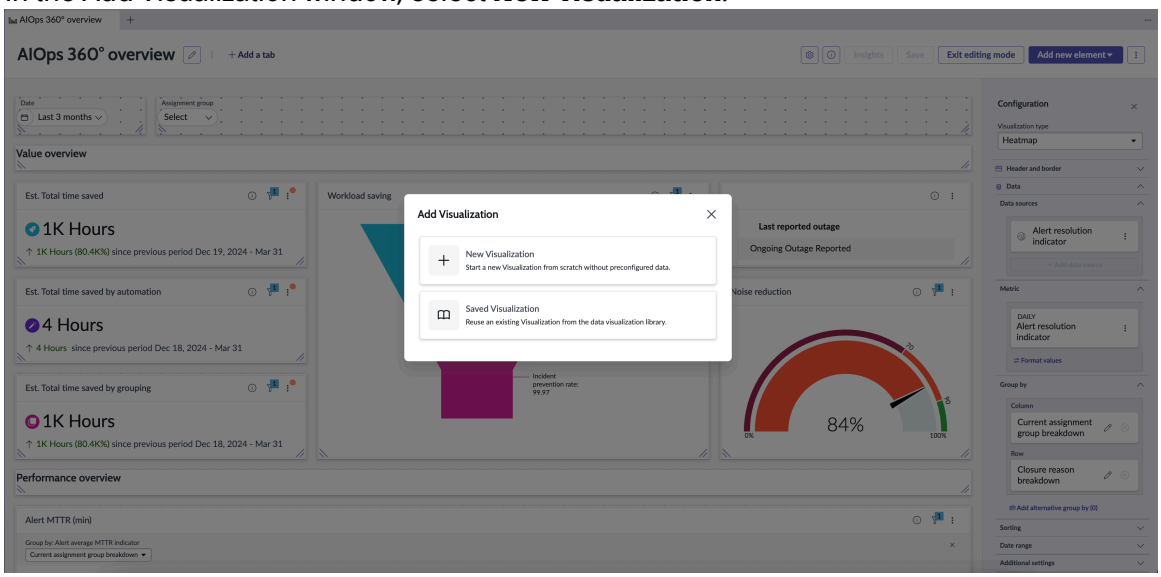

6. From the **Visualization type** field, select **Funnel**.

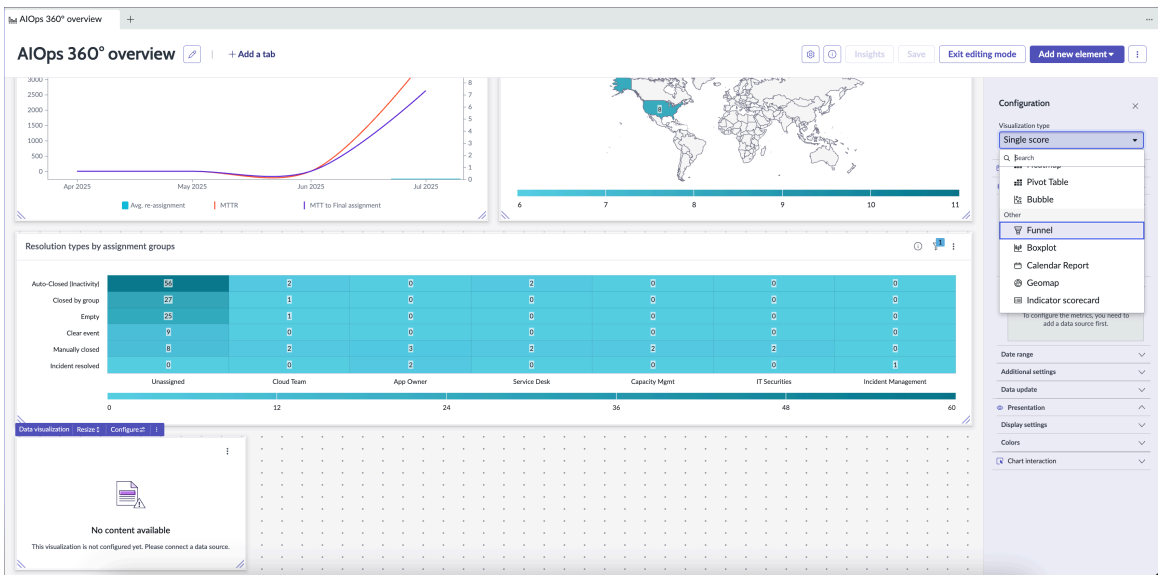

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The **Funnel** visualization type is currently available only to customers who have the ITOM plugin installed.

7. Under **Data sources**, select **+ Add data source**.
   There are two kinds of data sources avail for the **Funnel** visualization type: Table and Indicator. Visualizations based on table data show a snapshot of the state of that data. Indicator visualizations reflect data over time based on data collection jobs. You can configure funnel stages using indicators in one of the following ways:

   You can configure funnel stages using indicators in one of the following ways:

   - Option 1: Use a single indicator or table with a breakdown

     Select a single indicator or table, and then define the breakdown using the Group by field. The breakdown values are used as the funnel stages.

     Note: If you select only one indicator without a breakdown, the visualization displays an error: `Invalid configuration. A 'Group By' must be selected for the visualization to render.`

   - Option 2: Use multiple indicators without a breakdown

     Select multiple indicators without setting a breakdown. In this case, each indicator's value represents a stage in the funnel.

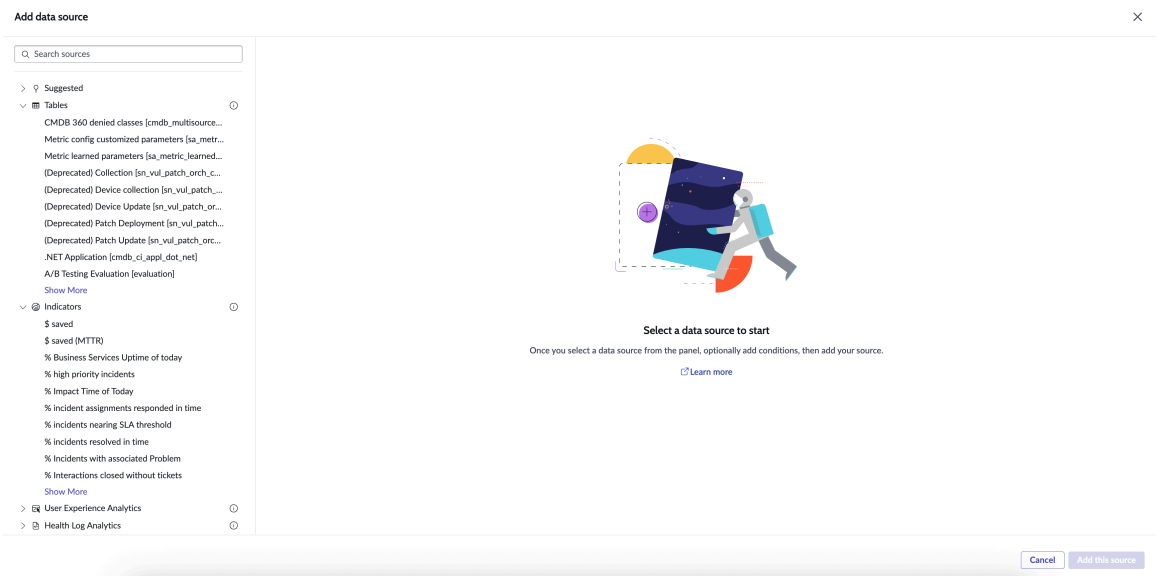

8. Select the required data source and select **Add this source**.

   When you select a table or indicator data source for a funnel visualization, refer to the following topics for the available data configuration options:

   - Table data options for funnel visualizations
   - Indicator data options for funnel visualizations

9. Select **Save**.

Table data options for funnel visualizations

When you select a table data source for the funnel visualization, the following data configuration options are available.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data configuration options

Field	Description
Predefined conditions	Preset filters to quickly narrow down data based on common criteria.
+ Add custom conditions	Custom filters based on user-defined criteria.
Header and border	
Show header	The visualization header, including title and icons.
Show header separator	Option to display a line separating the header from the rest of the component.
Chart title	Text shown as the title of the chart.
Description	Optional explanatory text shown below the title.
Wrap title	Option to wrap long titles onto multiple lines.
Show border	Option for border visibility setting for the widget container.
Header background color	Background color setting for the header area.
Data	
Data sources	Available inputs that supply data to the visualization.
Metric	
Label	Custom name for the selected metric. By default, the label matches the data source name.
Select an aggregation	Aggregation type applied to the metric (e.g., sum, average). The default is Count, which displays the number of records selected. **Note:** Select **Count Distinct** to count only unique records. For example, when generating a report of users with one or more roles from a list, you count each user only once. Without this option, users with multiple roles are counted more than once.
Field	Source field from the data set for the selected metric. **Note:** The data appears in a format that matches the field type. For example, an integer field like **Priority** is shown as a number. A duration field like **Business duration** from the [Incident] table is displayed in days, hours, and minutes. If function fields were created in Report Designer for the table, they are available in the drop-down list. For more information, see Report on function fields.
Format values	Formatting style for displaying metric values. Select to open the Formatting values dialog. For more information, see Value formatting in reports.
Group	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data configuration options (continued)

Field	Description
Field for incident	Field used specifically to group incident data.
Max number of groups	Limit on the number of groups shown in the visualization.
Show others	Option to group excess data into an **Others** category.
Data update	
Follow filter	Follows filters set at the page level. When selected, the visualization reflects workspace filters. When not selected, the visualization ignores all page-level filters.
Show filter icon	Option to display the filter icon and the number of filters affecting the visualization. If the check box is not selected, the icon and filter count are hidden.
Refresh after being away for X min	Trigger based on time spent away from the dashboard within the same workspace, including the Platform Analytics experience. Defines the number of minutes after which the visualization automatically refreshes when you return. Do not confuse this setting with **Scheduled repetition** on inline dashboards, which refreshes the entire dashboard only while it remains open.
Info messages	
Set custom message when no data	Custom text displayed when no data is available for the visualization.
Colors	
Set color type	Color scheme applied to different data values or ranges.
Chart interaction	User actions supported within the chart, such as hover or select.

Indicator data options for funnel visualizations

When you select an indicator data source for the funnel visualization, the following data configuration options are available.

You can select a single indicator and use the **Group by** field to calculate stage values. Alternatively, select multiple indicators without a breakdown, where each indicator's value represents a funnel stage.

Data configuration options

Field	Description
Breakdown	Provides a way to segment data using fields like Alert Severity or Alert Status. Each condition includes a **Breakdown** field, an **Operator**, and an **Element** that defines the target value. All

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data configuration options (continued)

Field	Description
	conditions are applied using AND logic, meaning data must meet every specified criterion to be included in the visualization.
Header and border	
Show header	Option for visibility setting for the widget header.
Show header separator	Option to display a dividing line below the header.
Chart title	Text shown as the title of the chart.
Description	Optional explanatory text shown below the title.
Wrap title	Option to wrap long titles onto multiple lines.
Show border	Option for border visibility setting for the widget container.
Header background color	Background color setting for the header area.
Data	
Data sources	Available inputs that supply data to the visualization.
Metric	
Label	Custom name for the selected metric. By default, the label matches the data source name.
Select an aggregation	Aggregation type applied to the metric (e.g., sum, average). The default is Count, which displays the number of records selected. **Note:** Select **Count Distinct** to count only unique records. For example, when generating a report of users with one or more roles from a list, you count each user only once. Without this option, users with multiple roles are counted more than once.
Field	Source field from the data set for the selected metric. **Note:** The data appears in a format that matches the field type. For example, an integer field like **Priority** is shown as a number. A duration field like **Business duration** from the [Incident] table is displayed in days, hours, and minutes. If function fields were created in Report Designer for the table, they are available in the drop-down list. For more information, see Report on function fields.
Format values	Formatting style for displaying metric values. Select to open the Formatting values dialog. For more information, see Value formatting in reports.
Group	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data configuration options (continued)

Field	Description
Field for incident	Field used specifically to group incident data.
Max number of groups	Limit on the number of groups shown in the visualization.
Show others	Option to group excess data into an **Others** category.
Data update	
Follow filter	Follows filters set at the page level. When selected, the visualization reflects workspace filters. When not selected, the visualization ignores all page-level filters.
Show filter icon	Option to display the filter icon and the number of filters affecting the visualization. If the check box is not selected, the icon and filter count are hidden.
Refresh after being away for X min	Trigger based on time spent away from the dashboard within the same workspace, including the Platform Analytics experience. Defines the number of minutes after which the visualization automatically refreshes when you return. Do not confuse this setting with **Scheduled repetition** on inline dashboards, which refreshes the entire dashboard only while it remains open.
Info messages	
Set custom message when no data	Custom text displayed when no data is available for the visualization.
Colors	
Set color type	Color scheme applied to different data values or ranges.
Chart interaction	User actions supported within the chart, such as hover or select.

Event Management

ServiceNow® Event Management is a robust application that helps keep your IT systems healthy by spotting problems quickly and fixing them. It collects events from different sources, figures out what's causing the issues, and converts them into alerts. These alerts are then analyzed, grouped, and acted upon to resolve issues and maintain system health.

Event Management seamlessly integrates with various monitoring tools via connectors and listeners, enabling comprehensive data collection and correlation. Available as a separate subscription from the main ServiceNow platform, Event Management allows organizations to tailor their usage to specific needs, ensuring effective and efficient incident management without unnecessary overhead.

https://player.vimeo.com/video/1120212844?h=131956e0de&badge=0&autopause=0&player_id=0&app_id=58479

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore Explore the various types of alert grouping and the rules for correlation	Configure Setup Event Management	Use View alert information and monitor service health
Integrations Learn how you can integrate Event Management with other monitoring systems	Event processing Understand the analysis and management of events	Reference Get details about domain separation in Event Management

Dependencies

For a list of Event Management plugin dependencies, see Plugins or applications installed with ITOM AIOps.

Troubleshoot and get help

- Search the Known Error Portal for known error articles
- Ask or answer questions in the Event Management community
- Contact ServiceNow Support
- Video: Getting started with Event Management
- Videos: Getting started with ITOM
- Developer training
- Developer documentation

Exploring Event Management

Explore Event Management to understand its overview, process flow, user roles, and benefits for comprehensive IT issue monitoring and resolution.

Event Management provides comprehensive monitoring, analysis, and remediation of IT issues by effectively managing various components within an IT environment. These components include discovered services, application services, dynamic CI groups, and alert groups.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Discovered Services: Defined by interrelated Configuration Items (CIs) from the CMDB, a discovered service is identified through Service Mapping. It includes a service map with mapping relationships, an impact tree showing outage severity, active and related alerts, and CI properties. This service information is displayed on dashboards, the Alerts list, and the Events list.
- Application Services: Created by selecting specific CIs, application services allow for targeted monitoring and management. For more details, refer to the Application Services documentation.
- Dynamic CI Groups: These are collections of CIs grouped based on shared criteria, such as location. Dynamic CI groups help populate application services, simplifying management.
- Alert Groups: Alert groups organize sets of alerts to streamline maintenance and management, making it easier to respond to IT issues efficiently.

Process flow

Event Management receives external events and generates alerts based on predefined rules. The MID Server polls external event tracking tools and sends data to Event Management for storage and processing. Events are stored in the Event [em_event] table, and alerts are created by matching event rules. Alerts are then transformed and enriched with additional content, accumulated if thresholds are met, and mapped to specific fields. The system searches for matching message keys to update existing alerts or create new ones, associating related events under a single alert. Alerts are bound to specific Configuration Items (CIs) for root cause analysis. For more information, see Event Management process flow.

Users

Role title [name]	Description
Admin [evt_mgmt_admin]	Has read and write access to all Event Management features to configure Event Management. **Note:** Exercise caution with the evt_mgmt_admin role, as it can be elevated to the admin role. A user with the evt_mgmt_admin role has the ability to add and modify scripts that run on a global scope. Ensure proper access control. With this role, the user can create and/or update the following scripts: • Alert correlation rules • Alert management rules • Maintenance rules • Advanced scripts • Event field mapping • Pre- and post-binding scripts
Operator [evt_mgmt_operator]	Manages alerts, including closing and acknowledging them. Oversees the overall Event Management process, ensuring events are properly categorized, prioritized, and routed for resolution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description
Team Operator [evt_team_operator]	Manages Event Management operations within a specific team as defined in the **Assignment Group** field. This role allows the operator to read and write alerts exclusively for their assigned team. Additionally, the operator can make configuration changes specific to their team, including updates to Alert Automation and the Integrations Launchpad. **Note:** The evt_team_operator role must be assigned to an assignment group to view and manage alerts for that group. If the role is created but not associated with any groups that have alerts assigned, the operator cannot see any alerts.
User [evt_mgmt_user]	Manages the lifecycle of alerts, including performing basic operations such as viewing and acknowledging them.

Benefits

- Rapid Issue Detection: Quickly identifies and highlights potential IT issues.
- Efficient Alert Handling: Aggregates and correlates alerts for streamlined management.
- Automated Actions: Initiates automatic remediation processes to speed up issue resolution.
- Comprehensive Monitoring: Integrates with multiple tools for a complete system overview.
- Root Cause Analysis: Offers tools to identify underlying causes of issues.
- Customizable Rules: Tailors event and alert management rules to specific needs.
- Reduced Downtime: Minimizes system downtime with prompt problem resolution.
- Enhanced Visibility: Real-time dashboards offer insights into system health.
- Cost Efficiency: Lowers operational costs by preventing prolonged issues.

Event Management process flow

Event Management collects, analyzes, and converts events into alerts, enabling efficient tracking and remediation.

Event Management receives external events and generates alerts based on event and alert management rules. Events are sent directly to your instance via email server, script, SNMP trap, or web service API. The corresponding alerts appear on dashboards for tracking and remediation purposes.

As the computer, software, or service generates events, the MID Server polls the external event tracking tool. The MID Server, maintaining a connection to Event Management, sends the information to your instance for storage, processing, and remediation.

The instance stores events in the Event [em_event] table and attempts to generate alerts based on predefined rules and event mappings. Regardless of whether an alert is generated, the original event is available for review and remediation. Alerts are generated according to the following process flow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Match Event Rule: Find the best matching event rule for an event.

   A rule is matched if the source of the event matches the source specified in an existing rule. Additionally, a rule is matched if the event matches the optional rule filter and the event additional_info value matches the rule Additional Information filter. A rule without any filter is ignored, such as when the source filter or Additional Information filter is missing. If multiple rules are defined for the same type of event, use the rule order to determine the sequence of rule application.

2. Ignore Rule: If the rule **Ignore** check box is selected, no alert is generated. However, the event is still available for review and remediation.
3. Apply Transforms:
   - If transforms have been defined, apply them.
   - If compose parameters are set, apply the additional content to display to the user in the alert.
4. Threshold Accumulation: If Active in the threshold section is selected, accumulate all events until the threshold is met, then generate a single alert for the events.
5. Event Field Mapping
   - Search for an event field mapping even if there was no event rule.
   - If an event field mapping is found, apply the mapping information.
   - If the event has no severity after the event transformations, retain the event for reference purposes and do not generate an alert.
6. Alert Generation
   - Search the Alert [em_alert] table for a matching message key.
   - If a matching message key exists, update the alert according to the event information.
   - If a matching message key does not exist, create an alert.
   - If another event has the same matching key, associate the events under a single alert.
   - For root cause analysis purposes, bind the alert to a specific Configuration Item (CI).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event workflow

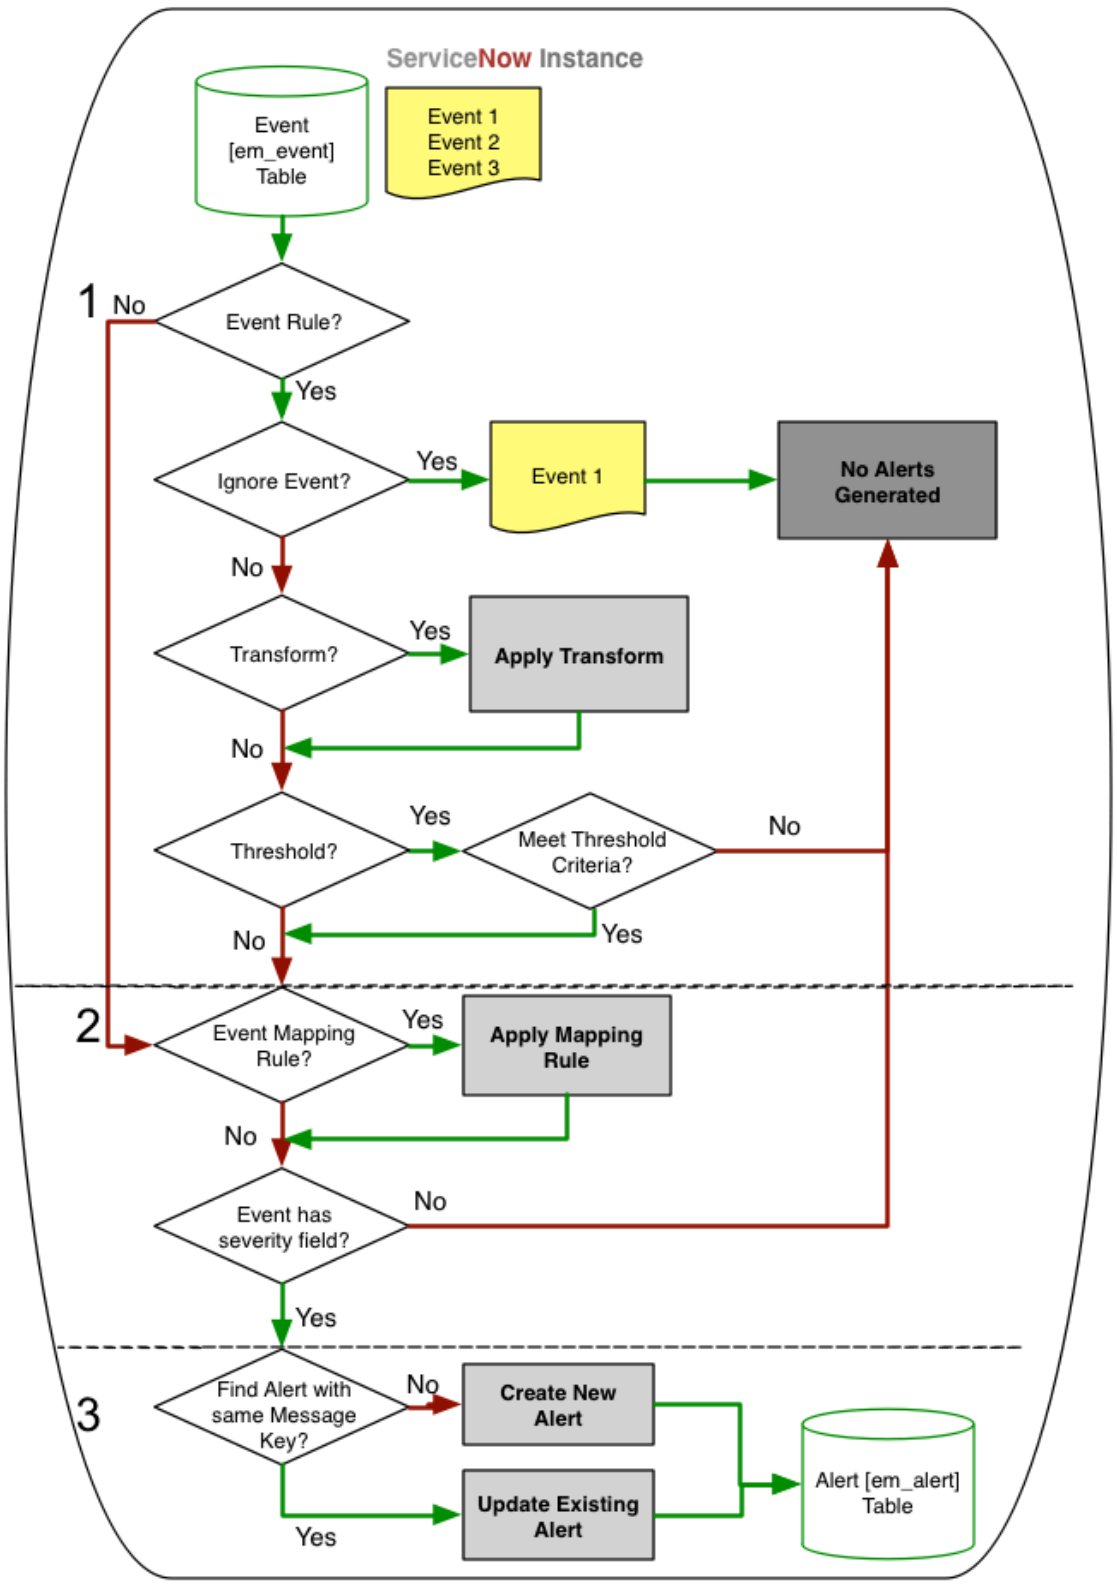

Event Management architecture

Event Management architecture integrates data collection, processing, and alerting into a unified system for streamlined IT issue detection and resolution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Architecture

As events occur on various systems, the MID Server connector instance sends them to the main instance. Event Management generates alerts, applies alert management rules, and prioritizes alerts for remediation and root cause analysis. This information can be viewed on dashboards, the alert list in Alert Intelligence, or from a service map.

Configuring Event Management

Event Management administrators administer events, manage and monitor alerts, aggregate alerts, and work review and monitor services' status with the Operator Workspace service monitor.

Event Management operators find alerts, analyze them, and take action to resolve the underlying issue.

Configure Event Management using ServiceNow Otto for Setup

The ITOM Configuration console gives administrators a single place to complete all Event Management setup steps — from installing plugins to configuring alert automations.

The console contains two modules: Discovery and Event Management.

Discovery guides you through the end-to-end setup of your IT infrastructure scanning. It covers assigning administrators, deploying and validating MID servers, configuring IP and cloud credentials, scheduling discovery runs, and reviewing schedules. It also includes diagnostics to identify discovery errors and detect anomalies, ensuring your environment is fully validated. It also includes diagnostics to identify discovery errors and detect anomalies. This ensures your environment is fully validated before you promote your configuration to production.

Event Management helps you take control of your IT alerts from start to finish. It begins with assigning administrators, setting up integration accounts, and installing the necessary integrations to connect your monitoring tools. From there, you can review and enrich incoming alerts, create incidents automatically, and reduce noise by grouping, ignoring, or delaying alerts. The platform also provides structured training for your operations team on key tools such as Service Operations Workspace, Express List, and Service Dashboard. Finally, built-in dashboards help you report on outcomes and measure progress before promoting your configuration to production.

Access Event Management configuration console

Access the IT Operations Management configuration console to set up Event Management and configure AIOps capabilities.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
   The Admin Home page opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

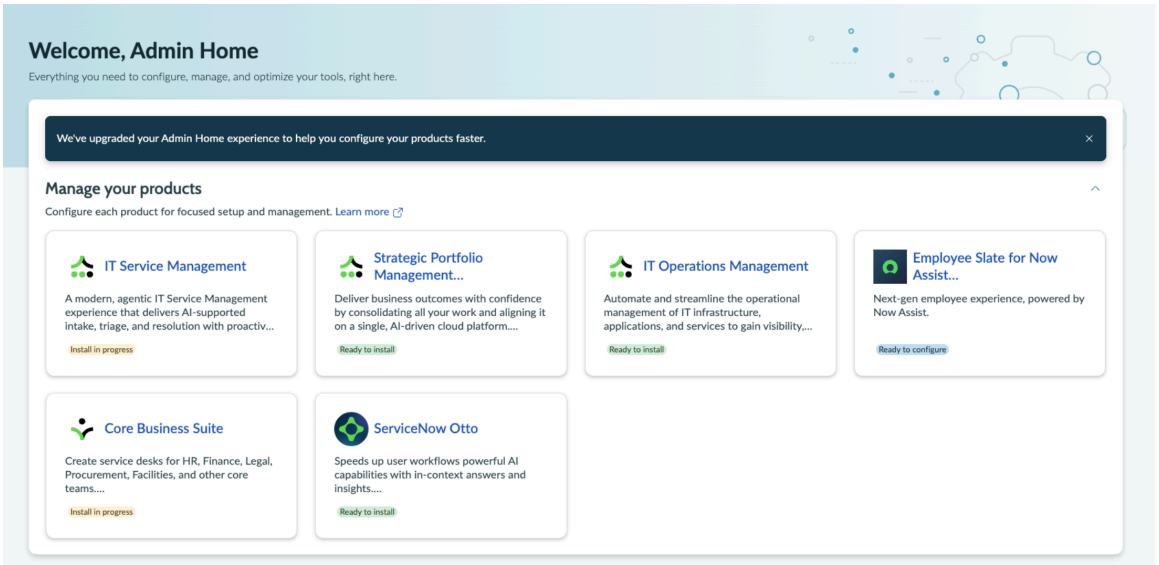

2. In the IT Operations Management tile, select **View product overview**.
3. In the IT Operations Management page, select **Configure**.

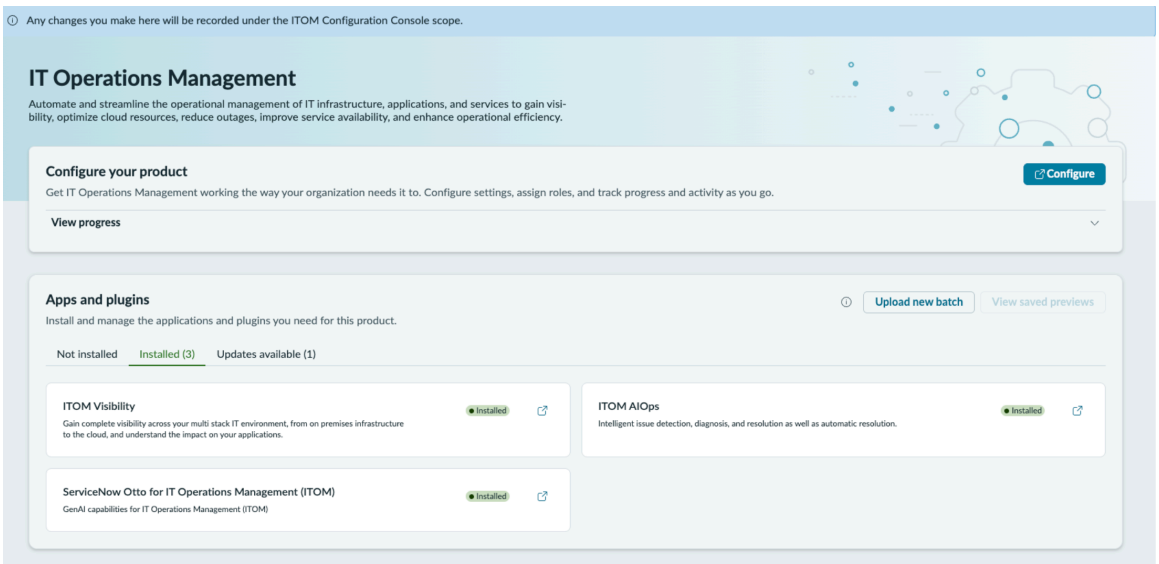

The Configure IT Operations Management page opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

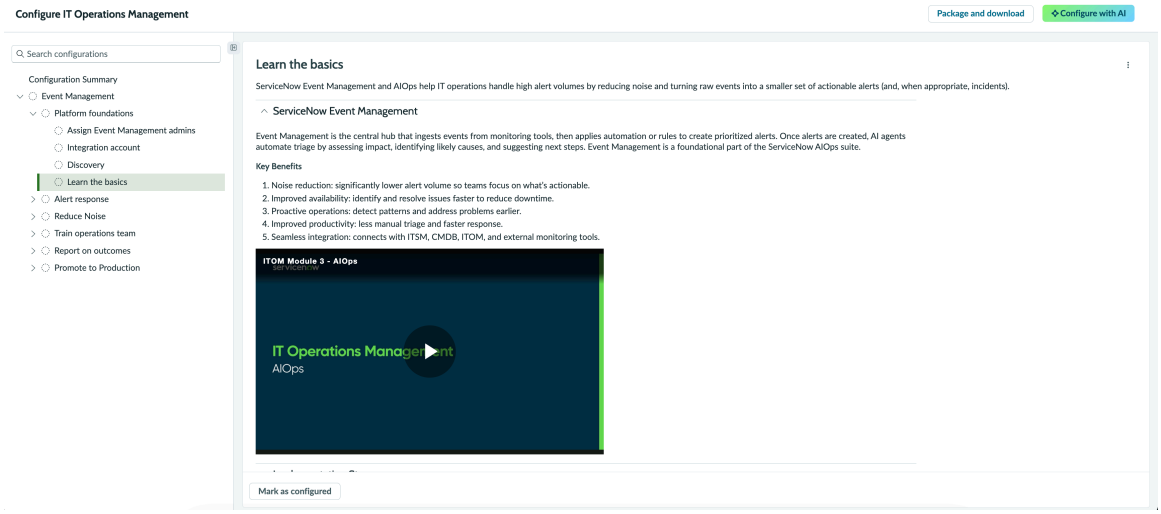

Note: The **Configure with AI** UI action appears only after you install both the ServiceNow Otto for ITOM and ServiceNow Otto plugins. This action lets you set up AI-powered alert handling.

Access Configuration Console from AIOps Configuration Center

Access ITOM Configuration Console to set up Event Management and configure AIOps capabilities.

Before you begin

Verify that the ITOM AIOps and ServiceNow Otto for ITOM plugins are installed.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under **Setup** > **Get started** section, select **Event Management Onboarding**.

   The Configure IT Operations Management page opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

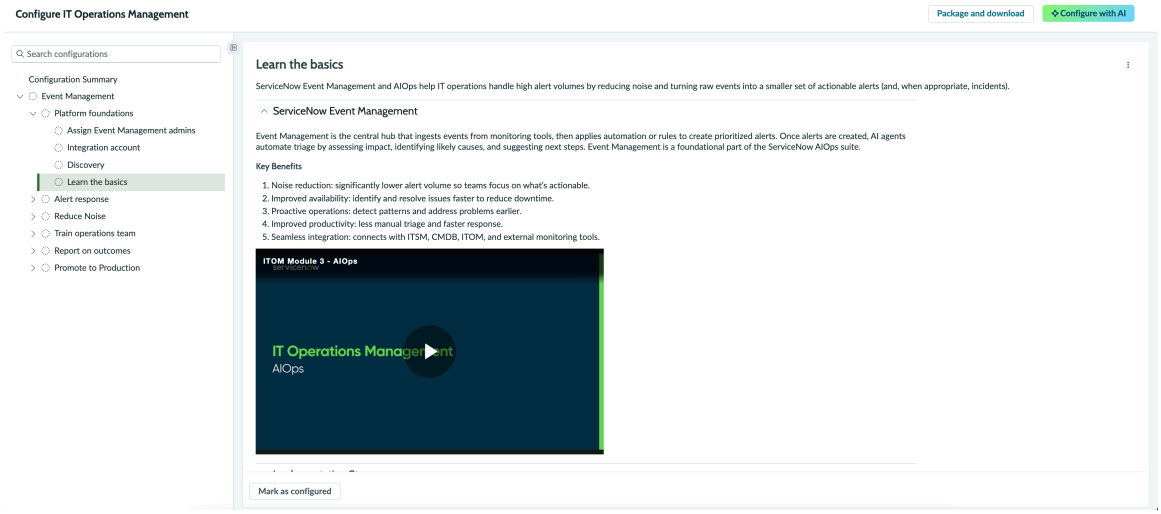

Note: The **Configure with Otto** UI action appears only after you install both the ServiceNow Otto for ITOM and ServiceNow Otto plugins. This action lets you set up AI-powered alert handling.

Assign Event Management admin role to group

Assign the evt_mgmt_admin role to users who will be in charge of configuration and operational control.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin

About this task

The evt_mgmt_admin role provides groups with the necessary permissions to configure and manage Event Management settings, including alert processing, incident creation, and operational controls.

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **Assign Event Management admins**.
4. Select **Add group**.
   By default, the **Assign role to existing group** is selected.
5. **Optional:** If you want to create your group, select **Create new group** and perform the following steps:
   a. In the **Group name** field, enter the name of the group.
   b. In the **Description** field, enter a description.
   c. In the **Users** field, select the users to add to the group.
   d. Select **Save**.
6. In the **Group** field, select the group to assign the evt_mgmt_admin role to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Select **Save**.
   The role is assigned to the selected group.
8. To complete the setup, select **Mark as configured**.

Delete group from Event Management admin role

Delete group from the evt_mgmt_admin role when you no longer need them.

Before you begin

Verify that you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **Assign Event Management admins**.
4. Select and open the group you want to delete.
5. On the Edit group page, select **Delete Group**.
6. In the confirmation window, select **Delete Group**.
   The group is removed from the list of Assign Event Management admins page.

Delete admin role from group

Delete the evt_mgmt_admin role from a group when you no longer need them.

Before you begin

Verify that you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **Assign Event Management admins**.
4. Select and open the group from which you want to remove the evt_mgmt_admin role.
5. On the Edit group page, select **Remove role from group**.
6. In the confirmation window, select **Remove role**.
   The evt_mgmt_admin role is removed from the specific group.

Create integration account

Create a dedicated account with the evt_mgmt_integration role for third-party monitoring systems to push events to ServiceNow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin

About this task
A dedicated integration account ensures third-party monitoring systems have secure, stable, and appropriately scoped access to push events into ServiceNow without relying on personal user accounts.

Note: Create an integration account user only if you plan to use basic authentication. This step is not required if you're using API keys or OAuth. If you prefer those methods, refer to the corresponding documentation for setup instructions.

Procedure
1. Navigate to **Configuration Summary** > **Event Management** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **Integration account**.
4. Select **Create account**.
   The Create integration account page opens.
5. In the **User ID** field, enter the account id or the name of the account.
6. In the **Password** field, enter the password for the integration account.
7. Select **Create account**.
   The integration account is created.
8. To complete the setup, select **Mark as configured**.

Install Discovery

Configure Discovery to lay the groundwork for AIOps success by establishing a key foundation for your team.

Before you begin
Verify that the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins are installed.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin

About this task
Discovery is the source of truth for IT operations, automatically identifying, mapping, and cataloging infrastructure components and relationships to enable faster root cause analysis and reduced operational noise.

Procedure
1. Navigate to **Configuration Summary** > **Event Management** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **Discovery**.
4. Select **Configure Discovery**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the Discovery plugin is installed, the system takes you to the Assign Discovery Admins page where you can assign users to the discovery admin role (discovery_admin). If the Discovery plugin is not installed, you're directed to the IT Operations Management page. Install the ITOM Visibility plugin, which also installs the Discovery plugin automatically.

5. To complete the setup, select **Mark as configured**.

Learn the basics

ServiceNow® Event Management and AIOps help IT operations teams manage high alert volumes by reducing noise and turning raw events into prioritized, actionable alerts. This section covers core concepts, implementation stages, CMDB considerations, prerequisites for setup, and resources to help you get started quickly and confidently.

ServiceNow Event Management

Event Management is the central hub that ingests events from monitoring tools, applies automation and rules to create prioritized alerts, and uses AI agents to automate triage. It reduces noise, improves availability, and seamlessly integrates with ITSM, CMDB, ITOM, and external monitoring tools — making it a foundational part of the ServiceNow AIOps suite.

Implementation Stages

Deploy Event Management using a phased approach, starting with a single high-value service and progressing through four stages — basic alert response, noise reduction, faster MTTR, and proactive operations.

Should I Wait for CMDB

Deploy Event Management and improve your CMDB in parallel to show value quickly. Modern observability tools already provide useful event metadata. Use Alert Automations to improve alert deduplication and grouping without waiting for full CMDB population. Then you can strengthen results by adding Discovery and Service Mapping over time to improve CI/service context.

Prerequisites for Setup

Before implementation, gather credentials for your monitoring tools and submit any necessary access requests to Network and Security teams to avoid delays. If your setup requires pull-based integrations, Agent Client Collector, or Health Log Analytics, a MID Server will also need to be deployed and configured in advance.

How to get help

A curated set of resources is available to support your implementation journey — including training courses, workshop presentations, implementation guides, official documentation, community forums, and video content. You can also reach out to your account team or raise a ticket through the ServiceNow Support Portal at any time.

- Event Management Fundamentals – ServiceNow University course to build a strong foundation in setup, operations, and optimization.
- Event Management Workshop Presentation - slide deck for onboarding and staff training
- Event Management Implementation Guide - Detailed guidance to implement Event Management
- Event Management Documentation - official reference for configuration, administration, and troubleshooting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Event Management Forum – Post questions, share knowledge, and join discussions.
- ServiceNow YouTube Channel - Official video content, demos, and event recordings.

Install integrations

Access Integrations Launchpad to configure and install integrations for Event Management monitoring tools.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Alert response**.
2. Expand **Alert response**.
3. Select **Install integrations**.
4. Select **Install integrations in Integration Launchpad**.
   The Integrations Launchpad page opens.
5. Follow the steps in Integrations Launchpad in Service Operations Workspace for ITOM that apply to your integration installation.
6. To complete the setup, select **Mark as configured**.

Review alerts

Access the Express List interface to review and manage alerts in your Event Management system.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you are in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

About this task

When reviewing alerts, validate the following to ensure your event pipeline is healthy and alerts are actionable:

- Check CI binding. Alerts should populate the cmdb_ci field. If there is no CI, your binding rules need tuning.
- Verify severity spread. If everything lands as Critical or Info, your mapping needs work.
- Confirm that clear events close alerts. Trigger a condition, then resolve it at the source. The alert should auto-close.
- Spot-check key fields. Ensure node, source, resource, and message_key are meaningfully populated (not blank or generic).
- Review Error events. Check em_event status = Error for parsing or mapping issues.
- Check the alert count. Make sure it matches in both your source system and in ServiceNow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Alert response**.
2. Expand **Alert response**.
3. Select **Review alerts**.
4. Select **Review alerts in Express List**.
   Express List opens with a list of active alerts.

   Note: To get more information on alerts so you can more efficiently monitor systems and services, resolve alerts, evaluate the alert impact, track issues, and report incidents, see Express List in the Service Operations Workspace for ITOM.

5. To complete the setup, select **Mark as configured**.

Alert enrichment automations

Configure alert enrichment rules to automatically add context and metadata to incoming alerts for better analysis and response.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

About this task

Alert enrichment rules automatically enhance incoming alerts with additional context such as configuration item details, service mappings, and operational metadata. This enrichment helps operations teams understand the business impact and take appropriate action more quickly.

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Alert response**.
2. Expand **Alert response**.
3. Select **Enrich alerts**.
   A list of enrich alert automations appear.
4. Select **+Create automation**.
   The Enrich alerts page opens.
5. Follow the steps in Create Enrich automation.
6. To complete the setup, select **Mark as configured**.

Create incidents

Configure automation rules to automatically create incidents from alerts that require immediate attention and resolution.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you are in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task
Incident creation rules automatically generate incident records from alerts that meet specific criteria. This automation ensures that critical issues are promptly escalated to the appropriate teams for resolution while reducing manual effort in the incident management process.

Procedure
1. Navigate to **Configuration Summary** > **Event Management** > **Alert response**.
2. Expand **Alert response**.
3. Select **Create incidents**.
   A list of respond alert automations appears.
4. Select **+Create automation**.
   The Respond to alert page opens.
5. Follow the steps in Create Respond automation.
6. To complete the setup, select **Mark as configured**.

Alert ignore automations

Configure ignore rules to automatically suppress alerts that are not actionable or relevant to your operations team.

Before you begin
Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you are in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

About this task
Ignore rules help reduce alert noise by automatically suppressing alerts that match specific criteria. This allows operations teams to focus on actionable alerts while filtering out known non-critical events, maintenance notifications, or false positives.

Procedure
1. Navigate to **Configuration Summary** > **Event Management** > **Reduce Noise**.
2. Expand **Reduce Noise**.
3. Select **Ignore alerts**.
   A list of ignore alert automations appears.
4. Select **+Create automation**.
   The Ignore alerts page opens.
5. Follow the steps in Create Ignore automation.
6. To complete the setup, select **Mark as configured**.

Alert group automations

Configure grouping rules to automatically consolidate related alerts into single actionable items for more efficient incident management.

Before you begin
Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

About this task

Alert grouping rules automatically combine related alerts that represent the same underlying issue or event. This reduces alert volume and helps operations teams focus on resolving the root cause rather than managing multiple individual alerts for the same problem.

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Reduce Noise**.
2. Expand **Reduce Noise**.
3. Select **Group alerts**.
   A list of group alert automations appears.
4. Select **+Create automation**.
   The Group alerts page opens.

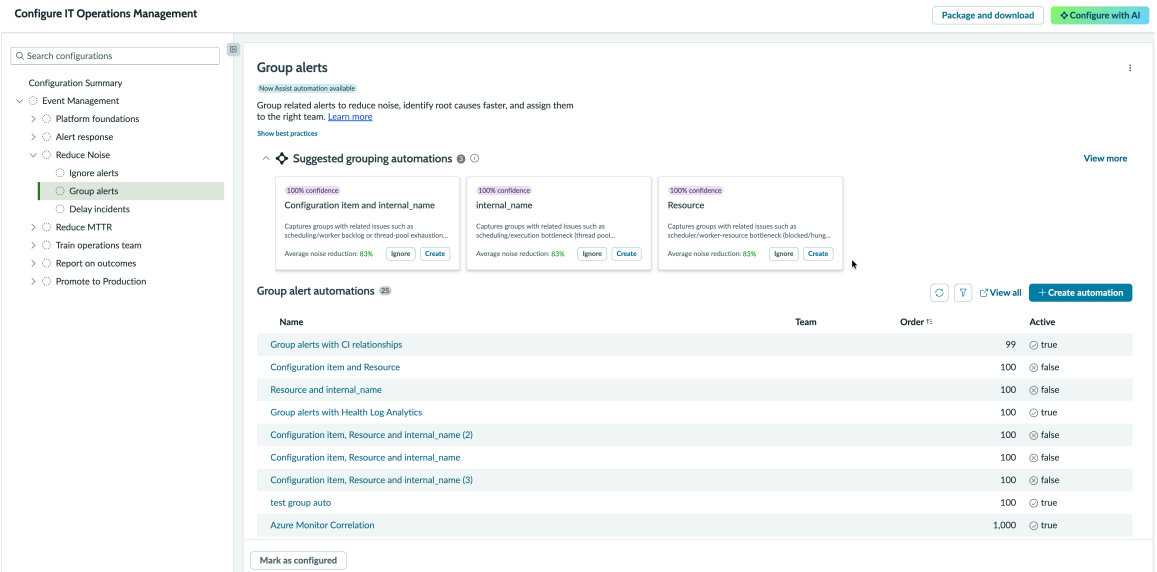

You can use the suggested grouping automations to create your grouping automation. If you have the ServiceNow Otto for IT Operations Management (ITOM) plugin installed, select **Configure with AI** from the top-right of the page and ask the ServiceNow Otto chatbot to create an automation. The created automation appears in the list of automations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

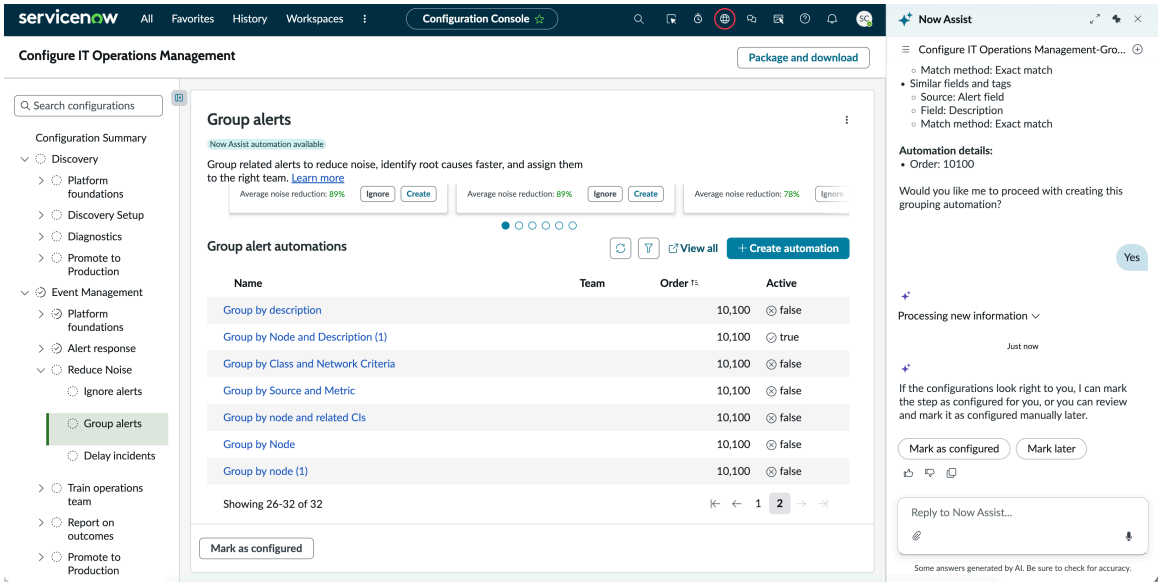

For more information on group automation, see Create Group automation.

5. To complete the setup, select **Mark as configured**.

Delay incidents

Configure delay rules to postpone incident creation for alerts that may resolve automatically or represent transient conditions.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you're in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

About this task

Delay incident rules help reduce unnecessary incident creation by waiting for a specified period before generating incidents from alerts. This is particularly useful for flapping conditions or transient issues that may resolve automatically, preventing ticket creation for self-resolving problems.

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Reduce Noise**.
2. Expand **Reduce Noise**.
3. Select **Delay incidents**.
   A list of respond alert automations appears.
4. To create an automation and set a wait time before incident creation, select **+Create automation**.
   The Respond to alerts page opens.
5. Follow the steps in Create Respond automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: To modify the wait time in an existing automation, open the automation and update the wait time under **Create incident and other response actions**. Ensure you select **Create Incident (Advanced)** — the delay configuration form is only available through this option.

6. To complete the setup, select **Mark as configured**.

Run remediation actions

Run remediation actions automatically or trigger them manually from the Express List to speed up alert investigation and resolution.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you are in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

About this task

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Reduce MTTR**.
2. Expand **Reduce MTTR**.
3. Select **Remediation actions**.
   A list of respond alert automations appears where the **Response subflow** field in the Respond alert automation is set to anything other than **Create incident (advanced)**.
4. Select **+Create automation**.
5. Follow the steps in Create Respond automation.
6. To complete the setup, select **Mark as configured**.

Training operations team

Get your L1/L2 operators up to speed with Event Management. This section walks through the key concepts, tools, and workflows they need to confidently triage alerts, understand service impact, and take action from day one.

Event Management Operator Essentials course

Designed for IT operations professionals, this course covers the full alert lifecycle in Event Management. You learn to navigate the Service Operations Workspace and manage alerts using Express List and AIOps Dashboards. You also use Alert Automation and AI to reduce noise and accelerate resolution. You gain hands-on experience with the four-stage automation pipeline (Ignore, Enrich, Group, and Respond) and earn a digital badge upon completion. Select the **Event Management Operator Essentials** tile to open and start the course.

Service Operations Workspace

An integrated environment where IT operations and service management converge — alerts, incidents, services, and operational context all in one place. Key highlights include a shift-start landing page for quick situational awareness and shared context (CIs, changes, related records) visible without tool-switching. Playbooks and remediation actions are built in, and cross-team alignment makes acknowledgements and escalations immediately visible to teammates.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Express List

A live, continuously updating alert queue designed for active triage. Operators are coached to prioritize by business impact over raw severity, work alert groups rather than individual alerts, and use AI alert analysis for plain-language summaries. They also acknowledge purposefully, escalate quickly when stuck, and document work notes to avoid repeated investigation by the next tier.

Service Dashboard

The business-impact lens on what's firing. While the Express List shows what is alerting, the Service Dashboard shows what it means for the business. It surfaces service health tiles, business criticality ratings, CI topology maps, and service dependencies. This transforms 'I have 15 alerts' into actionable service-level awareness for operators and leadership alike.

Dashboards

Access AIOps dashboards in Service Operations Workspace to monitor the health, performance, and availability of your IT resources in real-time.

To get to the AIOps dashboards in Service Operation Workspace page, navigate to **Configuration Summary** > **Event Management** > **Report on outcomes** > **Dashboards**.

AIOps dashboards

The Dashboards section surfaces three purpose-built AIOps dashboards inside Service Operations Workspace, each serving a different reporting need:

- AIOps 360° Overview: A comprehensive view of AIOps activities and their overall impact, including time saved through automation and grouping, workload savings, and incident prevention rates.
- AIOps Value Realization: Focused on business value metrics, helping stakeholders see the tangible ROI of AIOps adoption.
- AIOps Operational: A monitoring view of AIOps system behavior and operations, useful for admins tracking how the system is performing day to day.

To open a dashboard, select its tile.

Learn how to deploy Event Management to production

Understand the process and considerations for deploying Event Management configurations from development to production environments using update sets.

Deploying Event Management configurations to production requires careful planning and execution to maintain system stability and performance. Update sets provide a controlled method for moving configurations between environments.

Deployment planning

Before deploying to production, verify that all configurations work correctly in your development or staging environment. Test alert automation rules, integration accounts, and role assignments to confirm they function as expected.

Update set process

Update sets capture configuration changes made in your development environment and allow you to transfer them to production. This includes Event Management automation rules, user roles, integration settings, and custom configurations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The deployment process automatically records changes as update sets, so you don't need to create them manually. To deploy, export the recorded update set and import it into your production environment. Always review the update set contents before deployment.

Update set: included and excluded components

Category	Components
Included in update sets	• Custom tables and fields • Form and list layouts, form sections • Business rules, client scripts, UI actions, UI policies, UI pages, UI macros, and UI scripts • Script includes and workflows • System properties (except those tagged as Private) • Reports and views • Access Control Lists (ACLs)
Not included in update sets	• Transactional data records (incidents, problems, change requests, knowledge articles) • Users, groups, and group memberships • Configuration Items (CIs) and CMDB data • Scheduled jobs and schedulesHomepages (unless manually added) • Email and notification logs • Number counter records (sys_number_counter) • System properties tagged as Private • Plugin-specific elements and business rules in deactivated applications

Deployment considerations

- Schedule deployments during maintenance windows to minimize impact on operations
- Back up your production environment before applying updates
- Test automation rules in production with a limited scope before full activation
- Verify that integration accounts and credentials are properly configured in production
- Confirm that all required plugins are installed in the production environment
- Review role assignments to confirm they align with your production security model

Post-deployment verification

After deployment, monitor your Event Management system to confirm that alerts are processed correctly and automation rules function as expected. Check integration connections and verify that incidents are created appropriately.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Additional resources

For detailed guidance on update sets and deployment procedures, refer to the ServiceNow documentation on application deployment and change management processes.

Assign Event Management operator role to group

Assign the evt_mgmt_operator role to groups to enable operations teams to work with alerts and manage Event Management workflows.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you are in the Configure IT Operations Management page.

Role required: evt_mgmt_admin or evt_team_operator

About this task

The evt_mgmt_operator role provides the necessary permissions for operations teams to view alerts, manage incidents, and use Event Management tools. Assign this role to groups that contain your L1 and L2 operations staff who will be working with alerts on a daily basis.

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Platform foundations**.
2. Expand **Platform foundations**.
3. Select **Assign Event Management operators**.
4. Select **Add group**.
   By default, the **Assign role to existing group** is selected.
5. **Optional:** If you want to create your group, select **Create new group** and perform the following steps:
   a. In the **Group name** field, enter the name of the group.
   b. In the **Description** field, enter a description.
   c. In the **Users** field, select the users to add to the group.
   d. Select **Save**.
6. In the **Group** field, select the group to assign the evt_mgmt_operator role to.
7. Select **Save**.
   The role is assigned to the selected group.
8. To complete the setup, select **Mark as configured**.

Assign AIOps NOC managers to group

Assign the evt_aiops_manager role to the group that is responsible for managing NOC operators and AIOps AI specialists.

Before you begin

Verify you have installed the ITOM AIOps and ServiceNow Otto for IT Operations Management (ITOM) plugins.

Ensure you are in the Configure IT Operations Management page.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The evt_aiops_manager role provides the permissions that operations leaders need to oversee alert handling, manage team assignments, and configure Event Management tools. Assign this role to groups that contain your operations managers and team leads who supervise L1 and L2 staff and own alert-handling outcomes.

Procedure

1. Navigate to **Configuration Summary** > **Event Management** > **Promote to Production**.
2. Expand **Promote to Production**.
3. Select **Assign AIOps NOC managers**.
4. Select **Add group**.
   By default, the **Assign role to existing group** is selected.
5. **Optional:** If you want to create your group, select **Create new group** and perform the following steps:
   a. In the **Group name** field, enter the name of the group.
   b. In the **Description** field, enter a description.
   c. In the **Users** field, select the users to add to the group.
   d. Select **Save**.
6. In the **Group** field, select the group to assign the evt_aiops_manager role to.
7. Select **Save**.
   The role is assigned to the selected group.
8. To complete the setup, select **Mark as configured**.

Request Event Management

AIOps Experience plugin (`sn_sow_aiops`) requires a separate subscription and must be activated by ServiceNow personnel. This plugin includes the required Event Management components.

Before you begin

Role required: admin

About this task

To purchase a subscription, contact your ServiceNow account manager. The account manager can arrange to have the plugin activated on your organization's production and subproduction instances, generally within a few days.

If you don't have an account manager, decide to delay activation after purchase, or want to evaluate the product on a subproduction instance without charge, follow these steps.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Install Event Management

Retrieve the most updated apps for the Event Management application (com.glideapp.itom.snac) in the ServiceNow® Store. Periodically check the ServiceNow® Store for new app versions.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Role required: admin

About this task

The following items are installed with Event Management:

- Roles
- Scheduled jobs
- Tables

For more information, see Components installed with Event Management.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the ITOM AIOps application using the filter criteria and search bar.

   If you cannot find the application, request it from the ServiceNow Store.
3. In the Application installation dialog box, review the application dependencies.

   Dependent plugins and applications are listed if they will be installed, are currently installed, or need to be installed. If any plugins or applications need to be installed, you must install them before you can install Event Management.
4. **Optional:** If demo data is available and you want to install it, select the **Load demo data** check box.

   (Optional) Demo data comprises sample records that describe application features for common use cases. Load demo data when you first install the application on a development or test instance.

   Important: If you don't load the demo data during installation, it's unavailable to load later.
5. Select **Install**.

Enhance Event Management performance

The Event Management Performance Accelerator plugin helps maintain Event Management performance at a high level. This plugin is optional.

Before you begin
Role required: admin

About this task
The *em_alert_history* and *em_impact_status* tables can grow very large, negatively impacting Event Management performance. The Event Management Performance Accelerator plugin includes fix scripts that populate missing data, update key fields to maintain data accuracy, and create indexes to enhance performance. These actions help optimize large datasets and reduce the impact on query speed and overall system responsiveness.

When upgrading from an earlier version with more than 5 million records in either the *em_alert_history* or *em_impact_status* tables, you must activate this plugin manually. When there are fewer than 5 million records, it activates automatically. You can customize this number by modifying the *evt_mgmt.plugin_activation.table_max_size* property.

To evaluate the Event Management Performance Accelerator plugin on a subproduction instance, install the plugin in Application Manager. For instructions, see Install an app or plugin.

Procedure

Activate the Event Management Performance Accelerator plugin.

Event Management setup

After activating Event Management, set it up to receive and process events, and generate and analyze alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event Management setup without using guided setup

Set up Event Management by completing these tasks in the following order:

1. Configure a MID Server to receive and process events via the MID Server.
2. Configure the MID Web Server extension.
3. Configure Configure Event Management connectors.
4. Configure event field mappings and Binding alerts to CIs to manage alert generation.
5. Alert management rules for resolving alerts, perform, and CI remediation for alert management.
6. Request Service Mapping and get a top-down discovery to receive CI relationships for software and hardware.
7. Configure impact calculation for services, to establish priority for alert resolution.
8. Configure alert groups to consolidate related alerts.
9. Configure the Predictive Intelligence plugin (com.glide.platform_ml) to enable machine learning and finding similar alerts.
10. Configure any other general tasks that appear in this section as appropriate.

Note: Event Management does not support creating incidents on remote instances.

Event Management configuration using Setup Hub

ITOM configuration console provides a sequence of tasks to help you configure Event Management. For more information, see Configure Event Management using ServiceNow Otto for Setup.

Event Management during a platform upgrade

During a platform upgrade Event Management jobs whose **Upgrade safe** flag is marked as `true` remain running.

During an upgrade of the platform, the Event Management connectors are working and continue to retrieve events. The events are being processed, transformed to an alert using event rules and if there is a suitable alert management rule, a task is created. However, the impact calculation is not supported during the platform upgrade. When applying configuration changes or upgrades, all impacted MID Servers restart.

Note: Event Management jobs that started running before the platform upgrade commenced continue to run during the upgrade.

The following Event Management jobs remain running as their **Upgrade safe** flag is marked as `true` :

- Event Management - Connector execution
- Event Management - Update stuck connect
- Event Management - Alert Priority Queue
- Event Management - Close flapping alerts
- Event Management - Close threshold alert
- Event Management - Evaluate Scoped Alert Rules Management0
- Event Management - Maintenance Calculator
- Event Management - Process events

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

During the platform upgrade, all other Event Management jobs wait for the platform upgrade to finish.

Note: During an Event Management plugin upgrade, all Event Management jobs do not work.

MID Web Server

The MID Web Server is part of the common infrastructure of the MID Server.

The MID Web Server extension provides options for authentication and data security, and runs for as long as it is enabled. The extension starts a web server on the MID Server to handle web requests from external systems. The raw data is pushed to the extension, either from a client or with a customized script. The data is collected by the MID Server and is transmitted to the ServiceNow instance.

The MID Web Server extension supports MID Server clusters that are configured for failover. When selecting a MID Server cluster option, an algorithm determines which MID Server in the cluster runs the extension. The extension can run on only one MID Server. If the primary MID Server goes down, the extension automatically starts to run on the secondary MID Server.

Related topics

Configure the MID Web Server extension

Configure a secure MID Web Server extension

Configure the MID Web Server extension

The MID Web Server is a MID Server extension that enables developing REST APIs to send events and metrics to the MID Server. The extension is leveraged by other MID Server extensions, such as Metric Intelligence, MID WebService Event Listener, and the Agent Client Collector websocket endpoint extension.

Before you begin

- Deploy and start a MID Server.
- Enable the MID Web Server [com.snc.sa.mid.webserver] plugin.

Role required: agent_admin.

About this task

Agent Client Collector supports the following authentication strategies:

- **Keybased**: For backwards (pre-Brazil release) compatibility only. See Configure key-based MID Web Server authentication.
- **API key**: See Configure MID Web Server API key authentication.
- **mTLS**: See MID Web Server and agent mTLS Authentication.

Procedure

1. Navigate to **All** > **MID Server** > **Extensions** > **MID Web Server**.
2. In the MID Web Server Contexts list, click **New**.
3. Fill in the values on the Event Management MID Web Server extension form.
4. Click **Save** to save the data.
5. Restart the MID Web Server extension to apply your changes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Navigate to **All** > **MID Server** > **Extensions** > **MID Web Server** and select the MID Server you just created.
7. In the **Related Links** section, click **Start** to start the extension.

Configure a secure MID Web Server extension

Configure a TLS listener for extra security and encryption of data transferred to and from the MID Web Server extension. Access a private key and certificate (or certificate chain) from a Java keystore. The keystore is used during the TLS handshake where the MID Web Server and the client acknowledge each other.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **MID Server** > **Extensions** > **MID Web Server**.
2. Select **New**.
3. In the **Authentication Type** field, select an authentication type to be used.
   - Keybased – See Configure key-based MID Web Server authentication
   - Basic – Standard username/password authentication
   - API key – See Configure MID Web Server API key authentication
   - mTLS – See MID Web Server and agent mTLS Authentication
4. Select either the MID unified keystore or the web server keystore.
   The MID Web Server unified keystore is the most efficient option. To enable backward compatibility and store keypairs in an additional keystore, use the Web server keystore. The Web server keystore allows you to use the MID unified keystore as well.

Option	Description
MID unified keystore	**a.** Select the **Use MID Unified Keystore** check box **b.** Enter an alias name in the **Keystore Certificate Alias** field or leave the field empty to use the keystore's default alias. For more information, see the **Install custom certificates in the MID Server unified key store** section in MID Server unified key **Important:** Leaving the **Keystore Certificate Alias** field empty causes connecting clients to skip certificate verification. Instead, the MID Server uses its own self-signed certificate.
Web server keystore	**a.** Clear the **Use MID Unified Keystore** check box. **b.** Specify an alias name in the **Keystore Certificate Alias** field or leave the field empty if you have a single alias in that keystore.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	c. If the password to the private key of the re quired alias is different from the keystore password, specify the alias in the **Keystore Key Password** field. **d.** Enter the keystore password in the **Key store Password** field.

Related topics

Connect the agent to the MID Server using mTLS

Configure key-based MID Web Server authentication

Provide added security to your MID Web Server extension by using key-based authentication. Generate an authentication token to be sent in the Authorization header of incoming client requests.

Before you begin

Note: This procedure is only for compatibility with releases prior to Brazil. For details on the procedure in the Brazil release for configuring the MID Web Server, see Configure the MID Web Server extension.

- Deploy and start a MID Server.
- Configure a MID Web Server extension and select Keybased as the authentication type. For details, see Configure the MID Web Server extension.

Role required: agent_admin

Procedure

1. Using a script program, create a token by constructing a string using defined elements of the HTTP/HTTPS request (HTTP verb, content type header, and request path, received from the client accessing the MID Web Server extension endpoints).
2. Create an HMAC (Hash Message Authentication Code) of the string by signing the generated string with the auto-generated secret key that is displayed in the **Secret Key**.
   This key is unique per context.
3. Provide information to send this authentication token in the request Authorization header.

Item	Value
Path to a web service API for sending raw data	URL format: `https://<MID Server IP address>:<port number>/api/mid/sa/metrics` Use a port number that matches one of the port numbers set up in the Web Server extension. Example: `http://10.10.10.10:8097/api/mid/sa/metrics`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Item	Value
Request type	POST
Date format	`yyyy-MM-dd'T'HH:mm:ss.SSS'Z'` For example: 2016-06-08T20:54:58.917Z
Content-Type	application/json

Use the following request elements to generate the required string: HTTP-Verb, Content-Type, Date, and request path. Specify these elements and place them in this order:

- HTTP-Verb + "\n" +
- Content-Type + "\n" +
- Date + "\n" +
- Request-Path

For this example, the request string is:

`POST\napplication/json\n2016-06-08T20:54:58.917Z\n/api/mid/sa/metrics`

For the timestamp requirement, a valid timestamp that uses HTTP date header is required for authenticating the request. Ensure that the timestamp is within 15 minutes of the MID Server.

Example: How to generate the HMAC of the string that uses defined elements of the HTTP/HTTPS request, using Java

```
package sample;
import com.glide.util;
import java.security.SignatureException;

import javax.crypto.Mac;
import javax.crypto.spec.SecretKeySpec;

public class AuthUtil {

private static final String HMAC_SHA1_ALGORITHM = "HmacSHA1";

/***
 * Generates base64-encode the HMAC(Hash Message Authentication
 Code) of input data
 *
 * @param data
 * @param key
 * @return
 * @throws java.security.SignatureException
 */
public static String signData(String data, String key) throws
 java.security.SignatureException {
 String result;
 try {
  // get an hmac_sha1 key from the raw key bytes
  SecretKeySpec signingKey = new SecretKeySpec(key.getBytes(),
 HMAC_SHA1_ALGORITHM);
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
  // create hmac_sha1 Mac instance and initialize with the
 signing key
  Mac = Mac.getInstance(HMAC_SHA1_ALGORITHM);
  mac.init(signingKey);

  // compute the hmac on input data bytes
  byte[] rawHmac = mac.doFinal(data.getBytes("UTF-8"));

  // base64-encode the hmac
  result = Base64.encode(rawHmac);

 } catch (Exception e) {
  throw new SignatureException("Failed to generate HMAC : " +
 e.getMessage());
 }
 return result;
}
}
```

Configure MID Web Server API key authentication

Authenticate incoming requests from clients to the MID Web Server extension using API key authentication. API authentication is a secure and simple way to authenticate your request. You can create or modify a MID Web Server API key.

Before you begin

Create a MID Web Server extension to use API Key authentication, as described in Configure the MID Web Server extension. Ensure that you select **API Key** in the extension's **Authentication Type** field.

Role required: event_mgmt_admin

About this task

By default, the maximum number of API keys you can create on an instance is 10. To change this number, modify the *mid.webserver.max.api.keys* property value on the System Properties page.

When working in the global domain while Domain Separation is enabled, API keys are available to all extensions, regardless of their domain. When working in a specific domain, API keys are available to extensions in that domain, its parent domains, and the global domain.

You can configure API keys to expire on a specified date.

When creating or updating an extension to use API key authentication, the system checks for an available API key. If none is available, the system creates one.

Procedure

1. Navigate to **All** > **MID Server** > **Profiles and Deployments** > **MID Web Server API Key**.
2. Either create a new API key or modify an existing API key.
   - To create a new API key:
     a. Select **New**.
     b. On the MID Web Server API Key form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MID Web Server API Key form

Field	Description
Name	Name for the API key.
Active	Option for activating the API key.
Authentication key	Read-only encoded value of the authentication key. No value appears until after you configure the key.
Expires	Option for setting an expiration date for the API key. Expired, deactivated, or deleted API keys are not accepted by the web server extension.

c. Select **Submit** to save the API key.

The API key appears on the Mid Web Server API Keys page.

- To modify an existing API key:

a. Navigate to **MID Server** > **Profiles and Deployments** > **MID Web Server API Key**.

b. Select an API key on the page.

c. In the **Related Links** section, select **View API Key**.

d. Copy the API key that appears in the MID Web Server dialog box.

When connecting to a MID Web Server extension configured with API Key authentication, place an API key in the Authorization header. The API key must be one that the extension has access to. Use the following format:

`Key <API_KEY>`

where `<API_KEY>` is the value of the API key displayed after selecting **View API Key**.

e. Select **Submit**.

MID Web Server and agent mTLS Authentication

Mutual authentication using the Transport Layer Security protocol (mTLS) is a secure, certificate-based authentication scheme. With mTLS, the server (the MID Web Server extension) and the client (the agent) authenticate each other.

TLS authentication occurs when a client (such as a web browser) authenticates the server it's accessing by validating the certificate presented by the server. The client verifies that the certificate (or certificate chain) is signed by a trusted certificate authority (CA). The client must have access to a bundle of CA certificates, which it uses to validate the server certificate.

When mTLS is configured, the client presents a certificate or certificate chain to the server. The server authenticates the client in the same way that the client authenticates the server (mutual authentication).

In addition to verifying the certificate signature, some TLS implementations also check the revocation status of the certificate. In these cases, the Online Certificate Status Protocol (OCSP) is used. In this protocol, the verifying party sends a request to the OCSP responder set up by the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

CA that issued the presented certificate. (The responder's address is typically embedded in the CA-issued certificates). The response indicates whether the presented certificate is revoked.

Before enabling mTLS, you must first create keys and certificates, install the .pem file on your MID Server, and connect the agent to your MID Server.

Related topics

Configure the MID Web Server extension

Configure mTLS authentication for a MID Web Server

Create keys and certificates

Create keys and certificates in your root directory to enable Transport Layer Security (TLS) setup. TLS setup is required before you can configure mTLS on the MID Web Server and agent.

Before you begin

- Ensure that there are no keys installed in the MID unified keystore by running the following command:

  ```
  bin/scripts/manage-certificates.(sh/bat) -l
  ```

  When there are no keys installed, the output is: `defaultsecuritypairhandle`

- Invalidate your MID Server.

  If additional keys exist in the output, reinstall those keys after invalidating the MID Server.

- Ensure that the MID Server is connected to the instance.
- Select a directory in which you want to create certificates, to be called the *root* directory.
- For Windows environments only: Install Win32/Win64 OpenSSL with all dependencies and ensure it is on the system PATH.

Note: This procedure includes commands for both CentOS 7/Linux and Windows Server environments. Select the commands relevant for your host operating system. If working with another Linux distribution, adapt the commands as needed for your specific OS.

Role required: agent_client_collector_admin

Procedure

1. In a Linux environment:

   a. Create subdirectories for your certificates in your root directory.

   ```
   mkdir -p labca labmid labacc;
   ```

   b. Generate the EC key pair.

   ```
   openssl ecparam -list_curves;
   openssl ecparam -out labca/ec-labcakey.pem -name prime256v1
    -genkey;
   ls labca/;
   ```

   The generated output is the `ec-labcakey.pem` file.

   c. Create the root CA certificate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
openssl ecparam -in labca/ec-labcakey.pem -text -noout;
openssl req -x509 -new -nodes -key labca/ec-labcakey.pem
 -sha512 -days 365 -out labca/labcacert.pem -subj
 "/C=<country>/ST=<state>/L=<location>/O=<organization>
 Lab/OU=<organization unit>/CN=<cn abbreviation>";
openssl verify labca/labcacert.pem;
```

Generated output:

- `labca/labcacert.pem: C = <country>, ST = <state>, L = <location>, O = <organization> Lab, OU = <organization unit>, CN = <cn abbreviation>`
- `OK`

d. Import the CA certificate into the system trust store.

```
sudo cp -a
 labca/labcacert.pem /etc/pki/ca-trust/source/anchors/;
sudo update-ca-trust extract;
openssl verify labca/labcacert.pem
```

Generated output: `labca/labcacert.pem: OK`

e. Obtain the fully qualified domain name (FQDN) for the host.

```
hostname --all-fqdns;
```

f. Generate the server certificate signing request (CSR) and sign it with the CA.

Enter the FQDN as the `<hostname>` value in the following commands. If more than one FQDN value is returned, use the value with the following format: `hostname.domain.domain.com`.

```
openssl req -new -newkey rsa:4096
 -keyout labmid/rsa-labmidkey.pem
 -sha512 -nodes -out labmid/mid.csr -subj
 "/C=<country>/ST=<state>/L=<location>/O=<organization>/OU=<o
rganization unit>/CN=<hostname>";
openssl x509 -req -days 365 -in labmid/mid.csr -CA
 labca/labcacert.pem -CAkey labca/ec-labcakey.pem
 -CAcreateserial -extensions SAN -extfile
 <(cat /etc/pki/tls/openssl.cnf <(printf
 "\n[SAN]\nsubjectAltName=DNS:<hostname>")) -out
 labmid/mid.crt;
```

Generated output:

```
Signature ok
subject=/C=<country>/ST=<state>/L=<location>/O=<organizatio
n>/OU=<organization unit>/CN=<mid server host fqdn>
Getting CA Private Key
```

g. Combine the key and certificate files into one file, named `mid.pem`.

```
cat labmid/rsa-labmidkey.pem labmid/mid.crt > labmid/mid.pem;
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. In a Windows environment:

   a. Create subdirectories for your certificates in your root directory.

   ```
   mkdir labca labmid labacc
   ```

   b. Generate the EC key pair.

   ```
   openssl ecparam -list_curves
   openssl ecparam -out labca\ec-labcakey.pem -name prime256v1
    -genkey
   dir labca
   ```

   The generated output is the `ec-labcakey.pem` file.

   c. Create the root CA certificate.

   ```
   openssl ecparam -in labca\ec-labcakey.pem -text -noout
   openssl req -x509 -new -nodes -key labca\ec-labcakey.pem
    -sha512 -days 365 -out labca\labcacert.pem -subj
    "/C=<country>/ST=<state>/L=<location>/O=<organization>
    Lab/OU=<organization unit>/CN=<cn abbreviation>"
   openssl verify labca\labcacert.pem
   ```

   Generated output: `C=<country>, ST=<state>, L=<location>, O=<organization> Lab, OU=<organization unit>, CN=<cn abbreviation>`

   d. Import the CA certificate into the system trust store.

   ```
   certutil -addstore -f "Root" labca\labcacert.pem
   openssl verify -CAfile labca\labcacert.pem
    labca\labcacert.pem
   ```

   Generated output: `labca/labcacert.pem: OK`

   e. Obtain the fully qualified domain name (FQDN) for the host.

   ```
   net config workstation | findstr /C:<Full Computer name>
   ```

   The `Full Computer name` value returned is the FQDN to use as the `<hostname>` later in this procedure.

   f. Create the SAN (Subject Alternative Name) configuration file.

   ```
   copy "%ProgramFiles%\OpenSSL-Win64\bin\cnf\openssl.cnf"
    openssl-san.cnf
   echo.>>openssl-san.cnf
   echo [SAN]>>openssl-san.cnf
   echo subjectAltName=DNS:<hostname>>>openssl-san.cnf
   ```

   Adjust the source `openssl.cnf` path to match your OpenSSL installation location if it differs from the default.

   g. Generate the server certificate signing request (CSR) and sign it with the CA.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Enter the FQDN as the `<hostname>` value in the following commands. If more than one FQDN value is returned, use the value with the following format: `hostname.domain.domain.com`.

```
openssl req -new -newkey rsa:4096
 -keyout labmid\rsa-labmidkey.pem
 -sha512 -nodes -out labmid\mid.csr -subj
 "/C=<country>/ST=<state>/L=<location>/O=<organization>/OU=<o
rganization unit>/CN=<hostname>"
openssl x509 -req -days 365 -in labmid\mid.csr -CA
 labca\labcacert.pem -CAkey labca\ec-labcakey.pem
 -CAcreateserial -extensions SAN -extfile openssl-san.cnf
 -out labmid\mid.crt
```

Generated output:

```
Signature ok
subject=/C=<country>/ST=<state>/L=<location>/O=<organizatio
n>/OU=<organization unit>/CN=<mid server host fqdn>
Getting CA Private Key
```

Note: With OpenSSL 3.x (the Win32/Win64 installer), the `Getting CA Private Key` line may not appear in the output, but `Signature ok` and the correct `subject=` line are still present. To confirm the certificate was signed correctly, verify `labmid\mid.crt` directly:

```
openssl x509 -in labmid\mid.crt -noout -subject -dates
```

This command shows the subject you specified and valid `notBefore`/`notAfter` dates.

h. Combine the key and certificate files into one file, named `mid.pem`.

```
copy /b labmid\rsa-labmidkey.pem+labmid\mid.crt
 labmid\mid.pem
```

What to do next

Set up the MID Web Server with a .pem file.

Set up the MID Web Server with a .pem file

Install the .pem file into the MID unified keystore and set up the MID Web Server to enable configuring mTLS on your MID Web Server and agent.

Before you begin

Copy the `labmid/mid.pem` file (created in the Create keys and certificates procedure) to your MID Web Server host directory.

Note: This procedure includes commands for both CentOS 7/Linux and Windows Server environments. Select the commands relevant for your host operating system. If working with another Linux distribution, adapt the commands as needed for your specific OS.

Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In a Linux environment:

   a. Access the agent subdirectory under the MID Server's host directory.

   b. Import the .pem file into the MID unified keystore.

   ```
   bin/scripts/manage-certificates.sh -a midwss /<path>/mid.pem;
   bin/scripts/manage-certificates.sh -l;
   ```

   Replace `<path>` with the directory where you copied the mid.pem file.

   The relevant part of the output is: `"defaultsecuritypairhandle,midwss"`

   c. Validate your MID Server on the instance.

   d. Select the MID Server on the instance and select **Setup ACC listener**.
   A new MID Web Server is created.

   e. Navigate to the newly created MID Web Server record
   (*ecc_agent_ext_context_webserver*) on the instance.

   f. Select the MID Web Server and set the Keystore certificate alias value to `midwss`.

   g. Select **Save**.

   h. Restart the MID Web Server.

2. In a Windows environment:

   a. Access the agent subdirectory under the MID Server's host directory.

   b. Import the .pem file into the MID unified keystore.

   ```
   bin\scripts\manage-certificates.bat -a midwss <path>\mid.pem
   bin\scripts\manage-certificates.bat -l
   ```

   Replace `<path>` with the directory where you copied the mid.pem file.

   The relevant part of the output is: `"defaultsecuritypairhandle,midwss"`

   c. Validate your MID Server on the instance.

   d. Select the MID Server on the instance and select **Setup ACC listener**.
   A new MID Web Server is created.

   e. Navigate to the newly created MID Web Server record
   (*ecc_agent_ext_context_webserver*) on the instance.

   f. Select the MID Web Server and set the Keystore certificate alias value to `midwss`.

   g. Select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

h. Restart the MID Web Server's Windows service using one of the following methods:

- Use the Services app on the Windows host.
- Use the command line with the MID Server's configured service name:

```
net stop AgentClientCollector
net start AgentClientCollector
```

What to do next

Connect the agent to the MID Web Server using TLS.

Connect the agent to the MID Web Server using TLS

Connect the agent to the MID Web Server to enable configuring mTLS on your MID Web Server and agent.

Before you begin

Ensure that you have installed the .pem file and set up the MID Web Server. For details, see Set up the MID Web Server with a .pem file.

Note: This procedure includes commands for both CentOS 7/Linux and Windows Server environments. Select the commands relevant for your host operating system. If working with another Linux distribution, adapt the commands as needed for your specific OS.

Role required: agent_client_collector_admin

Procedure

1. In a Linux environment:

a. Add the `labcacert.pem` file to your agent host's trust store.

```
sudo cp
 -a /<path>/<to>/labcacert.pem /etc/pki/ca-
trust/source/anchors/;
sudo update-ca-trust extract
openssl verify /<path>/<to>/labcacert.pem
```

The generated output is: `/<path>/labcacert.pem: OK`

b. Configure the `acc.yml` file to use TLS.

i. Set the *`insecure-skip-tls-verify`* property to `false`.

ii. Set the *`backend-url`* property to use the MID Server's FQDN.

```
backend-url="wss://<mid server fqdn>:<mid web server
 port>/ws/events"
```

c. Restart the agent.

```
systemctl restart acc;
```

d. Verify in the logs that the agent is connected to the MID Server.

2. In a Windows environment:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Add the `labcacert.pem` file to your agent host's trust store by importing it into the Windows Certificate Store.

```
certutil -addstore -f "Root" <path>\labcacert.pem
openssl verify -CAfile <path>\labcacert.pem
 <path>\labcacert.pem
```

The generated output is: `<path>\labcacert.pem: OK`

b. Configure the `acc.yml` file (`C:\ProgramData\servicenow\agent-client-collector\config\acc.yml`) to use TLS.

i. Set the *`insecure-skip-tls-verify`* property to `false`.

ii. Set the *`backend-url`* property to use the MID Server's FQDN:

```
backend-url:
 - "wss://<mid server fqdn>:<mid web server port>/ws/events"
```

c. Restart the agent using one of the following methods:

- Command line:

```
net stop AgentClientCollector
net start AgentClientCollector
```

- Services app: Select and hold the **Agent Client Collector** entry and select **Restart**.

d. Verify in the logs that the agent is connected to the MID Server.

Logs are located at: `C:\ProgramData\ServiceNow\agent-client-collector\log\acc.log` (or the location set by the *`log-file`* configuration parameter).

What to do next

Configure mTLS authentication for a MID Web Server.

Configure mTLS authentication for a MID Web Server

Enhance security in your MID Web Server extension by enabling mTLS authentication.

Before you begin

Ensure that you have enabled Transport Layer Security (TLS) on the agent. For details, see Connect the agent to the MID Server using mTLS.

Ensure that the *`insecure-skip-tls-verify`* parameter in the `acc.yml` configuration file is set to `false`. For details on the `acc.yml` file, see Configuration file options.

Note: This procedure includes commands for both CentOS 7/Linux and Windows Server environments. Select the commands relevant for your host operating system. If working with another Linux distribution, adapt the commands as needed for your specific OS.

Role required: agent_client_collector_admin

About this task

The MID Web Server extension searches the following locations (in the specified order) to access the truststore location and password:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Truststore location: The *mid.webserver.truststore.path* JVM system property.

  If that property is empty, the extension retrieves the location from the *javax.net.ssl.trustStore* JVM system property.

  If no location is specified, the truststore location defaults to the absolute path of the `cacerts` file of the JRE running the MID Server.

- Truststore password: The **Truststore Password** field on the extension form in the instance.

  If that field is empty, the system retrieves the password from the *javax.net.ssl.trustStorePassword* JVM system property.

  If no location is specified, the password defaults to *changeit*.

Procedure

1. In a Linux environment:

   a. Navigate to the root folder of your MID Server.

   b. Run the following command to add your certificate to the MID trust store:

   ```
   ./jre/bin/keytool -importcert
    -file /etc/pki/ca-trust/source/anchors/labcacert.pem
    -destkeystore ./jre/lib/security/cacerts -alias mtlsca
   ```

   c. Enter `changeit` when prompted for a password.

   d. Select **yes** on the confirmation message window to indicate that you trust the certificate.

   e. Run the following command to verify that your certificate was successfully added to the MID trust store:

   ```
   ./jre/bin/keytool -list -keystore ./jre/lib/security/cacerts
    -alias mtlsca
   ```

   f. Enter `changeit` when prompted for a password.

   g. In the MID Server wrapper override configuration file (`wrapper-override.conf`, located in the `conf` directory under the MID Server's home directory), configure the revocation of client certificates using the *mid.webserver.cert.revocation.check.enabled* property.

   If you have a custom internal certificate, set it to `false` by adding the following line to `conf/wrapper-override.conf` on the MID Server:

   ```
   wrapper.java.additional.4=-Dmid.webserver.cert.revocation.che
   ck.enabled=false
   ```

   When enabling this property (`true`), configure the *mid.webserver.ocsp.responder.url* property with the OCSP responder URL. This value overrides any URL embedded in the certificate.

   h. If you have changed properties in the wrapper override configuration file, restart the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

i. Access the MID Server record on your ServiceNow instance and change the value of the **Authentication type** field to **mTLS**.

j. Restart the MID Web Server.

k. Verify that the MID Web Server and websocket endpoint are up and running.

2. In a Windows environment:

a. Navigate to the root folder of your MID Server.

b. Run the following command to add your certificate to the MID trust store:

```
jre\bin\keytool.exe -importcert -file <path>\labcacert.pem
 -destkeystore jre\lib\security\cacerts -alias mtlsca
```

c. Enter `changeit` when prompted for a password.

d. Select **yes** on the confirmation message window to indicate that you trust the certificate.

e. Run the following command to verify that your certificate was successfully added to the MID trust store:

```
jre\bin\keytool.exe -list -keystore jre\lib\security\cacerts
 -alias mtlsca
```

f. Enter `changeit` when prompted for a password.

g. In the MID Server wrapper override configuration file (`wrapper-override.conf`, located in the `conf` directory under the MID Server's home directory), configure the revocation of client certificates using the *`mid.webserver.cert.revocation.check.enabled`* property.

- If you have a custom internal certificate, set it to `false` by adding the following line to `conf\wrapper-override.conf` on the MID Server:

```
wrapper.java.additional.4=-Dmid.webserver.cert.revocation.ch
eck.enabled=false
```

- When enabling this property (`true`), configure the *`mid.webserver.ocsp.responder.url`* property with the OCSP responder URL. This value overrides any URL embedded in the certificate.

h. If you changed properties in the wrapper override configuration file, restart the MID Server using one of the following methods:

- Windows Services app
- Command line using the MID Server's configured service name:

```
net stop AgentClientCollector
net start AgentClientCollector
```

i. Access the MID Web Server record on your ServiceNow instance and change the value of the **Authentication type** field to **mTLS**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

j. Restart the MID Web Server.

k. Verify that the MID Web Server and websocket endpoint are up and running.

What to do next

Connect the agent to the MID Server using mTLS.

Connect the agent to the MID Server using mTLS

Before configuring mTLS authentication on the agent, you must run a series of commands that enable configuring Transport Layer Security (TLS) authentication.

Before you begin

Ensure that you have performed the following tasks:

1. Create keys and certificates
2. Set up the MID Web Server with a .pem file
3. Connect the agent to the MID Web Server using TLS
4. Configure mTLS authentication for a MID Web Server
5. Connect the agent to the MID Server using mTLS

Role required: agent_client_collector_admin

Procedure

1. Generate a key for your agent.

```
openssl ecparam -out labacc/acc.key -name prime256v1 -genkey
```

2. Generate a certificate request for your agent, according to the following format:

```
openssl req -new -key labacc/acc.key -out labacc/acc.csr -subj
 "/C=<country>/ST=<state>/L=<location>/O=<organization>/OU=<org
anization unit>/CN=<cn abbreviation>/emailAddress=<email
 address>"
```

   For example:

```
openssl req -new -key labacc/acc.key -out labacc/acc.csr
 -subj "/C=US/ST=NC/L=Raleigh/O=ServiceNow/OU=ITOM
 Lab/CN=acclinux/emailAddress=john.smith@servicenow.com"
```

3. Generate a signed certificate for your agent.

```
openssl x509 -req -days 365 -in labacc/acc.csr -CA
 labca/labcacert.pem -CAkey labca/ec-labcakey.pem
 -CAcreateserial -extensions client -out labacc/acc.crt
```

4. In the `labacc` folder, copy your key and certificate files to the agent host's virtual machine.

```
cp ./acc.key <agent host config folder>
cp ./acc.crt <agent host config folder>
```

   The configuration folder paths depend on your OS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Linux: `/etc/servicenow/agent-client-collector/`
- Windows: `C:\ProgramData\servicenow\agent-client-collector\config\`
- macOS: `/Library/Application\ Support/servicenow/agent-client-collector/`

5. Navigate to the location where you copied the `.key` and `.crt` files from.
6. Run the following commands to update the files' read permissions and enable the agent to read from them.
   - Linux:
     a. `chown servicenow:servicenow acc.key`
     b. `chmod 0400 acc.key`
   - Windows:
     a. Select and hold (or right-click) the `.key` file and select the **Security** tab from the file properties.
     b. Add the *`servicenow`* user to the list of users with read access to the file.
   - macOS:
     a. `chown _servicenow:_servicenow acc.key`
     b. `chmod 0400 acc.key`

   Note: The agent must be run with the default servicenow user.

7. In the acc.yml configuration file, add the following parameters to specify the path to your key and certificate files.

```
key-file: "<path to acc.key file>/acc.key"
cert-file: "<path to acc.cert file>/acc.crt"
```

8. Restart the agent.
   - Linux: `systemctl restart acc`
   - Windows:
     a. Open the Services application.
     b. Select and hold (or right-click) the Agent Client Collector entry and select **Restart**.
   - macOS:
     a. `launchctl unload -w /Library/LaunchDaemons/com.sn.acc.plist`
     b. `launchctl load -w /Library/LaunchDaemons/com.sn.acc.plist`
9. Check the logs to verify that the agent was able to connect to the instance using TLS.

   Agent Client Collector logs are accessible from:

   - Linux: `/var/log/servicenow/agent-client-collector/acc.log`
   - Windows: The location specified by the `log-file` configuration flag. The default location is: `C:/ProgramData/ServiceNow/agent-client-collector/log/acc.log`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the C : /ProgramData directory is hidden (such as on older Windows versions), change the Explorer filter to show hidden elements.

- macOS: The default location is: /Library/Application\ Support/servicenow/agent-client-collector/log/acc.log

Related topics

MID Web Server and agent mTLS Authentication

Event Management configuration preferences

Preferred settings of properties and general configuration.

Use the Known Error Portal and the Community to further help you find information issues.

General preferences

Self-health

By default, the self-health monitoring feature is not enabled. To enable it, navigate to **Event Management** > **Settings** > **Properties** and select **Yes** for the *Enable Event Management self-health monitoring* (evt_mgmt.self_health_active) property. Use this feature to monitor and track many Event Management features.

Note: CIs used in the self-health service are created in the CMDB.

Use the following settings to help with preventing performance degradation.

Topic	Details
Business rules	• Avoid writing business rules for event [em_event] tables, as they do not run in the current default REST URL that is used for event injection. • Business rules that are written for alert [em_alert] tables must be highly efficient or they may result in performance degradation. Instead of writing a business rule, consider whether it is more appropriate to write a job. An inefficient business rule can cause incident creation for an alert to fail and the alert impact calculation to fail. • Do not write async business rules for alert tables. • Business rules must not change the **Category** field on event [em_event] tables.
Scaling up	Check the average event processing time before scaling up event throughput when first starting with Event Management. Do this check after an initial flow of events and all rules are in place. If processing time takes over a few milliseconds per event, determine the cause for the processing slowdown before continuing to scale. Performance duration can be checked in the Performance Statistics [sa_performance_statistics] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Topic	Details
Configure for large-scale environments	• Set the **Enable multi node event processing** (*evt_mgmt.event_processor_enable_multi_node*) property to **Yes**. Enable multi-node in production environments and set values based on the size of the deployment and expected event rate. • Set the **Number of scheduled jobs processing events** (*evt_mgmt.event_processor_job_count*) property to 4. • If you are sending events from a custom source, verify that events have**Message Key** or **Source**, **Node**, **Type**, and **Resource** data.
Latency issues for receiving events	Check the following settings: • Verify that the **Bucket** field in the event [em_event] table is set to a value that is greater than zero (*0*). • Navigate to **System Scheduler** > **Scheduled Jobs** and search for *- process events**. Check that all *- process events** jobs exist according to the **Number of scheduled jobs processing events** (*event_processor_job_count*) property configuration. Verify that the **State** is *Running* or *Ready*. If the state is *Queued* or *Error*, set the job state to *Ready*.
Archive events	• Avoid changing the default retention time for events. • To log events for a longer time, create an archive table and a job that copies new events to it. Do this by scheduling a job to regularly back up events [em_event] to a custom table. • Do not extend table rotation by adding more days.

Event integration

SNMP traps

- Use a monitoring tool to send SNMP traps, rather than sending them directly from devices.
- To avoid having to rewrite event rules, upload MIBs prior to defining the event rules.

Web service API

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Using a web service API for integration can reduce the number of event rules needed. This action avoids having to transform events (prepared data is sent in an event to the instance).
- Use dedicated credentials for integration. Optionally, designate credentials specific to each event source.

CloudWatch

Use dedicated credentials for integrating CloudWatch with ServiceNow.

Email

Use email only if the source has a low volume and other options are not available, such as, running a script or forwarding an SNMP trap.

Event rules

Configuration settings when creating event rules:

- Write Event Rules to apply to the broadest number of events possible. More specific rules can then be created as necessary and should use a lower-order value.
- If a more general rule can achieve the same outcome, avoid writing Event Rules that apply only to a certain subset of events.
- When Event Rules are applied to events, no changes are made to the original event. All processing occurs in memory, so use the **Processing Notes** field and/or use the **Check Process of Event** UI action link to troubleshoot.
- If you change a rule/transform that has existing mapping rules, you should review and retest with events that are either actual or simulated.
- Ensure that the **From** field value exactly matches a string in the JSON in the *additional_info* field of an event. This matching happens when a rule has been configured based on information in a MIB file. If the MIB file is not uploaded, the JSON for the SNMP trap shows varbinds (variable binding) with dotted names, instead of the translated name in the MIB. The event field mapping rule then fails to be applied.
- Establish a consistent naming convention. A common convention is: *<customer acronym>.<Event Source>.<Description>*. For example, *ACME.OEM.Normalize*
- If two Event Rules have similar conditions set, use the **Order** field to control which Event Rule runs.
- Use Event Rules to associate an alert with a CI.

Additional settings for constructing event rules:

Desired result	Required activity
Effective de-duplication and enabling efficient parallel event processing	Populate the **Source**, **Node**, **Type**, **Resource**, **Metric Name** fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Desired result	Required activity
CI binding	• Bind to host - by populating the **Node** field and optionally **CI identifiers**. • Binding to application and / or device – by populating the **CI Identifiers** field and the **Additional information** field.
Alert Correlation, using alert aggregation	Populate the **Resource** and **Metric Name** fields. **Note:** If CI is also bound, Alert Correlation is improved.

Custom event fields

Include additional fields in the **Additional information** field of the event only.

Do not add additional fields to an event by adding a custom field to the event [em_event] table.

Do not add columns to the event [em_event] table.

For information about how to include additional fields in events, see Custom alert fields.

De-Duplication

Configuration settings for de-duplication.

- The **message_key** field is used for De-Duplication. If reliable message key values are not provided with the source event, it is important to have a well-defined plan for constructing these identifiers.
- If the message key is not defined, then the default message key is *<Source + Node + Type + Resource + Metric Name>*.
- The guideline is to have the event source populate the *<Source + Node + Type + Resource + Metric Name>* fields out-of-the-box and populate the message key. This action enables a better distribution of event processing among instance workers and nodes.
- If the source event does not have values for these fields, make sure to populate them using transform rules. This action does not affect event processing, but is used for de-duplication. Populate as many of these fields as possible before they are sent to the instance. This action provides better distribution of events over the processor workers, and therefore better throughput and scale.

CI binding

- Where possible, always attempt to bind an alert to a CI.

  Note: Although Event Rules are defined on events, CIs are bound to alerts that result from those events, and are not bound to the event.

- To bind a host, machine or any device with an IP, populate the event **Node** field with a unique host name, FQDN, IP, or MAC address.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If other identifiers are necessary to identify a host, then populate the ci_identifiers field with a JSON format. The JSON format must contain the **CMDB** field name and value to perform the match.

Note: The event **Node** field must be populated from an event rule or populated with a unique host name from the source before the event is inserted.

- The primary binding strategy is to use the **Node** field. If the **Node** field is not pre-populated in the event, it can be populated using event rules.

Alert settings

Alert lifecycle

General alert functionality:

- An alert is opened when an event is not ignored or its threshold is exceeded by an event rule. De-duplication must not identify the event as belonging to an existing alert.
- An alert is closed when a closing event is sent on the same message key, or the alert is closed manually.
- An alert is reopened if an opening alert that has the same message key is sent within the timeframe defined in properties (default is one hour).
- If an alert is opened and closed at a high rate, as defined in properties, it becomes flapping. When this opening and closing rate stops, the alert goes out of flapping state.
- If an incident is opened from an alert, that alert remains open as long as the incident remains open. By default, when either the incident or the alert is closed the other is closed as well. This behavior can be configured using properties.
- Do not close an alert when creating a corresponding incident.
- Do not delete an open alert. Close an alert first and then delete it.
- Use `Acknowledge` to denote that the alert is known, and can temporarily be ignored.
- Do not use `Acknowledge` to mark an alert as needing attention.
- Do not create alerts in any of these states:
  - `Closed`
  - `OK`
  - `Open`
- The `evt_mgmt.alert_auto_close_interval` property automatically closes alerts after the specified period. Do not specify 0, as this value disables the feature and may lead to performance degradation.
- Do not create alerts in OK state. In some monitoring systems OK denotes that an issue has been resolved, while in other monitoring systems OK is used to denote events that are not of operational significance. For the former case, use `Clear` instead of OK using a Mapping Rule. For the latter case, have an Ignore rule, unless the events are of specific value.

Alert management rules

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Use alert management rules to define automated responses to alerts, such as opening incidents, running subflows, and launching applications or URLs. For more information, see Alert management rules for resolving alerts.
- Use alert filters to specify which alerts the rule applies to. For more information, see Create an alert management rule.
- Use the **Order** field to control which alert management rule runs first when rules have similar conditions. For more information, see Create an alert management rule.
- Use subflows to customize alert handling. For example, you can resolve alerts or notify teams. For more information, see Create a custom subflow for alerts.
- A scheduled job applies alert management rules to new or updated alerts every 11 seconds. If an alert management rule does not start immediately, allow 10–15 seconds before you start troubleshooting. For more information, see Alert management rules for resolving alerts.

Business rules

- Business rules created on alert tables should not take more than a few milliseconds. In place of using a business rule, consider if the same functionality can be achieved using a job.
- Do not use business rules to associate an alert with a CI. Use event rules to do binding instead of using business rules.

Planning

- Organize event source configuration of filters, modules, and so on, into multiple parallel efforts, rather than in serial.
- Validate processed event formats to ensure that data that is parsed is aligned with desired results.
- Test production events in a non-production environment. Integrate with non-production element managers and ServiceNow instances. If non-production element managers are not available, send events from element managers to both production and non-production environments.

Services and dashboard

- Use Service Groups to group services into logical groups to reduce the number of services displayed on the Service Health dashboard.
- Import manually built service maps.

Metric Intelligence collector logs and files

Metric Intelligence collector logs and files are located under the path `$(MID_SERVER_DIR)/agent`. Use these logs and files for troubleshooting and monitoring purposes.

Location of Metric Intelligence collector logs and files

Log or file	Path
PowerShell metric collector log file	Logs/retrieve_metrics{connector instance ID}.log
PowerShell output file	work/metrics/metrics_output_{connector instance ID}.txt
PowerShell input file	work/metrics/parameters_{connector instance ID}.txt

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric Intelligence performance can be checked in the MID Server log file when the *mid.log.level* MID Server parameter is in debug mode.

Metric Intelligence performance numbers are available in the Performance Statistics [sa_performance_statistics] table. To view the performance numbers, filter the Performance Statistics list for **Metric Collector**.

Related topics

Rotate event and alert table for cleanup

Domain separation and Event Management

Domain separation is supported in Event Management. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

How domain separation works in Event Management

When domains are separated in Event Management, users can only see and manage alerts and events in their own (tenant) domain.

A domain column is present for Event Management tables that are provided in the base system. The column displays the name of the domain to which the event or alert belongs.

Note:

If the domain column does not display in the list, click (Update Personalized List) and add the required column. In addition, you can add a column that displays the domain path.

For information about domain separation limitations in Event Management, see Configure Event Management domain separation.

Related topics

Domain separation for service providers

Configure Event Management domain separation

You can configure Event Management for **domain separation** to create logically defined domains that limit unauthorized access to data. When domains are separated in Event Management, users can only see and manage alerts and events in their own (tenant) domain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: evt_mgmt_admin and evt_mgmt_integration

About this task
Activate the **Domain Support – Domain Extension Installer** plugin and configure the MID Server for Event Management.

The following Event Management features have limited domain separation support.

Features with limited domain separation support

Feature	Support
Event – alert flow	Supported. Separation is based on the domain user that sent events. User access is required for the credentials of the sending API events or in the configuration of the MID Server reporting events. In a multi-domain environment, each MID Server can serve only one domain according to the integration user that it uses. In the configuration of the connector instance, make sure that the MID Server have same domain as Connector instance. To configure pull connectors to support custom domain separation, see Configure domain separation for pull connector events.
Impact calculation	Supported. Segregation is based on the manner in which CIs are segregated.
Application services	Partially supported. Segregation is based on the manner in which CIs are segregated.
Dynamic CI groups	Partially supported. Segregation is based on the manner in which CIs are segregated. **Note:** The discovery process does not segregate CIs by domain.
Remediation	Supported. While editing alert management rules, users can only apply relevant workflows. For more information on domain separation in the Flow Designer, see Domain separation and Flow Designer.
Alert Aggregation	Supported. In domain-separated environments, alert groups are created only for alerts in the same domain.

Procedure
1. If it is not already active, activate the Domain Support – Domain Extension Installer plugin.
2. Configure the connector instance to be in the same domain as the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event Management Integrations

An event is a notification from one or more monitoring tools that indicate something of interest has occurred, such as a log message, warning, or error.

Events are processed and alerts are generated based on the following factors:

- Event collection configuration via a MID Server, script, SNMP trap collector, or email message.
- An Event rule, alert binding, or event field mapping configuration for processing events from various sources.

When a package of events that contains events with a custom (non-supported) state is processed, the events with the custom state are not processed.

Integration with external events

Event Management receives external events via an integration or scripts. If you are using a script to collect events, no configuration is required. All other integrations require configuration.

When an integration is configured, Event Management can collect events that are sent directly to the ServiceNow AI Platform instance or via the MID Server, SNMP trap or email.

The integration with external systems can be initiated from a ServiceNow AI Platform instance using a pull connector, or from an external system using a push mechanism.

Integration via a ServiceNow AI Platform instance

The ServiceNow AI Platform instance includes push connectors that help external monitoring systems to integrate with Event Management directly.

You can configure the authentication in the ServiceNow AI Platform instance and update the external monitoring systems to use endpoints to send the events for the push connectors that are available. For more information, see Push Connectors.

Integration via MID Server

The MID Server facilitates communication and movement of data between external applications, data sources, services, and the ServiceNow AI Platform instance.

Connector operation	Description
Pull	You can configure connector instances to use connector definitions (script instructions). This enables the MID Server to connect to external servers to obtain event information using a pull operation.
Push	You can configure external systems to push event information to the ServiceNow AI Platform instance via MID Server. For example, you can configure the MID WebService Event Collector to push event messages to the MID Server. For more information, see Pushing events to the MID Server using web service API.

Standard event form

A standard event form must be used by all events, regardless of how they arrive at the ServiceNow AI Platform instance. For more information, see Event field format for event collection.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event Management Connectors domain personalization

Create events in different domains for all Event Management connectors using just a single connector instance by personalizing domain separation of Event Management connectors.

How the domain is identified

In a ServiceNow instance, events are created within the logged-in or MID Server user domains. However, you can create events in a different domain by providing metadata values in the event HTTP request, including the URL, headers, request body (payload), and connector parameters.

The logged-in user or the MID Server user must have access to the provided domain. The connector resolves the domain based on the provided identifier and creates events in that domain.

System properties provide the default metadata that sets the domain-separated table to identify the domain. For more information, see Domain properties installed with Event Management.

Connector Exceptions

You cannot personalize domain separation for the following pull connectors:

- SCOM
- IBM Netcool (supports only IBM Netcool_v2)
- HPOM
- Hyperic

You cannot personalize domain separation for the following push connectors:

- SAP SolMan Transform Script
- bmcTrueSight Transform Script
- Transform Generic Events MID
- TransformEvents_xmlSample

Related topics

Domain separation and Event Management

Configure domain separation for pull connector events

Configure pull connectors to personalize domain separation of events so you can use them to create events in domains other than the user's currently logged-in or MID Server domain.

Before you begin

Make sure that you have the following setup:

- You have read access to the domain information table for the evt_mgmt_integration role (default table: core_company). For more information, see Configure an ACL.
- You enabled the *evt_mgmt.connector_enable_custom_domain_separation* system property .
- You installed the following plug-ins on your ServiceNow instance:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Event Management *com.glideapp.itom.snac* plug-in
- Event Management connectors *sn_em_connector* plug-in
- Domain Support - Domain Extensions Installer *com.glide.domain.msp_extensions.installer* plug-in

You can download the plug-ins from the ServiceNow Store.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances** > **.**
2. On the **Connector Instances** list, deactivate the connector instance of the connector for which you want to configure domain separation by selecting it and clearing the **Active** check box.
3. Select **Update**.
4. Navigate to **All** > **Event Management** > **Integrations** > **Connector Definitions**.
5. From the **Connector Definitions** list, select your connector definition.
6. Determine whether all events should be created in a provided domain or determined by a domain provided in the payload of each event.

   a. In the **Connector Parameters** table, double-click **Insert a new row**

   b. Choose the event creation domain definition location.
   - To have all events be created in the provided domain, enter `payloadDomainInfoFieldValue` in the **Name** column.
   - To have events be created in the domain provided in the payload of the event, enter `payloadDomainInfoFieldName` in the **Name** column.

   Note: If both the *payloadDomainInfoFieldValue* and the *payloadDomainInfoFieldName* are provided, the domain provided in the payload of the event takes precedence.

   c. Select **Update**.
7. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
8. On the **Connector Instances** page, select your connector instance from the list.
9. In the Connector Instance Values list, provide a value in the **Value** field of the parameters.

Parameter	Value
payloadDomainInfoFieldValue	The value that will be used to match the do main record. For example, ServiceNow.
payloadDomainInfoFieldName	The key in the payload that has the domain identifier. For example, if the payload key **com pany** holds the domain identifier, then the value in the **Value** column would provide the company information. For example: **company** {

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Value
	"source": "Solarwinds", "host": "127.0.0.1", "company": "ServiceNow" } **Note:** Dot-walking isn't supported.

10. Select **Save**.
11. Select **Test Connector** to verify the connection between the MID Server and the connector.
    - If the test is successful, a confirmation message appears.
    - If the test fails, a message provides instructions to correct the problem.

    On a successful test, an onscreen message appears. If the test fails, follow the onscreen message instructions to correct the problem and then run another test.

    Note: Use a network tool such as *ping* to verify whether the credentials are correct and network connection from the MID Server to the connector server.

12. After the connector is verified, select the **Return to Connector Instance** link.
13. Activate the connector instance by selecting **Active** and **Update**.
14. **Optional:** Configure the domain for events by providing domain metadata through the connector parameters.

    For example, to use the User [sys_user] table for a single connector instance in the **Connector Definitions** form, create parameters with the following names, and add their values in the **Connector Instance** form.

    - *connectorDomainInfoTableName* - sys_user

      Note: The table must be domain separated.

    - *connectorDomainInfoColumnName* - user_name
    - *connectorDomainIdColumnName* - sys_domain
    - *connectorDomainPathColumnName* - sys_domain_path

    For more information, see Connector domain metadata.

Result

The Event Management connector is set up for domain separation.

Configure domain separation for push connector events

Configure push connectors to create events in domains other than the user's currently logged-in or MID Server domain.

Before you begin

Make sure that you have the following setup:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You have read access to the domain information table for the evt_mgmt_integration role (default table: core_company). For more information, see Configure an ACL.
- You enabled the *evt_mgmt.connector_enable_custom_domain_separation* system property.
- You installed the following plug-ins in the ServiceNow AI Platform® instance:
  - Event Management (*com.glideapp.itom.snac*)
  - Event Management connectors (*sn_em_connector*)
  - Domain Support - Domain Extensions Installer (*com.glide.domain.msp_extensions.installer*)

    You can download the plug-ins from the ServiceNow Store.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Push Connectors**.
2. Open a push connector.
3. On the **Push Connector** screen, scroll down to the **Push Connector Configurations** list and select **New**.

   Note: If this list is not displayed, you must create a push connector configuration parameter. For more information, see Create a push connector configuration parameter.

4. In the **Push Connector Configurations** screen, determine whether events should be created in the provided domain or the domain provided in the payload of each event.

Domain creation location	Parameter and values
The provided domain	○ **Field name**: payloadDomainInfoFieldValue ○ **Value**: domain-identifier

The domain provided in the payload of each event	○ **Field name**: payloadDomainInfoFieldName
	○ **Value**: field-name-in-payload
	Example: If the payload key “company” holds the domain identifier, then “company” is the value.
	“company”
	{
	“source”: “Solarwinds”,
	“host”: “127.0.0.1”,
	“company”: “ServiceNow”
	}

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If both the *payloadDomainInfoFieldValue* and the *payloadDomainInfoFieldName* parameters are provided, then the *payloadDomainInfoFieldName* takes precedence.

5. Select **Submit**.

Result

Events are created in the domain referred to by the domain identifier.

Note: The table that stores domain information records (for example, core_company) must be domain separated. For more information, see Domain-separate a custom table.

Related topics

Integrate with push connectors

Connector domain metadata

Configure Event Management connectors

Event Management provides many connectors to pull or push events from external devices. Connectors are available from the ServiceNow store as well as from third parties. You can also create custom connectors.

Provided connectors (pull)

This table lists the connectors that are provided with Event Management.

Note:

- The Azure Metrics and AWS CloudWatch Metrics connectors display in the base system only when the Operational Intelligence plugin (com.snc.sa.metric) is activated.
- All connectors support external credential storage solutions. For details on configuring external credential storage, see External credential storage configuration.

Connectors	Supported version	Description
Azure Monitor		Send alert state changes to the Azure Portal.
HPOM	08.60.005	Receive events from the HP Operations Manager (HPOM).
Hyperic	5.8.4.0	Receive events from the VMware vRealize Hyperic server.
Logicmonitor	156	Receive events from the Logicmonitor server.
IBM Netcool	8.1.0	Receive events from IBM Netcool/OMNIbus Object Servers and Impact Servers.
Icinga2	2.4.1	Receive events from the Icinga2 server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connectors	Supported version	Description
Apache Kafka Consumer Connector	• 3.0.0 • 3.2.0	Receive events from the Apache Kafka server.
ServiceNow Cloud Observability	2.1.1 or higher	Receive events from the ServiceNow Cloud Observability platform.
NagiosXI	5.5.2	Receive events from the Nagios XI server.
NNMi	10.30.653	Receive events from the NNMi server.
OBM, also known as OMi v2	10.01, 10.10, 10.11, 10.12, 10.60, 10.61, and 11.01	Receive alerts from the OBM server.
OMi	10.01, 10.10, 10.11, 10.12, 10.60, and 10.61	Receive alerts from OMi server.
Op5	7.3.15	Receive alerts from an OP5 Monitor.
Opsview	5.4.0	Receive alerts from an Opsview Monitor
PRTG	21.3.71	Receive alerts from a Paessler PRTG Network Monitor source.
SAP Solution Manager	7.2 740 SP16	Receive alerts from an SAP Solution Manager Monitor source.
SCOM	• 2007 – version 6.1.7221.0 • 2012 - version 7.1.10226.0 • 2016 - version 7.2.117190	Alert and metric collection from the Microsoft System Center Operations Manager (SCOM).
SolarWinds	• NPM 12.0 • SAM 6.2.4	Receive events from the SolarWinds monitor.
vCenter	6.5	Receive events from your VMware vSphere environment.
vRealize	VMware vRealize Operations Manager Supported version: 8.10.0	Receive events from the vRealize Operations Log and Event Management servers.
Zabbix	5.0.9	Receive alerts from the Zabbix server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provided connectors (push)

This table lists the connectors that are provided with Event Management. Generic push connectors are also provided, one for event collection through the MID Server and the other through the instance.

Before you begin, ensure the Event Management Connectors plugin is installed in the ServiceNow AI Platform instance. You can download the plugin from the ServiceNow Store website.

Connectors	Description	Supported authentication
AWS	Use the connector to process Amazon Web Services (AWS) alarms as events. The event collection is through the instance.	N/A
Azure	Use OAuth authentication to process Azure platform event messages. The event collection is through the instance.	• Integrate Azure Monitor with OAuth authentication • Integrate Azure Monitor with basic authentication
Configure Azure Monitor Bi-directional connector	Use the Azure Monitor Pull connector to send information from ServiceNow Event Management to the Azure Portal.	Basic authentication
Microsoft Azure MID Push Connector	Collect JSON formatted event messages sent from Microsoft Azure. The collection is through the MID Server.	Basic authentication
BMC TrueSight	Collect JSON formatted event messages sent from BMC TrueSight Operations Management (TrueSight). The event collection is through the MID Server.	Basic authentication
New Relic	Collect JSON formatted event messages sent from New Relic. The event collection is through the instance.	Basic authentication
Dynatrace	Collect JSON formatted event messages sent from Dynatrace. The event collection is through the instance.	• Integrate Dynatrace with OAuth authentication • Integrate Dynatrace with basic authentication
Datadog	Collect JSON formatted event messages sent from Datadog. The event collection is through the instance.	• Integrate Datadog with OAuth authentication • Integrate Datadog with basic authentication

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connectors	Description	Supported authentication
Catchpoint	Configure the Catchpoint connector to receive Catchpoint events. Event Management can process the Catchpoint events as ServiceNow events.	Basic authentication
Google Cloud Platform (GCP)	Collect JSON formatted event messages sent from GCP. The event collection is through the instance.	Basic authentication
Google Cloud Platform (GCP) MID Push Connector	Collect JSON formatted event messages sent from GCP. The event collection is through the MID Server.	Basic authentication
Logicmonitor	Collect JSON formatted event messages sent from Logicmonitor. The event collection is through the instance.	Basic authentication
Logicmonitor MID Push Connector	Collect JSON formatted event messages sent from Logicmonitor. The event collection is through the MID Server.	Basic authentication
Oracle Cloud Push Connector	Use basic authentication to process Oracle Cloud alarms. The event collection is through the instance.	Basic authentication
Prometheus	Collect JSON formatted event messages sent from Prometheus. The event collection is through the instance.	Basic authentication
Sentry	Configure the Sentry connector to receive Sentry events. Event Management can process the Sentry events as ServiceNow events.	Basic authentication
SNMP traps	Process SNMP traps as events.	Basic authentication
ThousandEyes	Collect JSON formatted event messages sent from ThousandEyes. The event collection is through the instance.	• Integrate ThousandEyes with OAuth authentication • Integrate ThousandEyes with basic authentication
Email	Configure an inbound email action to send email	Basic authentication

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connectors	Description	Supported authentication
	notifications when events and alerts are triggered.	
Oracle Enterprise Manager trap collection	Configure the SNMP trap connector to receive Oracle Enterprise Manager Cloud Control SNMP traps. Event Management can process these traps as events.	Basic authentication
Grafana	Collect JSON formatted event messages sent from Grafana. The event collection is through the instance.	Basic authentication
EIF listener	Configure the EIF listener context to receive EIF events. Event Management can process these EIF events as ServiceNow events.	Basic authentication
Integrate Honeycomb events	Configure the Event Management environment for the collection of events from Honeycomb by authenticating Honeycomb as a data source. In the Honeycomb platform, set your ServiceNow AI Platform as the rest endpoint using a webhook.	Basic authentication
Integrate Instana events	Configure Instana with Event Management by adding a standard webhook in the Instana console.	Basic authentication
Integrate Scout APM events	Enable the collection of events from Scout APM by authenticating Scout APM as a data source to integrate it with Event Management	Basic authentication
Integrate Panopta as a data source	Configure the Panopta cloud-based monitoring solution with Event Management.	Basic authentication
URL format for generic JSON This URL can be used as-is and requires an event rule to be configured.	Configure this generic JSON target URL and configure an event rule to collect events:`https://<<INSTANCE>>/api/sn_em_connector/em/inbound_event?source=genericJson`	N/A

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow store - connectors

Connectors are available from the ServiceNow store. Download the required content from ServiceNow Store.

Third party connectors

Connectors are also available from third parties. For further information, contact your sales representative for available connectors and listeners.

Third party connectors	Link
AppDynamics Service Model Integration Use this integration to enable your instance to have AppDynamics Application Performance Monitoring data fed into Event Management.	AppDynamics Service Model Integration
Juniper Mist Use this integration to managed service providers and large enterprises the ability to integrate real-time events and alerts from Juniper Mist with ServiceNow platform's Event Management. This capability provides a standardized approach to transform and persist Juniper Mist events in the Events Management table. The events data in ServiceNow allows customers to create their own incidents based on default or user defined alert rules without interacting with the Mist dashboard.	https://store.servicenow.com/$appstore.do#!/store/application/e89296eb1b95561019a8a753604bcb24/1.0.0?referer=%2Fstore%2Fsearch%3Flistingtype%3Dallintegrations%25253Bancillary_app%25253Bcertified_apps%25253Bcontent%25253Bindustry_solution%25253Boem%25253Butility%25253Btemplate%25253Bgenerative_ai%25253Bsnow_solution%26q%3DJuniper%2520Mist&sl=sh
NetScout - nGeniusONE Alert integration The NETSCOUT integration module enables nGeniusONE to generate enhanced alerts to ServiceNow ITOM with a contextual launch capability for service triage. ServiceNow ITOM users can launch queries into nGeniusONE from any other alert with that context. nGeniusONE dashboards and reports will support the visualization of those alert conditions to and from ServiceNow ITOM	NetScout - nGeniusONE Alert integration
Panopta Integration Integrate with this application for one-directional transfer of events from Panopta to Event Management.	Panopta integration
Splunk Integration Integrate the Splunk Add-on with ServiceNow to enable Splunk users to create incidents and events in Event Management. For more information about the configuration	Splunk Integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Third party connectors	Link
of this integration, see Splunk Add-on for ServiceNow.	
ScienceLogic Integrate with this application for one-directional transfer of events from ScienceLogic to Event Management	ScienceLogic Integration
SevOne Integrate with this application for one-directional transfer of events from SevOne to Event Management	SevOne Integration
Site24X7 Integrate with this application for one-directional transfer of events from Site24x7 to Event Management	Site24x7 Integration

Create a custom pull connector

You can create a customized pull connector that requires a script, connector definition, and connector instance, to retrieve events on behalf of an event source.

Before you begin

Role required: evt_mgmt_admin

About this task

A custom connector script can make remote API calls for a new event source to send events to the ServiceNow instance. You can create a script, and then add it as part of a new Event Management connector definition and connector instance. This procedure automatically creates a JavaScript code file using the name that you enter in the **Name** field as the filename. The JavaScript custom code that accomplishes these actions must:

- Connect to an event monitoring tool.
- Retrieve events from an event monitoring tool.
- Send events to the Event [em_event] table using a web service API. See REST API.

Each connector definition is specific to an event source vendor. The connector definition specifies the MID Server script include that pulls events from the external event source. In addition, the connector definition specifies what connector instance value parameters are needed to connect to the external event source host.

Procedure

1. Create a custom MID Server script include.
   This example uses Groovy (deprecated). In place of a script include in Groovy, it is recommended to use JavaScript.

   a. Implement these methods:

   - `@Override OperationStatus testConnection()`
   - `@Override OperationStatus execute()`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Design the class to extend the *ThirdPartyConnector*.

Example

```
public class HypericConnector extends ThirdPartyConnector
```

c. Import platform classes for event creation, sending, logging, and third-party connector base classes.

Example

```
package com.service_now.mid.probe.tpcon.test

import com.glide.util.Log
import com.service_now.mid.MIDServer
import com.service_now.mid.probe.event.IEventSender
import com.service_now.mid.probe.event.SNEventSender
import com.service_now.mid.probe.tpcon.OperationStatus
import com.service_now.mid.probe.tpcon.OperationStatusType
import com.service_now.mid.probe.tpcon.ThirdPartyConnector
import com.snc.commons.eventmgmt.Event
```

d. Connect to the data collector, such as SCOM or VMware vRealize Hyperic (Hyperic), with the API provided by the collector.

e. Use the *context* parameters to set the event field values that came from a connector instance.

Example

```
def authString = (context.username +
 ':'+ context.password).getBytes().encodeBase64().toString()
def urlStr = 'http://'+context.host + ':'
 + context.parameters.port + '/' + apiFunction
```

f. Implement the *execute* function in the script.
It reads events from the connector using its API to create Event objects and send them to the ServiceNow instance. For example:

Example

```
GPathResult alertsResponse1 =
 readAlerts('hqu/hqapi1/alert/find.hqu?begin=' + lastSignatureStr +
'&end='+ tillStr +'&count=1000000&severity=1&inEscalation=false&notFixed=false');
Event event = new Event()
event.emsSystem = context.name
event.source = "Hyperic"
event.description = it.@reason
event.type = it.@name
...
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
IEventSender eventSender =
 MIDServer.getSingleton(SNEventSender.class)
    for (Event event  : list) {
        eventSender.sendEvent(event)
     }
```

2. Navigate to **MID Server** > **MID Server Script Files** and create a script.
3. Specify the **Parent** field as *Groovy*, complete the form as appropriate, and then click **Submit**.
4. Navigate to **Event Management** > **Integrations** > **Connector Definitions** and create a connector definition.
5. In the **Groovy script to run** field, select the MID Server script file, complete the form as appropriate.
   In addition to **username** or **host**, you can add any other parameter, for example, **port**, and then click **Submit**.
6. Navigate to **Event Management** > **Integrations** > **Connector Instances** and create a connector instance.
7. In the **Connector definition** field, select the connector definition, complete the form as appropriate, and then click **Submit**.
8. To confirm or debug the script, use debug printouts from Groovy to the MID Server log.
9. To monitor incoming events using the custom connector instance, navigate to **ECC** > **Queue** and filter on **ConnectorProbe**.
10. Create a custom connector definition.
    a. Navigate to **All** > **Event Management** > **Integrations** > **Connector Definitions**.
    b. Click **New**.
    c. In the **Name** field, type a descriptive name for the connector definition.
    d. In the **Script type** field, select **Javascript**.
    e. Right-click the form header and select **Save**.
    The **Javascript to run** field is automatically populated with the name of the JavaScript template file.
11. To edit the default JavaScript file, click the information icon ( ) to the right of the JavaScript file name and follow these steps:
    a. Specify the information required for the **testConnection** function.
    This function tests the connector definition to verify if the connection to the target is valid.
    b. Specify the information required for the **execute** function.
    This function retrieves the information from the external source.
    c. Specify the information required for the **retrieveKpi** function.
    This function retrieves the metric data from the external source.
12. **Optional:** You can replace the default JavaScript with your JavaScript code.
13. Fill in the fields, as appropriate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector Definition form

Field	Description
Default schedule	The number of seconds between attempts to receive events.
Bi-directional	Specify to enable the external monitoring system to be updated if the alert is changed. There is default implementation for SCOM.
Alert field identifier	Specify the alert field that, if changed manually, causes the external monitoring system to be updated. **Note:** This field appears only if **Bi-directional** is selected.
Collect metrics	Select to enable the collection of metrics. **Note:** Metrics are collected only for SCOM external event sources.
Metrics collection default schedule (seconds)	Specify the number of seconds of information that must be collected. The default is 10 seconds. **Note:** This field appears only if **Collect metrics** is selected.
Connector Parameters	Specify the parameters to enable communication with an event server.
Connector Definition to MID Server Capabilities	The MID Server name to process events from the event server. If not specified, an available MID Server is used.

14. Click **Submit**.

Configure a pull connector

Configure a pull connector to schedule the frequency of event collection.

Before you begin

Before starting this procedure:

- Locate or define a connector definition.
- Create the credentials to connect to the event source.
- If the connector instance is to be used for the collection of Metric Intelligence metric data, the Metric Intelligence plugin must be activated.

Role required: evt_mgmt_admin

About this task

You can use a connector instance to control the location and manner in which events arrive from external sources. You can optionally select to collect Metric Intelligence information from the external sources that you are connecting to. The Metric Intelligence feature is for the analysis of trends, threshold definitions, and anomaly detection. You can run a test to confirm that connector instance parameters let the MID Server receive events from a supported external event source to the Event [em_event] table. The test also confirms:

- The MID Server is running.
- The host is running on the IP address that is in the **Host IP** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Both the MID Server and the host are running in the same domain.
- The connector instance value fields are valid.
- A connection can be made to the event source, using API, to retrieve events.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New**.
3. Fill in the fields, as appropriate.

Connector Instance Form

Field	Description
Name	A unique name for the connector instance record, such as the name of event source host.
Description	Any optional information that the administrator wants to use to identify this record.
Assignment Group	Designates which team is responsible for managing, maintaining, and responding to issues related to this connector instance.
Connector definition **Note:** Metric Intelligence connectors (such as the AWS CloudWatch and AzureConnector), appear in the base system only if the Metric Intelligence plugin (com.snc.sa.metric) is activated.	The vendor and protocol used to gather events from the external event source. Select the connector definition that matches the source of external events. The options available in the base system are, for example: ○ AWS CloudWatch ○ AzureConnector ○ Datadog ○ HPOM ○ IBM Netcool ○ NagiosXI ○ NNMi ○ OMi ○ SCOM ○ SolarWinds ○ Op5 ○ PRTG ○ vCenter ○ Zabbix For a list of the connector definitions available in the base system, see Configure Event Management connectors.
Host IP	The IP address of the event source host. The system uses this IP address to select the appropriate MID Server for communicating with the event source host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Credential	The record from the Credentials [discovery_credentials] table containing valid credentials to the event source host.
Event collection last run time	The date and time of the most recent event import. The value of this field is automatically populated.
Last event collection status	The status of the last import. The value of this field is automatically populated. ○ **None** - A valid connection has not yet been established. ○ **Success** - A successful connection was established. ○ **Error** - A connection was established. However the external event source was not updated.
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from the external event source. This value cannot be lower than the minimum schedule property, which by default is 120 seconds.
Last error message	Last error message received by the connector. If the test connector fails, an error message appears in this field.
Active	This option appears only after the form has been saved. Do not select this check box until after you test the connection between the MID Server and the connector. Select this check box to enable pulling events from this external event source.
Bi-directional	Select to invoke the bi-directional option. This option is available only when the specified connector definition has been configured with bi-directional values. When selected, this option enables the bi-directional exchange of values to-and-from the external event source. The **Last bi-directional status** option appears only when this option is selected. This option appears only if the connector definition that is selected is SCOM.
Last bi-directional status	The value of this field is automatically populated. The status values are: ○ **None** - A valid connection has not yet been established. ○ **Success** - A successful connection was established. ○ **Error** - A connection was established. However the external event source was not updated. This option appears only if the connector definition that is selected is SCOM.
Metrics collection	Designates collection of Metric Intelligence metric data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: The metrics connector supports working against the SCOM database (MSSQL) that is configured to work with SSL. This option appears only if the connector definition supports metric collection, such as SCOM.
Metrics collection last run time	The date and time of the most recent metric data collection. The value of this field is automatically populated.
Last metrics collection status	Status of the metric data collection activity. The value of this field is automatically populated. This option appears only if the connector definition that is selected supports metric collection (such as SCOM) and if the **Metrics collection** option is selected.
Metrics collection schedule (seconds)	The time, in seconds, to repeat the metric data collection scheduled job. This option appears only if the connector definition that is selected supports metric collection (such as SCOM) and when the **Metrics collection** option is selected.
Metrics database host	Enter either the IP address or the host name of the metrics database host. This option appears only if the connector definition that is selected supports metric collection (such as SCOM).
Connect using a named instance	Select to specify the MSSQL named instance to connect to. This option appears only if the connector definition that is selected is SCOM.
Metrics database named instance	The name of the metrics database instance. This option appears only if the connector definition that is selected supports metric collection (such as SCOM).
Metrics database port	The port used by the metric database. The connection is made using JDBC. The base system value is 1,433. This option appears only if the connector definition that is selected supports metric collection (such as SCOM).
Database login with Windows authentication	Perform database login with the credentials of the log-on user that is defined on the MID Server service. This option appears only if the connector definition that is selected is SCOM.
Metrics collection running	When selected, indicates that the Metric Intelligence collection is running.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	This option appears only if the connector definition that is selected supports metric collection (such as SCOM).
Metrics database credential	Credentials for the metric database. This option appears only if the connector definition that is selected supports metric collection (such as SCOM).
Connector Instance Values	The related list containing connection parameters for the event source host. The list of parameters depends on the selected connector definition.
MID Servers for Connectors	The name of theMID Server to process events from the event server.

4. Right-click the form header and select **Save**.
5. Click **Test connector** to verify that the MID Server can communicate with the external server host.
6. If the test fails, follow the instructions on the page to correct the problem and then run another test. The description of the reason for failure is shown in the MID Server log, according to the type of error, such as Events, Metrics, or Bi-Directional.
7. After a successful test, select the **Active** check box and then click **Update**.
8. Confirm that Event Management processes events using this connector instance by manually creating an event.

Related topics

Configure Event Management connectors

Configure a default MID Server for connectors

You can set a default MID Server for connectors to ensure that there is always a MID Server available to receive external events.

Before you begin

Role required: evt_mgmt_operator

About this task

The default MID Server must have connectivity to all external event sources.

Procedure

1. Enter `sys_properties.list` in the navigation filter to open the System Property [sys_properties] table.
2. Select the `mid.server.connector_default` system property.
3. On the form, fill in the fields.

Task Template form

Field	Description
mid.server.connector_default	Default MID Server for connectors

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Determines the MID Server connectors to use when no MID Server is specified. Must match a MID Server name. ○ **Type**: select `string` from the list ○ **Value**: enter the name of an existing MID Server, for example, SNC MID Server ○ **Location**: System Property [sys_properties] table

Note: You can reset the default MID Server by clearing the *`mid.server.connector_default`* property or by updating the value of this property to a different MID Server. This property is not automatically reset in all scenarios.

4. Click **Update**.

Configure the Dynatrace metrics connector instance

Configure the Dynatrace connector instance to receive Metric Intelligence raw data from the Dynatrace server.

Before you begin

To activate metric collection, ensure that the MID Server that retrieves metrics is configured with the Metric Intelligence extension and that the extension is in **Started** mode. See Manually configure the Metric Intelligence extension.

- Create an access token in Dynatrace with Read Metrics (metrics.read) scope permissions, to enable fetching metric data. This token is used in the connector to call the Dynatrace REST API.
- Configure the Service graph connector for Dynatrace. This is configured in the instance and discovers Dynatrace CIs. For details on this Service graph connector, see Configure Service Graph Connector for Observability - Dynatrace using guided setup.

Role required: evt_mgmt_admin

About this task

By default, the Dynatrace service graph connector discovers CIs and supports metric collection for the following entity types:

- HOST
- APPLICATION
- SERVICE
- PROCESS_GROUP_INSTANCE

Configure other entity metrics for collection when creating event rules for CI binding.

Procedure

1. Navigate to **Event Management** > **Integrations** > **Connector Instances**.
2. Select **New** and create a connector instance.
   For details on the connector instance fields displayed on the page, see Dynatrace connector instance form.
3. Select and hold (or right-click) the form header and select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The connector instance values are added to the form and the parameters that are relevant to the connector appear.

4. **Optional:** In the Connector Instance Values section, you can edit the values of the mandatory Dynatrace parameters.
   For details on the connector instance value parameters, see Dynatrace connector instance value parameters.
5. In the MID Servers for Connectors section, add a MID Server for the Dynatrace connector to use.
6. In the Related Links section, select **Configure metrics to collect**.
   Only metrics with the value **Active=true** are collected by the instance.
7. Configure the metrics you want the instance to collect with the value **Active=true**.
   Metrics with the value **Active=false** are not collected by the instance.
8. **Optional:** Select **New** to create new metrics.
   Configure the fields on the Connector instance metrics - New record (sn_em_connector_instance_metric) page.

Connector instance metrics page

Field	Description
Metric name	Display name for the metric to be configured in the ServiceNow instance.
Metric connector instance	Set by default to **Dynatrace connector instance**. All metrics are collected for the Dynatrace connector instance.
Metric type	The name of the Dynatrace entity type which the metric belongs to.
Metric ID	The Dynatrace metric ID of the metric being collected.
Active	Select for the metrics to collect.

   - Ensure that the CIs related to the metrics exist in the instance.
   - Create an event rule to perform CI binding, as needed. For details, see Create or edit an event rule.
9. Select **Test Connector** to verify the connection between the MID Server and the connector.
10. After a successful test, select the **Active** check box and then select **Update**.

Configure the Nagios metrics connector instance

Configure the Nagios metric connector instance to receive Metric Intelligence raw data from the Nagios server.

Before you begin

To activate metric collection, ensure that the MID Server that retrieves metrics is configured with the Metric Intelligence extension and that the extension is in **Started** mode. See Manually configure the Metric Intelligence extension.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Create basic authentication credentials on ServiceNow instance with Nagios SQL database username and password.
- ServiceNow does not support discovery for Nagios. Therefore, create the CI entry manually in the respective table using the hostname configured on the Nagios platform.

Role required: evt_mgmt_admin

About this task

By default, CI binding will occur for the following entities:

- HOST HTTP Endpoint: The event rule applies to the Nagios metric HTTP if you are running a check that includes command_name as service_http.
- MySQL Instance: The event rule applies to the Nagios metric MySQL if you are running a check that includes command_name as mysql.

If you are running any other checks on HTTP endpoints or MySQL, you can edit the existing event rule to add the command_name. For other entities, configure event rules based on CI binding requirements.

Procedure

1. Navigate to **Event Management** > **Integrations** > **Connector Instances**.
2. Select **New** and create a connector instance.
   For details on the connector instance fields displayed on the page, see Nagios connector instance form.
3. Select and hold (or right-click) the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. **Optional:** In the Connector Instance Values section, you can edit the values of the mandatory Nagios parameters.
   For details on the connector instance value parameters, see Nagios connector instance value parameters.
5. In the MID Servers for Connectors section, add a MID Server for the Nagios connector to use.
6. Ensure that the CIs related to the metrics exist in the instance.
7. Create an event rule to perform CI binding, as needed.
   For details, see Create or edit an event rule.
8. Select **Test Connector** to verify the connection between the MID Server and the connector.
9. After a successful test, select the **Active** check box and then select **Update**.

Configure the Zabbix metrics connector instance

Configure the Zabbix Metric connector instance to receive Metric Intelligence raw data from the Zabbix server.

Before you begin

To activate metric collection, ensure that the MID Server that retrieves metrics is configured with the Metric Intelligence extension and that the extension is in **Started** mode. See Manually configure the Metric Intelligence extension.

Create basic authentication credentials on ServiceNow instance with Zabbix server username and password.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

ServiceNow does not support discovery for Zabbix, so CI entries must be manually created in the appropriate table using the hostname configured in the Zabbix platform.

- Create an Event Rule to bind specific CIs.
- Users can add tag information to Zabbix items, and the relevant tags will be included in the metric-to-CI records. These tags can also be used as filter conditions in the Event Rule for CI binding.

Procedure

1. Navigate to **Event Management** > **Integrations** > **Connector Instances**.
2. Select **New** and create a connector instance.
   For details on the connector instance fields displayed on the page, see Zabbix connector instance form.
3. Select and hold (or right-click) the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. **Optional:** In the Connector Instance Values section, you can edit the values of the mandatory Zabbix parameters.
   For details on the connector instance value parameters, see Zabbix connector instance value parameters.
5. In the MID Servers for Connectors section, add a MID Server for the Zabbix connector to use.
6. Select **Test Connector** to verify the connection between the MID Server and the connector.
7. After a successful test, select the **Active** check box and then select **Update**.

Note: ServiceNow supports only unsigned float and unsigned integer metric types for Zabbix.

Configure the SolarWinds metrics connector instance

Configure the SolarWinds connector instance to receive Metric Intelligence raw data from the SolarWinds server.

Before you begin

To activate metric collection, ensure that the MID Server that retrieves metrics is configured with the Metric Intelligence extension and that the extension is in **Started** mode. See Manually configure the Metric Intelligence extension.

Configure the Service graph connector for SolarWinds. This is configured in the instance and discovers SolarWinds CIs. For details on configuring this connector, see Service Graph Connector for SolarWinds.

Role required: evt_mgmt_admin

About this task

By default, the Solarwinds Metric Connector Collects metrics from the following tables:

- Orion.CPULoad
- Orion.NPM.InterfaceTraffic
- Orion.VolumePerformanceHistory
- Orion.VolumeUsageHistory
- Orion.ResponseTime

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Event Management** > **Integrations** > **Connector Instances**.
2. Select **New** and create a connector instance.
   For details on the connector instance fields displayed on the page, see Solarwinds connector instance form.
3. Select and hold (or right-click) the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. **Optional:** In the Connector Instance Values section, you can edit the values of the mandatory SolarWinds parameters.
   For details on the connector instance value parameters, see Solarwinds connector instance value parameters.
5. In the MID Servers for Connectors section, add a MID Server for the SolarWinds connector to use.
6. In the Related Links section, select **Configure metrics to collect**.
   Only metrics with the value **Active=true** are collected by the instance.
7. Configure the metrics you want the instance to collect with the value **Active=true**.
   Metrics with the value **Active=false** are not collected by the instance.
8. **Optional:** Select **New** to create new metrics.
   Configure the fields on the Connector instance metrics - New record (sn_em_connector_instance_metrics_solarwinds) page.

Connector instance metrics page

Field	Description
Metric name	Name for the metric to be configured in the ServiceNow instance.
Metric connector instance	Set by default to **SolarWinds connector instance**. All metrics are collected for the SolarWinds connector instance.
Additional columns	(Optional) List of additional columns to fetch for CI-related information. For example: DateTime, Node, Resource. If **DateTime** is not provided, the current time is used as the metric timestamp. If **Node** and **Resource** are provided, they are populated accordingly. To reference fields from a parent table, use dot-walking format, such as `Interface.InterfaceName`. If a field should populate a specific standard field (like **Resource**), use an alias with the same name. For example, to use `InterfaceName` as the resource, specify it as `Interface.InterfaceName AS Resource`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Table Name	Name of the table from which metrics are collected.

- Ensure that the CIs related to the metrics exist in the instance.
- Create an event rule to perform CI binding, as needed. For details, see Create or edit an event rule.

9. Select **Test Connector** to verify the connection between the MID Server and the connector.
10. After a successful test, select the **Active** check box and then select **Update**.

> **Note:** Ensure that **Registration Mode** is enabled for the SolarWinds source in the Metric Sources table.

Configure event collection from HP OMi

Configure the HP Operations Manager (OMi) connector instance to receive alerts from the HP OMi server.

Before you begin

Supported versions: 10.01, 10.10, 10.11, 10.12, 10.60, and 10.61.
The OMi connector is not supported on versions below version 10.

Create Basic Auth type of credential to connect to OMI.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a connector instance with the following details:

Field	Value
Name	Descriptive name for the OMi connector.
Description	Description for the use of the OMi event collection instance.
Host IP	Specify the OMi IP address.
Credential	Click in the **Credential** field. Select a Basic Auth credential from the list or click the search icon. Either select the required credentials from the list or click **New** and create the required credentials using Basic Auth type. If you create the credentials, save them using a unique and recognizable name, for example OMiOPS.
Event collection last run time	Last run time value. Automatically updated.
Last event collection status	Last run time status. Automatically updated.
Event collection schedule (seconds)	Frequency in seconds that the system checks for new events from OMi. The default value is 120 seconds.
Last error message	Last error message. Automatically updated.
Connector Definition	Select **OMi**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

3. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.

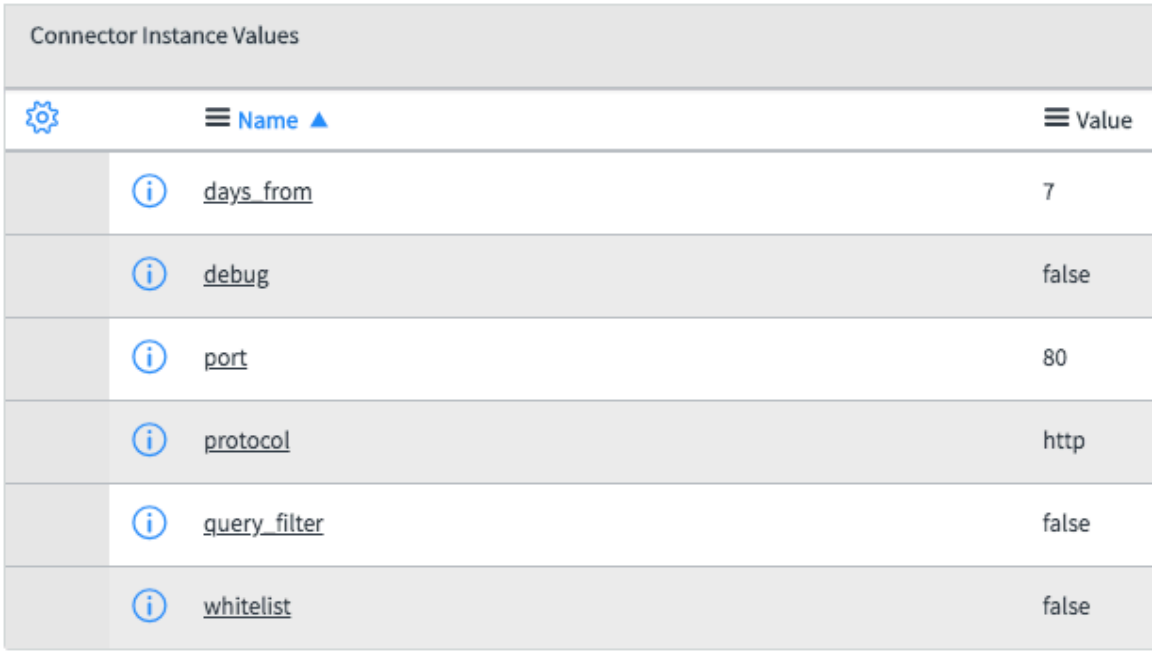

4. In the Connector Instance Values section, specify the OMi values.

   a. **days_from**.
   Specify the number of days for which events must be collected on the first collection cycle. Default 7.

   b. **debug**.
   Default `false`. To enter debug mode, specify `true`.

   c. **port**.
   Default 80.

   d. **protocol**.
   Default `http`.

   e. **query_filter**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Default value is false. If you do want to use this extra filter to query events, specify the event attributes. For example, `assigned_user = "admin" AND title = "My Title"`. The query filter criteria are:

- A filter property that specifies an event attribute, such as, `related_ci`
- A supported operator, for example, `OR` or `AND`
- A value for the filter property, for example, `admin` or `integrator`

Everything from the query_filter field undergoes URL encoding and is then passed as a parameter to the OMi API call. For detailed information on how to build the filter query, see "Filtering by Event Attributes: query" in the HP OMi Extensibility Guide. For example, to query the events for a list of events that are assigned to the admin user, specify:

```
event_list?query=assigned_ user%20EQ%20"admin"
```

To nest query_filter event attributes, specify the hierarchy using square brackets "[ ]".

f. inclusion list.
Default value:`false`. Specify a comma-separated list to include this feature to add attributes that should be collected and added to the Event Management **Additional information** field.

5. Right-click the form header and select **Save**.
6. Click **Test Connector** to verify the connection between the MID Server and the connector.
7. If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the HP OMi server.

8. After a successful test, select the **Active** check box and then click **Update**.

Configure event collection from HPOM

Configure the HPOM connector instance to receive events from HP Operations Manager (HPOM).

Before you begin

Role required: evt_mgmt_admin
Supported version: 08.60.005.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create an HPOM connector instance with the following details:

Field	Value
Name	Specify a unique name for the HPOM connector instance.
Host IP	Specify the HPOM IP address.
Credential	Click in the **Credential** field. Select a credential from the list or click the search icon. Either select the required credentials from the list or click **New** and create the required credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	If you create the credentials, save them using a unique and recognizable name, for example, HPOMOPS.
Schedule (seconds)	The frequency in seconds that the system checks for new events from HPOM.
Description	Type a description for the use of the HPOM event collection instance.
Connector definition	The vendor and protocol used to gather events from the external event source. Select the **HPOM** connector definition.
Last error message	The last error message is automatically updated.
Last run time	The last run time value is automatically updated.
Last run status	The last run time status is automatically updated.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

3. Right-click the form header and click **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. In the Connector Parameters section, specify the values of the mandatory HPOM parameters.
   - **driver**. Specify the driver to be used to make the call. This driver is specific to the type of database you are calling into.
     - For an IBM DB2 Universal driver, specify: com.ibm.db2.jcc.DB2Driver
     - For an Oracle driver, specify: oracle.jdbc.OracleDriver
     - For a Microsoft SQL driver, specify: com.microsoft.sqlserver.jdbc.SQLServerDriver
     - For a MySQL driver, specify: com.mysql.jdbc.Driver
     - For a Sybase driver, specify: com.sybase.jdbc3.jdbc.SybDriver
   - **url**. Specify a URL. The JDBC protocol uses a connection string to establish the authentication and other parameters to establish a connection between the client and the database. Each database has its own connection string format. In the following URLs, replace the variables, for example, <username> and <password>, with your values.
     - For an IBM DB2 Universal URL, specify: [jdbc:db2://sysmvs1.stl.ibm.com:<port number>/<database name>:user=<username>;password=<password>]

       For example: jdbc:db2://sysmvs1.stl.ibm.com:5021/regionsdb:user=example;password=wlrs21.
     - For an Oracle URL, specify: [jdbc:oracle:thin:@<IP address>:<port number>:<database name>]

       For example: jdbc:oracle:thin:@172.31.255.255:4028:OracleMain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For a Microsoft SQL URL, specify: [jdbc:sqlserver://<IP address>;user=<username>;password=<password>]

  For example: jdbc:sqlserver://localhost;user=example;password=w3lrs2.

- For a MySQL URL, specify: [jdbc:mysql://<IP address>/database?user=<username>;password=<password>]

  For example: jdbc:mysql://localhost/database?user=example;password=65wlrs.

- For a Sybase URL, specify: [jdbc:sybase:Tds:<IP address>:<port number>/<database name>?user=<username>;password=<password>]

  For example: jdbc:sybase:Tds:172.31.255.255:4100/SYBMAIN?user=example;password=wlrs1m.

5. Use a network tool, such as ping, to verify that the HPOM server is running and the IP address matches the value in the *Host IP* field.
6. Click **Test connector** to verify the connection between the MID Server and the HPOM connector.
7. If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the HPOM server.

8. After a successful test, select **Active** and then click **Update**.

Related topics

Configure a pull connector

Configure event collection from Hyperic

Configure the Hyperic connector instance to receive events from the VMware vRealize Hyperic server.

Before you begin

Role required: evt_mgmt_admin
Supported version: 5.8.4.0.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a Hyperic connector instance with the following details:

Field	Value
Name	Specify a unique name for the Hyperic connector instance.
Host IP	Specify the Hyperic IP address.
Credential	In the Credentials form, click **New** and create the required credentials. Save the credentials using a unique and recognizable name, for example, HypericOPS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Schedule (seconds)	The frequency in seconds that the system checks for new events from Hyperic.
Description	Type a description for the use of the Hyperic event collection instance.
Connector definition	The vendor and protocol used to gather events from the external event source. Select the **Hyperic** connector definition.
Last error message	The last error message is automatically updated.
Last run time	The last run time value is automatically updated.
Last run status	The last run time status is automatically updated.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

3. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. In the Connector Parameters section, specify the values of the mandatory Hyperic parameters.
   a. **port** Default value 7080.
   b. **protocol** Default value http.
5. Click **Test connector** to verify the connection between the MID Server and the Hyperic connector.
6. If the test fails, follow the instructions that are issued by the error to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the Hyperic server.

7. After a successful test, select the **Active** check box and then click **Update**.

Related topics

Configure a pull connector

Configure event collection from IBM Netcool

Configure the IBM Netcool_V2 connector to receive events from IBM Netcool/OMNIbus Object Servers and Impact Servers. The IBM Netcool_V2 connector uses REST API calls and is bidirectional.

Before you begin

For the IBM Netcool_V2 connector:

- Role required: evt_mgmt_admin
- Perform connector validation on IBM NetCool/OMNIbus version 8.1.0.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

To use IBM Netcool_V2 connector, configure a connector instance for the MID Server.

Note: The IBM Netcool connector is deprecated and only IBM Netcool_V2 connector is supported.

Procedure

1. Navigate to **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a connector instance with the following details:

Field	Value
Name	Specify a unique name for the Netcool connector instance.
Description	Type a description for the use of the Netcool event collection instance.
Host IP	Specify the Netcool IP address.
Credential	In the Credentials form, click **New**, and create the required credentials. Save the credentials using a unique and recognizable name, for example, NetcoolOPS.
Event collection last run time	The last run time value is automatically updated.
Last event collection status	The last run time status is automatically updated.
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from Netcool. The default value is 120 seconds.
Last error message	The last error message is automatically updated.
Connector definition	Select IBM Netcool_V2 to gather events from the external event source.
MID Servers (MID Server for Connectors section)	Specify a MID Server that is up and is valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section. If no MID Server is specified, an available MID Server that has a matching IP range is used.

3. Right-click the form header and click **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. In the Connector Parameters section, specify the value of the required Netcool parameters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector	Required parameters
IBM Netcool_V2	○ *additional_alert_columns*: com ma separated list of columns to include in event additional information from net cool alert. Specify here if the desired columns are not already part of event. ○ *debug*: Boolean flag to enable debug logs. Default: false. ○ *initial_sync_in_days*: The number of days of data fetched on first pull. The de fault value is 3, but can be changed accord ing to your requirements. ○ *message_key_delimiter*: Separate hi erarchical parts of a message key, enabling structured namespacing. ○ *port*: The port number on which IBM Net cool is running. For example: 8080. ○ *Protocol*: Protocol to use when connect ing to the Netcool server. Default: https.

5. Verify that Netcool servers are up and running:
   For the IBM Netcool_V2 connector, use this REST API call to verify that the Netcool API is up and running: `http://hostname:port/objectserver/restapi/alerts/status`
6. Right-click the form header and click **Save**.
7. Click **Test connector** to verify the connection between the MID Server and the Netcool connector.
8. If the test fails, follow the instructions that are issued by the error to correct the problem and then run another test.
9. After a successful test, select the **Active** check box and then click **Update**.

Related topics

Configure a pull connector

Configure event collection from an Icinga2 connector

Configure the Icinga 2 (Icinga) connector instance to receive events while monitoring your network resources.

Before you begin

Role required: evt_mgmt_admin

Supported version: 2.4.1.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New**.
3. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector Instance form

Field	Description
Name	Descriptive and unique name for the Icinga connector.
Description	Description for the use of the Icinga event collection instance.
Connector definition	Name of the required connector definition. Select **Icinga2**.
Host IP	IP address where Icinga is installed.
Credential	Permission to connect to Icinga. Click the search icon in the **Credential** field. Use Basic Authentication credentials. Either select the required credentials from the list or click **New** and create the required credentials on the Credentials form. If you create the credentials, save them using a unique and recognizable name, for example, ICINGA2CHCK.
Event collection last run time	Last run time value. Automatically updated.
Last event collection status	Last run status. Automatically updated.
Event collection schedule (seconds)	Frequency, in seconds, that the system checks for new events from Icinga. The default value is 120 seconds.
Last error message	Last error message. Automatically updated.

4. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.

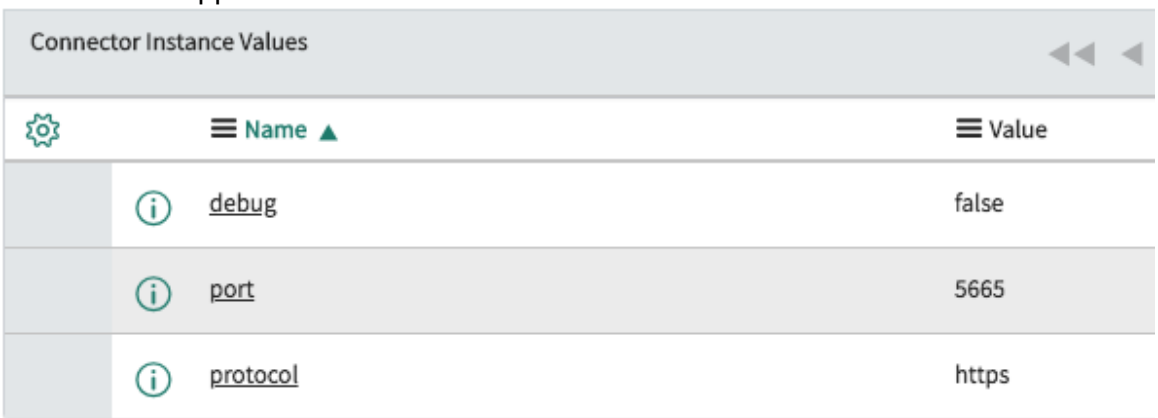

   a. debug: Display debug messages. Default value: false. Specify true to see debug messages.

   b. logPayloadForDebug: Display payload related debug messages. Default value: false. Specify true to see payload related debug messages.

5. In the Connector Instance Values section, verify and where required, modify the default connector instance values.

Field	Description
debug	Display debug messages. Default value: `false`. Specify `true` to see debug messages.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
port	Number of the connector port. Default value: `5665`.
protocol	Type of protocol. Default protocol type: `https`.
alert_type	Provide the Icinga alert type details. There are 2 types of an alert in Icinga2: SOFT and HARD SOFT: Connector instance will pull a soft alert. HARD: Connector instance will pull a hard alert.

6. **Optional:** In the MID Servers for Connectors section, specify a MID Server that is up and is valid.
   If no MID Server is specified, an available MID Server that has a matching IP range is used.

   You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.
7. Right-click the form header and select **Save**.
8. Test the connection between the MID Server and the Icinga connector.
   a. Click **Test Connector**.
      If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

      Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the external monitoring tool.

   b. After a successful test, select the **Active** check box.
9. Click **Update**.

Configure event collection from Logicmonitor

The Logicmonitor pull connector sends information from Event Management to Logicmonitor. It sends responses received from a Push connector in a bi-directional environment to Logicmonitor.

Before you begin
Role required: evt_mgmt_admin

About this task
Starting from the Xanadu release, the OOTB (Out-Of-The-Box) rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the**Apply additional matching rules**check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the connector.

Note: This feature applies only to active event rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to**All** > **Event Management** > **Event Connectors (Pull)** > **Connector Instances**.
2. Click**New**.
   The**Connector Instance New Record**page appears.
3. Enter values in the required fields, which are indicated by a red asterisk.

Field	Description
Name	Descriptive name for the connector.
Connector definition	Search for and select**Logic Monitor**.
Credential	Select**Basic Auth**for basic authentication. Pro vide access_id and access_key from Logic monitor's token based API authentication in the credentials as the basic authentication username and password.
Host IP	The Logicmonitor IP address.

When adding bi-directional functionality for a Push connector, the Connector Instance name must match the Source Instance field of the event generated by that source.

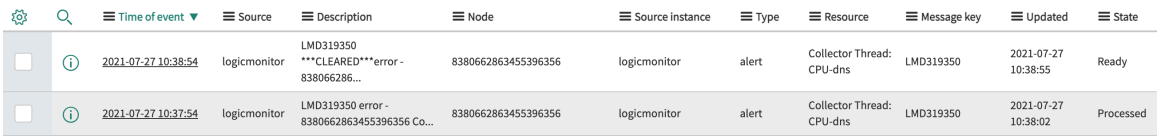

Time of event	Source	Description	Node	Source instance	Type	Resource	Message key	Updated	State
2021-07-27 10:38:54	logicmonitor	LMD319350 ***CLEARED***error - 838066286...	8380662863455396356	logicmonitor	alert	Collector Thread: CPU-dns	LMD319350	2021-07-27 10:38:55	Ready
2021-07-27 10:37:54	logicmonitor	LMD319350 error - 8380662863455396356 Co...	8380662863455396356	logicmonitor	alert	Collector Thread: CPU-dns	LMD319350	2021-07-27 10:38:02	Processed

The source of events is**logicmonitor**so you should create an instance with the name`logicmonitor`.

All > Active = true

Name	Active	Event collection last run time	Connector definition	Last event collection signature	Last event collectio
logicmonitor	true	2021-07-27 11:21:41	Logic Monitor		Success

Note: The Update queue business rule in the Alert table identifies each manual update of the alert and updates the connector queue accordingly. By default, changes to all alert fields are tracked.

4. Select the **Bi-directional** and **Active** check boxes, and then click **Save**.
   The LogicMonitor pull connector does not retrieve events or populate the Events [em_event] table.

Configure event collection from NNMi

Configure the HP Network Node Manager i (NNMi) connector instance to receive events while monitoring your network resources.

Before you begin

Supported version: NNMi 10.30 and later, including NNMi 25.5.

Note: The connector supports NNMi versions 10.30.653 and 25.5. If you use another version and encounter issues, contact ServiceNow Support.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New**.
3. On the form, fill in the fields.

Connector Instance form

Field	Description
Name	Descriptive and unique name for the NNMi connector.
Description	Description for the use of the NNMi event collection instance.
Connector definition	Name of the required connector definition. Select **NNMi**.
Host IP	IP address where NNMi is installed.
Credential	Permission to connect to NNMi. Click the search icon in the **Credential** field. Either select the required credentials from the list or click **New** and create the required credentials on the Credentials form. If you create the credentials, save them using a unique and recognizable name, for example, NNMiCHCK. The selected credentials use basic authentication.
Event collection last run time	Last run time value. Automatically updated.
Last event collection status	Last run status. Automatically updated.
Event collection schedule (seconds)	Frequency, in seconds, that the system checks for new events from NNMi. The default value is 120 seconds.
Last error message	Last error message. Automatically updated.

4. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.

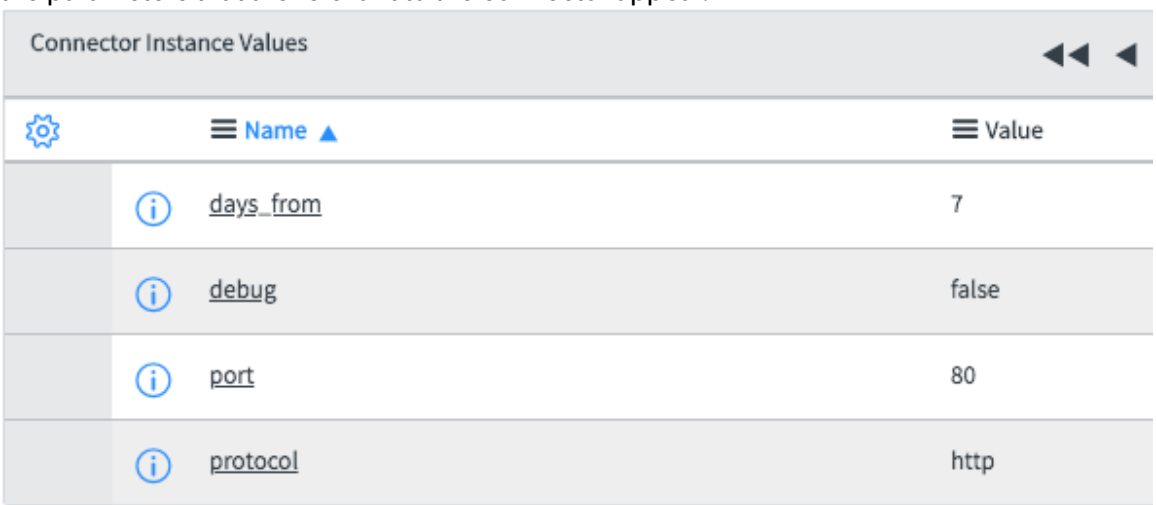

5. In the Connector Instance Values section, verify and where required, modify the default connector instance values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
days_from	Number of days for which events must be collected at the first collection cycle. Default value: `7` days. **Note:** The maximum number of past events that can be collected is 3,000. The first time that the connector runs, it collects all events (up to the 3,000 limit) that were created from the `days_from` date until the current date.
debug	Display debug messages. Default value: `false`. Specify `true` to see debug messages.
port	Number of the connector port. Default value: `80`.
protocol	Type of protocol. Default protocol type: `http`.

6. **Optional:** In the MID Servers for Connectors section, specify a MID Server that is up and is valid.
   If no MID Server is specified, an available MID Server that has a matching IP range is used.

   You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.
7. Right-click the form header and select **Save**.
8. Test the connection between the MID Server and the NNMi connector.

   a. Click **Test Connector**.
   If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the external monitoring tool.

   If you encounter the error `"TypeError: Cannot read property 'inc:getIncidentsResponse' from undefined`, refer to the Known Error article KB0262185 for the updated script and instructions.

   b. After a successful test, select the **Active** check box.
9. Click **Update**.

Configure event collection from NagiosXI

Configure the NagiosXI connector instance to receive events from the Nagios Core monitor.

Before you begin

Supported version: 5.5.2.
The NagiosXI connector instance requires a credential that lets the instance access NagiosXI accounts. You can use an existing credential or create a new one. For more information, see Create Nagios XI server credentials.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New**.
3. On the form, fill in the fields.

Field	Description
Name	Descriptive name for the Nagios connector.
Description	Description for the use of the NagiosXI event collection instance.
Host IP	Specify the Nagios XI IP address.
Credential	Select the credential with basic authentication that you created for this connector. For more information, see Create Nagios XI server credentials. Ensure that the user password contains the NagiosXI user API key, for example 04lquEPqf4JimWCm8RWbJokOpW8LYBUfEvJp9OSHSRYe4QDrHPFndYbV
Event collection last run time	The last run time value is automatically updated.
Last event collection status	The last run time status is automatically updated.
Event collection schedule (seconds)	The frequency, in seconds, that the system checks for new events from Nagios. The default value is 120 seconds.
Last error message	The last error message is automatically updated.
Connector Definition	Select **NagiosXI**.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

4. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
5. In the Connector Instance Values section, specify the Nagios values.
   - **initial_sync_in_days**. Specify the number of days for which events must be collected at the first collection cycle. Default value `7` days.
   - **port**. Default value `80`.
   - **protocol**. The default protocol type is `http`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **time_zone**. The default time zone is `GMT`; update this value to the Nagios timezone. For example, if Nagios is sending the events according to the EDT time zone, update this value to `GMT - 04 : 00` (abbreviations such as **EDT** are not supported).
- **severity_by_state_type**. The default value is `false`. When set to `true`, the state_type is used along with the state to determine the severity which creates less critical events as shown in the following mapping table.

Mapping Table

_state_type(Connect definition parameter)	State (Nagios JSON payload attribute)	State_type (Nagios JSON payload attribute)	ServiceNow Severity
false	2	0/1	1 (Critical)
false	1	1	1(Critical)
false	1	0	2(Major)
false	0	0/1	5(OK)
true	0	0/1	5(OK)
true	1	0	4 (Warning-soft -> Warning)
true	2	0	2(Critical-soft -> Major)
true	3	0	4(Unknown-soft-> Warning)
true	1	1	3(Warning-hard -> Minor)
true	2	1	1(Critical-hard -> Critical)
true	3	1	3(Unknown-hard -> Minor)

6. Right-click the form header and select **Save**.
7. Click **Test Connector** to verify the connection between the MID Server and the connector.
8. If the test fails, follow the instructions that the error issues to correct the problem and then run another test.
   If this message appears: `Connection test failed: Failed to connect to Nagios on test connector. Invalid API Key`, then enter the API Key for the specific user to be able to read the Nagios events.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the Nagios Core monitor.

9. After a successful test, select the **Active** check box and then click **Update**.

Create Nagios XI server credentials

Create credentials to access Nagios XI server.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Click **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Basic Auth Credential**.
4. On the form, fill in the fields.

Basic Auth Credentials form

Field	Description
Name	Enter a unique and descriptive name for this credential. For example, you might call it `Nagios authentication`.
User name	Name to identify the user.
Password	Nagios XI user API key. For example:`04lquEPqf4JimWCm8RWbJokOpW8LYBUfEvJp`
Active	Option to enable the use of this credential.
Order	The order (sequence) in which the platform tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order. Default value: `100`

5. Click **Submit**.

Result
The credential for use with the Nagios XI server connector is created.

Related topics
Credentials and connection information

Configure event collection from OBM

Configure the Operation Bridge Manager (OBM), also known as OMi v2, connector instance to receive alerts from the OBM server. The OBM connector script OBM v2 will be available after installing the Event Management Connectors app.

Before you begin
Role required: evt_mgmt_admin

Supported versions: 10.01, 10.10, 10.11, 10.12, 10.60, 10.61, and 11.01.

About this task
Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Event Connectors (Pull)** > **Connector Instances**.
2. Click **New** and create a connector instance with the following details:

Field	Value
Name	Descriptive name for the OBM connector.
Description	Description for the use of the OBM event collection instance.
Host IP	Specify the OBM IP address.
Credential	Select the required credentials from the list or click **New** and create the required credentials. If you create the credentials, save them using a unique and recognizable name, for example OBMOPS.
Event collection last run time	The last run time value is automatically updated.
Last event collection status	The last run time value is automatically updated.
Event collection schedule (seconds)	Frequency in seconds that the system checks for new events from OBM. The default value is 120 seconds.
Last error message	The last error message is automatically updated
Connector Definition	Select **OBM V2**. **Note:** This script will be available after installing the Event Management connector scoped app.
MID Servers (MID Server for Connectors section)	Specify the MID Server to run the OBM connector instance. You can configure several MID Servers. If the first server is down, the next MID Server If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section. If no MID Server is specified, an available MID Server that has a matching IP range is used.

3. Right-click the form header and select **Save**.
   The parameters that are relevant to the OBM connector appear.
4. In the Connector Instance Values section, specify the OBM values.

   a. **days_from**.
   Specify the number of days for which events must be collected on the first collection cycle. The default value is `7` days and data can only be pulled a maximum of 7 days from the OBM system.

   b. **debug**.
   The default value is `false`. To enter debug mode, change the value to `true`.

   c. **port**.
   The default value is `80`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. **protocol**.
The default value is `http`.

e. **query_filter**.
The default value is false. If you do want to use this extra filter to query events, specify the event attributes. For example, `assigned_user = "admin" AND title = "My Title"`. The query filter criteria are:

- A filter property that specifies an event attribute, such as, `related_ci`
- A supported operator, for example, OR or AND
- A value for the filter property, for example, `admin` or `integrator`

Everything from the query_filter field undergoes URL encoding and is then passed as a parameter to the OBM API call. For detailed information on how to build the filter query, see "Filtering by Event Attributes: query" in the OBM Extensibility Guide. For example, to query the events for a list of events that are assigned to the admin user, specify:

`event_list?query=assigned_ user%20EQ%20"admin"`

To nest query_filter event attributes, specify the hierarchy using square brackets "[ ]".

f. **inclusion list**.
The default value is `false`. Specify a comma-separated list to include this feature to add attributes that should be collected and added to the Event Management **Additional information** field.

5. Right-click the form header and select **Save**.
6. Click **Test Connector** to verify the connection between the MID Server and the connector.
7. If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify that the OBM server is running and the IP address matches the value in the **Host IP** field.

8. After a successful test, select the **Active** check box and then click **Update**.

Configure OP5 or OP5_v2 connector

Configure the OP5 or OP5_v2 Monitor connector instance to receive alerts from an OP5 Monitor source.

Before you begin

The OP5_v2 connector is supported in Event Management connectors version 2.3.0 - August 2022, available from the ServiceNow Store.

Supported version:

- OP5: 7.3.15
- OP5_v2: 8.3.8

Role required: evt_mgmt_admin

About this task

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified),

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a connector instance with the following details:

Field	Value
Name	Descriptive name for the OP5 or OP5_v2 connector.
Description	Description for the use of the OP5 or OP5_v2 event collection instance.
Host IP	Specify the OP5 or OP5_v2 IP address.
Credential	Either select the required credentials from the list or click **New** and create the required credentials. If you create the credentials, save them using a unique and recognizable name, for example, OP5OPS.
Event collection last run time	The last run time value is automatically updated.
Last event collection status	The last run time status is automatically updated.
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from OP5 or OP5_v2.
Last error message	The last error message is automatically updated.
Connector definition	Select **OP5** or **OP5_v2**.
Connector Instance Values section	The parameters that are specific to **OP5** or **OP5_v2** display here when the form is saved.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

3. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. In the Connector Instance Values section, specify the OP5 or OP5_v2 values.

   a. **days_from**.
   Specify the number of days for which events must be collected at the first collection cycle. Default value 7 days.

   b. **debug**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Displays debug messages. Default value: false. Specify true to see debug messages

c. **logPayloadForDebug**.
Displays payload related debug messages. Default value: false. Specify true to see payload related debug messages.

d. **port**.
Default value `443`.

e. **state_types**.
Specify filters for event collection based on the **state_types** attribute. You can filter events based on state type. Specify the type, for example, `soft` or `hard`.

5. Right-click the form header and select **Save**.
6. Click **Test Connector** to verify the connection between the MID Server and the connector.
7. If the test fails, follow the instructions that the error issues to correct the problem and then run another test.
8. After a successful test, select the **Active** check box and then click **Update**.

Configure Opsview_v2 connector

Configure the Opsview_V2 connector instance to receive alerts from an Opsview Monitor source.

Before you begin

Ensure you have the supported version of the Opsview_V2 connector: 6.7.0.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a connector instance with the following details:

Field	Value
Name	Descriptive name for the Opsview connector.
Description	Description for the use of the Opsview event collection instance.
Host IP	Specify the Opsview IP address.
Credential	Either select the required credentials from the list or click **New** and create the required credentials. If you create the credentials, save them using a unique and recognizable name, for example, OpsviewOPS. **Note:** User should have the CONFIGUREVIEW role with access to 'VIEW ALL, CHANGE NONE'.
Event collection last run time	The last run time value is automatically updated.
Last event collection status	The last run time status is automatically updated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from Opsview.
Last error message	The last error message is automatically updated.
Connector definition	Select **Opsview_v2**.
Connector Instance Values section	The parameters that are specific to Opsview_V2 display here when the form is saved.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

3. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. In the Connector Instance Values section, specify the Opsview_v2 values.

   a. **days_from**.

   Specify the number of days for which events must be collected at the first collection cycle. Default value 7 days.

   b. **debug**.
   Display debug messages. Default value: false. Specify true to see debug messages.

   c. **logPayloadForDebug**.
   Display payload related debug messages. Default value: false. Specify true to see payload related debug messages.

   d. **eventtype_filter**.
   Specify filters for event collection based on the eventtype attribute. Specify one of these values 0,1,2,3.

   e. **port**.

   Default value is 80.

   .

   f. **protocol**.
   The default protocol type is HTTPS.

5. Right-click the form header and select **Save**.
6. To verify the connection between the MID Server and the connector, select **Test Connector**.
7. If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If this message appears: `Connection test failed: Failed to connect to Opsview on test connector. Invalid API Key`, then enter the API Key for the specific user to be able to read the Opsview events.

Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the Opsview Monitor source.

8. After a successful test, select the **Active** check box and then select **Update**.

Configure PRTG connector

Configure the PRTG connector instance to receive alerts from a Paessler PRTG Network Monitor source.

Before you begin

Role required: evt_mgmt_admin
Supported version: 21.3.71.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a connector instance with the following details:

Field	Value
Name	Descriptive name for the PRTG connector.
Description	Description for the use of the PRTG event collection instance.
Host IP	Specify the PRTG IP address.
Credential	Either select the required credentials from the list or click **New** and create the required credentials. If you create the credentials, save them using a unique and recognizable name, for example, PRTGOPS.
Event collection last run time	The last run time value is automatically updated.
Last event collection status	The last run time status is automatically updated.
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from PRTG.
Last error message	The last error message is automatically updated.
Connector definition	Select **PRTG**.
Connector Instance Values section	The parameters that are specific to **PRTG** display here when the form is saved.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
4. In the Connector Instance Values section, specify the PRTG values.

   a. **add_columns**.
   Specify which columns to be returned by the connector. Separate the names with a comma and without a space between the names. The output is stored in the **Additional information** field of the event. Default value `false`. If `false` is specified, only the default list of columns is retrieved: parent, name, type, objid, tags, message, status, priority, datetime, actions, baselink, basetype, modifiedby.

   Note: Only columns names that are applicable to the **Message** content type can be specified.

   b. **date_format**.
   The format date and time. Default format M/d/yyyy/ h:mm:ss a. If you receive an event whose date is in a different format, modify this value to match the format of the incoming event. If you do not, the event will not be processed correctly.

   For example, an event might arrive on June 27, 2019 at 11:25 AM with a listed date of *`2019/06/27/ 11:25:00 a`*. Modify the **date_format** value to *`yyyy/M/d/ h:mm:ss a`* to match the format of the received event.

   In the **date_format**, `a` represents AM, and `p` represents PM.

   c. In the **days_from** field, specify the number of days for which events must be collected at the first collection cycle.
   Default value `7` days.

   d. **debug**.
   Default value `false`.

   e. **port**.
   Default value `80`.

   f. **protocol**.
   The default protocol type is `http`.

   g. **prtg_api_fetched_events_num**.
   Maximum number of events, per query, that can be fetched from the PRTG client into the MID Server. Default number: `50,000`.

   h. **time_zone**.
   The default time zone is the `GMT` time zone; update this value to the PRTG time zone. For example, if PRTG is sending events according to the EDT time zone, update this value to `GMT - 4` (abbreviations such as EDT are not supported).

5. Right-click the form header and select **Save**.
6. Click **Test Connector** to verify the connection between the MID Server and the connector.
7. If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If this message appears: `Connection test failed: Failed to connect to PRTG on test connector. Invalid API Key`, then enter the API Key for the specific user to be able to read the PRTG events.

Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to PRTG.

8. After a successful test, select the **Active** check box and then click **Update**.

Apache Kafka Consumer Connector

The Apache Kafka Consumer connector instance enables you to create events from messages collected from the Apache Kafka topic as a JSON payload that contains essential information in a data block.

Apache Kafka is a distributed publish-subscribe messaging system that receives data from disparate source systems and makes the data available to target systems in real time.

You can use this connector to create ServiceNow events from Kafka messages. Messages in the Kafka topic must be in JSON format.

Like other message broker systems, Kafka facilitates the asynchronous data exchange between processes, applications, and servers.

Kafka software runs on one or more servers. Each node, called a broker, in a Kafka cluster helps producer applications write data to topics, which are where records are stored and published. Consumer applications read from topics. Topics are divided into partitions to make them more manageable and Kafka guarantees strong ordering for each partition.

Configure the Apache Kafka Consumer connector

Configure the Apache Kafka Consumer connector instance to create events from streaming messages collected by the Apache Kafka connector.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Select **New**.
3. On the form, fill in the fields.

Connector instance form

Field	Value
Name	Descriptive and unique name for the Kafka Consumer connector.
Description	Description to be used by the Kafka Consumer event collection instance.
Connector definition	Name of the required connector definition, which in this case should be Kafka Consumer.
Host IP	The host IP.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Note: This field must contain a value to complete the creation process so the placeholder 1.1.1.1 should be entered as a temporary value.
Event collection last run time	This field is automatically set to the last runtime value.
Last event collection status	This field is automatically set to the last runtime status.
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from the Kafka Consumer. The default value is 60 seconds.
Last error message	This field is automatically set to the last error message.

4. For PLAINTEXT or SASL_PLAINTEXT, in the **Credential** field, create a Basic authentication credential.
   Use the **Kafka SSL credential** type for a new Kafka pull connector instance with **SASL_PLAINTEXT** authentication. Go to step 5.

   Use the **Kafka SSL credential** type when the Kafka server's **sasl.enabled.mechanisms** or **sasl.mechanism** parameter's value don't equal **PLAIN**. Go to step 5.

   a. Select the Search icon next to the **Credential** field.

   b. On the Credentials screen, select **New**.

   c. From the list of available credentials, select **Basic Auth Credentials**.

   d. On the Basic Authentication screen, give information for either PLAINTEXT authentication, which doesn't require credentials, or SASL_PLAINTEXT authentication, which requires credentials.
      - For PLAINTEXT authentication, in the **Name** field, type any value and leave the **User Name** and **Password** fields empty.
      - For SASL_PLAINTEXT authentication, give a name for the credential in the **Name** field, preferably prefaced by SASL. Provide the user name and password In the **User Name** and **Password** fields.

   e. Select **Submit**.

5. For a SASL_PLAINTEXT, SSL, or SASL_SSL authentication credential, in the **Credential** field, create a Kafka SSL credential.

   a. Select the Search icon next to the **Credential** field.

   b. On the Credentials screen, select **New**.

   c. From the list of available credentials, select **Kafka SSL Credentials**.

   d. On the Kafka SSL Authentication screen, fill in the Kafka SSL credentials fields.
      For the Kafka SSL credentials fields descriptions, see Kafka SSL credentials fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you clear the **Disable hostname verification** field, then in the **Additional Kafka consumer properties** field, enter the value of the *ssl.endpoint.identification.algorithm* parameter the same as of the Kafka server parameter value. For example:

```
ssl.enabled.protocols=TLSv1.2,TLSv1.1,TLSv1;
sasl.jaas.config=org.apache.kafka.common.security.plain.Pl
ainLoginModule required username=<username>
 password=<password>;
sasl.enabled.mechanisms=PLAIN;ssl.endpoint.identification.
algorithm=;
```

e. Select **Submit**.

6. Right-click the form header and select **Save**.
7. Verify the connector instance values based on your Kafka setup and the JSON payload from the Kafka topic.

   Change the connector instance values if necessary.

Connector instance Values table

Field	Description
authentication_type	The Kafka Consumer authentication type. The currently supported values are: ○ PLAINTEXT (no authentication) (Default) ○ SASL_PLAINTEXT (basic authentication – user name and password)
bootstrap_servers	The servers that establish the connection with the Kafka cluster. Values in this field should be in the form host1:port1, host2:port2, and so on (comma-separated). This field is required.
consumer_group_name	The Consumer group name. If the same Kafka topic is being consumed from two different instances, use different consumer group names so all events are captured for both instances. This field is required.
time_of_event_field	The name of the field in the JSON payload/ message that includes the time of the event.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	If a field in the JSON payload/message includes the time of the event, that value should be the value for this parameter. If no time field value is given, then the received time of the event in the ServiceNow instance is set as time_of_event. The time_of_event field should be a first-level field in the JSON payload/message. Nested field names aren't allowed.
timezone	The time zone of the time_of_event field in the JSON payload/message. The timezone value isn't used when the time_of_event field is empty. For example, if the time_of_event in the payload/message is in the IST time zone, the value would be GMT+0530. Default: GMT.
date_format	The date_time format of the time_of_event field in the JSON payload/message. If no time field value is given, then the time when the event was received in the instance in GMT time is used. The date_time value isn't used when the time_of_event field is empty. Default: yyyy-MM-dd HH:mm:sss.
debug	Displays debug messages. Default value: false. Specify true to see debug messages.
logPayloadForDebug	Displays payload related debug messages. Default value: false. Specify true to see payload related debug messages.
topic	The Kafka topic from which the messages are fetched.

8. In the **MID Servers for Connectors** section, specify a MID Server that is up and is valid.

   Note: You can configure several MID Servers. If the first server is down, the next MID Server is used. If that MID Server isn't available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the **MID Servers for Connectors** section.

   If you do not specify a MID Server, an available MID Server that has a matching IP range is used.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. Right-click the form header and select **Save**.
10. Test the connection between the MID Server and the Kafka Consumer connector.
    a. Select **Test Connector**.
    b. If the test connection fails, verify whether the credential is valid, and that the network is connected from the MID Server to the Kafka broker.

    Note: Kafka topic name validation occurs only in **Test Connector** validation.

11. After a successful test, make the connector instance active by selecting the **Active** check box.
12. Select the **Update** button.

What to do next
After the connector is created, you must map the fields. For more information, see Map Kafka message payload attributes to alert fields.

Map Kafka message payload attributes to alert fields

Map Kafka message attributes to alert fields to make alerts based on the messages more meaningful. Use event field mapping to map Kafka severity values to appropriate ServiceNow values.

Before you begin
Role required: evt_mgmt_admin

About this task
Use event field mappings to provide more comprehensive information in an event alert. Substitute values from the event field mapping rule into the event. Transform information in events to populate specified alert field values and compose alert fields from various values to generate alerts for tracking and remediation. If needed, use both event field mappings and alert rules.

The following table compares a message received from a topic and a flattened payload in the event **Additional info** field. The message received from a topic is flattened as a single field and appears in the event **Additional info** field. Use these flattened messages to create event rules and mappings.

Actual versus flattened payload

Actual Payload inserted into a topic	Flattened payload in the Additional Info field

<pre>{
"version": "v0.1.0",
"name": "fibsrc",
"id": "mnop6789",
"trigger_description": "Triggers on duration_ms goes above 5K",
"status": "TRIGGERED",</pre>
	<pre>{
	"flattened.alert_type": "on_true",
	"flattened.description": "Current value (12.296815 k) greater than threshold value (5 k)",
	"flattened.id": " mnop6789",</pre>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Actual versus flattened payload (continued)

Actual Payload inserted into a topic	Flattened payload in the Additional Info field

<pre>{
"alert_type": "on_true",
"summary": "Triggered: fibsrc",
"description": "Current value (12.296815 k) greater than threshold value (5 k)",
"operator": "greater than",
"threshold": 5000,
"groups": [
{
"Group": "GroupA",
"Result": 12296.815478000368
}
],
"trigger": {
"Result": 12296.815478000368
}
}</pre>
	<pre>"flattened.kafkaKey": "null",
	"flattened.kafkaOffset": "28",
	"flattened.kafkaPartition": "0",
	"flattened.kafkaTopicName": "topic123",
	"flattened.name": "fibsrc",
	"flattened.operator": "greater than",
	"flattened.status": "TRIGGERED",
	"flattened.summary": "Triggered: fibsrc",
	"flattened.trigger_description": "Triggers on duration_ms goesa above 5K",
	"flattened.version": "v0.1.0",
	"flattened.version": 5000,
	"flattened.groups.0.Group": "GroupA",
	"flattened.groups.0.Result": 12296.815478000368,
	"flattened.trigger.Result": 12296.815478000368
	}</pre>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. To map the Kafka message attributes to the predefined alert fields to generate alerts by creating event field mappings.
   a. Navigate to **All** > **Event Management** > **Rules** > **Event Field Mapping**.
   b. In the Event Field Mapping window, select **New**.
   c. In the Name field, enter a name for the field mapping.
   d. In the Source field, enter **Kafka Consumer**.
   e. In the Mapping type field, select the appropriate mapping type
   f. For a list of the available mapping types, see Event Field Mappings.
2. To map the Kafka message attributes to the predefined alert fields to generate alerts by creating event rules.
   a. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
   b. Select **New** .
   c. In the Name field, enter a name for the event rule.
   d. In the Source field, enter **Kafka Consumer**.
   e. Select **Event Filter**.
   f. Enter the appropriate filter according to your requirement.
   g. Select **Transform and compose alert output**.

      When creating event rules, the event fields are displayed in the Event Rules Transform and Compose Alert Output tab. Because the Kafka message varies from user to user, you must transform the event fields to the alert fields to see the details in the alert.

      Note: Because the severity value is required to create an alert, you must map the **Severity** field using event rules or event field mappings. If the Severity field is empty, the event isn't processed and enters into an error state.

   h. Map the fields of the event and **Additional Info** field to the fields of the alert.
   i. Select **Binding**.
   j. Enter the appropriate binding according to your requirement.
   k. Select **Submit**.
      For more information, see Create or edit an event rule.
3. In the All Events screen, check whether any of the events from Kafka are in an error state due to a missing event rule and provide missing rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **All** > **Event Management** > **All Events**.

b. In the All Events table, open any Kafka consumer event that shows Error as the state in the **Severity** column.

   The message received from the Kafka consumer can vary. The connector cannot know the attributes that will be received. You must map the **Severity** field or other event fields to the fields in the message to ensure that alert fields are correctly populated.

c. On the All Events screen, create a new event rule by selecting **Create Event Rule**.
   For more information on creating event rules, see Event Rules.

   If no event field mappings exist to handle the new type of event, create event field mappings.
   For more information, see Create event field mappings.

Result

Your Apache Kafka consumer connector is configured. When new messages arrive into the Kafka topic, they are received into the ServiceNow instance as events, and alerts are then created from those events.

Configure the Kafka metrics connector instance

Configure the Kafka metric consumer connector instance to read message send to Kafka server over topic.

Before you begin

To activate metric collection, ensure that the MID Server that retrieves metrics is configured with the Metric Intelligence extension and that the extension is in **Started** mode. See Manually configure the Metric Intelligence extension.

- For PLAINTEXT in the Credential field, create a Basic authentication credential.

  Select the Search icon ( ) next to the Credential field.

- ServiceNow does not support discovery for Kafka. Therefore, create the CI entry manually in the respective table.

Role required: evt_mgmt_admin

Procedure

1. On the Credentials screen, select **New**.
2. From the list of available credentials, select **Basic Auth Credentials**.
3. On the Basic Authentication screen, give information for PLAINTEXT authentication, which doesn't require credentials.
   For PLAINTEXT authentication, in the **Name** field, type any value and leave the **User Name** and **Password** fields empty.
4. Select **Submit**.
5. For a SASL_PLAINTEXT, SSL, or SASL_SSL authentication credential, in the Credential field, create a Kafka SSL credential.

   a. Select the Search icon ( ) next to the Credential field.

   b. On the Credentials screen, select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. From the list of available credentials, select **Kafka SSL Credentials**.

For the Kafka SSL credentials fields descriptions, see Kafka SSL credentials fields.

Note: If you clear the **Disable hostname verification** field, then in the **Additional Kafka consumer properties** field, enter the value of the *ssl.endpoint.identification.algorithm* parameter the same as of the Kafka server parameter value. For example:

```
ssl.enabled.protocols=TLSv1.2,TLSv1.1,TLSv1;
sasl.jaas.config=org.apache.kafka.common.security.plain.Pl
ainLoginModule required username=<username>
 password=<password>;
sasl.enabled.mechanisms=PLAIN;ssl.endpoint.identification.
algorithm=;
```

Configure SAP Solution Manager connector

Configure the SAP Solution Manager (Solman) connector instance to enable communication between the SAP Solution Manager and Event Management.

Before you begin

Role required: admin
Supported versions: Tested on SAP Solution Manager 7.2. Might also work on other versions.

About this task

This configuration enables you to update the SAP Solution Manager with the statuses that have been updated by ServiceNow. For general information about the SAP Solution Manager, see SAP configurations for the SAP Solution Manager connector.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a connector instance.
3. Fill in the fields on the form.

Connector instance form

Field	Description
Name	Descriptive name for the SAP Solman connector.
Description	Description for the use of the SAP Solman event collection instance.
Active	Option that activates the connector instance. Select this option only after running a successful test.
Connector definition	The vendor and protocol used to gather events from the external event source. Select **SAP Solman**.
Host IP	The IP address where SAP Solman is installed.
Credential	Required credentials entered either by selecting them from the results of a search through the search icon or by clicking **New**. or creating them. If you create credentials, save them using a unique and recognizable name, such as SAPSOLMANOPS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Event collection last run time	The last run time value. Automatically updated.
Last event collection status	The last run time status. Automatically updated.
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from SAP Solman.
Bi-directional	Ensure that you select this option to invoke the bi-directional exchange of values to-and-from the external event source. The **Last bi-directional** status option appears only when selecting this option.
Last bi-directional status	This field appears only when **Bi-directional** is selected. This field is automatically populated with the most recent bi-directional status.
Last error message	Last error message. Automatically updated.
Connector Instance Values section	The parameters that are specific to SAP Solman, automatically displayed when saving the form.
MID Servers (MID Server for Connectors section)	MID Server that is up and valid. You can configure several MID Servers so that if the first server is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order in which their details are entered into the MID Server for Connectors section. If no MID Server is specified, an available MID Server that has a matching IP range is used.

4. Right-click the form header and select **Save**.

   The connector instance values are added to the form and the parameters that are relevant to the connector appear.

5. In the Connector Instance Values section, specify the default connector instance values.
   - **Client**: Default value = `001`
   - **Port**: Default value = `8000`
6. Right-click the form header and select **Save**.
7. Click **Test Connector** to verify the connection between the MID Server and the connector. If the test fails, follow the instructions in the error message and then run another test.

   Note: Use a network tool such as *`ping`* to verify the credentials and the network connectivity from the MID Server to the SAP Solman monitor.

8. After a successful test, select the **Active** check box and click **Update**.

Related topics

Use the SAP Solution Manager Push connector

SAP configurations for the SAP Solution Manager connector

Configure your SAP environment to work with the ServiceNow Event Management platform so you can use the SAP Solution Manager connector.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When integrating your SAP environment with Event Management, the SAP Solution Manager sends events through a MID Server to Event Management, which generates an alert from the event.

After the MID Server processes events, you can perform the following actions in the SAP Solution Manager:

- Add comments
- Confirm alerts
- Set status to **Externally Managed**

The following diagram illustrates the flow of data between SAP and Event Management.

SAP Solution Manager connector

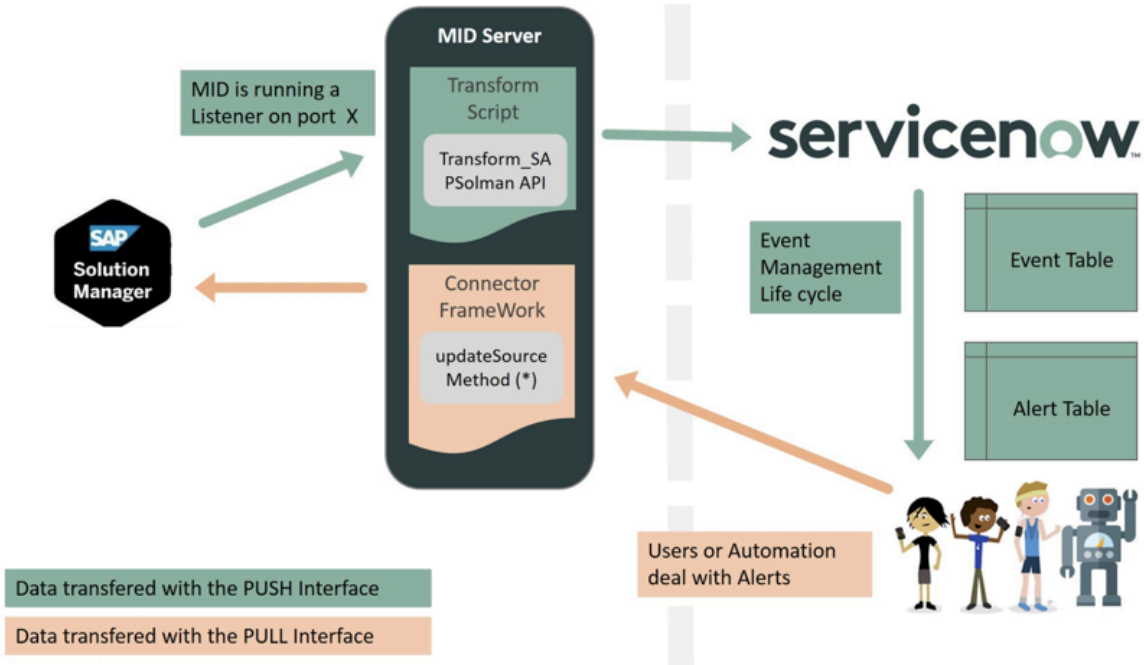

Data moves between SAP and Event Management using both Push and Pull interfaces.

- **Pull interface**:

  The Connector Update Queues table [em_connector_update_queue] lists manual updates of alerts and sends them to the SAP Solution Manager.

  The table's business rule identifies each manual update of alerts in Event Management and updates the connector queue with these changes. By default, changes to all alert fields are tracked. For information on adjusting this setting, see the SAP Solution Manager documentation.

  The Event Management - Queue connector processor scheduled job removes alert changes from the queue and sends them to the MID Server. By default, the system performs this process in batches of 1,000 alerts. You can customize this batch size by setting the `evt_mgmt.max_update_source_records` property.

  Note: You must define the Pull interface to enable bi-directional functionality in the SAP Solution Manager connector. For details, see the SAP Solution Manager documentation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Push interface**: Because SAP can only push data to an external source, the integration uses the Push interface to store alerts sent from SAP to Event Management.

  SAP alerts create events in Event Management that use both event rules and CI identification rules to create Event Management alerts.

  After Event Management creates alerts, a business rule listens for a change to the alert. If it detects a change, Event Management updates the SAP Solution Manager alerts.

Related topics

Event Management Integrations

Use the SAP Solution Manager Pull connector

Use the SAP Solution Manager Pull connector

The SAP Solution Manager Pull connector sends information from Event Management to the SAP Solution Manager. The Pull connector sends responses received from a Push connector in a bi-directional environment.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Event Connectors (Pull)** > **Connector Instances**.
2. Click **New**.
   The **Connector Instance New Record** page appears.
3. Enter values in the required fields, which are indicated by a red asterisk.

Field	Description
Name	Descriptive name for the connector.
Connector definition	Search for and select **SAP Solution Manager**.
Credential	Select **Basic auth** for basic authentication.
Host IP	The SAP Solution Manager IP address.

The Connector Update Queues table [*em_connector_update_queue*] lists manual updates of alerts and sends an event to the SAP Solution Manager, using the Pull interface.

Note: When adding bi-directional functionality for a Push connector, the Connector Instance name must match the Source Instance field of the event generated by that source. The **Connector Instance** field value is a combination of `SID + _ + Mandant`.

For example, the following image displays an event generated by the instance source **PSM_001**. To enable sending information back, you would have to name your connector instance **PSM_001**.

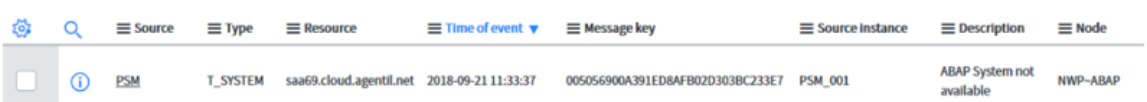

4. Set the following values for the connector instance:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Client**: Represents the client value in SAP, usually 001.
- **Port**: The SAP Solution Manager WebService port. Default value = 8000.

The following table lists the actions you can perform in Event Management, and the resulting updates occurring in SAP Solution Manager.

Event Management actions and SAP Solution Manager updates

Action in Event Management	Result in SAP Solman
Close an alert.	Confirms an alert in the alert inbox. The alert is removed.
Select the **Maintenance** setting in an alert.	Adds a comment indicating that the alert is in Maintenance mode.
Clear the **Maintenance** setting in an alert.	Removes the Maintenance mode comment.
Select the **Acknowledged** setting in an alert.	Adds a comment indicating that the alert is in the Acknowledged mode.
Clear the **Acknowledged** setting in an alert.	Removes Acknowledged mode comment.

Note the following behaviors:

- The Update queue business rule in the Alert table identifies each manual update of the alert and updates the connector queue accordingly. By default, changes to all alert fields are tracked.
- The Event Management — Queue Connector Processor job removes alert changes from the queue and sends them to the MID Server. By default, Event Management performs this dequeue process in batches of 1,000 alerts. You can configure the batch size by setting the `evt_mgmt.max_update_source_records` property.

Use the SAP Solution Manager Push connector

The MID Server web service Event Collector enables you to collect alerts sent from the SAP Solution Manager through event stream notification capabilities. The interface runs both a Push and a Pull interface to interact directly with the SAP Solution Manager.

Before you begin

- Deploy and start the MID Server.
- Configure and start the MID Web Server.
- Configure and start the MID Server web service Event Collector.

Role required: evt_mgmt_admin

About this task

SAP Solution Manager sends JSON formatted event messages to the MID Server. The MID Server parses the collected alert messages and transforms them using the SAP Solman Transform [*TransformEvents_SAPSolman*] script.

JSON formatted messages are collected using basic authentication. For details on other supported authentication methods, see Configure the MID Web Server extension.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Push connectors**.
2. Select the **SAP Solman Transform Script** script.
3. Use the following URL to push alert messages from SAP Solution Manager to the MID Server: `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=TransformEvents_SAPSolman`, where the following are true:
   - **MID_Server_IP** is the IP address of the MID Web Server extension
   - **MID_Web_Server_Port** is the listening port of the MID Web Server extension

What to do next

Create an Remote Function Call (RFC) in the SAP Solution Manager and install a transport, as described in Configure RFC in SAP Solution Manager.

Related topics

Configuring MID Servers

Configure the MID Web Server extension

Configure the MID WebService Event Collector Context

SAP Solution Manager setup configurations

As part of the SAP Solution Manager setup, you must perform several configurations to enable SAP Solution Manager to interact with Event Management.

Configure RFC in SAP Solution Manager

As part of enabling communication with Event Management, you must create a Remote Function Call (RFC) in the SAP Solution Manager and install a transport.

Before you begin

Role required: evt_mgmt_admin

Download the following .psm files and import them into the SAP Solution Manager system.

- https://downloads.docs.servicenow.com/resource/enus/cad/K900249.PSM
- https://downloads.docs.servicenow.com/resource/enus/cad/R900249.PSM

Procedure

1. On the SAP UI, enter the transaction code `SM59` and create a new RFC with the following parameters on the **Technical Settings** tab:

RFC parameters

SAP Parameter	Value
RFC Destination	`MID_EVENT_COLLECTOR2`. Enter this value exactly because the custom SAP BADI `Z_ALRT_REACTION_IMPL` calls this RFC destination by name.
Target Host	Name of the MID Server.
Service No	Listener port on the MID Server
Path Prefix	The URL of the MID Server transform script: `/api/mid/em/inbound_event?Transform=TransformEvents_SAPSolman`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Technical Settings tab

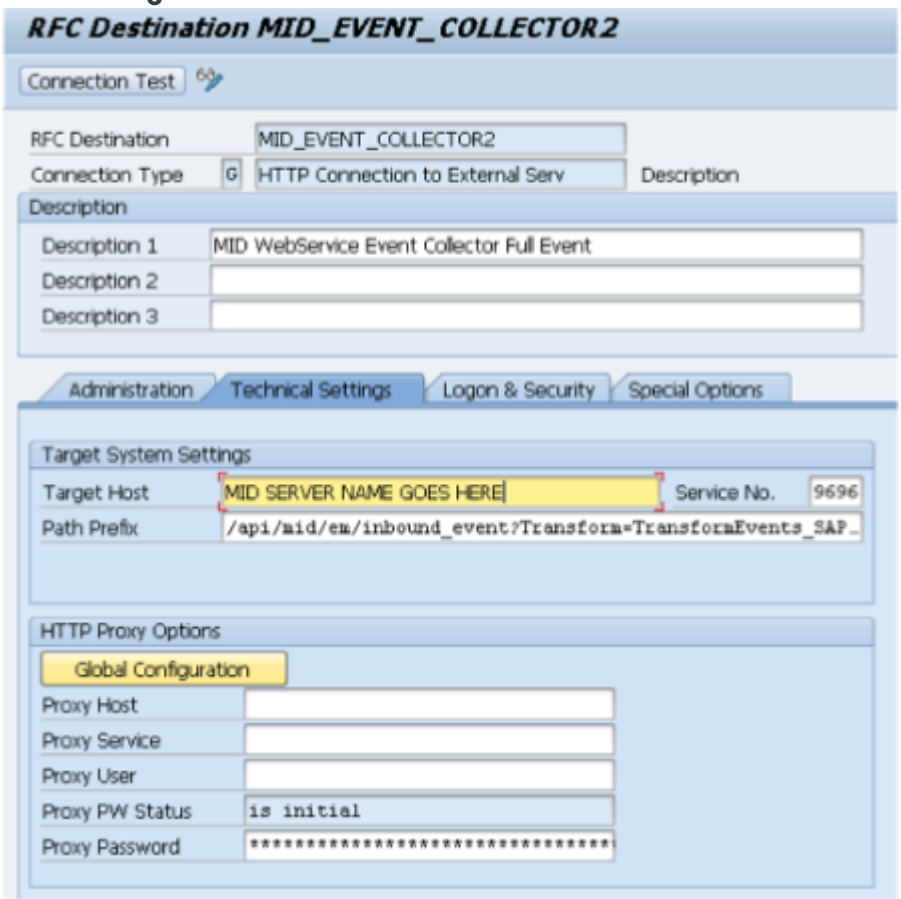

Note: The **RFC Destination** field value must be `MID_EVENT_COLLECTOR2`.

2. On the **Logon & Security** tab, select **Basic Authentication** and enter your Event Management logon credentials in the **User** and **Password** fields.

Logon & Security tab

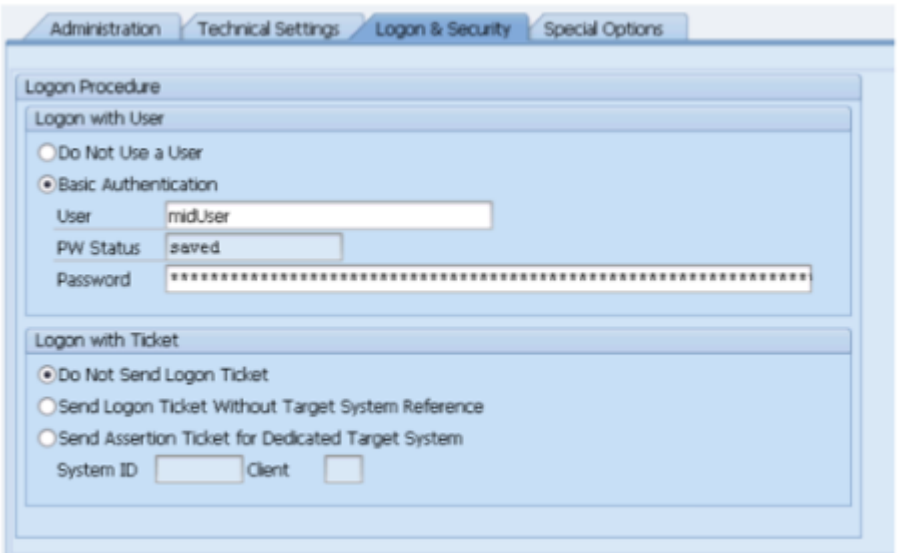

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A custom SAP BADI (Business Add-in) named Z_ALRT_REACTION_IMPL is created and stored as a workbench transport request. After installation, this BADI calls the RFC to connect to Event Management.

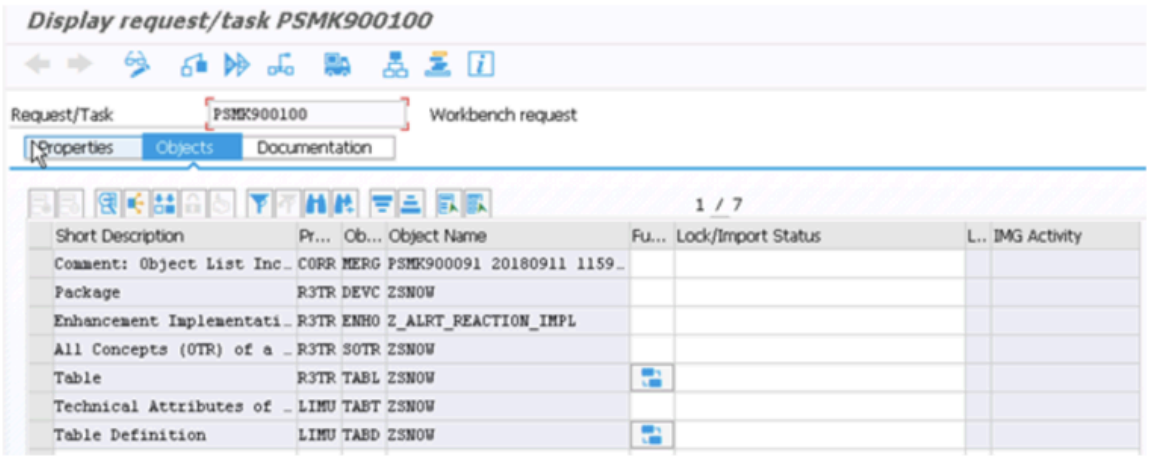

Configure SAP to forward alerts to your BADI

Configure SAP to forward alerts to the BADI (Business Add-in) you create for the integration.

Before you begin

Role required: evt_mgt_admin

About this task

Third-party alerts get forwarded globally. You can configure the alert reaction only for a specific monitoring template.

Procedure

1. On the SAP UI, navigate to **Application Operations** > **System Monitoring**.
2. In the ribbon at the top of the System Monitoring configuration page, select **2 Configure Infrastructure** > **2.3 Default Settings**.
   The **Default Settings** page appears.

Default Settings page

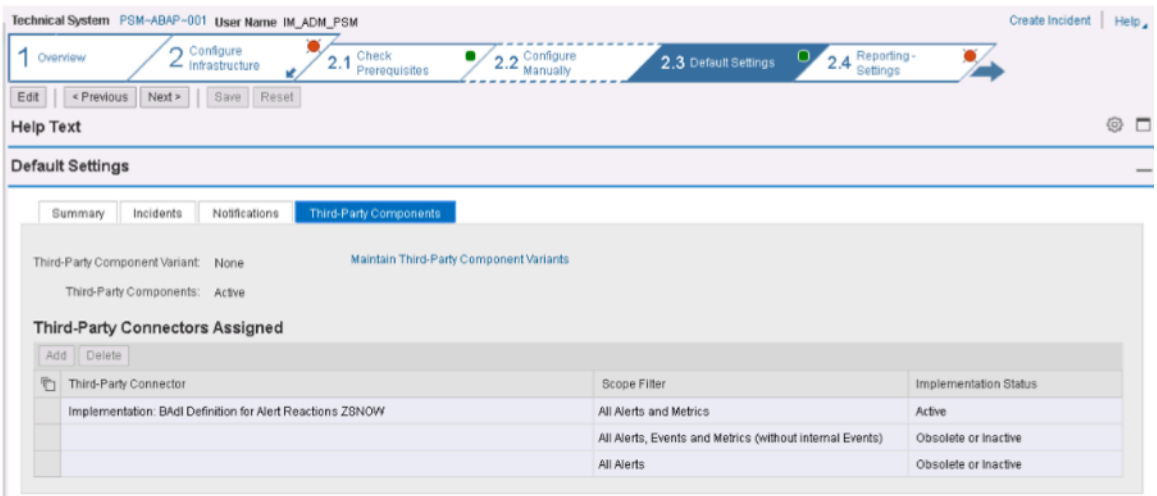

3. Configure and activate a third-party component.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Enter edit mode by clicking the **Edit** button.

b. Select the **Third-Party Components** tab.

c. Click **Add** and add the following value as the third-party connector: `Implementation: BADI Definition for Alert Reactions ZSNOW with scope = All Alerts and Metrics`.

d. Click **Save** to activate third-party components globally.

4. Define the objects to which the monitoring template will be assigned.

a. In the ribbon at the top of the page, select **5 Define Scope**.
The Define Scope page appears.

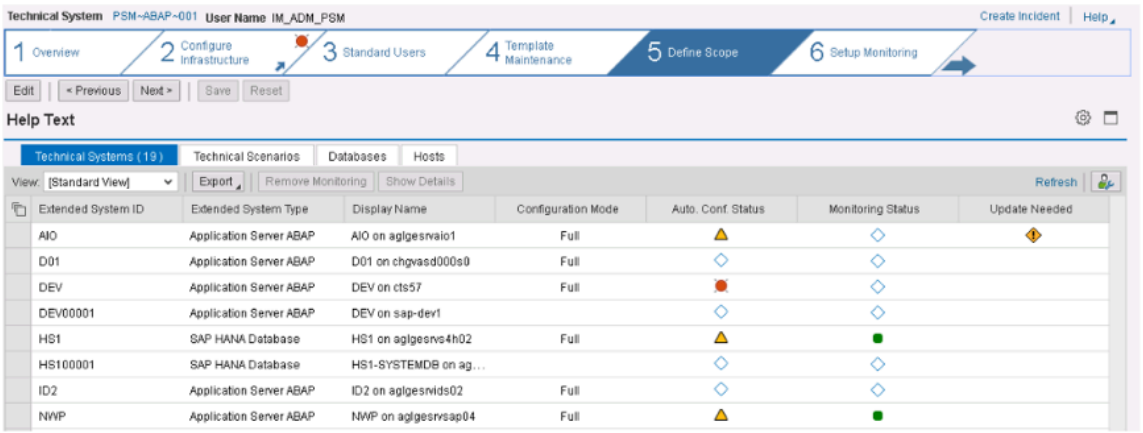

b. Select the managed object to which you want to assign the monitoring template.
The monitoring template is assigned to the selected objects.

5. Reactivate the monitoring of templates.

a. Select **6 Setup Monitoring**.

b. Search for and locate entries with **Setup Status = Reconfigure Required** and assign the monitoring template to those entries.

c. Select **Apply** and **Activate**.
When the jobs finish, verify that the **Assignment Status** is green.
All templates and active alerts assigned to the selected objects use the new forwarding method.

What to do next

View alerts generated in SAP Solution Manager to see any pressing issues, as described in View alerts in the SAP Solution Manager inbox.

Configure the alert inbox in SAP Solution Manager

Configure SAP solution manager to retrieve alerts in the system by enabling the alert inbox with the SOA manager.

Before you begin

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to SOA MANAGER in your solution manager: `/sap/bc/webdynpro/sap/appl_soap_management?sap-client=XXX&sap-language=EN#`
2. Select **Simplified Web Service Configuration**, then validate the form to activate the service. You must use simplified mode else the resulting web services is different with other modes and not compatible with the connector.

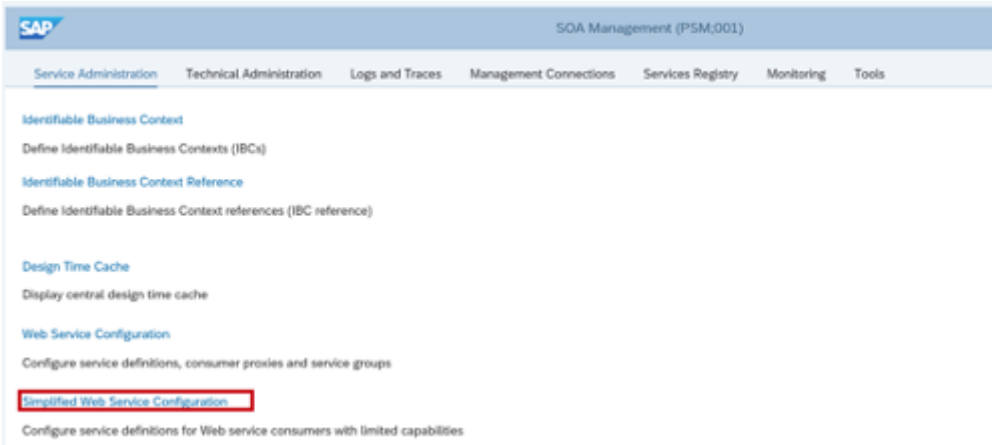

3. Select **Search**.
4. Check the row with the creation type **Simple consumer related**.
   If anything goes wrong, check Configuring a Service Provider (Simplified)

What to do next

View alerts generated in SAP Solution Manager to see any pressing issues, as described in View alerts in the SAP Solution Manager inbox.

View alerts in the SAP Solution Manager inbox

You can view alerts generated in SAP Solution Manager to see any pressing issues. All alerts also forward to Event Management automatically.

Before you begin

Role required: evt_mgt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. On the SAP UI, open Solution Manager and select the **System Alerts** tab.

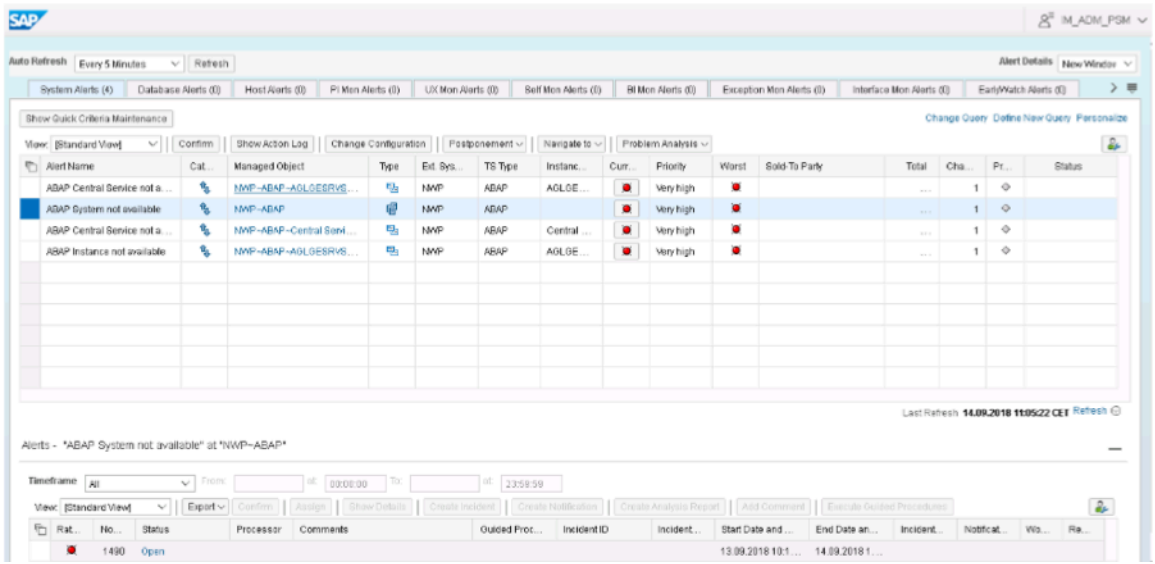

The system alerts are listed on the page.

2. Navigate to (**Event Management** > **All Alerts)** to view the alerts in Event Management.

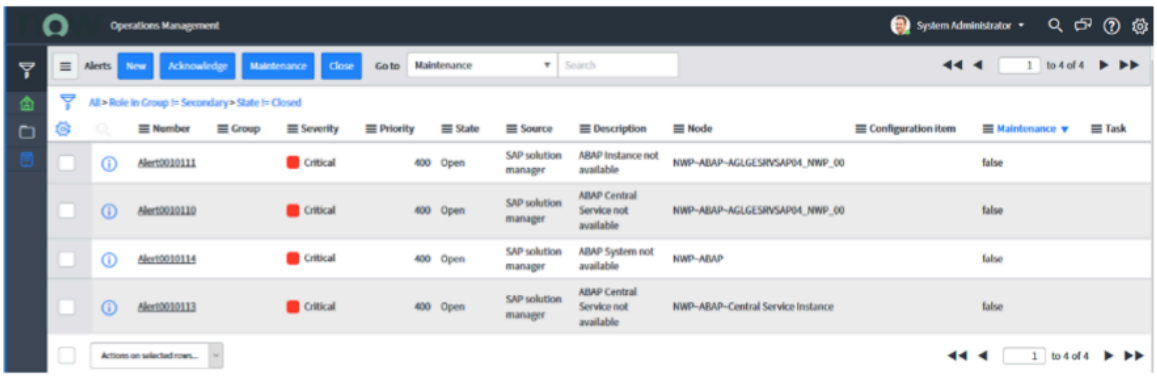

The following table describes the field mapping between SAP Solution Manager and Event Management for alerts.

Field mapping for SAP Solution Manager and Event Management

SAP Field	Event Management Field	Description
severity	severity	SAP uses a number from 1 (less critical) to 9 (more critical)
message_key	message_key	SAP message group ID, messageGUID
description	description	Event description
source	source	SAP system ID
managed_object_type	type	The type of the object, such as Host or Instance
node	node	Object name

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SAP Field	Event Management Field	Description
managed_object_name	resource	Solution manager node name
N/A (there is no value provided by SAP. Event Management generates the indicated field.)	source instance	The event class must match the connector pull instance name, with the structure <Source>_<Client>.
source	additional_info.sid	SAP system ID
rating	additional_info.rating	Event color in SAP
resource	additional_info.mandant	SAP mandant
metric_name	metric_name	Metric name for the SAP event
time_of_event	time_of_event	Time of event in SAP

Example: High CPU Usage on SAP Host

Scenario: An alert is generated in SAP Solution Manager when the CPU usage on a critical host exceeds 90% for more than 10 minutes.

SAP Solution Manager Alert

Field	Value
Severity	8 (Critical)
Message Key	CPU_Overload_12345
Description	High CPU usage on SAPHost01
Source	SAP_Prod
Type	Host
Node	SAPNode_01
Time	2024-07-31T14:23:00Z

Here's how the alert is mapped in Event Management:

Mapped Event Management alert

Field	Value
Severity	8
Message Key	CPU_Overload_12345
Source	SAP_Prod
Type	Host
Node/Resource	SAPNode_01
Source Instance	SAP_Prod_100
SID	SAP_Prod
Rating	Red

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Mapped Event Management alert (continued)

Field	Value
Mandant	100

View and update your SAP Solution Manager certificate

View your SAP Solution Manager certificate, and update the certificate if necessary.

Before you begin
Role required: evt_mgmt_admin

Procedure

1. Connect to SAP on client 00.
2. Enter the `strust` transaction.
3. In the Trust Manager UI, open the SSL client.
4. Import the SSL certificate and save it.
5. Use the `smicm` transaction to enable the system to work with both HTTP and HTTPS clients.
6. Navigate to **Administration** > **ICM** > **Exit Hard** > **Global**.

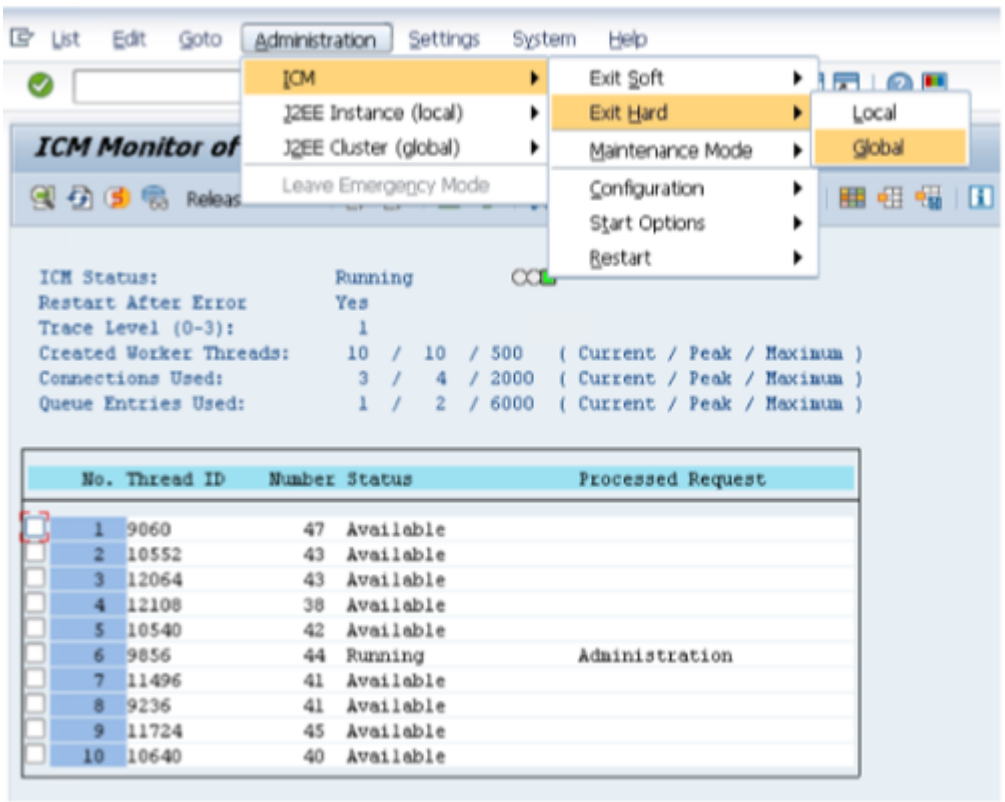

The ICM module restarts.

View the SAP interface log

View the SAP interface log to help you debug issues.

Before you begin
Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

Select a transaction in the SAP UI.
The ZSNOW interface debugging table appears.

SAP Solution Manager transaction codes

The transaction code abbreviations that you can use in the SAP interface when working with the SAP Solution Manager connector.

SAP Solution Manager transaction codes

Transaction Code	Description
ALRTINBOX	Displays the alert inbox.
SOAMANGER	Opens and activates the *Sap_Solman_Alert_Inbox* web service.
SRT_UTIL	View the SOA manager logs.
SOLMAN_WORKCENTER	Opens the SAP Solution Manager interface in a browser.
SM59	Displays the entire output interface.
SE38	Executes the Advanced Business Application Programming (ABAP) code. (The equivalent ServiceNow code is *ZTEST_SNOW_CONNECTOR*.)
MAI_TOOLS	Displays the alarm tool.
SE16	Displays the interface debugging table (ZSNOW).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure alert collection from SCOM

Alerts from the Microsoft System Center Operations Manager (SCOM) are collected using the SCOM connector instance.

Before you begin

Role required: evt_mgmt_admin

About this task

The SCOM connector instance requires configuration before receiving alerts and Metric Intelligence raw data from the SCOM server. You can limit the collection of SCOM alerts to only those alerts that belong to the SCOM group that you specify.

Procedure

1. Configure the SCOM connector instance.
2. Configure the Microsoft System Center Operations Manager (SCOM) connector instance to receive alerts and Metric Intelligence raw data from the SCOM server.
3. Limit collected SCOM alerts to specific SCOM groups.

Configure the SCOM connector instance

Configure the Microsoft System Center Operations Manager (SCOM) connector to receive alerts and Metric Intelligence raw data from the SCOM server. SCOM event collection and metric collection are handled by two separate connector definitions: SCOM (for alerts and bi-directional exchange) and SCOM Metrics (for Metric Intelligence raw data).

Before you begin

Common prerequisites (both connectors):

- The MID Server is running on Windows.
- The MID Server resides in the same domain as the SCOM server.
- The MID Server uses the same time zone as the SCOM server.

Additional prerequisites for SCOM event collection:

- The MID Server is running with a user that has local admin permissions to enable the MID Server to run PowerShell.
- The MID Server user has read access to the SCOM API.
- The MID Server has .NET framework version 3.5 (for SCOM 2007/2012) or .NET 4.6 or higher (for SCOM 2016/2019/2022/2025).
- If Bi-directional is selected, ensure that PowerShell version 3.0 is installed on Windows and the MID Server is running with a user that has local admin permissions.

Additional prerequisites for SCOM metric collection:

The MID Server that retrieves metrics is configured with the Metric Intelligence extension and the extension is in Started mode. See Manually configure the Metric Intelligence extension.

If "Database login with Windows authentication" is selected, the MID Server service must run with a user having read access to the SCOM database (OperationsManagerDW).

To configure this:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. In the local services, right-click the MID Server service and select Properties.
2. In the Log On tab, verify that "This account" is selected with the details of the Windows domain user that has read access to the SCOM database.

If "Database login with Windows authentication" is NOT selected, you need a Windows credential in the credential store that has read access to the OperationsManagerDW database.

Upgrading from a previous release: If you are upgrading from a release where a single SCOM connector handled both events and metrics, the upgrade automatically disables metric collection on the old SCOM connector. The upgrade also deprecates the old metric-related parameters and event rules. Your existing SCOM connector continues to handle event collection as before — no action needed for events. To resume metric collection after the upgrade, create a new SCOM Metrics connector instance.

Role required: evt_mgmt_admin

About this task

There are two connector definitions for SCOM. The SCOM (Events) connector collects alerts using PowerShell and native .NET SCOM client executables. It requires DLL files from the SCOM server to be uploaded to the MID Server. It also supports bi-directional alert exchange. The SCOM Metrics connector collects Metric Intelligence raw data using JavaScript running JDBC directly against the OperationsManagerDW database. It does not require any DLL files and does not collect events or support bi-directional exchange.

You can also use this setup for the SCOM managed instance on Azure.

This connector has the *logPayloadForDebug* log parameter enabled, which logs event and metric payloads from the source system. Once debugging is complete, set this parameter to **false** to prevent overloading the system.

Supported versions:

- 2007 – version 6.1.7221.0
- 2012 – version 7.1.10226.0
- 2016 – version 7.2.117190 and 7.3.13261.0
- 2019 – version 10.19.10050.0
- 2022 – version 10.22.10118.0

Procedure

1. Configure the SCOM (Events) connector instance.

   Follow these steps to collect SCOM alerts and optionally enable bi-directional alert exchange. This connector uses PowerShell and native .NET SCOM client executables.

   a. On the SCOM server, download the following files to a local computer.

Version	SCOM path and library names
SCOM 2012R2 or SCOM 2012	`%ProgramFiles%\Microsoft System Center 2012 R2 or 2012\Operations Manager\Server\SDK Binaries`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Version	SCOM path and library names
	▪ Microsoft.EnterpriseManagement.Core.dll ▪ Microsoft.EnterpriseManagement.Operation ▪ Microsoft.EnterpriseManagement.Runtime.d
SCOM 2007	%ProgramFiles%\System Center Operations Manager 2007\SDK Binaries ▪ Microsoft.EnterpriseManagement.Operation ▪ Microsoft.EnterpriseManagement.Operation **Note:** Don't append 2007 to the Microsoft.EnterpriseManagement.Operatic file.
SCOM 2016	%ProgramFiles%\Microsoft System Center 2016\Operations Manager\Server\SDK Binaries ▪ Microsoft.EnterpriseManagement.Core.dll ▪ Microsoft.EnterpriseManagement.Operation ▪ Microsoft.EnterpriseManagement.Runtime.d **Note:** The MID Server must be installed with .NET 4.6 or higher.
SCOM 2019	%ProgramFiles%\Microsoft System Center 2019\Operations Manager\Server\SDK Binaries ▪ Microsoft.EnterpriseManagement.Core.dll ▪ Microsoft.EnterpriseManagement.Operation ▪ Microsoft.EnterpriseManagement.Runtime.d **Note:** The MID Server must be installed with .NET 4.6 or higher.
SCOM 2022	▪ Microsoft.EnterpriseManagement.Core.dll ▪ Microsoft.EnterpriseManagement.Operation ▪ Microsoft.EnterpriseManagement.Runtime.d **Note:** The MID Server must be installed with .NET 4.6 or higher.
SCOM 2025	▪ Microsoft.EnterpriseManagement.Core.dll ▪ Microsoft.EnterpriseManagement.Operation ▪ Microsoft.EnterpriseManagement.Operation ▪ Microsoft.EnterpriseManagement.Runtime.d

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Version	SCOM path and library names
	Note: The MID Server must be installed with .NET 4.6 or higher.

Note: A connection cannot be made to SCOM 2012/SCOM 2016/SCOM 2019/SCOM 2022 and SCOM 2007 from the same MID Server. To work with both SCOM 2012/SCOM 2016/SCOM 2019/SCOM 2022 (any of them) and SCOM 2007 in your instance, do the following before uploading files to your instance:

i. Append .2012 to the end of the `Microsoft.EnterpriseManagement.OperationsManager.dll` file, so it becomes: `Microsoft.EnterpriseManagement.OperationsManager.dll.2012`. Do this for 2012/2016/2019/2022.

ii. Append .2007 to the end of the `Microsoft.EnterpriseManagement.OperationsManager.dll` file found in the 2007 path.

Using these modified file names enables the relevant MID Server to load the applicable DLL when both SCOM versions are deployed. Do not append 2007 to the `Microsoft.EnterpriseManagement.OperationsManager.Common.dll` file (for SCOM 2007).

At any given time, you can define connectors of either SCOM 2007, SCOM 2012/2016, or SCOM2019. Since SCOM 2012 and SCOM 2016 use the same library files, they are able to work at the same time. Although, SCOM 2019 uses the different SCOM DLL files so you cannot configure SCOM 2012, 2016, and 2019 all together. Each SCOM connector can use the database of another connector, for metric support.

b. Navigate to **MID Server** > **JAR Files**.

c. Select **New** and add a separate record for the SCOM version for each of the DLL files that you downloaded from the SCOM server.

i. In the **Name** field, specify the SCOM version and an identifier to make the name unique, for example `2012 1`.

If you are using SCOM 2016, specify 2012 as the version.

ii. Select the paper clip icon in the form header and then attach one of the appropriate DLL files that you downloaded.

iii. Select **Submit**.

d. Repeat step 3, creating a separate record for each of the remaining DLL files.

Ensure that you have a unique identifier after the SCOM version for each file that you attach, for example `2012 2`.

e. Navigate to **Event Management** > **Integrations** > **Connector Instances**.

f. Click **New** and create a new connector instance.

For details on the connector instance fields displayed on the page, see SCOM connector instance form.

g. Select and hold (or right-click) the form header and select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The connector instance values are added to the form and the parameters that are relevant to the connector appear.

h. In the Connector Instance Values section, you can edit the values of the mandatory SCOM parameters.

i. login_with_windows_authentication Default value: **false**.

Set to **true** to enable SCOM event collection and the bi-directional exchange of event values to work with Windows authentication. When invoking this value, ensure that you do the following on the MID Server:

1. Navigate to the list of local services, select and hold (or right-click) the MID Server service, and select **Properties**.
2. In the **Log On** tab, ensure that **This account** is selected with the details of the user in the Windows domain with read access to the SCOM database.

ii. *scom_date_format*: Default format: M/d/yyyy/ h:mm:ss a

If you receive an event whose date is in a different format, modify this value to match the format of the incoming event. If you do not, the event does not process correctly.

For example, an event arrives on June 27, 2019 at 11:25 AM with a listed date of *2019/06/27/ 11:25:00 a*. Modify the *scom_date_format* value to *yyyy/M/d/ h:mm:ss a* to match the format of the received event.

In *scom_date_format*, a represents AM, and p represents PM.

iii. *scom_initial_sync_in_days*: Default value: 7.

iv. *scom_version*: It is mandatory to specify the SCOM version.

Select from 2025, 2022, 2019, 2016, 2012, or 2007.

v. *stop_isMonitorAlert_bidirectional*: Boolean.

Set to **true** to stop only SCOM monitoring alerts bi-directionally. Stopping SCOM monitoring alerts can prevent false alarms on the SCOM console when attempting to update the alerts. When set to **false**, SCOM monitoring alerts continue running bi-directionally for all types of SCOM alerts. Default value = false.

vi. *scom_alert_bulk_size*: Default value: 500. Number of alerts fetched per SCOM collection cycle.

vii. *debug_enabled*: Default value: false. Set to true to enable debug logging on the MID Server.

i. Select and hold (or right-click) the form header and select **Save**.

j. Restart the MID Server service to copy the files.

k. Select **Test connector** to verify the connection between the MID Server and the connector.

If the test fails, correct the problem by following the instructions issued by the error and then run another test.

Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to SCOM.

l. After a successful test, select the **Active** check box and then select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

The default binding rules that contain SCOM as the external source, that applies to IT alerts and Metric Intelligence raw data, are the following SCOM Management Packs:

- All OS Management Packs
- MS SQL Server
- IIS

m. After a successful test, select the **Active** check box and then select **Update**.

If bi-directional is configured, the bi-directional exchange of values to-and-from the external event source is enabled. These scenarios describe the default bi-directional functionality for SCOM connectors:

- When an alert is resolved in SCOM, it is auto-closed in ServiceNow. However, it is updated irrespective of the bi-directional feature because during each collection cycle, all alert changes are updated.
- When an alert is manually closed in ServiceNow, it is auto-closed in SCOM. If the alert state is changed to **Reopen**, SCOM is also updated.
- When an incident is created and associated to an alert in ServiceNow, SCOM receives the incident number as a ticket ID. However, the state of the incident is not available on SCOM. Therefore when the incident is resolved in ServiceNow, SCOM is not updated as the incident number remains the same. When the alert is associated with a new incident, the new incident number is updated in SCOM.

2. Configure the SCOM Metrics connector instance.

Follow these steps to collect Metric Intelligence raw data from SCOM. This connector connects directly to the SCOM OperationsManagerDW database using JDBC. It does not require SCOM DLL files on the MID Server. Note: This connector collects metrics only. It does not collect events or support bi-directional exchange. To collect SCOM alerts, configure the SCOM (Events) connector.

a. Navigate to **Event Management** > **Integrations** > **Connector Instances**.

b. Click **New** and create a new connector instance.

c. In the **Connector definition** field, select **SCOM Metrics**.

d. Fill in the following fields:

Metrics database host	The IP address or hostname of the server hosting the OperationsManagerDW database.
Connect using a named instance	Select this if the SQL Server uses a named instance instead of a port number.
Metrics database port	The port used by the database. Default is 1433. This field is hidden if "Connect using a named instance" is selected.
Metrics database named instance	The SQL Server named instance name. This field is visible only when "Connect using a named instance" is selected.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Database login with Windows authentication	Select this to use the MID Server service account's Windows credentials for database access. When selected, the database credential field is hidden.
Database credential	Visible only when Windows authentication is NOT selected. Select a Windows credential from the credential store that has read access to the OperationsManagerDW database. This field is mandatory when Windows authentication is disabled.
Metrics collection schedule (seconds)	Frequency in seconds to repeat metric collection. Default is 10.

e. Right-click the form header and select **Save**.

f. In the Connector Instance Values section, configure the following parameters:

Field	Description
scom_db_name	The name of the SCOM data warehouse database. Default is OperationsManagerDW. Change this only if your database has been renamed.
metric_chunk_size	Default value is 50000. Maximum number of metric records fetched per batch.
debug_enabled	Default is false. Set to true to enable debug logging on the MID Server.
log_payload	Default is false. Set to true to log raw metric payloads for debugging. Set back to false after debugging is complete to avoid overloading the system.

g. Right-click the form header and select **Save**.

h. Select **Test connector** to verify the database connection.

The test runs a simple query against the OperationsManagerDW database to confirm connectivity.

i. After a successful test, select the **Active** check box and select **Update**.

Limit collected SCOM alerts to specific SCOM groups

Limit the collection of SCOM alerts to only those alerts that belong to the specified SCOM group.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Use the SCOM Authoring workspace and Operations console to define a SCOM group and ensure that the group contains the required computers or instances.
   For further information about SCOM groups, see the SCOM documentation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Create a role that has a scope of the SCOM group.
3. Add the new role to the SCOM user.
4. Assign the role to the user defined in the credentials given to the SCOM connector instance.
5. Remove all other roles from the user.

Result
Only alerts from the specified SCOM group arrive at the instance.

SCOM metric event rules

The base system comes with Microsoft System Center Operations Manager (SCOM) metric event rules. SCOM metric event rules bind event metrics to configuration items (CIs).

When an event's metrics match the event filter conditions, the metrics bind to the indicated CI. The event rules that come with the base system are:

Base system metric event rules

Rule Name	CI type
SCOM Metrics - Windows Server	cmdb_ci_win_server
SCOM Metrics - IIS Web Server	cmdb_ci_microsoft_iis_web_server
SCOM Metrics - SQL Server Reporting Services	cmdb_ci_db_mssql_reporting
SCOM Metrics - MS SQL Database	cmdb_ci_db_mssql_database
SCOM Metrics - MSFT SQL Instance	cmdb_ci_db_mssql_instance

You can create a custom event rule to bind metrics with other CI types. For details, see Create or edit an event rule.

If the event rule filter doesn't match, the default rule, *`SCOM Metrics - Default`* (CI type = *`cmdb_ci_hardware`*) is invoked. The order number of the default event rule must be higher than the order number of any SCOM event rules. This includes both base system rules and custom rules. This ensures that all other rules are checked for matching before invoking the default rule. For details on event rules, see Event rules.

Configure event collection from SolarWinds monitor

Configure the SolarWinds monitor connector instance to receive events from the SolarWinds monitor.

Before you begin
Supported versions:

- SAM 6.2.1
- NPM 2024.1

Note: The SolarWinds monitor connector instance can import events from:

- Network Performance Monitor (NPM)
- SolarWinds Service & Application Monitor (SAM)

The SolarWinds monitor connector instance requires a credential that lets the instance access SolarWinds monitor accounts. You can use an existing credential or create a new one. The

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

provided user must have access to the SolarWinds API to receive events from the SolarWinds connector.

Role required: evt_mgmt_admin and evt_mgmt_integration

About this task

This connector has the *debug* and *logPayloadForDebug* log parameters enabled. The *debug* parameter logs debug messages, such as calls made to the source system to retrieve events. The *logPayloadForDebug* parameter logs event and metric payloads from the source system. Once debugging is complete, set this parameter to **false** to prevent overloading the system. To use the legacy XML format, set the *useXMLResponse* parameter to **true**.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New**.
3. On the form, fill in the fields.

Field	Description
Name	Type a descriptive name for the SolarWinds monitor connector.
Description	Type a description for the use of the SolarWinds monitor event collection instance.
Connector Definition	Select **Solarwinds_v2**.
Host IP	Specify the SolarWinds monitor IP address.
Credential	Select the credential with basic authentication that you created for this connector. For more information, see Create SolarWinds monitor credentials.
Event collection last run time	The last run time value of the scheduled job. This field is updated automatically.
Last event collection status	The last event collection status. This field is updated automatically.
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from the SolarWinds monitor server.
Last error message	The last error message received by the connector is automatically updated.

4. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.
5. In the Connector Instance Values section, specify these fields:

Field	Value
initial_sync_in_days	Specify how many days the initial pull must contain to retrieve events. The default value is 7 days.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
nodes_custom_properties	Specify the SolarWind monitor property names that are associated with the relevant node. This information is added to the **Additional Information** field of the event in JSON format. The SolarWinds monitor custom properties are user-defined fields, for example, country, building, or serial number, that you can associate with monitored network objects. Separate multiple entries with a comma.
port	Specify the value of the port. By default, the value is 17774. **Note:** For version 2023.1 and above, the port is 17774.
queryAlertTable	A boolean toggle that determines the source table for event queries. Setting the value to **false** queries the Orion.Event table, while setting it to **true** queries the Orion.AlertHistory table. Default: false

6. **Optional:** Name of a MID Server.
   If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section. The privileges that are required for the user account that is used for the MID Server to collect events from SolarWinds monitor is that of evt_mgmt_integration.
7. Click **Test connector** to verify the connection between the MID Server and the SolarWinds monitor connector.
   If the test fails, follow the instructions that are issued by the error to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to SolarWinds monitor.

8. After a successful test, select the **Active** check box and then click **Update**.

Related topics

Configure a pull connector

Create SolarWinds monitor credentials

Create a Basic Auth credential in ServiceNow to store the SolarWinds user name and password that the SolarWinds monitor connector uses to access the SolarWinds API.

Before you begin

Confirm you have a local SolarWinds Orion account for Basic authentication (username + password) with rights to query the SolarWinds Information Service (SWIS) REST API. Create this account in the Orion Web Console; it is the account ServiceNow uses to poll alerts.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Click **New**.
3. Select **Basic Auth Credential**.
4. On the form, fill in the fields.

Basic Auth Credentials form

Field	Description
Name	Enter a unique and descriptive name for this credential. For example, you might call it `SolarWinds authentication`.
User name	Enter the user name of the SolarWinds account that the connector uses to access the SolarWinds API.
Password	Enter the password for the SolarWinds account.
Active	Option to enable the use of this credential.
Order	The order (sequence) in which the platform tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order. Default value: `100`

5. Click **Submit**.

Result
The credential for use with the SolarWinds monitor connector is created.

Related topics

Credentials and connection information

Configure event collection from vCenter

Configure the VMware vCenter Server (vCenter or vCenter_V2) connector instance to receive events from your VMware vSphere environment.

Before you begin
Events collected by this connector are events that match the vCenter ESX Servers event rule. Supported version: 7.0.2.00500.

Role required: evt_mgmt_admin

About this task
Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

This connector has the *debug* and *logPayloadForDebug* log parameters enabled. The *debug* parameter logs debug messages, such as calls made to the source system to retrieve events. The *logPayloadForDebug* parameter logs event and metric payloads from the source system. Once debugging is complete, set this parameter to **false** to prevent overloading the system. To use the legacy XML format, set the *useXMLResponse* parameter to **true**.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New**.
3. On the form, fill in the fields.

Connector Instance form

Field	Description
Name	Descriptive and unique name for the vCenter connector.
Description	Description for the use of the vCenter event collection instance.
Connector definition	Name of the required connector definition. Select **vCenter** or **vCenter_V2** (vCenter_V2 is located in the Event Management connectors scope).
Host IP	IP address where vCenter is installed.
Credential	Permission to connect to vCenter. Click the search icon in the **Credential** field. Either select the required credentials from the list or click **New** and create the required credentials on the Credentials form. If you create the credentials, save them using a unique and recognizable name, for example, vCenterCHCK.
Event collection last run time	Last run time value. Automatically updated.
Last event collection status	Last run status. Automatically updated.
Event collection schedule (seconds)	Frequency, in seconds, that the system checks for new events from vCenter. The default value is 120 seconds.
Last error message	Last error message. Automatically updated.

4. Right-click the form header and select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The connector instance values are added to the form and the parameters that are relevant to the connector appear.

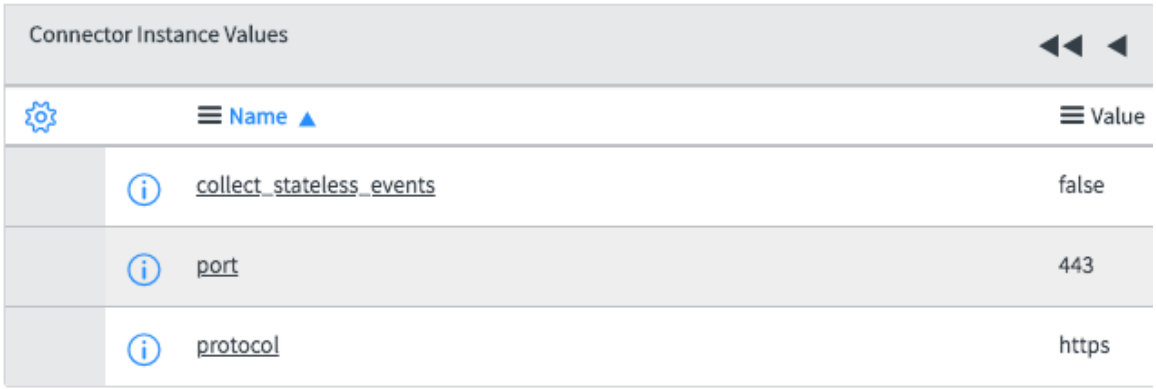

5. In the Connector Instance Values section, specify the vCenter values.

Field	Value
collect_stateless_events	Events that have a status (stateful) are collected when this field is set to `false` (default). These events are of type AlarmStatusChangedEvent and AlarmClearedEvent. If this field is set to `true`, events of other types are collected as well.
port	Specify the value of the port. Default value: `443`.
protocol	Specify the protocol. Default value: `https`.

6. **Optional:** In the MID Servers for Connectors section, specify a MID Server that is up and is valid.
   If no MID Server is specified, an available MID Server that has a matching IP range is used.

   You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

7. Right-click the form header and select **Save**.
8. Test the connection between the MID Server and the vCenter connector.

   a. Click **Test Connector**.
   If the test fails, follow the instructions that the error issues to correct the problem and then run another test.

   > **Note:** Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the external monitoring tool.

   b. After a successful test, select the **Active** check box.

9. Click **Update**.

Configure event collection from vRealize

Configure the VMware vRealize Operations (vRealize or vRealize_V2) connector instance to receive events from the vRealize Operations Log and Event Management servers. vRealize uses basic authentication. vRealize_V2 uses token-based authentication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: evt_mgmt_admin

Supported version: 8.10.0.

Note: Connect to vRealize using a local account with at least Read-only vRealize Operations permission. You can use an existing credential or create a new one.

About this task

This connector has the *debug* and *logPayloadForDebug* log parameters enabled. The *debug* parameter logs debug messages, such as calls made to the source system to retrieve events. The *logPayloadForDebug* parameter logs event and metric payloads from the source system. Once debugging is complete, set this parameter to **false** to prevent overloading the system. To use the legacy XML format, set the *useXMLResponse* parameter to **true**.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New** and create a vRealize Operations connector instance with the following details:

Field	Value
Name	Specify a unique name for the vRealize Operations connector instance.
Host IP	Specify the vRealize Operations IP address.
Credential	Select the credential with basic authentication that you created for this connector. vRealize_V2 uses token-based authentication providing that the user can generate the token. For more information, see Create vRealize credentials.
Schedule (seconds)	The frequency in seconds that the system checks for new events from vRealize Operations.
Description	Type a description for the use of the vRealize Operations connector.
Connector definition	The vendor and protocol used to gather events from the external event source. Select **vRealize** or **vRealize_V2** (vRealize_V2 is located in the Event Management connectors scope).
Last error message	The last error message is automatically updated.
Last run time	The last run time value is automatically updated.
Last run status	The last run time status is automatically updated.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

3. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the Connector Parameters section, specify the values of the mandatory vRealize Operations parameters.

   a. **port**: Default value = `443`.

   b. **protocol**: Default value = `https`.

   c. **days_from**:
   Number of days worth of data to be pulled on the first event collection run. Default value = `14`. The **days_from** option is available only when using vRealize_V2.

   d. **useXMLResponse**:

   Optional. Controls the content type the connector uses to communicate with the vRealize Operations API. When set to false (default), the connector sends requests and parses responses as JSON. Set to true to use the legacy XML format. Default value = false. The useXMLResponse option is available only when using vRealize_V2.

5. Click **Test connector** to verify the connection between the MID Server and the vRealize Operations connector.
6. If the test fails, follow the instructions that are issued by the error to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to vRealize.

7. After a successful test, select **Active** and then click **Update**.

Related topics

Configure a pull connector

Create vRealize credentials

Create credentials to access vRealize.

Before you begin
Role required: evt_mgmt_admin

About this task
Create Basic Auth type of credential to connect to vRealize.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Click **New**.
3. Select **Basic Auth Credential**.
4. On the form, fill in the fields.

Basic Auth Credentials form

Field	Description
Name	Enter a unique and descriptive name for this credential. For example, `NewvRealize`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
User name	Name to identify the user. ○ If the username is internal in vRealize, use it as is. ○ If vRealize is integrated with Active Directory (AD), use the format `DOMAIN\username@sourceauth`. For example, `local\administrator@All vCenter Servers`.
Password	Password to use this credential.
Active	Option to enable the use of this credential.
Order	The order (sequence) in which the platform tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order. Default value: `100`

5. Click **Submit**.

Result
The credential for use with the vRealize connector is created.

Related topics
Credentials and connection information

Configure event collection from Zabbix server
Configure the Zabbix server connector instance to receiving alerts from the Zabbix server.

Before you begin
Supported versions: from 3.0.0 up to 7.0.0
The Zabbix server connector instance requires a credential that lets the instance access Zabbix server accounts. You can use an existing credential or create a new one.

Role required: evt_mgmt_admin

About this task
This connector has the *`logPayloadForDebug`* log parameter enabled, which logs event and metric payloads from the source system. Once debugging is complete, set this parameter to **false** to prevent overloading the system.

Procedure
1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Click **New**.
3. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	Descriptive name for the connector.
Host IP	IP address for the Zabbix server.
Credential	Select the credential with basic authentication that you created for this connector. For more information, see Create Zabbix server credentials.
Schedule (seconds)	Frequency, in seconds, that the system checks for new events from the Zabbix server.
Description	Description for the use of the Zabbix server event collection instance.
Connector Definition	Name of the connector definition. Select **Zabbix**.
Last error message	Last error message. Automatically updated.
Last run time	Last run time value. Automatically updated.
Last run status	Last run time status. Automatically updated.
MID Servers (MID Server for Connectors section)	Optional. Name of a MID Server. If no MID Server is specified, an available MID Server that has a matching IP range is used. In the MID Servers for Connectors section, specify a MID Server that is up and valid. You can configure several MID Servers. If the first is down, the next MID Server is used. If that MID Server is not available, the next is selected, and so on. MID Servers are sorted according to the order that their details were entered into the MID Server for Connectors section.

4. Right-click the form header and select **Save**.
   The connector instance values are added to the form and the parameters that are relevant to the connector appear.

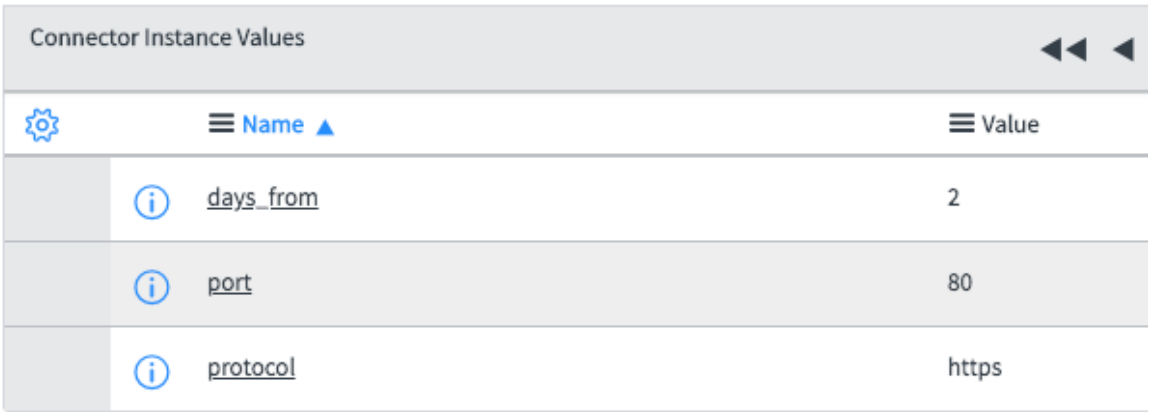

5. Verify and modify the values, as needed.
   In the Connector Instance Values section, the default connector instance values appear.

Fields in the Connector Instance Values section

Field	Description
days_from	Number of days for which events must be collected at the first collection cycle. Default value: 2

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: The maximum number of past events that can be collected is 3,000. The first time that this connector runs, the number of past events is counted in the past from the current date.
port	Port number. Default: `80`
protocol	Protocol type. Default: `https`

6. Right-click the form header and select **Save**.
7. Click **Test Connector** to verify the connection between the MID Server and the connector.
8. If the test fails, follow the instructions that are issued by the error to correct the problem and then run another test.

   Note: Use a network tool, such as ping, to verify credential correctness and network connectivity from the MID Server to the Zabbix server.

9. After a successful test, select **Active** and then click **Update**.

Create Zabbix server credentials

Create credentials to access Zabbix server.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Click **New**.
3. Select **Basic Auth Credential**.
4. On the form, fill in the fields.

Basic Auth Credentials form

Field	Description
Name	Enter a unique and descriptive name for this credential. For example, you might call it `Zabbix authentication`.
User name	Name to identify the user.
Password	Password to use this credential.
Active	Option to enable the use of this credential.
Order	The order (sequence) in which the platform tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	order number (or none), the instance tries the credentials in a random order. Default value: `100`

5. Click **Submit**.

Result
The Basic Auth credential for use with the Zabbix server connector is created.

Related topics

Credentials and connection information

Configure a Datadog work item connector instance

Configure a Datadog work item to aggregate related alerts and signals from Datadog into a trackable operational problem.

Before you begin

- Verify you already have a Datadog account with API access and Case Management in use.
- Role required: evt_mgmt_operator

About this task
ServiceNow represents a work item as an alert. Each related alert is retained individually, enabling grouping and correlation across multiple resources and monitoring systems. Alerts aren't limited to those Datadog grouped. As a result, ServiceNow identifies broader incidents, reduces duplicate alerts, and provides a more comprehensive operational view.

Note the following before you start:

- If you already use the Datadog push connector, use the same push connector URL to send work items. Configure the Datadog Case Management notification rule and the case payload template described on the push connector page.
- The push connector receives the work item and creates the work item alert and its tracking records. The pull connector then fetches the correct work item title, retrieves related alerts, and drives bi-directional close-back.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Push Connectors**.

   Note: Create this push connector to auto-create the pull connector with the matching name. The connector names must match for bi-directional and fetch related alerts.

2. Select a **Datadog Push Connector** instance used to send work items.
3. Select the **Configure Work Item Alerts** button.

   You're redirected to the pull connector instance form.

4. Fill in the fields.

Field	Description
Name	Name of the connector. The pull connector in stance name must match the push connector instance name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Description	Any optional information that you want to use to identify this record.
Host IP	Use your Datadog site's API host. For example, `api.datadoghq.eu`
Credential	The ServiceNow credential record that stores your Datadog API key or access token. The authentication method is controlled by the `useAccessToken` connector instance value. For more information, see the Connector instance value parameters table.
Active	This option appears only after the form is saved.
Bi-directional	Select to enable bi-directional exchange of values to and from the external event source. This option is available only when the connector definition has bi-directional values configured. After a work item alert is closed in ServiceNow, the corresponding case is automatically closed in Datadog.
Last bi-directional status	The value of this field is automatically populated. This option appears only when the **Bi-directional** option is selected.The status values are: ○ None - A valid connection has not yet been established. ○ Success - A successful connection was established. ○ Error - A connection was established. However the external event source was not updated.

Connector instance value parameters:

Parameter	Description
useAccessToken	Controls the authentication method. By default, this value is false. The **Credential** option must store the Datadog API key and you must set **application_key**. The connector sends the `DD-API-KEY` and `DD-APPLICATION-KEY` headers. When this value is true, the **Credential** option must store the Datadog OAuth access token. The connector sends `Authorization: Bearer <token>` header (`application_key` is not used).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
application_key	The application key of the Datadog connec tor. By default, this value is empty which indi cates Basic authentication.
fetchCaseRelatedAlerts	Controls whether related alerts are fetched for the work item. The default is true. Set it to false to disable.
caseApiPath	Set as `/api/v2/cases/`.
eventApiPath	Set as `/api/v2/events/`.

5. Right-select the form header and select **Save**.
6. Select **Test connector** to verify the connection.
7. After a successful test, select the **Active** check box and then select **Update**.

Related topics

Configure a push connector

Configure a push connector

You can configure listeners and connectors to push event information to the instance or MID Server.

Event Management receives external events through a push connector or script. If you are using a script to collect events, no configuration is required. All other methods of collecting events require configuration.

After event collection configuration, Event Management can collect events that are sent directly to the instance, or via the MID Server, SNMP trap, or email.

Note: Integrating with push connectors is only available by installing the Event Management Connector plugin from the ServiceNow Store.

Create integration user

Follow the steps in Creating users to create a user and grant the [evt_mgmt_integration] role to the user.

Note: Verify that the **Identify Type** of the user is **Machine**.

Standard event form

All events must use a standard event form, regardless of how they arrive at the instance. For more information, see Event field format for event collection.

Event collection

Event Management facilitates communication and movement of data between external applications, data sources, services, and the ServiceNow AI Platform instance.

Event collection modes	Description
Provided by the base system	Listeners and connectors provided with the Event Management application:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event collection modes	Description
	• AWS CloudWatch, see Integrate AWS platform as a data source. • Azure event messages, see Integrate Azure Monitor as an authenticated data source. • BMC TrueSight events, see Event collection from BMC TrueSight and BMC TrueSight_v2. • Datadog events, see Integrate Datadog platform events. • Dynatrace events, see Integrate Dynatrace platform events. • Google Cloud Platform (GCP) events, see Integrate Google Cloud Platform (GCP) events • New Relic events, see Integrate New Relic platform events. • Generic event collection through the MID Server. • Generic event collection through the instance. Where event messages are received in XML, you can base your script on the example XML transform script provided with the base system.
Via MID WebService	You can configure the MID WebService Event Collector to push event messages to the MID Server. For more information, see Pushing events to the MID Server using web service API.
Create your own content	You can configure your own listener: • Via MID WebServer API. See Pushing events to the instance using web service API. • Via Custom Transformation. See Event collection from custom payloads.
Via SNMP	SNMP listeners provided with the Event Management application: Generic SNMPv2. See Configure event collection for SNMP traps. In addition, you can create your own custom listener. See Integrate with push connectors.
Via Email	Email listener. See Configure event collection from email.

Related topics

Push connector instance form

Configure ServiceNow Cloud Observability event collection

Integrate ServiceNow Cloud Observability with Event Management by adding a standard webhook in the ServiceNow Cloud Observability platform. Download the Event Management Connector plugin from the ServiceNow Store so you can integrate with ServiceNow Cloud Observability.

Before you begin

- Role required: evt_mgmt_integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Authenticate ServiceNow Cloud Observability as a data source to enable Event Management to collect events from ServiceNow Cloud Observability. In your ServiceNow Cloud Observability platform, use a standard webhook to set your ServiceNow AI Platform instance as the rest endpoint.

Procedure

1. Create a CI.

   Note: Discovery for ServiceNow Cloud Observability services is not supported. Create a CI manually in a ServiceNow instance (cmdb_ci_service_auto table) to enable binding.

   a. Navigate to **All** > **Configuration** > **Application Service**.

   b. Select **New**.

   c. In the form, enter the name.
   The name must be the service name provided in the ServiceNow Cloud Observability platform.

   d. From the Operational Status list, set the operational status to **Operational**.

   e. Select **Next**.

   f. Select **Done**.

2. In your ServiceNow Cloud Observability platform, create a notification destination.

   a. From the navigation bar, click#**Alerts**.

   b. Select the#**Notification destinations** #tab.

   c. Select#**Create a destination**.

   d. Select **webhook**.

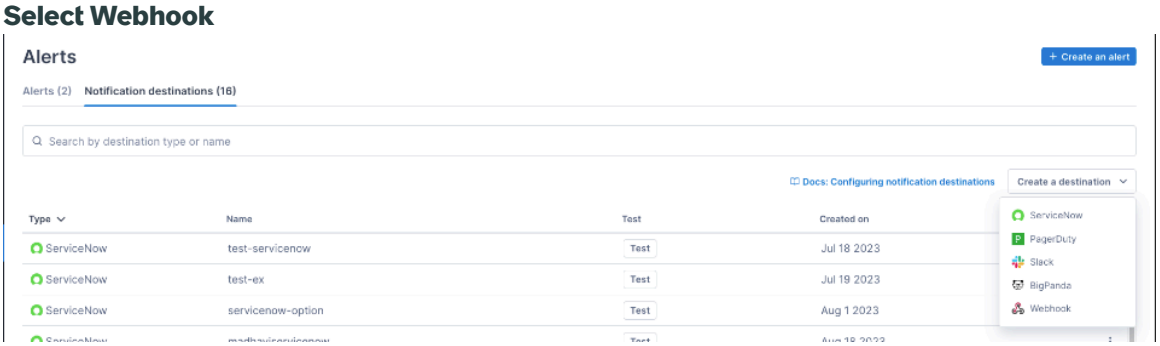

   e. In the Name field, enter a name for the webhook.

   f. In the URL field,#provide the webhook URL.
   This is the default #format of the URL to push events from#ServiceNow Cloud Observability to the#ServiceNow AI Platform Instance. Use this format: `https//`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
&lt;username&gt;:&lt;password&gt;@&lt;instance-name&gt;.service-now.com/
api/sn_em_connector/em/inbound_event?source=lightstep
```

g. Select **New Header**.

h. In the Headers field, enter the **service-name** and define its **value**.

Note: In a ServiceNow instance, CI binding occurs only if the **service-name** header value matches the name fields of the CI you created in the**cmdb_ci_service_auto table**.

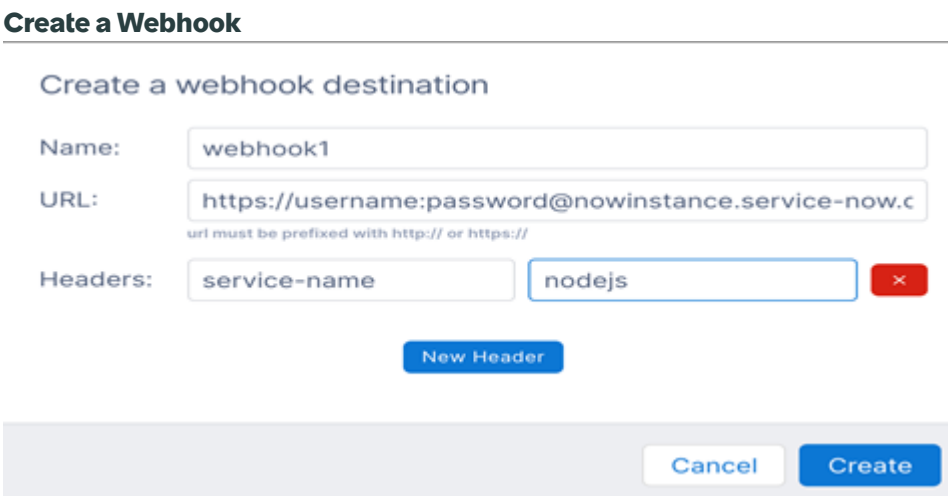

i. Click **Create**.

3. In the ServiceNow Cloud Observability platform, configure an alert.

a. From the navigation bar, click#**Alerts**.

b. Select#**Create an Alert** and enter an alert name.

c. In the Alert Configuration area, define the threshold.

d. In the Notification Rules area, in the Send Notifications to fields, select the notification destination as **webhook**.

e. Search for the destination name for the webhook that you created in step 2.

f. Select **Save**.
The default severity for stream based alerts is **Critical**. You can change the default in the #Push Connector Configuration area of an ServiceNow AI Platform instance. The valid severity values are:

- 1- Critical
- 2- Major
- 3- Minor
- 4- Warning
- 5- Info

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure a MID Server push listener for event collection

The MID Server supports the collection of event messages using the MID Web Server to collect data from external sources. The data is transformed (parsed) to the format required for event fields in the event [em_event] table and transmitted to the instance.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_admin

About this task

The MID Web Server extension enables external clients to push metric data and events to the MID Server. See Configure the MID Web Server extension. Use the MID WebService Event Collector to enable a URL method to push event messages to the MID Server. See Configure the MID WebService Event Collector Context.

Configure a MID Server with either *`EventManagement`* or *`All`* in the Supported Applications List.

To ensure uninterrupted services, consider also configuring a failover MID Server cluster.

Procedure

1. Navigate to **All** > **MID Server** > **Servers**.
2. Double-click the MID Server that you want to configure, or click **New** to create a new MID Server.
3. In the **Supported Applications** section, click **Edit**.
   a. Click **Edit**.
   b. In the slushbucket, either select **Event Management** or **ALL** and click the **>** add button.
   c. Click **Save**.
4. To ensure that the MID Server is validated, under Related Links, select **Validate**. For more information, see Validate a MID Server.
5. Click **Update**.

What to do next

Configure the MID WebService Event Collector to enable a URL method to push event messages to the MID Server. For more information, see Configure the MID WebService Event Collector Context.

Related topics

- Configure key-based MID Web Server authentication
- Configure the MID Web Server extension
- Configure the MID WebService Event Collector Context

Pushing events to the MID Server using web service API

Configure the MID WebService Event Collector to provide a URL method to push event messages to the MID Server.

The MID WebService Event Collector, by default, enables you to connect to an event generator and send event messages through the MID Server, using a URL in the format: `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/jsonv2`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The JSON v2 format is the same as the format that clients use to send event messages to the instance. This URL provides good performance.

For more information, see Configure the MID WebService Event Collector Context.

Note: Event Management accepts events in the timezone of the MID Server through which the listener has sent the event message. If you require a different timezone, you can configure the script as required, see Integrate with push connectors.

URL to push custom payloads

The MID Server can also receive a custom payload in JSON, XML, or plain text format. In this case, the MID Server transforms the event messages using the script include that you provide, and then sends the events to the instance. The URL in this case is:

`http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/inbound_event?Transform={Transform_script_name}` For an example of a script include, see the default `TransformEvents_xmlSample` script include.

For an example of how the MID WebService Event Collector transforms JSON formatted event messages, see Event collection from BMC TrueSight and BMC TrueSight_v2.

To send events in generic JSON format, you can use the URL in this format:

```
http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/inbound_event?Transform=TransformEvents_MidJsonGeneric
```

In this case, the `TransformEvents_MidJsonGeneric` MID script include that is provided in the base instance transforms the event messages into the event fields. All fields are added as **Additional Information** fields into the ServiceNow event that is generated.

Note: The URL in the format `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/{transform_script_name}`is also supported.

The REST API URL for validating connectivity is: `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/ping`

Script include to transform collected events

Use the required MID Server script include to transform (parse) collected event messages and populate them into the mapped event field. The MID Server script include name is composed of a prefix and suffix. The prefix is the mandatory text `TransformEvents_` and the suffix is the transform name on the URL, for example, `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/xmlSample` . The composed script name is therefore `TransformEvents_xmlSample`.

You must supply a script include that receives the event messages in a text variable. The script must have a `transform()` function that performs the transformation and prepares the return array of Event objects. When the Event Management plugin is activated, you are provided with the `TransformEvents_xmlSample` sample include script that transforms events from XML format.

Note: Event fields that are not identical to fields in the event table are saved in the **Additional information** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event message headers

Specify the relevant request header value for the **Content-Type** field according to the format of the event message.

Request header values for the Content-Type

Message format	Content-Type value
JSON	application/json
XML	application/xml
text	text/plain

Authentication and data security options

Options for authentication and data security:

- **Authentication type**, which can be set to either `API Key`, `mTLS`, or `Basic`.
- **Secure Connection**, which lets you choose whether incoming and outgoing data is secured when transmitted. If you choose the advanced secured option, it requires that you obtain a certificate from a well-known certificate authority, and then provide the **Keystore Certificate Alias** and the **Keystore Password**.
- **Use MID Unified Keystore**, which provides encryption protection and enables you to install a custom certificate in the MID Server unified keystore. For details, see Install custom certificates in the MID Server unified key store.
- Users must supply a script that receives the JSON, XML, or text formatted event messages in a text variable. The script must have a transform() function that performs the transformation and prepares the return array of event objects. The `TransformEvents_xmlSample` include script that transforms events from XML format is provided as an example when the Event Management plugin is activated.

For more information about authentication and data security, see Configure the MID Web Server extension.

Related topics

Pushing events to the instance using web service API

Pushing events to the instance using web service API

You can use a web service interface, supported by ServiceNow, that operates on the JSON object as the data input and output format.

Before you begin

Role required: evt_mgmt_integration

About this task

To insert records in the event [em_event] table with a single call, use this web service API:
`https://<instancename>.service-now.com/api/global/em/jsonv2`

Note: Business rules on the event table are not invoked when this URL is used.

Do not add additional fields to an event by adding a custom field to the event table [em_event]. However, additional fields should be included in the **Additional information** [additional_info] field of the event. For more information about how to include additional fields in events, see Custom alert fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information about troubleshooting performance issues with inbound integrations, see Troubleshooting inbound integrations performance.

The URL in the format: `http://<instancename>.service-now.com/em_event.do?JSONv2&sysparm_action=insertMultiple` is also supported in case you do want to activate Business rules. However, the performance of the first URL is superior.

To work with the URL that ends with `=insertMultiple`, add the following properties, with their respective values, to the MID Server:

- "mid.probe.event.queue.compress" - value: false
- "mid.probe.event.bulk_size" - value: 100
- "mid.probe.event.endpoint.url" - value: em_event.do?JSONv2%26sysparm_action=insertMultiple

Procedure

1. Send the request with these headers:

Parameter	Type	Description
Accept	String	The acceptable type for this message. The default value is `application/json`.
Content-Type	String	The content type for this message. The default value is `application/json`.
POST	String	The request type is POST, with one or more trailing records.

2. One or more events in JSON format can be sent as the payload of the web service call. Event fields that should be populated are:

Variable	Description
source	The name of the event source type. For example, SCOM or SolarWinds.
event_class	Specific instance of the *source*. For example, `SCOM 2012 on 10.20.30.40`
node	The **node** field should contain an identifier for the Host (Server/Switch/Router/etc.) that the event was triggered for. The value of the **node** field can be one of the following identifiers of the Host: ○ Name ○ FQDN ○ IP ○ Mac Address If it exists in the CMDB, this value is also used to bind the event to the corresponding ServiceNow CI.
resource	Populate this field with the device or application name. Use this field if the event refers to a device (such as Disk, CPU, or Network Adapter) or to an application or service running on a Host. Examples: `Disk C:\`, `Nic 001`, or `Trade web application`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Variable	Description
metric_name	Name of the metric that triggered the alert. For example, `Used Memory` or `Total CPU utilization`.
type	The type of event. This type might be similar to the **metric_name** field, but is used for general grouping of event types.
message_key	This value is used for de-duplication of events. For example, there might be two events for the same CI, where one event has CPU of 50% and the next event has CPU of 99%. Where both events must be mapped to the same ServiceNow alert, they should have the same message key. The field can be left empty, in which case the field value defaults to `source+node+type +resource+metric_name`. The **message_key** should be populated only when there is a better identifier than the default.
severity	Severity of the event. ServiceNow values for severity range from 1 – Critical to 5 – OK, with the severity of 0 – Clear. Original severity values should be sent as part of the additional information.
additional_info	This field contains a JSON string (stringified JSON) with key/value pairs representing supplemental information not mapped to standard event fields. Examples include IDs of objects in the event source, event priority (if different from severity), and assignment group information. Values not already in JSON key/value format are normalized to stringified JSON when the event is processed.
time_of_event	Time when the event occurred on the event origin. The format is: yyyy-MM-dd HH:mm:ss The specified time must be according to GMT.
resolution_state	Optional – To indicate that an event has been resolved or no longer occurring, some event monitors use 'clear' severity. Others use a 'close' value for severity. This field is used for monitors that use the 'close' value. Valid values are `New` and `Closing`.

3. To create multiple records with a single call, trigger the event web service using the following URL:

   Replace the <instance name> variable with the name of the required instance.

   `https://<instancename>.service-now.com/api/global/em/jsonv2`

   Example
   Example showing the payload for two events that are sent in a single web service call:

```
{
  "records": [
    {
      "source": "SCOM",
      "event_class": "SCOM 2012 on scom.server.com",
      "resource": "D:",
      "node": "name.of.node.com",
      "metric_name": "Percentage Logical Disk Free Space",
      "type": "Disk space",
      "severity": "4",
      "description": "The disk D: on computer
 V-W2K8-abc.abc.com is running out of disk space. The value
 that exceeded the threshold is 38% free space.",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      "additional_info": "{\"scom-severity\":
 \"Medium\", \"metric-value\": \"38\", \"os_type\":
 \"Windows.Server.2008\"}"
    }
},
    {
      "source": "SCOM",
      "event_class": "SCOM 2012 on scom.server.com",
      "resource": "MSSQL-database-name",
      "node": "other.node.com",
      "metric_name": "DB Allocated Size (MB)",
      "type": "Database Storage",
      "severity": "3",
      "description": "High number of active connections for
 MSSQL-database-name running on name.of.node.com. Active
 connections exceed 5000.",
      "additional_info": "{\"scom-severity\":
 \"Medium\", \"metric-value\": \"38\", \"os_type\":
 \"Windows.Server.2008\"}"
      }
    }
  ]
}
```

4. Trigger the event web service using the following URL:
   Replace <instancename> with the name of the required instance. `https://<instancename>.service-now.com/api/global/em/jsonv2`

   Example
   Example showing the payload for one event that is sent in a single web service call:

```
{
  "records": [
    {
      "source": "SCOM",
      "event_class": "SCOM 2007 on scom.server.com",
      "resource": "C:",
      "node": "name.of.node.com",
      "metric_name": "Percentage Logical Disk Free Space",
      "type": "Disk space",
      "severity": "4",
      "description": "The disk C: on computer
 V-W2K8-dfg.dfg.com is running out of disk space. The value
 that exceeded the threshold is 41% free space.",
      "additional_info": "{\"scom-severity\":
 \"Medium\", \"metric-value\": \"38\", \"os_type\":
 \"Windows.Server.2008\"}"
      }
    }
  ]
}
```

   Example
   Example: cURL request:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
{
  curl -X POST
 "https://<instancename>.service-now.com/api/global/em/jsonv2"
 \
  --header "Content-Type: application/json" \
  --user "username:password" \
  --data '{
    "records": [
      {
        "source": "SCOM",
        "event_class": "SCOM 2007 on scom.server.com",
        "resource": "C:",
        "node": "name.of.node.com",
        "metric_name": "Percentage Logical Disk Free Space",
        "type": "Disk space",
        "severity": "4",
        "description": "The disk C: on computer
V-W2K8-dfg.dfg.com is running out of disk space. The value
that exceeded the threshold is 41% free space.",
        "additional_info": "{\"scom-severity\":
\"Medium\", \"metric-value\": \"38\", \"os_type\":
\"Windows.Server.2008\"}"
      }
    }
   ]
  }'
```

Related topics

Pushing events to the MID Server using web service API

Configure the MID WebService Event Collector Context

Configure the MID WebService Event Collector Context to provide a URL method to push event messages from an external source to the MID Server.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_admin

About this task

The default format of the URL to push event messages from an external source to the MID Server is `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/jsonv2`. This URL provides good performance.

From an external source, to push event messages that are not in jsonv2 format, the format of the URL is: `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/inbound_event?Transform={Name_of_Transform_Script}`, where the *`{Name_of_Transform_Script}`* variable is the full name of the script and always begins with the text: `TransformEvents_`.

For example, assume the following values:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- {MID_Server_IP}: 10.118.69.27
- {MID_Web_Server_Port}: 8097
- Transform script name: EventsToProcess

The URL to be used is therefore: `http://10.118.69.27:8097/api/mid/em/inbound_event/TransformEvents_EventsToProcess`

Note:

- The URL in the format `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/{transform_script_name}` is also supported.
- The date format for events is yyyy-M-d h:mm:ss.

  If you receive an event whose date is in a different format, you must use a `{transform_script_name}` that is appropriate for the incoming event's date format. If you do not, the event will not be processed correctly.

  For example, an event might arrive on June 27, 2019 at 11:25 AM with a listed date of *`2019/06/27/ 11:25:00 a`*. Use a `{transform_script_name}` with a date format of *`yyyy/MM/dd/ HH:mm:ss a`* to match the format of the received event.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **MID WebService Event Listener**.
2. In the MID WebService Event Collector Contexts list, click **New**.
3. On the form, fill in the fields.

MID WebService Event Collectors Context form

Field	Description
Name	A unique name for this collector for easy identification.
Short description	Enter a brief, meaningful description of this collector.
MID Web Server Extension	Specify and then start the MID Web Server extension. The supported authentication methods are listed in the **Authentication Type** field of the MID Web Server extension. For information about how to configure a MID Web Server extension, see Configure the MID Web Server.
Status	This field is auto-populated with the status of the MID Web Server extension. This field is blank until the MID Web Server extension is started. After issuing a command to the MID Web Server extension, one of the following values is displayed: ○ **Started**: The collector is running. ○ **Stopped**: The collector is not running. ○ **Offline**: The MID Server is down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ **Error**: The collector failed with an error (the error message is displayed in **Error Message**). ○ **Warning**: A run-time exception has occurred.
Execute on	**Specific MID Server** or **Specific MID Server Cluster**, as defined on the specified MID Web Server extension. **Note: Specific MID Server Cluster** is only for internal processing of ECC queue record and is not for external API load balancing.
MID Server	The **Specific MID Server** or **Specific MID Server Cluster**, as defined on the specified MID Web Server extension.
Executing on	The name of the MID Server on which the MID Web Server extension is running.

4. Right-click the form heading and click **Save**.
5. Under **Related Links**, click **Start** to start the collector.

Commands in the Event Management Context form

Related Link	Description
Start	If it is not running, start the collector. This action verifies that a web service API endpoint with the Event Management application is running on the MID Server.
Stop	Stops the running collector on the configured MID Server. If the collector is not running, no action is taken.
Restart	Stops, then starts the collector on the configured MID Server.
Update parameters	Sends updated parameters to the collector. Parameters are also updated when the Event Management MID Server context extension is updated. If you click this control when the collector is not running, no update is made.

Example:
Showing the use of the URL to transform JSON v2 formatted event messages

Assume that JSON v2 formatted event messages are sent to the MID Server. When using the `jsonv2` URL, there is no need to use a script include.

Data for the JSON v2 example

Field	Value
MID_Server_IP	10.218.64.27
MID_Web_Server_Extension_Port	8097
Event message format	jsonv2

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Replace the variables in the default format of the URL `http://<my-instance>.service-now.com/api/global/em/jsonv2`with values from the preceding table:`http://10.218.64.27:8097/api/global/em/jsonv2`

Example showing the URL to push messages not in jsonv2 format

The format of the URL to push event messages from an external source that are not in jsonv2 format is `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/inbound_event/Transform={Name_of_Transform_Script}`. The {Name_of_Transform_Script} variable is the full name of the script and always begins with the text: TransformEvents_. Specify the script name as the `Transform` header parameter with the prefix `TransformEvents_`.

For this example, assume that the script name is EventsToProcess, the URL is therefore:`http://10.138.64.27:8097/api/mid/em/inbound_event/TransformEvents_EventsToProcess`

Example showing JSON v2 formatted event messages

Note: When copying and pasting the text following, hidden characters might also be copied and can cause unexpected results.

```
curl -v -H "Accept: application/json" -H "Content-Type:
 application/json" -X POST --data "{
    \"records\":
    [ {
         \"source\" : \"Simulated\",
        \"node\" : \"nameofnode\",
        \"type\" : \"High Virtual Memory\",
        \"resource\" : \"C:\",
        \"severity\" : \"5\",
        \"description\" : \"Virtual memory usage exceeds 98%\",
        \"ci_type\":\"cmdb_ci_app_server_tomcat\",
        \"additional_info\":\"{\\\"name\\\":\\\"My
 Airlines\\\"}\"
      },
      {
      \"source\" : \"Simulated\",
      \"node\" : \"01.myairlines.com\",
      \"type\" : \"High CPU Utilization\",
      \"resource\" : \"D:\",
      \"severity\" : \"5\",
      \"description\" : \"CPU on 01.my.com at 60%\"
      }
   ]
}" -u UserName:Password
 http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/jsonv2
```

Example:
Example showing JSON v2 formatted event messages with MID Web Server API key

Note: When copying and pasting the text following, hidden characters might also be copied and can cause unexpected results.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
curl --location -g --request POST
 'http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/json
v2' \
--header 'Accept: application/json' \
--header 'Content-Type: application/json' \
--header 'Authorization: key <mid_webserver_api_key>' \
--data-raw '{
   "records":
    [ {
         \"source\" : \"Simulated\",
        \"node\" : \"nameofnode\",
        \"type\" : \"High Virtual Memory\",
        \"resource\" : \"C:\",
        \"severity\" : \"5\",
        \"description\" : \"Virtual memory usage exceeds 98%\",
        \"ci_type\":\"cmdb_ci_app_server_tomcat\",
        \"additional_info\":\"{\\\"name\\\":\\\"My
 Airlines\\\"}\"
     },
     {
     \"source\" : \"Simulated\",
     \"node\" : \"01.myairlines.com\",
     \"type\" : \"High CPU Utilization\",
     \"resource\" : \"D:\",
     \"severity\" : \"5\",
     \"description\" : \"CPU on 01.my.com at 60%\"
     }
   ]
}'
```

Example: Configuring external API integrations with MID Server failover

To ensure high availability for external API integrations, deploy an external load balancer that distributes traffic across multiple MID servers. Note that the MID cluster web server configuration handles internal processing only (through ecc_queue) and can't manage inbound external API traffic.

Recommended architecture for external API integrations

Field	Value
MID Server 1	mid-1.internal:8080
MID Server 2	mid-2.internal:8080
External Load Balancer	mid-cluster.example.com:443
API endpoint format	jsonv2

Reference the external load balancer hostname in your webhook URL rather than individual MID Server addresses. Replace the variables in the default format: `https://<load-balancer>/api/mid/em/inbound_event?Transform=<format>` with values from the preceding table: `https:///mid-cluster.example.com/api/mid/em/inbound_event?Transform=jsonv2`.

This configuration distributes incoming requests across both MID servers. If one server becomes unavailable, the load balancer automatically routes traffic to the remaining operational server without service interruption.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Configure the MID Web Server extension
- Event collection from BMC TrueSight
- Event collection from Microsoft Azure Monitor
- Event collection from Google Cloud Platform

Event collection from custom payloads

The MID WebService Event Collector enables you to collect event information from custom payloads in JSON, XML, or plain text format.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_admin

About this task

The MID Server transforms the collected event messages and populates the Event table fields (em_table) in the instance. The collection of formatted event messages is described in this procedure using basic authentication. For information about supported authentication methods, see Configure the MID Web Server extension.

The format of the required MID Server URL is: `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/inbound_event?Transform={Transform_script_name}`. The name of the MID Server script include is composed by appending a suffix to the default `TransformEvents_` prefix. For information about the collection of JSON v2 event messages, see Configure the MID WebService Event Collector Context.

Note: The URL in the format `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/{transform_script_name}` is also supported.

Procedure

1. Configure the MID WebService Event Collector, see Configure the MID WebService Event Collector Context.
2. Start the MID WebService Event Collector.

Example:

Transformation of XML formatted event messages using the custom payload URL

Assume that XML formatted event messages are sent to the MID Server. Use this example to return an array of event objects from the collected event messages. The name of the MID Server script include is composed by appending a suffix to the default `TransformEvents_` prefix. For the purposes of this example, the user supplied the xmlSample script include. Using these details, the name of the MID Server script include is `TransformEvents_xmlSample`. The MID Server transforms the collected event messages by parsing the messages using the script include and then transmitting them to the instance.

Data for the XML example

Field	Value
MID_Server_IP	10.218.64.27

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data for the XML example (continued)

Field	Value
MID_Web_Server_Extension_Port	8097
transform_script_suffix _name	xmlSample

Replace the variables in the URL with the values from the above table:
`http://10.218.64.27:8097/api/mid/em/xmlSample`

Note: When copying and pasting the text below, hidden characters might also be copied and can cause unexpected results.

Example showing XML formatted event messages:

```
<records>
    <event>
    <source>My Source</source>
    <node>host1</node>
    <type>type1</type>
    <severity>3</severity>
    <description>Virtual memory usage exceeds 98%</description>
    </event>
    <event>
    <source>My Source</source>
    <node>host2</node>
    <type>type2</type>
    <severity>2</severity>
    <description>Virtual memory usage exceeds 90%</description>
    </event>
    </records>
```

Related topics

Configure the MID Web Server extension

Configure the MID WebService Event Collector Context

Integrate with push connectors

Integrate with a push connector to connect to an external event source. Push connectors process the collected event messages and transform them to the required event format.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed.

Role required: evt_mgmt_admin

Note: If you are upgrading a ServiceNow AI Platform instance from Paris or earlier, you can use legacy listener transform scripts. See Use legacy listener transform scripts for further information.

About this task

The push connectors accept event messages that are generated by external event sources. Configure a connector to listen to an external event source. Using a custom push connector, send the event messages through either the MID Server or the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can use this generic JSON target URL to collect events:`https://<<INSTANCE>>/api/sn_em_connector/em/inbound_event?source=genericJson`. This URL can be used as-is and requires that an event rule be configured.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Push Connectors**.
2. Click **New** or click the push connector that you want to modify, for example, AWS or Azure.
3. On the form, fill in the fields.

Field	Description
Name	Unique name for this push connector for easy identification.
Description	Enter a description for the use of the push connector.
Header name	Specify the name of the header that you want the value for, such as Set-Header.
Header value	The value of the header name that you want the value for, such as Set-Header.
Type	Select to send events either `Instance` or `MID`. In each case, use the URL of the selected type. See the description of the **URL parameter value** field.
Active	Select to enable pulling events from this external event source.
Order	Order in which an event rule is evaluated when multiple rules are defined for the same type of event. Event rules are evaluated in ascending order.
URL parameter value	○ Send events to the instance. Use the URL in this format for all connectors: `https://<<INSTANCE>>/api/sn_em_connector/em/inbound_event?source={URL_parameter_value}` ○ Send events through the MID Server. Use the URL in the format:`http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/em/inbound_event?Transform={Transform_script}` **Note:** Sending events to the MID Server requires the prior configuration of the MID WebService Event Listener extension.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the Script section:
   - If the value selected for the **Type** field is `MID`, the **Transform script** field appears. In this field, specify or search for the name of the MID script include. The script accepts event messages from the required external event source and parses them into the required event format. Use this naming convention for the script: `TransformEvents_<your source>`
   - If the value selected for the **Type** field is **Instance**, the **Script** editor appears. In the **Script** editor, enter the customized script that accepts event messages that the required external event source generates and that the script parses into the required event format.

Example
This example shows the fields that have been transformed, being added to an event form.

```
(function process(/*RESTAPIRequest*/ request, body) {
 /*Function that receives a JSON object, adding all its fields
 to the Additional information object. The field name is a
 concatenation of the field key and the parent field key if it
 exists.*/
 function updateAdditionalInfo(event,
 field,jsonObject,additionalInfo) {
        for (var key in jsonObject) {
            var newKey = key;
            if (field != "") {
                newKey = field + '_' + key;
            }
   // You can do some transformation here and set fields on the
 event
   //if(key == "MySource")
   //   event.source = jsonObject[key];
            additionalInfo[newKey] = jsonObject[key];
        }
    }

    try
 {
        gs.info("TransformEvents_generic received body:" +
 body);
  var jsonObject = JSON.parse(body);
        var event = new GlideRecord('em_event');
  event.source = "GenericJson"; //TODO: Need to define
        event.event_class = "GenericJsonClass"; //TODO: Need to
 define
        event.severity = "5";

        var additionalInfo = {};
        updateAdditionalInfo(event,
 "",jsonObject,additionalInfo);
  /*Iterates over Additional information JSON object and
 adds all nested objects' fields as fields of the Additional
 information object*/
        var notDone = true;
        while (notDone) {
            notDone = false;
            for (var key in additionalInfo) {
                if
 (Object.prototype.toString.call(additionalInfo[key]) ==
 '[object Object]') {
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                    notDone = true;
                    updateAdditionalInfo(event,
 key,additionalInfo[key],additionalInfo);
     additionalInfo[key] = "";
                 }
             }
         }
  gs.info("TransformEvents_generic generated additional
 information:" + JSON.stringify(additionalInfo));
        event.additional_info = JSON.stringify(additionalInfo);
        event.insert();
 }
 catch(er){
  gs.error(er);
  status=500;
  return er;
 }
 return "success";
})(request, body);
```

Use legacy listener transform scripts

Use legacy listener transform scripts when upgrading a ServiceNow AI Platform instance from Paris or earlier. These scripts can be run as part of existing integrated systems, or in unison with Event Management push connectors.

You can use listener transform scripts to ensure that previous integrations are supported. For example, if you are using the `https://<instance-name>.service-now.com/api/global/em/inbound_event` endpoint in various existing integrations, you can continue to use this endpoint in a Quebec or later ServiceNow AI Platform instance.

If you want to use the Event Management push connectors, download the plugin from the ServiceNow Store website. You can then access the push connectors in the ServiceNow AI Platform instance, and the `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event` endpoint is also created.

Use the legacy endpoint with push connectors

You can continue to use the `https://<instance-name>.service-now.com/api/global/em/inbound_event` endpoint with the push connectors, by deactivating the relevant transform scripts. To deactivate a script, navigate to **Event Management** > **Integrations** > **Listener Transform Scripts**. Then click the relevant script in the displayed list, and clear the **Active** check box in the form.

Use the legacy endpoint and deactivate all listener transform scripts

You can also continue to use the `https://<instance-name>.service-now.com/api/global/em/inbound_event` endpoint but deactivate all listener transform scripts to use the push connectors instead. To deactivate all scripts, navigate to **System Properties** > **All Properties**, locate the *sn_em_connector_listener_table_for_InboundEvent* property, and set its value to **True**.

Migrate custom listener transform scripts to the push connectors module

If you have created custom listener transform scripts, they must be migrated to the push connectors module if you want to use the new `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event` endpoint. To migrate the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

custom script, navigate to **Event Management** > **Integrations** > **Listener Transform Scripts**, click the relevant listener transform script, and select the displayed script. Then navigate to **Integrations** > **Push Connectors**, click **New**, and paste in the script you copied from the listener transform script.

Integrate AWS platform as a data source

Integrate Amazon Web Services (AWS) with Event Management. To add AWS platform as a data source, configuration is required in the AWS platform.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Download the Event Management Connectors plugin from the ServiceNow Store.

Role required: evt_mgmt_admin

About this task

When an AWS platform alarm arrives, Event Management:

- Extracts information from the original AWS platform alarm to populate the required event fields and inserts the event into the database.
- Captures the content in the **additional_info** field.

The AWS platform push connector is located in **Event Management** > **Integrations** > **Push Connectors**. In the Push Connectors page, click **AWS Events Push Connector**.

Note: The AWS push connector that is provided in the base system handles AWS CloudWatch alarms only and not CloudWatch events created through event rules. To handle Simple Notification Service (SNS) alarms that are other than AWS CloudWatch, create a new push connector or customize the AWS push connector.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

Procedure

1. In the AWS platform console, select **Simple Notification Service** .
   If an SNS topic does not exist, create one.
2. Under the SNS topic, create a new subscription.

   a. Take **Topic ARN** from the topic that you created.
   The Amazon Resource Name (ARN) is necessary for binding an Event Management alert to a CI.

   b. Set **Protocol** to: *https*.

   c. In the ServiceNow instance, create a non-interactive user account or identify an existing non-interactive ServiceNow user for this integration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Assign the user only the `evt_mgmt_integration` role. Configure the password according to the password policy for your ServiceNow instance.

d. Set **Endpoint** to: `https://<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=aws`

If AWS platform multi-factor authentication (MFA) is enabled, when signing in to the AWS platform website, you're prompted for the user name and password, as well as an authentication code from the AWS platform MFA device of the user.

3. Wait until the subscription changes from `Pending` to `Confirmed` and the subscription ARN is populated.
This can take a few minutes.
4. Create alarms in AWS platform to send to Event Management and link the alarms to the SNS topic that you created.
Raw message delivery must not be enabled in AWS alarm. These following event rules are provided with the base system:

Event rule	Description
AWS host binding	Bind AWS platform alarms, on either the host or virtual machine (VM), to the host Hardware CI.
AWS - General	Bind AWS platform alarms on the Load Balancer (LB) to the Cloud Load Balancer CI, Amazon Relational Database Service (RDS) to the Cloud Database CI, and Host or Virtual Machine (VM) to the Virtual Machine Instance CI. **Note:** To bind events from the Host or Virtual Machine to the Virtual Machine Instance CI, disable the AWS host binding rule.

Note: For the push connector to handle the incoming event, you must sent a request header **x-amz-sns-message-type** with value **Notification**.

This example shows a JSON string for AWS VM binding:

```
{
"Type" : "<type>",
"MessageId" : "<message ID #>",
"TopicArn" : "<topic arn>",
"Subject" : "ALARM: \"<alarm type>\" in <location>",
"Message" : "{\"AlarmName\":\"<alarm
 name>\",\"AlarmDescription\":<alarm
 description>,\"AWSAccountId\":\"<account ID
 #>\",\"NewStateValue\":\"ALARM\",\"NewStateReason\":
\"Threshold Crossed: <threshold and datapoint which
 it crossed>.\",\"StateChangeTime\":\"<date and
 time>\",\"Region\":\"<location>\",\"OldStateValue\":\"OK\",
\"Trigger\":{\"MetricName\":\"CPUUtilization\",\"Namespace\":
\"<namespace>\",\"Statistic\":\"<statistic
 name>\",\"Unit\":null,\"Dimensions\":[{\"name\":
\"InstanceId\",\"value\":\"<dimension
 value>\"}],\"Period\":<period
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 number>,\"EvaluationPeriods\":<number of evaluation
 periods>,\"ComparisonOperator\":
\"GreaterThanOrEqualToThreshold\",\"Threshold\":<threshold
 number>}}",
"Timestamp" : "<date and time>",
"SignatureVersion" : "<version number>",
"Signature" : "<signature>",
"SigningCertURL" : "<URL>",
"UnsubscribeURL" : "<URL>"
}
```

5. To enable multiple member accounts to forward CloudWatch or EventBridge events to a centralized account, where a single SNS topic delivers them to the ServiceNow Event Ingestion endpoint, you need to make configuration on AWS side.
   To configure it, see https://repost.aws/knowledge-center/cloudwatch-cross-account-sns.
   This design simplifies integration, avoids duplicating SNS topics across accounts, and provides a single entry point for all monitoring events flowing into ServiceNow.

Related topics

Event field format for event collection

Integrate with push connectors

Integrate AWS with REST API key token

Integrate using an API key to establish secure communication and automate data exchange via REST API. This simplifies integration, enabling seamless access to services and enhancing operational efficiency.

Before you begin

Role required: evt_mgmt_integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API Key**.

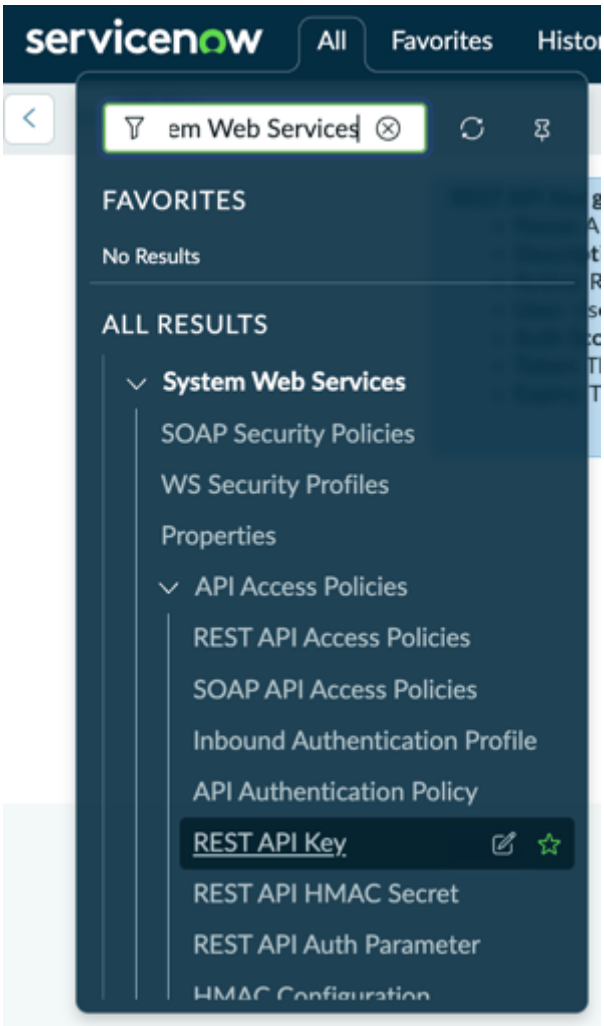

2. Select **New**.
3. On the form, fill the fields.

> **Note:** The user needs to be created with the evt_mgmt_integration role. Create a new one on the sys_user table to use here to facilitate the right access.

API Key

Field	Description
Name	Name to identify the REST API Key
Description	Description for the REST API Key.
Active	Status of the REST API Key.
User	User associated to the REST API Key. Use the look-up icon to select the user.
Auth Scope	Option to add auth scope to manage the authority of the REST API Key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Token	The REST API key generated by the ServiceNow AI Platform. Copy the key to use as part of the REST API call within the Header or Query parameter.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration.

4. Unlock Auth Scope and add: UserAccount.
5. Use the form menu and select **Save**.

   The system generates a token and saves it in the **Token** field. To see the token, use the lock icon and copy the contents display below the field. This is your query parameter (or header) value when your other system sends a REST API request to ServiceNow.

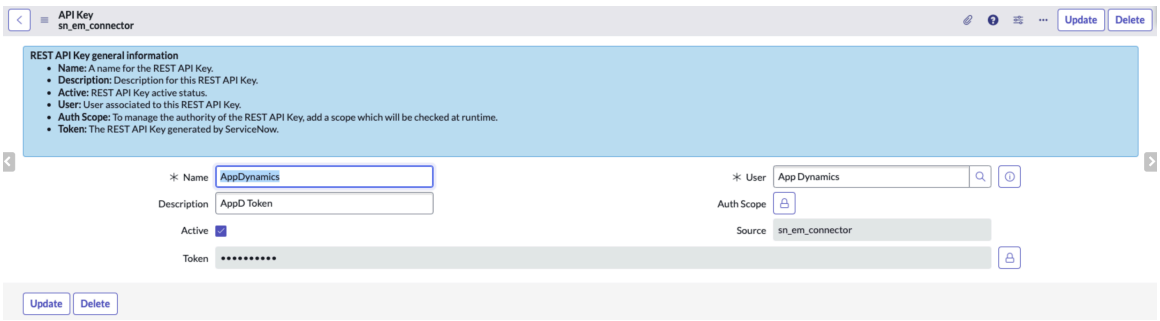

6. In the AWS environment, under **Simple Notification Service**, create a new topic and then create a new subscription by selecting the topic that you have created.
7. In the Create Subscription window, enter the ServiceNow endpoint value in the **Endpoint** field. Format of the api-key value must be: `https://[INSTANCENAME].service-now.com/api/sn_em_connector/em/inbound_event?source=[SOURCENAME]&sys_id=[SYSID OF PUSH CONNECTOR INSTANCE]&x-sn-apikey=[AUTH TOKEN]`.

   The [AUTH TOKEN] value is the system-generated token displayed in the **Token** field.

Integrate Azure Monitor as an authenticated data source

Integrate Microsoft Azure with Event Management by adding the Azure Monitor as an authenticated data source.

You can configure the Event Management environment for the collection of events from Azure Monitor by setting your ServiceNow AI Platform instance as the rest endpoint.

Once the endpoint is configured, when an Azure Monitor alert message arrives, Event Management:

- Authenticates the Azure Monitor alert message with the relevant ServiceNow user, using OAuth configuration or a standard webhook.
- Extracts information from the original Azure Monitor alert message to populate required event fields and inserts the event into the ServiceNow AI Platform database.
- Captures specified content in the **Additional Information** field of the event form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What authentication is used

There are two methods of authentication:

- OAuth authentication: Provides enterprise-grade authentication to keep your enterprise environment safe. Authentication is performed using Azure Monitor V1 or V2 access tokens. For more information, see Integrate Azure Monitor with OAuth authentication.
- API key authentication: Integrate using API key to secure communication between Azure and ServiceNow. For more information, see Integrate Azure with REST API key token.
- Basic webhook authentication: Provides a basic standard of authentication, without the need for Azure Active Directory. This authentication can be especially useful for distributed small teams, such as SRE or DevOps teams. For more information, see Integrate Azure Monitor with basic authentication.

What to know before you begin

You can use your integrated Azure Monitor as a data source only after you have verified the following:

- For both methods of authentication, the relevant ServiceNow sys_user is assigned the evt_mgmt_integration role.
- The Event Management Connectors plugin is installed in the ServiceNow AI Platform instance. You can download the plugin from the ServiceNow Store website.
- Azure Cloud Discovery must be performed to ensure that the created alerts are bound to the configuration items in the ServiceNow AI Platform. For more information, see Discovery for Microsoft Azure.

Event Rules and Event Field mappings

These event rules and event field mappings are provided with the base system:

Module	Description
Event Rules	**Azure Monitor**: A general event rule to handle all Azure Monitor events.
Event Field Mappings	**Azure Monitor - ci_type**: To map ci_type of events based on resourceType field. A base set of mapping pairs are provided.

These are the mappings provided with the base system in **Azure Monitor - ci_type**:

Transform Value Pairs

Event Mapping Pairs

Key	Value
microsoft.compute/virtualmachines	cmdb_ci_vm_object
microsoft.documentdb/databaseaccounts	cmdb_ci_cloud_database
microsoft.sql/servers	cmdb_ci_cloud_database
microsoft.web/sites	cmdb_ci_cloud_webserver
microsoft.cache/redis	cmdb_ci_cloud_database
microsoft.documentdb	cmdb_ci_cloud_database
microsoft.sql/servers/databases	cmdb_ci_cloud_database

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can add new mapping pairs to the Event Field Mapping - **Azure Monitor - ci_type** as per the requirement, to map events to the respective ci_type based on resourceType.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the connector.

Note: This feature applies only to active event rules.

If you want to send alert state changes on the ServiceNow instance from the ServiceNow alerts to the Azure Portal, you need to enable the Azure Monitor Bi-directional connector. For more information, see Configure Azure Monitor Bi-directional connector.

Severity mapping from Azure severity to ServiceNow event severity

Azure severity condition	ServiceNow event severity
When an Azure alert monitorCondition is Fired	
Azure Sev0	ServiceNow Critical (severity "1")
Azure Sev1	ServiceNow Major (severity "2")
Azure Sev2 and Sev3	ServiceNow Warning (severity "4")
Azure Sev4	ServiceNow OK (severity "5")
When an Azure alert monitorCondition is resolved	
Any Azure severity	ServiceNow CLEAR (severity "0")

Azure Monitor Issue integration

An Azure Issue is a unified case that aggregates related alerts and signals from Azure Monitor into a single, trackable operational problem. It preserves investigation context and serves as a durable record for incident management Incident Management, so you can resolve problems faster instead of tracking scattered alerts.

Before you begin

Verify you already have an Azure Monitor Workspace.

Role required: evt_mgmt_admin

About this task

ServiceNow represents an Azure Issue as a single alert. View the alerts that Azure grouped into the Issue in the **Related Records** tab from the Express List side panel or the alert record in Service Operations Workspace. ServiceNow retains each related alert as an individual alert. This approach lets ServiceNow group and correlate alerts across multiple resources and monitoring systems. ServiceNow is not limited to the alerts that Azure grouped into the Issue. As a result, ServiceNow identifies broader incidents, reduces duplicate alerts, and provides a more comprehensive operational view. For this reason, ServiceNow represents an Azure Issue as a single alert instead of an alert group.

Note the following before you start:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you already use the Azure Push Connector to send Microsoft Azure alerts to ServiceNow, use the same push connector URL to send Azure Issues.
- If you're configuring Azure alerts/issues for the first time, use the Azure Push Connector URL to create an action group in the Azure portal. Use the following URL in the Azure action group secure webhook field:

```
https://<instance_name>.service-now.com/api/sn_em_connector/em/inbound_event?source=azuremonitor
```

- If you have different tenants with different Azure Service Principal credentials (per tenant or per subscription), create separate Azure Push Connector instances. Use them to push Azure alerts or issues into your ServiceNow instances.

You can integrate Azure Monitor issues using Integration Launchpad. For more information, see Configure Azure Monitor Issue in Integrations Launchpad.

Procedure

1. In the Azure Monitor portal, perform the following steps:
   a. Navigate to your Azure Monitor Workspace.
   b. Navigate to **Settings** > **Action groups** and select the existing action group that you use to send Azure alerts to ServiceNow.

      This creates an alert, and in turn, the Alert Management rule creates an incident. If you're configuring Azure issues for the first time, first create an action group by following the steps in Integrate Azure Monitor with OAuth authentication.

   c. If you need related alerts and resource information for an issue in the Azure portal under the **Alerts** and **Resources** tab and for bi-directional synchronization, configure the pull connector by performing the steps as mentioned below.
2. In the ServiceNow AI Platform instance, configure a pull connector if you have used a push connector to send Azure issues by perform the following steps:
   a. Navigate to the push connector you're using by going to **All** > **Push connectors**.
   b. Select the Azure Monitor push connector.
   c. Select **Configure Issue Alerts** on the push connector form.

      You're redirected to the pull connector instance form.

   d. If the Azure Monitor pull connector is already activated, verify your credentials and that the instance is active, then select **Test Connection** to validate.
   e. If the Azure Monitor pull connector is not activated yet, configure it as follows:
      i. Leave the pull connector instance name as `Azure Monitor`.
      ii. In the **Credentials** field, enter your Azure Service Principal credential.
      iii. Leave the **Host** field as `1.1.1.1` as it will not be used.
      iv. Select the **Active** check box to activate the connection.
      v. Set *`fetchIssueRelatedAlerts`* to control whether the system fetches related alerts for the issue.

         The default is true. Set it to false if you want to disable this behavior.

      vi. Set *`fetchIssueRelatedAlertsLimit`* to the maximum number of related alerts the system fetches for the issue. The default is 10.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vii. Select **Test Connection** to validate the connection.

If Azure bi-directional connector is enabled, when an Azure issue is closed in ServiceNow, it automatically closes in the Azure portal as well.

When an Azure issue is closed in ServiceNow, it automatically closes in the Azure portal as well.

3. In the ServiceNow AI Platform instance, configure a pull connector if you have used a push connector instance to send Azure issues by perform the following steps:
   a. Navigate to the push connector instance you're using in ServiceNow by navigating to **All** > **Push connector Instances**.
   b. Select the Azure Monitor push connector instance that was used to send Azure Issues in to ServiceNow instance.
   c. Select **Configure Issue Alerts** UI action present in the push connector form.

      You're redirected to the pull connect instance form.

   d. In the pull connector creation form, perform the following steps:

      i. Leave the pull connector instance name as is.

         The pull connector instance name must match with the push connector instance name.

      ii. In the **Credentials** field, enter your Azure Service Principal credential.
      iii. Leave the **Host** field as `1.1.1.1` as it will not be used.
      iv. Select the **Active** check box to activate the connection.
      v. Set *`fetchIssueRelatedAlerts`* to control whether the system fetches related alerts for the issue.

         The default is true. Set it to false if you want to disable this behavior.

      vi. Set *`fetchIssueRelatedAlertsLimit`* to the maximum number of related alerts the system fetches for the issue. The default is 10.
      vii. Select **Test Connection** to validate the connection.

         If Azure bi-directional connector is enabled, when an Azure issue is closed in ServiceNow, it automatically closes in the Azure portal as well.

Integrate Azure Monitor with basic authentication

Integrate Microsoft Azure with Event Management by adding a standard webhook in Azure Monitor.

Before you begin

Confirm that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Confirm you have created a user with an **Identify Type** of **Machine** and the evt_mgmt_integration role.

Roles required: evt_mgmt_admin and web_service_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Configure the Event Management environment for the collection of events from Azure by authenticating Azure Monitor as a data source. In your Azure Monitor portal, set your ServiceNow AI Platform instance as the rest endpoint using a standard webhook.

On the Azure portal, alert correlation rules are defined through the *Correlate alerts* setting within Alert Processing Rules. When *Correlate alerts* is assigned on the Azure portal, the Azure Monitor alerts received on the ServiceNow® instance within 60 minutes are grouped using tag based alert clustering.

Procedure

1. In the ServiceNow AI Platform, create a user account or identify an existing ServiceNow user for this integration.
   Also ensure the relevant ServiceNow sys_user is assigned the evt_mgmt_integration role. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.
2. In your Azure Monitor portal, update the Azure Monitor alert rules to use an action group with a standard webhook:

   a. Create an action group with *action type=webhook*, and provide the URL that includes the user ID and password of the ServiceNow user in the endpoint.
   For example:

   ```
   https://<username>:<password>@<instance-name>.service-now.com/
   api/sn_em_connector/em/inbound_event?source=azuremonitor
   ```

   > **Note:** If you have multiple tenants on the Azure portal and you want to use bidirectional functionality to acknowledge, close, or reopen an alert on the Azure Portal, then the URL for the webhook has to be in the following format: `https://<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=azuremonitor&event_class=<new_connector_instance_name>`. The <new_connector_instance_name> is the new Azure pull connector instance that you need to create for each tenant with other Azure service principal credentials. To know how to create the pull connector, see Configure Azure Monitor Bi-directional connector.

   b. In the Webhook section, make sure that **Yes** is selected for the **Enable the common alert schema** option.

Result

Alerts start flowing from the Azure Monitor into the Event Management plugin. The plugin extracts information from the original Azure Monitor alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to **All Events** to see the events.

What to do next

If you want to send alert state changes on the ServiceNow instance from the ServiceNow alerts to the Azure Portal, you need to enable the Azure Monitor Bi-directional connector. For more information, see Configure Azure Monitor Bi-directional connector.

Integrate Azure with REST API key token

Integrate using an API key to establish secure communication and automate data exchange via REST API. This simplifies integration, enabling seamless access to services and enhancing operational efficiency.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: evt_mgmt_integration

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API Key**.

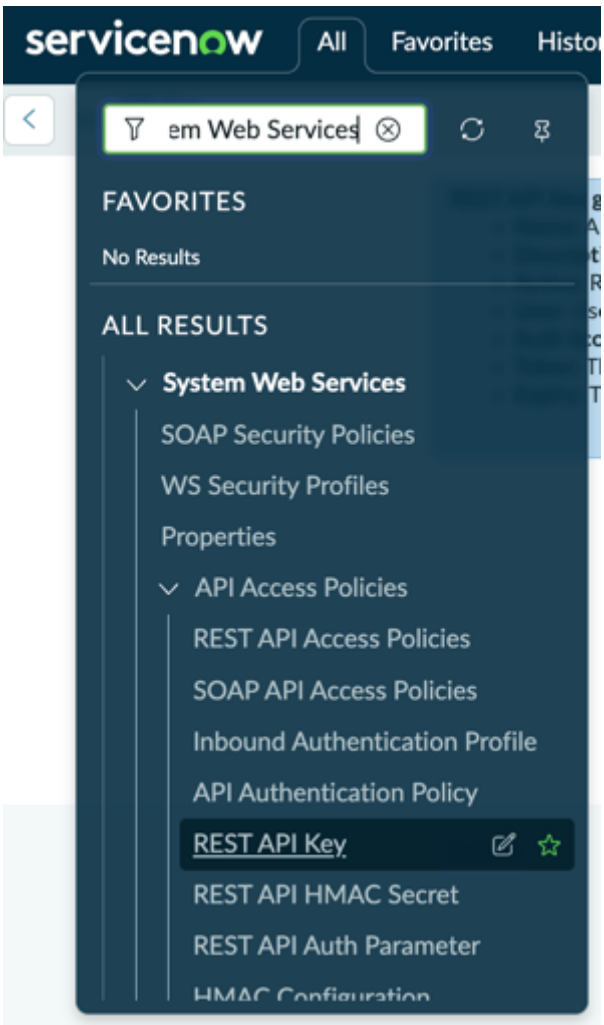

2. Select **New**.
3. On the form, fill the fields.

> **Note:** The user needs to be created with the evt_mgmt_integration role. Create a new one on the sys_user table to use here to facilitate the right access.

API Key

Field	Description
Name	Name to identify the REST API Key
Description	Description for the REST API Key.
Active	Status of the REST API Key.
User	User associated to the REST API Key. Use the look-up icon to select the user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Auth Scope	Option to add auth scope to manage the authority of the REST API Key.
Token	The REST API key generated by the ServiceNow AI Platform. Copy the key to use as part of the REST API call within the Header or Query parameter.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration.

4. Unlock Auth Scope and add: UserAccount.
5. Use the form menu and select **Save**.

   The system generates a token and saves it in the **Token** field. To see the token, use the lock icon and copy the contents display below the field. This is your query parameter (or header) value when your other system sends a REST API request to ServiceNow.

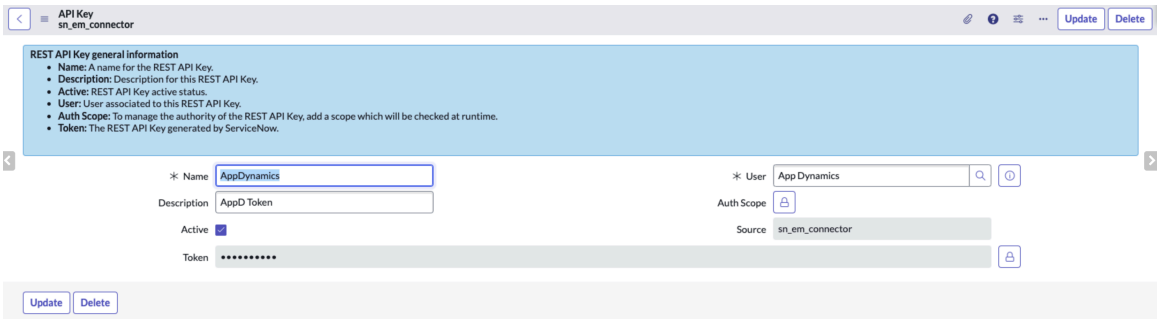

Integrate Azure Monitor with OAuth authentication

Integrate Microsoft Azure with Event Management by authenticating Azure V1 or V2 tokens in the Azure Monitor.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

On the Azure portal, alert correlation rules are defined through the *`Correlate alerts`* setting within Alert Processing Rules. When *`Correlate alerts`* is assigned on the Azure portal, the Azure Monitor alerts received on the ServiceNow® instance within 60 minutes are grouped using tag based alert clustering.

Roles required: evt_mgmt_admin, web_service_admin, and oauth_admin

About this task

Configure the Event Management environment for the collection of events from Azure Monitor. In your Azure Monitor portal, set your ServiceNow AI Platform instance as the rest endpoint using V1 or V2 tokens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Azure Monitor portal:

   a. Perform app registration and expose an API.
   For more information about registering an app and exposing an API in Azure, see https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-configure-app-expose-web-apis.

   b. Create an action group with a secure webhook, and provide the rest endpoint as `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=azuremonitor`.

   > **Note:** If you have multiple tenants on the Azure portal and you want to use bidirectional functionality to acknowledge, close, or reopen an alert on the Azure Portal, then the URL for the secure webhook has to be in the following format: `https://<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=azuremonitor&event_class=<new_connector_instance_name>`. The <new_connector_instance_name> is the new Azure pull connector instance that you need to create for each tenant with other Azure service principal credentials. To know how to create the pull connector, see Configure Azure Monitor Bi-directional connector.

   For servicenowservices.com instances, the rest endpoint has to be `https://<instance-name>.servicenowservices.com/api/sn_em_connector/em/inbound_event?source=azuremonitor&x-sn-apikey=<api-key>`.

   For more information about adding a secure webhook to an action group, see https://docs.microsoft.com/en-us/azure/azure-monitor/platform/action-groups.

   c. Navigate to **Alerts** > **Manage Alert Rules**.

   d. In the Secure Webhook section, make sure that **Yes** is selected for the **Enable the common alert schema** option.

   e. Add the action group with the secure webhook to an alert rule.

2. In the Azure Monitor portal, verify which Azure token is in use by your registered application.

   a. Navigate to **App Registration**, and select the registered application.

   b. In the Manage section, click **Manifest**.

   c. In the editor screen, select the **Microsoft Graph App Manifest** tab and locate the *requestedAccessTokenVersion* parameter or select the **AAD Graph App Manifest** tab and locate the *accessTokenAcceptedVersion* parameter.
   If the value of *requestedAccessTokenVersion* or *accessTokenAcceptedVersion* is **2**, the integration must use V2 tokens.

   If the value of *requestedAccessTokenVersion* or *accessTokenAcceptedVersion* is **1** or **null**, the integration must use V1 tokens.

3. In your ServiceNow AI Platform instance, ensure that the ServiceNow user is assigned with the correct Application (client) ID or Application ID URI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Also ensure the relevant ServiceNow sys_user is assigned the evt_mgmt_integration role.

a. Navigate to **System Security** > **Users and groups** > **Users**.

> **Note:** To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

b. Verify that the **Source** field for the ServiceNow sys_user is populated with the correct Application (client) ID or Application ID URI, as defined in the Azure Monitor portal.
If the application is using an Azure V1 token, the **Source** field must be populated with the Application ID URI of the registered application. If the application is using an Azure V2 token, the **Source** field must be populated with the Application (client) ID of the registered application.

If the **Source** field is not displayed, change the form layout to display this field. Click the context menu icon ( ) and select **Configure > Form Layout**. Move **Source** to the Selected list.

> **Note:** If your ServiceNow AI Platform instance version is earlier than Paris Patch 2 or Orlando Patch 9, the **Source** field value must be enclosed in square brackets. For example, `[api://azuretest]`.

4. In your ServiceNow AI Platform instance, navigate to **Event Management** > **Integrations** > **Azure OAuth Config**.
5. Click **Azure OAuth OIDC Entry** and enter the `Client ID`.
   If the registered application is using an Azure V2 token, the `Client ID` should be the same as the `Application (client) ID` of the app registrations defined in your Azure Monitor portal.

   If the registered application is using an Azure V1 token, the `Client ID` should be the same as the `Application ID URI` of the app registrations defined in your Azure Monitor portal (the exposed API).
6. In the **OAuth OIDC Provider Configuration** field, click the info icon ( ).
7. In the OIDC Provider Configuration window, click **Open Record**.

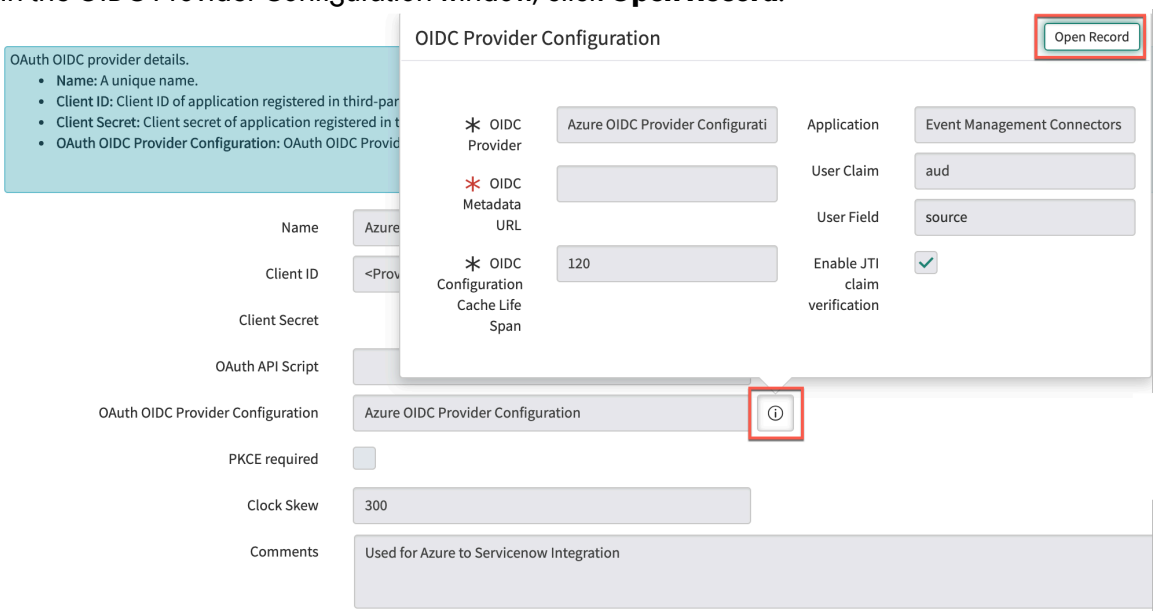

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Enter the `OIDC Metadata URL` according to the Azure token used by the registered application, as shown in the following table.

Azure token	OIDC Metadata URL
V2	In the OIDC Provider Configuration form, add the following URL to the **OIDC Metadata URL** field: `https://login.microsoftonline.com/<tenant-id>/v2.0/.well-known/openid-configuration` Ensure that `<tenant-id>` is replaced with the correct Azure Tenant ID. **Note:** If attempting to integrate Azure Gov Cloud with the ServiceNow Azure Monitor Connector, remember to change the claim value to `f1f34126-d4ef-40e1-ad4b-bf5d47b4860d`.
V1	In the OIDC Provider Configuration form: ○ In the Claim Name column, change `azp` to `appid`. 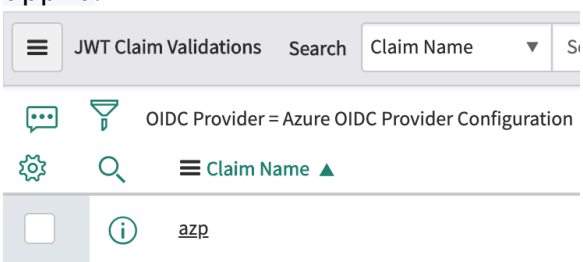 **Note:** If attempting to integrate Azure Gov Cloud with the Service Now Azure Monitor Connector, re member to change the claim value to `f1f34126-d4ef-40e1-ad4b-bf5d47b4860d`. ○ In the **OIDC Metadata URL** field, add the following URL: `https://login.microsoftonline.com/<tenant-id>/.well-known/openid-configuration` Ensure that `<tenant-id>` is replaced with the correct Azure Tenant ID.

Result

When an alert is created in Azure Monitor as part of the alert rule, the notification is sent to the ServiceNow AI Platform using the secure webhook endpoint. In your ServiceNow AI Platform instance, navigate to **All Events** to see the events. If you want to send alert state changes on the ServiceNow instance from the ServiceNow alerts to the Azure Portal, you need to enable the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Monitor Bi-directional connector. For more information, see Configure Azure Monitor Bi-directional connector .

Configure Azure Monitor Bi-directional connector

The Azure Monitor Pull connector sends information from ServiceNow Event Management to the Azure Portal. The pull connector sends the alert state changes from the ServiceNow environment to the Azure Portal.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

When an alert is triggered in the Azure Portal a Secure Webhook is created to send the alert notifications into ServiceNow. The Azure Monitor Connector will receive the alert notification as an event in the ServiceNow instance.

The Azure Monitor Bi-directional connector is supported in the following ServiceNow AI Platform versions:

- Quebec Patch 9 or later.
- Rome Patch 3 or later.
- San Diego.

Role required: evt_mgmt_admin

About this task

When you configure the Azure Monitor Bi-directional Connector, the bi-directional exchange of values to and from the external event source is enabled. When an alert is **Acknowledged**/**Closed**/**Reopened** on the ServiceNow instance, the state changes in the Azure Portal alert using the bi-directional connector.

These scenarios describe the default bi-directional functionality for the Azure Monitor connector:

- When an alert is closed in the Azure Portal, it is automatically closed in ServiceNow. However, it is updated irrespective of the bi-directional feature because closing the event received from the Azure Portal will close the alert.
- When an alert is manually closed in ServiceNow, all the associated Azure alerts will be closed in the Azure Portal.
- If the alert state is changed to **Reopen** in ServiceNow, all the associated Azure alerts will be opened in the Azure Portal.
- When an alert is **Acknowledged** in ServiceNow, all the associated Azure alerts will be acknowledged in the Azure Portal.

Note: For changing the alert state on the Azure Portal, Mid Server gets access token connecting to https://login.microsoftonline.com/ and to change the alert state on the Azure portal Mid Server calls changestate API on URL https://management.azure.com/ as described in https://learn.microsoft.com/en-us/rest/api/monitor/alertsmanagement/alerts/change-state?view=rest-monitor-alertsmanagement-2023-07-12-preview&tabs=HTTP .

Procedure

1. For a single tenant, perform the following steps:
   - Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
   - Select **Azure Monitor**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- In the **Name** field, verify that the name is same as the name in the Azure push connector from where ServiceNow receives Azure alerts.
- Select the **Active** check box.

  For this configuration, the **Host IP** field has no functional impact and is not used, so a dummy IP address (1.1.1.1) is provided.
- In the **Credential** field, add the Azure Service Principal credentials that has authorization to perform actions.

  To change the alert state, the service principal must have the Microsoft.AlertsManagement permissions and the Monitoring Contributor role.
- Select or add the MID Server to be used for this connector.

  If alerts on the MS Azure portal are not updating as expected, check the MID Server logs for errors.

2. For multiple tenants, perform the following steps:
   - Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
   - Select **New**.
   - In the **Name** field, enter the name of <new_connector_instance_name> that you have mentioned in the URL for secure webhook at Integrate Azure Monitor with OAuth authentication.
   - In the Host IP, enter any dummy IP address such as 1.1.1.1.
   - In the **Credential** field, add the Azure Service Principal credentials of that tenant that has authorization to perform actions.
   - Select the **Bi-directional** check box.
   - Select or add the MID Server to be used for this connector.
   - Select **Submit**.
   - Select the **Active** check box.
   - Select **Update**.

What to do next

Multiple alerts from the Azure Portal may be mapped to a single ServiceNow alert by de duplication. The mapping between the ServiceNow alert and Azure alerts will be maintained in the sn_em_connector_event_data table. The retention of records will depend on the alert's state:

- If the Azure Bi-directional Connector is not active, the records will be deleted in 2 days.
- When the Azure Bi-directional Connector is active:
  - If the alert is in a closed state, then it will retain the data for 7 days. This duration can be customized by the sn_em_connector.eventdata_closedevent_interval system property.
    - If the closed alert is re-opened in 7 days then it will reopen the corresponding closed alerts from the Azure Portal.
    - If the closed alert is re-opened after 7 days then it will not reopen the corresponding closed alerts from the Azure Portal.
  - If the alert is in an open state, then it will retain the data for 30 days. . This duration can be customized by the sn_em_connector.eventdata_openevent_interval.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To close or acknowledge an open ServiceNow alert within 30 days, you can close or acknowledge all corresponding Azure alerts on the Azure Portal.
- If you want to close or acknowledge an open ServiceNow alert after 30 days, you cannot close or acknowledge the corresponding Azure alerts on the Azure Portal.

Event collection from Microsoft Azure Monitor

The MID WebServer Event Collector enables you to collect JSON formatted event messages sent from the Microsoft Azure portal.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Ensure that you:

- Deploy and start the MID Server. See MID Server configuration .
- Configure and start the MID Web Server. See Configure the MID Web Server extension.
- Configure and start the MID WebService Event Collector. See Configure the MID WebService Event Collector Context.

Role required: evt_mgmt_admin

About this task

On the Azure portal, alert correlation rules are defined through the *`Correlate alerts`* setting within Alert Processing Rules. When *`Correlate alerts`* is assigned on the Azure portal, the Azure Monitor alerts received on the ServiceNow® instance within 60 minutes are grouped using tag based alert clustering.

JSON formatted event messages are sent from Microsoft Azure. The MID Server transforms the collected event messages by parsing them using the TransformEvents_MidAzureMonitor script include, located here: **Event Management** > **Integrations** > **PushConnectors**. In the Push Connectors page, click **Azure Monitor Mid Push Connector**.

The default format of the URL to push event messages from Microsoft Azure to the MID Server is `http://<MID_Web_Server_User>:<MID_Web_Server_Password>@<MID_Server_IP>:<MID_Web_Serv api/mid/em/inbound_event?Transform=TransformEvents_MidAzureMonitor`.

Variables in the default URL

Variable	Description
MID_Server_IP	IP address of the MID Web Server Extension.
MID_Web_Server_Port	Listening port of the MID Web Server Extension.
MID_Web_Server_User	Username for the MID Web Server Extension.
MID_Web_Server_Password	Password of the user of the MID Web Server.

The following procedure describes the collection of JSON formatted event messages using basic authentication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Azure Monitor Portal, verify that the MID Server IP is accessible from the Azure portal and the MID rest point is accessible from Microsoft Azure.
   Alternatively, install the MID Server on an Azure VM. Provide the Azure VM public address in the URL and allow the mid_web_Server_port in **Azure VM** > **Networking Settings** > **Inbound port rules**.
2. Create an action group with a webhook and provide the rest endpoint as `http://<MID_Web_Server_User>:<MID_Web_Server_Password>@<MID_Server_IP>:<MID_Web_Se api/mid/em/inbound_event?Transform=TransformEvents_MidAzureMonitor`

   Note: For more information about adding a webhook to an action group, see Create and manage action groups in the Azure portal on the Microsoft documentation site.
3. In the Webhook section, make sure **Yes** is selected for the **Enable the common alert schema** option.
4. Navigate to **Alerts** > **Manage Alert Rules**.
5. Add the action group with the webhook to an alert rule.

What to do next

If you want to send alert state changes on the ServiceNow instance from the ServiceNow alerts to the Azure Portal, you need to enable the Azure Monitor Bi-directional connector. For more information, see Configure Azure Monitor Bi-directional connector.

Event collection from BMC TrueSight and BMC TrueSight_v2

The MID WebService Event Collector enables you to collect JSON formatted event messages sent from BMC TrueSight Operations Management (TrueSight), previously known as BMC ProactiveNet Performance Management (BPPM). It uses event stream notification capabilities.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Ensure that you:

- Deploy and start the MID Server. See MID Server configuration .
- Configure and start the MID Web Server. See Configure the MID Web Server extension.
- Configure and start the MID WebService Event Collector. See Configure the MID WebService Event Collector Context.

Role required: evt_mgmt_admin

About this task

JSON formatted event messages are sent from BMC TrueSight. The MID Server transforms the collected event messages by parsing them using the TransformEvents_bmcTrueSight script include, located here: **Event Management** > **Integrations** > **Listener Transform Scripts**. In the Listener Transform Scripts page, click **TransformEvents_bmcTrueSight**.

The default format of the URL to push event messages from BMC TrueSight to the MID Server is `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=TransformEvents_bmcTrueSight`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The default format of the URL to push event messages from BMC TrueSight_v2 to the MID Server is `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=TransformEvents_bmcTrueSight_v2`.

Variables in the default URL

Variable	Description
MID_Server_IP	IP address of the MID Web Server Extension.
MID_Web_Server_Port	Listening port of the MID Web Server Extension.
MID_Web_Server_User	Username for the MID Web Server Extension.
MID_Web_Server_Password	Password of the user of the MID Web Server Extension. However, encode the password in BASE64. For example, if the MID Web Server Extension password is `admin`, in this case use `YWRtaW4=` as the password for MID_Web_Server_Password. To determine the BASE64 value, you can use the utility at this website: https://www.base64encode.org/

The following procedure describes the collection of JSON formatted event messages using basic authentication. For more information about supported authentication methods, see Configure the MID Web Server extension.

Procedure

1. In BMC TrueSight, set the MID WebService Event Collector as the target for BMC TrueSight notification, using these parameters:

   Variables in the generic BMC URL

Variable	Description
serverHost	Name of the BMC TrueSight server.
port	Number of the BMC TrueSight port.

2. In the message body and generic BMC URL, replace the parameters with the values in the preceding table and connect to BMC TrueSight.
   Use the BMC Publish-Subscribe REST approach to create a stream event subscription and send event messages in JSON format. The generic format of the BMC URL to create a stream event subscription is: `http|https://<serverHost>:<port>/bppmws/api/stream/subscribe?dataType=event`. For example: `http://TrueSight1server:40/bppmws/api/stream/subscribe?dataType=event`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

Example of mapping event messages to event fields in the Event (em_event) table

Event field	Event message
node	mc_host
type	mc_object_class
metric_name	mc_parameter
resource	mc_object
severity	severity
time_of_event	date_reception
description	msg

Example:
Example, showing the message body.

```
{
"action": "create",
"data": {
"mc_abstracted": [],
"itsm_operational_category3": "",
"mc_object_class": "CPU",
"source": "MySource",
"adapter_host": "",
"itsm_operational_category2": "",
"itsm_location": "",
"itsm_operational_category1": "",
"mc_abstraction": [],
"mc_tool_id": "",
"mc_original_priority": "PRIORITY_1",
"event_handle": "3172",
"mc_modhist": [],
"mc_tool_class": "",
"mc_origin": "",
"mc_tool_uri": "",
"mc_service": "",
"mc_smc_type": "",
"mc_object_owner": "",
"mc_object": "CPU_0",
"mc_notes": [],
"itsm_model_version": "",
"mc_origin_key": "",
"mc_tool_address": "",
"mc_ueid": "mc.pncell_bmcswetsom96.1829b8ed.0",
"mc_bad_slot_names": [],
"date": "20161114081525.000000-300",
"itsm_item": "",
"mc_timeout": "0",
"mc_priority": "PRIORITY_1",
"itsm_type": "",
"itsm_manufacturer": "",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
"mc_tool_key": "",
"itsm_product_name": "",
"mc_smc_impact": "NOT_ELECTED",
"mc_parameter": "MEMFreeMem",
"mc_bad_slot_values": [],
"itsm_incident_id": "",
"mc_incident_report_time": "0",
"mc_event_model_version": "1.1.00",
"mc_relationships": "0/0",
"CLASS": "EVENT",
"mc_using_organization": "",
"mc_tool_sev": "",
"mc_parameter_threshold": "",
"mc_smc_effects": [],
"mc_smc_priority": "0",
"itsm_company": "",
"date_reception": "1479129325",
"mc_smc_alias": "bmcswetsom96_2",
"pn_detail_diag_count": "0",
"mc_effects": [],
"mc_object_uri": "",
"mc_date_modification": "1479129325",
"mc_incident_time": "0",
"status": "OPEN",
"msg": "Free Memory Low",
"pn_detail_diag": "0",
"mc_host": "bmcswetsom97",
"repeat_count": "69",
"mc_host_address": "",
"mc_tool_suggestion": "",
"mc_using_organization_id": "",
"mc_collectors": ["1.1", "3.1.1", "4.1", "5.1"],
"mc_account": "",
"mc_long_msg": "",
"mc_owner": "",
"mc_history": ["pncell_bmcswetsom96:3172"],
"itsm_category": "",
"mc_host_class": "",
"mc_tool_rule": "",
"mc_local_reception_time": "1479129325",
"mc_associations": [],
"mc_relation_source": "",
"mc_propagations": [],
"mc_client_address": "172.16.173.131",
"itsm_incident_status": "",
"mc_tool_time": "0",
"mc_origin_class": "",
"mc_tool": "",
"mc_action_count": "0",
"mc_cause": "0",
"mc_origin_sev": "",
"mc_host_id": "0",
"mc_event_relations": [],
"mc_operations": [],
"mc_event_category": "CONFIGURATION_MANAGEMENT",
"pn_device_name": "",
"mc_arrival_time": "1479129325",
```


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
"severity": "OK",
"mc_notification_history": [],
"mc_smc_causes": [],
"mc_acl": [],
"mc_original_severity": "CRITICAL",
"duration": "0",
"mc_location": "Unknown",
"mc_smc_id": "",
"mc_parameter_value": "20",
"administrator": "",
"mc_event_subcategory": "OTHER",
"mc_parameter_unit": ""
},
"subscriptionId": "af18f6ba-a2d3-4574-81c8-7f366ca70ebd",
"type": "EVENT"
}
```


Metric collection from OpenTelemetry (Otel) metrics

The MID WebService metric Collector enables you to collect JSON and protobuf formatted metrics sent from OpenTelemetry (Otel).

Before you begin

Ensure that you have installed the Event Management Connectors (sn_em_connector) plugin on the ServiceNow AI Platform instance and the ServiceNow version is Australia or above.

Ensure that you:

- Deploy and start the MID Server. See MID Server configuration.
- Configure and start the MID Web Server. See Configure the MID Web Server extension.
- To activate metric collection, ensure that the MID Server that retrieves metrics is configured with the Metric Intelligence extension and that the extension is in Started mode. See Manually configure the Metric Intelligence extension.
- Enable the rest listener.
- Install the Discovery and Service Mapping Patterns plugin if you want Kubernetes CI tables.

Role required: evt_mgmt_admin

About this task

Protobuf and JSON formatted metrics are sent from Opentelemetry. The MID Server transforms the collected metrics by parsing them using the *`TransformMetrics_MidOpenTelemetry`* script include, located at **Event Management** > **Integrations** > **PushConnectors**. In the Push Connectors page, select **OpenTelemetry**.

The default format of the URL to push event messages from the Otel to the MID Server is: `http://{MID_Server_IP}:{MID_Web_Server_Port}/api/mid/sa/inbound_metrics?transform=TransformMetrics_MidOpenTelemetry`

Variables in the default URL

Variable	Description
MID_Server_IP	IP address of the MID Web Server Extension.
MID_Web_Server_Port	Listening port of the MID Web Server Extension.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Variables in the default URL (continued)

Variable	Description
MID_Web_Server_User	Username for the MID Web Server Extension.
MID_Web_Server_Password	Password of the user of the MID Web Server.

Procedure

1. CONFIGURE RECEIVERS: HOSTMETRICS RECEIVER - SYSTEM METRICS COLLECTION

Add the following hostmetrics receiver under the "receivers:" section to collect system metrics:

```
receivers:
  hostmetrics:
    collection_interval: 10s
    scrapers:
      cpu:
      memory:
      disk:
      filesystem:
      network:
      processes:
      paging:
```

Note:

- Collection_interval set to 10 seconds for frequent metric updates.
- Adjust the interval based on your monitoring requirements (higher values reduce overhead).
- Enable only the scrapers you need to minimize resource usage.

2. CONFIGURE PROCESSORS
TRANSFORM PROCESSOR - REMOVE UNNECESSARY FIELDS

a. Add the following transform processor under the "processors:" section to remove unnecessary fields from the payload:

```
transform:
  metric_statements:
    - context: scope
      statements:
        - set(name, "")
        - set(version, "")
```

b. MEMORY LIMITER PROCESSOR - PREVENT MEMORY ISSUES

Add the following memory_limiter processor under the "processors:" section to prevent out-of-memory errors:

```
memory_limiter:
  limit_mib: 5
  spike_limit_mib: 1
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- limit_mib set to 5 MB of memory usage
- spike_limit_mib adjusted to 1 MB to handle temporary spikes
- Adjust these values based on your system's available memory

c. FILTER PROCESSOR - SELECT SPECIFIC METRICS

Add the following filter processor under the "processors:" section to include only selected metrics:

```
filter/minimal:
  metrics:
    include:
      match_type: strict
      metric_names:
        # CPU Metrics
        - system.cpu.time
        - system.cpu.load_average.1m
        - system.cpu.load_average.5m
        - system.cpu.load_average.15m
        # Memory Metrics
        - system.memory.usage
        - system.paging.usage
        # Disk Metrics
        - system.filesystem.usage
        - system.disk.io
        - system.disk.io_time
        # Process Metrics
        - system.processes.count
        - system.processes.created
        # Paging Metrics
        - system.paging.operations
        # Network Metrics
        - system.network.io
        - system.network.packets
        - system.network.errors
        - system.network.dropped
        - system.network.connections
        # MySQL Metrics
        - mysql.operations
        - mysql.locks
        - mysql.uptime
        - mysql.threads
        # HTTP Check Metrics
        - httpcheck.duration
        - httpcheck.status
        - httpcheck.error
        # Prometheus node_exporter Metrics
        - node_cpu_seconds_total
        - node_memory_MemTotal_bytes
        - node_memory_MemAvailable_bytes
        - node_disk_read_bytes_total
        - node_disk_written_bytes_total
        - node_disk_io_time_seconds_total
        - node_network_receive_bytes_total
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
        - node_network_transmit_bytes_total
        # K8s Cluster Metrics
        - k8s.node.condition_ready
        - k8s.node.allocatable_cpu
        - k8s.node.allocatable_memory
        - k8s.deployment.available
        - k8s.deployment.desired
        # Kubelet Stats Metrics
        - k8s.node.cpu.usage
        - k8s.node.memory.working_set
        - k8s.pod.cpu.usage
        - k8s.pod.phase
        - k8s.pod.memory.working_set
        - k8s.container.cpu.usage
        - k8s.container.memory.working_set
        - k8s.container.restarts
```

Note: Histogram metrics are not supported by ServiceNow MID Server.

Even if histogram metrics are sent to the MID Server, the ServiceNow transform script will automatically ignore them. However, as a best practice, it is recommended to filter out histogram metrics at the collector level to optimize network bandwidth and reduce unnecessary data transmission.

To exclude histogram metrics, add the following filter:

```
filter/drop_histograms:
    metrics:
      metric:
        - 'type == METRIC_DATA_TYPE_HISTOGRAM'
```

3. CONFIGURE EXPORTERS

Choose one of the following authentication methods based on your requirements:

a. OPTION A: API KEY AUTHENTICATION

Add the following exporter under the "exporters:" section:

```
otlphttp/opentelemetry:
  metrics_endpoint:
 "http://{mid_server_ip}:{port}/api/mid/sa/inbound_metrics?tr
ansform=TransformMetrics_MidOpenTelemetry"
  headers:
    Authorization: "Key {api_key}"
  timeout: 30s
  compression: gzip
  tls:
    insecure: true
  sending_queue:
    enabled: true
    num_consumers: 4
    queue_size: 100
  retry_on_failure:
    enabled: true
    initial_interval: 5s
    max_interval: 30s
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Replace {mid_server_ip}, {port}, and {api_key} with your actual values.

b. OPTION B: BASIC AUTHENTICATION

For Basic Authentication, you need to configure both an extension and an exporter.

i. Configure Basic Auth Extension

Add the following under the "extensions:" section:

```
extensions:
  basicauth/mid:
    client_auth:
      username: "otel"
      password: "Otel@123"
```

Note: Replace username and password with your actual credentials.

ii. Configure Basic Auth Exporter

Add the following under the "exporters:" section:

```
otlphttp/opentelemetryBasic:
  metrics_endpoint:
 "http://{mid_server_ip}:{port}/api/mid/sa/inbound_metrics?t
ransform=TransformMetrics_MidOpenTelemetry"
  auth:
    authenticator: basicauth/mid
  timeout: 30s
  compression: gzip
  tls:
    insecure: true
  sending_queue:
    enabled: true
    num_consumers: 4
    queue_size: 100
  retry_on_failure:
    enabled: true
    initial_interval: 5s
    max_interval: 30s
```

Note: Replace {mid_server_ip} and {port} with your actual values.

4. CONFIGURE SERVICE PIPELINE

Add the configured processors and exporters to the metrics pipeline under the "service:" section:

```
service:
  pipelines:
    metrics:
      receivers: [<your_configured_receivers>]
      processors: [memory_limiter, transform, filter/minimal]
      exporters: [otlphttp/opentelemetry,
 otlphttp/opentelemetryBasic]
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- Replace <your_configured_receivers> with your actual receiver names (e.g., hostmetrics).
- Include only the exporters you have configured (API Key or Basic Auth, or both).
- Processor order matters: memory_limiter should be first, followed by transform and filter.
- Ensure the processor names match those defined in Step 2.
- If using histogram filtering, add filter/drop_histograms after filter/minimal.

5. ADDITIONAL CONFIGURATION OPTIONS
By default, OpenTelemetry sends data in Protocol Buffers (protobuf) format.

To send data in JSON format instead, add the following header to your exporter configuration:

```
headers:
  Content-Type: "application/json"
```

COMPLETE CONFIGURATION EXAMPLE.

Below is a complete example configuration file combining all components:

```
receivers:
  hostmetrics:
    collection_interval: 10s
    scrapers:
      cpu:
      memory:
      disk:
      filesystem:
      network:
      processes:
      paging:

processors:
  memory_limiter:
    limit_mib: 5
    spike_limit_mib: 1

  transform:
    metric_statements:
      - context: scope
        statements:
          - set(name, "")
          - set(version, "")

  filter/minimal:
    metrics:
      include:
        match_type: strict
        metric_names:
          - system.cpu.time
          - system.cpu.load_average.1m
          - system.memory.usage
          - system.network.io
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
  filter/drop_histograms:
    metrics:
      include:
        metric:
          - 'type != METRIC_DATA_TYPE_HISTOGRAM'

extensions:
  basicauth/mid:
    client_auth:
      username: "otel"
      password: "Otel@123"

exporters:
  otlphttp/opentelemetry:
    metrics_endpoint:
 "https://10.0.0.100:8443/api/mid/sa/inbound_metrics?transform=
TransformMetrics_MidOpenTelemetry"
    headers:
      Authorization: "Key your_api_key_here"
    timeout: 30s
    compression: gzip
    tls:
      insecure: false
    sending_queue:
      enabled: true
      num_consumers: 4
      queue_size: 100
    retry_on_failure:
      enabled: true
      initial_interval: 5s
      max_interval: 30s

service:
  extensions: [basicauth/mid]
  pipelines:
    metrics:
      receivers: [hostmetrics]
      processors: [memory_limiter, transform, filter/minimal,
 filter/drop_histograms]
      exporters: [otlphttp/opentelemetry]
```

Integrate Dynatrace platform events

Integrate Dynatrace with Event Management by adding Dynatrace as an authenticated data source.

What authentication is used

There are two methods of authentication:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- OAuth authentication: Provides enterprise-grade authentication to keep your enterprise environment safe. Authentication is performed using Dynatrace access tokens. For more information, see Integrate Dynatrace with OAuth authentication.
- Basic webhook authentication: Provides a basic standard of authentication, without the need for Dynatrace Active Directory. This authentication can be especially useful for distributed small teams, such as SRE or DevOps teams. For more information, see Integrate Dynatrace with basic authentication.

Note: When a problem is reopened in Dynatrace, the original start time is retained and not updated to the time of reopening. This behavior is intentional and reflects how Dynatrace manages the problem lifecycle. Problems reopened within the defined reopening period (typically 30 minutes) are treated as a continuation of the original issue. The start time continues to reflect the timestamp of the earliest triggering event, even after reopening.

Do not use the evidence details timestamp as a substitute for the problem start time. The evidence timestamp refers to individual contributing events within the problem. It is part of the Root Cause Analysis (RCA) and does not indicate the start or duration of the overall problem.

Starting from the Xanadu release, the base system rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the connector.

Note: This feature applies only to active event rules.

Integrate Dynatrace with basic authentication

Integrate Dynatrace with Event Management by adding a standard webhook in the Dynatrace console.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Ensure that configuration items for the hosts managed by Dynatrace exist in the ServiceNow AI Platform instance. These CIs can be physical or virtual and can be either manually created or discovered via IP discovery or Cloud Discovery.

Confirm you have created a user with an **Identify Type** of **Machine** and the evt_mgmt_integration role.

Roles required: evt_mgmt_integration and web_service_admin

About this task

Configure the Event Management environment for the collection of events from Dynatrace by authenticating Dynatrace as a data source. In your Dynatrace console, set your ServiceNow AI Platform instance as the rest endpoint using a standard webhook.

Procedure

1. In your Dynatrace console, define Host Naming rules:
   a. Navigate to **Settings** > **Process and contextualize** > **Naming** > **Host Naming**.
   b. Define the Host Naming rules for each cloud provider (Azure/AWS/GCP) to uniquely identify a CI from ServiceNow.
   This ensures the ServiceNow CIs can be uniquely identified from the payload received from Dynatrace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You do not need to create a Host Naming rule for VMware machines because Dynatrace manages them as physical servers.

2. Define anomaly detection rules:
   a. Navigate to **Settings** > **Analyze and alert** > **Alerts** > **Hosts**.
   b. In the **Hosts** tab, define rules on when to create alerts on the managed hosts.
3. Define the integration settings:
   a. Navigate to **Settings** > **Analyze and alert** > **Notifications** > **Problem notifications**.
   b. In the Set up custom integration form, add the Webhook URL: `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=dynatrace`

      Note:

      For basic authentication, select **Create basic authorization header**. You can see the username and password fields for the user.

   c. Enter the user name and password of the relevant user.

      Note: Ensure the evt_mgmt_integration role is assigned to the selected user. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

   d. In the Custom payload section, add the following payload structure for the events that will be generated.

      Ensure that ImpactedEntities and ProblemDetailsJSONv2 are passed as JSON objects, not strings.

```
{
  "ImpactedEntities": {ImpactedEntities},
  "ImpactedEntity": "{ImpactedEntity}",
  "PID": "{PID}",
  "ProblemDetailsHTML": "{ProblemDetailsHTML}",
  "ProblemDetailsJSONv2": {ProblemDetailsJSONv2},
  "ProblemDetailsMarkdown": "{ProblemDetailsMarkdown}",
  "ProblemDetailsText": "{ProblemDetailsText}",
  "ProblemID": "{ProblemID}",
  "ProblemImpact": "{ProblemImpact}",
  "ProblemSeverity": "{ProblemSeverity}",
  "ProblemTitle": "{ProblemTitle}",
  "ProblemURL": "{ProblemURL}",
  "State": "{State}",
  "Tags": "{Tags}"
}
```

4. Define the integration settings for the Grail problem event, which ServiceNow® now supports through its new payload.
   a. Log in to the Dynatrace instance and, from the left-hand panel, select **Workflow**.
   b. On the Workflow page, select **+ Workflow** to create a new workflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Create a blank workflow.

d. In the Select a trigger section, select **Problem trigger**.

e. In the middle workflow section, on the problem trigger tile, add the **Send ITOM event** task.

f. Select **+ Create a new connection**.

g. In the **Connection name** field, enter a name.

h. In the **ServiceNow ITOM Event API URL** field, paste the URL in this format:`https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=dynatrace`.

i. In the **Type** field, select **Basic Authentication**, and then select **Create basic authorization header**. The username and password fields appear.

j. Enter the user name and password of the relevant user.

> **Note:** Ensure the `evt_mgmt_integration` role is assigned to the selected user. For proper authentication, use the least-privileged user with the `evt_mgmt_integration` role rather than a highly privileged user.

k. Save the workflow and select **Run** to execute the workflow.

Result

Alerts start flowing from the Dynatrace console into the Event Management plugin. The plugin extracts information from the original Dynatrace alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to **All Events** to see the events.

> **Note:** By default, host binding is enabled for Dynatrace events for all providers (Azure/AWS/Google). If all hosts in the environment are discovered using Cloud Discovery by providing credentials and discovered resources are in the cmdb_ci_vm_object list, then the VM binding may not occur. To resolve this, you must enable the **Dynatrace - General** event rule. For further information about Event rules, see Event rules.

Integrate Dynatrace with OAuth authentication

Integrate Dynatrace with Event Management by authenticating Dynatrace V1 or V2 tokens in the Dynatrace Monitor.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

The Event Management plugin must be installed on the ServiceNow AI Platform instance.

Verify Configuration Items for the hosts managed by Dynatrace exist in ServiceNow. These CIs can be physical or virtual, and can be either manually created or discovered using IP discovery or Cloud discovery.

Role required: evt_mgmt_integration

About this task

Dynatrace integration with OAuth authentication is available starting from the Washington DC platform release or later.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Configuration on the ServiceNow instance.

   a. Create *glide.oauth.inbound.client.credential.grant_type.enabled* system property to use Client Credentials grant type for OAuth inbound integrations.

      - Type: true|false
      - Value: true

      Note: An administrator must create this system property because the evt_mgmt_admin role lacks permission to do so.

   b. Create an OAuth Entity endpoint for clients to access the instance.

      i. Navigate to **All** > **System OAuth** > **Application Registry** and then select **New**.

      ii. On the interceptor page, select **Create an OAuth API endpoint for external clients** and then fill in the form.

      iii. To create an OAuth client record, select **Submit**.

   c. Add the OAuth Application User and the Default Grant Type to the created OAuth client record.

      i. On the OAuth Entity form, add the **OAuth Application User** and the **Default Grant Type** fields to use the Client Credentials grant type for OAuth inbound integrations.

      ii. Open the created OAuth client record.

      iii. From the page header, select the **More** options icon.

      iv. Select **Configure** > **Form Layout**

      v. On the Form Layout page, add **OAuth Application User and Default Grant Type** from the list of fields.

      vi. **Save** or **Update** the form.

      vii. For the **OAuth Application User** field, select the user.

      The user must have the evt_mgmt_integration role that will be used to create the events sent by Dynatrace.

      Note: To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

      viii. Set **Default Grant Type** field value to **Client Credentials**.

2. In the Dynatrace console, define the integration settings.

   a. Navigate to **Settings** > **Integration** > **Problem notifications** > **Set up notifications**.

   b. In the Set up custom integration form, webhook URL.

   ```
   https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=dynatrace
   ```

   c. Enable **Use OAuth 2.0 for authentication**.

   d. In the OAuth 2.0 credentials section, perform the following steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

i. Provide Access token URL as `https://<instance_name>.service-now.com/oauth_token.do`.

ii. Provide the value for **Client ID** with the client id of the OAuth Entity Client created in ServiceNow instance.

iii. Provide “Client secret” with the client secret of the OAuth Entity Client created in ServiceNow instance.

iv. Select **Save**.

e. In the Custom payload section, add in the following payload structure for the events that will be generated. Ensure that `ImpactedEntities` and `ProblemDetailsJSONv2` are passed as JSON objects, not strings.

```
{
  "ImpactedEntities": {ImpactedEntities},
  "ImpactedEntity": "{ImpactedEntity}",
  "PID": "{PID}",
  "ProblemDetailsHTML": "{ProblemDetailsHTML}",
  "ProblemDetailsJSONv2": {ProblemDetailsJSONv2},
  "ProblemDetailsMarkdown": "{ProblemDetailsMarkdown}",
  "ProblemDetailsText": "{ProblemDetailsText}",
  "ProblemID": "{ProblemID}",
  "ProblemImpact": "{ProblemImpact}",
  "ProblemSeverity": "{ProblemSeverity}",
  "ProblemTitle": "{ProblemTitle}",
  "ProblemURL": "{ProblemURL}",
  "State": "{State}",
  "Tags": "{Tags}"
}
```

Result

Alerts start flowing from Dynatrace agents into the Event Management plugin. The plugin extracts information from the original Dynatrace alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to All Events to see the events.

Note: By default, host binding is enabled for Dynatrace events for (AWS/Azure/GCP) providers. If all hosts in the environment are discovered using Cloud Discovery by providing credentials and discovered resources are in the cmdb_ci_vm_object list, then the VM binding may not occur. To resolve this, you must enable the Dynatrace - General event rule. For further information about Event rules, see Event rules.

Integrate New Relic platform events

Integrate New Relic with Event Management by adding a standard webhook in the New Relic old and new consoles.

Before you begin

Verify that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_integration

About this task

Configure the Event Management environment for the collection of events from New Relic by authenticating New Relic as a data source. Set your ServiceNow AI Platform instance as the rest

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

endpoint using a standard webhook. Use procedure step 1 for the new console. Use procedure step 2 for the old console.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

Procedure

1. For the new console:

   a. Create a destination.

      i. Navigate to **Alerts & AI** > **Destinations**.

      ii. In the **Add destinations** section, select the **Webhook** channel type.

      iii. In the **Webhook name** field, enter the webhook name.

      iv. In the **Endpoint URL** field enter: `https://<InstanceName>.service-now.com/api/sn_em_connector/em/inbound_event?source=newrelic`

      v. In the **Basic Auth** section, enter your user name and password.

      Note: Verify the evt_mgmt_integration role is assigned to the selected user. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

      vi. Select **Save destination**.

   b. Create a policy.

      i. Navigate to **Alerts & AI** > **Alert Conditions & Policies**.

      ii. Select **New Alert policy** and define a policy.

      iii. Select **Set up Notifications**.

      iv. In the **Add Channel** section select **Webhook**, and then select the destination that you created in the previous substep.

      v. In the **Payload template** section, paste the following JSON payload.

```
{ "account_id": {{nrAccountId}}, "account_name": {{json
 accumulations.tag.account.[0]}}, "closed_violations_count":
 {{closedIncidentsCount}}, "open_violations_count":
 {{openIncidentsCount}}, "condition_description":
 "{{escape accumulations.conditionDescription.[0]}}",
 "condition_family_id":
 {{accumulations.conditionFamilyId.[0]}}, "condition_name":
 "{{escape accumulations.conditionName.[0]}}",
 "current_state": {{json state}}, "details":
 {{json issueTitle}}, "duration": {{#if
 issueDurationMs}}{{issueDurationMs}}{{else}}0{{/if}},
 "event_type": "INCIDENT", "incident_acknowledge_url":
 {{json issueAckUrl}}, "incident_id": {{json issueId}},
 "incident_url": {{json issuePageUrl}}, "metadata":
 { {{#if locationStatusesObject}}"location_statuses":
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
{{json locationStatusesObject}},{{/if}} {{#if
accumulations.metadata_entity_type}}"entity.type":
{{json accumulations.metadata_entity_type.[0]}},{{/if}}
{{#if accumulations.metadata_entity_name}}"entity.name":
{{json accumulations.metadata_entity_name.[0]}},{{/if}}
"section": "metadata" }, "owner": {{json owner}},
"policy_name": {{json accumulations.policyName.[0]}},
"policy_url": {{json policyUrl}}, "runbook_url": {{json
accumulations.runbookUrl.[0]}}, "severity": {{json
priority}}, "targets": [ { "id": "{{labels.targetId.[0]}}",
"name": "{{#if accumulations.targetName}}{{escape
accumulations.targetName.[0]}}{{else
if entitiesData.entities}}{{escape
entitiesData.entities.[0].name}}{{else}}N/A{{/if}}",
"link": "{{issuePageUrl}}", "product":
"{{accumulations.conditionProduct.[0]}}", "type": "{{#if
entitiesData.types.[0]}}{{entitiesData.types.[0]}}{{else}
}N/A{{/if}}", "labels": { {{#each
accumulations.rawTag}}"{{escape @key}}": {{#if
this.[0]}}{{json this.[0]}}{{else}}"empty"{{/if}}{{#unless
@last}},{{/unless}}{{/each}} } } ], "timestamp":
{{updatedAt}}, "violation_callback_url": {{json
issuePageUrl}}, "violation_chart_url": {{json
violationChartUrl}} }
```

vi. Select **Send test notification** to get the test event in your ServiceNow instance.

vii. Select **Save message**.

viii. Select **Activate workflow** to activate the workflow.

c. Create alert conditions to apply to the policy:

i. Navigate to **Alerts & AI** > **Alert Conditions & Policies**.

ii. Select **Create alert condition**, define the required alert details, and add it to the policy you created in the previous step.

iii. Select **Save Condition**.

2. For the old console.

a. Create a notification channel.

i. Navigate to **Alerts & AI** > **Notification channels**.

ii. In the **Channel details** tab, select the **Webhook** channel type.

iii. In the **Base Url** text box, enter: `https://<InstanceName>.service-now.com/api/sn_em_connector/em/inbound_event?source=newrelic`

iv. In the Basic Auth section, enter the user name and password of the relevant user.

> **Note:** Verify the evt_mgmt_integration role is assigned to the selected user. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

v. In the **Custom Payload** section, select **Add Custom Payload**, and select the **JSON** payload type.

vi. Paste the following JSON payload into the **Custom Payload**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
{ "account_id": "$ACCOUNT_ID", "account_name":
 "$ACCOUNT_NAME", "closed_violations_count_critical":
 "$CLOSED_VIOLATIONS_COUNT_CRITICAL",
 "closed_violations_count_warning":
 "$CLOSED_VIOLATIONS_COUNT_WARNING",
 "condition_description": "$DESCRIPTION",
 "condition_family_id": "$CONDITION_FAMILY_ID",
 "condition_name": "$CONDITION_NAME", "current_state":
 "$EVENT_STATE", "details": "$EVENT_DETAILS",
 "duration": "$DURATION", "event_type":
 "$EVENT_TYPE", "incident_acknowledge_url":
 "$INCIDENT_ACKNOWLEDGE_URL", "incident_id": "$INCIDENT_ID",
 "incident_url": "$INCIDENT_URL", "metadata":
 "$METADATA", "open_violations_count_critical":
 "$OPEN_VIOLATIONS_COUNT_CRITICAL",
 "open_violations_count_warning":
 "$OPEN_VIOLATIONS_COUNT_WARNING", "owner": "$EVENT_OWNER",
 "policy_name": "$POLICY_NAME", "policy_url": "$POLICY_URL",
 "runbook_url": "$RUNBOOK_URL", "severity": "$SEVERITY",
 "targets": "$TARGETS", "timestamp": "$TIMESTAMP",
 "violation_callback_url": "$VIOLATION_CALLBACK_URL",
 "violation_chart_url": "$VIOLATION_CHART_URL" }
```

b. Create a policy.

i. Navigate to**Alerts & AI** > **Policies**.

ii. Select**New Alert policy**and define a policy.

iii. In the**Notification Channels**section, select**Add**, select**Webhook**, and then select the channel you created in the previous step.

c. Create alert conditions to apply to the policy:

i. Navigate to**Infrastructure** > **Settings** > **Alerts**.

ii. Select**Create alert condition**, define the required alert details, add it to the policy you created in the previous step.

iii. Select**Create**.

Result

Alerts start flowing from the New Relic console into the Event Management plugin. The plugin extracts information from the original New Relic alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to**All Events**to see the events.

Note: By default, host binding is enabled for New Relic events for all providers (Azure/AWS/GCP). If all hosts in the environment are discovered using Cloud Discovery by providing credentials and discovered resources are in the cmdb_ci_vm_object list, then the VM binding may not occur. To resolve this, you must enable the**New Relic - General**event rule. For further information about Event rules, see Event rules.

Integrate New Relic with REST API key token

Integrate using an API key to establish secure communication and automate data exchange via REST API. This simplifies integration, enabling seamless access to services and enhancing operational efficiency.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: evt_mgmt_integration

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API Key**.

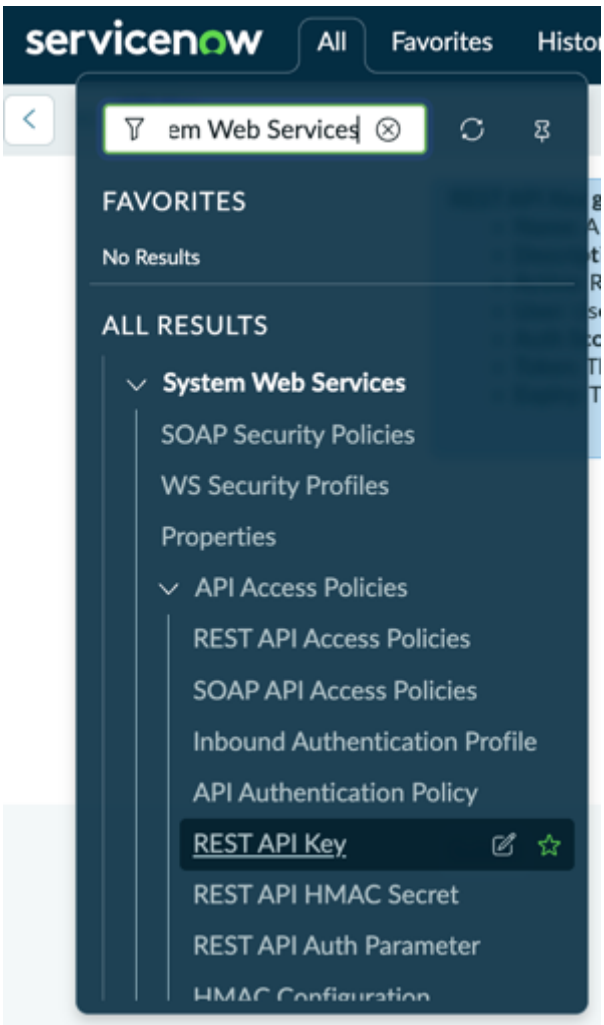

2. Select **New**.
3. On the form, fill the fields.

> **Note:** The user needs to be created with the evt_mgmt_integration role. Create a new one on the sys_user table to use here to facilitate the right access.

API Key

Field	Description
Name	Name to identify the REST API Key
Description	Description for the REST API Key.
Active	Status of the REST API Key.
User	User associated to the REST API Key. Use the look-up icon to select the user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Auth Scope	Option to add auth scope to manage the authority of the REST API Key.
Token	The REST API key generated by the ServiceNow AI Platform. Copy the key to use as part of the REST API call within the Header or Query parameter.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration.

4. Unlock Auth Scope and add: UserAccount.
5. Use the form menu and select **Save**.

   The system generates a token and saves it in the **Token** field. To see the token, use the lock icon and copy the contents display below the field. This is your query parameter (or header) value when your other system sends a REST API request to ServiceNow.

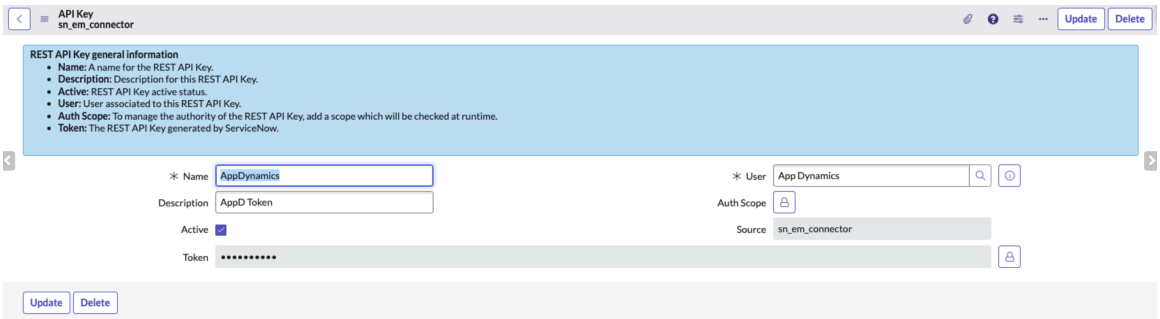

Integrate Catchpoint events

Integrate Catchpoint with Event Management by adding an alert webhook in the Catchpoint platform.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_admin

About this task

Configure the Event Management environment for the collection of events from Catchpoint by authenticating Catchpoint as a data source. In the Catchpoint platform, set your ServiceNow AI Platform as the rest endpoint using a webhook.

Procedure

1. In the Catchpoint platform, create an alert webhook destination.

   a. Login to the Catchpoint platform.

   b. Navigate to **Catchpoint** > **Setting** > **API** > **Alert Data Webhook** to set the webhook endpoint.

   c. In the API Alert Webhook, provide the following details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

API Alert Webhook form

Field	Value
Name	Any name based on your requirement.
URL	The default format of the URL to push events from Catch Point to the ServiceNow AI Platform instance, `https://<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=catchpoint`.
On Failure Alert	The email address for the API Alert webhook.
Notification Trigger	Select based on your requirement. The default value is 3.
Add the following headers	▪ Authorization: **Basic {Base 64 | username:password}** ▪ Content-Type: **application/json**
Format	Select **Template** and add the following template: `{ "Source": "CatchPoint", "severity": "${notificationLevelId}", "alertType": "${AlertType}", "time_of_event": "${alertCreateDateUtc(YYYY-MM-DD HH:MI:SS)}", "monitorType": "${Switch(${MonitorTypeId},'0', 'IE','2', 'Object','3','Emulated','8', 'Ping','9', 'Tracert','10', 'DNS-Traversal','11', 'Ping-TCP','12', 'DNS-Exp', '13','DNS-Direct', '14','Tracert-UDP', '15', 'Port-TCP','16', 'FTP','17', 'Data-Push-API','18', 'Chrome','19', 'Playback', '20', 'Playback-Mobile', '21', 'SMTP', '22', 'Port-UDP', '23', 'Ping-UDP', '24', 'Streaming', '25', 'API', '26', 'Mobile', '27', 'SFTP', '28', 'SSH', '29', 'Tracert-TCP')}", "testName": "${TestName}", "testId": "${TestId}",`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value

	"testUrl": "${TestUrl}",
	"testPath": "${testPath}",
	"testLink": "${testLink}",
	"productName":
	"${ProductName}",
	"clientId": "${ClientId}",
	"productId": "${productId}",
	"nodeName":"${nodeDetails("${
	nodeName}")}",
	"nodeClientServerAddress":"${
	nodeDetails(\"${NodeClientAdd
	ress},\"),}",
	"nodeServerAddress":
	"${nodeDetails(\"${nodeServe
	rAddress},\"),}",
	"Labels" : "${testLabels}",
	"addtionalInformation":
	"Test initial URL:
	${testUrl}\n\nAlert type:
	${AlertType}\n\nScatter
	Chart:
	${scatterplotChartURL}\n\nWa
	terfall Chart:
	${waterfallChartURL}\n\nTest
	properties:
	${testLink}\n\nTest time
	of run / Test time of
	alert (Central Time):
	${reportDateLocal(YYYY-MM-DD
	HH:MI:SS)} /
	${alertCreateDateLocal(YY
	YY-MM-DD HH:MI:SS)}"}

2. For the HTTP Endpoint monitoring tests, in the ServiceNow instance, create a CI.

> **Note:** The Catchpoint connector supports multiple types of monitoring tests such as Endpoint, Networking, and others. The following steps apply only to the HTTP Endpoint monitoring tests.

   a. Navigate to **Configuration** > **CI Class Manager**.

   b. Click **Hierarchy** and search for **HTTP(S) Endpoint**.

   c. In the CI Type field, select **cmdb_ci_endpoint_http**.
   If the services are not modeled, then you should manually add the CIs under the [cmdb_ci_endpoint_http] table. Also, the manual endpoint should be removed before the automatic discovery of services. For other type tests CI bindings, you can create Event rules based on their requirements.

Result

Alerts start flowing from the Catchpoint connector into the Event Management plugin. The following table shows how the Catchpoint severity mapping translates into the ServiceNow mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Severity levels between Catchpoint and ServiceNow

Catchpoint severity	ServiceNow severity
0	Warning
1	Critical
2/3	Clear

Integrate Datadog platform events

Integrate Datadog with Event Management by adding a standard webhook in the Datadog console.

What authentication is used

There are two methods of authentication:

- OAuth authentication: Provides enterprise-grade authentication to keep your enterprise environment safe. Authentication is performed using Datadog Monitor V1 or V2 access tokens. For more information, see Integrate Datadog with OAuth authentication.
- Basic webhook authentication: Provides a basic standard of authentication, without the need for Datadog Active Directory. This authentication can be especially useful for distributed small teams, such as SRE or DevOps teams. For more information, see Integrate Datadog with basic authentication.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

Integrate Datadog with basic authentication

Integrate Datadog with Event Management by adding a standard webhook in the Datadog console.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

- Role required: evt_mgmt_integration
- Confirm you have created a user with an **Identify Type** of **Machine** and the evt_mgmt_integration role.
- The Event Management plugin must be installed on the ServiceNow AI Platform instance.
- Verify Configuration Items for the hosts managed by Datadog exist in ServiceNow. These CIs can be physical or virtual, and can be either manually created or discovered using IP discovery or Cloud discovery.

About this task

Configure the Event Management environment for the collection of events from Datadog. In your Datadog console, set your ServiceNow AI Platform instance as the rest endpoint using a standard webhook.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Datadog console, add cloud provider tags:

   a. Navigate to **Menu** > **Infrastructure** > **Hosts**.

   b. Click **Inspect** for the VM on which the Datadog agent is installed.

   c. Click **Edit Tags** and add the *cloud_provider* tag.
   The tag should be suffixed with the value of **aws**, **azure**, or **gcp** to identify the AWS, Microsoft Azure, or Google Cloud Platform (GCP) provider options. For example: cloud_provider:aws

   d. Click **Save Tags** to save the new tag.

   Note: Datadog adds the *cloud_provider* tag by default when the Azure and Google integration is installed on the Datadog agent. Tags can also be added at the cloud provider level. Datadog collects all the tags defined at the VM level on the cloud.

2. In the Datadog console, add a webhook:

   a. Navigate to **Menu** > **Integration** > **Integrations**, and search for webhooks.

   b. Create a webhook that includes the user ID and password of the ServiceNow user in the endpoint.
   For example, use `https://<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=datadog`.

   Note: Verify that the selected user is assigned the evt_mgmt_integration role. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

   c. Add the required payload structure.

   - Add in the following payload structure for monitoring alerts.

```
{
    "body": "$EVENT_MSG",
    "last_updated": "$LAST_UPDATED",
    "event_type": "$EVENT_TYPE",
    "title": "$EVENT_TITLE",
    "date": "$DATE",
    "org": {
        "id": "$ORG_ID",
        "name": "$ORG_NAME"
    },
    "id": "$ID",
    "alert_id": "$ALERT_ID",
    "alert_metric": "$ALERT_METRIC",
    "metric_namespace": "$METRIC_NAMESPACE",
    "alert_priority": "$ALERT_PRIORITY",
    "alert_transition": "$ALERT_TRANSITION",
    "alert_status": "$ALERT_STATUS",
    "alert_title": "$ALERT_TITLE",
    "alert_type": "$ALERT_TYPE",
    "host_name": "$HOSTNAME",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    "priority": "$PRIORITY",
    "tags": "$TAGS",
    "alert_scope": "$ALERT_SCOPE"
}
```

- Add in the following payload structure for monitoring work items.

```
{
  "id": "$CASE_ID",
  "type": "$CASE_TYPE",
  "status": "$CASE_STATUS",
  "priority": "$CASE_PRIORITY",
  "title": "$EVENT_TITLE",
  "body": "$EVENT_MSG",
  "assignee": {
    "id": "$CASE_ASSIGNEE_ID",
    "name": "$CASE_ASSIGNEE_NAME",
    "email": "$CASE_ASSIGNEE_EMAIL"
  },
  "attributes": $CASE_ATTRIBUTES,
  "correlated_alerts": $CASE_CORRELATED_ALERTS,
  "date_created": "$DATE",
  "date_updated": "$LAST_UPDATED",
  "tags": "$TAGS"
}
```

ServiceNow auto-detects the case payload and represents it as a work Item alert (message key `Datadog_Case_<case_id>`).

3. In the Datadog console, define rules:
   a. Navigate to **Menu** > **Monitors** > **Manage Monitors**, and configure rules on when to create a problem from Datadog on the managed hosts.
   b. In the **Notify your team** section, select the webhook on which the problem should be notified.

   Similarly, create a rule for case or work items that triggers on events such as created, priority changed, status changed, closed and so on. It then notifies your ServiceNow webhook.

Result

Alerts start flowing from Datadog agents into the Event Management plugin. The plugin extracts information from the original Datadog alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to **All Events** to see the events.

For case or work items, while creating push connectors select the **Configure Work Item Alerts** button. You are redirected to the Datadog work Items pull connector instance form, where you configure related-alert fetch and bi-directional synchronization. For more information, see Configure a Datadog work item connector instance.

Note: By default, host binding is enabled for Datadog events for (AWS/Azure/GCP) providers. If all hosts in the environment are discovered using Cloud Discovery by providing credentials and discovered resources are in the cmdb_ci_vm_object list, then the VM binding may not occur. To resolve this, you must enable the **Datadog - General** event rule. For further information about Event rules, see Event rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Configure a push connector

Configure a Datadog work item connector instance

Integrate Datadog with OAuth authentication

Integrate Datadog with Event Management by authenticating Datadog V1 or V2 tokens in the Datadog Monitor.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

The Event Management plugin must be installed on the ServiceNow AI Platform instance.

Verify Configuration Items for the hosts managed by Datadog exist in ServiceNow. These CIs can be physical or virtual, and can be either manually created or discovered using IP discovery or Cloud discovery.

Role required: evt_mgmt_integration

About this task

Datadog integration with OAuth authentication is available starting from the Washington DC platform release or later.

Procedure

1. Configuration on the ServiceNow instance.

   a. Create *`glide.oauth.inbound.client.credential.grant_type.enabled`* system property to use Client Credentials grant type for OAuth inbound integrations.

      - Type: true|false
      - Value: true

      Note: An administrator must create this system property because the evt_mgmt_admin role lacks permission to do so.

   b. Create an OAuth Entity endpoint for clients to access the instance.

      i. Navigate to **All** > **System OAuth** > **Application Registry** and then select **New**.

      ii. On the interceptor page, select **Create an OAuth API endpoint for external clients** and then fill in the form.

      iii. To create an OAuth client record, select **Submit**.

   c. Add the OAuth Application User and the Default Grant Type to the created OAuth client record.

      i. On the OAuth Entity form, add the **OAuth Application User** and the **Default Grant Type** fields to use the Client Credentials grant type for OAuth inbound integrations.

      ii. Open the created OAuth client record.

      iii. From the page header, select the **More** options icon.

      iv. Select **Configure** > **Form Layout**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

v. On the Form Layout page, add **OAuth Application User and Default Grant Type** from the list of fields.

vi. **Save** or **Update** the form.

vii. For the **OAuth Application User** field, select the user.

The user must have the evt_mgmt_integration role that will be used to create the events sent by Datadog.

Note: To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

viii. Set **Default Grant Type** field value to **Client Credentials**.

2. In the Datadog console, add cloud provider tags.
   a. Navigate to **Menu** > **Infrastructure** > **Hosts**.
   b. Select **Inspect** for the VM on which the Datadog agent is installed.
   c. Select **Edit Tags** and add the cloud_provider tag.

   The tag must be suffixed with the value of aws, azure, or gcp to identify the AWS, Microsoft Azure, or Google Cloud Platform (GCP) provider options. For example: cloud_provider:aws.

   d. Select **Save Tags**.

   Note: Datadog adds the cloud_provider tag by default when the Azure and Google integration is installed on the Datadog agent. Tags can also be added at the cloud provider level. Datadog collects all the tags defined at the VM level on the cloud.

3. In the Datadog console, add a webhook.
   a. Navigate to **Menu** > **Integration** > **Integrations** and search for webhooks.
   b. On the webhook tile, select **Configure**.
   c. Select the **Auth Methods** tab and then select **New Auth Method**.
   d. Enter the name and select the protocol as **OAuth 2.0 Client Credentials**.
   e. Provide Access token URL as `https://<instance_name>.service-now.com/oauth_token.do`.
   f. Provide the value for **Client ID** with the client id of the OAuth Entity Client created in ServiceNow instance.
   g. Provide “Client secret” with the client secret of the OAuth Entity Client created in ServiceNow instance.
   h. Select **Save**.
   i. Create webhook URL as

   `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=datadog`

   j. Select the OAuth Method created in the previous steps in AuthMethod dropdown.
   k. ▪ Add in the following payload structure for monitoring alerts.

```
{
    "body": "$EVENT_MSG",
    "last_updated": "$LAST_UPDATED",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    "event_type": "$EVENT_TYPE",
    "title": "$EVENT_TITLE",
    "date": "$DATE",
    "org": {
        "id": "$ORG_ID",
        "name": "$ORG_NAME"
    },
    "id": "$ID",
    "alert_id": "$ALERT_ID",
    "alert_metric": "$ALERT_METRIC",
    "metric_namespace": "$METRIC_NAMESPACE",
    "alert_priority": "$ALERT_PRIORITY",
    "alert_transition": "$ALERT_TRANSITION",
    "alert_status": "$ALERT_STATUS",
    "alert_title": "$ALERT_TITLE",
    "alert_type": "$ALERT_TYPE",
    "host_name": "$HOSTNAME",
    "priority": "$PRIORITY",
    "tags": "$TAGS",
    "alert_scope": "$ALERT_SCOPE"
}
```

- Add in the following payload structure for monitoring work items.

```
{
  "id": "$CASE_ID",
  "type": "$CASE_TYPE",
  "status": "$CASE_STATUS",
  "priority": "$CASE_PRIORITY",
  "title": "$EVENT_TITLE",
  "body": "$EVENT_MSG",
  "assignee": {
    "id": "$CASE_ASSIGNEE_ID",
    "name": "$CASE_ASSIGNEE_NAME",
    "email": "$CASE_ASSIGNEE_EMAIL"
  },
  "attributes": $CASE_ATTRIBUTES,
  "correlated_alerts": $CASE_CORRELATED_ALERTS,
  "date_created": "$DATE",
  "date_updated": "$LAST_UPDATED",
  "tags": "$TAGS"
}
```

ServiceNow auto-detects the case payload and represents it as a work Item alert (message key `Datadog_Case_<case_id>`).

Result

Alerts start flowing from Datadog agents into the Event Management plugin. The plugin extracts information from the original Datadog alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to All Events to see the events.

For case or work items, while creating push connectors select the **Configure Work Item Alerts** button. You are redirected to the Datadog work Items pull connector instance form, where you configure related-alert fetch and bi-directional synchronization. For more information, see Configure a Datadog work item connector instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: By default, host binding is enabled for Datadog events for (AWS/Azure/GCP) providers. If all hosts in the environment are discovered using Cloud Discovery by providing credentials and discovered resources are in the cmdb_ci_vm_object list, then the VM binding may not occur. To resolve this, you must enable the Datadog - General event rule. For further information about Event rules, see Event rules.

Create a rule for case or work items that triggers on events such as created, priority changed, status changed, and closed. The rule then notifies your ServiceNow webhook.

Related topics

Configure a push connector

Configure a Datadog work item connector instance

Integrate Google Cloud Platform (GCP) events

Integrate Google Cloud Platform (GCP) with Event Management by adding a standard webhook in the GCP console.

Before you begin

- Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.
- The Event Management plugin must be installed on the ServiceNow AI Platform instance.
- Verify Configuration Items for the hosts managed by GCP exist in ServiceNow. These CIs can be physical or virtual, and can be either manually created or discovered using IP discovery or Cloud discovery.
- The Event Management Connector plugin supports GCP alert data in JSON 1.2 format only.
- Role required: evt_mgmt_integration

About this task

Configure the Event Management environment for the collection of events from GCP. In your GCP console, set your ServiceNow AI Platform instance as the rest endpoint using a standard webhook.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

Procedure

1. In the GCP console, add a webhook.

   a. Navigate to **Monitoring** > **Alerting**, and click **Edit Notification Channels**.

   b. In the Webhooks section, click **Add New**.

   c. Select the **Use HTTP Basic Auth** check box, and enter the user name and password of the relevant ServiceNow user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

> **Note:** Ensure that the selected user is assigned the evt_mgmt_integration role. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

d. Enter the endpoint URL.
For example, use`https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=googlemonitor`.

2. In the GCP console, create an alerting policy.
   a. Navigate to **Monitoring** > **Alerting**, and select **Create Policy**.
   b. Configure policy conditions that initiate alerts when the conditions are violated.
   c. Select the webhook on which the problem should be notified.

> **Note:** If you need to directly open an incident in the GCP console from the alert, navigate to **Quick Response (in Alert)** > **Show Incident in GCP Console**.

Result
Alerts start flowing from GCP into the Event Management plugin. The plugin extracts information from the original GCP alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to**All Events** to see the events.

GCP does not send values such as the severity level in the payload. The default severity is Minor, which can be changed in the*`Push Connector Configuration`* section of **Push Connectors** > **Google Monitor Push Connector**. The valid values of severity are**1- Critical**,**2- Major**,**3- Minor**,**4- Warning**, and**5- Info**.

> **Note:**
>
> To map the GCP hosts to the**cmdb_ci_server**record instead of**cmdb_ci_vm_object**, disable the**Google Monitor - General**event rule created for GCP. For further information about event rules, see Event rules.

Event collection from Google Cloud Platform (GCP)

The MID WebService Event Collector enables you to collect JSON formatted event messages sent from Google Cloud Platform (GCP).

Before you begin
Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Ensure that you:

- Deploy and start the MID Server. See MID Server configuration .
- Configure and start the MID Web Server. See Configure the MID Web Server extension.
- Configure and start the MID WebService Event Collector. See Configure the MID WebService Event Collector Context.

Role required: evt_mgmt_admin

About this task
JSON formatted event messages are sent from Google Cloud Platform. The MID Server transforms the collected event messages by parsing them using the TransformEvents_MidGoogleMonitor script include, located here: **Event Management** >

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integrations > **PushConnectors**. In the Push Connectors page, click **Google Monitor Mid Push Connector**.

The default format of the URL to push event messages from the GCP to the MID Server is `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=TransformEvents_MidGoogleMonitor`.

Variables in the default URL

Variable	Description
MID_Server_IP	IP address of the MID Web Server Extension.
MID_Web_Server_Port	Listening port of the MID Web Server Extension.
MID_Web_Server_User	Username for the MID Web Server Extension.
MID_Web_Server_Password	Password of the user of the MID Web Server.

The following procedure describes the collection of JSON formatted event messages using basic authentication.

Procedure

1. Verify that the MID Server IP is accessible from the Google Cloud Platform (GCP) console and that the MID rest point is accessible from GCP.
   Alternatively, install the MID Server on a GCP VM and provide the VM external IP address in the URL. In the GCP console, create a firewall rule.
   a. Navigate to **VPC Network** > **Firewall** > **Create Firewall**.
   b. Set the direction to **ingress** to allow the mid_web_server_port to connect.
   c. Add the GCP VM that the MID Server is installed on to the **Applicable to instances** list.
2. In the GCP console, add a webhook.
   a. Navigate to **Monitoring** > **Alerting** and click **Edit Notification Channels**.
   b. In the Webhooks section, click **Add New**.
   c. Select the **Use HTTP Basic Auth** check box, and enter the username and password of the MID Webserver User.

      Note: Ensure that the selected user is assigned the evt_mgmt_integration role. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

   d. Enter the endpoint URL.
      For example, use `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=TransformEvents_MidGoogleMonitor`
3. In the GCP console, create an alerting policy.
   a. Navigate to **Monitoring** > **Alerting** and click **Create Policy**.
   b. Configure the policy conditions that initiate alerts when the conditions are violated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Select the webhook notification channel on which the problem should be notified.

d. Select the **Notify on incident closure** check box to send the closure events to the MID Server and save the alert policy.

Note: If you need to directly open an incident in the GCP Console from the alert, navigate to **Quick Response (in Alert)** > **Show Incident in GCP Console**. For more information about creating an alert policy, see Managing alerting policies on the Google Cloud documentation site.

Result

GCP does not send values such as the severity level into the payload. The default severity is Minor, which can be changed in the *Push Connector Configuration* section of **Push Connectors** > **Google Monitor MID Push Connector**. The valid values of severity are **1- Critical**, **2- Major**, **3- Minor**, **4- Warning**, and **5- Info**.

Integrate Grafana events

Integrate Grafana with Event Management by adding Grafana as an authenticated data source.

What authentication is used

The methods of authentication are as follows:

- Basic webhook authentication: Provides a basic standard of authentication, using username and password. For more information, see Integrate Grafana events with basic authentication.
- OAuth authentication: Provides enterprise-grade authentication to keep your enterprise environment safe. Authentication is performed using access tokens. For more information, see Integrate Grafana with OAuth authentication.
- API key authentication: Integrate using API key to secure communication between Grafana and ServiceNow. For more information, see Integrate Grafana with REST API key token.

Integrate Grafana events with basic authentication

Integrate Grafana with Event Management by adding a standard webhook in the Grafana console.

Before you begin

- Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.
- The Event Management plugin must be installed on the ServiceNow AI Platform instance.
- Confirm you have created a user with an **Identify Type** of **Machine** and the evt_mgmt_integration role.
- The CI binding happens through the IP, instances should be discovered as servers to initiate the binding. CIs discovered through cloud or CIs in **cmdb_ci_vm_object** don't get bounded as Grafana doesn't provide any information for object IDs.
- Role required: evt_mgmt_integration

About this task

Configure the Event Management environment for the collection of events from Grafana. In your Grafana console, set your ServiceNow AI Platform instance as the rest endpoint using a standard webhook.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can also configure Grafana as an event push connector from the Integrations Launchpad. For more information, see Configure an event push connector.

Note: Perform steps 1 and 2 for Grafana v7.4.3. For Grafana v9.3.2. skip to step 3.

Procedure

Grafana v7.4.3

1. In the Grafana console, create a notification channel.
   a. Navigate to **Alerting** > **Notification Channels**.
   b. Click **New Channel** and select the type as **Webhook**.
   c. Enter a URL.
   For example, use `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=grafana`.
   d. Under the **Optional WebHook** settings, select the HTTP Method, **POST**, and enter the user name and password of the relevant ServiceNow user.
      Note: Ensure that the selected user is assigned the evt_mgmt_integration role. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.
   e. Click **Test** to verify the connection.
2. In the Grafana console, create alerting rules.
   a. Click **Edit** on the required panel and then **Create Alert**.
   b. In the **Alert** tab, add the following tags:

Tags	Description
Severity	(This tag is optional and has to be lower cased.) Severity with which the alert has to be created.
Count	(This tag is required, otherwise the alerts don't close automatically in ServiceNow.) The number of devices, hosts, or servers monitored by this alert rule.

   c. Add the notification channel under **Notifications** in the alert rule.
   d. Click **Save**.
   e. Navigate to **Alerting** > **Alert Rules** to access the created rule.

Grafana v9.3.2.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the Grafana console, create a contact point.

   a. Navigate to **Alerting** > **Contact Points**.

   b. Select **New Contact Point**, enter a name, and select the type as **Webhook**.

   c. Enter a URL.
   For example, use `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=grafana`.

   d. In the **Optional WebHook** settings, set the HTTP Method as **POST**, and enter the user name and password of the relevant ServiceNow user.

      Note: Ensure that the selected user is assigned the evt_mgmt_integration role. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

   e. Click **Save** and **Test** to verify the connection.

4. In the Grafana console, create a notification policy.

   a. Navigate to **Alerting** > **Notification Policies**.

   b. Select **New Policy**, add matchers.
   Matchers are labels that direct alert rules to the contact point.

   c. Select your contact point.

   d. Select **Save Policy**.

5. In the Grafana console, create alert rules.

   a. Navigate to **Alerting** > **Alert Rule**.

   b. Select **New Rule**.

   c. Set the Query and Alert Condition.

   d. Set the data source, metrics, and the necessary filters.

   e. Set the alert evaluation behavior.

   f. Enter the rule name.

   g. In the summary and annotations section, add a summary or a description.

   h. From the scroll down, select **+Add New**.

   i. Enter the node, resource, and metric information.
   This information is required for the node, resource, and metric_name fields for the event.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Use the value string to pass these values. For example, a node can be {{ $values.C0.Labels.instance }}.

j. Under custom labels, provide the required labels to direct your notifications to the appropriate contact point.

Note: You can also pass the severity under custom labels. The same is used for events. The available severity options are: critical, major, minor, warning.

k. Click **Save**.

Result

Alerts start flowing from Grafana into the Event Management plugin. The plugin extracts information from the original Grafana alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to **All Events** to see the events.

Note:

- For Grafana v7.4.3, you may see extra clear events coming into ServiceNow which aren't bound to an alert. Ignore them as they're created to overcome the issue of Grafana not sending information in the clear notification.
- A severity parameter is added which is used only for Grafana 8 and above. The provided value is used as a severity for an event if the severity isn't passed in custom labels.

Integrate Grafana with OAuth authentication

Integrate Grafana with Event Management with ServiceNow using OAuth authentication.

Before you begin

- Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.
- The Event Management plugin must be installed on the ServiceNow AI Platform instance.
- The CI binding happens through the IP, instances should be discovered as servers to initiate the binding. CIs discovered through cloud or CIs in **cmdb_ci_vm_object** don't get bounded as Grafana doesn't provide any information for object IDs.
- Role required: evt_mgmt_integration

About this task

Configure the Event Management environment for the collection of events from Grafana. In your Grafana console, set your ServiceNow AI Platform instance as the rest endpoint using a standard webhook.

Perform the following steps for Grafana v9.3.2 to send event to ServiceNow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Configuration on the ServiceNow instance.

   a. Create *glide.oauth.inbound.client.credential.grant_type.enabled* system property to use Client Credentials grant type for OAuth inbound integrations.

      - Type: true|false
      - Value: true

      Note: An administrator must create this system property because the evt_mgmt_admin role lacks permission to do so.

   b. Create an OAuth Entity endpoint for clients to access the instance.

      i. Navigate to **All** > **System OAuth** > **Application Registry** and then select **New**.

      ii. On the interceptor page, select **Create an OAuth API endpoint for external clients** and then fill in the form.

      iii. To create an OAuth client record, select **Submit**.

   c. Add the OAuth Application User and the Default Grant Type to the created OAuth client record.

      i. On the OAuth Entity form, add the **OAuth Application User** and the **Default Grant Type** fields to use the Client Credentials grant type for OAuth inbound integrations.

      ii. Open the created OAuth client record.

      iii. From the page header, select the **More** options icon.

      iv. Select **Configure** > **Form Layout**

      v. On the Form Layout page, add **OAuth Application User and Default Grant Type** from the list of fields.

      vi. **Save** or **Update** the form.

      vii. For the **OAuth Application User** field, select the user.

      The user must have the evt_mgmt_integration role that will be used to create the events sent by Grafana.

      Note: To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

      viii. Set **Default Grant Type** field value to **Client Credentials**.

2. In the Grafana console, create a contact point.

   a. Navigate to **Alerts &IRM** > **Alerting** > **Contact Points**.

   b. Select **New Contact Point**, enter a name, and select the type as **Webhook**.

   c. Enter a URL.
   For example, use `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=grafana`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. In the **Optional WebHook** settings, set the HTTP Method as **POST**.

e. In **HTTP Config**, select **+Add**.

f. In **OAuth2**, select **+Add**.

g. Provide token URL as: `https://<instance_name>.servicenow.com/oauth_token.do`.

h. Provide the value for **Client ID** with the client id of the OAuth Entity Client created in ServiceNow instance.

i. Provide **Client secret** with the client secret of the OAuth Entity Client created in ServiceNow instance.

j. Select **Save Contact point**.

3. In the Grafana console, create a notification policy.

a. Navigate to **Alerting** > **Notification Policies**.

b. Select **New Policy**, add matchers.
Matchers are labels that direct alert rules to the contact point.

c. Select your contact point.

d. Select **Save Policy**.

4. In the Grafana console, create alert rules.

a. Navigate to **Alerting** > **Alert Rule**.

b. Select **New Rule**.

c. Set the Query and Alert Condition.

d. Set the data source, metrics, and the necessary filters.

e. Set the alert evaluation behavior.

f. Enter the rule name.

g. In the summary and annotations section, add a summary or a description.

h. From the scroll down, select **+Add New**.

i. Enter the node, resource, and metric information.
This information is required for the node, resource, and metric_name fields for the event.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Use the value string to pass these values. For example, a node can be {{ $values.C0.Labels.instance }}.

j. Under custom labels, provide the required labels to direct your notifications to the appropriate contact point.

Note: You can also pass the severity under custom labels. The same is used for events. The available severity options are: critical, major, minor, warning.

k. Click **Save**.

Result

Alerts start flowing from Grafana into the Event Management plugin. The plugin extracts information from the original Grafana alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to **All Events** to see the events.

A severity parameter is added which is used only for Grafana 8 and above. The provided value is used as a severity for an event if the severity isn't passed in custom labels.

Integrate Grafana with REST API key token

Integrate using an API key to establish secure communication and automate data exchange via REST API. This simplifies integration, enabling seamless access to services and enhancing operational efficiency.

Before you begin

Role required: evt_mgmt_integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API Key**.

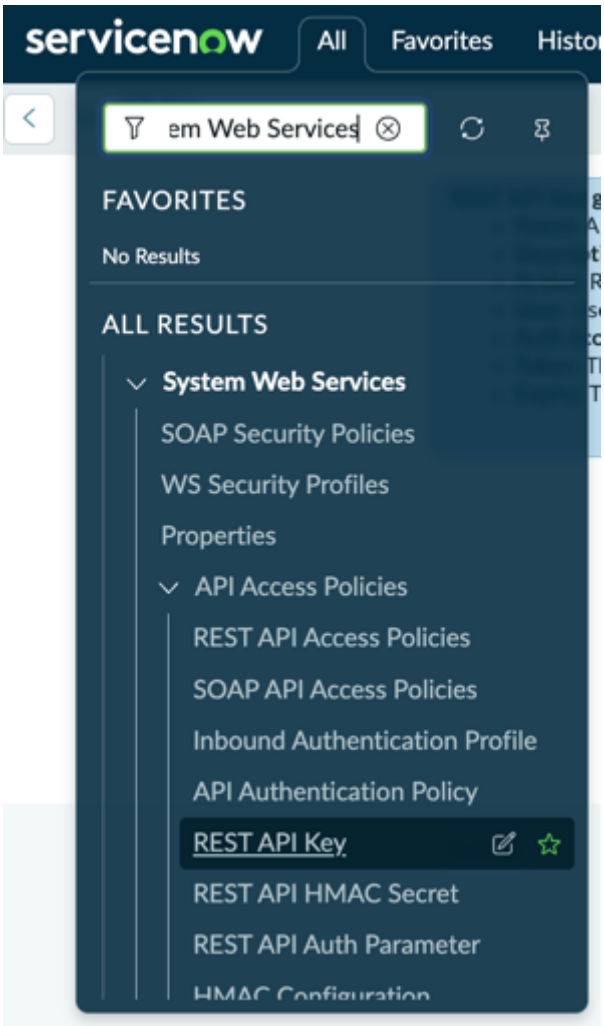

2. Select **New**.
3. On the form, fill the fields.

> **Note:** The user needs to be created with the evt_mgmt_integration role. Create a new one on the sys_user table to use here to facilitate the right access.

API Key

Field	Description
Name	Name to identify the REST API Key
Description	Description for the REST API Key.
Active	Status of the REST API Key.
User	User associated to the REST API Key. Use the look-up icon to select the user.
Auth Scope	Option to add auth scope to manage the authority of the REST API Key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Token	The REST API key generated by the ServiceNow AI Platform. Copy the key to use as part of the REST API call within the Header or Query parameter.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration.

4. Unlock Auth Scope and add: UserAccount.
5. Use the form menu and select **Save**.

   The system generates a token and saves it in the **Token** field. To see the token, use the lock icon and copy the contents display below the field. This is your query parameter (or header) value when your other system sends a REST API request to ServiceNow.

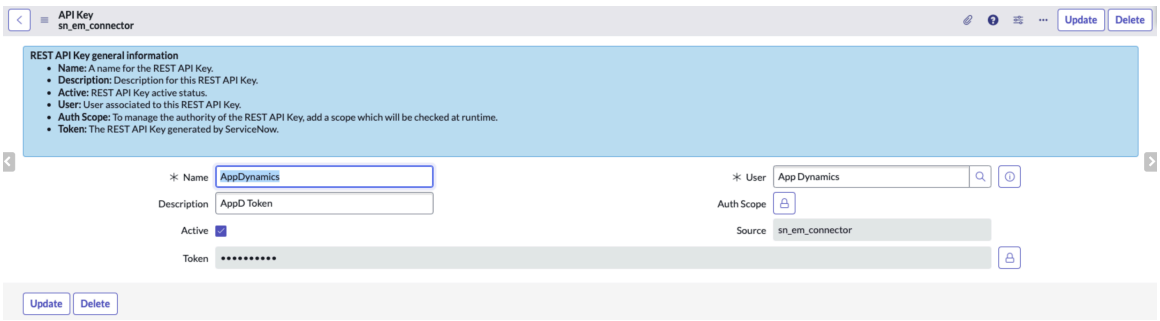

Integrate Honeycomb events

Integrate Honeycomb with Event Management by creating a webhook and configuring it as a trigger in the Honeycomb platform.

Before you begin

Discovery for Honeycomb services and datasets is not supported. For event/CI binding to work, manually create a CI in the ServiceNow instance.

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance. For more information, see Install Event Management

Role required: evt_mgmt_admin

About this task

Configure the Event Management environment for the collection of events from Honeycomb by authenticating Honeycomb as a data source. In the Honeycomb platform, set your ServiceNow AI Platform as the rest endpoint using a webhook.

Procedure

1. In the Honeycomb platform, create an integration of *`type=webhook`*.

   a. Navigate to **Usage** > **Integrations**.

   b. Select the **Add Integration** button.

   c. Enter values in the **Name** and **Webhook URL** fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When enting the Webhook URL, use the following format: `https://<USERNAME>:<PASSWORD>@<INSTANCE_NAME>.service-now.com/api/sn_em_connector/em/inbound_event?source=honeycomb`

d. In the **Provider** field, select **Webhook**.

e. Select **Add** to complete webhook creation.

2. Add the webhook to a trigger.

a. Navigate to **Triggers** > **Create New Trigger**.

b. Add a name and description for the trigger in the relevant fields.

c. In the **Alerts** section, configure the threshold and frequency in the relevant fields.

d. In the **Recipients** section, select **Add Recipient** and select the webhook created in step 1.

e. Select **Save Trigger**.

Note: To view the trigger in the Honeycomb platform, navigate to **Quick Response (in Alert)** and select **Show Trigger in Honeycomb**.

Result

Alerts flow from the Honeycomb connector into the Event Management plugin. The plugin extracts information from the original Honeycomb trigger message to populate the required event fields, and inserts the event into the database. You can view the events in your ServiceNow AI Platform instance by navigating **All Events**.

Honeycomb does not send severity information in the trigger message. The default severity for the Honeycomb alert is *3 - Minor*, which can be changed in the Push Connector Configuration section of **Push Connectors** > **Honeycomb Push Connector** . The valid severities are: *1 - Critical*, *2 - Major*, *3 - Minor*, *4 - Warning*, and *5 - Info*.

Integrate Instana events

Integrate Instana with Event Management by adding Instana as an authenticated data source.

What authentication is used

The methods of authentication are as follows:

- Basic webhook authentication: Provides a basic standard of authentication, using username and password. For more information, see Integrate Instana events.
- OAuth authentication: Provides enterprise-grade authentication to keep your enterprise environment safe. Authentication is performed using access tokens. For more information, see Integrate Instana with OAuth authentication.

Integrate Instana events

Integrate Instana with Event Management by adding a standard webhook in the Instana console.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Confirm you have created a user with an **Identify Type** of **Machine** and the evt_mgmt_integration role.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Instana console, navigate to **Settings** > **Events & Alerts** > **Alert Channels**.
2. Open the **Add Alert Channel** dropdown and select the **Generic Webhook** option.
3. In the **Webhook URLs** field, create an alert channel by entering the following URL:
   `https://<instance>.service-now.com/api/sn_em_connector/em/inbound_event?source=instana`
4. To authenticate the ServiceNow URL, enter the following in the **Custom HTTP Request Headers**:

Field	Value
Key	Authorization
Value	Basic <base64_encoded username:password>

5. Configure Instana to create alerts for event types or events.

   a. Navigate to **Settings** > **Events & Alerts** > **Alerts**.

   b. Select **New Alert**.

   c. In the **Events** section, select one of the following:
      - Alert on Event Type(s)
      - Alert on Event(s)

   d. When creating alerts for event types:
      i. Select the type of events for which you want alerts to be created, such as **Critical Issues**, **All Incidents**, **Warning issues**, and so forth.
      ii. In the **Alerting** section, select the **servicenow** alert channel.

   e. When creating alerts for events:
      i. Select the events for which you want alerts to be created.
      ii. In the **Alerting** section, select the **servicenow** alert channel.

   When Instana creates an alert, the alert is sent to the ServiceNow instance. Alerts created for the host entity are bound to the Server (cmdb_ci_server) table, and alerts created for the website entity are bound to the HTTP(S) (cmdb_ci_endpoint_http) table.

   Add other entity types in the Instana - ci_type Event Field Mapping table to ensure that the entity binds with the correct CI in the ServiceNow instance.

6. To view the event in the Instana source system, you can access it from a ServiceNow instance:
   a. Navigate to **All** > **Event Management** > **Service Operations Workspace**.
   b. Select **All Alerts**.
   c. Select the relevant alert.
   d. On the **Details** tab, locate the **Actions** section on the right side of the page and in the **Show Event on Instana UI** cell, select **Launch app**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integrate Instana with OAuth authentication

Integrate Instana with Event Management with ServiceNow using OAuth authentication.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_admin

Procedure

1. Configuration on the ServiceNow instance.

   a. Create *glide.oauth.inbound.client.credential.grant_type.enabled* system property to use Client Credentials grant type for OAuth inbound integrations.

      - Type: true|false
      - Value: true

      Note: An administrator must create this system property because the evt_mgmt_admin role lacks permission to do so.

   b. Create an OAuth Entity endpoint for clients to access the instance.

      i. Navigate to **All** > **System OAuth** > **Application Registry** and then select **New**.

      ii. On the interceptor page, select **Create an OAuth API endpoint for external clients** and then fill in the form.

      iii. To create an OAuth client record, select **Submit**.

   c. Add the OAuth Application User and the Default Grant Type to the created OAuth client record.

      i. On the OAuth Entity form, add the **OAuth Application User** and the **Default Grant Type** fields to use the Client Credentials grant type for OAuth inbound integrations.

      ii. Open the created OAuth client record.

      iii. From the page header, select the **More** options icon.

      iv. Select **Configure** > **Form Layout**

      v. On the Form Layout page, add **OAuth Application User and Default Grant Type** from the list of fields.

      vi. **Save** or **Update** the form.

      vii. For the **OAuth Application User** field, select the user.

      The user must have the evt_mgmt_integration role that will be used to create the events sent by Grafana.

      Note: To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

      viii. Set **Default Grant Type** field value to **Client Credentials**.

2. In the Instana console, navigate to **Settings** > **Events & Alerts** > **Alert Channels**.

3. Open the **Add Alert Channel** drop-down list and select the **Generic Webhook** option.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the **Webhook URLs** field, create an alert channel by entering the following URL: `https://<instance>.service-now.com/api/sn_em_connector/em/inbound_event?source=instana`.
5. Enable the **Enable OAuth** toggle switch to authenticate the ServiceNow URL.
   a. In the **Client ID** field, provide the client id value of the OAuth Entity Client created in ServiceNow instance.
   b. In the **Client Secret** field, provide the client secret value of the OAuth Entity Client created in ServiceNow instance.
   c. In the **Token URL** field, enter `https://<instance_name>.servicenow.com/oauth_token.do`
   d. To test the integration, select **Test channel**
   e. Select **Save**.
6. Configure Instana to create alerts for event types or events.
   a. Navigate to **Settings** > **Events & Alerts** > **Alerts**.
   b. Select **New Alert**.
   c. In the **Events** section, select one of the following:
      - Alert on Event Type(s)
      - Alert on Event(s)
   d. When creating alerts for event types:
      i. Select the type of events for which you want alerts to be created, such as **Critical Issues**, **All Incidents**, **Warning issues**, and so forth.
      ii. In the **Alerting** section, select the **servicenow** alert channel.
   e. When creating alerts for events:
      i. Select the events for which you want alerts to be created.
      ii. In the **Alerting** section, select the **servicenow** alert channel.

   When Instana creates an alert, the alert is sent to the ServiceNow instance. Alerts created for the host entity are bound to the Server (cmdb_ci_server) table, and alerts created for the website entity are bound to the HTTP(S) (cmdb_ci_endpoint_http) table.

   Add other entity types in the Instana - ci_type Event Field Mapping table to ensure that the entity binds with the correct CI in the ServiceNow instance.
7. To view the event in the Instana source system, you can access it from a ServiceNow instance:
   a. Navigate to **All** > **Event Management** > **Service Operations Workspace**.
   b. Select **All Alerts**.
   c. Select the relevant alert.
   d. On the **Details** tab, locate the **Actions** section on the right side of the page and in the **Show Event on Instana UI** cell, select **Launch app**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integrate ServiceNow Cloud Observability Events

Integrate ServiceNow Cloud Observability with Event Management by adding a standard webhook in the ServiceNow Cloud Observability platform. Download the Event Management Connector plugin from the ServiceNow Store so you can integrate with ServiceNow Cloud Observability.

Before you begin

- Role required: evt_mgmt_integration

About this task

Authenticate ServiceNow Cloud Observability as a data source to enable Event Management to collect events from ServiceNow Cloud Observability. In your ServiceNow Cloud Observability platform, use a standard webhook to set your ServiceNow AI Platform instance as the rest endpoint.

Note: Ensure the evt_mgmt_integration role is assigned to the selected user. To ensure proper authentication, use the least privileged user with the evt_mgmt_integration role, rather than a high privileged user.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

Procedure

1. Create a CI.

   Note: Discovery for ServiceNow Cloud Observability services isn't supported. Create a CI manually in a Service Instance (cmdb_ci_service_auto) table to enable binding.

   a. Navigate to **All** > **Configuration** > **Application Service**.

   b. Select **New**.

   c. In the form, enter the name.
   The name must be the service name provided in the ServiceNow Cloud Observability platform.

   d. From the Operational Status list, set the operational status to **Operational**.

   e. Select **Next**.

   f. Select **Done**.

2. In your ServiceNow Cloud Observability platform, create a notification destination.

   a. From the navigation bar, select **Alerts**.

   b. Select the **Notification destinations** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Select **Create a destination**.

d. Select **ServiceNow** or **webhook**.
Use the ServiceNow destination to connect only with the ServiceNow platform. If you have any issues with the ServiceNow destination, then use the **webhook** destination option.

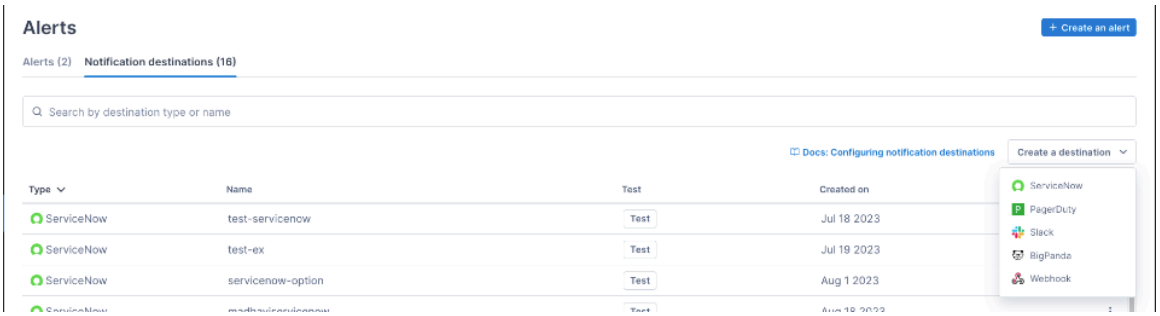

3. In your ServiceNow Cloud Observability platform, create a ServiceNow notification destination.

a. In the **Name** field, enter a name for the ServiceNow destination.

b. In the **URL** field, provide the ServiceNow instance URL.
This is the default format of the URL to push events from ServiceNow Cloud Observability to the ServiceNow AI Platform instance. Use this format: `https//<instance-name>.servicenow.com/api/sn_em_connector/em/inbound_event?source=lightstep`

c. In the **Username** field, enter your ServiceNow user name.

d. In the **Password** field, enter your ServiceNow user password.

Create a ServiceNow destination

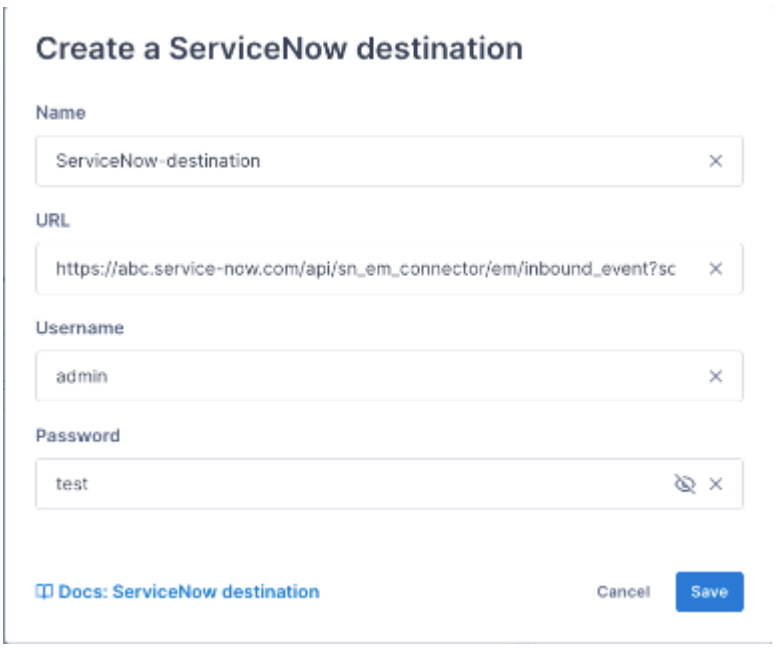

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: In a ServiceNow instance, CI binding occurs only if the name of the service in the ServiceNow Cloud Observability platform value matches the name fields of the CI you created in the *cmdb_ci_service_auto* table.

e. Select **Save**.

4. **Note:** If you have any issues with the ServiceNow destination in the previous step, then use the **webhook** destination option.

In your ServiceNow Cloud Observability platform, create a **webhook** notification destination.

a. In the **Name** field, enter a name for the webhook destination.

b. In the **URL** field, provide the webhook URL.
This is the default format of the URL to push events from ServiceNow Cloud Observability to the ServiceNow AI Platform instance. Use this format: `https//<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=lightstep`

c. Select **Payload Template**.

d. Select the **ServiceNow Event Management** option.

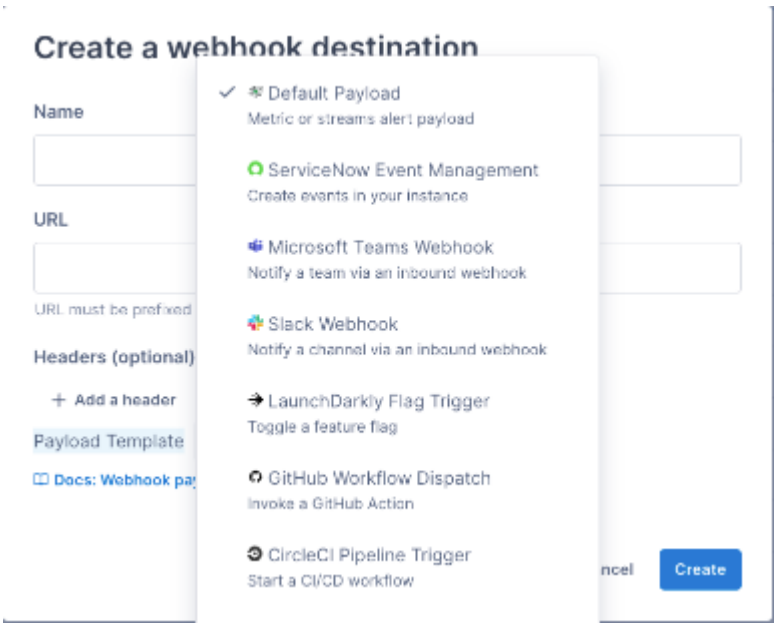

Note: In a ServiceNow instance, CI binding occurs only if the name of the service in the ServiceNow Cloud Observability platform value matches the name fields of the CI you created in the *cmdb_ci_service_auto* table.

e. If you selected the **Default Template** option in the previous substep, then select **New Header** and in the **Headers** field, enter the **service-name** and define its **value**.

Note: In a ServiceNow instance, CI binding occurs only if the **service-name** header value matches the name fields of the CI you created in the **cmdb_ci_service_auto table**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

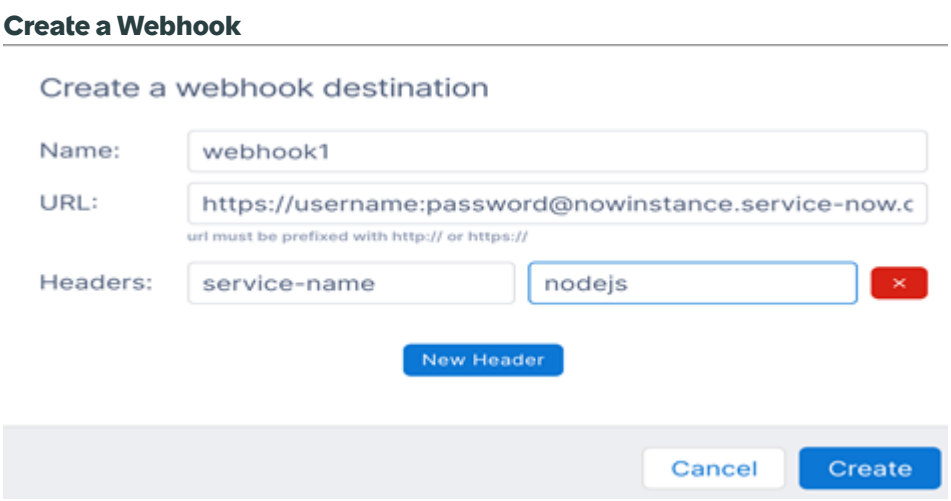

f. Select **Create**.

5. In the ServiceNow Cloud Observability platform, configure an alert.

a. From the navigation bar, select **Alerts**.

b. Select **Create an Alert** and enter an alert name.

c. In the Alert Configuration area, define the threshold.

d. According to the destination that you created, in the **Notification Rules** area, in the **Send Notifications to** field, select your notification destination type and search for the destination name.

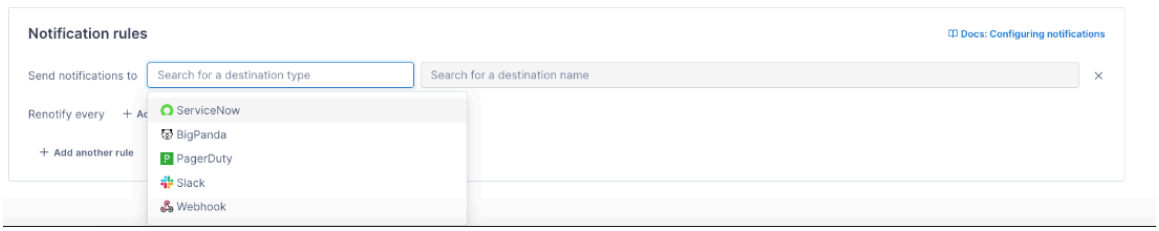

e. Select **Save**.
The default severity for stream-based alerts is **Critical**. You can change the default in the Push Connector Configuration area of an ServiceNow AI Platform instance. The valid severity values are:

- 1- Critical
- 2- Major
- 3- Minor
- 4- Warning
- 5- Info

Integrate Logicmonitor events

Integrate Logicmonitor with Event Management to send events into ServiceNow by adding a webhook using Basic Authentication, it will also be available with bi-directional functionality.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.
- The Event Management plugin must be installed on the ServiceNow AI Platform instance.
- Role required: evt_mgmt_integration

About this task

Configure the Event Management environment for the collection of events from Logicmonitor. In your Logicmonitor console, set your ServiceNow AI Platform instance as the rest endpoint using a standard webhook.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

Procedure

1. In the Logicmonitor console, create a custom HTTP Delivery integration.
   a. Navigate to **Settings** > **Integrations** and click **Add** > **Custom HTTP Delivery Integration**.
   b. In the Manage Custom HTTP Delivery Integration form, fill in the **Name** and **Description** field and select the **Use the same URL and data to notify on various alert activity** check box.
   c. Specify when to send a notification on alert activities by selecting the **Acknowledged**, **Cleared**, and **Escalated/De-escalated** check boxes as needed.

      Note: If you are using the bi-directional Connector update, do not select the **Acknowledged** check box. Doing so can trigger a loop back to the Alert and close the Alert record unintentionally.

   d. In the HTTP method menu, select **HTTP POST**.
   e. Fill in the **Username** and **Password** field, and enter a URL with the format as, `https://{Instance-name}.service-now.com/api/sn_em_connector/em/inbound_event?source=logicmonitor`.
   f. Under the Alert Data section, select **Raw** and select the format as **JSON**.
   g. Add the custom payload provided in the description of the connector.
   h. Click **Test Alert Delivery** to verify the connection or **Save**.
2. When you're finished creating the integration, create an escalation chain for the alert settings.
   a. Navigate to **Settings** > **Alert Settings** > **Escalation Chains** and click **Add**.
   b. In the Manage Testing chain form, fill in the **Name** and **Description** field and select the rate at which the alerts should be notified.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Under the Stages section, select the Add icon to add a stage and click **Save**.

d. In the new stage, select the Add icon to search for the user.

e. Add the integration as the Contact Method and click **Save**.

3. After creating the escalation chain, configure alert rules for the chain you created.

a. Navigate to **Settings** > **Alert Settings** > **Alert Rules**, and click **Add**.

b. On the Add Alert window, configure the conditions that initiates the alerts when conditions are violated.

c. Under the Escalation chain menu, select the escalation chain you created.

d. Click **Save**.

Note: To map the CI to the **cmdb_ci_vm_object** record instead of the **cmdb_ci_server**, enable the event rule created for Logicmonitor. The Logicmonitor's collector name, which comes in as a hostname in the payload, should be the same as the object_id of the virtual machine discovered in the **cmdb_ci_vm_object** record.

Event collection from Logicmonitor

The MID WebService Event Collector enables you to collect JSON formatted event messages from the Logicmonitor.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Ensure that you:

- Deploy and start the MID Server. See MID Server configuration .
- Configure and start the MID Web Server. See Configure the MID Web Server extension.
- Configure and start the MID WebService Event Collector. See Configure the MID WebService Event Collector Context.

Role required: evt_mgmt_admin

About this task

JSON formatted event messages are sent from Logicmonitor. The MID Server transforms the collected event messages by parsing them using the TransformEvents_MidLogicmonitor script include, located here: **Event Management** > **Integrations** > **PushConnectors**. In the Push Connectors page, click **Logicmonitor Mid Push Connector**.

The default format of the URL to push event messages from Logicmonitor to the MID Server is `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=TransformEvents_MidLogicmonitor`.

Variables in the default URL

Variable	Description
MID_Server_IP	IP address of the MID Web Server Extension.
MID_Web_Server_Port	Listening port of the MID Web Server Extension.
MID_Web_Server_User	Username for the MID Web Server Extension.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Variables in the default URL (continued)

Variable	Description
MID_Web_Server_Password	Password of the user of the MID Web Server.

Procedure

1. In the Logicmonitor console, create a custom HTTP Delivery integration.
   a. Navigate to **Settings** > **Integrations** and click **Add** > **Custom HTTP Delivery Integration**.
   b. In the Manage Custom HTTP Delivery Integration form, fill in the **Name** and **Description** field and select the **Use the same URL and data to notify on various alert activity** check box.
   c. Specify when to send a notification on alert activities by selecting the **Acknowledged**, **Cleared**, and **Escalated/De-escalated** check boxes as needed.
   d. In the HTTP method menu, select **HTTP POST**.
   e. Fill in the **Username** and **Password** field, and enter a URL with the format as, `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=TransformEvents_MidLogicmonitor`.
   f. Under the Alert Data section, select **Raw** and select the format as **JSON**.
   g. Add the custom payload provided in the description of the connector.
   h. Click **Test Alert Delivery** to verify the connection or **Save**.
2. When you're finished creating the integration, create an escalation chain for the alert settings.
   a. Navigate to **Settings** > **Alert Settings** > **Escalation Chains** and click **Add**.
   b. In the Manage Testing chain form, fill in the **Name** and **Description** field and select the rate at which the alerts should be notified.
   c. Under the Stages section, select the Add icon to add a stage and click **Save**.
   d. When you create the new stage, select the Add icon again to search for the user. Add the integration as the Contact Method and click **Save**.
3. After creating the escalation chain, configure alert rules for the chain you created.
   a. Navigate to **Settings** > **Alert Settings** > **Alert Rules**, and click **Add**.
   b. On the Add Alert window, configure the conditions that initiates the alerts when conditions are violated.
   c. Under the Escalation chain menu, select the escalation chain you created.
   d. Click **Save**.

> **Note:** To map the CI to the **cmdb_ci_vm_object** record instead of the **cmdb_ci_server**, enable the event rule created for Logicmonitor. The Logicmonitor's collector name, which comes in as a hostname in the payload, should be the same as the object_id of the virtual machine discovered in the **cmdb_ci_vm_object** record.

Integrate Oracle Cloud Infrastructure alarms

Integrate Oracle Cloud Infrastructure (OCI) alarms with Event Management to send events into ServiceNow by adding a https subscription using Basic Authentication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.
- Oracle Cloud Discovery must be performed to ensure that the created alerts are bound to the configuration items in the ServiceNow AI Platform. For more information, see Oracle Cloud Infrastructure Discovery.
- Role required: evt_mgmt_integration

About this task

The Oracle Cloud Notifications service enables you to set up communication channels for publishing messages using topics and subscriptions. When a message is published to a topic, the Notifications service sends the message to all of the topic's subscriptions.

When an Oracle Management Cloud alert is raised, becomes severe, or clears, you may want to be actively notified through an https custom URL to the ServiceNow instance.

The Oracle Cloud connector creates events in ServiceNow®.

You can configure the Event Management environment for the collection of alarms from Oracle Cloud by setting your ServiceNow AI Platform instance as the https subscription.

Procedure

1. In the Oracle Cloud console, set up a https subscription.

   a. On the Create Subscription form, under the Configure Subscription section in the Protocol list, select **HTTPS (Custom URL)**.

   Note: When the subscription is created, the password will not be shown in clear text.

   b. In the URL field, enter a URL with the format: `https://<username>:<password>@<instancename>.service-now.com/api/sn_em_connector/em/inbound_event/oraclecloud`.
   For example, `https://oracleevents:oracleevents@instance1.service-now.com/api/sn_em_connector/em/inbound_event/oraclecloud`.
   When the subscription is created with the ServiceNow instance URL, the Oracle Cloud Push connector will receive a payload that will confirm the subscription and is moved to an **Active** state in the Oracle Cloud console.

2. In the Oracle Cloud console, you can create an alarm definition and set the notification destination with the topic containing the https ServiceNow subscription.

   a. Navigate to **Observability & Management** > **Alarm Definitions**.

   b. Click **Create Alarm** or select an existing alarm that you want to set the topic containing the https ServiceNow subscription.

   c. In the Destinations section in the Topic list, select the topic with the https ServiceNow subscription.
   When an alarm is triggered in the Oracle Cloud console, the same alarm will be received in the Oracle Cloud Push connector in the ServiceNow instance and an event will be created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

Depending on the Oracle Cloud alarm resource, the alert will bind to the CI using the following Event Field Mapping:

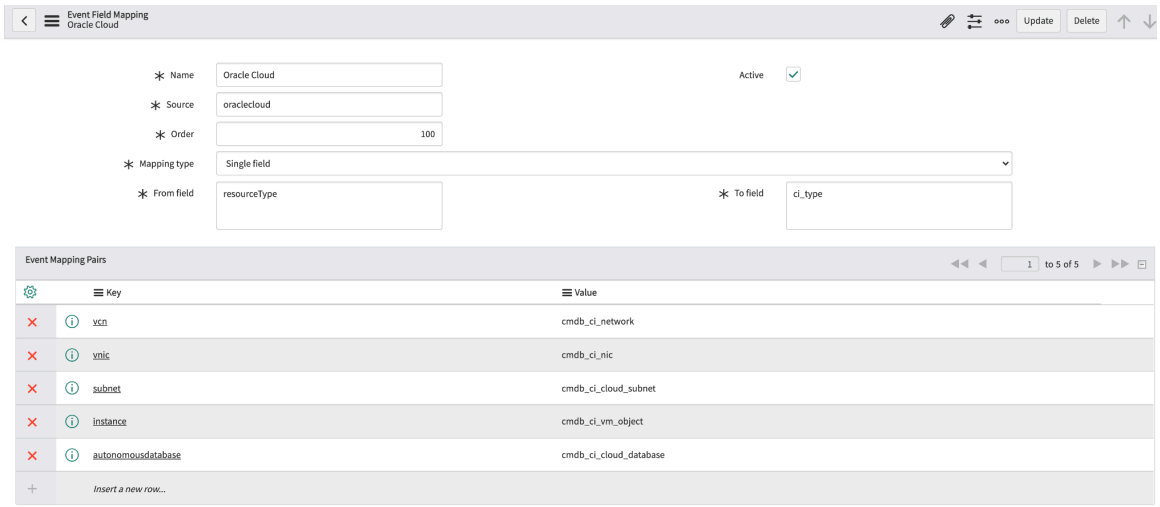

You can create other event rules or modify the Event Field Mapping to map to other resource types.

Once the https subscription is configured and an Oracle Cloud alarm triggers a message in Event Management the message:

- Authenticates the Oracle Cloud alarm message with the relevant ServiceNow user given in the https subscription.
- Extracts information from the original Oracle Cloud alarm message to populate the required event fields and inserts the event into the ServiceNow AI Platform database.
- Captures specified content in the **Additional Information** field of the event form.

Integrate Prometheus events

Integrate Prometheus with Event Management by adding a standard webhook in Prometheus's Alert Manager.

Before you begin

- Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.
- Verify Configuration Items for the hosts managed by Prometheus exist in ServiceNow. These CIs can be physical or virtual, and can be either manually created or discovered using IP discovery or Cloud discovery.
- Role required: evt_mgmt_integration

About this task

Configure the Event Management environment for the collection of events from Prometheus.

Starting from the Xanadu release, the OOTB (Out-Of-The-Box) event rules provided with the connector, which you have not previously used (i.e., neither activated, deactivated, nor modified), will now have the **Apply additional matching rules** check box set to true. Previously, this check box was disabled. This change allows you to execute more event rules or automation using the same filter conditions for the events.

Note: This feature applies only to active event rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Prometheus instance, add a webhook.
   a. Navigate to Alertmanager.yml and add a webhook entry under the recievers.
   b. Enter the endpoint uRL. For example, use `https://<instance-name>/service-now.com/api/sn_em_connector/em/inbound_event?source=prometheus`.
   c. Type `basic_auth` underneath *`http_config`*.
   d. Add a username and password under *`basic_auth`*, then add the created webhook name in the receiver under *`route`*.
2. In the Prometheus instance, create an alert rule.
   a. Create a file for the alert rule and add its' path inside prometheus.yml under *`rules_files`*.
   b. Create the alert rule as per the prometheus format using PromQL.
   c. Restart AlertManager and Prometheus to reflect the changes.

Result

Alerts start flowing from Prometheus into the Event Management plugin. The plugin extracts information from the original Prometheus alert message to populate the required event fields and inserts the event into the database. In your ServiceNow AI Platform instance, navigate to **All Events** to see the events.

If severity isn't received in the payload, the default severity is Major, which can be changed in the *`Push Connector Configuration`* section of **Push Connectors** > **Prometheus Push Connector**. The valid values of severity are **1- Critical**, **2- Major**, **3- Minor**, **4- Warning**, and **5- Info**.

Note:

To map hosts to the *`cmdb_ci_vm_object`* record instead of *`cmdb_ci_hardware`*, enable the **Prometheus VM Binding** event rule created for Prometheus. For further information about event rules, see Event rules.

Integrate Sentry events

Integrate Sentry with Event Management by adding a standard webhook in the Sentry platform.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Discovery for Sentry services are not supported. You must create a CI manually in the [cmdb_ci_service_auto] table to enable binding by completing the following steps:

1. From the navigation bar, go to **Configuration/Application Service**.
2. Provide the project name you are monitoring on the Sentry application.
3. Set the operational status to **Operational**.
4. Click **Next** then **Done**.

Role required: evt_mgmt_admin

About this task

Configure Event Management for the collection of events from Sentry by authenticating Sentry as a data source. In your Sentry platform, use a standard webhook to set your ServiceNow AI Platform as the rest endpoint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Sentry console, create an internal integration.

   a. Navigate to **Settings** > **Developer Settings**.

   b. Click **New Internal Integration**.

   c. In the Internal Integration Details section, on the Sentry Integration Details page, fill in the fields.

   Sentry Integration Details

Fields	Description
Name	Name of your integration.
Webhook URL	The default format of the URL to push events from Sentry to the ServiceNow AI Platform instance. For example, `https://<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=sentry`.
Alert Rule Action	Option to enable the integration. **Note:** When enabled, the integration becomes available in the Issue Alert and Metric Alert rules in Sentry. The alerts will be sent to the ServiceNow instance.
Schema	Schema for your UI components. This field should be left empty.
Overview	Description for the integration.
Authorized JavaScript Origins	Auth Tokens from the browser to enable the origins of the pages making the requests. This field should be left empty.

   d. In the Permissions section, under Issue & Event, specify the level of access you want. You can select **Read**, **Read & Write**, or **Admin**.

   e. In the Webhooks section, select the **Issue** check box.

   f. Click **Save Changes**.

2. In the Sentry console, create alert rules and configure them to receive alerts for issues from Sentry to the ServiceNow instance.

   a. Navigate to **Alerts** > **Create Alert Rule**.

   b. Select the alert type and click **Set Conditions**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The Issues alert type creates issue alerts. The other alert options create metric alerts.

c. On the Alert Rule page, create either an issue alert or a metric alert.

Alert Rule

Option	Description
Issue alert	**i.** Enter the environment, team, and alert name for your alert. **ii.** Under Set conditions, enter the **When** and **If** conditions that will trigger the alert. **iii.** For the **Then** condition, select **Send a notification via an integration** and select your internal integration from the list. **iv.** Set the action interval for the alert.
Metric alert	**i.** Add a filter to your transactions. **ii.** Select the function and time interval. **iii.** Set the thresholds that will trigger the alert. **iv.** Add the actions to be performed when the thresholds defined are met. **v.** Add your internal integration from the list to the list of actions to receive the alerts. **vi.** Add a rule name and team for the alert.

d. Click **Save Rule**.

Integrate Scout APM events

Enable the collection of events from Scout APM by authenticating Scout APM as a data source to integrate it with Event Management.

Before you begin

The Event Management plugin (com.glideapp.itom.snac) must be installed.

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Discovery for Scout APM applications is not supported. To enable binding, you must manually create a CI in a ServiceNow instance in the Service Instance [cmdb_ci_service_auto] table.

Role required: evt_mgmt_admin

Procedure

1. Enable Scout APM monitoring as an application service on your ServiceNow instance.

a. Navigate to **Configuration** > **Application Service**.

b. Provide the application name that you are monitoring.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Set the operational status to Operational.

d. Select **Next**

e. Select **Done**.

2. Log in to the Scout APM console.

3. Create a notification channel.

a. Navigate to **Organization** > **Alerts & Notifications**.

b. Select **New Webhook**.

c. Provide the name of the webhook and the webhook URL.
Use the following URL format: `https://<USERNAME>:<PASSWORD>@<INSTANCE_NAME>.service-now.com/api/sn_em_connector/em/inbound_event?source=scout`

d. Validate whether the webhook is operational by selecting **Send Test Alert**.
A new event appears on your ServiceNow instance in the Events [em_event] table with the data source as Scout APM.

4. In the Organization section, create a new notification group.

a. On the **Alerts & Notifications** page, navigate to **Organization** > **Notification Groups**.

b. Add a name for the group and select the webhook channel you created.

c. Select the **Create Notification Group** button.

5. In the application section, create alert conditions and configure them to receive alerts for notifications from Scout APM to the ServiceNow instance.

a. Navigate to **Applications** > **Alert Conditions**.

b. Select the type of alert condition.
For more information, see the Scout APM documentation.

c. Select **Create Alert condition**.

d. In the Alert Conditions table, select the notification group you created.
The alert condition is displayed on the Scout console.

6. Configure error notifications to receive alerts for errors from the Scout APM to the ServiceNow instance.

a. Navigate to **Application** > **Error Notifications**.

b. Add a name and description for the notification.

c. Select the notification group that you created.
The Scout APM notification channel is configured with the error notifications.

d. Select the **Add Notification Group** button.
For any new error notifications in the Scout APM application, a new event will be generated in the Events [em_event] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
You can view all notifications on your ServiceNow instance in the Events [em_event] table with the data source as Scout APM.

Integrate Sumo Logic events

Use the Sumo Logic push connector to integrate Sumo Logic with Event Management by adding a standard webhook in the Sumo Logic platform.

Before you begin
Ensure that the Event Management connector's plugin is installed on the ServiceNow AI Platform instance.

Role required: evt_mgmt_integration

About this task
Configure the Event Management environment for the collection of events from Sumo Logic by authenticating Sumo Logic as a data source. In the Sumo logic site, set your ServiceNow AI Platform instance as the rest endpoint using a webhook.

Procedure
1. In the Sumo Logic site, create a notification destination.
2. Go to Monitoring and select **Custom Webhook**.
3. Fill in the form with the Name, Description, and the URL.

   Use this URL format to push events from Sumo logic to the ServiceNow Instance `.https://<username>:<password>@<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=sumologic`

   Note: Use a credential having the evt_mgmt_integration role as the username and password.

   In the Payload section, use the following template:

```
{
    "type": "{{MonitorType}}",
    "node": "{{ResultsJSON._sourceHost}}",
    "metric_name": "{{Name}}",
    "description": "monitor Query: {{Query}}\n
 Trigger Condition:{{TriggerCondition}}\n
 Trigger Value:{{TriggerValue}}\n Trigger Time
 Range:{{TriggerTimeRange}}\n View Trigger
 Query:{{TriggerQueryURL}}\n View Monitor:{{QueryURL}} \n
 Results:{{ResultsJSON}}",
    "name": "{{Name}}",
    "short_description": "{{Description}}",
    "MonitorType": "{{MonitorType}}",
    "Query": "{{Query}}",
    "QueryURL": "{{QueryURL}}",
    "ResultsJson": "{{ResultsJson}}",
    "NumQueryResults": "{{NumQueryResults}}",
    "Id": "{{Id}}",
    "DetectionMethod": "{{DetectionMethod}}",
    "TriggerType": "{{TriggerType}}",
    "TriggerTimeRange": "{{TriggerTimeRange}}",
    "time_of_event": "{{TriggerTime}}",
    "TriggerCondition": "{{TriggerCondition}}",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    "TriggerValue": "{{TriggerValue}}",
    "TriggerTimeStart": "{{TriggerTimeStart}}",
    "TriggerTimeEnd": "{{TriggerTimeEnd}}",
    "SourceURL": "{{SourceURL}}",
    "alertResponseUrl": "{{alertResponseUrl}}"
  }
```

Log Monitors Type events can be created without CI binding since you're receiving the node value as "node": "{{ResultsJSON._sourceHost}}" in the event payload from Sumo Logic. You can create events but you can't bind them since you don't get them as host names for the log monitor type. A push connector parameter was created - "create_log_monitorType_events" with a value of False. If the param value is False, the payload is ignored, and "log" events aren't created. If you change the param to True, events are created, and then you must define event rules for CI binding.

Refer to the following ServiceNow Severity Mapping with Sumo logic TriggerType table:

Severity Mapping

Sumo Logic Trigger Type	Severity
Critical	Critical
Warning	Minor
MissingData	Warning
ResolvedCritical/ResolvedWarning/ ResolvedMissingData	Clear

Integrate Sumo Logic with REST API key token

Integrate using an API key to establish secure communication and automate data exchange via REST API. This simplifies integration, enabling seamless access to services and enhancing operational efficiency.

Before you begin

Role required: evt_mgmt_integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API Key**.

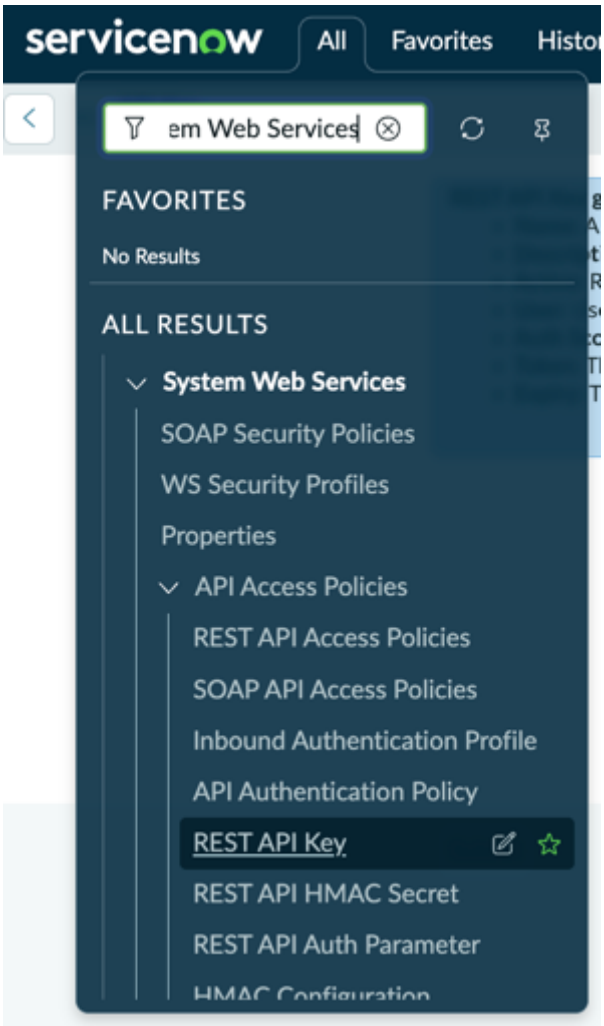

2. Select **New**.
3. On the form, fill the fields.

> **Note:** The user needs to be created with the evt_mgmt_integration role. Create a new one on the sys_user table to use here to facilitate the right access.

API Key

Field	Description
Name	Name to identify the REST API Key
Description	Description for the REST API Key.
Active	Status of the REST API Key.
User	User associated to the REST API Key. Use the look-up icon to select the user.
Auth Scope	Option to add auth scope to manage the authority of the REST API Key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Token	The REST API key generated by the ServiceNow AI Platform. Copy the key to use as part of the REST API call within the Header or Query parameter.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration.

4. Unlock Auth Scope and add: UserAccount.
5. Use the form menu and select **Save**.

   The system generates a token and saves it in the **Token** field. To see the token, use the lock icon and copy the contents display below the field. This is your query parameter (or header) value when your other system sends a REST API request to ServiceNow.

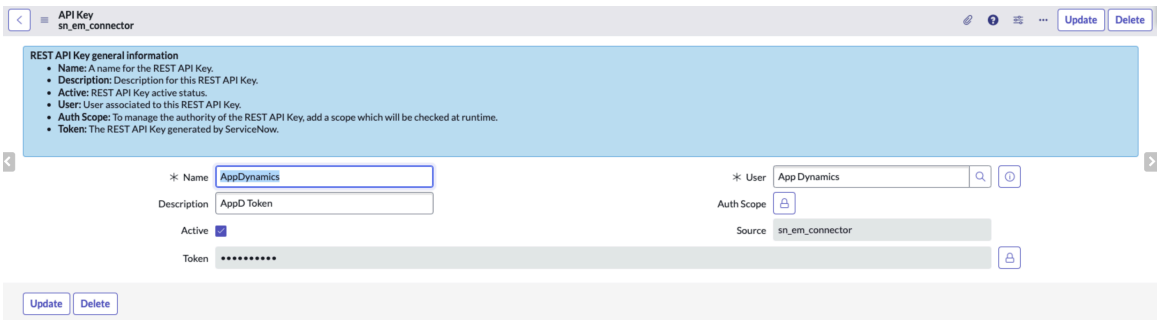

Integrate ThousandEyes platform events

Integrate ThousandEyes with Event Management by adding a standard webhook in the ThousandEyes console.

What authentication is used

There are two methods of authentication:

- OAuth authentication: Provides enterprise-grade authentication to keep your enterprise environment safe. Authentication is performed using ThousandEyes Monitor V1 or V2 access tokens. For more information, see Integrate ThousandEyes with OAuth authentication.
- Basic webhook authentication: Provides a basic standard of authentication, without the need for ThousandEyes Active Directory. This authentication can be especially useful for distributed small teams, such as SRE or DevOps teams. For more information, see Integrate ThousandEyes with basic authentication.

Integrate ThousandEyes with basic authentication

Integrate ThousandEyes with Event Management by adding a webhook in the ThousandEyes platform.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

Discovery for ThousandEyes services is not supported, so you must create a CI manually in the ServiceNow instance for event binding to work.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Confirm you have created a user with an **Identify Type** of **Machine** and the evt_mgmt_integration role.

Role required: evt_mgmt_admin

About this task

Configure the Event Management environment for the collection of events from ThousandEyes by authenticating ThousandEyes as a data source. In your ThousandEyes platform, set your ServiceNow AI Platform instance as a rest endpoint using a webhook.

Procedure

In the ThousandEyes platform, create a notification destination.

a. From the navigation bar, navigate to **Alerts** > **Alert Rules**.

b. Click on any existing alert rule and select the **Notification** tab.

c. Select **Edit Webhook** and enter the following information.

Edit Webhooks

Field	Description
Name	Enter a name for the webhook.
URL	Provide the webhook URL. **Note:** The default format of the URL to push events from ThousandEyes to the ServiceNow instance is `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=thousandeyes`.
Auth Type	Select **Basic**.

d. Click **Add New Webhook**.

Result

The default severity for the ThousandEyes alert is **Major -2**, which can be changed in the *Push Connector Configuration* section of **Push Connectors** > **ThousandEyes Push Connector**. The valid values of severity are **1- Critical**, **2- Major**, **3- Minor**, **4- Warning**, and **5- Info**.

The default timezone of the ThousandEyes alert is GMT, if you want to send the alert in a different timezone, it can be changed in the *Push Connector Configuration* section of **Push Connectors** > **ThousandEyes Push Connector**. The valid value format is **GMT+01:00**, **GMT +02:00**, **GMT-05:30**, **GMT-02:00**, etc.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integrate ThousandEyes with REST API key token

Integrate using an API key to establish secure communication and automate data exchange via REST API. This simplifies integration, enabling seamless access to services and enhancing operational efficiency.

Before you begin

Role required: evt_mgmt_integration

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API Key**.

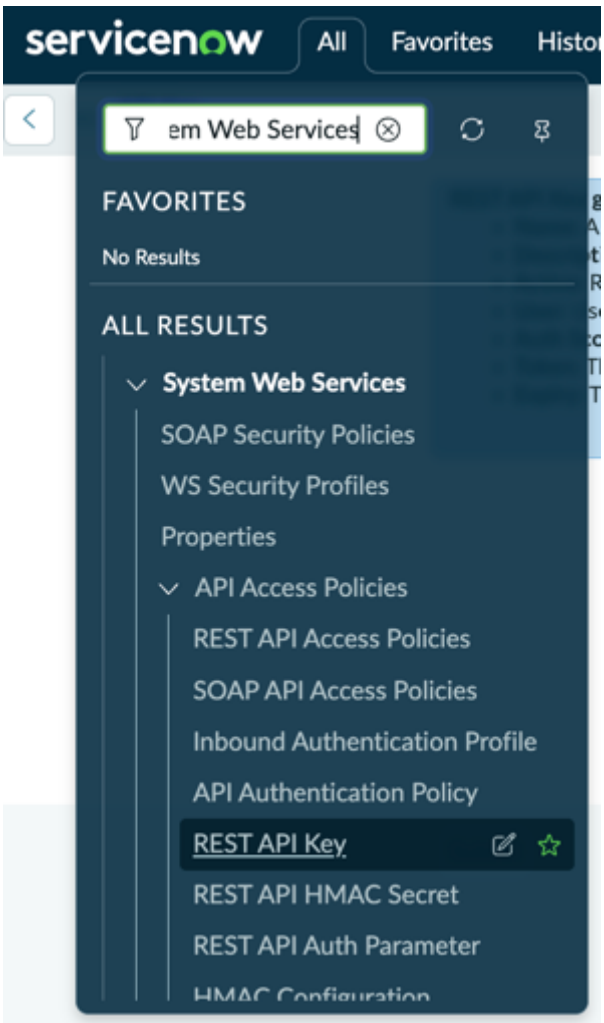

2. Select **New**.
3. On the form, fill the fields.

> **Note:** The user needs to be created with the evt_mgmt_integration role. Create a new one on the sys_user table to use here to facilitate the right access.

API Key

Field	Description
Name	Name to identify the REST API Key
Description	Description for the REST API Key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Active	Status of the REST API Key.
User	User associated to the REST API Key. Use the look-up icon to select the user.
Auth Scope	Option to add auth scope to manage the authority of the REST API Key.
Token	The REST API key generated. Copy the key to use as part of the REST API call within the Header or Query parameter.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration.

4. Unlock Auth Scope and add: UserAccount.
5. Use the form menu and select **Save**.

   The system generates a token and saves it in the **Token** field. To see the token, use the lock icon and copy the contents display below the field. This is your query parameter (or header) value when your other system sends a REST API request to ServiceNow.

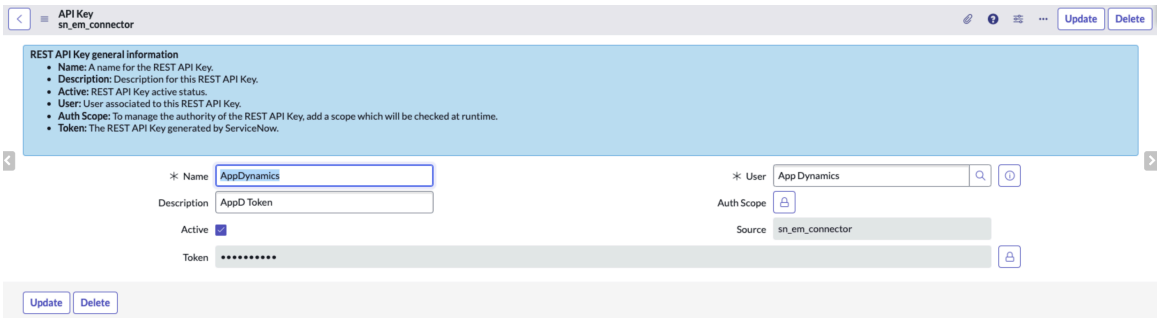

Result

Now you are ready to start sending events. In your monitoring tool you can send events either with the header or the query `param: x-sn-apikey`. For example:

```
https://[INSTANCENAME].service-now.com/api/sn_em_connector/em/inbound_event?source=[SOURCENAME]&sys_id=[SYSID OF PUSH CONNECTOR INSTANCE]&x-sn-apikey=[AUTH TOKEN]
```

Integrate ThousandEyes with OAuth authentication

Create credentials in the instance and configure the ThousandEyes webhook.

Before you begin

Role required: web_service_admin

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry** and then select **New**.
2. Select **Create an OAuth API endpoint for external clients**.
3. Copy the automatically generated field **Client ID**.
   The **Client ID** is used when configuring the ThousandEyes SeviceNow integration.
4. Unlock the **Redirect URL** field and assign the value: **https://app.thousandeyes.com/webhooks-oauth-callback/**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Configure the Access and Refresh token lifespans by providing a value in seconds.
6. In ThousandEyes, navigate to **Manage** > **Alert Rules**.
7. Open an alert rule and select the **Notifications** tab.
8. Under **Webhooks**, select **Configure Webhooks** or **Edit webhooks**, and then select **Add New Webhook**.
9. Complete the following fields:

   Name

   Enter a name for the webhook.

   URL

   ```
   https://<instance-name>.service-now.com/api/
   sn_em_connector/em/inbound_event?source=thousandeyes
   ```

   Auth Type

   Select **OAuth**.

   Auth URL

   ```
   https://<instance-name>.service-now.com/oauth_auth.do
   ```

   Client ID

   Paste the Client ID generated in the ServiceNow instance.

10. Select **Get Token**.
11. When the ServiceNow authorization page opens, sign in and select **Allow**.
12. Test and save the webhook.

Integrate Panopta as a data source

Integrate the Panopta cloud-based monitoring solution with Event Management. To add Panopta as a data source, configuration is required in Panopta.

Configure Panopta

To prepare Panopta as a datasource, follow the procedure described in this link: How do I integrate with ServiceNow.

When the Event Management integration is configured in Panopta, add the integration to any of your Panopta Alert Timelines so that the ServiceNow instance receives incident events accordingly.

Optional Event Management configuration

Configure an alert action rule for launch in context of Panopta. For more information, see Launch web application from alert and Create or edit an alert action rule.

Related topics

Configure Event Management connectors

Integrate Event Integration (EIF) format event connector

Use the push connector that allows events to be forwarded from products, generally from IBM, that support the Event Integration (EIF) format to a MID Server.

Before you begin

Ensure that the Event Management Connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Ensure the MID Server is assigned with `EventManagement` or `ALL applications`.

Role required: evt_mgmt_integration

About this task

Configure the Event Management environment for the collection of events from products that support the EIF format to a MID Server.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **EIF listener context**.
2. Click **New**.
3. Create a new EIF listener context by completing the EIF listener context form:

   a. On the form, fill in the fields.

EIF listener context form

Field	Description
Name	Unique name for the EIF listener context.
Short Description	Short description of the MID Server extension execution context.
Port	Port for the EIF listener to receive EIF events. The default value is **9998**. However, if the port is already in use, you can use any other free port instead.
Status	The status of the EIF listener. This field is blank until the listener is run. After the listener has run, the value is automatically populated to one of the following statuses: ▪ Started: The listener is running. ▪ Stopped: The listener is not running. ▪ Offline: The MID Server is down. ▪ Error: The listener failed with an error.
Execute on	The MID Server location for running the listener.
MID Server	The name of the designated MID Server when you select **Specific MID Server** in the **Execute on** field.
Executing on	The name of the MID Server on which the extension is running. This field is automatically populated with the name of the MID Server even if it's down. If you stop the extension, this field is empty.

   b. On the form, you can select the related links on the EIF listener context form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Commands available on the EIF listener context form

Related Link	Description
Start	Select this option if the listener is not running. The listener is started on the configured MID Server and port number.
Stop	Select this option to stop the running listener on the configured MID Server.
Restart	Select this option to stop then start the listener on the configured MID Server and port number.
Update Parameters	Select this option if you want to make changes to the listener configuration of the current running listener. If the parameters are valid, the listener disconnects and reconnects with new parameters.
Test Parameters	Select this option to test if the port is a valid number.

4. When you're finished, right-click the form header and select **Save**.
5. Click **Update**.
6. If you want to check the incoming EIF traffic, complete the following steps:
   a. Navigate to **MID Server** > **Servers**.
   b. Select the MID Server the EIF extension is running.
   c. Under the **Configuration Parameters** tab, add the **mid.log.level** properties.

What to do next

If the MID Server fails to start, you can review the MID Server log errors.

EIF events warning severity

If the EIF event payload has a warning severity, it will be mapped differently in the ServiceNow instance.

EIF events severity	ServiceNow events severity mapping
UNKNOWN	Warning
Warning	Minor
Minor	Major
Critical	Critical
Fatal	Critical

Create a push connector configuration parameter

Some connectors have no parameters that are shipped out of the box. Therefore, it’s necessary to create and add the parameter to the push connector configurations list in the form layout before adding the new parameter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: evt_mgmt_admin, personalize_form

Procedure
1. Navigate to **All** > **Event Management** > **Integrations** > **Push Connectors**.
2. Open a push connector.
3. On the **Push Connector** form, right-click on the form header and select **Configure** > **Form Layout**.
4. In the **Available** list, search for **Push Connector Configurations** > **Push Connector**, and move it to the **Selected** list.
5. Select **Save**.
6. In the Push Connector Configurations [sn_em_connector_push_config] table, double-click on **Insert a new row** to add a new push connector configuration parameter.
7. Select **Update**.

Result
Push Connector Configurations list added to push connector form layout. You can now create push connector parameters.

Configure a push connector instance

Configure a push connector instance to receive the events from an external system.

Before you begin
Ensure that you've defined a push connector definition.

Role required: evt_mgmt_admin

About this task
You can use a connector instance to receive the events coming from external sources. You can create the following type of push connector instances:

- Instance type
- MID type

Procedure
1. Navigate to **All** > **Event Management** > **Integrations** > **Push Connector Instances**.
2. Select the connector instance that you want to configure, or select **New** to create a push connector instance.
3. On the **Push Connector Instance** form, fill in the fields.

   For more details, see Push connector instance form.
4. Select **Submit** or **Update**.

Result
The push connector instance form is saved and now you can receive the events from an external system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure event collection for SNMP traps

The SNMP listener runs on the MID Server, which acts as a collection endpoint for SNMP traps. The MID Server sends the traps to the ServiceNow instance for further processing as an event by Event Management.

Before you begin

Role required: agent_admin

Note:

- To view the MID SNMP Trap Listener in the UI: agent_admin
- To perform other operations, any one of the following roles is required: agent_admin, evt_mgmt_integration, or evt_team_operator

Ensure that the MID Server is assigned with either `EventManagement` or `ALL applications`.

About this task

Use the **MID SNMP Trap Listeners** to listen for SNMP traps. In this configuration, the specified MID Server acts as a collection endpoint for SNMP traps.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **MID SNMP Trap Listeners**.
2. Select **New**.
3. Complete the SNMP Trap Collector Context form.

Option	Description
Name	Unique name for the SNMP Trap Collector.
Short description	Short description of the MID Server extension execution context.
SNMP version	Specify the SNMP version: ○ v1 and v2c ○ v3
Minimum Security Level	The valid security level values are: ○ None - No security level value set. ○ noAuthnoPriv - Communication without au thentication and privacy ○ authNoPriv - Communication with authenti cation and without privacy. ○ authPriv - Communication with authentica tion and privacy. **Note:** This field appears only when `v3` is selected in the **SNMP version** field.
UDP Port	UDP port for SNMP Trap Collector to receive SNMP Traps. The default value is *1,162*. However, if the port is already in use, you can use any other free port instead.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Status	The status of the trap listener. This field is blank until the listener is run. After the listener has run, the value is automatically populated to one of these statuses: ○ **Started:** The listener is running. ○ **Stopped:** The listener is not running. ○ **Offline:** The MID Server is down. ○ **Error:** The listener failed with an error.
Execute on	The MID Server location for running this listener. Select from specific MID Server. The recommended setting is **Specific MID Server**. Network hardware typically has to be configured to send to a specific IP address.
MID Server	Specify the name of the designated MID Server when you select **Specific MID Server** in the **Execute on** field.
EngineID	Displays only when **SNMP version=v3**. ○ When sending an SNMP Inform message, include the default value of this field in the message. ○ When sending an SNMP Trap message, include a value other than the default value of this field in the message.
Executing on	The name of the MID Server on which the extension is running. This field shows the name of the MID Server even if the MID Server is down. If you stop the extension, this field is empty. This field is automatically populated.

4. Right-click the form header and select **Save**.
5. In the **Credentials for SNMPs** section, specify the required SNMP v3 credentials that the SNMP listener uses for authentication.

   Note: The **Credentials for SNMPs** section appears only when `v3` is selected in the **SNMP version** field.

6. In the **Related Links** section, click **Test parameters**.
7. Click **Update**.
8. **Optional:** If the MID Server fails to start, review the MID Server log errors.

   (Optional) When working with object identifiers (OIDs), a trailing OID counter can affect event rule matching. To remove the trailing counter value from translated OIDs, navigate to **MID Server** > **Properties** and set the **mid.em.snmp_oid_key_counter_cut.enabled** property to

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

true. The property removes the trailing numeric value only when it can be identified as the OID counter. If additional numeric segments remain in the translated OID (for example, because MIB translation is incomplete), the counter is not removed. The system cannot determine which numeric values are significant and which represent the counter.

What to do next

In **Related Links**, these actions may be run against the SNMP traps listener:

Commands available in the SNMP trap listener form

Related Link	Description
Start	If it is currently not running, starts the listener. The listener is started on the configured MID Server and port number.
Stop	Stops the running listener on the configured MID Server. If the listener is not running, no action is taken.
Restart	Stops, then starts the listener on the configured MID Server and port number.
Test parameters	Tests that the port is a valid number.
Update parameters	If you make changes to the listener configuration, use this option to update the parameters of the currently running listener. First, the parameters are tested for validity. If the parameters are valid, the listener disconnects and reconnects with the new parameters.

Note: Updates to the SNMP trap listener take effect only after clicking **Restart** or **Update parameters**.

Related topics

SNMP probe MIB modules

Configure the SNMP traps listener to receive OEM traps

Configure the SNMP traps listener to receive OEM traps

Configure the SNMP trap listener to receive traps from Oracle Enterprise Manager (OEM).

Before you begin

Role required: evt_mgmt_admin

About this task

Event Management can process Oracle Enterprise Manager Cloud Control SNMP traps as events. For the SNMP trap listener to receive traps from OEM, you must designate the MID Server that runs the SNMP trap listener as a recipient of the trap. Upon receiving an OEM trap, the MID Server sends the trap to the ServiceNow instance for further processing as an event by Event Management.

The OEM 12c Trap event rule, Oracle EM MIB, and associated alert action rules are provided out-of-the-box, as well as these alert action rules:

- `Oracle EM Launch Target Status`
- `Oracle EM Launch View Event`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **MID SNMP Traps Listener**.
2. Click **New**.
3. Fill in the fields, as described in Configure event collection for SNMP traps.
   For the **SNMP version** field, select: `v1 and v2c`.

   The SNMP version field on the trap listener form has two options: v1 and v2c (default) and v3.The OEM 12c Trap event rule specifically processes SNMPv1-style traps (it relies on v1 fields like enterprise, generic_trap, and specific_trap). So selecting v1 and v2c is correct for OEM.

4. Verify the setup.
   - Check the listener is running: After submitting the trap listener record, look at the MID Server log for a message like "TrapListenerExtension: started. Listening on port <port>". This confirms the MID Server is listening.
   - Verify events are being created: Navigate to Event Management > All Events (or the em_event table) and filter by source = Oracle EM. If OEM is sending traps and the listener is working, events should appear here.
   - Check the event rule matched: Open an event and verify the event_class field shows "Trap From Enterprise 111" and the OEM-specific fields (like oraEMNGEvent_Severity, oraEMNGEventHost_Name, etc.) are populated in additional info.
   - Check alert creation: Navigate to Alerts and verify alerts are created from the OEM events, with the "View Target Status" and "View Event" launch-in-context actions available.

   Some basic troubleshooting:
   - Port binding error: If the MID Server log shows "IOException starting the SNMP listener", the port is likely already in use. Choose a different port or check for conflicts.
   - No events appearing: Enable MID Server debug logging — the listener logs full trap details (PDU, variable bindings, OIDs) when debug is on. This helps confirm whether traps are reaching the MID Server at all.
   - Authentication failures (v3 only): The log will show "Authentication failure" messages with the sender address and SNMP USM error code. Common causes: wrong credentials, mismatched engine ID, or security level mismatch.
5. Click **Submit**.

> **Note:**
>
> Ensure all required SNMP MIB files are uploaded on the instance. For more information, see Load a MIB module.
>
> .

What to do next

In Oracle Enterprise Manager Cloud Control, configure the MID Server as a trap listener target.

Configure SNMP Trap collection for high availability

For SNMP traps, the MID Server requires failover configuration for the trap listener.

You can configure two MID Servers for failover. Because the SNMP Trap listener on the MID Server receives inbound traffic on the IP address and port of the MID Server, you configure each MID Server to receive the same SNMP traps. Only one MID Server is active at any point in time, therefore Event Management does not receive duplicate traps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure message keys to spread SNMP object identifiers

By default, most SNMP trap events are processed by a single Event Management processing job. This can negatively effect event processing. Configure message keys on the MID Server to ensure that more than one processing job is invoked, ensuring optimal SNMP trap performance.

Before you begin
Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **MID Server** > **Servers**.
   The **MID Servers** page opens.
2. Select the relevant MID Server and on the MID Server page, select the **Configuration Parameters** tab.
3. Click **New**.
4. In the **Parameter name** field, select *`mid.snmp.event.oid.keys`*.
5. In the **Value** field, enter a comma separated list of OIDs to be used for the message key.
6. Click **Submit**.

   Note: The OID trap counter may cause problems activating event rules. To remove the trap counter from the OID, navigate to **MID Server** > **Properties** and set the **mid.em.snmp_oid_key_counter_cut.enabled** property to **true**.

7. To set conditions in a script that determine the message key to be used based on the SNMP traps received:
   a. Navigate to **MID Server** > **Script Includes**.
   b. Locate and select **SNMPTrapKeyFilter**.
   c. In the **Script** field, set conditions for which OIDs are to be used based on the SNMP trap event received.
   d. Select the **Active** check box to activate the script.
   e. Click **Update**.

Configure event collection from email

Configure an inbound email action to send email notifications when events and alerts are triggered.

Before you begin
Role required: evt_mgmt_admin

About this task
When you add an inbound email action to send email, Event Management uses inbound email actions to process inbound email in the same manner as any external event. The email content is used to generate and update events, use rules to generate alerts, or change event severity. The severity is updated after the impact calculation and stored in the em_impact_status table.

Procedure

1. Navigate to **All** > **System Properties** > **Email Properties**.
2. In the **Inbound Email Configuration** section, in the **Email receiving enabled** option, select **Yes** to enable the collection of events from email.
3. Navigate to **System Policy** > **Email** > **Inbound Actions**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In the list of Inbound Email Actions records, search for and open the default **create event** form.

a. On the **Inbound Email Actions** form, specify these values.

Field	Value
Active	Select the check box to enable the email notification.
Target table	Select **Event [em_event]**.
Action type	Select **Record Action**.

b. Select the **Description** tab and enter a description for the email message.

c. Select the **Actions** tab and select these **Field Actions**.

- **[Severity] [To] [<the event severity>]**
- **[Type] [From email] [<portion of the email>]**

  For example, you can select **Subject** as the portion of the email to use.

4. To customize the inbound parameters and map the text to event variables, add a **Script**.

Example
For example:

```
current.source = 'email';
current.event_class = 'email';
current.description = email.body_text;
current.time_of_event = new GlideDateTime();
current.insert();
```

5. From the instance, send an email that contains matching data, and then confirm that the event or alert information appears in Event Management.
For more information sending email messages to create an instance, see Inbound email actions.

Note: Ensure that the user who sends the email has the role **evt_mgmt_integration**.

6. Select **Submit**.

Result

The inbound email is sent to the em_event table and regular Event Management processes continue, for example, event rules and alert management rules run.

Event forwarding

Accelerate the event processing testing life cycle by forwarding a stream of events from your ServiceNow production environment to your non-production environment.

Use event forwarding to forward events from one instance to another, for example, from a production to a non-production instance. Event forwarding enables you to test and evaluate event rules, event field mappings, alert management rules, and alert correlation. This approach has no impact on your production environment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Not all monitoring sources can send events to multiple target instances. In such cases, you can configure a scheduled job to periodically forward the event stream. The scheduled job forwards events from your ServiceNow instance that is connected to the monitored source to other instances.

Set up event forwarding

Create an event forwarding configuration record to enable events to flow from one ServiceNow instance to another instance. Forwarding events to multiple target instances requires creating separate configuration records for each target instance.

Before you begin

You must have a credential with the evt_mgmt_integration role, which enables you to create events in the target instance. If you don't have this credential, see Create basic auth server credentials for information on how to create it.

Role required: evt_mgmt_admin

About this task

The source instance forwards events to a target instance. The source instance must have a user with the evt_mgmt_integration role, a password for the credential, and a new event forwarding record. The target instance receives events from the source. The target instance must have a credential with the same user name and password as the source instance.

The event forwarding record contains details about the target ServiceNow instance URL to which events are to be forwarded along with the credentials to connect it.

You must create a separate configuration record for each target instance to forward events to multiple target instances. Events are then sent to each target instance with an active event forwarding configuration.

Procedure

1. On your source instance, navigate to **All** and enter `sn_em_connector_event_sync_configuration.list` in the filter navigator.
2. On the **Event Sync Configurations** page, select **New**.
3. On the **Event Sync Configuration** page, in the **Instance url** field, enter the target ServiceNow URL to receive events in the format `https://<your_instance>.service-now.com`.
4. In the **Credentials** field, provide the credentials for the target instance by either entering them or searching for them.
   - To enter your credentials, start typing and then select the credential from the drop-down list.
   - To search for your credentials, select the Lookup using list icon ( ) to select your credential from the credentials lookup table.
5. Select **Save**.
6. Select **Test Connection** to validate the connection to the target ServiceNow instance.
   - If the connection succeeded, a confirmation message displays.
   - If the connection didn't succeed, check your credentials and the instance URL.
7. When validated, select the **Active** check box to activate the event forwarding to the target instance and select **Submit**.
   To forward events, you must also activate the **Event Management - Event Forwarding Job** scheduled job.
   The configuration is created and is listed on the **Event Sync Configurations** page.
8. For multiple targets, repeat the steps for a separate configuration record for each instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Periodically run an event forwarding job

Periodically run an event forwarding job

Activate the event forwarding job to periodically send events to all target instances with active event forwarding configurations.

Before you begin
Role required: evt_mgmt_admin

About this task
Event forwarding requires both an active event forwarding configuration record and an active **Event Management - Event Forwarding Job** scheduled job.

You can customize the *event - forwarding* job behavior through related system properties. For more information, see Event forwarding properties.

Note: To minimize the performance impact, the scheduled job sends events not older than the latest hour when it ran for the first time after setting up event forwarding.

Procedure

1. Navigate to **All** > **Scheduled jobs** > **Event Management - Event Forwarding Job**.
2. On the **Scheduled Script Execution** page, select the **Active** check box.
3. In the **Repeat Interval** field, set the appropriate time interval for the job to run periodically.

   The default value is for the job to run every minute.
4. Select **Update**.

Result
The event forwarding job runs periodically based on the repeat interval that you set. Events are sent to all target ServiceNow instances that have active event forwarding configuration records.

If the admin account is disabled, configure the user of the scheduled job. By default, it's admin, but in this case, it must be changed.

Note:

If the event forwarding to the instance failed due to any reason, the **Last Event Forwarding status** field displays `Failed` with a corresponding error message displayed in the **Last Event Forwarding error** field. The target instance generates a self-health monitoring event during the next successful event forwarding indicating the failure of event forwarding for the specified time range.

Create basic auth server credentials

Create credentials to access a ServiceNow instance.

Before you begin
Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Select **New**.
3. Select **Basic Auth Credential**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. On the form, fill in the fields.

Basic Auth Credentials form

Field	Description
Name	Enter a unique and descriptive name for this credential.
User name	Name to identify the user.
Password	Password to use this credential.
Active	Option to enable the use of this credential.
Order	The order (sequence) in which the platform tries this credential while it attempts to log in to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order. Default value: `100`

5. Select **Submit**.

Result
The Basic Auth credential for use with the server instance is created.

Processing Events

Event processing is the process of taking events or streams of events, analyzing them and taking automatic action. The process includes viewing events, event binding, event rules and event field mapping.

An event is a notification from one or more monitoring tools that indicate something of interest has occurred, such as a log message, warning, or error. Events are processed and alerts are generated based on the following factors:

Event collection configuration via a MID Server, script, SNMP trap collector, or email message.

An Event rule, alert binding, or event field mapping configuration for processing events from various sources. When a package of events that contains events with a custom (non-supported) state is processed, the events with the custom state are not processed.

For more information, see Exploring Event Management.

View events

Event Management tracks individual events to manage external systems. These events are notifications from monitoring tools indicating occurrences of interest, like log messages, warnings, or errors. Event Management gathers events from external sources and stores them in the Event [em_event] table, offering a list of raw incoming events.

Before you begin
Role required: evt_mgmt_admin, evt_mgmt_operator, evt_mgmt_user, or evt_mgmt_integration

Note: Ensure that you do not change the **Classification** field value in event [em_event] tables, either manually, by script, or by event rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The event monitoring tool generates the values of the source and resource fields. Event Management implementers can define event types and register nodes to help uniquely identify incoming events and create alerts for the specific needs of the enterprise. Event Management uses this information to determine whether to create an alert or update an existing one.

An event source may generate duplicate events with the same identifying information. For events with the same identifying information, Event Management uses the time interval between events to determine whether the events represent an existing issue or new issue.

Additional fields should be included in the **Additional information** field of the event. Do not add additional fields to an event by adding a custom field to the event [em_event] table. For more information about how to include additional fields in events, see Custom alert fields.

Note: Avoid writing business rules for event [em_event] tables, as they can result in performance degradation.

Procedure

1. Navigate to **All** > **Event Management** > **All Events**.

   The All Events list displays the following columns.

All Events list

Column	Description	Populated by
Time of event	The time that the event occurred, in the network node time zone. The system supports time resolution in seconds.	External event monitoring tool
Source	Event monitoring software that generated the event, such as SolarWinds or SCOM. Maximum length: 100 characters.	External event monitoring tool
Description	Reason for event generation. Shows extra details about an issue. For example, a server stack trace or details from a monitoring tool. Maximum length: 4000 characters.	External event monitoring tool
Node	Node name, fully qualified domain name (FQDN), IP address, or MAC address that is associated with the event, such as IBM-ASSET. Maximum length: 100 characters.	External event monitoring tool
Type	The metric type to which the event is related, such as Disk or CPU, which is used to identify an event record from which alerts are created. Maximum length: 100 characters.	External event monitoring tool
Resource	Node resource that is relevant to the event. For example, Disk C, CPU-1, the name of a process, or service. Maximum length: 100 characters.	External event monitoring tool
Metric Name	Unique name that describes which metrics are collected and for which this alert has been created.	External event monitoring tool
Message Key	Unique event identifier to identify multiple events that relate to the same alert. If this value is empty, it is generated from the **Source**, **Node**, **Type**, **Resource**, and **Metric Name** field values. Maximum length: 1024 characters.	External event monitoring tool
State	The status of the event:	Event

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description	Populated by
	○ **Ready**: Event has been received and is waiting to be processed. ○ **Queued**: Event is queued by the event processor job. ○ **Processed**: Event was successfully processed. ○ **Error**: Failure occurred while processing the event. For example, the event collection method or event **Severity** is blank. ○ **Ignored**: Value is not in use.	
Severity	Mandatory. The options are typically interpreted as follows: ○ **Critical**: Immediate action is required. The resource is either not functional or critical problems are imminent. ○ **Major**: Major functionality is severely impaired or performance has degraded. ○ **Minor**: Partial, non-critical loss of functionality or performance degradation occurred. ○ **Warning**: Attention is required, even though the resource is still functional. ○ **OK**: An alert is created. The resource is still functional. ○ **Clear**: No action is required. An alert is not created from this event. Existing alerts are closed.	Event
Alert	If an alert was created as a result of the event, this field contains the unique ID that Event Management generates to identify the alert.	A matching event rule or event field mapping

2. To review a single event, click the required row in the **Time of event** column.
3. To review the event rules that processed the event, in the **Related Links** section click **Check process of event**.
   For further ways to see which event rules were applied to the event, you can display the `Event rule` and `Processing Notes` fields in the Events list.

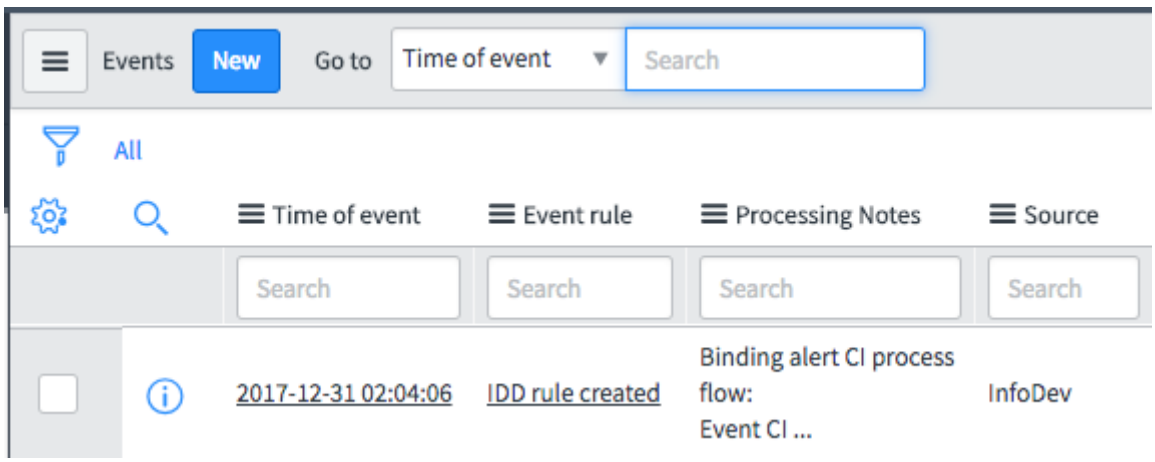

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Optional event fields to display the applied event rule

Column	Description
Matching Event rule	The event rule that matches the current event rule definitions. This information is primarily used to create event rule recommendations for events that currently are not covered by rules. If event rules have been modified, this information can be retrieved from the Processing Notes field, described below. **Note:** This field does not show the history of event rules that were applied to the event.
Processing Notes	Information related to the binding to CI type and the event rules that were applied to the event.

4. To create an event rule for an event that is not associated with an event rule:

   a. Click the event.

   The event opens in the Event form.

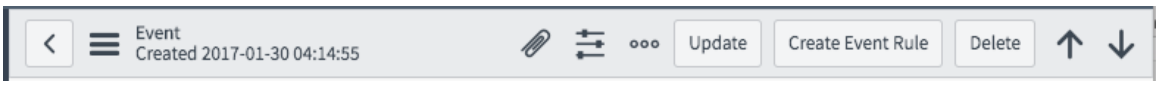

   b. Click **Create Event Rule**.
   The fields of the new Event rule are populated with data from the event. You can use the event information, for example, when composing fields for alerts.

Team-based integrations in Event Management

Team-based integrations empower teams to optimize event processing within Event Management to enhance efficiency and operational effectiveness.

Team-based integrations provides the following capabilities:

- Assign teams to particular connectors, providing a structured approach to integration efforts.
- Enable each team to define and apply custom rules, enabling a more targeted and straightforward handling of event processing.
- Route alerts automatically to the appropriate team members responsible for the integrated systems based on integration ownership and configuration item (CI) ownership.
- Enable seamless integration and automation through the automatic assignment of alerts to appropriate team members based on integration ownership and configuration item (CI) ownership.

Customize event rules, event mapping, and the precedence of team rules for team-based integrations. For more information, see Team-based integration properties.

Related topics

Event rules

Event field mapping configuration

Alert assignment groups for teams

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event rules

Use event rules to generate alerts for tracking and remediation. Event rules are stored in the Event Rule [em_match_rule] table. Configure and customize event rules to manage events and alert generation.

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

Event rules do not change the event records in the Event table. Changes to event data are stored in the ServiceNow instance memory.

For a workspace-based interface for Event Management automation, see Alert automation in Service Operations Workspace for ITOM.

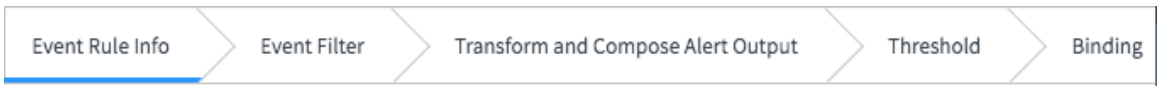

Use the Event rule designer to create and configure event rules. You can use the default event rules or event rules that you have created to:

- Apply an event rule filter to determine whether the rule applies to an event.
- Apply a transform with **Event Match Fields** and optional **Event Compose Fields** to format the alert text.
- Apply a threshold to create custom alerts for rapidly recurring events.
- Automatically create or close alerts.
- Bind alerts to CI information from the CMDB.
- Run multiple sequential rules defined for the same event by selecting the **Apply additional matching rules** check box. The event rules run in ascending order as defined in the **Order** field. Event rules applied to assignment groups only run after the global rules have run.

Create or edit an event rule

You can create event rules to generate alerts for tracking and remediation. Use team-based integrations in event rules to make sure that connector ownership and execution of rules give precedence to general rules. Teams can maintain consistency and hierarchy while offering flexibility and customization options.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

View the list of available event rules on the **Event Rules** page to determine whether you want to create or edit an event rule.

You can create rules that:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Transform information in events to populate specified alert field values and compose alert fields from various values.
- Configure threshold rules that create or close alerts only when the incoming matching events exceed the specified threshold.
- Bind alerts to CIs using CI identifiers.

Options to create the rule are:

- Create an event rule and assign event fields for alert generation.
- Create a rule from an existing event or group of events that don't have a rule. In this case, the event fields are copied to the Event Match Fields section of the rule.
- Edit an existing event rule.
- For Team-based integrations, select an assignment group.

- Run multiple sequential rules defined for the same event by selecting the **Apply additional matching rules** check box. The event rules run in ascending order as defined in the **Order** field. Event rules applied to assignment groups only run after the global rules have run.

You can refresh an existing event rule with new event data. For more information see, Refresh event rules.

Note:

- Event rules that aren't configured to perform any action are skipped. Therefore, if the rule isn't configured as ignore, threshold, or binding, it's important to specify either the match or the composed fields.
- Make sure that you don't change the **Classification** field value in event [em_event] tables, either manually, by script, or by event rule.

Note: To automatically filter out irrelevant alerts or transform and standardize alert data for better response, you can also use Ignore automation and Enrich automation.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules** and take one of the following actions.

Option	Description
Create an event rule from an existing event	**a.** Select the link for unassociated events or grouped events that aren't mapped to the rules. 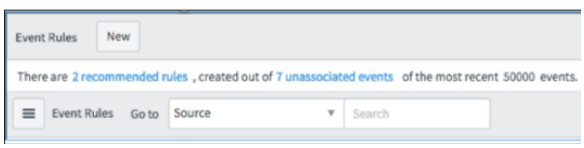 **b.** Select the event that you want to use for cre ating the rule. The event fields are copied to the Event Field Rules section of the rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Edit an existing event rule	In the event rule list, select the required event rule to be modified. The event rule opens in the event rule designer where you can modify the values of the fields. Select **Save and Upgrade** 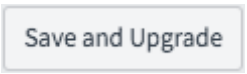 to modify the rule when the following banner message appears and you want to convert the event rule. `Rule cannot be viewed in the event rule designer. To modify the rule click 'Save and Upgrade'.`
Create an event rule	Select **New**.

2. Ensure that **Active** 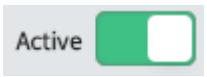 is selected.
   When the rule is deactivated, Event Management finds and applies another event rule. An alert is still created for the event unless *`Ignore`* is selected in another applicable rule or when configuring the filter for this event rule.
3. Enter a unique and meaningful name and fill in the form.

Event Rule Info form

Field	Description
Source	Category to which this matching rule applies. The mapping rule only applies to events with the same event class value. If this value is empty, apply the rule to all events.
Order	Order in which an event rule is evaluated when multiple rules are defined for the same type of event. Event rules are evaluated in ascending order.
Description	Type additional information that describes the event rule.
Apply additional matching rules	Select to apply additional event matching rules according to the Order field. The last rule with binding settings sets the CI binding. When selected, the Thresholds tab is inactive.
Assignment group	For team-based integrations, select an assignment group. If no assignment group is defined in the event rule, then this event rule is considered as a global rule. When the rules are running – first the global rules run and then the rules that belong to the assignment group that the event's source instance belongs to.

4. **Optional:** Define the event rule using these Event Rule Designer features.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Event Filter	Define a filter to restrict to which events the event rule must apply. See Filter the events that an event rule applies to.
Transform and compose alert output	Configure the customization of alert content. See Configure an event rule to customize alert content.
Threshold	Create or close alerts according to the specified threshold. See Set a threshold to suppress alert generation.
Binding	Configure event rules to automatically bind alerts to CI information from the CMDB. See Binding alerts to CIs.

5. Select **Save**, **Submit**, or **Update**.

Related topics

- Pattern matching
- Use event input information
- Configure an event rule to customize alert content
- Refresh event rules
- Simulate event processing
- Filter the events that an event rule applies to

Use event input information

The **Event Input** pane that is included in the steps to create an event rule provides a reference to the information that you can use when configuring an event rule.

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

In the **Transform and compose alert output** pane, you can also drag the information into the required fields to customize alert content. The fields in the right-hand pane appear as a reference to which fields, expressions, and their values are available when working in the Event Rule designer. The fields that display in the Transform and compose alert output left-hand workarea become the fields in the generated alert form. In addition, you can define alert tags that are specific for customization.

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

General components and navigation

The Event Input pane consists of these sections:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Event Additional Info
- Expressions
- Event Raw Info

The Event Input pane displays in all Event Rule designer sections, except for the Event Rule Info section. The fields in the right-pane can be modified only in the Transform and compose alert output section. In all the other sections, the Event Input pane is read-only and available for reference purposes.

Note: If **Accessibility enabled** 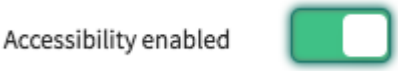

is selected in System Settings, standard edit mode is not available .

Edit modes

Mode	Description
Standard edit mode	Text selected in the left-hand workarea cannot be edited. Undo is not available. In this mode, you can only mark the text or values and transform it by selecting a field name.
Regex edit mode	Edit text in the left-hand workarea by selecting it, as required. After you have edited the text, you cannot go back to the standard mode. **Undo** is not available. However, **Reset** is available, returning the field to its original value. You can compose text using regular expression (regex) format conventions, for example, type (.*) to create a field name. You cannot use field values like: type "$" to include the value of a field and specify the name of the field in parentheses {}. You can include special characters in a field by "escaping" them, that is, by prefixing the character with a "\". For example, to include a "\" in your text, enter "\\".

Event additional info

The Event additional info section displays default event form fields. It also displays fields that were created to display in alerts and added to the **Additional information** section of the alert.

When configuring text in the fields in the left-hand workarea, you can drag field names and/or their value from this section into the fields.

Expressions

The Expressions section is available only if an additional information field was created. You can create expressions from event raw data or from additional information data.

1. Click in a field in the **Event Input** pane to perform transformation actions.
2. Create an expression by transforming information and associating it with a field name. Either select a field name from the suggested list, or create a field name, as required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click in a field in the **Event Input** pane. In the left-hand workarea, select the required information.

Create expressions by transforming information and associating it with a field name. Select a field name from the suggested list, or you can create a field name, as required.

Event raw info

The Event raw info section is always available. Its contents are the fields from the Event form.

To transform text, click a field in the right-panel to open it in the work area on the left-panel, in the Mark Expressions section. You can select text in the **Mark Expressions** section where the selected text becomes the value of a newly created field. Name the new field. This new field is added to the Expressions section.

Note: When creating regex expression for input, ensure that you create input with the variables in sequential order. While variables can be skipped for inclusion, the selected order must be maintained.

Controlling how fields appear:

- Use the exclusion list [em_event_rule_black_list] table to hide fields from the right-pane.
- Use the em_event_rule_order form to order the position of each field.
- Fields that are not in the exclusion list and also not part of the em_event_rule_order table still appear in the event rule designer. They appear in the right-panel and in the **Transform and Compose Alert Output** tab.

Related topics

Create or edit an event rule

Configure an event rule to customize alert content

Filter the events that an event rule applies to

Set a threshold to suppress alert generation

Filter the events that an event rule applies to

Define a filter to restrict to which events the event rule must apply. Configure the filter by providing a set of conditions that each event must match to be either excluded or included from applying to the event rule.

Before you begin

Ensure that a name is specified for the event rule.
Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

To construct the required condition statement, you can use event fields as well as additional fields that are defined in the **Additional Information** field of the event. Use these fields to configure conditions that can exclude events. Alternatively, you can filter incoming events to ensure that only those events that meet the required conditions are included.

Note: Event rules filter is case-sensitive.

While working in the Event Filter tab:

- In the Event rule list, if you selected a recommended rule or an existing rule, the filter fields are populated with values from these events. You can use these filter fields to compose filter conditions.
- You can click another tab (for example, **Threshold**), work in that area and then return to work in the Event Filter area without losing information.
- Use the right-panel as a source of reference for which fields are available to be used when composing filter conditions.

Note: Event filters created on additional_info fields only perform a string comparison and not a numeric comparison, even if the information entered into the value field appears as if it is numeric, as depicted in the following graphic:

Note: To automatically filter out irrelevant alerts or transform and standardize alert data for better response, you can also use Ignore automation and Enrich automation.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Click **New** or open an event rule.
3. Click **Event Filter**.

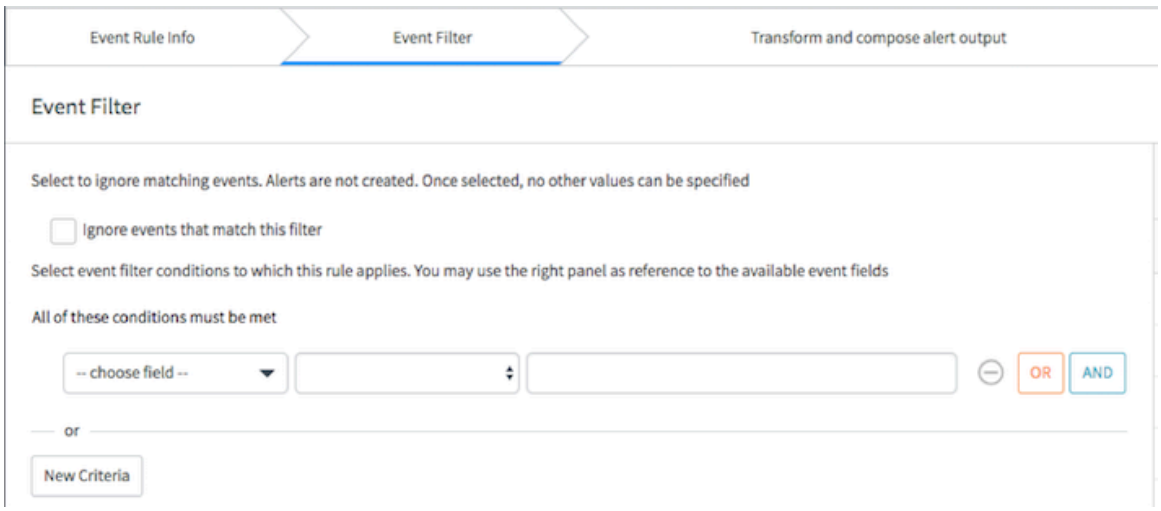

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. To ignore events, select **Ignore events that match this filter**.
   All events that match this event rule are ignored. You can add more conditions to the filter by clicking either **OR** or **AND**.

   Note: Even if an event is ignored by an event rule, the event is still recorded in the Event [em_event] table. Select this option to inhibit an alert from being created and be associated with the event.

5. To configure conditions to which the event rule must apply, click **New Criteria**.
   You can add further conditions by clicking either **OR** or **AND**.

   Note: The **%** character is not supported in filter conditions.

6. You can add field names to the Event Rule Simple Field Black Lists [em_event_rule_black_list] table to either show or hide fields from the filter.

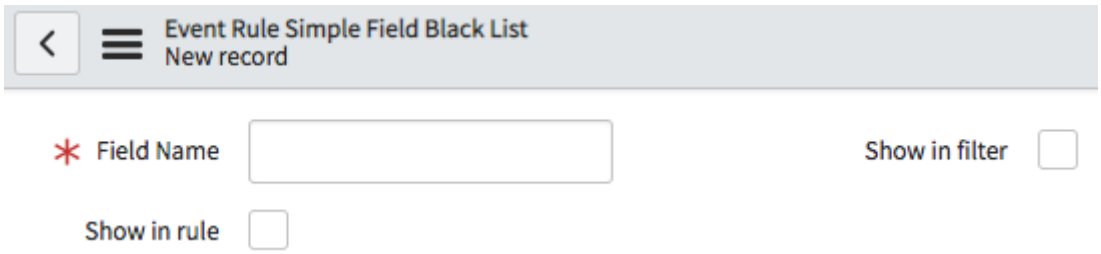

   - To remove fields from the filter and the right panel, add a new black list field with **Show in rule** not selected.
   - To remove fields only from the filter and still enable it to appear in the right panel, add an exclusion list field with **Show in rule** selected.

7. Click either **Save** or **Submit**.

Example:
Example showing filter conditions that are specified to run a rule when the metric name starts with the word "disk". For this example, any resource is acceptable, as long as it is not "cpu" and the node name is "MyServer".

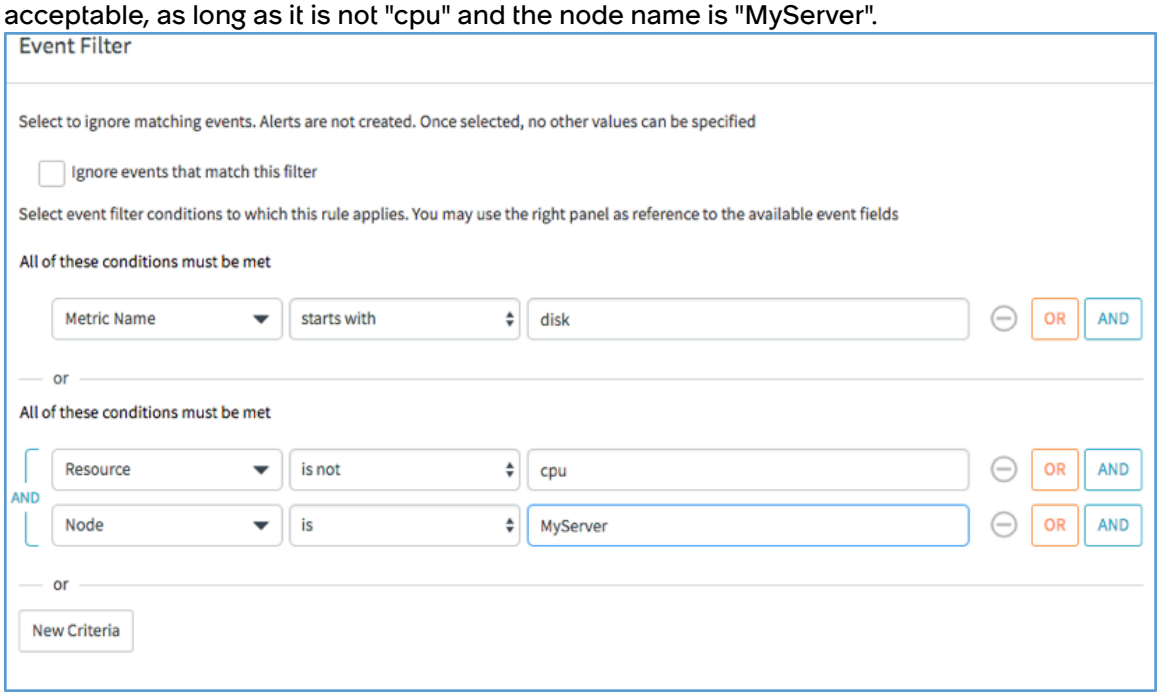

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example to filter events using time conditions. In the **Select Field or add New** field, select `Time of event`.

Specify the required time conditions, for example, anytime after yesterday.

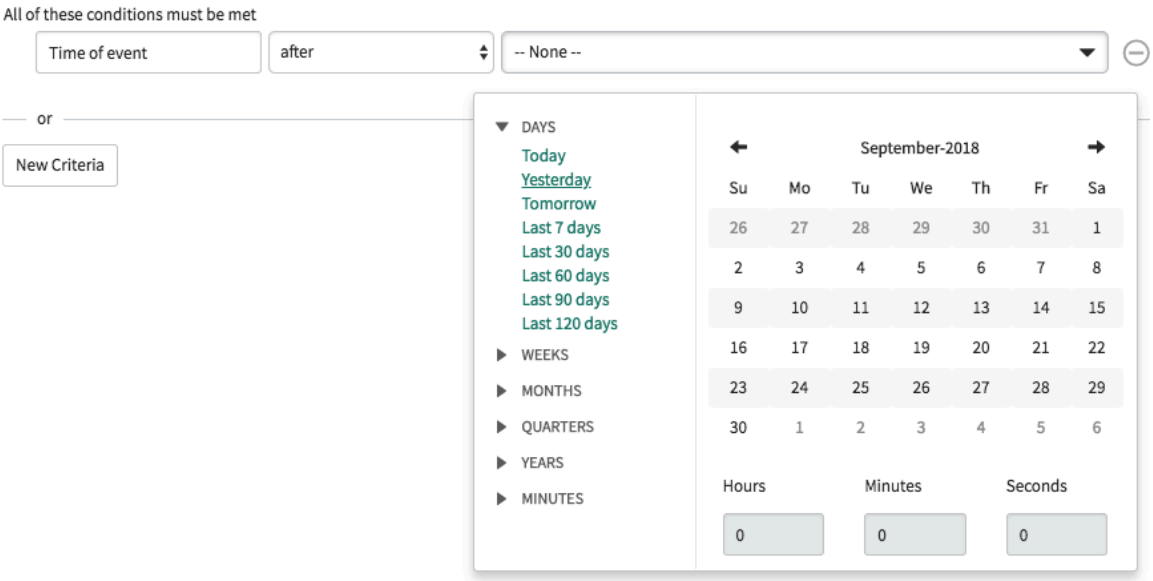

Related topics

- Create or edit an event rule
- Use event input information
- Configure an event rule to customize alert content
- Set a threshold to suppress alert generation

Find events that are not matched to rules

Find events that are not matched to any rules, and determine if it is necessary to create event rules to manage them.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

This task describes how to find events that are not matched to rules.

Note: Use Ignore automation and Enrich automation, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Near the top of the form, click the link for events or grouped events that are not mapped to rules.

Example

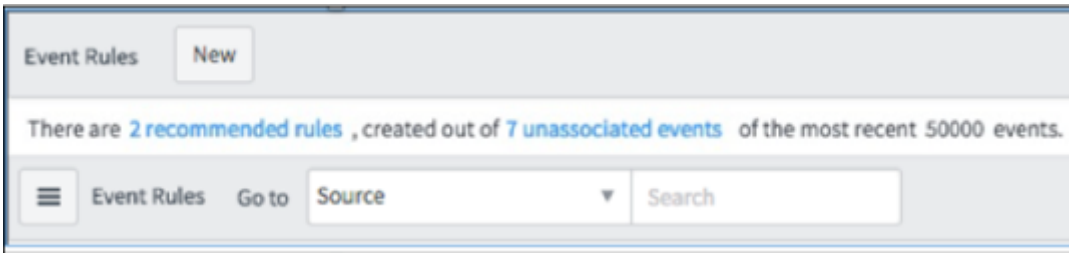

What to do next

Create an event rule. For example, event rules are useful for managing events that occur regularly.

View event rules

You can view all event rules on the Event Rules list.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

Note: Use Ignore automation and Enrich automation, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

Procedure

Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.

Event Rules list

Field	Description
Name	The event rule name.
Override default binding	Legacy binding criteria. If, in the **Bind** section criteria have not been specified or if the specified bind criteria are not matched, then the legacy binding criteria are considered.
Order	Order in which an event rule is evaluated when multiple rules are defined for the same type of event. Event rules are evaluated in ascending order.
Source	Event monitoring software that generated the event, such as SolarWinds or SCOM. Maximum length: 100 characters.
Updated	Update date and time for the rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can filter the Event Rule list to display the required subset of the information. However, if you create a favorite link after filtering, when the link is clicked the Event Rule list does not display the correct filter.

Related topics

View events

Configure an event rule to customize alert content

You can configure an event rule to customize alert content. You can customize the order of the fields and select which fields display. The fields in the left-hand work area of the **Transform and Compose Alert Output** section of an event rule are the fields that appear in the generated alert.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

You can customize alert content by adding the required event fields and modifying the text in the fields.

Note: When creating regex expression for input, ensure that you create input with the variables in sequential order. While variables can be skipped for inclusion, the selected order must be maintained. Use unique names within the Additional Info data to prevent conflicts with the predefined field variable names. Identical names cause issues in event rule regex. As a result, you may find that a saved regex is missing or unavailable when you later access it from the Expressions section.

Format alert text by using one or both of these methods:

Format alert text

Format method	Description
Transform	Create a field to display in an alert. The field and its value appear in the **Additional Information** section or in the **Alert Tags** section. In the **Event Input** pane, the new field is listed in the **Expressions** section.
Compose	Determine which text to display in the selected field.

Note:

- Configure which fields display in the **Event Input** pane by configuring the exclusion list [em_event_rule_black_list].

  A field entered into the exclusion list that already appears in an existing event rule is not removed from the rule. Instead, it is moved to the **Manual Attributes** section in the rule.

- Control the order of event form fields by configuring em_event_rule_order.

To automatically filter out irrelevant alerts or transform and standardize alert data for better response, you can also use Ignore automation and Enrich automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Click **New** or open an event rule.
3. Click **Transform and compose alert output**.
   In this section, you can use the **Transform** mode and the **Compose** mode.
4. To transform information in fields—standard mode:

   Note: In this step, the selection of text is made in standard edit mode. In this mode, undo is available.

   a. Click the required field in the Event Input pane. The value and contents of the selected field appear in the workarea on the left.

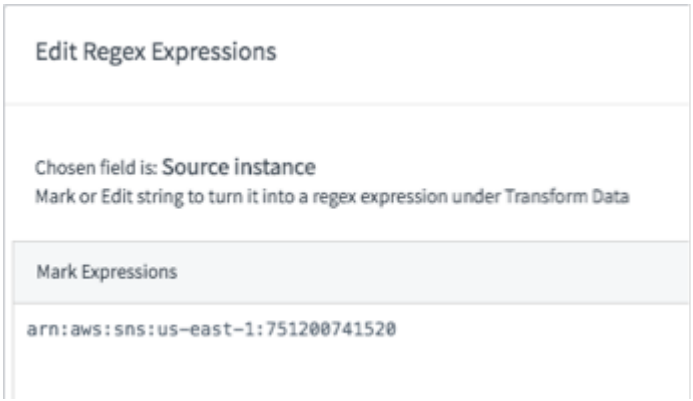

   b. Select standard edit mode .

   c. Using the pencil cursor, highlight the required text. You can double-click text in the field to select a word or select the required text and symbols.

   Note: If **Accessibility enabled** Accessibility enabled is selected in System Settings, standard edit mode is not available .

   d. Select a field name from the list or enter the required expression name. Fields that appear in the list are existing event fields, additional info fields or previously defined alert tags (fields prefixed with "t_"). The field name appears in the **Expressions** pane and the selected text is highlighted. Both the text and the expression name are highlighted in the same color. Each additional expression that you create, plus its value, is highlighted in a different color.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

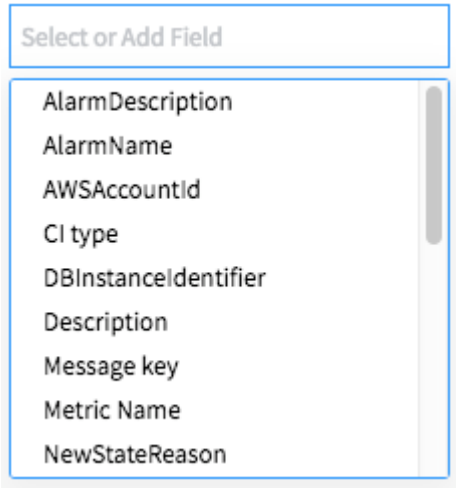

Note: Transform field names cannot contain a ‘ (single quote) character.

e. Click **Done**. The expression name appears in the **Expressions** section of the **Event Input** pane.

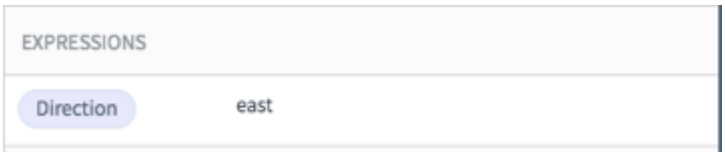

f. If you click an expression or its value in the **Event Input** pane, the details display in the workarea on the left.

5. To transform information in fields—regex mode:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- In regex mode, undo is not available.
- When building a regex expression from variables that are highlighted in the event, place the selected variables in the same order that they appear in the event field.
- When composing the event rule regex, specify the terms using an exact match of the input from the event fields. For example, for the text `vhxy2stlgwc_8888_hec03v012979`, the regex

  ```
  ^([^_]+)_
  ```

  looks as if it should catch the text `"vhxy2stlgwc"` that appears before the underscore. However, it does not work as this text is not an exact match. Use the full regex, as follows:

  ```
  ^([^_]+)_.*
  ```

- The Java Matcher parser is used for the regex engine. Ensure that you use a Java parser to test the regex.
- If the regex does not match the entire string when parsing data, the event rule will not capture the event, and no parsing takes place.
- The Java regex engine "matches" method is used. With this implementation, the regex engine attempts to match the entire input sequence against the specified pattern. Other matching methods, for example, the find method, are not supported.
- No regex flags are set by default, except the `Pattern.DOTALL` flag. With this implementation of Java regex, the DOT character "." matches any character, including a line terminator.
- Aside from `Pattern.DOTALL`, other regex flags cannot be changed.
- `Lookahead` is supported. However, `Lookbehind` is not supported.

a. Click the required field in the Event Input pane. The value and contents of the selected field appears in the workarea on the left.

b. Select regex edit mode .

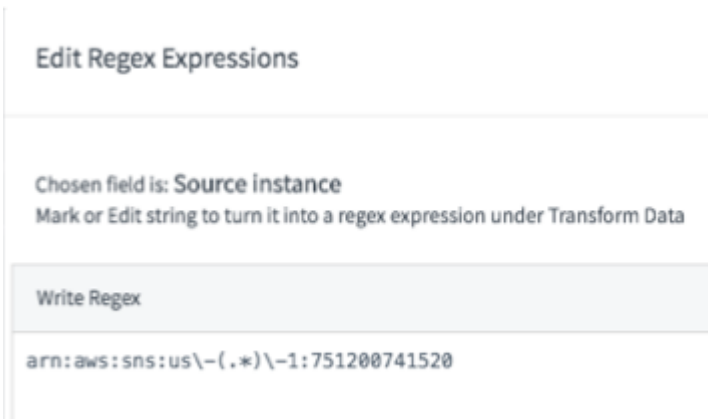

c. Select the required text and symbols 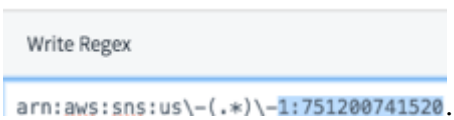

. Replace the highlighted text with (.*) to create a regex expression.

d. In the **Expressions** pane, click the **Select or Add Field** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

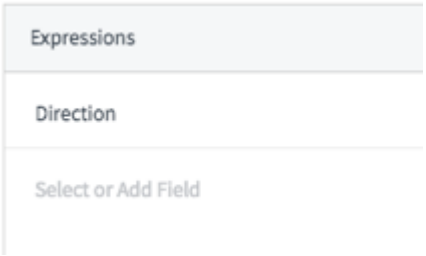

e. Select a field name from the list or enter the required expression name. Fields that appear in the list are existing event fields, additional info fields or previously defined alert tags (fields prefixed with "t_"). The field name appears in the **Expressions** pane.

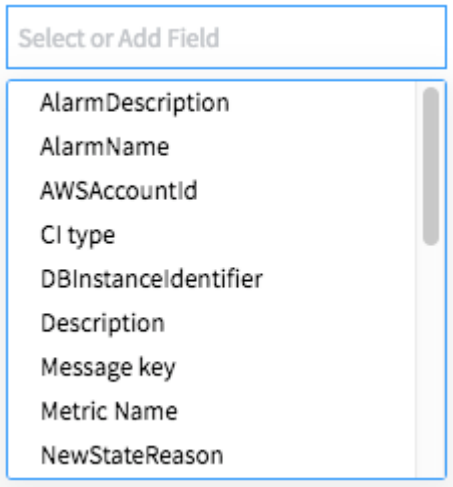

Note: Transform field names cannot contain a ‘ (single quote) character.

f. Click **Done**.

Note:

- When building a regex expression from variables that are highlighted in the event, place the selected variables in the same order that they appear in the event field.
- When composing the event rule regex, specify the terms using an exact match of the input from the event fields. For example, for the text vhxy2stlgwc_8888_hec03v012979, the regex ^([^_]+)_ looks as if it should catch the text “vhxy2stlgwc” that appears before the underscore. However, it does not work as this text is not an exact match. Use the full regex, as follows: ^([^_]+)_.*

6. To compose text, in the Transform and compose alert output section, click in the required field.

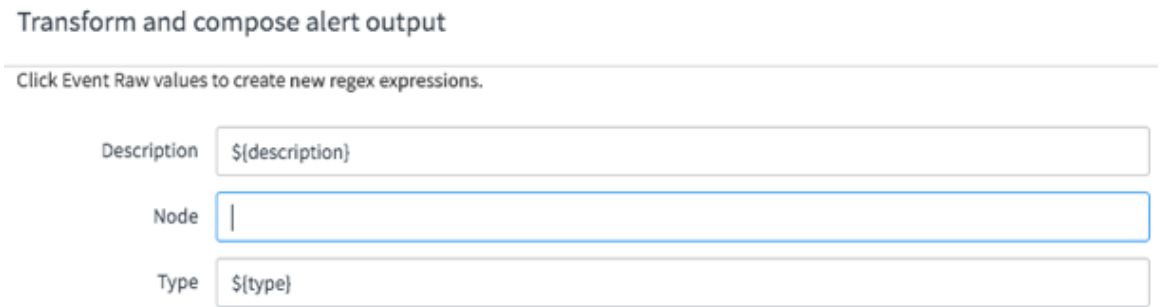

Compose text in the alert fields by performing any of these actions:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Drag a field from the Event Info pane into the required field.
- Use the Event Info pane as a reference of which fields are available. Write the field name details, using the format ${<fieldname>}, where "$" means "use the value". Between the brackets, the name of the event field replaces the <fieldname> variable.
- Freely enter text, spaces, or characters. For example, use spaces or underscore symbols (_) to make the text in the alert field readable.

7. To specify fields that must display in the generated alert, select the **Manual attributes** check box.

   Note: These fields, and their values, display in the **Additional Information** section or in the **Alert Tags** section for attributes prefixed with "t_".

   a. In the left-hand column, enter a value.

   b. In the right-hand column, either drag a value from the Add or Expression fields or type in a value.
   When entering a value with attributes prefixed with "t_", a drop-down list appears with all the previously defined "t_" attributes.

8. Click **Save** or **Update**.

What to do next
You can filter event rules. See Filter the events that an event rule applies to.

Related topics
Create or edit an event rule

Use event input information

Alert tags

Filter the events that an event rule applies to

Set a threshold to suppress alert generation

The event threshold is the rate upon which Event Management generates an alert. Receiving multiple events for a device over a short interval may warrant creating an alert, as the condition may be serious. However, receiving events over a longer interval may indicate a less serious situation which would not warrant creating an alert.

Before you begin
Role required: evt_mgmt_admin

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

You can configure the properties in an event rule to create alerts, suppress alert generation, or close existing alerts according to a specified threshold. The threshold is based on the value of event fields or number of occurrences over a specified period.

Note: **Field Name** can be the name of any numeric field in the **Additional information** field of the event. Therefore, if CPU is an additional information field for a specific event, then CPU can be used as a **Field Name**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Assume that you want to generate an alert when CPU utilization reaches or exceeds 80% three times with no more than 20 seconds between any two consecutive events. Create an event rule with these settings (an explanation for each value is given in parentheses):

- **Create Alert Operator**: *>=* (operator to determine whether utilization of **Field Name** reaches or exceeds the specified value)
- **Field Name**: *CPU* (events regarding high CPU usage)
- **Threshold Value**: *80* (percent)
- **Occurs**: *3* (three events occur where the CPU usage is at or above ">=" 80%)
- **Over(seconds)**: *20* (20 seconds or less between events)

To demonstrate how the preceding settings are evaluated, assume that the following events are received:

Scenario	Outcome
Reported elapsed time and the CPU usage for each event: • First event elapse time 20, CPU=85 • Second event elapse time 40, CPU=80 • Third event elapse time 60, CPU=70	In this scenario, no alert is generated since one event has a CPU utilization that is under 80%.
Reported elapsed time and the CPU usage for each event: • First event elapse time 20, CPU=85 • Second event elapse time 40, CPU=90 • Third event elapse time 70, CPU=95	In this scenario, an alert is not generated since the elapsed time in one event is over the specified 20 seconds.
Reported elapsed time and the CPU usage for each event: • First event elapse time 20, CPU=85 • Second event elapse time 40, CPU=95 • Third event elapse time 60, CPU=90	In this scenario, an alert is generated since in all events the elapsed time is within the specified time and the CPU usage is over 80%.

Note:

- When configuring an event rule to create or close alerts according to a threshold, events that arrive at the same second, as determined by the `time_of_event` field, are skipped. This occurs because these events are considered to be duplicates.
- When threshold conditions are configured, alerts close when the **Close Alert Operator** condition is met, and not when an event is received with **Severity = Clear** or **Resolution state = Closing**.

To automatically filter out irrelevant alerts or transform and standardize alert data for better response, you can also use Ignore automation and Enrich automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Create or open an event rule.
3. Click **Threshold**.
4. Select **Active**.
   If **Active** is not selected, the fields on this page do not appear.
5. In the **Create Alert Operator** field, select an operator.
   - If you select **Count**, specify the corresponding **Occurs** and **Over(seconds)** fields.
   - If you select any operator other than **Count**, the **Field name** and **Threshold value** fields appear. Specify the required values in these fields.

     Note: The value of the **Threshold value** property can be the name of any field in the **Additional information** of the event. For example, if `CPU` is a field in **Additional information** for a specific event, then `CPU` can be used as the **Threshold value**.

   - In the **Occurs** field, specify the required value.
   - In the **Over(seconds)** field, specify the required period.
6. To automatically close alerts, in the **Close Alert Operator** field, select an operator.
   Extra fields appear according to your selection.
   - If **Count** is specified as the operator for the **Create Alert Operator** field, then the selection in the **Close Alert Operator** field is either `None` or `Idle`. Specify the required value.
   - If you select **Idle**, configure the **Over(seconds)** field.
   - In the **Close Alert Operator** field, if you select an operator other than **Idle**, then configure the **Threshold value**, **Occurs**, and **Over(seconds)** fields.
7. Click **Save** or **Submit**.

Example:

To create an alert when a specific event occurs 5 times in 10 minutes, in **Threshold**:

1. In the **Field name** field, specify the name of any field that exists in the **Additional information** field in the event. The value of the field is irrelevant.
2. In the **Create Alert Operator** field, select `Count`.
3. In the **Occurs** field, specify `5`.
4. In the **Over** field, `600` (10 * 60 seconds).
5. Click **Save** or **Submit**.

To create an alert when a specific event occurs 5 times in 10 minutes with a metric value greater than 55. Assume that “metric_value” is a field in the **Additional information** of the event. Specify:

1. Select the *`Active`* check box.
2. In the **Create Alert Operator** field, select `>=`.
3. In the **Field name** field, specify **metric_value**.
4. In the **Threshold value** field, specify `55`.
5. In the **Occurs** field, specify `5`.
6. In the **Over** field, specify `600` (10 * 60 seconds).
7. Click **Save** or **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Create or edit an event rule
- Use event input information
- Configure an event rule to customize alert content
- Filter the events that an event rule applies to

Binding alerts to CIs

CI binding or linking is the process of finding and connecting a Configuration Item (CI) from the Configuration Management Database (CMDB) to an alert, using the logic defined in event rules. This helps ensure alerts are tied to the right IT components for better visibility and faster issue resolution.

Understanding CI binding

Binding alerts to CIs links incoming alerts to the correct Configuration Item (CI) representing a specific host, such as a computer, server, router, or virtual machine in your IT infrastructure. This ensures accurate alert tracking, simplifies troubleshooting by identifying the source of issues, and maintains a historical record of alerts tied to specific systems.

In IT Operations Management, the relationship between alerts and CIs plays a critical role in effectively managing services and infrastructure. A CI represents a component within your IT environment, such as a server, application, or database. Linking alerts to CIs ensures that alerts and incidents are directly associated with the affected component, enabling accurate impact analysis and faster resolution.

In dynamic and large-scale environments, it's often challenging to pinpoint which CI is responsible for triggering an alert. Alerts are generated from raw event data captured across multiple sources, but this data may not explicitly reference a CI. Without a reliable method to bind alerts to CIs, IT teams face the risk of:

- Poor visibility into service impact.
- Longer mean time to resolution (MTTR).
- Inefficient alert management processes.

Event rules provide a structured way to analyze incoming event data, enrich it with additional context, and map it to the correct CI. These rules act as a bridge, transforming raw event data into actionable alerts linked to CIs. For example, event rules can:

- Parse incoming event data: The process involves breaking down raw event data to extract important details, such as hostnames, IP addresses, or service tags. These details are then used to analyze the event—identifying its source, type, and impact—and to enrich the alert by adding relevant context, like affected services, priorities, or resolution steps, ensuring it's actionable for IT teams.
- Match patterns: Here, the extracted event details—such as hostnames, IP addresses, or service tags—are compared against entries in the CMDB to identify the corresponding Configuration Item (CI) based on predefined rules or logic. Once matched, the alert is linked to the CI, providing context about the affected IT component for better visibility and faster resolution.
- Bind the CI to the alert: At this point, the alert is automatically associated with the identified CI. It ensures that the alert is connected to the right IT component in the system, making it easier to track and manage, and helping teams quickly address the issue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Types of CI binding

Binding type	Description
Implicit Node binding (Default binding)	Binds alerts to host CIs by matching the **Node** field in the event with attributes like Name, FQDN, IP, or MAC address.
CI field matching binding	Binds alerts using the **CI Identifier** field, a JSON structure containing column names and values (e.g., Name, FQDN, IP, or MAC Address).
CI identification binding	Binds alerts to specific applications on hosts using event rules.
Process-level binding	Binds alerts to application CIs by matching process-related attributes.
Device-level binding	Binds alerts to device CIs by matching event data with attributes like hostname, IP address, or MAC address.

Key benefits

- Improved service visibility: Alerts linked to CIs provide a clear picture of which IT components are affected.
- Enhanced automation: Automated CI binding reduces manual effort, ensuring consistency.
- Faster resolution: Teams can quickly diagnose and resolve issues by focusing on the impacted CI.

Example use case

Scenario: A company's email server named MailServer-01 goes down.

Process:

1. Alert generation: An alert is triggered when MailServer-01 becomes unresponsive.
2. CI binding: Event Management automatically looks for MailServer-01 in the CMDB (Configuration Management Database).
   - If found, the alert is linked to this specific server (CI).
   - If not found, the alert remains unlinked until more information is provided.

Result: The alert now shows that MailServer-01 is the affected system.

- Impact: IT teams immediately know which server to investigate.
- Efficiency: They don't need to check other infrastructure components or systems, saving time and effort.

This direct link helps streamline troubleshooting by pointing to the exact problem source.

To enrich alerts by identifying the CI or extracting, composing, or tagging alert fields, you can also create an enrich automation in Service Operations Workspace. For more information, see Enrich automation.

Related topics

Simulate event processing

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Binding process flow

Learn the process of binding Configuration Items (CIs) to alerts. This includes handling event arrival, binding alerts using available fields when no node is present, and searching the CMDB for matching hosts. It also explains linking alerts to CIs based on host and CI type detection.

CI binding to alerts: Flow diagram

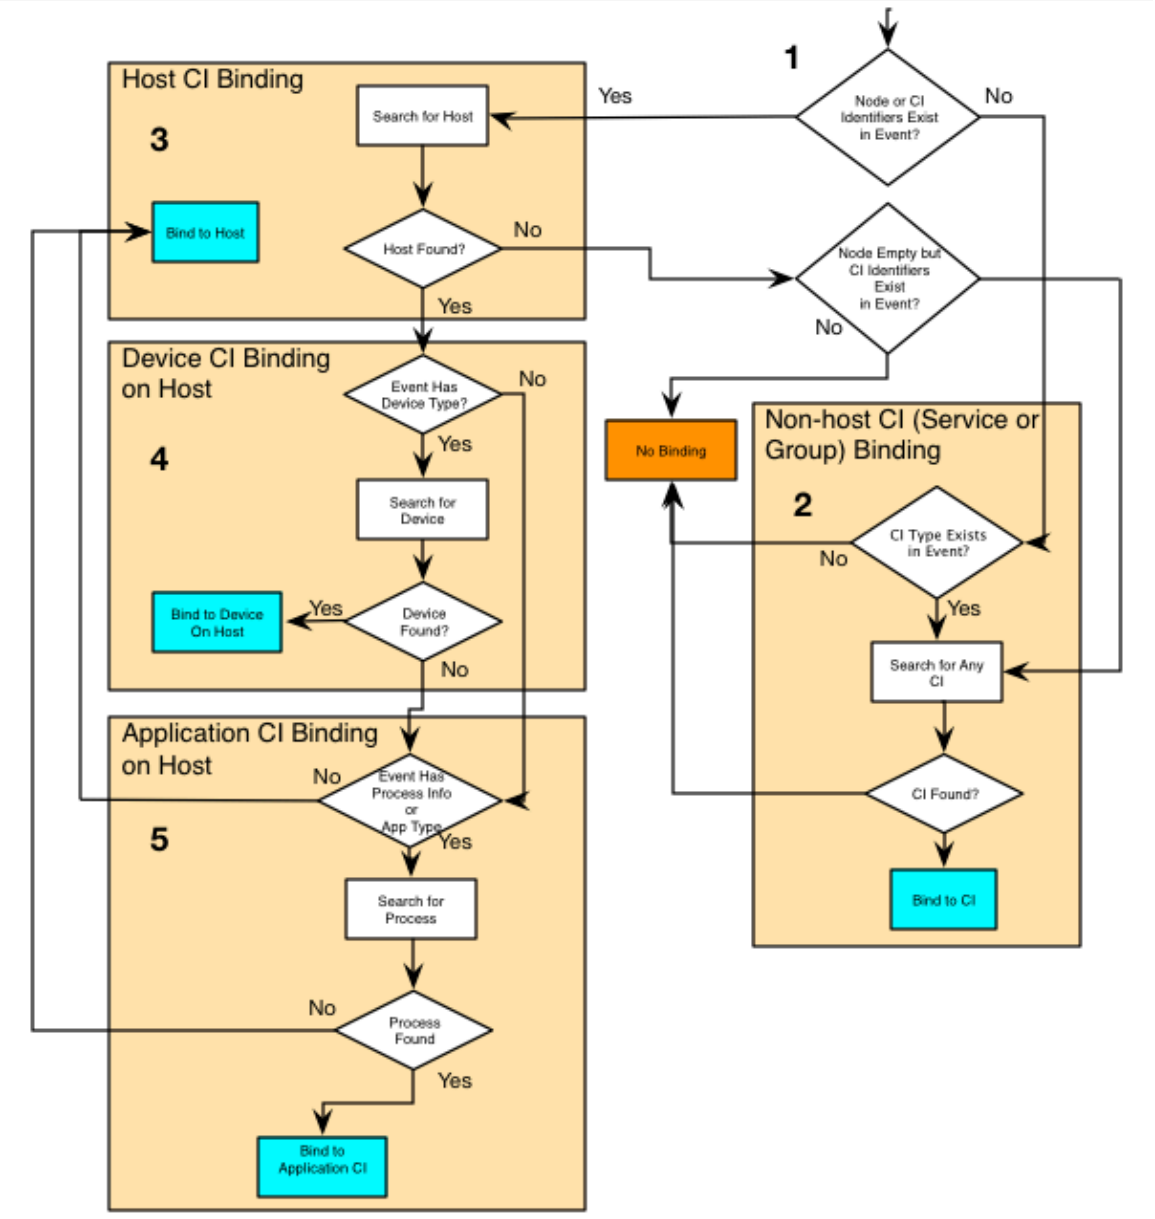

Process of linking CIs to alerts

1. Event Arrival: When an event is received, Event Management checks for node or CI identifiers.
2. Is Node Value Present?
   - Yes: If a node value is provided, Event Management searches the CMDB for a matching host.
   - No: If no node is found, the system attempts to bind the alert using other available details, such as:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Alert Type
- Additional Information
- CI Identifier

Note: For Kubernetes CIs and other non-host CIs, use an event rule with CI field matching instead of relying only on the Node value. Kubernetes CIs follow the same binding approach as other non-host CIs.

3. Host and CI Type Detection:
   - Both Found: If both a host and CI type are identified, the alert is linked to the corresponding device CI.
   - Host Only Found: If only a host is identified, the system links the alert to an application CI.

Additional Considerations

- Event Processing Notes: The event can include the binding process flow in its **Processing Notes** field, providing insights into the binding actions.
- Binding to Retired CIs: By default, events do not bind to CIs with a Retired status. To enable binding to these CIs, set the *evt_mgmt.ignore_retired_cis_in_binding* property to false.
- Specifying CI Statuses for Binding: You can configure the *evt_mgmt.install_status_list_to_ignore_in_binding* property by specifying CI status numbers to control which CIs are considered for binding. For example:

CI Statuses

Status Number	Status
1	Installed
2	OnOrder
3	InMaintenance
4	PendingInstall
5	PendingRepair
6	InStock
7	Retired
8	Stolen
100	Absent

Note: To specify multiple statuses, separate each status number with a comma.

Binding alerts to a specific host CI (default binding)

Binding alerts to Configuration Items (CIs) using the **Node** field or the **CI Identifier** field ensures accurate event association. By comparing an event record's **Node** or **CI Identifier** value, alerts are linked to the right system. This improves response, root cause analysis, and impact assessment by providing clear visibility into affected assets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How default binding works

The system can perform default binding using the **Node** field or the **CI Identifier** field from the event record.

When an event enters the system, a key field like **Node** is available in the event record. However, there is no **Node** field in the CI. Instead, the node value from the event is compared with various attributes in the host CI, such as Name, Fully Qualified Domain Name (FQDN), IP, or MAC address. If a match is found, the alert is linked to the corresponding CI. This is the default method for binding alerts to CIs.

The property *sa.active_operation_status* is used by the default binding mechanism, specifically when binding via IP/MAC using the node field in the event. When an event attempts to bind to a host through the node field (using IP/MAC), the system ignores hosts whose operational status is not included in this property and the operational status of the host must also not be null. By default, the property value is "1", which corresponds to Operational. You can extend the list to include additional statuses, and any status in the list is not ignored during binding.

Example: If *sa.active_operational_status* = "1,2", then both Operational and Non-operational hosts are considered valid for binding. The status attribute is defined by the following values:

- 1 – Operational
- 2 – Non-operational
- 3 – Repair in Progress
- 4 – DR Standby
- 6 – Retired

Note: In this binding process, the CI must be a host. Host CIs include Computers, Operating Systems (OS), Switches, Routers, or any CI type/class that extends the [cmdb_ci_hardware] table, meaning the CI type or class must be a hardware component.

If no match is found using node, the system looks at the **CI Identifier** field on the event record. The **CI Identifier** is a JSON structure containing column names and values for comparison (e.g., Name, Fully qualified domain name, IP, or MAC Address). The JSON is: {"column_name":"<column_value>"}. For example, if we want to bind to a CI identified by **serial_number** whose value is **Dell Latitude 7420 Laptop**, the JSON is: `{"serial_number":"Dell Latitude 7420 Laptop"}`. If a match is found, the alert is linked to the corresponding CI.

Difference between binding with Node vs. CI Identifier

In default binding, the system considers all fields in the **Additional Information** field and attempts to match their values with the CI table. The **CI Identifier** field allows specifying particular fields for matching, even if they are not present in **Additional Information**. This process uses a predefined JSON structure and applies only when the CI is a host.

Note: Even if the **Node** or the **CI Identifier** successfully binds the alert with the CI, event rules further determine how the binding occurs using the **Override default binding** check box available in Event rules. This ensures that alerts are bound correctly based on business logic, not just direct matches. This ensures that alerts are bound correctly based on business logic, not just direct matches.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example: Resolving a CI Using Node and Event Rules

Imagine a server (Server-123) in your network generates an event. The event record contains a **Node** field with the value "Server-123.example.com".

1. Matching with CI Attributes:
   a. The system checks the Node value against the CMDB.
   b. It compares "Server-123.example.com" with the FQDN, IP, MAC address, or Name of existing host CIs.
   c. If a match is found (e.g., FQDN in the CMDB is also "Server-123.example.com"), the alert is linked to that CI.
2. Applying Event Rules: Even if the Node resolves to Server-123, additional event rules might determine if the alert should be linked differently. For example, an event rule may specify that alerts from Server-123 should be linked to a parent CI (like a cluster) instead of the individual server.

Fallback CI binding fields

Use these **Additional Info** fields to help the system bind alerts to the correct Configuration Item (CI) when standard CI binding methods don't find a match.

- sn_fallback_binding_ci_id: Specifies the exact CI to bind to by providing the CI's `sys_id`
- sn_fallback_binding_ci_name: Specifies the name of the CI to search for and bind to. If multiple CIs have the same name, the system uses the first matching CI.
- sn_fallback_binding_ci_type: Used in conjunction with `sn_fallback_binding_ci_name`: Use this field together with sn_fallback_binding_ci_name to narrow the search when multiple CIs share the same name. This field specifies the CI class or type that the system should match.

The system evaluates these fallback fields only after it attempts all standard CI binding strategies and finds no matching CI. Store these fields (keys) in the Additional Info field of an event and populate them with relevant CI information. When standard binding methods fail, the system uses the values of these fields as fallback criteria for CI binding. For example, the Additional Info field might contain:

```
{"sn_fallback_binding_ci_name": "linux01"}
```

In this example, the system uses the value of `sn_fallback_binding_ci_name` (`linux01`) to find and bind the appropriate CI when standard CI binding strategies don't identify a match.

Overriding default binding

Overriding default binding allows flexibility in linking alerts to CIs beyond the default Node-based matching. This helps customize alert binding based on business needs, ensuring accurate associations and preventing mislinked alerts, improving incident management and root cause analysis.

By default, alerts are bound to CIs by matching the node value in the alert with attributes like Name, FQDN, IP, or MAC Address in the CI. However, if you want to change this behavior—either by using additional criteria or ignoring the node altogether—you can use the **Override Default Binding** check box.

This allows for two scenarios:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. When Binding Type = CI Field Matching
   - If the CI is a host, the Node is required for binding. For more information, see Bind host CIs using CI field matching.
   - If the CI is not a host, the Node must be empty. For more information, see Bind non-host CIs using CI field matching.
2. When Binding Type = CI Identification. For more information, see Bind alerts to CIs using CI identification.

This flexibility ensures alerts are linked correctly based on specific business needs.

Note: You can also use the Service Operations Workspace to define binding rules. For more information, see Create Enrich automation.

Bind host CIs using CI field matching

When CI Field Matching is used and the CI is a host, the Node value from the alert is used for binding. The system compares the Node with attributes like Name, FQDN, IP, or MAC Address in the CMDB to find a match. This ensures that alerts are correctly linked to the corresponding host CI.

Before you begin
Role required: evt_mgmt_admin

About this task
When an event enters the system, key field like **Node** is available on the event record. The **Node** field generally represents the name of the server. In CI there is no field like **Node**. However, the value of **Node** from the event is compared with different values in the host CI such as Name, Fully qualified domain name (FQDN), IP, or MAC Address. If a match is found, the alert is linked to the corresponding CI. This is the default binding of alerts to CIs. In this binding, the CI must be a host. Host CIs include Computers, OS, Switches, Routers—essentially any CI type or class that extends the [cmdb_ci_hardware] table.

There may be cases where no match is found because the column names in the event record and the table differ for the same item. In such cases, you can manually create an additional key-value pair with a name matching the table column, ensuring the matching process continues successfully. For information on how to create a manual field, see Bind CIs using CI field and column matching.

Note: Even if the node successfully binds the alert with the CI, event rules further determine how the binding occurs.

Procedure
1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Select **New** and complete the required fields of the event rule.
3. Select the **Binding** tab.
4. Select the **Override default binding** check box.
5. In the **Binding type** field, select **CI field matching**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

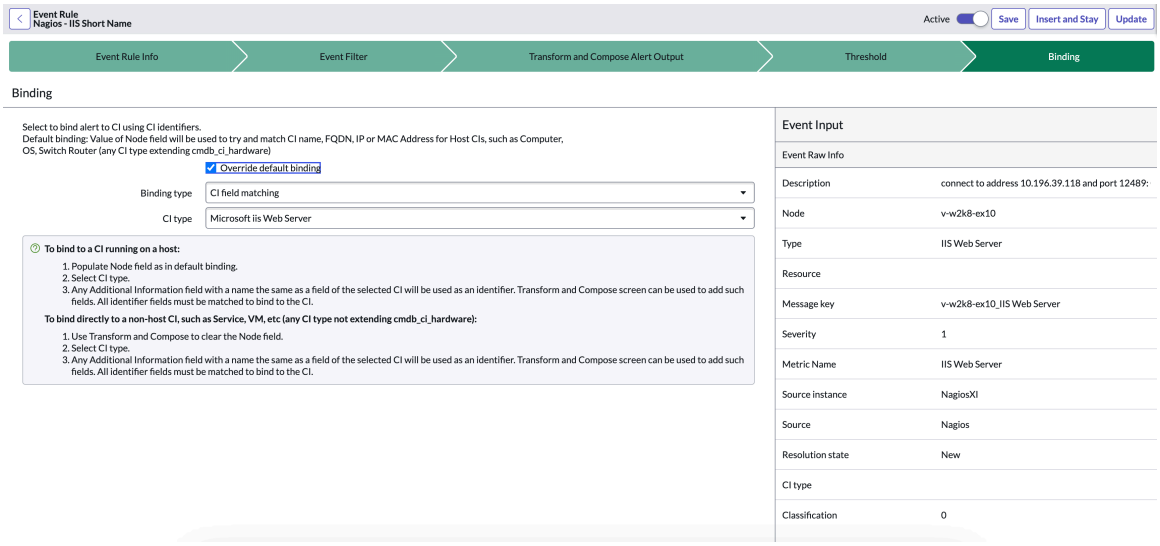

6. In the **CI type** field, select the appropriate type of hardware component (such as a server, router, or storage device) that the CI type represents.
   The **CI type** determines the specific CMDB table where the system searches for the matching CI. For example, if the **CI type** is **Network Adapter**, the system searches in the [cmdb_ci_network_adapter] table.
7. Select **Submit**.

> **Note:** You can also use the Service Operations Workspace to define binding rules. For more information, see Create Enrich automation.

Related topics

Bind non-host CIs using CI field matching

Bind alerts to a specific device

Bind each alert directly to the originating device to establish a clear source of impact. This ensures accurate troubleshooting, reduces noise from unrelated alerts, and helps teams focus on resolving the right issue faster.

Before you begin

Role required: evt_mgmt_admin

About this task

In IT environments, multiple devices like switches have many smaller components, such as ports. When something goes wrong, it's not enough to just say the switch has an issue—the system needs to pinpoint exactly which port is causing it. Binding alerts to the specific port that triggered the issue allows teams to act faster and fix the real issue without wasting time checking the entire switch or unrelated components.

The following example procedure shows the alert binding configuration for a switch named IP-Switch-1, ensuring alerts target the correct port. The event rule setup starts by enabling the **Override default binding** option and selecting **CI field matching** as the **Binding type**. Choosing **Switchport** as the **CI type** instructs the system to bind alerts to individual ports. The Binding Device Map defines how the system links the switch (cmdb_ci_ip_switch) to its ports (dscy_switchport) using the cmdb_ci reference field. This map guides the binding logic. Under **Configuration** > **Network** > **Switches**, selecting **IP-Switch-1** reveals its switchports and interface names. When an incoming event includes a matching interface name in the **Additional**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

information field, the system binds the alert directly to the correct port—ensuring precise and meaningful alert context.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Select **New**.

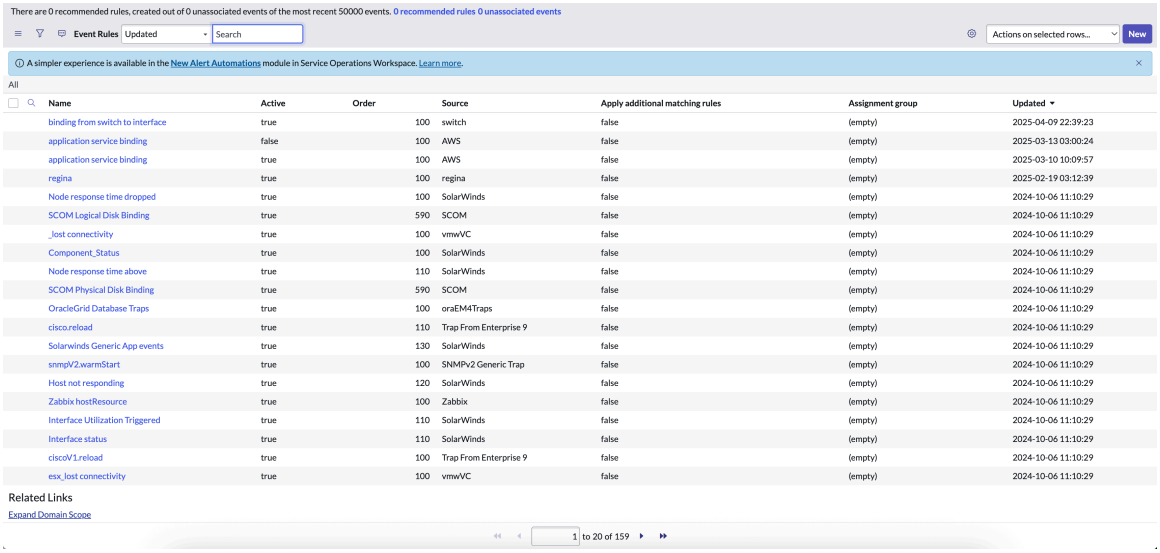

3. Select the **Binding** tab.
4. Select the **Override default binding** check box.
5. In the **Binding type** field, select **CI field matching**.
6. In the **CI type** field, select **Switchport**.
   The **CI type** determines the specific CMDB table where the system searches for the matching CI.

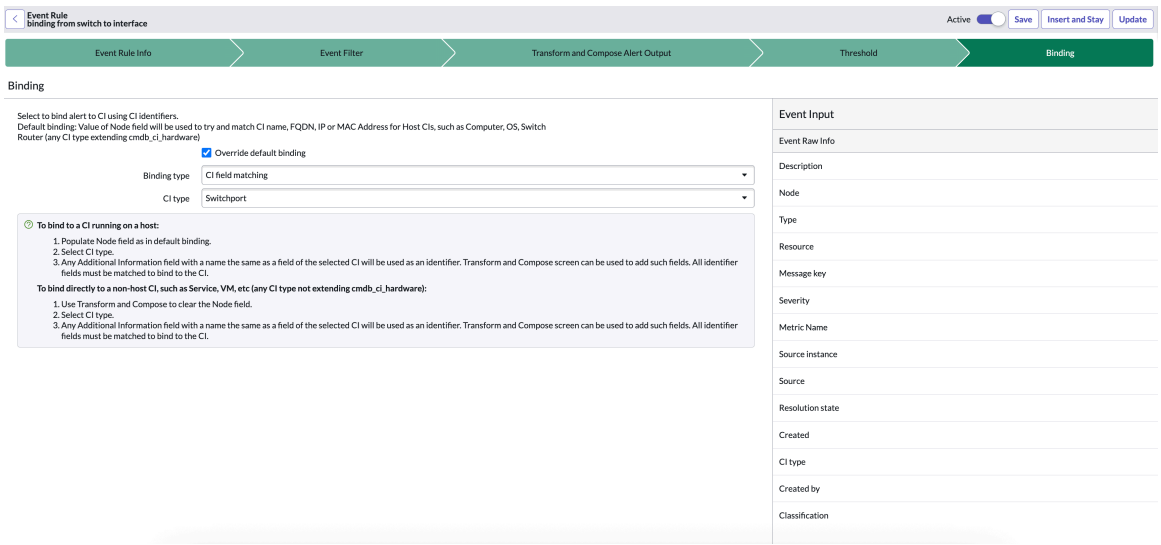

7. Navigate to **All** and search `em_binding_device_map.list`.
   The Binding Device Maps page opens. Here the column values **From ci type** is mapped to **To ci type** using the **Mapped column**. For example, **cmdb_ci_ip_switch** is mapped to the Switchport [dscy_switchport] table using the reference field **cmdb_ci**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

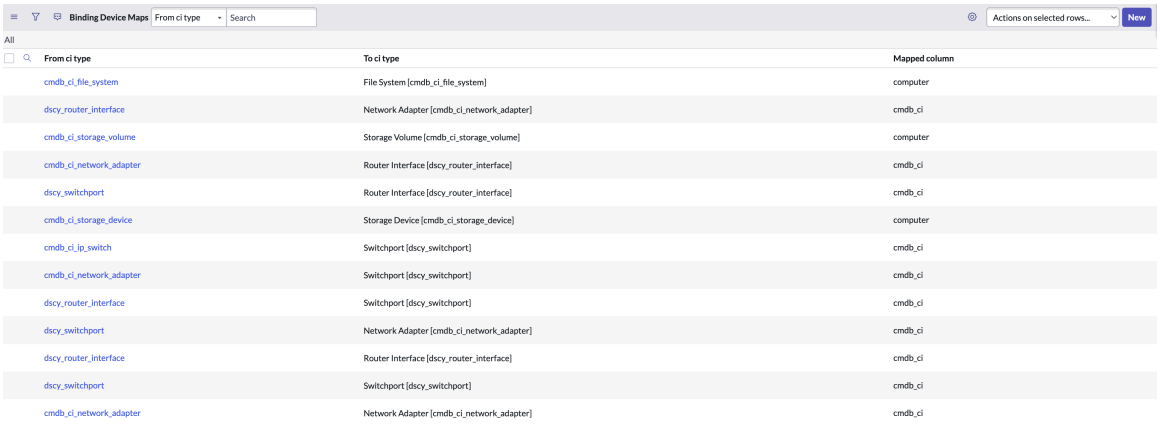

8. Navigate to **All** and search `switches` and select **Configuration** > **Network** > **Switches**. The IP Switches page opens.

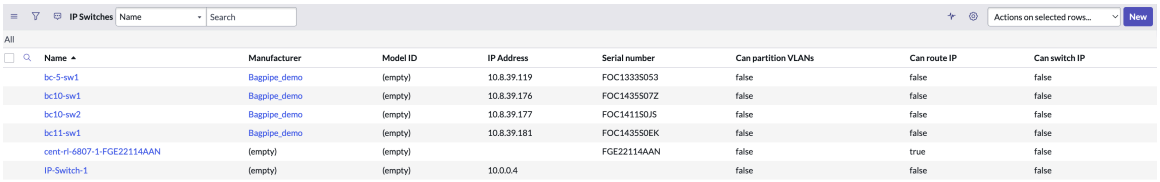

9. Select **IP-Switch-1** to view the details page.
10. Select the **Switchports** tab.
11. Navigate to **All** > **Event Management** > **All Events**.
12. Select the event with **Node** value as **IP-Switch-1**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In the following image, you can see the processing notes indicating that the binding has occurred between the alert and the matching device.

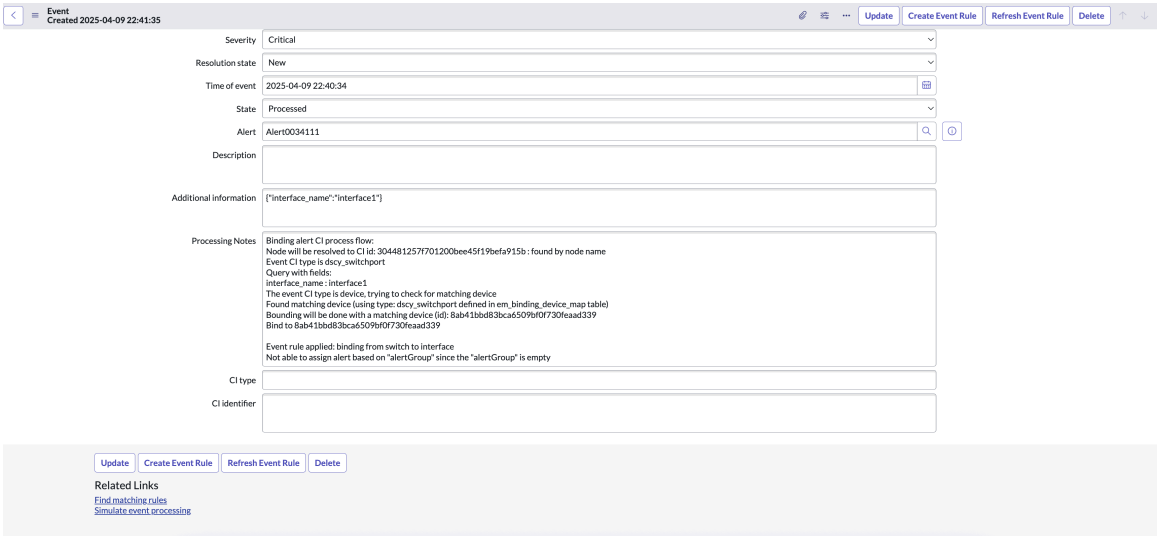

Bind alerts to a specific process

Bind specific server processes to their corresponding Configuration Items (CIs) in the CMDB to ensure accurate mapping and visibility. This binding is crucial for identifying service dependencies, reducing ambiguity from generic process names, and enabling effective monitoring. It supports faster alert resolution, impact analysis, and better alignment between infrastructure and application components in dynamic environments.

Before you begin

Role required: evt_mgmt_admin

About this task

Sometimes, when an alert (or event) comes into the system, it needs to be connected — or "bound" — to a Configuration Item (CI) in the CMDB. By default, the system binds alerts to host specified in the **Node** field of the event. Imagine a situation where you have a Windows server running multiple processes, like MSFT SQL Instances and SQL Server Analysis Services. The challenge is to bind an event to the specific process instance rather than just the host server. Multiple processes could have the same generic name, such as MSSQLSERVER, leading to ambiguity.

The following example procedure uses a Windows server as the host, MSFT SQL Instances as the CI class of the process, and MSSQLSERVER as the process name. The following steps assume that the event **Node** field provides the host name. The **Additional Information** field contains specific process details required for binding.

Alert-to-CI Binding via Process Mapping

Action	Steps
Set event rule: Add process name	Add the process name `sa_process_name` in the event rule.
Set event rule: Define CI type	Select the target CI type in the event rule. For example, MSFT SQL Instances.
Define process mapping	Navigate to the Process to CI Type Mapping table [em_binding_process_map] and add an entry mapping a CI type (e.g., MSFT

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert-to-CI Binding via Process Mapping (continued)

Action	Steps
	SQL Instances) to a process name (e.g., MSSQLSERVER).
Bind CI to alert	When an event is triggered, the system: **1.** Applies default binding to identify the host. **2.** Checks the CI type specified in the event rule. **3.** Looks up the Process to CI Type Mappings table [em_binding_process_map] for a matching process name and CI type. **4.** Searches the CI Relationships table [cmdb_rel_ci.list] for a Runs on::Runs relationship between the process and the host CI. **5.** Binds to a process when values from the **CI Type** column match the corresponding entries in the **Process** column in the Process to CI Type Mappings table [em_binding_process_map].

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Select **New**.
3. Select the **Transform and Compose Alert Output** tab and perform the following steps:

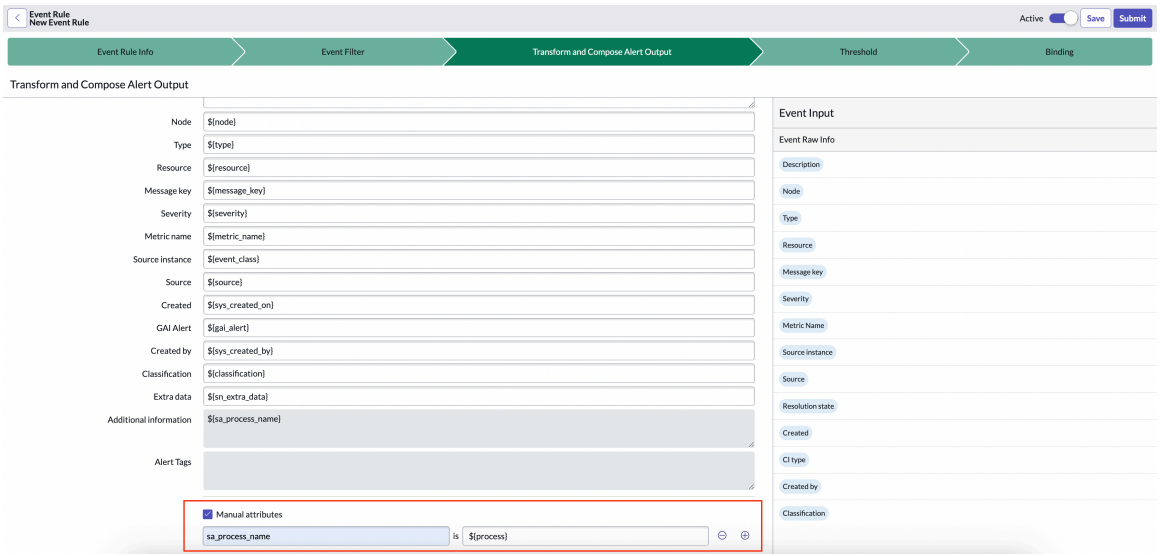

   a. Select the **Manual attributes** check box.
   b. Enter `sa_process_name` is `${process}`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

sa_process_name is a special variable name used to specify the name of the process to search for. ${sa_process_name} will appear in the **Additional information** field of the event. Instead of `${process}`, you can enter any other field name from which the variable `sa_process_name` derives its value.

4. Select the **Binding** tab.
5. Select the **Override default binding** check box.
6. In the **Binding type** field, select **CI field matching**.
7. In the **CI type** field, select **MSFT SQL Instances**.
   The **CI type** determines the specific CMDB table where the system searches for the matching CI.

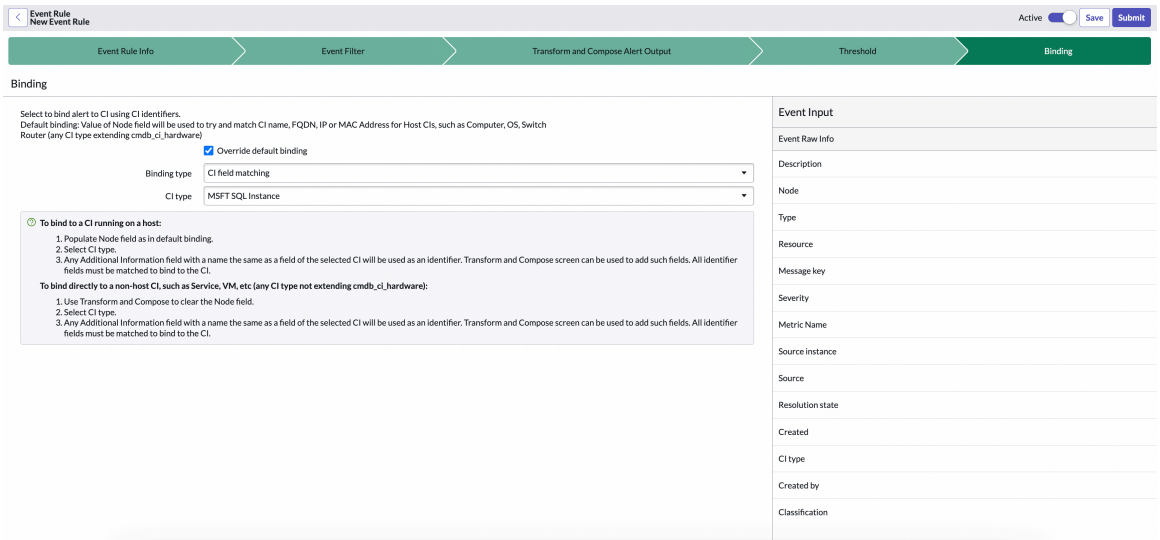

8. Navigate to **All** and search `em_binding_process_map.list`.

   The Process to CI Type Mappings page opens. Here the values from the **CI type** column are mapped to entries in the **Process** column. For example, **cmdb_ci_db_mssql_instance** is mapped to the **MSSQLSERVER** process.

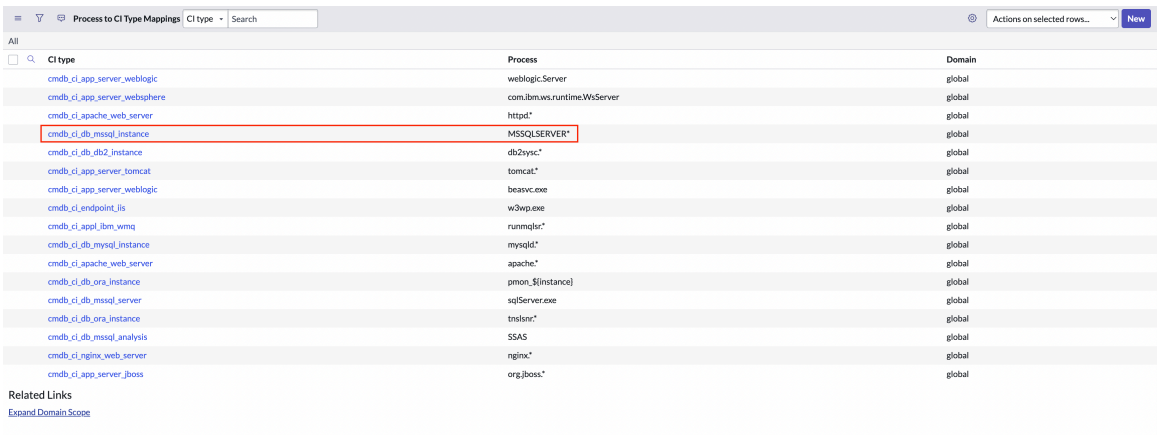

9. **Optional:** Navigate to **All** and search `cmdb_rel_ci.list`.
   The CI Relationships page opens.
10. **Optional:** Verify that the node is in a Runs on::Runs relationship with the appropriate process.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

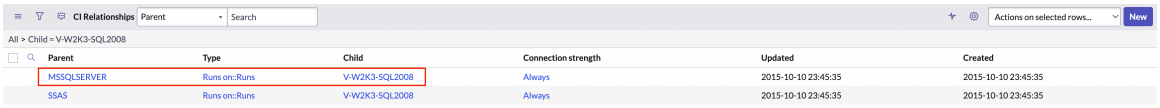

11. Navigate to **All** > **Event Management** > **All Events**.
12. Create an event with **Node** value as the name of the host such as Windows Server (V-W2K3-SQL2008).
    The **Additional information** field in the event has a key called **process** whose value is **MSSQLSERVER**.

    In the following image, you can see the processing notes indicating that the binding has occurred between the alert and the matching process.

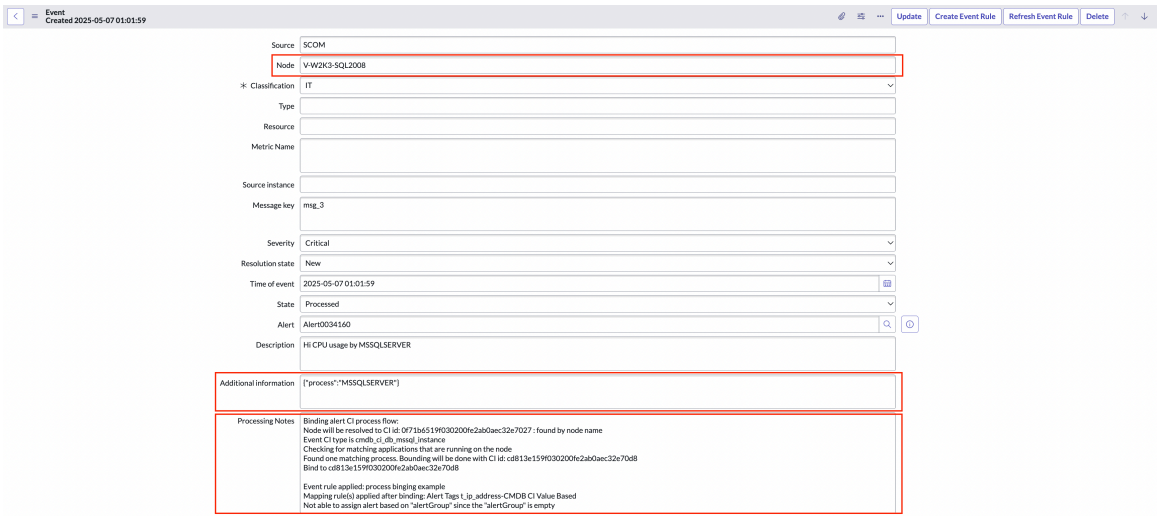

Bind non-host CIs using CI field matching

If no match is found using the **Node** field, the system uses the **CI identifier** field to match alerts with non-host CIs based on attributes like Name, FQDN, IP, or MAC. This ensures accurate alert association, improving visibility, troubleshooting, and root cause analysis for diverse infrastructure components.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Enrich automation in Service Operations Workspace, the updated way to transform and standardize alert data for better response.

If no match is found using the **Node** field, the system looks at the **Additional information** field of the alert. When you select a CI type, such as File System, the system automatically searches for a matching record in the [cmdb_ci_file_system] table. It uses the details provided in event rule record, specifically in the **Additional information** field, to refine the search. For example, if the **Additional information** field contains values like `{"mount_point": "/snap/amazon-ssm-agent/9565", "name": "/dev/loop0"}`, the system looks for a record in the [cmdb_ci_file_system] table that matches these values. If a match is found, the system binds the CI to the alert, ensuring accurate identification and association. Similarly, if the **CI type** is **Network Adapter**, the system searches in the [cmdb_ci_network_adapter] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

There may be cases where no match is found because the column names in the event record and the table differ for the same item. In such cases, you can manually create an additional key-value pair with a name matching the table column, ensuring the matching process continues successfully. For information on how to create a manual field, see Bind CIs using CI field and column matching.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Select **New** and complete the required fields of the event rule.
3. Select the **Binding** tab.
4. Select the **Override default binding** check box.
5. In the **Binding type** field, select **CI field matching**.

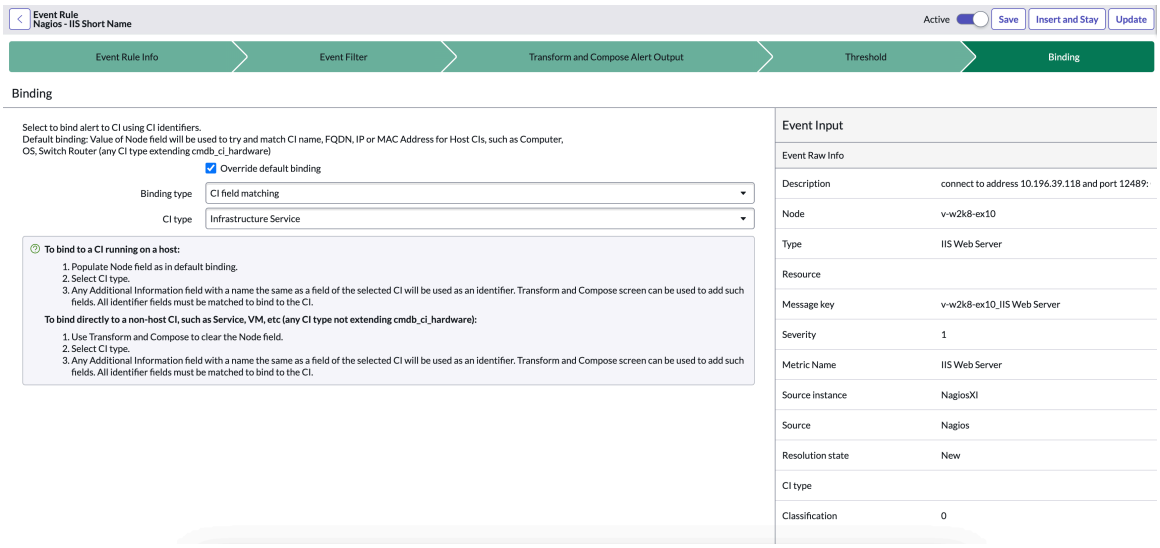

6. Select the **Transform and Compose Alert Output** tab.
7. Clear the **Node** field value.

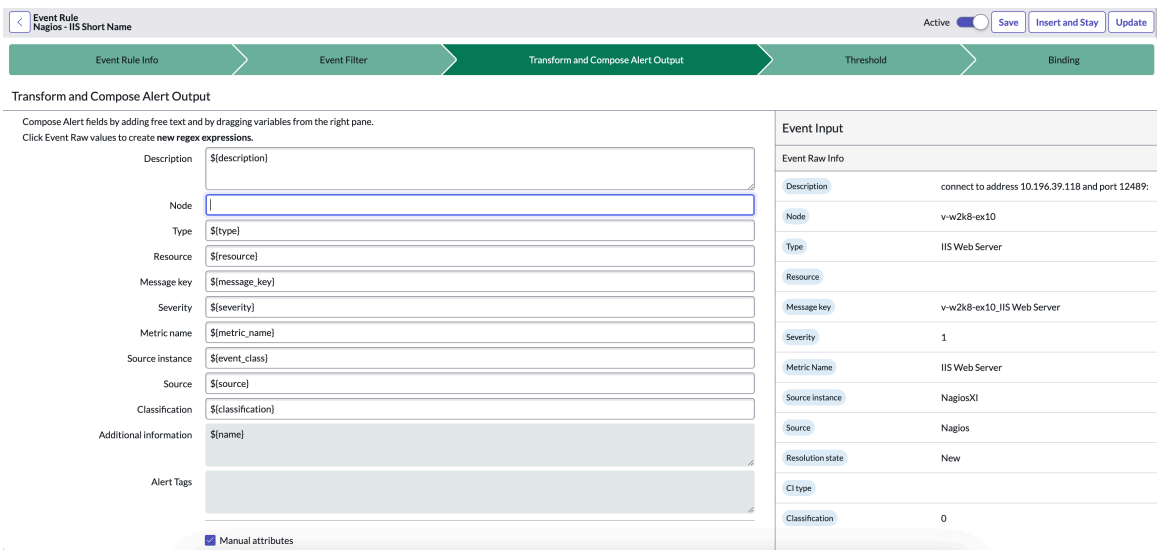

8. In the **CI type** field, select the non-host CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A non-host CI type is any CI type that does not extend the [cmdb_ci_hardware] table, such as a Service or VM. The **CI type** determines the specific CMDB table where the system searches for the matching CI.

9. Select **Save**.

Bind CIs using CI field and column matching

Bind CIs by matching event **Additional information** fields with CI attributes. If column names differ, manually create an additional key-value pair to align with the CI table, ensuring accurate CI association.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Enrich automation in Service Operations Workspace, the updated way to transform and standardize alert data for better response.

If no match is found using the **Node** field, the system looks at the **Additional information** field in the event record. There may be cases where no match is found because the column names in the event record and the table differ for the same item. In such cases, you can manually create an additional key-value pair with a name matching the table column and add it to the **Additional information** field, ensuring the matching process continues successfully.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Select **New** and complete the required fields, or open an existing event rule record to modify it.
3. Select the **Transform and Compose Alert Output** tab.
4. Select the **Manual attributes** check box.
5. In the **Field Name** and **Field Value** fields, enter the corresponding field name and specify the source field from which the value should be populated.

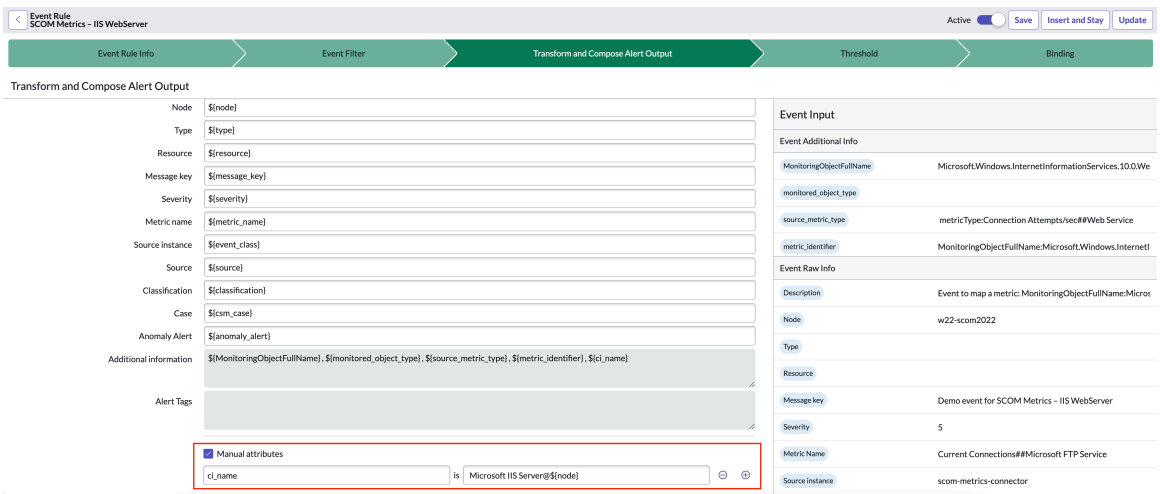

Use the add icon (⊕) icon to add multiple fields as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Select **Save**.
   The fields are created in the event rules record. However, when any alert is generated from an event using the event rule, the manually added fields appear in the **Additional information** field of the alert.

Bind alerts to CIs using CI identification

Bind alerts to specific applications on hosts using event rules to ensure accurate tracking and to improve issue resolution speed—leading to efficient remediation and better alignment of alerts with the right resources.

Before you begin

Role required: evt_mgmt_admin

About this task

Before going through the procedure, let's first understand how this binding works with an example. Suppose an out-of-the-box event rule, SCOM Metrics – IIS WebServer, generates an alert for a Microsoft IIS Web Server running on a host. Instead of binding the alert to just the host machine, you want it to be associated with the specific application (IIS Web Server) where the issue occurred. Here's how the system ensures accurate binding.

Let's say in the **Binding** tab, the **Class** and **Container Level 1** are defined along with their respective attributes, as shown in the image.

Binding tab

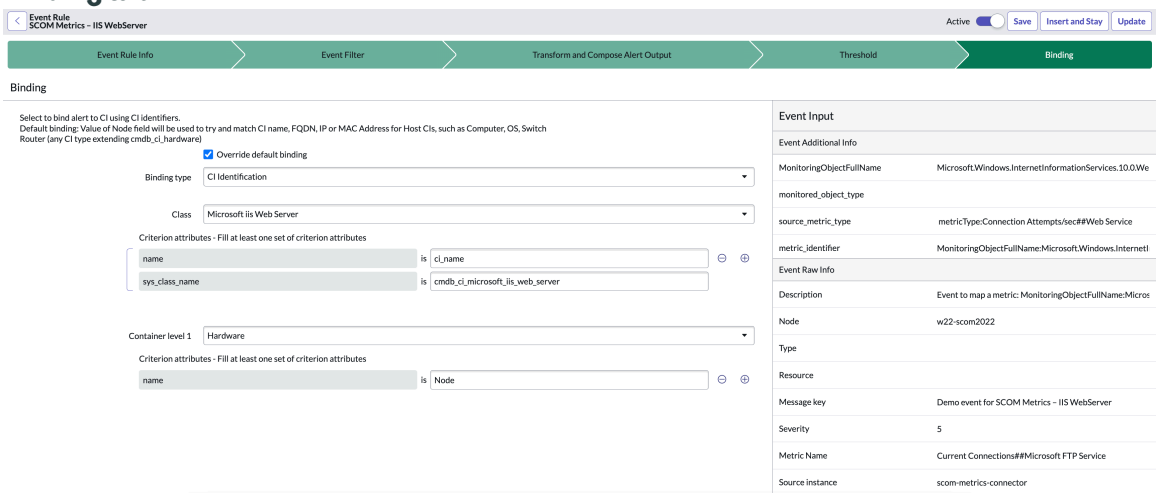

The binding procedure follows these steps:

1. Identifying the Host (Container Level 1)

   Since the Microsoft IIS Web Server runs on a host (in this case Hardware), the system first locates the host, which is defined as **Container Level 1**. The name field is used to identify the host, and its value is taken from the **Node** field of the event. Using this name, the system searches the CMDB to find a matching host CI.

2. Identifying the Microsoft IIS Web Server CI on the Host

   Once the host is identified, the system looks for a Microsoft IIS Web Server CI under that host with the following attributes:

   - name = ci_name
   - sys_class_name = cmdb_ci_microsoft_iis_web_server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Here, ci_name is derived from the key-value pair defined in the **Transform and Compose Alert Output** tab of the event rule. Assume that in the **Transform and Compose Alert Output** tab of the event rule, the **Manual attributes** check box is selected and the following key-value pair is already added:

- **Field Name**: `ci_name`
- **Field Value**: `Microsoft IIS Server@${node}`

This ensures that `ci_name` appears in the **Additional information** field of the alert and is used when configuring the attribute field of the CI class.

Transform and Compose Alert Output tab

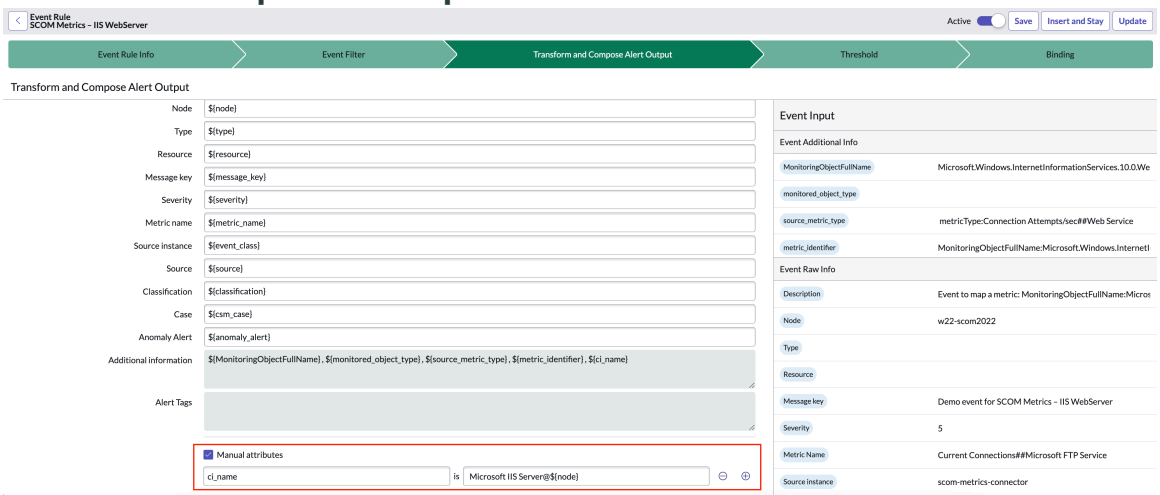

If a matching Microsoft IIS Web Server CI is found on the identified host, the alert is bound to the correct application CI rather than just the host. This ensures more precise incident tracking and resolution, linking the alert to the actual service affected instead of only the machine it runs on.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Select **New** and complete the required fields of the event rule.
3. Select the **Binding** tab.
4. Select the **Override default binding** check box.
5. In the **Binding type** field, select **CI Identification**.
6. In the **Class** field, select a CI class.
7. In the **Select an identifier** field, select an identifier for the class.

   An identifier for a class is a set of rules that helps the system recognize and match CIs correctly. Think of it like a fingerprint for each type of CI (such as a server, router, or application). The identifier uses key attributes—like a serial number, hostname, or IP address—to make sure that when new data comes in, it gets linked to the right CI instead of creating duplicates. For more information on CI identification rule, see Create a CI identification rule.

8. In the **Add value** field, add the exact value or specify the source field from which the value for the field must be populated.

   Add or remove fields using the add (⊕) and the delete (⊖) icons.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: While filling in the CI Class attributes, if you do not find the attribute you need, you can use the **Transform and Compose Alert Output** tab of the event rule to add the necessary key-value pair. For more information, see Configure an event rule to customize alert content.

The event rule transformation and composition is useful when relevant information exists in a different field but needs to be extracted or reassigned for proper alert processing. This is especially important when the required attribute is not directly available and must be manually defined.

For example, suppose the **Additional information** field of an alert contains the alert name under the **Node** field. However, for correct processing, we need this value to appear under the **Name** field instead. Since there is no existing name attribute in the event, we can manually create one using the **Transform and Compose Alert Output** tab in the event rule. Here's how:

a. Select the **Transform and Compose Alert Output** tab.

b. Select the **Manual attributes** check box.

c. Add the following key-value pair:

- **Field**: Name
- **Field Value**: Node

This configuration ensures that the system extracts the **Node** value and maps it to the name field, making it available for CI binding and alert correlation.

9. In the *Container level 1* area, specify the host.
10. In the **Select an identifier** field, select an identifier for the container.
11. In the **Add value** field, add the exact value or the field that is the source for the **Select an identifier** field of the container.

    Add or remove fields using the add (⊕) and the delete (⊖) icons.
12. Select **Save**.

Example: Binding alerts to non-host CIs

Bind alerts to an application service (a non-host CI) using event rules and event field mapping. This example demonstrates how to achieve this by leveraging the IP address associated with the service.

Before you begin

Role required: evt_mgmt_admin

About this task

When an event is generated, it may not always include the exact name of the Application Service it affects. Instead, it might provide details such as an IP address. In this example, suppose application services are identified by IP addresses. Since the goal is to bind the alert to an Application Service, you need to determine which service name corresponds to the given IP address.

To achieve this, follow these steps:

1. Match the correct Application Service: Override the default binding and select **Binding Type** as **CI field matching**. The binding process needs to associate the alert with an Application Service, not a host. In the event rule, set the **CI Type** to **Application Service** so that the system searches for the service in the Application Service table (cmdb_ci_service_auto).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Extract the IP address from the event description: The event might contain a message like `Service is down. IP address is: 196.1.1.24`. You need to extract the IP address from this description for further processing.

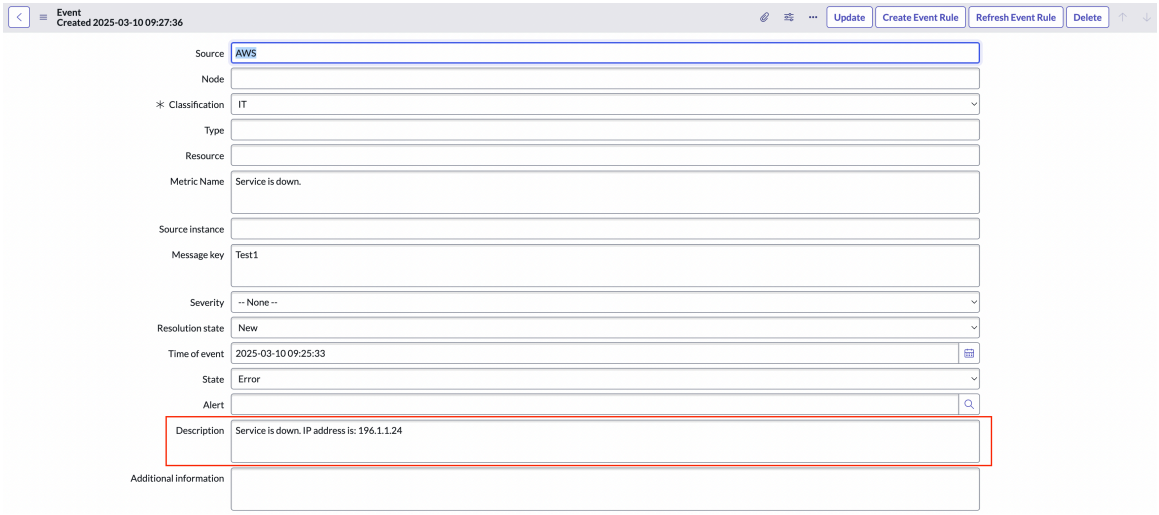

3. Translate the IP address into a service name using Event Field Mapping: Use Event Field Mapping to look up the extracted IP address in the CMDB and find the associated application service. Once a matching service is found, its name is stored in the **Name** field of the event.
4. Bind the alert to the correct Application Service: If the service name is identified, the alert's CI is populated with that service.

Procedure

1. Navigate to **All** > **Event Management** > **All Events**.
2. Open the specific event for which you want to create an event rule.
3. Select **Create Event Rule**.
   The new event rule page opens. Some fields in the event rule are populated based on the event record data, such as the Source field value.
4. In the **Event Rule Info** tab, in the **Name** field, enter `application service binding`.
5. Select the **Binding** tab.
   In this tab, select the CI type as a service (non-host) instead of a host. The goal is to bind alerts to an Application Service, not to a physical or virtual machine.
6. Perform the following steps:
   a. Select the **Override default binding** check box.
   b. In the **Binding type** field, select **CI field matching**.
   c. From the **CI type** list, select **Application Service**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

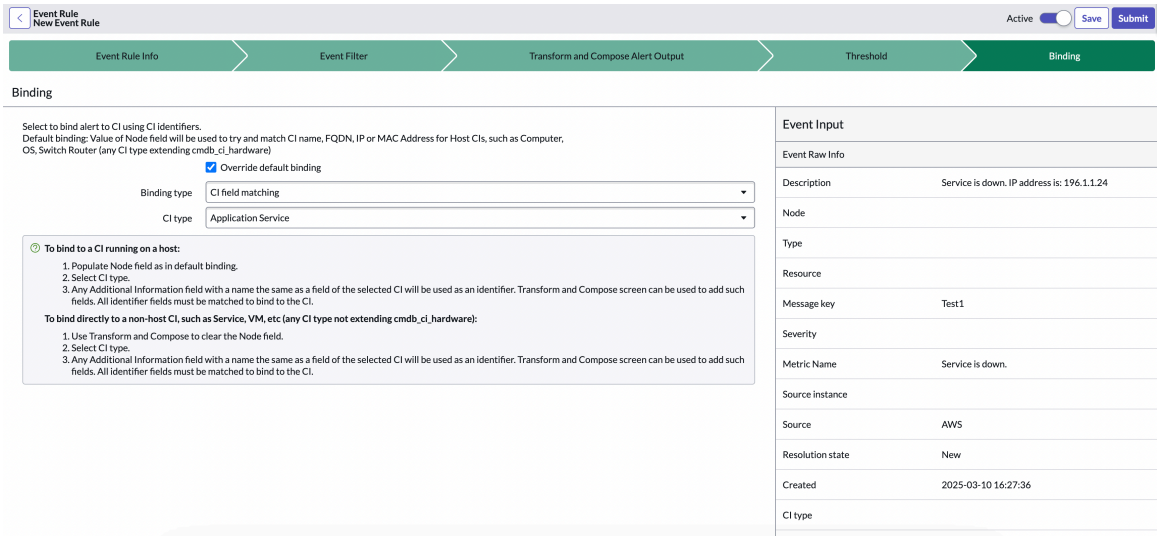

7. Select the **Transform and Compose Alert Output** tab.

   In this tab, you extract the IP address of the application service into a manually created field. Extracting the IP allows the system to look up the corresponding service in the Application Service table (cmdb_ci_service_auto). Without extracting the IP, the system wouldn't be able to determine which Application Service the event belongs to, making it difficult to bind the alert correctly.

8. Perform the following steps:

   a. Clear the **Node** field.

      When an event is processed, it often comes with multiple fields of information. The **Node** field typically contains the host or device name. However, the goal is to bind the alert to an Application Service rather than a host. Clearing the **Node** field ensures that the system focuses on identifying the correct service instead of mistakenly associating the alert with a host CI.

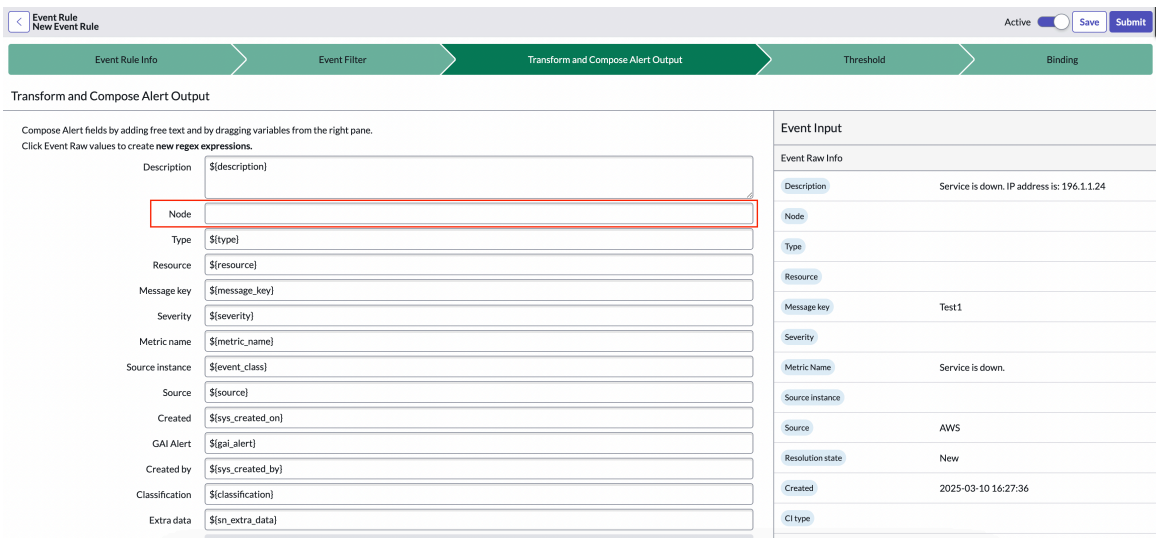

   b. Under **Event Input** > **Event Raw Info**, select **Description**.

      The Edit Regex Expressions window opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. In the Mark Expressions section, select the IP address, manually enter a field name (e.g., service_ip_address), and press Enter to save it.

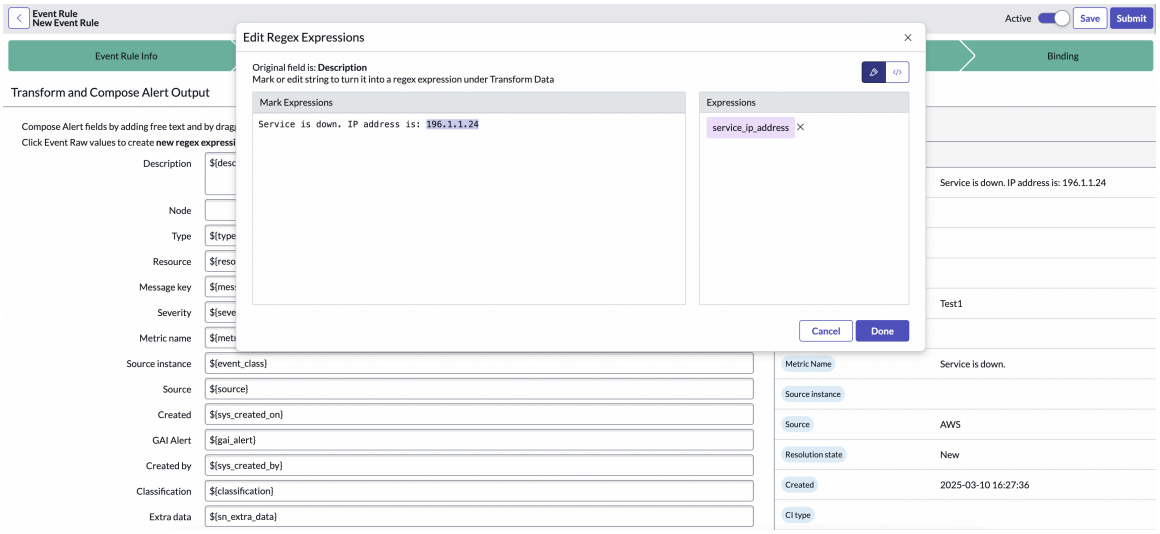

d. Select **Done**.

The IP address is extracted into the newly created **service_ip_address** field, which is added to the **Additional info** field of the alert. This allows it to be used later to identify the corresponding Application Service.

9. Select **Submit** in the event rule to save your changes.

10. Navigate to **All** > **Event Management** > **Rules** > **Event Field Mapping**.
In Event Field Mapping, configure the extracted IP address to map it to the corresponding service name. The event provides only an IP address, but alerts need to be associated with a service name (not just an IP). The system needs to map this IP to its corresponding Application Service so that teams can easily identify the impacted service.

11. Select **New**.
The Event Field Mapping window opens.

12. Perform the following:

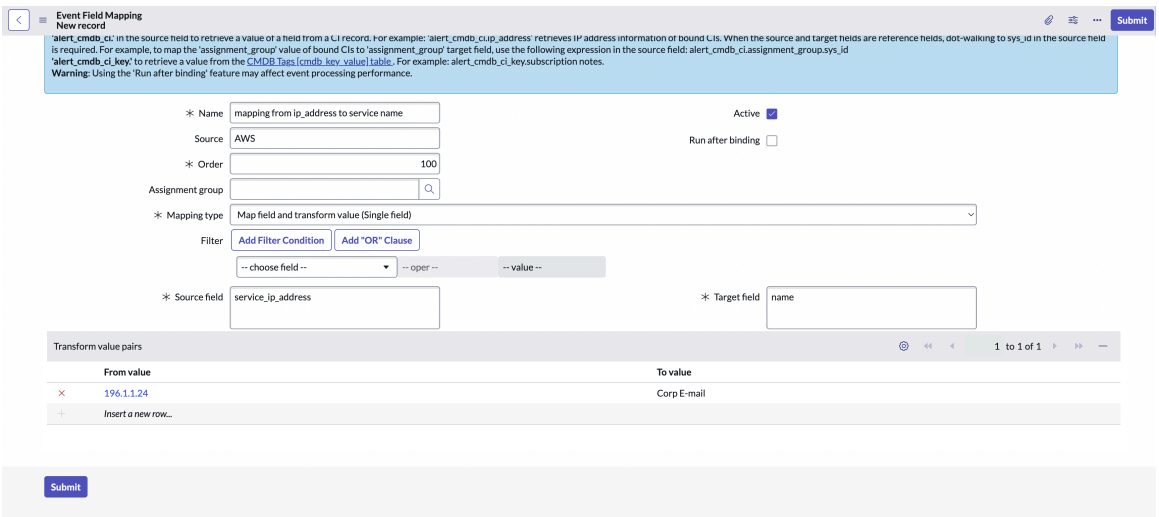

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Action
Name	Enter the name of the mapping. Example: mapping from ip_address to service name.
Source	Enter `AWS`. **Note:** The source is the same as the one in the Events record.
Mapping type	Select **Map field and transform value (Single field)**
Source field	Enter `service_ip_address`.
Target field	Enter `name`. The target field value is the column name in the Application Service table (cmdb_ci_service_auto) where the source field value is searched.
Transform value pairs	
From value	196.1.1.24
To value	Corp E-mail

13. Select **Submit**.

The system extracts the IP address of the service from the event and stores it in the **service_ip_address** field. The Application Service table (cmdb_ci_service_auto) does not have a column for searching by IP address. The Transform value pairs feature maps the IP address to a name—in this case, Corp E-mail—and saves it in the **name** field within the **Additional info** field of the alert. The **Target field** specifies where to search for the Corp E-mail in the Application Service table. The system looks for Corp E-mail in the **name** column of the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application Service table (cmdb_ci_service_auto). If a match is found, the alert is bound to the corresponding Application Service.

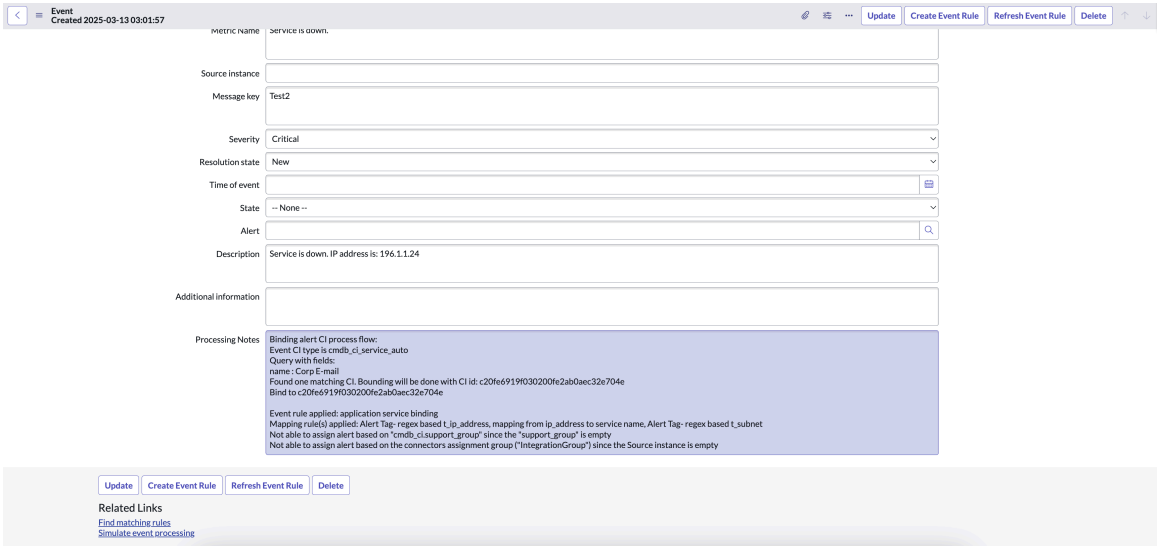

Example: Bind alerts to CIs using dynamic CI types

Use event field mapping to dynamically bind alerts to the appropriate CIs based on event attributes, eliminating the need for separate event rules for each CI type (also known as CI class). This approach simplifies configuration, improves accuracy, and enhances alert to CI binding.

Before you begin

Role required: evt_mgmt_admin

About this task

In general, for an event generated from a specific source mentioned in the **Source** field, the system applies the corresponding event rule. The override default binding option in an event rule specifies a specific CI class for binding. However, in a more dynamic use case, the CI class is not predefined and can change based on event attributes. For example, suppose a customer has multiple CIs with the same name, but each CI belongs to a different CI class. In the traditional event rule method, as the CI class (or the CMDB table) is defined for a CI, the mapping is incorrect because the system searches in the same CI class for all CIs with the same name and attributes, instead of looking in the specific CMDB table defined for each CI class.

To resolve this—in other words, to determine the correct CI class to search—the system uses the **CI type** field from the event to identify the CI class. By using an event field mapping that maps from event attributes to CI type, it becomes possible to know which CMDB table to search, ensuring that the CI is correctly mapped to the event.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Field Mapping**.

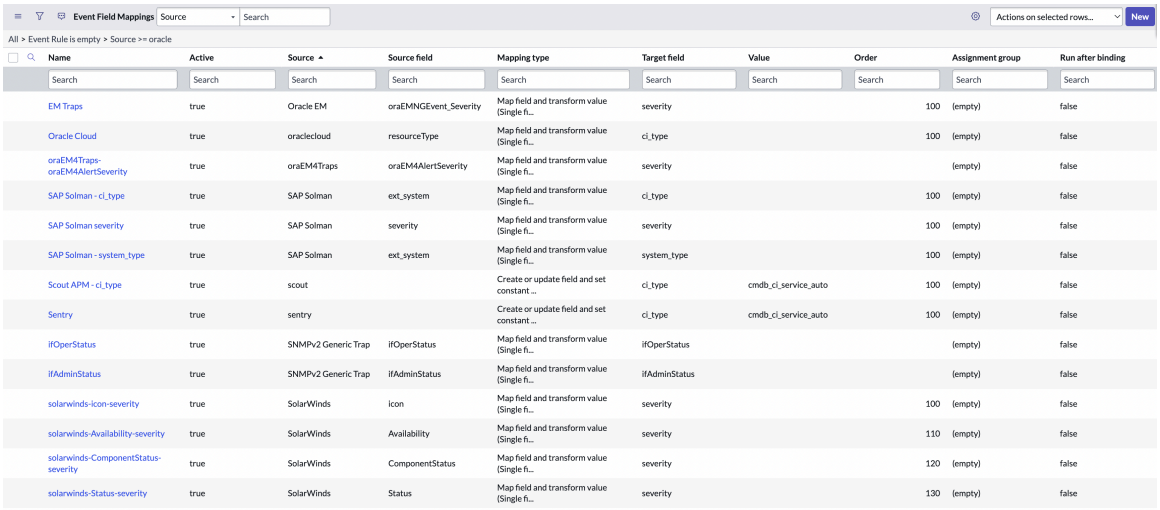

Name	Active	Source	Source field	Mapping type	Target field	Value	Order	Assignment group	Run after binding
EM Traps	true	Oracle EM	oraEMNGEvent_Severity	Map field and transform value (Single fi...	severity		100	(empty)	false
Oracle Cloud	true	oraclecloud	resourceType	Map field and transform value (Single fi...	ci_type		100	(empty)	false
oraEM4Traps-oraEM4AlertSeverity	true	oraEM4Traps	oraEM4AlertSeverity	Map field and transform value (Single fi...	severity			(empty)	false
SAP Solman - ci_type	true	SAP Solman	ext_system	Map field and transform value (Single fi...	ci_type		100	(empty)	false
SAP Solman severity	true	SAP Solman	severity	Map field and transform value (Single fi...	severity		100	(empty)	false
SAP Solman - system_type	true	SAP Solman	ext_system	Map field and transform value (Single fi...	system_type		100	(empty)	false
Scout APM - ci_type	true	scout		Create or update field and set constant ...	ci_type	cmdb_ci_service_auto	100	(empty)	false
Sentry	true	sentry		Create or update field and set constant ...	ci_type	cmdb_ci_service_auto	100	(empty)	false
ifOperStatus	true	SNMPv2 Generic Trap	ifOperStatus	Map field and transform value (Single fi...	ifOperStatus			(empty)	false
ifAdminStatus	true	SNMPv2 Generic Trap	ifAdminStatus	Map field and transform value (Single fi...	ifAdminStatus			(empty)	false
solarwinds-icon-severity	true	SolarWinds	icon	Map field and transform value (Single fi...	severity		100	(empty)	false
solarwinds-Availability-severity	true	SolarWinds	Availability	Map field and transform value (Single fi...	severity		110	(empty)	false
solarwinds-ComponentStatus-severity	true	SolarWinds	ComponentStatus	Map field and transform value (Single fi...	severity		120	(empty)	false
solarwinds-Status-severity	true	SolarWinds	Status	Map field and transform value (Single fi...	severity		130	(empty)	false

2. Select **Oracle Cloud**.

   This is an existing event field map used as an example. The event field map details page opens.

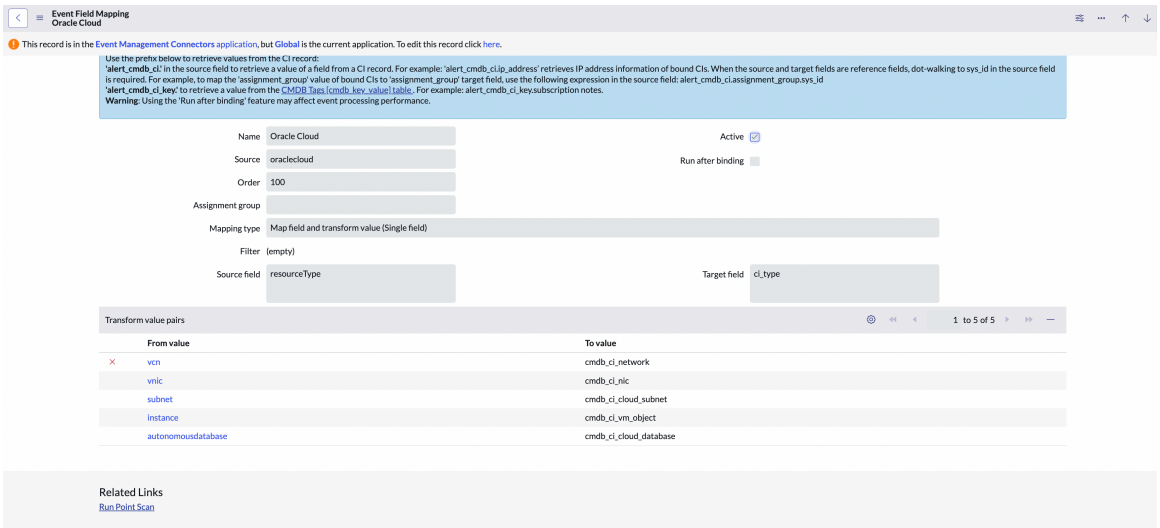

3. Review the event field mapping details page and understand the relevant fields and their corresponding values.

Event Field Mapping fields

Field	Description
Name	Event field map name.
Source	Event monitoring software that generated the event. In this case, oraclecloud.
Order	Number to define the order in which this action should be processed. Actions with lower numbers are processed first.
Mapping type	Select **Map field and transform value (Single field)**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Source field	**resourceType** is a field that is defined in the **Additional Information** field of the event defined by the Oracle Cloud system. The possible values of **resourceType** are listed in the **From value** column, while the corresponding CMDB tables for those values are specified in the **To value** column in the Transform value pairs section. If **resourceType** is **instance**, the system searches in the [cmdb_ci_vm_object] table for CI binding.
Target field	The **CI type** field (**ci_type**) in an event is a special field that provides the CI class name, indicating which CMDB table to search. The **CI type** field values are listed in the **To value** column in the Transform value pairs section.
Transform value pairs	
From value	Lists the possible values of the event field (e.g., `resourceType`) that are derived from the event data.
To value	Specifies the corresponding CMDB table associated with each value listed in the **From value** column. Indicates the CI type to bind the alert to.

View patterns for event group creation

Event groups are sets of events that do not have a matching event rule. You can view the patterns in a group of events to learn the impact of creating a rule based on the event source and description patterns.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

Events are grouped according to their source and description.

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules**.
2. Wait for the group calculation to complete and then click the created link.
   If there are not enough events, the calculation will not create groups.
3. In the list of events, select an event.
4. To change the pattern of an event, double-click the pattern, make your changes, and then click **Update**.

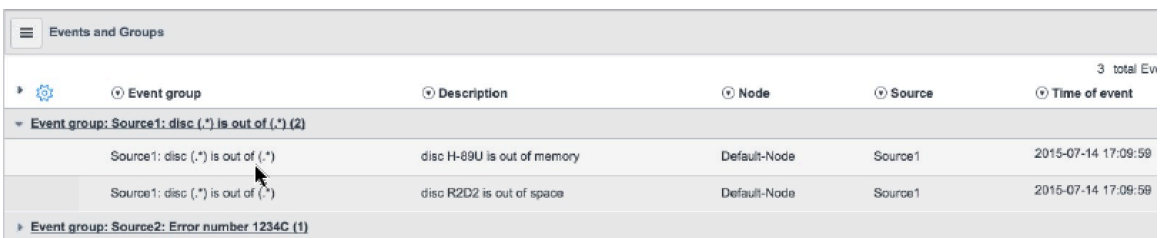

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Refresh event rules

Manually update event rules to reflect current event information because once an event rule is created, the **Event Additional info** and **Event Raw info** fields are not automatically updated.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

Procedure

1. Navigate to **All** > **Event Management** > **All Events**.
2. Select the event.
3. Select **Refresh Event Rule**.
4. Select the event rules to refresh.
5. Select **Refresh Selected**.

   If you selected multiple event rules, then the last selected item in the list opens in the Event Rule Designer.

   Note: If the current event rule regex expressions do not match the selected event after the refresh, the update fails and an error message appears specifying the problematic fields. This safeguard ensures that your previously defined event rule regex expressions are not broken due to the refresh.

Find rules that will be applied to an event

View the rules that will be applied on an event to determine how this event will be processed.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

Procedure

1. Navigate to **Event Management** > **All Events**, and then click an event number.
2. Under **Related Links**, click **Find matching rules**.
3. Review the information that appears in the blue box at the top of the page, and review all **Processing Notes** to find rules that apply to an event.

Related topics

View events

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Simulate event processing

Simulate event processing logic to understand which rules execute on an event and how event fields change.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: Use Ignore automation and Enrich automation in Service Operations Workspace, the updated way to automatically filter out irrelevant alerts, or to transform and standardize alert data for better response.

Procedure

1. Navigate to **All** > **Event Management** > **All Events**.
2. Open the event you want to simulate.
3. In the **Related Links** section, select **Simulate event processing.**
4. In the info message that appears at the top of the screen, select the link to the created simulated alert.
5. **Optional:** View which event rules and event field mapping rules were applied and the event field values by selecting the **Activities** tab.

Event field mapping configuration

Use Event field mappings rules to map values from specific fields to values in other fields.

Event Management stores event field mappings in the Event Field Mapping [em_mapping_rule] table. The mappings apply after event rule processing and prior to alert generation. The mapping values from the Event Mapping Pair [em_mapping_pair] table apply to the alert. The original event severity remains unchanged.

For example, if events come with the field "**org_severity**" with the values "*Low*, *Medium*, *High*" and you want the alert **Severity** to hold this value, create an event field mapping rule that maps the field **org_severity** to **Severity**, with these values:

Example showing org_severity mapped to Severity

Original org_severity value	Map to alert Severity value
Low	Warning
Medium	Major
High	Critical

Default event field mappings

Event Management provides default event field mappings for commonly used system monitoring tools. The Transform Value Pairs from event field mappings format the incoming event data for Event Management.

You can view the default event field mappings and mapping pairs by navigating to **Event Management** > **Rules** > **Event Field Mappings** and double-clicking **Name**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create event field mappings

Use event field mappings to map values from specific event fields to values in other fields to provide more comprehensive information in an alert. Use team-based integrations in event rules to make sure that connector ownership and execution of rules give precedence to global rules. Teams can maintain consistency and hierarchy while offering flexibility and customization options.

Before you begin

Role required: evt_mgmt_admin

About this task

Create the rule to match the event by its class and original values. Also, specify the new values to replace the original values in the event.

For users familiar with the classic Event Management experience, enrich automations create event rules but with an easier interface and with better team support. For more information, see Create Enrich automation.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Field Mapping**.
2. Select **New** or open an existing rule to edit and fill in the fields.

Event Field Mapping form

Field	Description
Name	Event field mapping name.
Source	Event monitoring software that generated the event, such as SolarWinds or SCOM. Maximum length: 100 characters.
Order	Number to define the order in which this action should be processed. Actions with lower numbers are processed first.
Assignment group	For Team-based integrations, select an assignment group. If no assignment group is defined in the event field-mapping rule, then this rule is considered global. When the rules are running – first the global rules run and then the rules that belong to the assignment group that the event's source instance belongs to.
Mapping type	Mapping mechanism that is used to change an event field value. ○ **Map field and transform value (Single field)**: Maps the **Source field** value to the matching value in the Transform value pairs [em_mapping_pair] table and populates the **Target field**. Source and Target fields may be event fields, additional info fields, or alert tags.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ **Create or update field and set constant value (Constant)**: Sets a constant value to a **Target field**. The **Target field** may be an event field, additional info field, or alert tag. ○ **Map field (Copy field)**: Copies the exact value of a Source field to the specified Target field. The Source and Target fields may be Event field, Additional info field, or alert tags. ○ **Map field and transform value using regex**: Generically maps the value from the Source field to the matching value in the Transform value pairs table and populate the **Target field**. Source and Target fields may be event, additional info, or alert tags fields. ○ **Map field using regex**: Copies the value from a generically defined **Source field** to the **Target field**. The Source field must be a field from the event additional info field. The Target field may be the field of event, additional info, or alert tag. With this type, you can generically find the name of the Source field by using a regular expression (regex). ○ **Map field and copy value from regex group**: Copies part of the Source field value to the Target field. Source and Target fields may be event fields, additional info fields, or alert tags. Define a regex to specify what part of the **Source field** should be copied. ○ **Advanced mapping using script**: Implements your own logic for event values transformation by using a script. When the 'Run after binding' check box is selected, the rule is executed after the CI binding phase. The dedicated script also has the sys_id of the CI in the input parameters. **Warning:** Complex scripts may affect event processing performance. ○ **Map field from reference table** Use this rule type to extract a **Source field** value from the selected **Reference table** by matching the **Field to match**value to the **Matching alert field** value and then mapping the **Source field** value to the **Target field**. The **Matching alert field** and the **Target field** may be event fields, additional info fields, or alert tags.
Filter	A condition for event field mapping. Select **Add Filter Condition** or **Add "OR" Clause** to configure multiple conditions. **Note:** The filter is case-sensitive.
Active	Activates or deactivates the event field mapping. If possible, find and apply another event field mapping rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Run after binding	○ When selected, the rule is executed after the CI binding phase. ○ Use the relevant prefix to retrieve values from the CI record: ▪ **alert_cmdb_ci.** to retrieve a field from the CI record. For example: Use 'alert_cmdb_ci.assignment_group' to populate (and transform) an alert 'assignment_group' field with the value of the CI's assignment group. Dot-walk to a related record is also possible. For example, use alert_cmdb_ci.asset.model ▪ **alert_cmdb_ci_key.** to retrieve a value from the CMDB Tags table. For example, alert_cmdb_ci_key.subscription notes. **Warning**: Using the 'Run after binding' feature may affect event processing performance. **Note:** This check box is displayed only for the following Mapping types: **Map field and transform value (Single field)**, **Map field (Copy field)**, **Map field and transform value using regex**, **Map field and copy value from regex group**, and **Map field from reference table**.

○ If you selected **Map field and transform value (Single field)**, fill in the fields as appropriate.

Map field and transform value (Single field) fields

Field	Description
Source field	Event field to map - includes all event and additional info fields. Note: For Additional info fields use: "additional_info.<field name>". When the 'Run after binding' check box is selected, you can use the **alert_cmdb_ci** and the **alert_cmdb_ci_key** prefixes to receive CI related data.
Target field	Event field where the mapping rule inserts or updates the value. When this field is identical to the **Source field**, the mapping rule updates the value in memory of the event field. Alert tags can also be defined by using the *<t_>* prefix.
Run after binding	▪ When checked, activates Event field mapping after CI binding. ▪ Use the following prefix to retrieve values from the CI record: ▪ **alert_cmdb_ci.** to retrieve a field from the CI record. For example, use 'alert_cmdb_ci.assignment_group' to populate (and transform) an

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	alert 'assignment_group' field with the value of the CI's assignment group. Dot-walk to a related record is also possible. For example, use alert_cmdb_ci.asset.model ▪ **alert_cmdb_ci_key.** to retrieve a value from the CMDB Tags table. For example, alert_cmdb_ci_key.subscription notes. **Warning**: Using the 'Run after binding' feature may affect event processing performance.

Key fields (Transform Value Pairs section)

Field	Description
Key	Value that the mapping rule searches for. Whenever the event field has this value, the mapping rule adds the value listed in the **Source field** field to the field listed in the **Target field**. This field is displayed when the **Mapping type** is **Map field and transform value (Single field)**. Select **+** to add more **Key** fields, as required.

Typical use case: To map the 'event_severity' field in the additional info field with possible values [such as WARNING, MINOR, CRITICAL, CLEAR, MAJOR] to the severity field, configure the following:

- Source field: 'event_severity'
- Target field: 'severity'
- Transform value pairs:
  - From value: 'WARNING', to value: '4'
  - From value: 'MINOR', to value: '3'
  - and so forth...

- If you selected **Create or update field and set constant value (Constant)**, fill in the fields as appropriate.

Create or update field and set constant value (Constant) fields

Field	Description
Target field	The **Target field** may be an an event field, additional info field, or alert tag. Typical use case: Set alert severity to 1 for all events that match the filter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Value	Value you want to use for the **To field**. This field is displayed when the **Mapping type** is **Constant**.

Typical use case: Set alert severity to 1 for all events that match the filter.

- If you selected **Map field (Copy field)**, fill in the fields as appropriate.

Map field (Copy field) fields

Field	Description
Source field	**Source field** value to be copied to the **Target field** field. The **Source field** may be an event field, additional info field, or alert tag.
Target field	Event field where the mapping rule inserts or updates the value. The **Target field** may be an event field, additional info field, or alert tag.
Run after binding	▪ When checked, activates Event field mapping after CI binding. ▪ Use the relevant prefix to retrieve values from the CI record: ▪ **alert_cmdb_ci.** to retrieve a field from the CI record. For example, use 'alert_cmdb_ci.assignment_group' to populate (and transform) an alert 'assignment_group' field with the value of the CI's assignment group. Dot-walk to a related record is also possible. For example, use alert_cmdb_ci.asset.model ▪ **alert_cmdb_ci_key.** to retrieve a value from the CMDB Tags table. For example, alert_cmdb_ci_key.subscription notes. **Warning**: Using the 'Run after binding' feature may affect event processing performance.

Typical use case: To copy the IP field from the event additional info field to the alert field node, use the following:

- 'Source field': ip
- 'Target field': node

- If you selected **Map field and transform value using regex**, fill in the fields as appropriate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Map field and transform value using regex

Field	Description
Source field	Source field value to be copied to the **Target field** field. The **Source field** may be an event field, additional info field, or alert tag.
Target field	Event field where the mapping rule inserts or updates the value. The **Target field** may be an event field, additional info field, or alert tag.
Run after binding	▪ When checked, activates Event field mapping after CI binding. ▪ Use the relevant prefix to retrieve values from the CI record: ▪ **alert_cmdb_ci.** to retrieve a field from the CI record. For example, use 'alert_cmdb_ci.assignment_group' to populate (and transform) an alert 'assignment_group' field with the value of the CI's assignment group. Dot-walk to a related record is also possible. For example, use alert_cmdb_ci.asset.model ▪ **alert_cmdb_ci_key.** to retrieve a value from the CMDB Tags table. For example, alert_cmdb_ci_key.subscription notes. **Warning**: Using the 'Run after binding' feature may affect event processing performance.

Typical use case: To map severity from the field color in the Additional info field to the Severity field, use:

- Source field: 'color'
- Target field: 'severity'
- Transform value pairs:
  - From value: '.*red', To value: '1'. This covers both 'dark red', 'light red', and other values containing 'red' and maps them all to '1'.
  - and so forth...

Key fields (Transform Value Pairs section)

Field	Description
Key	Value that the mapping rule searches for. Whenever the event field has this value, the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	mapping rule adds the value listed in the **Source field** field to the field listed in the **Target field**. This field is displayed when the **Mapping type** is **Map field and transform value (Single field)**. Select **+** to add more **Key** fields, as required.

- If you selected **Map field using regex**, fill in the fields as appropriate.

Map field using regex fields

Field	Description
Source field	Generically defined Source field value to be copied to the **Target field** field. The **Source field** must be a field from the event additional info field. With this type, you can generically find the name of the Source field by using a regular expression (regex).
Target field	Event field where the mapping rule inserts or updates the value. The **Target field** may be an event field, additional info field, or alert tag.

Typical use case: To map the 'region' field from the additional_info field that appears as 'tags.region' or 'tags|region' or 'cloud.region' to an alert tag *t_region*, use:

- 'Source field': *region
- 'Target field': t_region

- If you selected **Map field and copy value from regex group**, fill in the fields, as appropriate.

Map field and copy value from regex group fields

Field	Description
Source field	Generically defined Source field value to be copied to the **Target field** field. The Source field must be a field from the event additional info field. With this type, you can generically find the name of the Source field by using a regular expression (regex).
Target field	Event field where the mapping rule inserts or updates the value. The **Target field** may be an event field, additional info field, or alert tag.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Run after binding	▪ When checked, activates Event field mapping after CI binding. ▪ Use the relevant prefix to retrieve values from the CI record: ▪ **alert_cmdb_ci.** to retrieve a field from the CI record. For example, use 'alert_cmdb_ci.assignment_group' to populate (and transform) an alert 'assignment_group' field with the value of the CI's assignment group. Dot-walk to a related record is also possible. For example, use alert_cmdb_ci.asset.model ▪ **alert_cmdb_ci_key.** to retrieve a value from the CMDB Tags table. For example, alert_cmdb_ci_key.subscription notes. **Warning**: Using the 'Run after binding' feature may affect event processing performance.
Regex group expression	.*_.*_(.*)_.*

Typical use case: To extract the third octet from the IP and save it in the additional info *third_octet* field, use:

- 'Source field': ip
- 'Target field': third_octet
- 'Regex group expression': .*_.*_(.*)_.*

○ If you selected **Advanced mapping using script**, fill in the fields as appropriate.

Advanced mapping using script fields

Field	Description
Script	The code editor gives text editor support for inline scripts. The code editor has these features for the supported language services and Inline scripts: Syntax coloring, indentation, line numbers, and automatic creation of closing braces and quotes, Auto-suggestions, and auto-completions. Editing tips: ▪ To insert a fixed space anywhere in your code, press the Tab key. ▪ To indent a single line of code, select in the leading white space of the line and then press Tab. ▪ To indent one or more lines of code, select the code and then press Tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	▪ To decrease the indentation, press Shift +Tab. ▪ To remove one tab from the start of a line of code, select in the line and press Shift +Tab. ▪ To declare variables, use the var keyword so that they remain within the proper JavaScript scope.
Run after binding	When the 'Run after binding' check box is selected, the rule is executed after the CI binding phase. The dedicated script also has the sys_id of the CI in the input parameters. **Warning:** Complex scripts may affect event processing performance.

- If you selected **Map field from reference table**, fill in the fields as appropriate.

Map field from reference table fields

Field	Description
Reference table	The reference table. Select the table from where the field is to be mapped to the Target field in the event.
Field to match	The field from the Reference table to be matched with the Matching alert field from the event. It's used to filter out the Reference table records.
Source field	**Source field** value to be copied to the **Target field** field. The **Source field** is the field of the table that was selected as the Reference table.
Matching alert field	Event field to find matching Reference table records. The **Matching alert field** is matched with the **Field to match** field. The Matching alert field may be the Event field, Additional info field, or the Alert tag.
Target field	Event field where the mapping rule inserts or updates the value. The **Target field** may be an event field, additional info field, or alert tag.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Run after binding	▪ When checked, activates Event field mapping after CI binding. ▪ Use the relevant prefix to retrieve values from the CI record: ▪ **alert_cmdb_ci.** In the **Matching alert field**, to retrieve a value of a field from a CI record. For example: 'alert_cmdb_ci.ip_address' retrieves IP address information of bound CIs. When the **Matching alert field** and the **Field to match** fields are reference fields, dot-walking to sys_id in the **Matching alert field** is required. For example, to use the 'assignment_group' value of bound CIs to match with the **Field to match** field, use the following expression in the **Matching alert field**: alert_cmdb_ci.assignment_group.sys_id ▪ **alert_cmdb_ci_key.** to retrieve a value from the CMDB Tags table. For example, alert_cmdb_ci_key.subscription notes. **Warning**: Using the 'Run after binding' feature may affect event processing performance.

Typical use case: Extract 'short_description' from the 'pc_vendor_cat_item' records matching its 'name' field with the 'cat_item_name' field in the Additional info field of the event and map it to the 'description' field on the event in the following fields:

- 'Reference table': pc_vendor_cat_item
- 'Field to match': name
- 'Source field': short_description
- 'Matching alert field': cat_item_name
- 'Target field': description

3. Right-click the form header and select **Save**.
4. Select **Submit**.

Example:
The following values comprise a predefined rule that is applied to events in the *Trap From Enterprise 9* class. If the events contain the `snmpTrapOID` element with a value of `iso.org.dod.internet.private.enterprises.cisco.0.0`, the mapping rule changes the value to `reload` in the alerts. If the events contain the `snmpTrapOID` element, a value of `iso.org.dod.internet.private.enterprises.cisco.0.1`, the mapping rule changes the value to `tcpConnectionClose` in the alerts.

Field	Values
Name	cisco.snmpTrapOID

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Values
Source	Trap From Enterprise 9
Mapping type	Map field and transform value (Single field).
Source field	snmpTrapOID
Target field	snmpTrapOID

Transform Value Pairs	• Pair 1
	  ○ **Key**: iso.org.dod.internet.private.enterprises.cisco.0
	  ○ **Value**: reload
	• Pair 2
	  ○ **Key**: iso.org.dod.internet.private.enterprises.cisco.0
	  ○ **Value**: tcpConnectionClose

What to do next

Test an event field mapping by sending an event that includes a field that is present in the event field mapping.

Event identifiers

Event identifiers uniquely distinguish one event from another. Event Management uses these identifiers to determine whether to create a new alert or update an existing one.

By default, an event is uniquely identified by the following fields:

- **Message Key**
- **Domain**
- **Category**

Changing the **Category** field (for example, through a Background Script or Business Rule) may lead to unexpected system behavior.

If the **Message Key** is not provided, the system automatically generates one by concatenating the following fields:

- **Source**
- **Type**
- **Node**
- **Resource**
- **Metric Name**

If the above fields are not defined in the incoming event, you can define them using event rules.

Alert tags

Alert tags allow consolidation for all normalized fields and improve the admin experience to transform and normalize alert fields (key/value) enabling reuse of normalized fields across

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

different sources. This improves alert quality for correlation and provides more out-of-the-box TBAC (Tag Based Automatic Correlation) definitions.

Example: Field normalization in alert management

In some cases, the same information may exist under different field names across data sources. For instance, an incoming event from one source has the location stored under the field **region**, but in other source, the same information is captured under **impacted_region**. To avoid confusion and enable consistent grouping, filtering, or automation, we normalize these fields by creating a unified field — for example,**t_region**. This field pulls the value from **region** or from **impacted_region** depending what field is present in the event. This way, the normalized field **t_region** always holds the location value in a consistent manner, regardless of the original source field.

Alert tags field appears in the Alerts form. These tags are created by Event Rules and in Event mapping and are saved in the alert tags' table [em_alert_tags]. The naming convention used to create key/value pairs is t_<tag name>. This enables reusing tags in the event rules by allowing users to select tags that were previously defined. When new alert tags are defined, TBAC (Tag Based Alert Clustering) tags are automatically created. Using these TBAC tags, you can create new TBAC alert clustering definitions from the new source of tags.

Event field format for event collection

Event Management requires all events to use a standard form, regardless of how they arrive at the instance.

In the application navigation filter, enter `em_event.list`.

Event form

Name	Description
em_event.source [Source]	Event monitoring software that generated the event, such as SolarWinds or SCOM. Maximum length: 100 characters.
em_event.node [Node]	Node name, fully qualified domain name (FQDN), IP address, or MAC address that is associated with the event, such as IBM-ASSET. Maximum length: 100 characters.
em_event.type [Type]	Optional. The metric type to which the event is related, such as Disk or CPU, which is used to identify an event record from which alerts are created. Maximum length: 100 characters.
em_event.resource [Resource]	Node resource that is relevant to the event. For example, Disk C, CPU-1, the name of a process, or service. Maximum length: 100 characters.
metric_name [Metric Name]	The name of the metric that has been measured, such as `DB Disk Free Space (MB)`, `Disk Writes/sec`, or `Disk Write Bytes/sec`.
em_event.event_class [Source instance]	If the em_event.node field is not specified, it is mandatory for alerts to be created automatically. Values for the em_event.event_class field originate from either the source generating the events or by event rule. Name of the machine or software that generated the event. For example, SolarWinds on 10.22.33.44. Corresponding field display name is **Source Instance**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event form (continued)

Name	Description
em_event.message_key [Message key]	Unique event identifier to identify multiple events that relate to the same alert. If this value is empty, it is generated from the **Source**, **Node**, **Type**, **Resource**, and **Metric Name** field values. Maximum length: 1024 characters.
em_event.ci_type [CI type]	String containing information about CI type. Maximum length: 100 characters.
em_event.severity [Severity]	Event severity options are: • **Critical**: Immediate action is required. The resource is either not functional or critical problems are imminent. • **Major**: Major functionality is severely impaired or performance has degraded. • **Minor**: Partial, non-critical loss of functionality or performance degradation occurred. • **Warning**: Attention is required, even though the resource is still functional. • **OK**: An alert is created. The resource is still functional. • **Clear**: No action is required. An alert is not created from this event. Existing alerts are closed.
em_event.resolution_state [Resolution state]	Optional. If the field is empty, the resolution on corresponding alerts is still pending. Event state from the event source is either **New** or **Closing**. • **New**, the resolution on corresponding alerts is open. • **Closing** event state closes corresponding alerts.
em_event.time_of_event [Time of event]	Time that the event occurred in the source system. This field is a GlideDateTime field in UTC or GMT format. Maximum length: 40 characters.
em_event.state [State]	Current processing state of the event: • **Ready**: Event has been received and is waiting to be processed. • **Processed**: Event was successfully processed. • **Ignored**: Value is not in use. • **Error**: Failure occurred while processing the event. For example, the event collection method or event **Severity** is blank.
em_event.alert [Alert]	If an alert was created as a result of the event, this field contains the unique ID that Event Management generates to identify the alert.
em_event.description [Description]	Reason for event generation. Shows extra details about an issue. For example, a server stack trace or details from a monitoring tool. Maximum length: 4000 characters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event form (continued)

Name	Description
em_event.additional_info [Additional information]	Optional. A JSON string that gives more information about the event. The JSON data is supported for String values only, other value types are not supported. You must convert numbers to String values by enclosing them in double quotes. For example, this value is not supported: {"CPU":100 } while this value is supported: {"CPU":"100"}. Another example of a valid JSON string is: {"evtComponent":"Microsoft-Windows-WindowsUpdateClient","evtMessage":"Installation Failure: Windows failed. Error 0x80070490"}. This information can be used for third-party integration or other post-alert processing. Values in the **Additional information** field of an Event that are not in JSON key/value format are normalized to JSON key/value format when the event is processed. For example, assume that the following plain text is in the **Additional information** field "Connection instance is successful". When the event is processed, all this plain text becomes one JSON string and might not be useful within an alert. In the resultant alert, this string is in the **Additional information** field in JSON key/value format, containing the data: {"additional_content": "Connection instance is successful"}.
processing_notes [Processing Notes]	Display of the events processing log.

Custom alert fields

You can populate custom alert fields with data contained in **Additional information** field of the event.

Valuable data contained in **Additional information** fields of events can be useful, for example, for reporting. Alert fields are automatically populated from fields that have the same name in the event. This behavior holds true for **Additional information** event fields and for **Additional information** fields that Event Rule adds. Therefore, to populate a custom alert field with the value in **Additional information** fields, use the custom field name in the **Additional information** field. You can also use Event Rules for this purpose. Values in the **Additional information** field of an Event that are not in JSON key/value format are normalized to JSON format when the event is processed. For more information, see Configure an event rule to customize alert content.

The short description field of an alert is automatically written using the type, node, and description of the alert. To prevent the short description field from being over-written, open the *sys_properties* table, locate the *evt_mgmt.override_alert_short_description* property, change the value to `false`, and then click **Update**.

Depending on permissions, you may only be able to create fields with the **user_** prefix. If so, use Event Rules to create an **Additional information** field with the same name. To prevent some fields to be copied to the alert field, use the *evt_mgmt.alert_black_list_fields* property and add the field names that must be excluded. By default, the fields that are not copied are:

- message_key
- category
- additional_info

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- sys_updated_on
- sys_updated_by
- sys_created_by
- sys_created_on
- sys_mod_count
- sys_id

Testing and sending events

You can manually test and send events to confirm that Event Management properly manages events and generates alerts.

Before you begin

Role required: evt_mgmt_admin

About this task

For example, you can manually send events to:

- Confirm that the MID Server is using an event connector definition and instance to send events.
- Confirm that event rules, event field mapping, and other configurations do generate alerts.
- Track an operation or action that did not generate an event.

Procedure

1. Navigate to **All** > **Event Management** > **All Events**.
2. Click **New** and then fill in the fields.

Event form

Name	Description
Source	Event monitoring software that generated the event, such as SolarWinds or SCOM. Maximum length: 100 characters.
Node	Node name, fully qualified domain name (FQDN), IP address, or MAC address that is associated with the event, such as IBM-ASSET. Maximum length: 100 characters.
Type	The metric type to which the event is related, such as Disk or CPU, which is used to identify an event record from which alerts are created. Maximum length: 100 characters.
Resource	Unique event identifier to identify multiple events that relate to the same alert. If this value is empty, it is generated from the **Source**, **Node**, **Type**, **Resource**, and **Metric Name** field values. Maximum length: 1024 characters.
Metric Name	The name of the metric that has been measured, such as `DB Disk Free Space (MB)`, `Disk Writes/sec`, or `Disk Write Bytes/sec`.
Source instance	Name of the machine or software that generated the event. For example, SolarWinds on 10.22.33.44. Corresponding field display name is **Source Instance**.
Message key	Unique event identifier to identify multiple events that relate to the same alert. If this value is empty, it is generated from the **Source**, **Node**, **Type**, **Resource**, and **Metric Name** field values. Maximum length: 1024 characters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
Severity	The severity of the event. The options are typically interpreted as follows: ○ **Critical**: Immediate action is required. The resource is either not functional or critical problems are imminent. ○ **Major**: Major functionality is severely impaired or performance has degraded. ○ **Minor**: Partial, non-critical loss of functionality or performance degradation occurred. ○ **Warning**: Attention is required, even though the resource is still functional. ○ **Clear**: No action is required. An alert is not created from this event. Existing alerts are closed. ○ **OK**: An alert is created. The resource is still functional.
Resolution state	Optional. If the field is empty, the resolution on corresponding alerts is still pending. Event state from the event source is either **New** or **Closing**. ○ **New**, the resolution on corresponding alerts is open. ○ **Closing** event state closes corresponding alerts.
Time of event	Time that the event occurred in the source system. This field is a GlideDateTime field in UTC or GMT format. Maximum length: 40 characters.
State	The current processing state of the event: ○ **Ready**: Event has been received and is waiting to be processed. ○ **Processed**: Event was successfully processed. ○ **Ignored**: Value is not in use. ○ **Error**: Failure occurred while processing the event. For example, the event collection method or event **Severity** is blank.
Alert	If an alert was created as a result of the event, this field contains the unique ID that Event Management generates to identify the alert.
Description	Reason for event generation. Shows extra details about an issue. For example, a server stack trace or details from a monitoring tool. Maximum length: 4000 characters.
Additional information	A JSON string that gives more information about the event. The JSON data is supported for String values only, other value types are not supported. You must convert numbers to String values by enclosing them in double quotes. For example, this value is not supported: {"CPU":100 } while this value is supported: {"CPU":"100"}. Another example of a valid JSON string is: {"evtComponent":"Microsoft-Windows-WindowsUpdateClient","evtMessage":"Installation Failure: Windows failed. Error 0x80070490"}. This information can be used for third-party integration or other post-alert processing. Values in the **Additional information** field of an Event that are not in JSON key/value format are normalized to JSON key/value format when the event is processed. For example, assume that the following plain text is in the **Additional information** field “Connection instance is successful”. When the event is processed, all this plain text becomes one JSON string and might not be useful within an alert. In the resultant alert, this string is in the **Additional information** field in JSON key/value format, containing the data: {“additional_content”: “Connection instance is successful"}.

3. Click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View event processing statistics

Extract statistics from your instance to ensure that performance is not affected and extract metrics related to event processes to monitor event processing status.

Before you begin

Ensure that the *Event Management - Collect xmlstats* job is active. Configure this scheduled job by navigating to **System Definition** > **Scheduled Jobs**.

You can control the processing of event statistics using the *evt_mgmt.event_processing_stats_period* property. For information about configuring this property, see Configure statistics processing period.

Role required: evt_mgmt_operator

About this task

You can view key metrics from the instance to ensure that performance is not affected and monitor the status of the processing of events. For more information about viewing self-health monitoring statistics, see View Metric Intelligence self-health statistics.

Procedure

1. In a browser where the instance is running, run the XML Stats script by appending the following text to the instance URL: `xmlstats.do?include=XMLStatsEm`
2. Press **Enter**.
3. Monitor the event processing metrics that are reported to ensure that events are processed according to accepted benchmark values.

Result

Example: Sample output after running the XML Stats script

```
<xmlstats created="Tue Oct 23 04:01:25 PDT 2018"
 includes="XMLStatsEm" version="2">
<itom_event_management>
<event_processing>
<processed count="47" max_process_time="10"
 mean_process_time="5" record_time="2018-10-23 11:00:41"/>
<ready count="7" oldest_ready_event="37" record_time="2018-10-23
 11:00:41"/>
<error count="5" record_time="2018-10-23 11:00:41"/>
<total_created count="5" record_time="2018-10-23 11:00:41"/>
</event_processing>
</itom_event_management>
</xmlstats>
```

The event processing metrics in the report are described in the following table.

Note: The default processing cycle is events created in the last hour and updated in the last minute.

Event processing statistics

Metric	Description
Event processing statistics	
processed count=n	Number of events processed by the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event processing statistics (continued)

Metric	Description
max_process_time=n	Longest time taken by the instance to process an event.
mean_process_time=n	Average time taken by the instance to process the events in this cycle.
Event ready statistics	
ready count=n	Number of events that are in Ready state and are available to be processed. Default processing cycle: events created in the last hour and updated in the last minute.
oldest_ready_event=n	Longest time taken by an event to remain in Ready state. If this metric has a value that is bigger than your benchmark figure, investigate the event further to determine what is impeding the processing. Default processing cycle: events created in the last hour.
Other processing statistics	
error count=n	Number of events for which errors were reported. Default processing cycle: events created in the last hour and updated in the last minute.
total_created count=n	Overall number of events that were processed during this event processing cycle. Default processing cycle: events created in the last hour.

Enable processing of event process statistics

Enable the system property that switches on statistics processing for events to let the platform collect and analyze metrics such as event volume, processing flow, and handling time. This helps you monitor system performance, identify bottlenecks, and optimize event handling for faster and more efficient operations.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Select **New** and fill in the form.

System property fields

Field	Description
Name	Name of the property you are creating.
Description	Narrative that describes the property and its function.
Choices	Comma-separated values for a choice list. If you need a different choice list label and value, use an equal sign (=) to separate the label from the value. For example, Blue=0000FF, Red=FF0000,Green=00FF00 displays Blue, Red, and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Green in the list, and saves the corresponding hex value in the property value field.
Type	Data type, that is selected from a list (for example, integer, string, or true|false).
Value	Desired value for the property. When retrieving properties using the *gs.getProperty()* method, treat the results as strings. For example, a true|false property returns 'true' or 'false' (strings), not the Boolean equivalent.
Read roles	Roles that have read access to this property.
Write roles	Roles that have write access to this property.

3. Select **Submit**.

Related topics

Add a system property

Configure statistics processing period

Set the time period, in seconds, for collecting event processing statistics. For example, you can set a time period twice as long as the default 60 seconds to collect more statistics.

Before you begin

Role required: evt_mgmt_admin

About this task

Use the *evt_mgmt.event_processing_stats_period* property to configure the XMLStatsEm script time to collect event processing statistics. The default value of this property is 60 and the unit is second. Set the property values to run the script over a longer or shorter period. For example, set a shorter period of statistics collection so fewer CPU resources are used.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for the *evt_mgmt.event_processing_stats_period* property.
   Either configure the value of an existing *evt_mgmt.event_processing_stats_period* property to run the XMLStatsEm script over a shorter or longer time period or create a new property if it does not exist already.
   - If the property exists, set the property to the required time period to enable the XMLStatsEm script time to collect event processing statistics and click **Update**.
   - If the property does not exist, click **New**, fill in the form, and click **Submit**.

Related topics

Add a system property

Manage and monitor alerts

An alert is a notification for selected events that are considered to be important and require attention. Event Management generates alerts based on event rules.

You can monitor and resolve alerts in ways that are based on alert impact calculations and alert configuration and property settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert management rules for resolving alerts

You can configure Event Management to respond to alerts automatically. An alert management rule determines the required alert response, such as to open an incident, knowledge base article, open a task, launch remediation action.

Alert management rules provided with the base system as a store application (Alert Rules Management [sn_em_arm]) to help you respond to alerts. You can create filters to specify conditions for the rule so that the remedial action specified in the rule takes effect only when the conditions are met. For example, launch the required subflow or open an incident based on an alert. The alert's execution history is automatically updated to indicate the actions that were invoked.

Users with the evt_mgmt_admin role can use the alert management rule designer to create and customize alert management rules to act on specified alerts. Define rules with filters to determine which alerts the rule applies to. You can create rules to launch applications, URLs, subflows, remediation actions, or take other actions, such as to open an incident. For more information, see Create an alert management rule.

Users with the evt_mgmt_operator role can manually run alert management rules.

Alert management rule flow

The flow to create and run an alert management rule is:

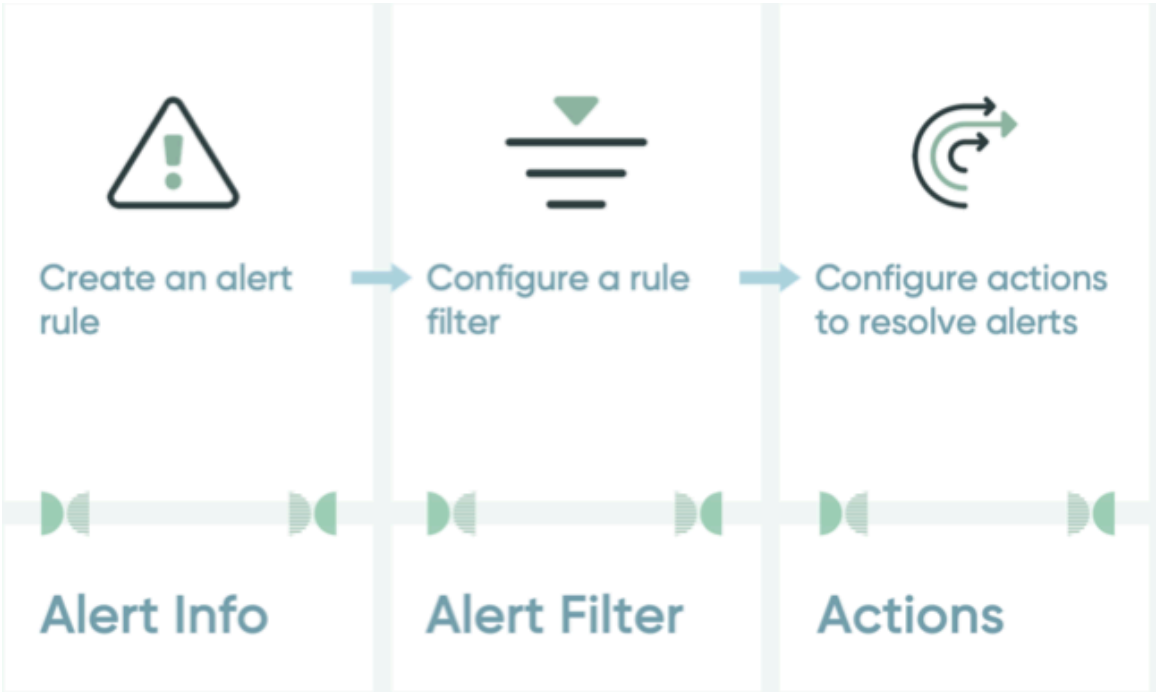

Alert management rule components

Component	Description
Alert Info	Configure a name and general information for the rule.
Alert Filter	Specify a filter to determine to which alerts the rule applies. You can specify the related list conditions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert management rule components (continued)

Component	Description
	Note: The fields that are not supported for alert filtering are: Overall Event Count, Priority, Priority Group, Priority Breakdown, Tags, and Impacted Services.
Actions	Specify the response to the alert, such as to run a subflow, perform remediation action, launch an application, or launch a URL in a browser.

How rules are applied to updated alerts

Alert management rules run on all updated open alerts. Rules don't run on closed alerts, even if they've been updated. The filters determine whether the rule's actions apply to the alert. For example, if a rule's condition indicates that an email message is sent when the alert severity changes to Major, the rule applies to an alert updated by a severity change from Warning to Major.

Use of filters and other actions

Filters ensure that the rule is invoked only when the configured condition occurs, and not for every update of the alert. For example, you can configure a rule so that updates that aren't relevant (such as a **Work notes** field update) don't cause the rule to run. As another example, a filter condition can specify that the alert management rule runs only when the alert severity is critical.

You can perform the following actions:

- Specify a filter that determines which alerts the rule applies to.
- In the Related List Conditions section of the form, configure additional conditions, for example, with an *Alert > Parent* relationship, to filter for any alerts that were received today.
- Respond to alerts. For example, by using subflows and workflows, create incidents for primary alerts with critical severity, or open a search engine in a browser to search for data according to the description field of the alert.
- Apply remediation. Remediation is based on Orchestration workflows that can be scripted to perform remediation tasks such as gathering system information or rebooting a server.

Note: For enhanced performance of *Event Management - Evaluate Scoped Alert Rules Management* scheduled jobs, use subflows instead of workflows.

Scheduled jobs that check alert management rules

Alert management rules are checked every 11 seconds by the default *Event Management - Evaluate Scoped Alert Rules Management0* scheduled job. The job then executes the required actions. For large-scale environments, you can add more than one job. Please contact Customer Service and Support.

Note: Only new users from Vancouver and up get two scheduled jobs: *Event Management - Evaluate Scoped Alert Rules Management0* and *Event Management - Evaluate Scoped Alert Rules Management1*. Users upgrading from previous family releases remain with a single scheduled job *Event Management - Evaluate Scoped Alert Rules Management0*.

Don't modify the *sn_em_arm.alert_management.num_of_jobs* property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, the alert grouping job (**Service Analytics group alerts using RCA/Alert Aggregation**) and the alert management (**Event Management - Evaluate Scoped Alert Rules Management0**) jobs run independently of each other. For more information about coordinating the alert response and the automated alert grouping, see Synchronizing alert response with automated alert grouping.

Migrate existing alert action rules

Existing alert action rules from an earlier release can be migrated to become alert management rules. You can modify an alert action rule only after migrating it to an alert management rule. For more information, see Migrate an alert action rule to an alert management rule.

Create an alert management rule

Create an alert management rule to track alerts and resolve them by determining the required response, for example, to open an incident or launch remediation action.

Before you begin

To enable remediation with a subflow, you can use a subflow that is available with the base system, or you can create your own subflow. For details, see Create a custom subflow for alerts.

Role required: evt_mgmt_admin, flow_designer

About this task

Note: Use Respond automation in Service Operations Workspace, the updated way to automate alert responses with an easier interface.

Use alert management rules to track and resolve alerts.

While working in the alert management rule designer, you can work in multiple sections without losing information in any section.

Note:

- Alert management rules that are not configured to perform any action are skipped and the rule is automatically set to inactive.
- Mark the alert with the **In Maintenance** status when:
  - an alert is bound to a CI by a user action, such as an alert management rule
  - and the CI is in the Maintenance state

Create alert management rules that:

- Locate other alert management rules that have relevance to the selected alert.
- Determine when the execution of the rule takes place.

Alert management rules do not necessarily complete in the order in which they are invoked.

You can configure alert management rules to:

- Automatically generate and link incidents, tasks, or knowledge articles to alerts.
- Automatically apply a remediation workflow or enable users to manually run remediation.
- Automatically construct a URL according to the value of specified fields in the alert.

To assist you, several alert management rules are provided with the base system. You can use them as presented or you can use them as examples to build custom alert management rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert management rules provided with the base system

Rule	Description	Active
Open sensor dashboard in PRTG	The sensor dashboard in the Paessler PRTG Network Monitor (PRTG) application opens.	Yes
Oracle EM Launch Target Status and View Events	Launch Oracle Enterprise Manager to view: • Target Status • Event for alerts from source Oracle EM	Yes
Drilldown to OMI	Drill down to the HP Operations Manager i (OMi) application.	Yes
Create Incident on Primary Critical Alert	Create an incident for primary critical alerts. The incident can be created automatically or manually.	No
Search Google for "description"	Open Google Search in a browser to search for data according to the description that appears in the alert.	Yes
Create Incident	Create an incident for all alerts that are not in maintenance state. The rule runs automatically on selective update.	No
Create Incident Manually	Manually create an incident for alerts that are not in maintenance state.	Yes
Create Major Incident Candidate	Create a major incident candidate for all alerts that are not in maintenance state and are not secondary alerts. A major incident candidate can be promoted to become a major incident.	No
Create Major Incident	Create a major incident for all alerts that are not in maintenance state and are not secondary alerts.	No

If your instance was upgraded from Kingston, the alert action rules that were provided with the Kingston base system are available to you. However, if you modified any of the rules, the changes made are not carried over.

Alert management rules run 5 seconds after an alert is updated, resetting the timer if updates occur within that window. This delay ensures remediation actions, such as incident creation, are triggered only when the issue is clear and stable, reducing duplicates and unnecessary noise. To change the default 5-second delay, create the `evt_mgmt.alert_rule_delay` property on the **All** > **System Properties** > **All Properties** and change the value. To know how to create a property, see Add a system property.

Procedure

1. Navigate to **Event Management** > **Rules** > **Alert Management Rules**.
2. Click **New** and then fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert Management Rule form

Field	Description
Name	Unique name for the rule.
Active	Check box for enabling the rule. If this check box is selected, you must specify: ○ in the **Alert Filter** section, an alert filter ○ in the **Actions** section, at least one of any of these actions: ▪ active subflow ▪ workflow ▪ quick response
Order	Order in which rules are evaluated when multiple rules are defined for the same alert. Alert management rules are evaluated in ascending order. The default value is `100`.
Multiple alert rules	Instruction about whether to search for additional rules: ○ **Search for additional rules**: Execute the current rule then continue and execute other matching rules in the order of rule priority, where the lower number has the higher priority. ○ **Stop search for additional rules**: Execute only the current rule for the alert that matches the filter. When selecting this option, a rule configured second in a rule hierarchy may run before the results of the first rule take effect.
Description	Descriptive text for the rule.
Assignment group	Assignment group that works on the alert. If no assignment group is defined in the alert rule, then this alert rule is considered as a global rule. When the rules are running – first the global rules run and then the rules that belong to the assignment group of the alert.

3. Click **Alert Filter** and specify conditions for alerts that this rule is applied to.

Alert Filter stage

Field	Description
Rule is activated when	Rule execution takes place when: ○ **Alert changes to filter**–content changes to the alert cause the alert to match the filter. If the filter is matched on following update of the alert, the rule is not applied. If the alert was closed and then reopens, at the next update of the alert and the filter is matched, the rule is applied. Thereafter, when there is an update of the alert, the rule is no longer applied.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ **Alert matches filter**–the content of the alert matches the filter. On following update of the alert and if the filter is matched, the rule runs and is applied to the alert. The rule remains applied for every matching update.
Alert filter	
Preview	Function to preview alerts that match the specified condition. A hyperlink shows how many alerts match the filter. 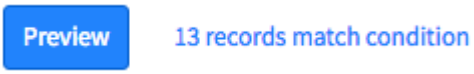 If you click the hyperlink, the browser opens another tab that lists alerts in the Alerts [em_alert_list] table. The list shows which alerts match the rule, including closed alerts. Alerts that have already been run by the rule are not marked in any way. You can click any alert to view further details.
Conditions	Conditions that, if fulfilled, cause the filter to be applied. For more information about building conditions, see Using the condition builder. To add another condition, click **New Criteria**. **Note:** The **Created on** condition is not invoked when running the **Event Management - Evaluate Alert Management Rules** job. Instead, use the **Updated on** condition, as the job detects alerts based on the time they were updated and not on the time they were created.
Related List Conditions	Conditions to include a relationship with another table in the filter. **a.** Click **Table** and select the required table. **b.** Specify the conditions for this filter. For more information about creating related lists, see Add related list conditions.

4. Click **Actions**.
   In this section, you can configure the following action types as a response to alerts or to remediate alerts:
   - Remediation Subflows: Execute a subflow provided with the base system.
   - Launch Applications: Open applications and browsers that you configure.

   Note: The Remediation Workflows option is deprecated. To enable flows to be triggered by alerts, use the Flow Designer.

5. **Optional:** In the Remediation Subflows section, follow these instructions to add subflows:

   a. Under **Subflow**, double-click the cell.

   b. Click the search icon.
   The list of subflows provided with the base system appears. For more information, see Event Management subflows in the base system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. From the subflow list, select a subflow.

d. Repeat, adding as many subflows as required.

(Optional)

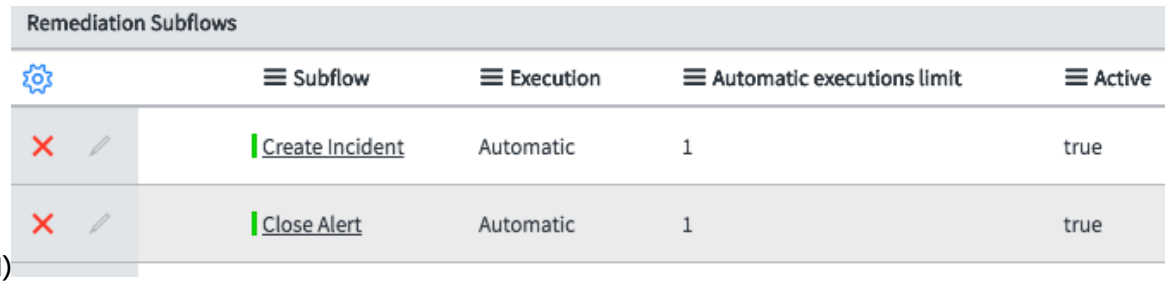

e. To specify when the subflow must be executed, double-click the cell under **Execution**.

Subflow execution options

Name	Description
Automatic	The subflow is executed automatically when the rule is matched.
Manual	Execute the subflow if required when the rule is matched.
Both	When the rule is matched, the subflow is executed automatically and you can optionally execute the subflow again manually.

f. Under **Automatic executions limit**, double-click the cell and enter the integer number of times to execute the subflow.
After the subflow has been executed the indicated number of times, it does not run anymore.

g. To enable the subflow to be executed, double-click the cell under **Active** and select `true`.
A link in the cell under **Link to Flow Designer** appears only after a subflow has been selected and the rule has been saved.

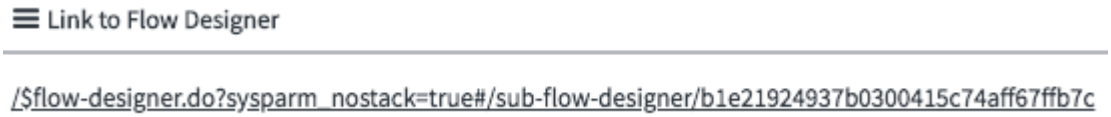

6. To add instructions to launch applications or to open browser windows, in the Launch Applications area:

a. Under **Display Name**, double-click the cell.
Specify a name for the link.

b. In the **URL** field, compose the URL using data from the alert in the format:`http://${source}.com:${port}/${cmdb_ci.name}`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Active** field is automatically updated.

Any URL-based action can utilize the alert parameters and the URLs can refer to wikis, messaging services, REST APIs, and so on.

7. Click **Submit**.

Result

The alert management rule is added to the list of available rules that can be used to resolve alerts.

Related topics

Apply a quick response in an alert

Event Management subflows in the base system

The subflows provided with the base system appear in the Remediation Subflows area of alert management rules.

Accessing the subflows

Navigate to **Event Management** > **Rules** > **Alert Management Rules** and click **New**. Click the **Actions** tab. In the Remediation Subflows area, double-click the **Insert a new row** field.

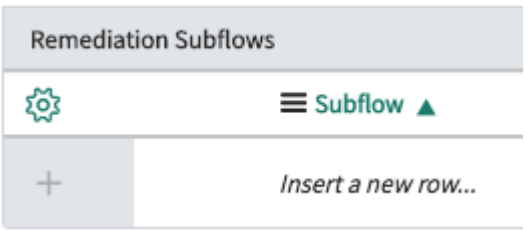

Click the search icon to add subflows. The list of subflows that are provided with the base system appears.

Subflows in the base system

Name	Description
Acknowledge Alert	Subflow to mark the alert as being Acknowledged. Acknowledge an alert to show that further attention is required.
Attach Knowledge Article (legacy)	Subflow to attach a knowledge article to the alert. This subflow is provided for instances that are migrated from legacy releases (prior to the London release). **Note:** Add the **Knowledge article** column to the Alert Management Rules [em_alert_management_rule] table, and select an article to attach to an alert when the rule executes.
Change Alert to Maintenance Mode	Subflow to mark the alert as being in Maintenance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subflows in the base system (continued)

Name	Description
Close Alert	Subflow to mark the alert as being Closed.
Create Incident	Subflow to create an incident. Fields from the alert are used to populate the matching fields in the incident that is created. **Note:** • If there is an existing incident that is attached to the alert, this subflow is not activated. • If the alert is in Maintenance, an incident is not created. • The alert management job runs even if the alert grouping job is not complete, if a specified time frame has passed. When this occurs, you can enable the *Avoid INTs on secondary alerts* rule to prevent incidents from being created for secondary alerts. This prevents duplicate incidents because an incident already exists for the primary alert when the *evt_mgmt.avoid_int_enabled* property is enabled.
Create Major Incident Candidate	Subflow to create a major incident candidate. Fields from the alert populate the matching fields in the major incident candidate that is created. A major incident candidate can be upgraded to become a major incident. **Note:** • If there is an existing incident that is attached to the alert, this subflow is not activated. • If the alert is in Maintenance, a major incident candidate is not created. • If the **Role in group** is **Secondary**, the major incident candidate is not created.
Create Major Incident from Alert	Subflow to create a major incident from alert. Fields from the alert are used to populate the matching fields in the major incident that is created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subflows in the base system (continued)

Name	Description
	Note: • If there is an existing incident that is attached to the alert, this subflow is not activated. • If the alert is in Maintenance, an incident is not created. • If the **Role in group** is **Secondary**, the major incident candidate is not created.
Create Major Incident with Impact	Subflow to create a major incident from an alert in which the **Impact** field is also taken as input. Fields from the alert are used to populate the matching fields in the major incident that is created. **Note:** • If there is an existing incident that is attached to the alert, this subflow is not activated. • If the alert is in Maintenance, an incident is not created. • If the **Role in group** is **Secondary**, the major incident candidate is not created.
Create Major Incident Candidate with Impact	Subflow to create a major incident candidate in which the **Impact** field is also taken as input. Fields from the alert populate the matching fields in the major incident candidate that is created. A major incident candidate can be upgraded to become a major incident. **Note:** • If there is an existing incident that is attached to the alert, this subflow is not activated. • If the alert is in Maintenance, a major incident candidate is not created. • If the **Role in group** is **Secondary**, the major incident candidate is not created.
Create Task (legacy)	This subflow uses a task template, if provided, or the *EventMgmtCustomIncidentPopulator*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subflows in the base system (continued)

Name	Description
	script for instances migrated from legacy releases (prior to the London release). If configured, apply the task template. **Note:** Add the **Task template** column to the Alert Management Rules [em_alert_management_rule] table, and select a task template and task to apply when the rule executes.
Overwrite Alert Template (legacy)	This subflow applies the alert template. This subflow is provided for instances that are migrated from legacy releases (prior to the London release). **Note:** Add the **Task type** column to the Alert Management Rules [em_alert_management_rule] table, and select an alert template to apply when the rule executes.

1. Select the subflow that you need.
2. To customize a subflow, see Create a custom subflow for alerts. This topic also describes the input parameters in a subflow.
3. To specify when the workflow must be executed, double-click the cell under **Execution**.

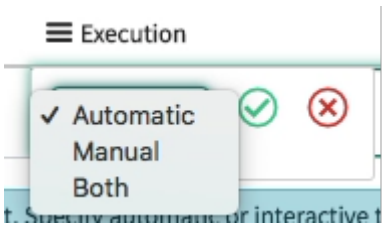

.

Create a custom subflow for alerts

You can create a subflow according to your requirements. For example, you can resolve alerts, notify teams, or run remediation actions.

Before you begin

Role required: evt_mgmt_admin, flow_designer

About this task

Copy an existing subflow, and then customize it. Publish the custom subflow so that it is available to be used to resolve any alert. For more information, see Create a subflow.

Note: If your ServiceNow instance uses domain separation, you must ensure that the new custom subflow that you created operates only on the data of the designated domain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Process Automation** > **Flow Designer**.
2. In the Flow Designer window, click **Subflows**.
3. Click **Create Incident**.
   The Create Incident subflow opens in Flow Designer in read-only format. The inputs in the Create Incident subflow are described in the following table.

Create Incident input parameters

Value	Description
alertRuleName	The name of the alert management rule that specifies this subflow, for example `IDD`.
alertRuleId	The sys_id of the alert management rule, for example, `1ae122e84f5f1300b28392918110c71e`.
alertGR	The alert record that matches the filter that is defined in the alert management rule, for example, `Alert0027931`
executionId	The ID of the execution to be retrieved, for example, `000182ebdb66630091187f1daf9619d4`
userName	The current user's name, for example, `admin`.
userDisplayName	The name that the current user specified to be displayed, for example, `System Administrator`.

For example, in an alert management rule that has processed alerts, in the Alert Executions area, click the **Link to execution** of the required alert.

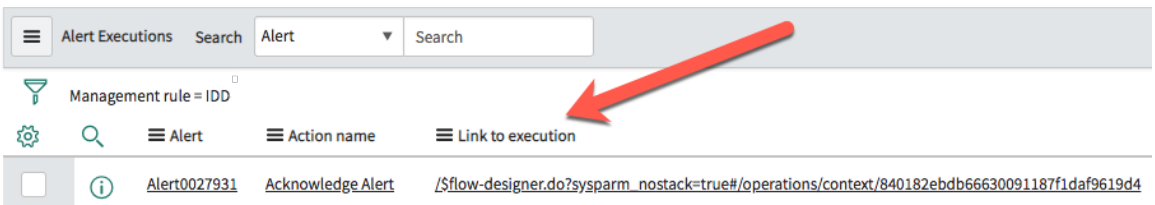

The Flow Designer opens. Click **Subflow Inputs and Outputs**. The input variable names of the subflow are displayed with their respective values.

4. In the subflow header, click more actions and select **Copy subflow**.
5. In the **New Subflow Name** field, enter the name of the new subflow.
6. To limit the application scope of the subflow, in the **Application** field, select the required application, otherwise select `Global`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Click **Copy**.
   The new subflow is created and is opened as a tab in the subflow designer editor. The inputs provided by the template appear.

   Inputs are fixed and cannot be added to.

   In the editor, you can delete subflow inputs or outputs, as required.

8. To delete an input, click .
   Inputs are fixed and cannot be added to.
9. Create subflow outputs by defining the names and data types.
   a. Click to add output.
   b. Define the name and the data type.
10. To save the subflow in draft status, click **Save**.
11. To save and publish the subflow so that it can be used in the alert management rule designer, click **Publish**.

Result

Include the modified subflow in an alert management rule to resolve alerts according to your customized requirements.

Related topics

Subflows

Alert executions information

Alert executions information provides a reference to the alert management rule actions that are performed. This information appears in the Alert Management Rule record only if an alert matches the filter in the rule and an action was performed. You can click any link in the Alert Executions list to open the referenced item.

Alert Executions

In the Alert Management Rule record, the Alert Executions list shows key related records for the rule.

Note: The Alert Executions list in the Alert Management Rules record shows alert execution information and lists the alert management rules that the selected alert matched.

Alert Executions list

Information in the Alert Executions list

The columns in the Alert Executions list are populated with details of the alerts that matched the alert management rule and the actions that ran. The listed alerts match the filter in the rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View alert execution information

You can click any link in the Alert Executions list to view the alert execution information of the referenced item. This information appears in the Alert Management Rule record only after an alert matched the filter in the rule and an action was performed.

Before you begin

Role required: evt_mgmt_admin

About this task

Among the information shown in the Alert Executions list is:

- The list of alerts that matched the filter of the rule.
- The tasks, including incidents, that were opened.
- Which remediation workflows ran and which subflows ran.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Alert Management Rules**.
2. In the Alert Management Rules list, click the rule.
   If an alert matched the filter in the rule and an action was performed, the Alert Executions list appears in the Alert Management Rules record. The list appears irrespective of the stage of the selected alert management rule.

Alert Executions list

Column	Description
Alert	Alert number of the alerts that matched the alert management rule.
Action name	Name of the action that was taken to resolve the alert (subflow or workflow).
Link to execution	Link to the action execution log that was created when the action was performed.
Related task	Link to the task that was created when the alert was resolved. For example, if the action for a rule is to create an incident, this field provides the incident number with a link to open the incident. A link is provided to the workflow remediation task that is created to monitor the remediation.
Log	Notice or information showing the result of the execution action, for example, how executions took place, "`1 out of 2 executions`". When the alert closes, the execution count number (the first number) resets to 0. This information does not indicate that the subflow was executed successfully, rather that the call to that subflow was successful and that the subflow is correctly configured.
Automatic run	Indication of how the action ran: automatically (true) or manually (false).
Created	Date and time that the entry was created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Migrate an alert action rule to an alert management rule

Existing alert action rules from an earlier release can be executed, but cannot be modified. Alert action rules that have been migrated become alert management rules and all the definitions of the alert action rules are migrated to the alert management rule format. Migrated rules can be modified.

Before you begin

Role required: evt_mgmt_admin

About this task

Existing alert action rules are active, read-only rules. An alert action rule can be modified only after it is migrated to an alert management rule.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Alert Action Rules**.
2. Select an alert action rule to migrate or delete.

Alert Action Rule record

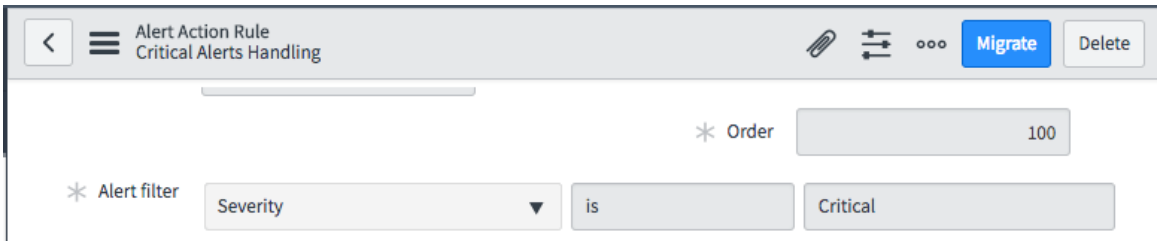

3. Click **Migrate**.
   If you clicked **Migrate**, the list of migrated rules can be found, together with new alert management rules, by navigating to **Event Management** > **Rules** > **Alert Management**.

   After a successful migration, all the definitions of the alert action rules are migrated to the alert management rule format.

Result

When an alert action rule has been migrated successfully, the alert action rule can no longer be modified. However, the new alert management rule can be modified.

What to do next

If the alert action rule that was migrated had a task template, you can customize the alert management rule designer to display the task template name.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. In the header of the alert management rule, click **Additional actions**

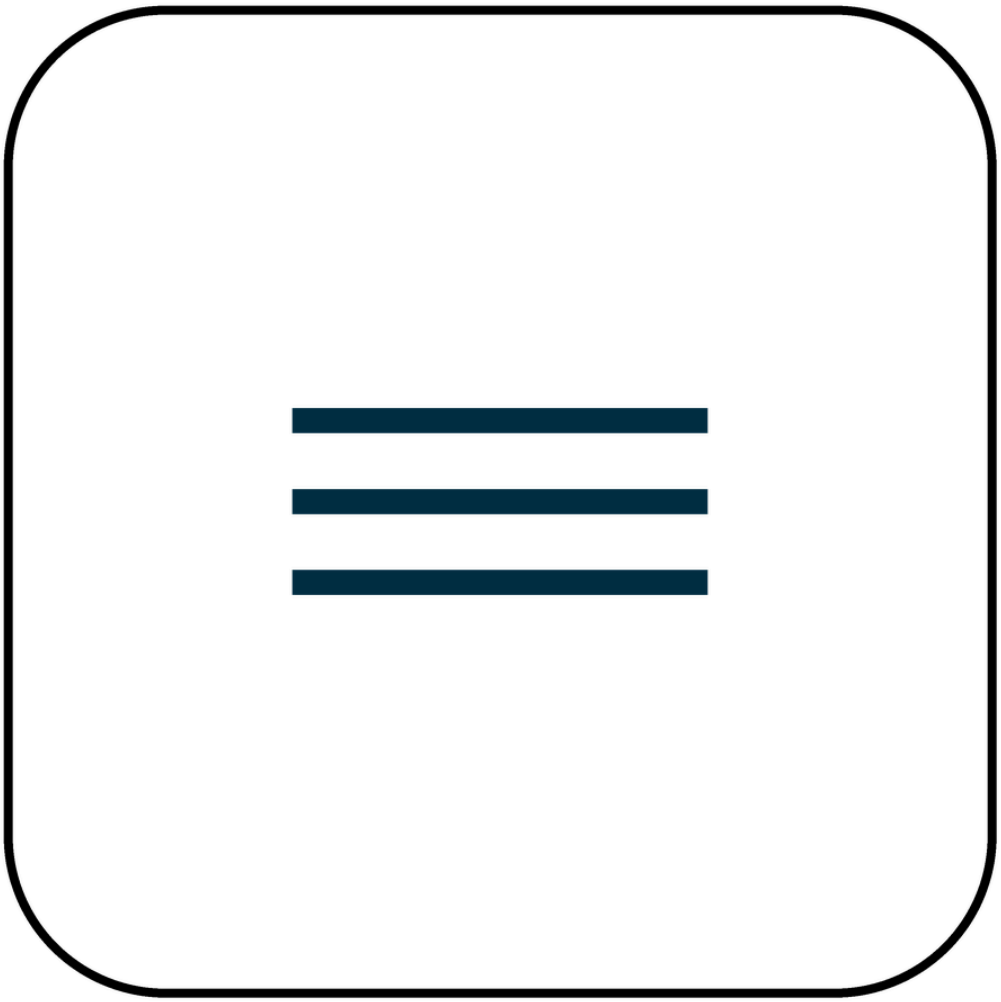

2. Select **Configure** > **Form Layout**.

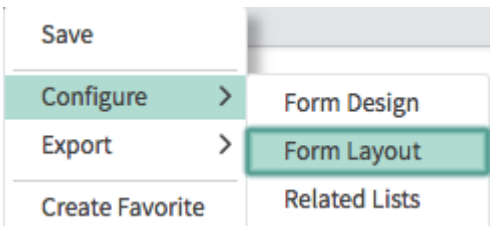

3. In the **Available** pane, select **Task template** and move it to the **Selected** pane. Position it as required.

4. In the **Task template** field, select the required template.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. In the Actions tab, in the Remediation Subflows area, select the Create Task (legacy) subflow to run the task template.

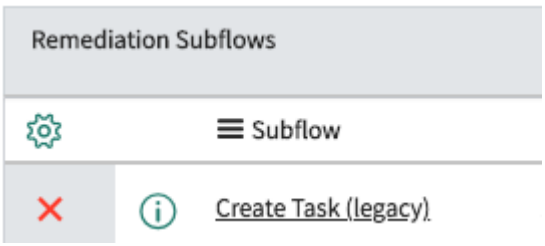

Related topics

Apply a quick response in an alert

Synchronizing alert response with automated alert grouping

Synchronize alert response with grouping by ensuring alert management jobs runs after alert grouping jobs—this prevents duplicate actions like incident creation on secondary alerts.

Before you begin

Role required: admin

About this task

By default, the alert grouping job (**Service Analytics group alerts using RCA/Alert Aggregation**) and the alert management job (**Event Management - Evaluate Scoped Alert Rules Management0**) run independently of each other. To synchronize the alert response with automated alert grouping (for example, to avoid creating an incident on a secondary alert), you must enable the property *evt_mgmt.avoid_int_enabled* to ensure that the Alert Management job is executed after the Alert Grouping job is completed. This property is accessible from the Event Management Properties page (**All** > **Event Management** > **Administration** > **Event Management Properties**). In addition, you must configure the relevant alert management rule (for example, for incident creation), to filter out all secondary alerts.

Note: If only one alert exists when the Alert Management job runs, an incident is created. When a second alert arrives later, the next Alert Grouping job forms a group, makes the new alert as secondary, and creates a new incident.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Event Management Properties**.
2. If you want to avoid incidents on secondary alerts, select the **Enable Avoid Incidents on Secondary alerts and wait for Grouping job to be executed** check box.
3. Select **Save**.
4. To filter out secondary alerts, navigate to **All** > **Event Management** > **Rules** > **Alert Management Rules**.
5. Select the relevant rule.
6. In the Alert Management Rule form, select the **Alert Filter** tab.
7. In the Condition area, select **And**.
8. Add a condition:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Role in Group**
- **is not**
- **Secondary**

9. Select **Update**.

Configure alert remediation actions

Run commands to perform alert remediation on remote Linux and Windows CIs.

Before you begin

Install the following plugins:

- IntegrationHub Enterprise spokes plugin
- ServiceNow IntegrationHub Action Step - PowerShell
- ServiceNow IntegrationHub Action Step - SSH

For Linux: Ensure that the remote Linux host is reachable using the MID Server IP address that runs action scripts.

For Windows:

- Ensure that Powershell v3.0 - 5.0 are supported and installed on the Windows target server.
- Enable Powershell remoting on the remote Windows target server.
- Ensure that the remote Windows target server is reachable using the FQDN from the MID Server that runs action scripts.

Add credentials to the credential aliases that come with the base system:

- linuxAdmin
- windowsAdmin

Role required: evt_mgmt_admin or flow_designer

Procedure

1. Access remediation actions in one of the following ways:
   - Quick actions in an alert, as described in Apply a quick response in an alert.
   - The Playbook tab in Service Operations Workspace, as described in Run actions to resolve alert issues .
2. Configure additional input in the alert's **Additional Information** field, using this format:

   ```
   {"remediation_action_resource":"<mandatory input>"}
   ```

   The mandatory input per action is described in the following table.

Action name	Mandatory input in Additional information field
Stop Windows Process	Process ID
Restart Linux Service	Service Name
Kill Linux Process	Process ID
Get Windows Processes	N/A

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action name	Mandatory input in Additional information field
Stop Windows Service	Service Display Name
Restart Windows Service	Service Display Name
Get Windows Services	N/A
Stop Linux Service	Service Name
Start Windows Service	Service Display Name
Suspend Windows Service	Service Display Name
Get Linux Services	N/A
Run Linux Top for Memory	N/A
Show Disk Space Used	N/A
Run Linux Top	N/A
Start Linux Service	Service Name
Show Disk Space Used	N/A

Result

After executing an alert action, the results are noted in the *Work Notes* field of the alert.

Related topics

- Troubleshooting WMI/Powershell issues on the remote machine
- WMI, PowerShell, and Windows Firewalls
- MID Server: troubleshooting WMI/Powershell issues - Credentials
- How to Enable/Test PowerShell Remoting on Windows Host

CI Remediation

Alert and configuration item (CI) remediations help troubleshoot and resolve underlying problems that generate alerts. Remediation is based on Orchestration workflows that can be scripted to perform remediation tasks such as gathering system information or rebooting a server.

Event Management provides a framework for configuring and implementing a system-wide remediation plan to resolve, minimize, or prevent problems that generate events and alerts. Configure alert management rules to remediate alerts and CI remediation rules to remediate a set of CIs.

Alert remediation

Use alert management rules to associate specific alerts with a remediation workflow. You can apply a workflow to alerts manually. Alternatively, you can configure the alert management rule that executes a workflow automatically for alerts that match the rule criteria. The workflow can also be scripted to request user approval before executing subsequent workflow tasks. The workflow is paused for user approval, even if it was automatically triggered.

CI remediation

Configure remediation more comprehensively by creating CI remediation rules to manually apply remediation to specific CIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remediation tasks

When a remediation workflow is executed, a remediation task is created to capture details such as the time that the workflow started to run, the alert that triggered the remediation workflow (if relevant), and the CI that was remediated. View these tasks to track remediation activities in the organization.

Related topics

Migrate an alert action rule to an alert management rule

Create or edit CI remediation

Create a CI remediation rule that lets users manually apply an Orchestration workflow for resolving issues with specific CIs associated with alerts. Define these CIs in the CI filter conditions of the rule.

Before you begin

Create the workflow to remediate CIs. In the workflow settings, select **Remediation Task [em_remediation_task]** in the **Table** field.

Then submit the workflow definition, and add the following conditions to the workflow properties:

- **[Run Remediation][is][true]**
- **[Workflow][is][<name of the new workflow that you just created>]**

  For example, if the name of the workflow is `CI remediation`, add this condition:

  [Workflow][is][CI remediation]

After you finish configuring the workflow, make sure that you publish it. For more details, see Create a workflow.

Role required: evt_mgmt_admin

About this task

A CI remediation rule associates a set of CIs that might experience problems with a remediation workflow. The remediation workflow can either resolve the underlying problem or help troubleshoot the problem that generated the alert. For example, you can proactively configure a workflow with the appropriate response actions for predictable alerts. For more details, see the Event Management Remediation short video.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **CI Remediations**.
2. Click **New**, or select a CI remediation to edit.
3. Fill in the fields, and click **Submit** or **Update**.

CI remediation form

Control	Description
CI filter	Conditions to match the CIs that you want to apply this remediation to.
Orchestration workflow	The remediation workflow to apply to CIs that match the filter conditions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
In service maps that are opened from the Event Management dashboard, this remediation can be applied to any CIs that match the filter conditions.

View remediation tasks

Event Management automatically creates a remediation task to capture every remediation that was applied to a CI or to an alert. It gives you an overall view of remediation activities in the organization.

Before you begin
Role required: evt_mgmt_user

Procedure

Navigate to **All** > **Event Management** > **Remediation Tasks**.

Column	Description
Number	Number assigned to this task.
Configuration item	Item or service affected by this task.
Priority	Priority that was assigned to this task, based on impact and urgency.
State	State of the task, which can be open or closed.
Assigned to	Person that this task was assigned to.
Short description	Description for the task.
Task type	This task type is a remediation task.

How alerts work with CIs in maintenance

When a CI is in maintenance, the impact tree, the service map, and Alerts tab are updated based on various factors.

Watch this brief video to learn about how alerts work with CIs.

A CI is in maintenance when:

- A change request is scheduled for the CI.
- The **Install Status** field on the CI record is set to **In Maintenance**.

Note: To customize how alerts work with CIs in maintenance, see Create maintenance rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How CIs in Maintenance appear in Event Management	Description and the optimal time to resolve the alert
CI with an active change request 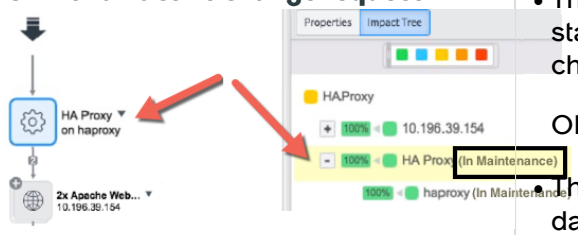 If an active change request is scheduled for the CI or if the **Install Status** of the CI is **In Maintenance**, all alerts on the affected CI are excluded from impact calculation. The Alerts tab also temporarily hides all corresponding alerts. The impact tree shows the CI in green with a note of **(In Maintenance)**. The impact tree and the service map temporarily show CIs in green. For a service, all alerts on CIs in the service are also hidden from the Alerts tab. The entire service is shown in green on the impact tree. For a host with an active change request, the host applications are considered as one unit. All child applications are treated in the same manner as the host until the change request is no longer active.	An active change request has the following values: • The **State** is **Scheduled** or **Implement**. • The current time is between the Planned start date and Planned end date on the change request. OR • The current time is between the Actual start date and Actual end date on the change request. You can monitor the progress of the change request. Wait until the change request moves to the **Review** or **Closed** state. Then you can address all alerts for the affected CI and any alerts that generated between start and end dates. Use the impact tree, topology, and Alerts tab to show the calculated impact severity. **Note:** The **Maintenance** check box for an alert is selected when the **Install Status** field on the CI record is **In Maintenance**. This **Maintenance** check box indicates that the alert must be hidden from the **Alerts** tab. When the maintenance job updates this field, the Updated field (sys_updated_on) changes for open alerts, but not for closed alerts.
CI with no active change request and not in maintenance When there is no active change request for a CI and when the CI is not in maintenance, impact calculation resumes. The impact tree, the service map, and Alerts tab show the calculated impact severity for alerts.	An inactive change request has the following values: • The **State** is **New**, **Assess**, **Authorize**, **Review**, **Close**, or has no changes at all. • The **State** is **Scheduled** or **Implement**. However, the current time is before or after the planned or actual date ranges on the change request. You can monitor the progress of the change request. When a CI has an inactive change request, you can address the corresponding alerts as appropriate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How CIs in Maintenance appear in Event Management	Description and the optimal time to resolve the alert
CI with no active change request and in maintenance If no active change request is scheduled for the CI and if the **Install Status** of the CI is **In Maintenance**, all alerts on this CI are excluded from impact calculation. The impact tree shows the CI in green with a note of **(In Maintenance)**. The impact tree and the service map temporarily show CIs in green.	All alerts on a CI with an inactive change request and **In Maintenance** has no impact calculation.

Related topics

Create maintenance rules

SLAs for application services and CIs

Event Management supports the creation of SLAs for application services and for CIs.

You can monitor and manage the quality of the application services offered by your organization. For example, you might want to set up an SLA that is triggered when the max alert severity of an application service or CI is **Critical**, and measure the time it takes to bring the application service or CI back to the **OK** state, when the conditions that caused the alert have been resolved.

SLAs and the Task table

You can create SLA definitions only for tables that extend the Task table. The Event Management application provides a table named Event Management SLA [em_ci_severity_task], which extends the Task table. Use this table in your SLA definitions to specify the severity level that should trigger and stop the SLA. During alert impact calculation, changes in the severity level of application services and CIs are automatically updated in the Event Management SLA table. Scheduled jobs keep the information in this table up to date.

How the Event Management SLA table is populated

The Event Management SLA table is populated differently for application services and CIs:

- For application services, the system automatically populates the Event Management SLA table when an application service is created or when its max severity is changed.
- For CIs, you must first identify which CIs can be made available for SLAs by creating an SLA configuration record. The system then automatically populates the Event Management SLA table with the CI max severity is changed.

Note: Duplicate CIs are not added to the Event Management SLA table even if the same CI matches more than one SLA configuration filter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create an SLA configuration for CIs

Create an SLA configuration from the Event Management application to determine which CIs are available for SLAs.

Before you begin
Role required: evt_mgmt_admin

About this task
The SLA configuration record is a filter on a table in the system. Select a table that has the CIs for which you want to create an SLA.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **SLA Configuration**.
2. Click **New**.
3. Fill in the form fields (see table).

SLA Configuration form fields

Field	Description
Name	Provide a descriptive name.
Table	Select a table with the CIs for which you want to create the SLA. You can also use the Configuration Item to Host View [query_based_service_basic_view] table, which is used by Dynamic CI Groups.
Filter	Configure the filter for the records.

4. Right-click the header and select **Save**.

Example:
The following is an example.

If you want to run an SLA on all Linux servers for major and critical alerts, first create an SLA configuration on the Linux Server [cmdb_ci_linux_server] table. Then create the SLA definition with this filter:

- `CI.class` is `Linux Server`
- `Severity` is `Major` or `Critical`

What to do next
Create an SLA definition on the CIs that match this SLA configuration.

Limit the records for the SLA configuration filter

If too many records are returned by the SLA configuration filter, you can add a property to set the maximum number of records.

Before you begin
Role required: evt_mgmt_admin

About this task
By default, the size of dynamic CI group records is limited to 10,000 elements. Exceeding this limit can affect system stability. If the filter configuration for a specified dynamic CI group returns too many records, you can limit the number of records. Set the value of the sa.qbs.max_num_of_cis property to limit the returned records. For more information, see Create a Dynamic CI Group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Enter `sys_properties.list` in the application navigator.
2. Add the following property:
   sa.qbs.max_num_of_cis
3. Set the value to an integer that represents the maximum number of records to return.

Create an SLA definition for a CI or application service

You can create SLA definitions for CIs and application services just as you can for other task records in the instance.

Before you begin
Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Service Level Management** > **SLA** > **SLA Definitions**.
2. Click **New**.
3. Fill out the fields on the SLA Definition form.
   For a description of each field, see Create an SLA definition for a CI or application service.
4. For the **Table** field, select **Event Management SLA [em_ci_severity_task]**.
5. Configure the **Start condition** filter by adding conditions, such as:

SLA filter conditions

For all CIs where the **Severity** is **Critical**		
Type	**is**	**CI** **Note:** Event Management SLAs work only on CI classes from Service Instance [cmdb_ci_service_auto].
Severity	**is**	**Critical**
For a specific service, such as Email, when the **Severity** is **Critical**		
Application Service	**is**	**Email**
Severity	**is**	**Critical**

6. Configure pause, stop and reset conditions.
7. Click **Submit**.

Alert lifecycle configuration

Event Management provides various modules, templates, and properties for configuring alerts and the actions that execute for these alerts.

Configure the alert active interval

The active interval property (**evt_mgmt.active_interval**) determines how Event Management handles a new event that is similar to events that appear on an existing closed alert. The system uses the active interval, event, and existing closed alert information to either reopen the alert and add the event information, or create a new alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: evt_mgmt_admin

About this task
When the new event is received and the alert is already closed, Event Management compares the time since the last alert update (**sys_updated_on**) and the active interval property value. Then, the system either updates the existing alert or creates a new alert (if the original alert is closed), as explained in the following table. If the alert is open, Event Management automatically updates it without checking the active interval property value.

Scenario	Result
The new event occurs before the active interval elapses for the existing closed alert.	The event is identified as a recurrence of an existing issue. The alert is updated with the event information and the alert is reopened.
The new event occurs after the active interval elapses for the existing closed alert.	A new alert is created because the event is identified as a new issue. The event information is added to the new alert.

Change the active interval by editing the **Active interval (in seconds), within which a new event reopens a closed alert** property value.

Procedure
1. Navigate to **All** > **Event Management** > **Administration** > **Event Management Properties**.
2. Change the number of seconds for the **Active interval (in seconds), within which a new event reopens a closed alert** property.
3. Click **Save**.

Configure alert flapping

Set flapping properties to determine when an alert enters and exits the flapping state. Flapping can indicate configuration problems (that is, thresholds set too low), troublesome services, or real network problems.

Before you begin
Role required: evt_mgmt_admin

About this task
Flapping occurs when the event source continues to generate events even after its associated alert has been closed. Flapping causes the resource's status to repeatedly fluctuate between the *OK* severity and a severity requiring attention, for example, *Critical*.

Procedure
1. Navigate to **All** > **Event Management** > **Administration** > **Event Management Properties**.
2. Edit the following properties, as appropriate.

Property	Description
evt_mgmt.update_alert_restricted_fields_elapsed_time	When the system receives events that do not update alerts, system fields are not updated by default. This occurs when the event data is identical to the existing alert. This helps re duce noise. For example, a "keep alive" event

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	might be sent every few minutes from a serv er. This confirms that a device is still connect ed, even if no new data is being exchanged. This property defines the minimum time inter val (in seconds) that must pass before an alert can be updated based on identical events. If an identical event is received within this in terval, the update is suppressed to avoid gen erating frequent, redundant notifications. By default, the interval is set to 86,400 seconds (24 hours).
evt_mgmt.flap_interval	Flap interval, in seconds, which is the time pe riod within which an alert enters the flapping state. An alert enters the flapping state when its cur rent **Flap Count** value reaches or exceeds the given *evt_mgmt.flap_frequency* prop erty value within the time period specified by the *evt_mgmt.flap_interval* property.
evt_mgmt.flap_frequency	Flap frequency, frequency an alert must reoc cur to enter the flapping state.
evt_mgmt.flap_quiet_interval	Flap quiet interval (in seconds), quiet time that must pass for an alert to exit the flapping state. An alert exits the flapping state when the time interval between the latest occur rence of the event and the **Flap last update time** of the alert reaches or exceeds the *evt_mgmt.flap_quiet_interval* prop erty value. An alert can also exit flapping in these conditions: ○ When an incoming event arrives after the *flap_quiet_interval* time has elapsed. The new alert generates and its state value depends on the incoming event severity. ○ After a background job, which runs every five minutes, verifies that no alert updates occurred and the *flap_quiet_interval* time has elapsed. If there are no updates, the alert shows the previous state value that it had before the background job ran.

3. Click **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can monitor the **Event Management** > **All Alerts** list for alerts that are in the flapping state. For details, see View alerts in the flapping state.

Related topics

Work with flapping alerts

View alerts in the flapping state

Alert priority

Determine the order in which to handle alerts according to the alert priority score. Multiple factors determine the alert priority score and this value changes with changes to the underlying factors.

Multi-factor decisions

The value of the alert priority score is a composite of the value of the category and its relative weight. You can configure the factors that make up the score, as described in the following sections, according to *`Category order`*, *`Category value mapping`*, and *`CI type weight factor mapping`*. The categories are:

- *`Accumulated categories`* - These categories accumulate items with different sub-types that may have a different status or weighing factor. For example, services that are factored according to different criticality or services that are factored according to cost.
- *`Choice-list categories`* - A list of items. You can map any of the items to a different value, thereby changing the priority for that item. For example, critical severity, which has the value of 1 by default, can be mapped to 4 to give it a higher weight in the calculated alert priority score.

List of categories

Category	Description
Accumulated categories	
Number of impacted business-services	This category is factored by the business-criticality of the service. An alert that has an impact on one service has a lower priority score than an alert that has an impact on many services. For example, if a router is down, then it might affect many services and a higher priority ensures that this alert is handled before handling another alert that has an impact on fewer services. An alert that has an impact on many non-critical services must be handled with priority. However, if an alert has an impact on only a few services, but they are critical services, then this alert has a higher priority.
CI type	This category distinguishes between various CI types. For example, an alert on a CI type that is a router or switch may have a higher priority than an alert on another CI type, like a server or application.
Number of Secondary alerts	Alerts that have characteristics similar to other alerts are usually grouped together, such as alerts with the same message key. A parent alert which has many secondary alerts has a higher priority than a parent alert that has only a few secondary alerts. All secondary alerts of a parent alert are counted. If a secondary alert is closed, it is no longer correlated to the parent alert. If a

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

List of categories (continued)

Category	Description
	parent alert is closed, all secondary alerts are no longer correlated and their priority value is calculated accordingly.
Choice-list categories	
Alert state	Priority score is not calculated for an alert that is in *Closed* state. The available states with their default values are: • Open = mapping value 1 • Reopen = mapping value 2 • Flapping = mapping value 3
Alert severity	The priority score reflects the level of seriousness of the underlying alert, for example, an alert that has severity of *Critical* has a higher weighted value than an alert with a severity of *Minor*. Priority is not calculated for an alert that has a severity of *OK*. The available levels are: • Critical • Major • Minor • Warning
Role	The priority score reflects the role of the alert. The priority score is higher for an alert that has a role of *Parent*, next is the role of *None*, and the least weight is given to an alert with a role of *Secondary*. A parent alert that has many secondary alerts has a higher value than a parent alert that has less secondary alerts. The available roles are: • Parent • None • Secondary

Structure of the alert priority score

Weighted value for each category

Each category has its own placement in the alert priority score, according to its order and limit. You can configure the order and limit, as described in Modify the alert priority score, below. There are only positive numbers.

All category values are mapped from the *em_alert_priority_category_mapping* table, except for the following:

- CI Type: Mapped from the *em_alert_priority_ci_type* table.
- Secondary: Value is the number of secondary alerts and is not mapped from any table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: To ensure optimal memory protection, alert groups with more than 1000 alerts may not have an accurate priority score.

Example constituents of alert priority score of 7302020.003

Category	Score
Services	(7.0, 1000000.0) For example, if there are 2 impacted services whose business criticality priority score is as follows: • Service1: Business criticality priority score = **3** • Service2: Business criticality priority score = **4** The sum of the business criticality priority scores is **7**, so the computed priority value is 7,000,000 (7 * 1,000,000) **Note:** The business criticality priority score is the mapped value (**Value after mapping**) in the application services entries of the Alert Priority Category Mappings table (em_alert_priority_category_mapping).
Severity	(3.0, 100000.0)
CI type	(20.0, 100.0)
Role	(2.0, 10.0)
Secondary	(0, 0.01)
State	(3.0, 0.001)

Using the scores from the preceding table, the alert priority score is computed as follows:

```
(Services value * 1,000,000) + (Severity value * 100,000) +
(CI type value * 100) + (Role value * 10) + (Secondary value) +
(State value * 0.001)
```

Replacing the values of the categories from the preceding table into the formula:

```
(7 * 1,000,000) + (3 * 100,000) + (20 * 100) + (2 * 10) + (0 *
0.001) + (3 * 0.001) = 7,302,020.003
```

In the **Priority Breakdown** area of the **More Information** tab, the alert priority value is displayed as 7302020.003. In the Alerts list, this value is displayed as 7302 (number of thousands).

Note: If you have customized business criticality values that differ from the out-of-box values, add an entry to the Alert Priority Category Mappings table [em_alert_priority_category_mapping].

Effect of category limits

Limits placed on categories enable overflow values to affect the next category that is ranked higher in priority order. Each accumulated category has a predefined limit. If the count goes above that limit, 1 is added to the next higher-order category. For example, if the number of CI types goes over the limit, then 1 is added to the next category (Severity) which is higher in priority: `yyy-zzz becomes yy(y+1) -000`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Display in All Alerts list

In the **All Alerts** list, the alert priority score value is displayed in the **Priority** column. For display purposes only, the actual calculated alert priority score is divided by 1000, while in the database, for calculation and sorting purposes, the full computed value is used.

The detailed computation of the alert priority of the selected alert is displayed in the **Priority Breakdown** area **More Information** tab.

Calculating the alert priority score

Calculating alert priority is based on all of a group's alerts, and not only the group's primary alert. Priority is calculated for alerts that match these conditions:

- State not-equal to **Closed**
- Severity not-equal to **OK**

Alerts that are queued for priority calculation are listed in the `em_alert_trigger_queue` table. In the Scheduled Jobs table, there are two `Event Management - Alert Priority Queue` jobs. These two jobs can run in parallel. You can use *`Insert and Stay`* to create another job. The `Event Management - Alert Priority Queue` job runs once every minute and calculates priorities in batches of 1000 open alerts. There are multiple factors that determine the value of the priority of an alert. A change in severity triggers an immediate update of the alert priority score. In this case, only the severity-part of the priority is changed, together with severity information in the **Priority Breakdown** area of the **More Information** tab.

Triggers that cause recalculation

Triggers that can cause the alert priority score to be recalculated are:

List of recalculation triggers

Trigger	Effect on the alert priority score
New alert	The generation of a new alert is a trigger to calculate the alert priority score.
Existing alert that changes from closed to any other state	Priority is not calculated for a closed alert, so if the alert state is changed from its current state of closed to open, reopen, or flapping, that change is a trigger to recalculate the alert priority score.
Role change, between primary, secondary, or none	Changes in role are a trigger for the alert priority score to be recalculated. For example: • An alert whose type has changed to `Primary`, triggers the alert priority score to be recalculated. • An alert whose type has changed to `None`, triggers the alert priority score to be recalculated. • An alert whose type has changed to `Secondary`, triggers the alert priority score to be recalculated.
CI type	When a CI becomes bound to alert, that change is a trigger to recalculate the alert priority score. This trigger to recalculate the priority occurs once only, when the CI is bound to the alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

List of recalculation triggers (continued)

Trigger	Effect on the alert priority score
Change of number of secondary alerts	A change in the number of secondary alerts triggers the alert priority score to be recalculated.
Severity	A change in severity is a trigger to recalculate the alert priority score. For example, changing from OK to `Major` triggers the score to be recalculated. Unlike the other triggers listed above, a change in severity triggers an immediate update of the severity part of the alert priority score calculation.

Modify the alert priority score

You can change the importance of some categories of the alert priority, by modifying their order and/or their weight, as described below. For example, if the CI type is higher in importance than the number of impacted services, you can change their respective order. As a result, the number of services is now multiplied by 100, while CI type is now multiplied by 1000000.

Note: The changes that you make, using this advanced procedure, change the default method of calculating the alert priority score. Alerts that might otherwise not have a high score, by changing these configurable values, change the way you determine the order in which to handle alerts.

Changing category order

Navigate to the *em_alert_priority_category_order* table. In the **Order** column, you can change the order of the required category.

Category value mapping

The *em_alert_priority_category_mapping* table shows the configuration value for each category choice. Each value in a drop-down list of categories can be remapped to a different value by configuring this table.

CI type weight factor mapping

Navigate to the *em_alert_priority_ci_type* table. You can add a new CI, for example, `DualCoreCPU` and in the **Priority** column, you can change the priority of the required category. In addition, you can edit existing Type and also change its priority value. This table is used to map the value of each CI type, for example, a mainframe that is CI_type: cmdb_ci_mainframe might have a priority of 80, while a server with CI type: cmdb_ci_server might have a priority of 60. The mapping enables you to customize the priority of the various Ci types.

Alert assignment groups for teams

Alert assignment groups assign alerts to the right teams promptly and automatically, improving overall incident management capabilities.

Admins determine the alert assignment group priority, making sure that alerts are routed to the appropriate team members.

The Alert assignment group is determined by the *evt_mgmt.alert_auto_assignment_field* property value. As an admin, you can choose to set the assignment group based on the alert fields using the property value and the order of the comma-separated fields by editing the property value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Assigning alert assignment group precedence
- Team-based integration properties
- Create or edit an event rule
- Create event field mappings

Assigning alert assignment group precedence

Assign alert assignment group precedence to make sure that alerts are routed to the appropriate team members.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** and search for `sys_properties.list`.
2. In the **Search** field, enter `evt_mgmt.alert_auto_assignment_field` and select the property record.
   For more information on the property, see Team-based integration properties.
3. In the **Value** field, enter the assignment group precedence.
   Set the assignment group precedence according to the order of the comma-separated fields such as cmdb_ci.support_group, IntegrationGroup.
4. Select **Update**.

Related topics

- Team-based integration properties

Configuring Express List views for users and user groups

Centrally control what users monitor by predefining views in Express List and assigning them to users and user groups.

As an Event Management admin, you can create Express List views to ensure that users focus on specific services, priorities, or alerts. Event Management enables you to define the filters, column order, and filter attributes for an Express List view and then assign the view to individual users and user groups.

Users belonging to a group can only see the predefined views you assigned to that group and their own saved views. If a user belongs to multiple groups, they can see the views you assigned to all of these groups. Users can switch between views and modify and save all views except their assigned default view. However, they can't override settings in any of their assigned views and modify or share their default view.

The filter conditions builder enables you to configure AND and OR conditions between attributes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Express List view filter conditions

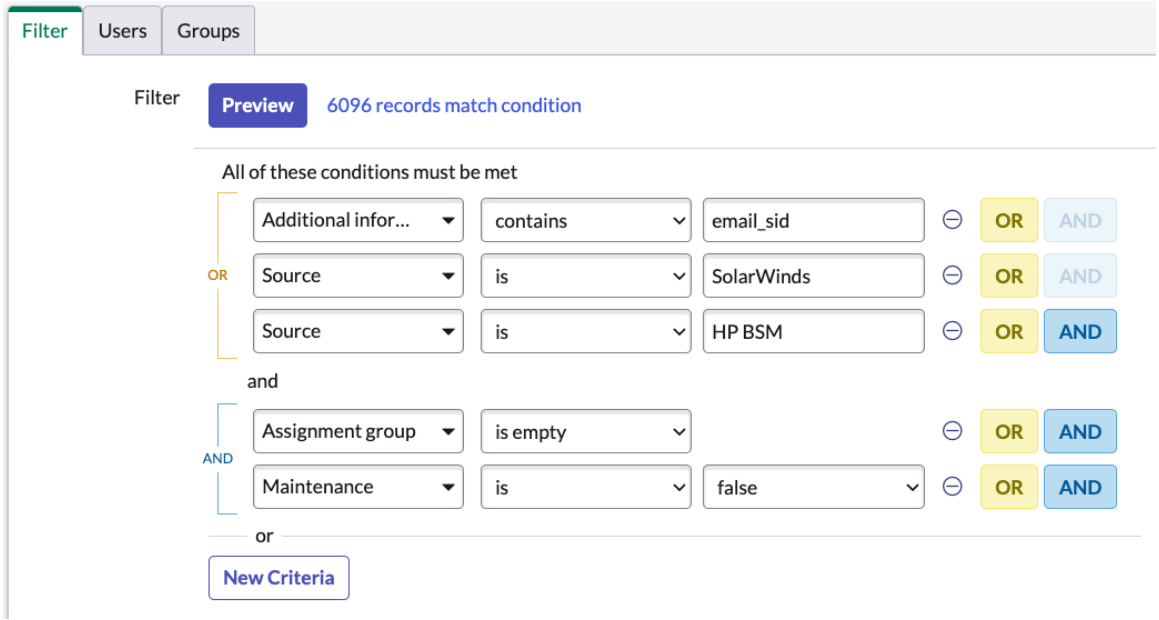

For example, the admin can configure conditions that allow the user to see only alerts that are not assigned to an assignment group or in maintenance. In the example, the alerts must also contain a value of email_sid and must be from one of the specified sources.

Related topics

Create a predefined Express List view for users

Identify an assigned Express List view

Create a predefined Express List view for users

Configure an Express List view for users to make sure that they focus on specific services, priorities, or alerts. You can set the filters, column order, and filter attributes for this view and assign it to individual users or user groups.

Before you begin

For more information about predefined Express List views, see Configuring Express List views for users and user groups.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Express List Views**.
2. Select **New**.
3. On the **Express List View** form, fill in the fields.

Express List View table

Field	Description
Name	Unique name of the Express List view.
Primary Sort By Field	The field by which to sort the Express List. Default: Initial event generation time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Secondary Sort By Field	Optional field by which to sort the results of the primary sort-by field.
Order	The order of this Express List view, if multiple views exist.
Order of Primary Sort By Field	The sorting order of the Primary Sort By Field: Ascending or Descending.
Order of Secondary Sort By Field	The sorting order of the Secondary Sort By Field, if set: Ascending or Descending.
Default Time Range	Time range selected by default when assigned users or groups open this Express List view. Default: **Last 24 hours**.
Default	Option to make this Express List view the default view for the assigned users or groups. If multiple default views exist, the value in the **Order** field determines their order.

4. Set filter conditions for the Express List view.

   a. In the **Filter** tab, select a field, an operator, and one or more values for a condition.

   > **Note:** The default field is State. When selecting fields for filter conditions, use only direct reference fields or single-level references and not nested references. Nested references in filter conditions aren't supported. For example, **Configuration#item.Class** #uses a reference field, **Configuration item**, that#contains #another reference field, **Class**.

   b. **Optional:** Add a filter condition or alert tag filter condition.

Task	Actions
Add a filter condition	i. Select either the AND or the OR operator. ii. Select a field, an operator, and one or more values for the condition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Task	Actions
Add an alert tag filter condition	i. In the Alert Tags Filter section, select **Insert a new row**. ii. Select the required tag, operator, and value for each tag filter. iii. Optionally add additional filter conditions by selecting **New criteria** and adding the tag filter condition.

c. View the number of results a query with the configured conditions would return by selecting **Preview**.

5. Assign users and groups to the Express List view.

Task	Procedure
Assign a user to the selected Express List view	a. In the **Users** tab, select **Insert a new row**. b. Select the search icon and search for the relevant user in the **Users** table. c. Select the user. d. Select the save icon.
Assign a user group to the selected Express List view	a. In the **Groups** tab, select **Insert a new row**. b. Select the search icon and search for the relevant group in the **Groups** table. c. Select the group. d. Select the save icon. **Note:** Users who belong to a group can only see the views that you assigned to that group. If a user belongs to multiple groups, they can see the views you con figured for all of these groups.

6. From the form menu, select **Save**.
   The following tabs display:
   - The **Express List Table Fields** tab shows the fields to be displayed as table columns in Express List.
   - The **Express List View Field Attributes** tab shows the filter attributes for the open alerts in Express List.
7. **Optional:** Add columns to the Express List.

   a. In the **Express List Table Fields** tab, select **New**.

   b. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Field	The column to add.
Express List View	Name of the view to which to add the column.
Minimum width	Minimum width of the column in pixels.
Order	Position of the column in the table.
Active	Option to display the column.

c. Select **Submit**.

8. **Optional:** Add filter attributes to customize this view for assigned users.
These filters are automatically applied to the Filters pane in Express List when users select this view.

a. In the **Express List View Field Attributes** tab, select **New**.

b. On the form, fill in the fields.

Field	Description
Field	Attribute to add to the Filter pane.
Express List View	Name of the view to which to add the attribute.
Order	Position of the attribute in the pane.
Active	Option to activate the attribute.

9. **Optional:** Exclude fields from the Express List view by deactivating them.

a. In the **Express List Table Fields** tab, select the information icon for the relevant field record.

b. Select **Open Record**.

c. On the field form, clear the **Active** check box.

d. Select **Update**

10. **Optional:** Exclude field attributes from the Express List view by deactivating them.

a. In the **Express List View Field Attributes** tab, select the information icon for the relevant field attribute record.

b. Select **Open Record**.

c. On the field form, clear the **Active** check box.

d. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can view the updated state of a field or a field attribute in the table by selecting the list controls icon and then selecting **Refresh List**.

11. When you're satisfied with the Express List view, select **Update**.

Assign users and groups to predefined Express List views

Assign individual users and user groups to preconfigured Express List views to make sure that they focus on specific services, priorities, or alerts.

Before you begin

For more information about predefined Express List views, see Configuring Express List views for users and user groups.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Express List Views**.
2. Select the predefined Express List view that you want to assign to a user or group.
3. Assign users or groups to the selected Express List view.

Task	Procedure
Assign a user to the selected Express List view	**a.** In the **Users** tab, select **Insert a new row...**. **b.** Select the search icon ( ) and search for the relevant user in the **Users** table. **c.** Select the user and then select the save icon ( ).
Assign a user group to the selected Express List view	**a.** In the **Groups** tab, select **Insert a new row...**. **b.** Select the search icon ( ) and search for the relevant group in the **Groups** table. **c.** Select the group and then select the save icon ( ). **Note:** Users belonging to a group can only see the predefined views you assigned to that group and their own saved views. If users belong to multiple groups, they can see the views you as signed to all of these groups.

4. Select **Update**.

Alert impact calculation

Impact calculation shows the magnitude of an outage on CIs, services, alerts, and alert groups. The system uses factors such as impact rules and CI relationships to calculate the severity

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

of a generated alert. The severity appears on the impact tree, application services maps, and dashboards.

Impact calculations are available for application services alert groups. The following factors are used to calculate the overall impact of an outage.

Note: Impact calculation is supported only for services of type Service Instance [cmdb_ci_service_auto] and services that extend this class. Impact is not calculated for Technology Management Service [cmdb_ci_service_technical] or Business Service [cmdb_ci_service_business] classes.

- Impact rules.
- Number of related active alerts.
- Past history of the affected CI.
- Relationships between CIs for a particular application service or application services.
- If the CI element includes a network or storage devices.
- Alerts on CIs in the maintenance state are excluded from impact calculation.

  Note:

  - CIs are considered to be in maintenance not only when an active change request is scheduled, but also when the **Status** field of the CI is set to **In Maintenance**.
  - When a child CI is put in maintenance, it also places the parent CI in maintenance.

- By default, impact is calculated for all operational application services. However, the system allows you to filter impact calculation by service class or by individual application service. For more information, see Add CMDB tables or classes for impact calculation and Add application services for impact calculation.

If there is a connection between services, the impact of one service on the other is also calculated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Impact calculations use information from various sources to set the alert severity

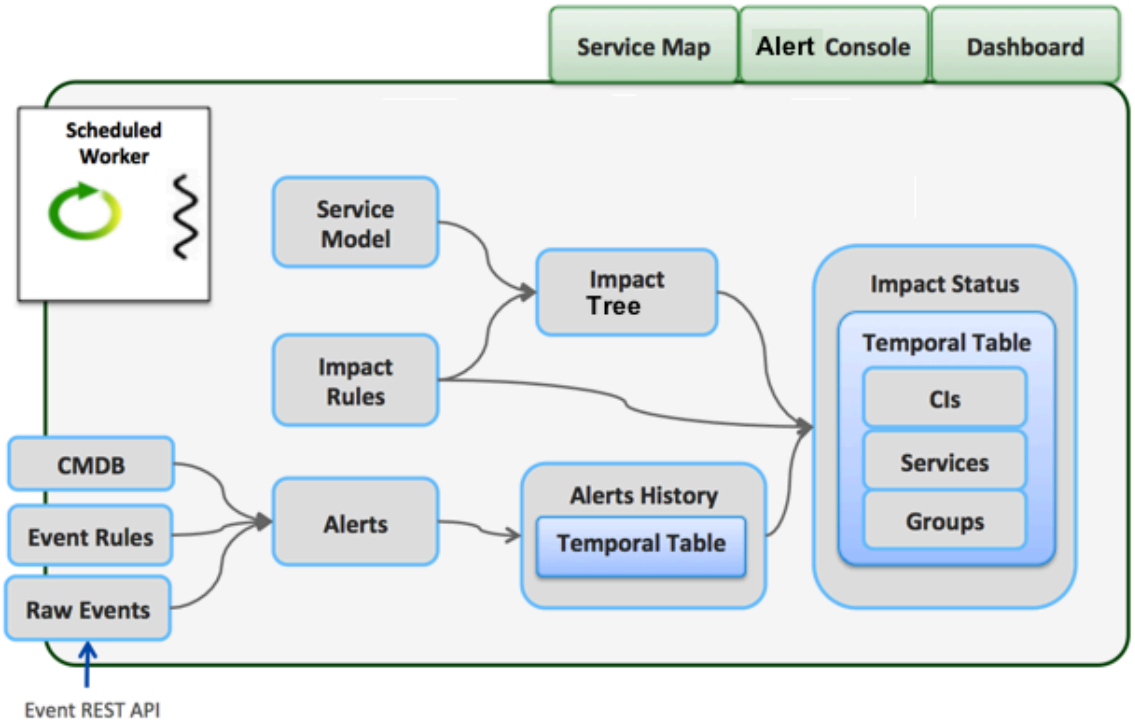

How impact is calculated

Impact calculation varies depending on the CI relationships for a application service or application services. Additional factors, such as change requests, network paths, storage paths, and related CIs all affect impact calculation.

Services

The following impact calculation flow operates for alerts where the outage does not affect a network or network storage. Event Management performs the following steps:

1. Create a service map. Use the Service Configuration Item Associations [svc_ci_assoc] and CI Relationships [cmdb_rel_ci] tables to create child-parent relationships in the application service or application services.
2. If there is no CMDB path from the service to the CI but an association appears in the svc_ci_assoc table, show a Depends-on relationship between the application service and the CI. Otherwise, show no connection.
3. For an application services, if the CIs assigned to the service are also connected to the service in the CMDB, the map keeps the hierarchy between CIs as they appear in CMDB. The CI service assignments appear in the Service Configuration Item Associations section of the Application service form. If there is no connection to the service in the CMDB, the CIs appear directly under the application services in the map.
4. Create the impact tree. Mark the magnitude of an outage by 100% down, 60% affected, 40% impaired, or 20% impaired. If the items in two or more clusters are affected, the impact is 100% down.

Change requests and the In Maintenance status

If an active change request is scheduled for the CI or if the **Install Status** of the CI is **In Maintenance**, all alerts on the affected CI are excluded from impact calculation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Alerts tab also temporarily hides all corresponding alerts. The impact tree shows the CI in green with a note of **(In Maintenance)**. The impact tree and the service map temporarily show CIs in green.

Note:

- CIs are considered to be in maintenance not only when an active change request is scheduled, but also when the **Status** field of the CI is set to **In Maintenance**.
- When a child CI is put in maintenance, it also places the parent CI in maintenance.

For a service, all alerts on CIs in the service are also hidden from the Alerts tab. The entire service is shown in green on the impact tree. For a host with an active change request, the host applications are considered as one unit. All child applications are treated in the same manner as the host until the change request is no longer active. For additional information, see How alerts work with CIs in maintenance.

Network paths

To account for network redundancy, Event Management uses a separate impact calculation. You can see network topology or path changes in the application service. The following impact calculation flow operates for alerts where a network path is affected. Event Management performs the following steps:

1. Create a application service map for the affected network.
   - Use the host ID and target IP information from the alert and the network path from the Network Paths [sa_network_paths] table.
   - Use the elements in the network path that derive from the Configuration Item [cmdb_ci] table. Also, use the elements that are associated to the path, from the Infra Path To Elements [sa_infra_path_assoc] table.
   - Set the relationships. The application CI has a Depends on::Used by relationship on an element in the path that is defined in the CI Relationship [cmdb_rel_ci] table. In the relationship, the application CI is the parent and the element in the network path is the child.
2. Calculate a separate severity for each regular element in the path. Each regular element in the path contributes its own severity to its ancestors up to the application CI where the path originated from.
3. Calculate all redundant elements in the path with the redundancy rule by reducing the severity on the impacted CIs by one level. For example, if the severity is *`Critical`*, the redundancy rule decreases severity by one level to *`Major`*.
4. Create the impact tree. Mark the magnitude of an outage by 100% down, 60% affected, 40% impaired, or 20% impaired. If the items in two or more clusters are affected, the impact is 100% down.

Storage paths

To account for storage device redundancy, Event Management uses a separate impact calculation. You can see impact tree updates when the network storage topology changes from the application service. Event Management performs the following steps for alerts that contain storage CIs:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Create a application service map for the affected storage device:
   - Use the storage device in the sa_fs_to_storage_path table. The storage device definition uses the file system information in the path.
   - Use the elements in the storage path that derive from the Configuration Item [cmdb_ci] table. Also, use the elements that are associated to the path from the Infra Path To Elements [sa_infra_path_assoc] table.
   - Set the relationships. The application CI has a Depends on::Used by relationship on an element in the path that is defined in the CI Relationship [cmdb_rel_ci] table. In the relationship, the application CI is the parent and the element in the storage path is the child.
2. Calculate a separate severity for each regular element in the path. Each regular element in the path contributes its own severity to its ancestors up to the original application CI the path.
3. Use the redundancy rule to calculate redundant elements in the path by reducing the severity on the impacted CIs by one level. For example, if the severity is *`Critical`*, the redundancy rule decreases by one level to *`Major`*.
4. Create the impact tree. Mark the magnitude of an outage by 100% down, 60% affected, 40% impaired, or 20% impaired. If the items in two or more clusters are affected, the impact is 100% down.

Related CIs

As alerts generate for a CI, additional impact calculations run for related CIs. For example additional impact calculations run for a application service dependency to a CI that is not actually part of the application service. These related CIs are not discovered as part of the service. Instead, the related CIs are specified by an infrastructure relationship definition.

The following impact calculation flow operates for alerts with CIs that have a dependency to related CIs which are considered outside the application service. Event Management performs the following steps:

1. Derive relationships between the application service CIs and related CIs. Use the relationships, impact rules, and other data from the Infrastructure Relations [em_impact_infra_rel_def] table.
2. Add related CIs to the impact tree and alert list on the Event Management dashboard.
   - Use data from the Infrastructure Relationship [em_impact_infra_rel_def] table to show containment links to the host.
   - Use the Impact Status [em_impact_status] and Alert History [em_alert_history] tables to determine the status.

Impact rules

Impact rules, which are used for impact calculation, estimate the magnitude or severity of an outage based on affected CIs.

The Impact Rule [em_impact_rule] table contains impact rules that show the applicable CIs, application services, and settings for impact. The following default impact rules are available.

Application Cluster Member

Determines how application cluster members affect the overall impact of the cluster. For example, if a three-member cluster requires **90% Influence** to set the severity for the entire cluster to **Major**, each member has **30% Influence** (90% divided by 3).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The severity of the entire cluster can only change to **Major** when all three members have a severity of **Major**.

You can configure different impact rules per cluster and thus the child CI impact propagation to the parent (for the same child CI) will be different. Therefore, you can manually create groups of Cis (aka manual clusters) and configure the impact rule at the cluster level for downstream towards the cluster children.

Example where the same child CI will propagate its impact to its parent cluster differently to each cluster

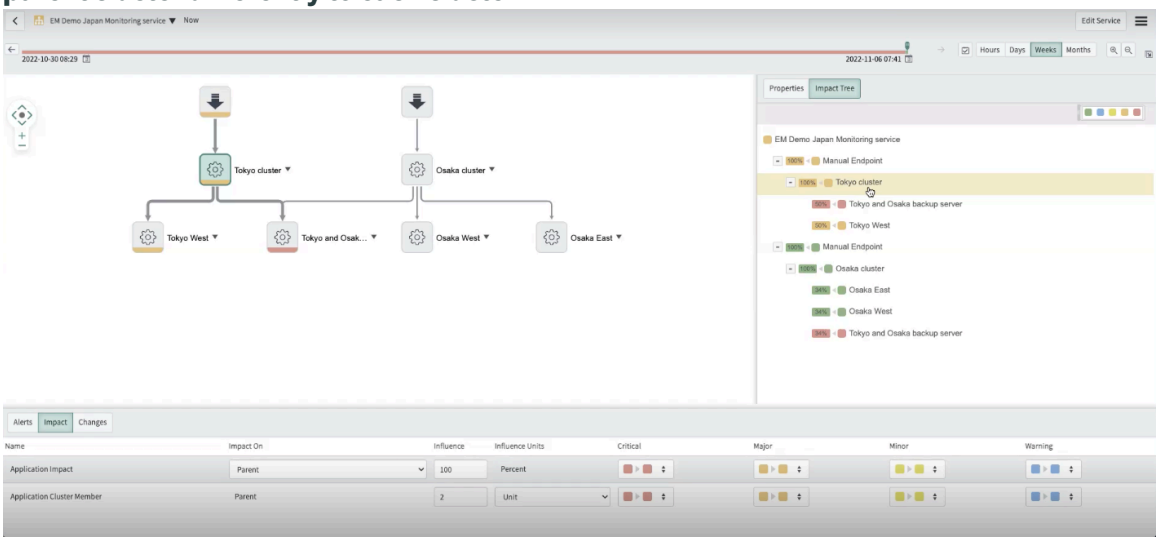

In the above example, there are two entry points. The Osaka cluster on the right-hand side has three CIs. The Tokyo cluster on the left-hand side has two CIs. The Tokyo and Osaka backup server has shared parents - Tokyo cluster and Osaka cluster. On the right-hand panel you can see the Impact Tree where the Tokyo cluster has two Application Cluster Members with 50% influence each and the Osaka cluster has three with 34% influence each.

For manual cluster configuration there are two rows: Application Impact and Application Cluster Member. The children are displayed since the Impact On field was selected as Parent and not Application Service. In the Application Cluster Member row the Influence field is configured to two. This implies that the minimum amount of children who fail (and that they propagate the failure upwards to their parents) are two. The Osaka cluster is configured to three. The percentage is different for the Tokyo and Osaka backup server for each cluster (50% and 34%). As you can see, even though the Tokyo and Osaka backup server failure is Red Critical, it influences the parents differently. The Osaka cluster remains green even though the Tokyo cluster failure is Orange Major.

Click a service or CI to see the alerts that are associated with it. For example, if you click the high-level application service, the alerts that are associated with it are displayed in the alert area under the Map View when you select **Alerts**. The alerts listed are those of the selected service. Alerts of child-services are listed when those services are selected.

The following impact fields are displayed when you select **Impact**.

Inclusion

Determines the impact on entities with a Contains relationship. This rule is read-only.

Infrastructure Dependencies

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Determines the definition of impact propagation for CIs in infrastructure relationships.

CI application service

Determines how impact applies to parent or child entities that are part of an application service.

CI Impact

Applies to application services. Determines the relationship between service members. The impact from child to parent CIs is always 100%. For example, the parent impact severity is derived from the child CI with the highest severity.

CI Parent in Application

Sets impact only on the parent entity.

Network Path

Determines how impact applies to parent or child entities that are part of a traditional network.

OS Cluster Member

Determines how host cluster members affect the overall cluster status based on a percentage or number of cluster members. For example, if a three-host cluster requires **60% Influence** to set the severity of **Major**, each member has **20% Influence** (60% divided by 3). The severity of the entire cluster can only change to **Major** when two or more cluster members have a severity of **Major**. The entire cluster is also considered to be down.

Storage Path

Determines how impact applies to parent or child entities that are part of a storage network.

Properties

In addition to configuring impact rules, you can configure properties for impact calculation.

Configure these properties, as appropriate:

Impact Calculation Properties

Property Name	Description
`evt_mgmt.impact_calculation.alert_grouping_support`	Enable alert grouping support.
`evt_mgmt.impact_maintenance.sleep_time_sec`	Minimum time in seconds for checking CI maintenance: checks both the Status field on the CI and any change request schedule for the CI.
`evt_mgmt.impact_calculation.alert_delay_msec`	The delay after creating or updating an alert, before it is used for impact calculation and grouping. Used to compensate for the late arrivals or slow business rules defined on the em_alert table. Default = 2000 msec. Used when alerts and events are processed one at a time (when the `evt_mgmt.max_objs_in_alert_query` property is not defined or is set to **1**).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Impact Calculation Properties (continued)

Property Name	Description
`evt_mgmt.impact_calculation.alert_copy_delay_when_alerts_are_processed_in_ba`	The delay after creating a updating an alert, before it is used for impact calculation and grouping. Used to compensate the late arrivals or slow business rules defined on em_alert table. Default = 30000 msec. Used in large customer environments with high traffic, when alerts and events are processed in batches (when the `evt_mgmt.max_objs_in_alert_query` property is set to a value greater than **1**).

For more information, see Event Management - Impact calculation [KB1157218].

Understand Service Maps

Service maps show active alerts for CIs and the relationships between CIs. By viewing this information, you can better understand the source of alerts and take remediation steps. The service map is available for all application services.

About Service Maps

A service map shows alerts with impacted CIs and CI interdependencies. For example, changes to a connection between a host and hypervisor appear on the service map. As service map definitions change, the service map, alert, and impact information updates accordingly. From the Quebec release, there is enhanced visibility. When the alert is bound to the entry point, the service map identifies an entry point problem.

Note: Alerts will not be displayed for the entry points on the service map.

You can open a service map from these places:

- From the Application services list, you can view service maps for application services.
- From the Monitored services list, you can view service maps for monitored services.

The following icons are used in service maps. The icon shapes are slightly different for application services.

Service map icons

Icon	Description
( )	Represents applications such as Microsoft IIS or SQL servers.
( )	Represents physical and VM computers and servers.
( )	Represents the network starting point. For example, Layer 3 devices appear toward the top of the map, and connected software and services appear near the end of the map.
( )	Shows the number of redundant CIs.
( )	Shows the workload between machines.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service map icons (continued)

Icon	Description
	The gray connector shows a relationship between CIs.
	Each CI with no active alerts box represents a network CI. A gray box represents a CI with no active alerts. Information about the CI is hidden.
	Hides multiple CIs that are designated as redundant.
	An impacted CI displays the color that represents the severity of the alert associated with the CI. • **Critical (red)**: Immediate action is required. The resource is either not functional or critical problems are imminent. • **Major (orange)**: Major functionality is severely impaired or performance has degraded. • **Minor (yellow)**: Partial, non-critical loss of functionality or performance degradation occurred. • **Warning (blue)**: Attention is required, even though the resource is still functional. • **OK (green)**: An alert is created. The resource is still functional. • **No color**: No active alerts.
	Represents a fiber channel, hard drives, or other data storage devices.
	Represents related web services for the network such as NGINX or JBoss web server.

Note: Session timeout settings don't apply to this screen. The session remains connected, even when there is no human interaction. If this setting is a concern, either log out or close the active tab in the browser.

Related topics

Alert impact calculation

View an alert impact on CIs in a service map

View the impact tree

The impact tree shows the relationships between CIs and the relative percentage impact for each child CI. This information is available for both discovered services and application services.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

The impact tree shows the result of impact rules on CI parent and child relationships. The tree represents a service map with CIs from the Service Configuration Item Associations [svc_ci_assoc] table. When related CIs are connected but not part of the service map, the impact for the related CIs can also appear in the impact tree.

The numbers or percentages on a parent CI summarize overall severity based on alerts from child CIs. The parent severity is based on impact rule computations.

You can also learn about the impact tree from the following video tutorial.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

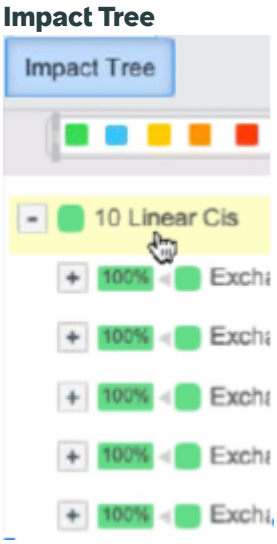

Procedure

1. Open the service map for the service from the Application service list.
   a. Navigate to **Event Management** > **Services** > **Application Services**.
   b. Click **View service** next to the service.
2. On the service map, click a CI that has a severity color.
3. Click the **Impact** tab.
4. From the **Impact Tree** tab in the side pane, expand the parent CI and click the CI that generated alerts.
5. Use the severity color to locate CIs that contain alerts:
   - **Critical (red)**: Immediate action is required. The resource is either not functional or critical problems are imminent.
   - **Major (orange)**: Major functionality is severely impaired or performance has degraded.
   - **Minor (yellow)**: Partial, non-critical loss of functionality or performance degradation occurred.
   - **Warning (blue)**: Attention is required, even though the resource is still functional.
   - **OK (green)**: An alert is created. The resource is still functional.
6. Do one or more of the following.

Option	Description
Adjust impact rules	**a.** Click the **Impact** tab. **b.** Review and adjust the impact rules as nec essary.
Remediate an alert	**a.** Click the **Alerts**. **b.** Right-click the alert and select **Run remedia tion**.

Network paths appearing in the impact tree can use a large amount of resources, which degrades system performance. To disable network paths from appearing

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

in the impact tree, navigate to **System Properties** > **All Properties** and set the *evt_mgmt.network_path_excluded* property to **true**.

To view impact tree using Service Operations Workspace, see View impact tree using Service Operations Workspace.

Adjust impact rules for a CI

Configure impact rules to customize the impact calculation for discovery services and manual services. The impact rules update the overall alert and show the impact on related CIs. When you change impact rules, the updates apply to alert severity in places such as the Event Management dashboard and Operator Workspace.

Before you begin

Role required: evt_mgmt_admin

About this task

You can view and adjust the impact rules of CIs from the application service map of a service.

Note: Event Management Dashboard is not supported for new instances in the Paris release. However, instances upgraded from a release prior to Orlando that use Event Management Dashboard can continue to do so.

Procedure

1. Open the application service map from either the Event Management dashboard or the Application service list.

Option	Description
From the Event Management dashboard	**a.** Navigate to **Event Management** > **Dash board**. **b.** Double-click the tile of the application ser vices.
From the Application Service list	**a.** Navigate to **Event Management** > **Services** > **Application Services**. **b.** Click **View Map** next to the application ser vices.

2. On the application service map, click a CI.
3. Below the map, click the **Impact** tab and then adjust the impact rules accordingly.

Impact Rules

Field	Description
Name	The name of the impact rule. ○ **Application Cluster Member**: Determines how application cluster members affect the overall impact of the cluster. For example, if a three-member cluster requires **90% Influence** to set the severity for the entire cluster to **Major**, each member has **30% Influence** (90% divided by 3). The severity of the entire cluster can only change to **Major** when all three members have a severity of **Major**. **Note:** To configure a manual cluster, see Configure a manual cluster

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ **Application Impact**: Determines how impact is applied to a specific CI, the overall application service, or parent/child entities within an application service. For the Application Impact rule, you can choose from the following Impact options: **Application Service** or **Parent**. ▪ If no CI is selected, the alert rule applies to the entire application service. ▪ If you select a specific CI, you can define whether the rule applies to the service or its parent. ○ **CI Parent in Application**: Sets impact only on the parent entity. ○ **Inclusion**: Determines the impact on entities with a Contains relationship. This rule is read-only. ○ **Infrastructure Dependencies**: Determines the definition of impact propagation for CIs in infrastructure relationships. ○ **Network Path**: Determines how impact applies to parent or child entities that are part of a traditional network. ○ **OS Cluster Member**: Determines how host cluster members affect the overall cluster status based on a percentage or number of cluster members. For example, if a three-host cluster requires **60% Influence** to set the severity of **Major**, each member has **20% Influence** (60% divided by 3). The severity of the entire cluster can only change to **Major** when two or more cluster members have a severity of **Major**. The entire cluster is also considered to be down. ○ **Storage Path**: Determines how impact applies to parent or child entities that are part of a storage network.
Impact On	The valid type of service for this impact rule. ○ **Application Service**: The impact rule applies to a application service. ○ **Parent**: The impact rule applies to the parent CI.
Influence	The value to allow the impact rule to set on the parent. Set a value on the parent. This field works with the **Influence Units** field.
Influence Units	The unit of measurement to show impact on the parent. Set the highest value on the parent CI. Set lower values on each child CI. When alerts occur for several child CIs, Event Management calculates the sum of the **Influence Units** of each affected child. If the sum exceeds the parent **Influence Units** value, the parent receives the highest severity in the set. This field works with the **Influence** field. ○ **Percent**: When the percentage exceeds the **Influence** value, apply the impact on the parent. ○ **Number**: When the number exceeds the **Influence** value, apply the impact on the parent. For example, you can set the parent **Influence Units** to **100%** and the child CIs to these values: ▪ Child 1: 34% ▪ Child 2: 34%

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	▪ Child 3: 34% ▪ Child 4: 70% When alerts for Child 3 and Child 5 have the Critical severity, the severity of the parent is set to Critical because the sum is greater than 100% (34% + 70% = 104%).
Impact when Critical	The alternative impact to use when the calculated severity is Critical. If the column to the right of **Impact when Critical** has a higher severity, update all right-most columns with the same or a lower severity as **Impact when Critical**. To make sure the topology and impact tree accurately show impact, always update each column accordingly. ○ **Critical**: Red (highest severity). ○ **Major**: Orange. ○ **Minor**: Yellow. ○ **Warning**: Blue (lowest severity).
Impact when Major	The alternative impact to use when the calculated severity is **Major**. If the column to the right of **Impact when Major** has a higher severity, update all right-most columns the same or a lower severity as **Impact when Major**. To make sure the topology and impact tree accurately show impact, always update each column accordingly. ○ **Critical**: Red (highest severity). ○ **Major**: Orange. ○ **Minor**: Yellow. ○ **Warning**: Blue (lowest severity).
Impact when Minor	The alternative impact to use when the calculated severity is **Minor**. Controls the field values to the right of this field. If the **Impact when Warning** column has a higher severity, use the **Impact when Minor** value for both fields. To make sure the topology and impact tree accurately show impact, always update each column accordingly. ○ **Critical**: Red (highest severity). ○ **Major**: Orange. ○ **Minor**: Yellow. ○ **Warning**: Blue (lowest severity).
Impact when Warning	The alternative impact to use when the calculated severity is **Warning**. To make sure the topology and impact tree accurately show impact, make sure that the columns to the left have higher severities. ○ **Critical**: Red (highest severity). ○ **Major**: Orange. ○ **Minor**: Yellow. ○ **Warning**: Blue (lowest severity).

What to do next

Review the changes on the impact tree. For example, if you changed a number or percentage influence for child CIs, review the impact tree updates accordingly.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View an alert impact on CIs in a service map

You can view service maps to see active alerts for CIs and the relationship between CIs. By viewing this information, you can better understand the source of alerts and take remediation steps. The service map is available for all application services.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

You can open a service map from these places:

- From the Application services list, you can view service maps for application services.
- From the Monitored services list, you can view service maps for monitored services.

Procedure

1. Open an application service map:
   a. Navigate to **Event Management** > **Services** > **Application Services**.
   b. Next to the service, click **View Map**.
2. Do one or more of the following.

Option	action
View alerts for a CI by type and severity	In the service map: **a.** Click a CI tile. **b.** Below the map, click the **Alerts** tab and re view the listed alerts.
View changes to a CI in a discovered ser vice	In the service map: **a.** Click a CI tile. **b.** Below the map, click the **Changes** tab.
Show alert bindings to CIs	In the service map: **a.** Click a CI tile. **b.** Click the additional actions menu icon to view the **Affected CIs** indicator. **c.** Below the map, click the **Affected CIs** tab.
View CI properties	In the service map: **a.** Click a CI tile. **b.** In the side pane, click the **Properties** tab and review information about the CI. **c.** If you want to view more detailed informa tion, scroll to the end of the pane and click **Detailed Properties**.
Show the impact rules for a CI	In the service map: **a.** Click a CI that has a severity. **b.** Below the service map, click the **Impact** tab and review the listed impact rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	action
Show alert details for networks or storage for a service	In the service map: **a.** Right-click a path between CIs. **b.** Select **Show network path**or **Show storage path**.
Display additional information for CIs	In the service map header: **a.** Click the additional actions menu icon **b.** You can select the **Map Indicators** for additional information that you want to view.
Remediate a CI	In the topology, right-click the CI, and then select **Remediation options**.

Related topics

Alert impact calculation

Understand Service Maps

Create an infrastructure relationship for related CIs

Infrastructure relationships show CIs that are connected to a application service but are not necessary parts of the service. Infrastructure relationships are only available for application services.

Before you begin

Role required: evt_mgmt_admin

About this task

When you create a CI infrastructure relationship, the information is stored in the Infrastructure Relations [em_impact_infra_rel_def] table. When alerts generate, the related CI accompanies the application service information on the Event Management dashboard and in the impact tree. Additional information for the related CIs only appears on a related dependency view map for the application service. The following default infrastructure relationships are available.

Default infrastructure relationships

Infrastructure relationship	Impact rule	Description
cmdb_ci_appl	OS Cluster Member	Shows alert impact between hardware and software applications.
cmdb_ci_esx_server	Infrastructure Dependencies	Shows alert impact between VCenter and ESX clusters.
cmdb_ci_kvm	Infrastructure Dependencies	Shows that alert impact on Linux Kernel-based Virtual Machine (KVM) connectivity.
cmdb_ci_vm_zones	Infrastructure Dependencies	Shows alert impact on the Solaris VM zones.

For example, based on of the cmdb_ci_vm_zones Infrastructure relationship definition, Event Management adds ZoneServer@mmp1 to the application service. The Containment rule manages impact severity on alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related CIs appear on the BSM

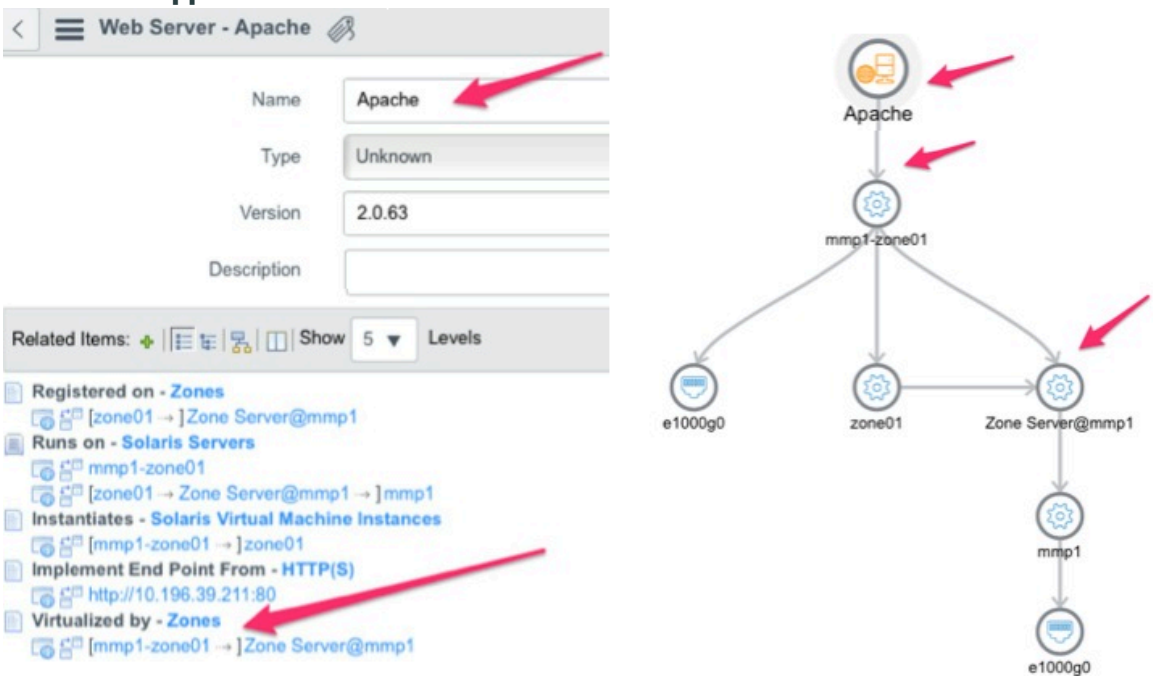

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Infrastructure Relations**.
2. Click **New**.
3. Fill in the fields, as appropriate.

Infrastructure Relations form

Field	Description
Child Type	The table that contains data about the child entity.
Parent Type	The table that contains data about the parent entity.
Relation Type	The relationship between the child and parent entities.
Impact Direction	The impacts direction to show on the application service map. ○ **From Child to Parent**: When an alert is regarding a child, show the impact on the parent. ○ **From Parent to Child**: When an alert is regarding a parent, show the impact on the child entity.
Impact Rule	The impact rule to calculate infrastructure relationships: ○ **OS Cluster Member** : Determines how host cluster members affect the overall cluster status based on a percentage or number of cluster members. For example, if a three-host cluster requires **60% Influence** to set the severity of **Major**, each member has **20% Influence** (60% divided by 3). The severity of the entire cluster can only change to **Major** when two or more cluster members have a severity of **Major**. The entire cluster is also considered to be down. ○ **Application Cluster Member** : Determines how application cluster members affect the overall impact of the cluster. For example, if a three-member cluster

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	requires **90% Influence** to set the severity for the entire cluster to **Major**, each member has **30% Influence** (90% divided by 3). The severity of the entire cluster can only change to **Major** when all three members have a severity of **Major**. ○ **Infrastructure Dependencies** : Determines the definition of impact propagation for CIs in infrastructure relationships. ○ **CI Application Service**: Determines how impact applies to parent or child entities that are part of an application service. ○ **CI Parent in Application**: Sets impact only on the parent entity. ○ **Inclusion**: Determines the impact on entities with a Contains relationship. This rule is read-only. ○ **Network Path**: Determines how impact applies to parent or child entities that are part of a traditional network. ○ **Storage Path**: Determines how impact applies to parent or child entities that are part of a storage network. ○ **CI Impact** : Applies to application services. Determines the relationship between service members. The impact from child to parent CIs is always 100%. For example, the parent impact severity is derived from the child CI with the highest severity.

4. Click **Submit**.

View impact of child service on parent service

In the map view of an application service, the severity of a child service is propagated to its parent service.

Before you begin

Role required:

- for application services, evt_mgmt_admin
- for discovered services, service_mapping_admin

About this task

In the map view of a discovered application service, you can see that the severity of a child service is propagated to the parent service.
An application service is affected by the severity of alerts that are directly associated with it and alerts on the CIs that are in the service. If the severity of a child service is higher than that of its parent service, the parent service shows the child service severity on the map view. In the map view, the **Impact Tree** list shows the CIs. The color of the CIs indicates the alert severity. The child service is also affected by application services that are under it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example where the parent service shows the severity of the child service

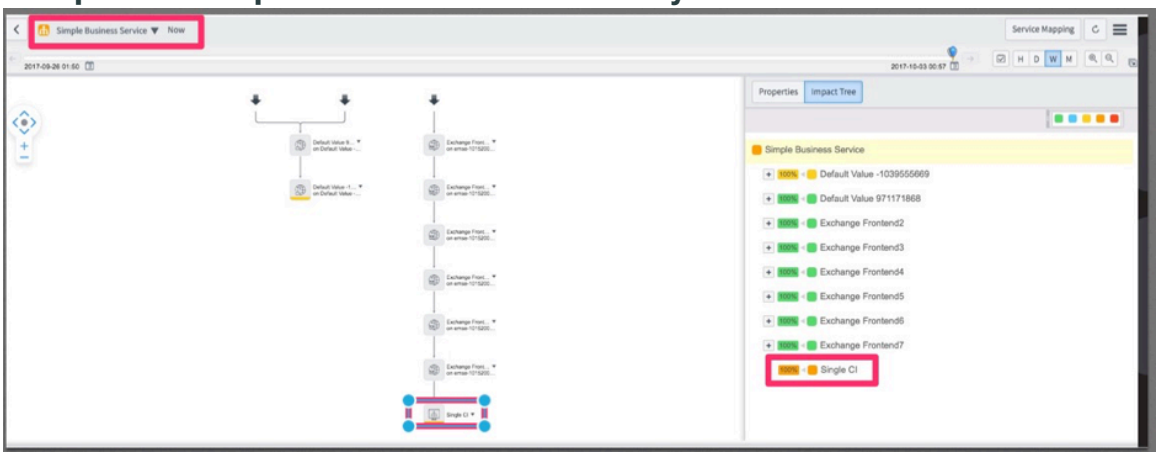

If a loop is created, for example, by connecting a parent service to a sub-service, this relationship can be viewed in the map view. The impact of this loop is not calculated and the Impact Tree does not display this relationship. For example, assume that there are three CIs, A, B, and C, where C is a sub-service to B and B is a sub-service to A. The relationship of the A connected service to C is displayed in the map view. However, this relationship is not shown in the Impact Tree.

Note: A message appears in the system error log in the following format, where the sys_id of the service and the node changes accordingly: `Found loop in ServiceId: [sysId], and node:[sysId]. Removing this relation to avoid impact loops`

For example, `Found loop in ServiceId:409d4fd59f030200fe2ab0aec32e705c, and node:df5cf6159f030200fe2ab0aec32e70a2. Removing this relation to avoid impact loops`

Click a service or CI to see the alerts that are associated with it. For example, if you click the high-level application service, the alerts that are associated with it are displayed in the alert area under the Map View. The alerts listed are those of the selected service. Alerts of child-services are listed when those services are selected.

Example of alerts listed under the Map View

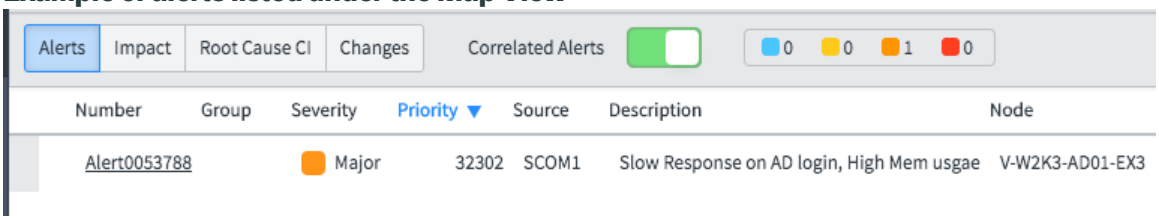

Procedure

Open an application service map:

a. Navigate to **Event Management** > **Services** > **Application Services**.

b. Select **View Service** next to the service.

c. On the right-hand panel, select **Impact tree**.

d. On the lower panel, select **Impact**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

In the service map, you can visually verify that the service health of the sub-service is rolled-up to the parent service. At the bottom of each icon of an impacted CI is a colored bar that represents the severity of the alert. This color bar also displays in the bottom bar of its parent service. For manual cluster configuration, you can modify the Application Impact and Application Cluster Member.

Related topics

- Transfer a map segment into another application service
- Configure a manual cluster
- Adjust alert impact while triaging an alert
- View an alert impact on CIs in a service map

Add CMDB tables or classes for impact calculation

Add the CMDB tables that contain application services to be considered during impact calculation, helping ensure accurate and relevant impact results.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

Organizations may have a vast number of CMDB tables—or classes—such as cmdb_ci_service_by_tags, cmdb_ci_service_auto and many others. However, not all of these are required for impact calculation. To streamline analysis and avoid processing irrelevant data, you can selectively add or remove only those CMDB classes that represent critical infrastructure or components. The procedure below shows how to add a CMDB class. If you want to remove any CMDB table from the list, see Remove CMDB tables or classes from impact calculation.

When you add a CMDB class, all application services related to that class are considered for impact calculation. However, if you do not want to include all services of a specific class and only want to select specific ones, you can do so on the Impact Filter Services page. For more information on how to include or exclude a service, see Add application services for impact calculation.

Procedure

1. Navigate to **All** and search for `em_impact_inclusion_class.list`.

   The Impact Inclusion Classes page opens, displaying a list of CMDB classes that are included in impact calculation.

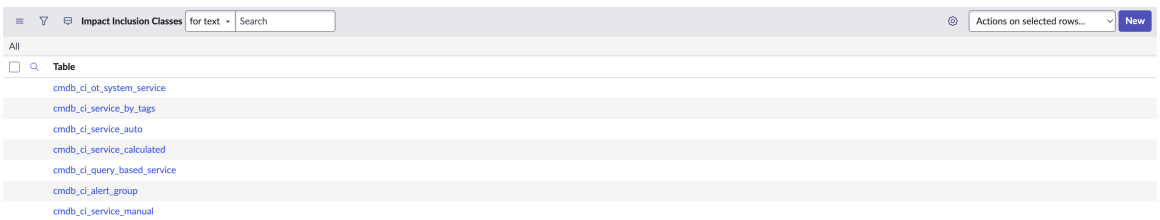

2. Select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

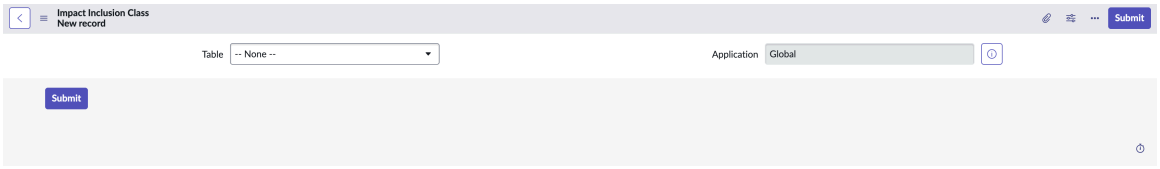

3. From the **Table** field, select a CMDB table.

> **Note:** The list contains only the CMDB tables supported for impact calculation, not all CMDB tables.

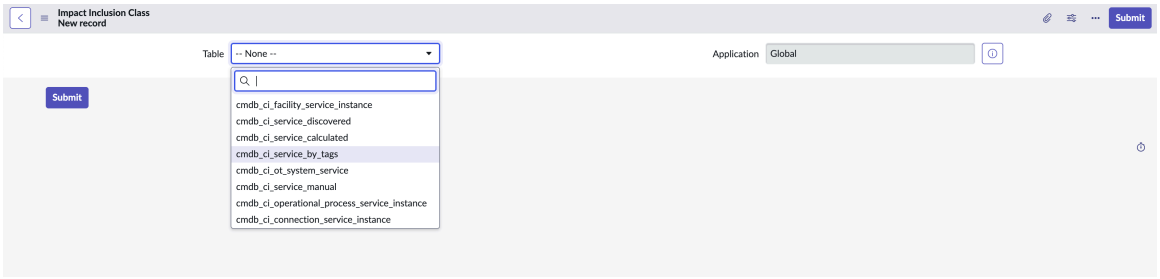

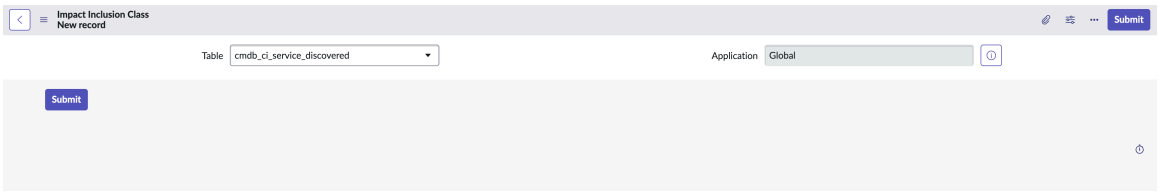

4. Select **Submit**.
   The CMDB table adds to the list of impact inclusion classes.

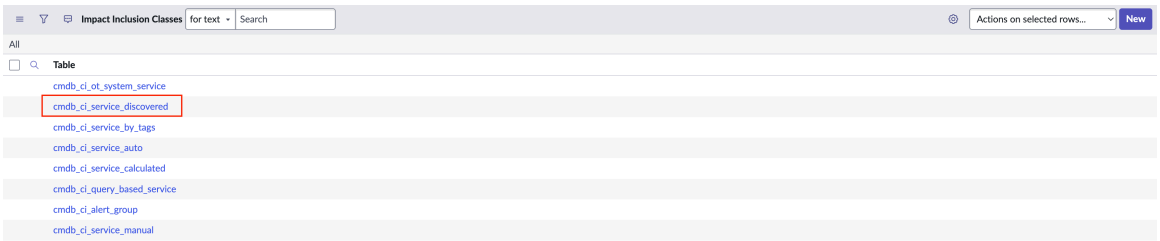

Remove CMDB tables or classes from impact calculation

Exclude unnecessary CMDB tables or classes from impact calculation to improve performance and focus on relevant data.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

Procedure

1. Navigate to **All** and search for `em_impact_inclusion_class.list`.

   The Impact Inclusion Classes page opens, displaying a list of CMDB classes that are included in impact calculation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

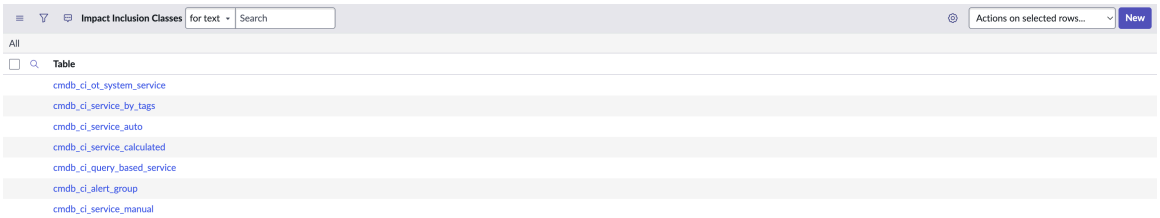

2. Select the CMDB table that you want to remove.
3. Select the Actions on selected rows menu and then select **Delete**.
4. In the confirmation window, select **Delete**.
   The CMDB table is removed from the list.

Add application services for impact calculation

Specify the application services that must be considered during impact calculation to ensure accurate service impact assessment.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

Organizations may have thousands—or even hundreds of thousands—of application services. If a CMDB class is added to the em_impact_inclusion_class.list table, then all application services related to that class are included in impact calculation. However, there are two possible scenarios:

- Scenario 1: CMDB class is added to em_impact_inclusion_class.list but you want only specific services to be considered for impact calculation. You can either:
  - Remove the unwanted services from the em_impact_filter_service.list, or
  - Set their status to false in the em_impact_filter_service.list.

  For information on how to remove the application service from the list, see Remove application services from impact calculation.
- Scenario 2: CMDB class is not added to em_impact_inclusion_class.list. So, none of the application services related to this class are considered for impact calculation. However, you want only specific application services from this class to be included, you can directly add those services to the em_impact_filter_service.list table and set their status to true without adding the CMDB class to the em_impact_inclusion_class.list table.

The following procedure shows how to add an application service directly to the em_impact_filter_service.list table so that it is considered for impact calculation.

Procedure

1. Navigate to **All** and search for `em_impact_filter_service.list`.

   The Impact Filter Services page opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

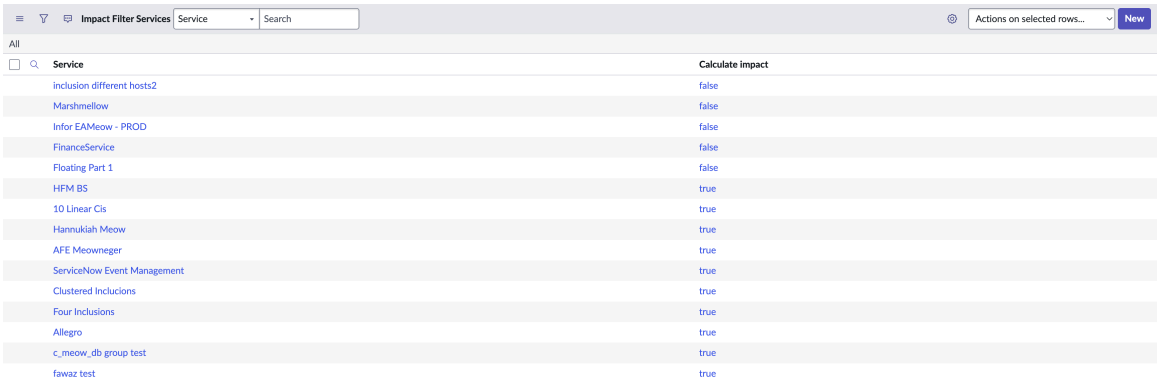

2. Select **New**.

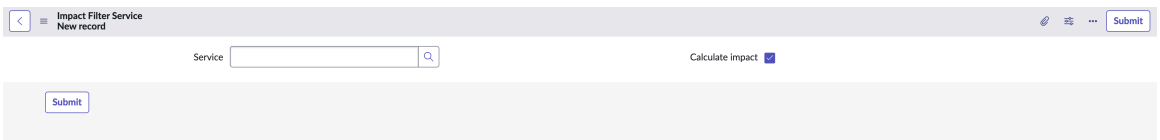

3. In the **Service** field, select the search icon ( ).

Note: By default, the **Calculate impact** check box is selected.

4. In the Service Instances page, select the service that you want to add for the impact calculation.

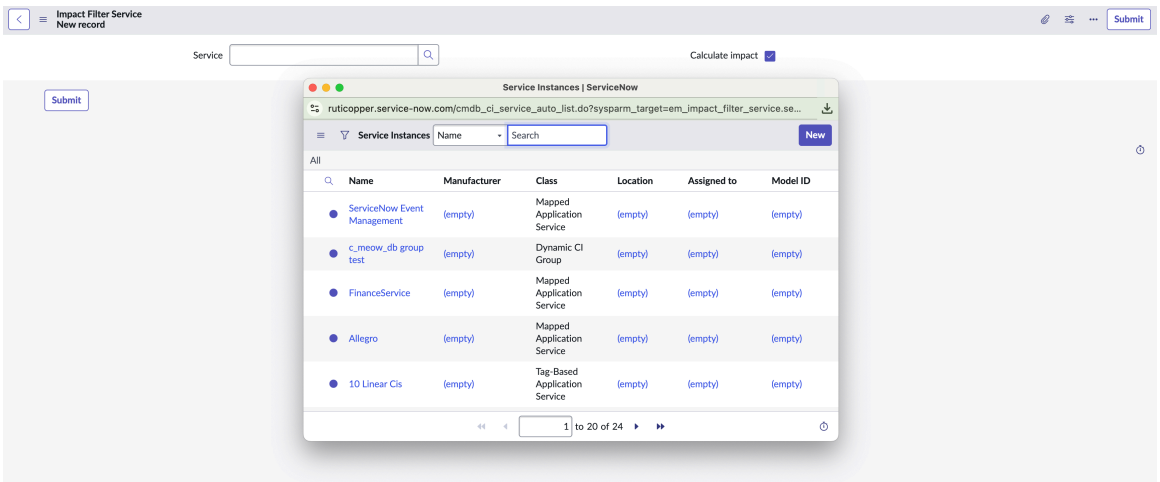

5. Select **Submit**.
The application service adds to the list of impact filter services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

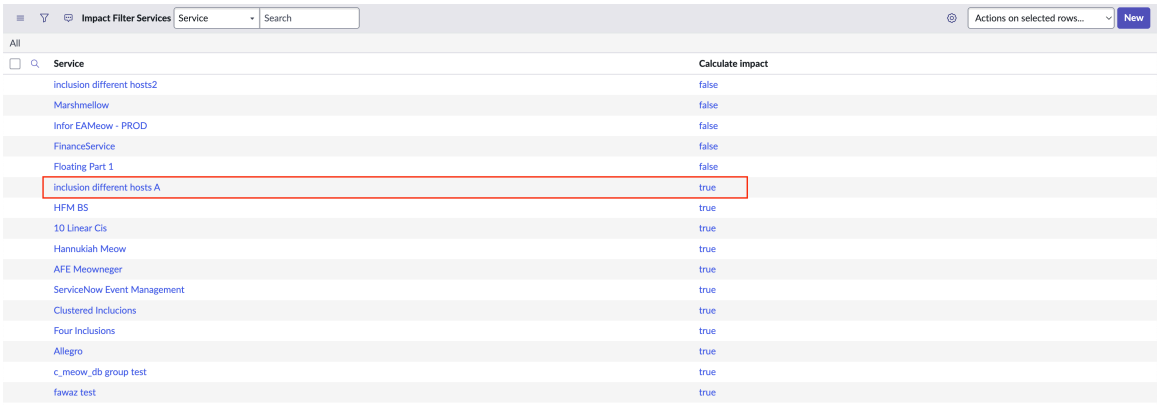

Note: Only the application services added to the Impact Filter Services list appear in the Service Dashboard.

Remove application services from impact calculation

Exclude specific application services from impact calculation to reduce noise and focus on critical components.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

Removing or deleting a record from the list won't necessarily exclude the application service from impact calculation if the CMDB class still exists in em_impact_inclusion_class.list. However, even if the class exists in em_impact_inclusion_class.list, the application service will not be considered for impact calculation if its status is set to **Calculate Impact** = **false** in em_impact_filter_service.list.

For the application service that you want to exclude from the impact calculation, add the application service with **Calculate Impact** = **false**, or change the **Calculate Impact** = **false**, if this record already exist. The procedure shows how to configure this setting.

Procedure

1. Navigate to **All** and search for `em_impact_filter_service.list`.

   The Impact Filter Services page opens.

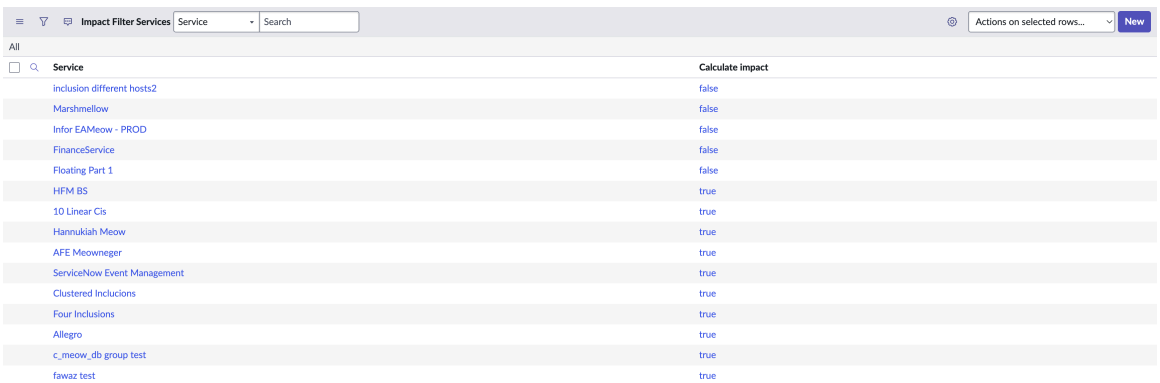

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. For the application service that you want to exclude from the impact calculation, change the **Calculate impact** status of the application service to **false**.
   This keeps the application service record in the list but excludes it from impact calculation.

Probable Root Cause Analysis (RCA)

Shorten the mean time to repair (MTTR) by discovering the root cause of an alert.

Probable root cause analysis (RCA) provides a list of probable root causes for alerts. This capability uses existing data in IT Service Management to help you determine the root cause of an issue.

RCA takes three factors into consideration: Alerts, configuration item (CI) topology, and change requests on the CIs. It extracts CIs from alerts or alert groups and analyzes the CI relationships in the CMDB. It then uses a set of configurable filters to consider relevant change requests on the CIs as possible root causes.

Change requests can be related to a time frame or to a CI, the application service it belongs to, or the software installed on it. For example, a CI must be upgraded or a security patch installed on it. A change request on one CI can affect other CIs.

RCA maps the alerts and change requests to the CIs. It proceeds to calculate the root cause score, a logical score used for sorting probable root causes. The calculation of this score is based on change requests on CIs, and on alerts on a crucial CI called the Topology Origin CI. The root cause score determines how RCA populates the probable root causes list.

Probable root cause analysis list

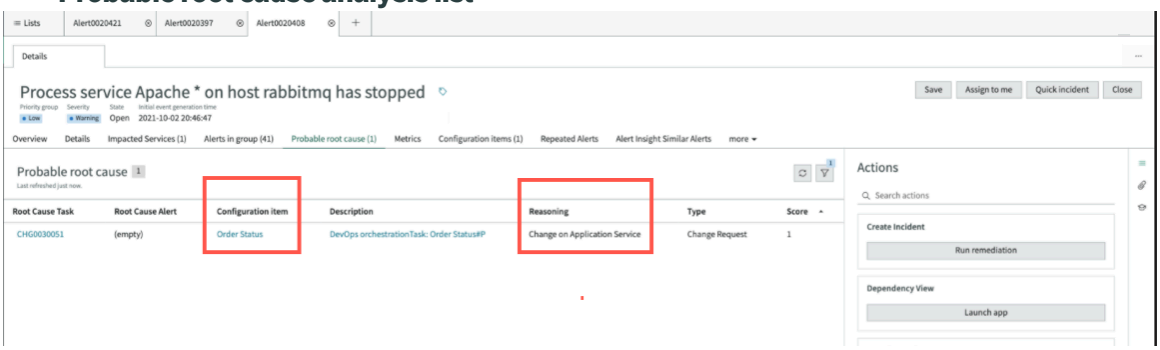

By default, the list shows the five probable root causes with the highest score. It lists the CIs related to each alert, and the reason RCA identified them as probable root causes.

Each probable root cause includes an explanation how RCA correlated the alert with the change request. The reason can indicate that the change request was applied directly to the CI that generated the alert. It can also indicate changes to a CI in the affected CIs list or to a related CI. Additionally, it can indicate changes to an application service or to software installed on the CI or related infrastructure. RCA identifies installed software by using the `cmdb_software_instance` table.

For alert groups, RCA uses topology analysis to identify upstream CIs that might explain cascading failures in the group. Changes on these topology origin CIs receive higher priority in the probable root cause score.

Scoring for probable root causes is determined by the following criteria, in the indicated order (lower number = higher priority):

1. Change on a CI that originates from topology.
2. Change on a related CI originates from topology.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Change on a CI.
4. Change on a related CI/change on an application service/change on software.
5. Alert on the CI that originates from topology.

To disable the Probable Root Cause Analysis feature, you must create the property *sa_analytics.disable_prc* and set the value to `true`. For more information on how to create a property, see Add a system property.

Customize RCA settings

Modify default settings that determine RCA behavior.

Before you begin

Role required: evt_mgmt_admin

About this task

Default filter settings control which change requests RCA considers as possible root causes. You may want to change certain default settings, for example to make RCA ignore standard change requests.

You can also change the number of probable root causes shown in the list. By default, the list shows the five probable root causes with the highest root cause score.

By default the RCA calculation is enabled. You can disable it by changing the default *sa_analytics.disable_prc* value to True.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Do one or both of the following, as needed.
   - Change default filter settings:
     a. Select **New**.
     b. In the **Name** field, specify the name of the property for which you want to change the default setting.

        The new property must have the same name as the related property in the table.

Property	Description	Type	Default value
sa_analytics.rca.filter_category	The categories of change requests that RCA considers as possible root cause.	string	" "
sa_analytics.rca.filter_type	The types of change requests that RCA considers as possible root cause.	string	"standard,normal,emergency"
sa_analytics.rca.filter_state	The state of change requests that	string	"-1,0,3,7,8,157"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description	Type	Default value
	RCA considers as possible root cause.		▪ Implement: -1 ▪ Review: 0 ▪ Closed: 3,7,8,157
sa_analytics.rca.filter_interval_in_hour	The time frame, in hours, during which change requests occurred that RCA considers as possible root cause.	int	168
sa_analytics.rca_use_legacy_query	Use legacy queries for RCA calculation.	Boolean	True
sa_analytics.rca.get_tasks_for_impacted_cis	Include impacted requests for impacted services in the RCA list.	Boolean	True
sa_analytics.rca_scoring_by_alert_proximity	Enables sorting of RCAs by time proximity to the alert.	Boolean	True
sa_analytics.rca_cmdb_relation_hops	Number of hops to search related configuration items.	int	1
sa_analytics.rca_time_filtering_by_planned_time	Enables filtering of change requests by planned time.	Boolean	True
sa_analytics.rca_time_filtering_by_actual_start_time	Enables filtering of change requests by implementation time.	Boolean	True
sa_analytics.rca.upper_filter_interval_in_hours	Determines the time upper bound for change requests filtering.	int	1
sa_analytics.rca.order_by_scoring_for_virtual_alerts	When calculating RCAs for virtual alert – order the records by scoring of PRCs of the corresponding secondaries.	Boolean	True
sa_analytics.rca_same_reasoning_use_newer	When 2 PRCs have the same reasoning, prioritize the one with the most recent change request during scoring calculations.	Boolean	True

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description	Type	Default value
sa_analytics.rca.prefer_cr_on_app_service	In case the service is the same change request on application service and on CI, prefer the one from application service.	Boolean	False

c. Enter a valid value for the property.

> **Note:** The value you set for the new property overrides the default value for the related property.

- Change the maximum number of probable root causes shown in the list:

a. Find the property *sa_analytics.rca_task_table_max_records*.

Property name	Description	Type	Default value
sa_analytics.rca_task_table_max_records	The maximum number of probable root causes shown in the Probable root causes list.	int	5

b. Replace the default value with a new valid value.

- Disable RCA by creating the *sa_analytics.disable_prc* property.

a. Select **New**.

b. Fill in the form with the following values.

Property name	Description	Type	Value
sa_analytics.disable_prc	Disabling the Probable root cause.	Boolean	True

3. Click **Submit**.

Alert similarity

Finding alerts that are similar to the alert that you are currently investigating can help save troubleshooting time by seeing how similar alerts were resolved.

Finding similar alerts

Alert similarity applies natural language processing (NLP) based on ServiceNow AI Platform machine learning (ML) as implemented in Predictive Intelligence.

ML logic filters out irrelevant information and provides further information to help you understand the cause of alerts. You can view solutions used to resolve similar alerts and the probable root cause.

ML helps find similar alerts based on text found in these alert record fields:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Configuration item.Class
- Metric name
- Description

Additional alert information

In the Alert Workspace, select the required alert and click **Insight** to view additional information about the alert.

In the Alerts list, select the required alert and use the Alert Insight tabs on the right side of the alert record to view additional information about the alert. For example, you can view the top recent alerts for the current CI and similar CIs to find the most suitable solution towards resolving the alert.

The alert insight tabs are:

- Impacted Services
- Flapping
- History
- Activities
- More Information
- Repeated Alerts
- Similar Alerts
- CI Incidents
- CI Change Requests
- CI Problems
- Related Incidents
- Related Change Requests
- Related Problems

For more information about the alert insight tabs, see Alert insight information.

Find similar alerts

You can find alerts similar to the alert currently being investigated. Save troubleshooting time by reviewing similar alerts to see how they were resolved.

Before you begin

Ensure that the Predictive Intelligence plugin is activated. To activate this plugin, see Activate a plugin.

Ensure that the `evt_mgmt.similarity_use_ml` property is set to `true`.

Role required: evt_mgmt_admin

About this task

Save time in resolving the current alert by reviewing how similar alerts were resolved. Provide resolution consistency by resolving similar alerts in the same way.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. In the list of alerts, click the alert record that you want to investigate.
3. In the alert record, click the **Similar Alerts** tab.
   Alerts that are similar to the alert that is being investigated are listed.

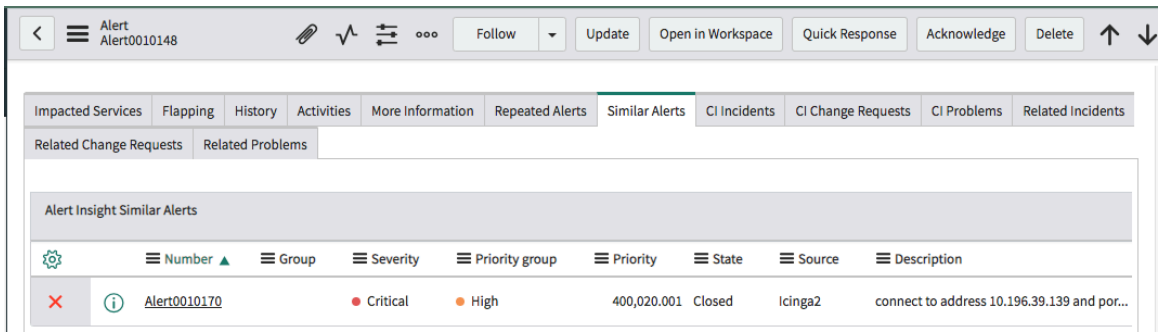

   By default, up to 10 similar alerts are listed. You can modify this value by setting the `evt_mgmt.similarity_max_similar_alerts_shown` property.

4. If you configured a threshold, in the Related links area, click **Show Similarities Details**.
   The Sys ID and confidence level of each similar alert is displayed. The Confidence level is an indication of how the similar alert matches the parameters that you set. As an example, a high confidence percentage value indicates a higher likelihood of similarity of the alerts to the values that you set.

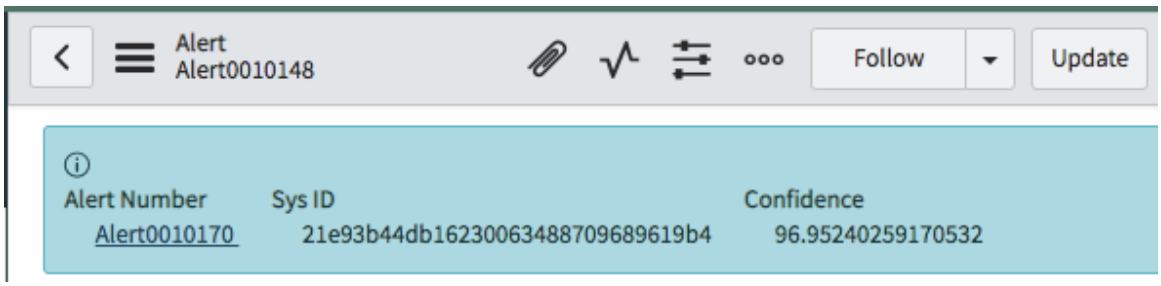

   You can modify the value of the threshold by setting the `evt_mgmt.similarity_use_threshold` property.

Similarity solutions

Similarity solutions enable you to use Machine Learning (ML) to compare the text in a resolved alert record to an open alert record to reuse its resolution approach.

Training a similarity solution

To train a similarity solution, you collect words to compile a collection that Machine Learning (ML) can use to compare text in alert fields. ML compares text in the **Short Description**, **Description**, **Source**, **Type**, **Resource**, and **Metric Name** fields in a resolved alert to see whether the words match words in an open alert. The resolved alert provides an example to show how the open alert can be resolved.

To train a solution, the filter must return at least one record. If your filter returns no records, update it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The preferred number of records for training a solution is between 30,000 records and 300,000. If you submit more than 300,000 records, the most recent 300,000 records are used to train the solution. Use only authentic records from the database.

- Ensure that the records you train are not too old and that they are relevant to your business needs. Keep the words in the collection current.
- Do not use hard-coded dates as filters because these filters are not updated when you retrain solutions unless you update them manually before every retraining. Instead, use relative date filters, for example, the last 3 months, last 6 months, or last 12 months.
- Perform training as needed until it provides an acceptable similarity solution. This practice provides you time to review and update your solution definition.

For more information on similarity solution concept, see Create and train a similarity solution.

Fields to include in the solution

Record the fields that are likely to contain words and phrases that help the system identify similar records for your solution.

The similarity fields that you select should be a subset of your input field selections. For example, if you select fields from incident records that are in Open state, do not select **Close note** as a similarity field. Because open records do not include **Close note** fields, the text cannot be similar.

The similarity fields are available to users when they create records.

About the similarity score

The similarity score is a measure from 0-100 of the degree of similarity between two alert records. Alert records that have a similarity score higher than the threshold that you specify is returned by the solution.

Review similarity examples and their scores using the **Show training progress** feature to determine whether to either increase or decrease the solution threshold. You can change the threshold value in the **Threshold for Similarity Score** field.

View training solution progress

Training times vary based on the number of records and classes within the training set. The more records and classes you use, the longer the training can take. For example, a data set containing 100,000 records and several hundred classes can take around five hours to complete.

To show the training solution progress, the ML solution automatically performs the following activities when you select **Show training progress** on the Solutions page. For more information, see View solution training progress.

Solution training activities

Activity	Description
Fetching files for training.	The system downloads the training records and sends them to the nearest training service.
Preparing the data.	The system removes duplicate records from the training set.
Training the solution.	The training service trains the solution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solution training activities (continued)

Activity	Description
Uploading the trained solution.	The training service uploads the solution as attachment records.

Create an Event Management similarity solution

Create and train a solution that applies machine learning to a collection of words to target and suggest similar alerts in your instance dataset. For example, you can compare the text in a resolved alert record to an open alert record to reuse its resolution approach.

Before you begin

Role required: evt_mgmt_admin

About this task

The encryption scheme that you use can affect whether you can train solutions.

- If your data is encrypted by Full Disc Encryption (FDE), you can train solutions.
- If the field you are using for training is encrypted using Field Encryption, ensure that the shared service worker user has the correct encryption module role.
- If you are using Edge Encryption, training is not supported.

Note: This feature supports only English language processing.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. Click **New**.
3. On the form, fill in the fields.

ML Solution Definition form

Field	Description
Solution Template	Use to define a word collection.
Label	Unique name for the similarity solution.
Name	Field that is auto-populated with a system-assigned name that is similar to the label name when the solution is created.
Table	Table that contains the records to use for the word collection, for example, the Incident table or the Asset table. When you assign a table value, a link appears that shows the number of records that match your current conditions.
Filter	Conditions that you want to apply to the training records. To train a solution, the filter must return at least one record. If your filter returns no records, update it.
Input Fields	Record the fields that contain the text and context that you want to include in your word collection. The field data type can be a string, reference, choice, or HTML, such as **Short description** and **Description**. **Note:** **Journal type** is not a supported data type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Good candidates for your input fields have text that is relevant to the solution. For example, if you are configuring a solution to find similar incident records, you can select the **Short Description**, **Description**, **Resolution notes**, and **Close notes** fields.
Training Frequency	Training is performed once by the system, by default. Because your data can age over time and degrade the accuracy of recommendations, consider invoking scheduled trainings once your solution definition is fairly stable.
Domain	On instances where domain separation is active, select the domain whose target records you want to use for your word collection. Create a separate similarity solution definition record for each domain whose field values you want to use for your word collection.
Similarity Fields	Record fields that are likely to contain words and phrases that help the system identify similar records for your solution. To change your **Similarity Field** choices, click the Lock icon ( ) to open the field and make your updates. Click the icon again to close the field and save your updates.
Similarity Window	The period in which you want to look for similar records. If you have a smaller number of total incoming records, a larger window might be best. However, if the window is too large, you may retrieve records that are not as useful. If the window is too small, you might not retrieve enough similar records. For example, if you are looking for similar incident records that are open, you can select **Last 1 day**. This selection targets the most recent records, many of which could still be open.
Similarity Window Filters	Filter conditions for your similarity window. These filters define the dataset conditions under which your similarity results are determined. For example, if you are looking for similar incident records that are open, you can filter your search by creating conditions such as: **[Incident state] [is] [In progress]** These filters are applied in addition to the **Input Field** filters.
Window Refresh Frequency	Frequency to refresh the similarity window. For example, if your window contains incident records that are open, you can select a refresh frequency value of **Every 15 minutes**. New incidents typically occur frequently throughout the day so this frequency increases the likelihood that newly opened records are included in the refresh. **Note:** If your similarity window is composed of records such as Knowledge article records, which are typically not created often, you can choose a larger refresh frequency such as **Every 1 day**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Click **Save** to save the solution definition record so you can return to it later, or click **Submit & Train** to create the solution definition record and train it.
5. If you submitted the solution for training, click **OK** on the **Training Activation** window to confirm.

Result

The system schedules the solution for training with the nearest training service. You and other users subscribed to the Predictive Intelligence Notifications category receive a notification when the training completes, including any errors that might have occurred in the training. When training is complete, the system uploads the solution as an Attachment record.

What to do next

Review the trained similarity solution examples in the **Similarity Examples** related link on your ML Solution Definition form.

View Event Management solution training progress

View solution training progress or statistics to determine whether a solution is available or how long the next training cycle might take to complete.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solutions**.
2. In the ML Solutions list, select the solution whose progress or statistics you want to view.
3. Click the **Show training progress** related link.
   The ML solution automatically performs activities that appear in the Training Progress pop-up window.

   The system shows a Training Progress pop-up window. In the following example, the progress failed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

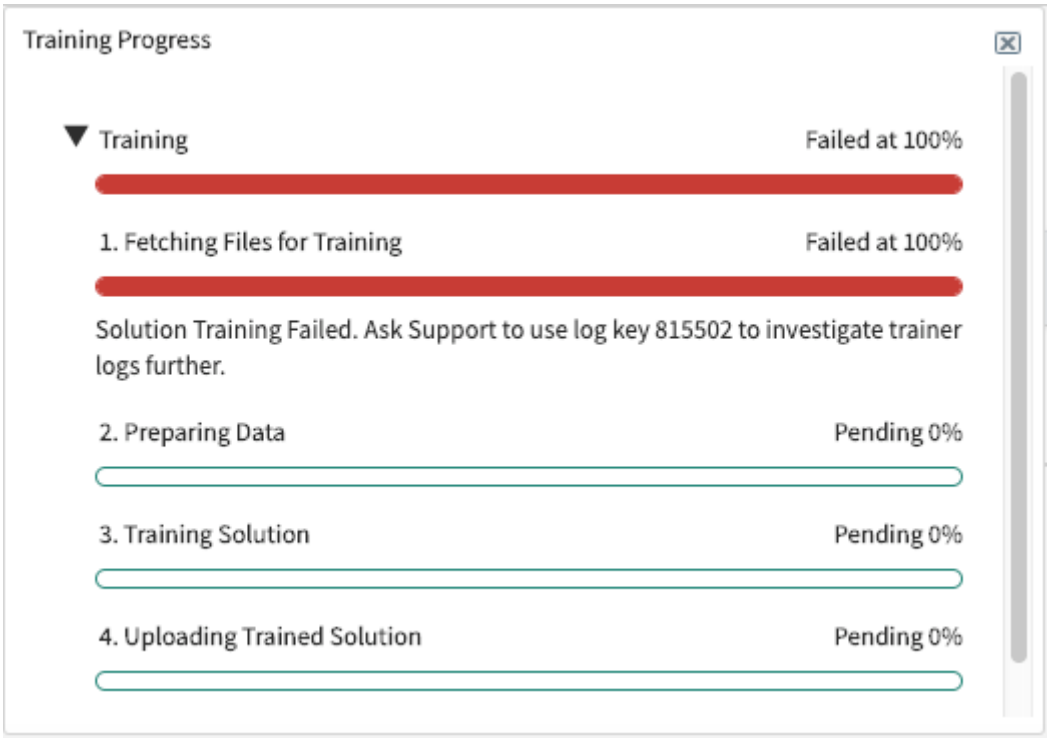

If the training is a success, all bars would be fully blue and the progress for each step would show 100%.

What to do next
See Review Event Management similarity examples.

Review Event Management similarity examples

Review the similarity examples and scores that the system provides during solution training to see how the selected alert record compares to existing alert records. For example, you can modify the similarity score threshold to increase the accuracy of your similarity recommendations.

Before you begin
Train a similarity solution in a ServiceNow datacenter. For more information, see Create an Event Management similarity solution.

Role required: evt_mgmt_admin

About this task
Adjust the current threshold to provide useful results. For example, a low similarity score has more coverage but less precision.

Procedure
1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solutions**.
2. In the ML Solutions list, locate your solution and click the Information icon (ⓘ).
3. Click **Open Record**.
4. Click the **Similarity Examples** related link.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

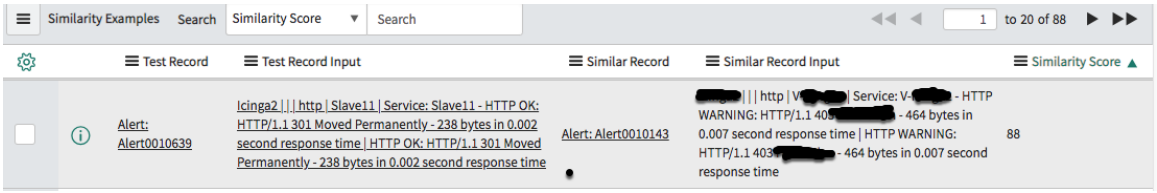

5. Review the similarity examples and their threshold scores to determine the accuracy and coverage levels you want applied to your solution similarity results.
   The higher the similarity score, the more precise it is and the less coverage it offers. The lower the similarity score, the more coverage it has and the less precision it offers.
6. Based on your review, determine whether to increase or decrease the similarity score threshold value for your similarity solution

   The solution returns only similarity results that have a score that is higher than the threshold. If the current threshold is not providing useful results, adjust the threshold value.
7. Navigate to **Predictive Intelligence** > **Similarity** > **Solutions**.
8. In the ML Solutions list, locate your solution and click the Reference icon (ⓘ).
9. Click **Open Record**.
10. In the Solution Statistics tab, enter a new value in the **Similarity Score Threshold** field.
    For example, suppose that the current score is 80. Based on your similarity example review, you might want to enter a higher score of 90. This returns fewer results but ones that are more similar to the alert record under review. A lower similarity score has more coverage but less precision.
11. Right-click in the form header and click **Save**.
    Your solution uses the new similarity threshold value. In the example, the solution returns similar results that have a score higher than 90.

Activate Event Management solution version

The system activates the most recent version of the solution, but you can activate any previously trained Event Management solution version if it is more appropriate.

Before you begin

Because you might have created more than one solution, you can choose to activate the solution that is most appropriate.
Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solutions**.
2. In the ML Solutions list, click the Reference icon (ⓘ) for the trained solution that you want to activate.
3. Click **Open Record**.
   You can assign a title to the solution using the `evt_mgmt_similarity_solution_title` property. The default value is `ml_x_global_alert_similarity`. This property changes the value in the **Solution Name** field.
4. In the solution record, click **Activate**.
   The system activates this solution version and deactivates any other solution version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Self-health monitors for Event Management

Use the Event Management self-health monitors to track Event Management features and resolve issues.

Monitor Event Management to detect and resolve issues

Enable the self-health feature to help detect and troubleshoot Event Management health issues. Some configuration is required before the self-health feature can display meaningful data. Once configured and the related scheduled jobs are enabled, the Service Operations Workspace displays data that is automatically collected and calculated on a recurring schedule.

Note: Event Management Dashboard is not supported for new instances in the Paris release. However, instances upgraded from a release prior to Orlando that use Event Management Dashboard can continue to do so.

The self-health feature creates CIs in the CMDB to represent the components of the self-health service. Use the provided monitors to monitor and track the health of services and Event Management features.

Event Management self-health can:

- Trigger alerts when there is an issue with Event Management.
- Display the **ServiceNow Event Management** application services in Service Operations Workspace. Use Service Operations Workspace to view alerts by application service, technical service, and alert groups. You can double-click a tile to view general information about the service.

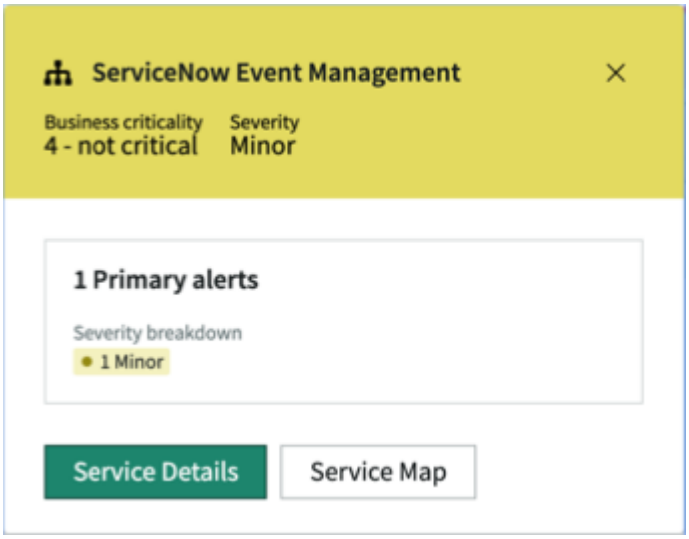

- Click **Service Details** to view details of the service.
- Click **Service Map** to view a map representation of the **ServiceNow Event Management** application services to see the relationships between CIs in the **ServiceNow Event Management** application services.

Self-health monitoring process workflow

After the *`Event Management – self-health monitoring`* property is enabled:

1. The **ServiceNow Event Management** application services is automatically created.
2. Jobs are automatically scheduled to run scripts that configure what must be monitored.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. The Monitoring States [em_monitor_state] table is updated after each event run that the configuration scripts trigger.
4. The alerts that are generated are monitored.

   The service displays the status of the monitored entity. For example, alerts are displayed in the Event Management dashboard, Service Operations Workspace or All Alerts list.

You can resolve the alerts according to the underlying issue by using **Quick Response** on the standard platform interface. For details, see Apply a quick response in an alert.

Jobs scheduled to process scripts

The jobs scheduled to run by default after the *Event Management — self-health monitoring* property is enabled are:

- The `Event Management - Insert Health Monitor` scheduled job which runs the scripts.
- The `Event Management - Update Health Monitor` scheduled job which runs the update scripts. This job runs once every minute.

Start or stop self-health monitoring

You can control the starting or stopping of the self-health monitor feature by configuring the self-health monitoring property. The first time that the self-health monitoring property is enabled, it automatically creates the ServiceNow Event Management application services.

Before you begin

Role required: evt_mgmt_admin

About this task

Configure the self-health monitoring property to monitor and track application services and Event Management.

Procedure

1. Navigate to **All** > **Event Management** > **Event Management Properties**.
2. Locate the *Enable Event Management self-health monitoring* (*evt_mgmt.self_health_active*) property.
   Select one of these values:

Enable Event Management self-health monitoring property

Field	Description
Enabled	The ServiceNow Event Management application services is created and set to monitor self-health. Events are received and alerts are generated.
Alerts Only	The ServiceNow Event Management application services and the CIs related to it are not created. Alerts are created but are not bound to a CI.
Disabled	The ServiceNow Event Management application services is not operational. Events are not received, and alerts are not generated.

3. Click **Save**.

Result

The ServiceNow Event Management application services is set according to the value selected for the *evt_mgmt.self_health_active* property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure the Event Management - Insert Health Monitor scheduled job

Determine what the Event Management - Insert Health Monitor scheduled job is to monitor. After the job runs, you can view the **ServiceNow Event Management** application services.

Before you begin

Ensure that the *Enable Event Management — self-health monitoring* property is enabled on the Event Management Properties page (**Event Management** > **Event Management Properties**).

Role required: evt_mgmt_admin

Procedure

1. Navigate to **System Definition** > **Scheduled Jobs**.
2. Locate and select the **Event Management - Insert Health Monitor** job.

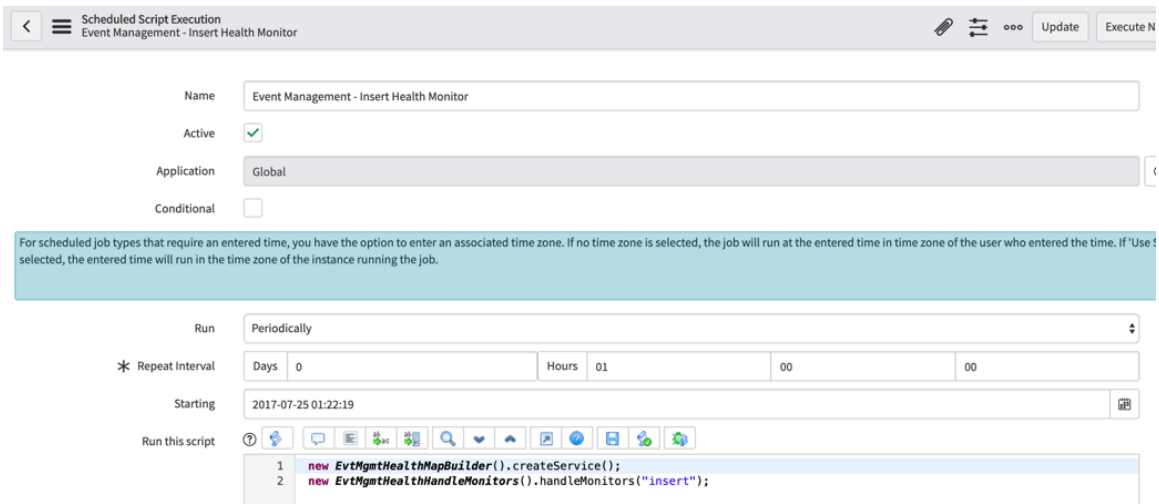

3. Modify the scripts indicated in the **Run this script** field to determine what the scripts are to monitor.

Monitor self-health with domain separation

Use domain separation to enable self-health to display Event Management health issues that are based on data, rules, and settings from the logged in user domain. The selected domain must not contain any child domains.

Before you begin

Ensure that the Domain Support — Domain Extension Installer (com.glide.domain.msp_extensions.installer) plugin is activated and that the MID Server you're connecting to is configured for Event Management.

Role required: evt_mgmt_admin

About this task

Self-health monitoring is domain aware. If the Domain Support — Domain Extension Installer (com.glide.domain.msp_extensions.installer) plugin has been activated, then self-health displays health issues based on data, rules, and settings from the logged in user domain. If rules and settings aren't defined for a child domain, then the settings of the parent are applied, recursively.

Note: To generate self-health monitoring events in the domain-separated instance, you must specify the domain of these generated events. This is done by specifying the `sys_id` of the required domain in the *evt_mgmt.domain_self_monitoring* property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When the Domain Support — Domain Extension Installer (com.glide.domain.msp_extensions.installer) plugin isn't activated, adding, or removing a connector or MID Server is immediately reflected in the map.

When the Domain Support — Domain Extension Installer (com.glide.domain.msp_extensions.installer) plugin is activated, updates are done periodically according to the interval specifications of the Event Management — Insert Health Monitor scheduled job.

Procedure

1. Navigate to **All** > **System Scheduler** > **Scheduled Jobs**.
2. In the **Schedule** list, locate and select the Event Management — Insert Health Monitor job.
3. In the Event Management — Insert Health Monitor job, under Related Links, select **Configure Job Definition**.
4. Set the required periodic update time in the **Repeat Interval** field by specifying the required values.
   The default interval is one hour.

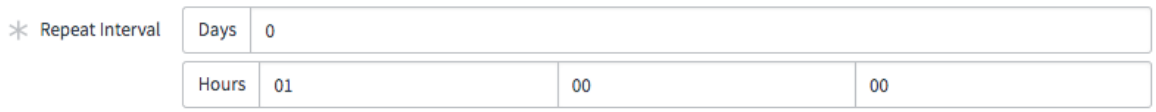

5. Select **Update**.

Result

Further configuration is required for CI creation, receipt of events, and alert generation. The `sys_id` of the domain must be specified in the **Value** field of the *`evt_mgmt.domain_self_monitoring`* property.

1. Navigate to **System Properties All Properties**.
2. Locate and select the *`evt_mgmt.domain_self_monitoring`* property.

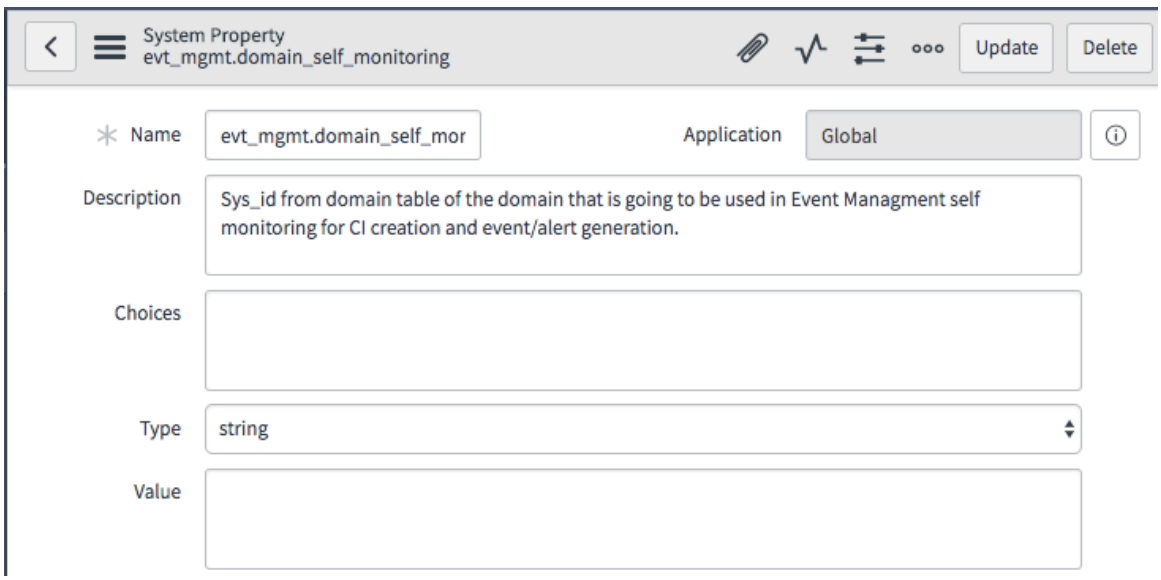

3. In the **Value** field, specify the `sys_id` of the domain.
4. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Domain separation setup and basic administration

Configure a self-health monitor

You can configure a self-health monitor to track Event Management components and see that they do not exceed the specified threshold.

Before you begin

Role required: evt_mgmt_admin

About this task

The Monitoring Configurations [em_monitor_conf] table lists the records that determine which Event Management components are being monitored. Each record contains the health monitor script that is relevant to what you are monitoring. If you intend to create a new health monitor script, you must configure a monitor to match the purpose of the custom script. You can configure the monitoring configuration records to track the items that are important to you. For example, you can modify the threshold aspect in the *`Connector's Idle state monitoring`* record, by specifying that the monitoring job frequency is 90 seconds instead of 120 seconds.

Note: Each of the monitoring configuration records can be disabled.

If a monitoring configuration specifies a threshold, the configuration value determines what alert severity to display. Threshold values specify the delay time values in units of minutes.

After the *`Event Management - self-health monitoring`* property is enabled, these jobs are automatically scheduled to run scripts that configure what must be monitored:

- The `Event Management - Insert Health Monitor` scheduled job runs the scripts. Determine what must be monitored by adding or modifying the script in this job, which runs once every hour.
- The `Event Management - Update Health Monitor` scheduled job runs the update scripts. This job runs once every minute.

Use the following procedure to configure whether a monitoring configuration can run and set the frequency of the monitoring job. You can specify the threshold values that, if exceeded, trigger an alert.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Self-Health Monitoring**.
2. From the Monitoring Configurations list, select a monitoring configuration record and modify the required field.
   On the form, fill in the fields.

Monitoring Configuration form

Field	Description
Active	Check box for enabling the monitoring configuration to be active.
Monitoring job frequency [seconds]	Frequency that the monitoring job is run.
Critical Threshold	Any point above the value is considered a critical consequence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Major Threshold	Any point above the value, but below the value of the **Critical Threshold**, is considered a major consequence.
Minor Threshold	If exceeded, the value is considered a minor consequence.
Warning Threshold	If the value is exceeded, a warning is issued.

3. Click **Update**.

Result
The values in the modified monitoring configuration record are compared to the actual values as they occur in events that are received. Alerts are created according to the configuration settings.

Example:
Example 1: Use the following procedure to modify the **Critical Threshold** field of the **Connector's Idle state monitoring** monitor from 30 seconds to 35 seconds. This field determines the number of seconds that pass before an alert is generated when the **Critical Threshold** idle state value of a connector is exceeded. In this example, the **Critical Threshold** field is modified from 30 seconds to 35 seconds.

1. Navigate to **Event Management** > **Settings** > **Self-Health configuration**.
2. From the Monitoring Configurations list, select **Connector's Idle state monitoring**.
3. In the **Critical Threshold** field, modify the value to `35`.
4. Click **Update**.

Result The threshold value of the **Critical Threshold** field is increased from 30 to `35`, thereby increasing the number of seconds that pass before an alert is generated for this threshold.

Example 2: This example describes the configuration of the severity threshold when using the **Connectors Status** monitor. The **Connector Status** monitor has a true (= value 1) or false (= value 0) configuration for the severity threshold. For this monitor, where a threshold has the value `1`, then that severity is displayed. Only one severity threshold can have the value`1`. This example assumes that the required severity displayed is `Minor`.

1. Navigate to **Event Management** > **Settings** > **Self-Health configuration**.
2. From the Monitoring Configurations list, select **Connectors Status**.
3. In the **Minor Threshold** field, specify the value `1`.
4. Click **Update**.

Result The monitor issues an alert if the threshold value of the **Minor Threshold** field is exceeded.

What to do next
Create a customized monitor. A template is provided to assist you to create script that can monitor the specified components. For more information, see Create a self-health monitor with a custom script.

Create a self-health monitor with a custom script

You can create a self-health monitor to use custom health monitor script to monitor specified Event Management components.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: evt_mgmt_admin

About this task
To augment the self-health monitors that are provided with the base system, you can create a self-health monitor to monitor the components that you specify. The health monitor scripts that are provided with the base system are listed in the following table.

Health monitor scripts provided with the base system

Health monitor script	Used by a monitor to:
Check delay in event processing	Monitor the duration of events [in minutes] that remain in Ready state and are not processed.
Connectors Monitor	Monitor: • Whether the time that a connector in idle state surpassed the threshold [in minutes] configured. • The active status of connectors.
Get Event Processing state	Monitor the time that events are in Ready state. Alerts are sent when this time exceeds the threshold.
MID Server Threshold Alerts	Monitor MID Server health. Events are sent when changes in MID Server state are detected.

You can create alert management rules to respond to and remediate the issues detected by self-health. Alerts are generated based on the specified severity and threshold. Navigate to **Event Management** > **Administration** > **Self-Health Monitoring** for the monitors that are provided with the base instance to monitor the issues detected by self-health, which are:

Monitor	Description
Connector's idle state monitoring	Monitor to verify whether any of the connectors were in idle state that surpassed the threshold [in minutes] that was configured. Health Monitor script selected: `Connectors Monitor`.
Connectors Status	Monitor to track the active status of the connectors. Select only one threshold to configure which severity to invoke using 1 to select the threshold field or 0 to ignore the threshold field. For more information, see Connector status example. Health Monitor script selected: `Connectors Monitor`.
Delay in event processing	Monitor to track the duration of events [in minutes] that remain in Ready state and were not processed. This duration period initiates an alert for the event processing CI with severity according to the configured thresholds [minutes].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Monitor	Description
	Health Monitor script selected: `Check delay in event processing`.
Event Processing job	Monitor the state of the event processing jobs. An alert is generated when the idle time of a job exceeds the configured threshold. Health Monitor script selected: `Get Event Processing state`.
MID Server MID Server Threshold Alerts	Monitor MID Server MID Server health. Events are sent when issues with MID Server availability and resource thresholds are detected. Health Monitor script selected: `MID Server Threshold Alerts`.

Note: To prevent a self-health monitor from detecting the connector status, at any time you can clear the **Active** check box of the monitor and then click **Update**.

Procedure

1. Navigate to **Event Management** > **Administration** > **Self-Health Monitoring**, and click **New**.
2. On the form, fill in the fields.

Monitoring Configuration form

Field	Description
Name	Descriptive and meaningful name for the monitor.
Description	Text narrative that describes the monitor.
Active	Check box to enable the monitor.
Monitoring job frequency [seconds]	Frequency that the monitoring job should run. Default: `120` seconds.
Last Run	Last run date and time. These values are automatically updated.
Script	Monitoring script. Click the search icon. The existing scripts are listed in the Health Monitor Scripts screen, or click **New** to create a custom health monitor script. For more information about creating a custom health monitor script, see step 3.
Note: Regarding all the following threshold configuration fields: All monitored values are compared against the thresholds. If a monitored value exceeds a threshold, an alert with the appropriate severity is issued.	
Critical Threshold	Any point above the value is considered a critical consequence.
Major Threshold	Any point above the value, but below the value of the Critical Threshold, is considered a major consequence.
Minor Threshold	If exceeded, the value is considered a minor consequence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Warning Threshold	If the value is exceeded, a warning is issued.
Additional information	Further information about the monitor.

3. On the Monitoring Configuration form, next to the **Script** field, click the search icon to open the Health Monitor Scripts form.
   To create a custom health monitor script, click **New**.
4. On the form, fill in the fields.

Health Monitor Script form

Field	Description
Active	Check box to enable the monitor.
Description	Text narrative that describes the custom script.
Name	Descriptive and meaningful name for the custom script.
Script	Custom script to run when called from a monitor. A script template is provided as a guide. You can use the Health monitor scripts provided with the base system as examples of how to author the script, and use the script editor tools. Ensure that the custom script includes: Run the monitor function, which is called according to the "monitoring job frequency" setting *`runUpdateHealthMonitorScripts(scriptId)`*. Add this function to run the monitoring logic. Specify the configuration function that is called every hour *`runInsertHealthMonitorScripts(scriptId)`*. Add the configuration logic and add a CI if needed. In the navigation pane search field, enter `em_monitor_state.list`. In the Monitoring States [em_monitor_state] table, specify the initialization of the monitor record in the required record. Use the helper functions defined in *`EvtMgmtHealthMonitorCommon`* and *`EvtMgmtEventAlertGenerator`* script includes: ○ EvtMgmtHealthMonitorCommon ○ healthMonitorCommon.shouldRun ○ healthMonitorCommon.calculateSeverity ○ EvtMgmtEventAlertGenerato ○ openAlert ○ closeAlert

You can use the Monitoring Configurations [em_monitor_conf] table for self-health monitoring configuration, the Health Monitor Scripts [em_monitor_scripts] table for scripts. Use the Monitoring States [em_monitor_state] table for displaying the results of the monitor.

These tables are initialized inside *`runInsertHealthMonitorScripts`*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information about script includes and APIs, see:Script includes that are installed with Event Management

5. Click **Submit**.

Result

To monitor the Event Management components that you specify, you can configure a self-health monitor that uses the custom health monitor script that you created.

View the Event Management self-health application services map

You can view Event Management application services maps to visualize configuration items (CIs) that comprise this service. The maps show the relations and connections between these CIs.

Before you begin

Role required: evt_mgmt_admin

The Event Management application services topology map only displays MID Servers that have passed validation and connectors that are active.

About this task

The Event Management components that are monitored by the self-health application services are listed in the following table:

Components monitored by the self-health application service

Component	What is monitored
Event connectors	For each connector: • Status: Success or Error • Delay: Elapsed time from 'last run time' (above threshold)
Event processing jobs	For each event processing job: • Number of actual jobs executed compared to the number of jobs configured to run. • Delay of any of the event processing jobs above the time set as threshold. • Duration of events in Ready state compared to the specified threshold. • Delay of event rule processing above the set threshold. If there is a delay, an alert is generated and the event rule is disabled.
MID Servers	MID Servers that have been validated are displayed. The MID Server status and the `Ecc_agent_issue` table are monitored. An event is triggered whenever this table is updated. The MID Server status is displayed, either **Up** or **Down**. For information about MID Server health, see MID Server heartbeat

Procedure

1. Navigate to **All** > **Event Management** > **Services** > **Application Services**.
2. Click the **ServiceNow Event Management** application services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. On the **Application Services** page, locate **ServiceNow Event Management** and click **View Service**.

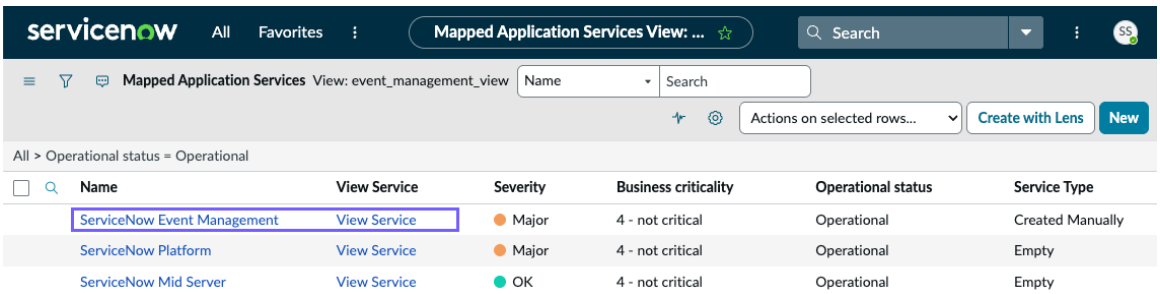

Displayed on the map are the related MID Servers and the Event Management object, as depicted in the following image; these MID Servers have an **UP** status. MID Servers are updated when they are created or deactivated.

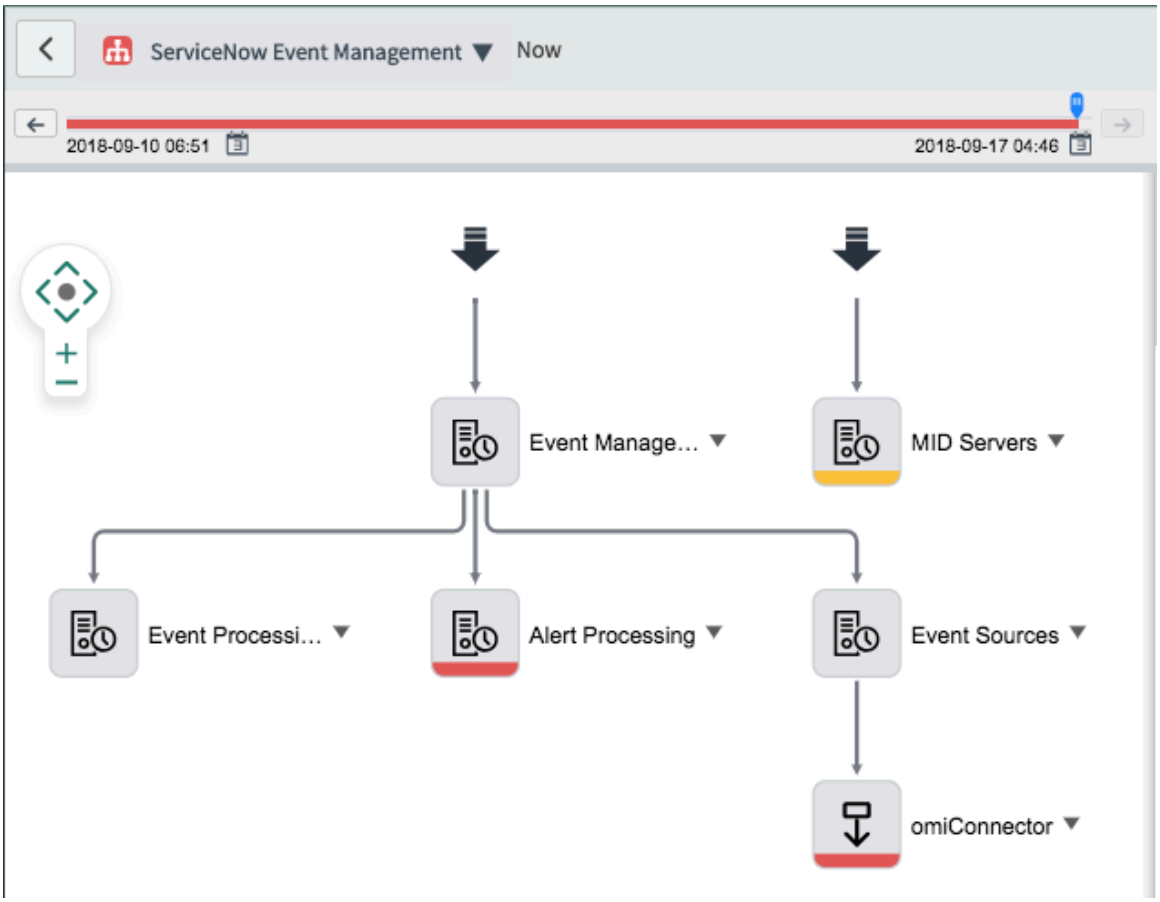

Under the Event Management object is Event Processing, Alert Processing, and Event Sources (for example, connectors). Only **Active** connectors are shown; when a connector is created, the connector is displayed only after it has been activated. The status of active connectors is updated and the status of non-active connectors is not displayed.

In the preceding image, the health of Event Management components is displayed. The omniConnector component has a red bar, indicating that a critical alert has occurred that affects this connector. In addition, this critical alert impacts the parent Alert Processing object, which as a result, has a red bar.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
You can select the configuration item that has a critical alert and change the impact from parent to application service.

Create maintenance rules

Use maintenance rules to mark CIs in maintenance status. When in maintenance status, these CIs are excluded from impact calculation.

Before you begin
Role required: evt_mgmt_admin

About this task
You can define rules to mark CIs that match the specified criteria as being in maintenance status. The marked CIs populate the Impact Maintenance CIs [em_impact_maint_ci] table.

Note: When running maintenance rules, the cmdb_ci status of matching CIs is not changed. However, matching CIs are flagged in the em_impact_maint_ci table by these rules and this status is considered for impact and alert calculations.

The field flagged by this rule as being in maintenance status is the CMDB status field.

Note: To activate a maintenance rule, select the **Active** field in the maintenance rule. To deactivate a maintenance rule, clear the **Active** field.

The maintenance rules provided with the base instance are:

Default maintenance rule	Description
CI in Change Window	Where the CI has an active change window, the matching CIs are marked as being in maintenance status. The rule runs a query against the change request [change_request] table to determine whether the rule is applied. All these conditions in the change_request table must be met: • **State** is one of these options: Scheduled, Implement, Work in Progress, or Open/New (state in (-2, -1, 1, 2)). • **Approval** is Approved (approval = 'approved'). • The change request window is active, that is, the current time is between **Planned start date** and **Planned end date**, or the current time is between **Actual start date** and **Actual end date**. • The change request record is not an on-hold record (on_hold='false'). **Note:** All these conditions must be present for the CI to be placed in maintenance status by this rule. For example, if the **State** of the change request approval status is `Change is waiting for`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Default maintenance rule	Description
	`approval`, then the change is not added to the em_impact_maint_ci table.
Maintenance status of CI	CIs whose CMDB status field is **In Maintenance** are flagged by this rule as being in maintenance status. By default, retired CIs are not included in CIs that are **In Maintenance**. To include retired CIs: **1.** Navigate to **All Properties**. **2.** Select **New**. **3.** Configure the *evt_mgmt.maintenance_rule_include_retired* property with **Value = true**. **Note:** Including thousands of retired CIs may decrease system performance.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Maintenance Rules**.

2. Click **New**.

3. Fill in the fields, as appropriate.

Column heading	Description
Name	The maintenance rule name.
Active	Select to activate the maintenance rule.
Advanced	Select to enable the optional script section to display.
Description	Information that describes this maintenance rule.
Flag CIs that run on this host	Select to flag all applications running on the recognized host as being in maintenance.
Table	Select the table that contains the CI that you require.
Filter	Specify how to select the data.
CI field name	Select the CI from the list. The list is populated according to your selection in the **Table** field. If a CMDB table, or a table derived from CMDB, was selected, specify `sys_id` for the **CI field name**. Otherwise, specify the required CI field that you want to use. See the examples.

4. Click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example:
Example of a maintenance rule using a CMDB table.
Assume that a company defines a CI as being in maintenance when the Operational Status of the CI is either **Repair in Progress** or **DR Standby**.

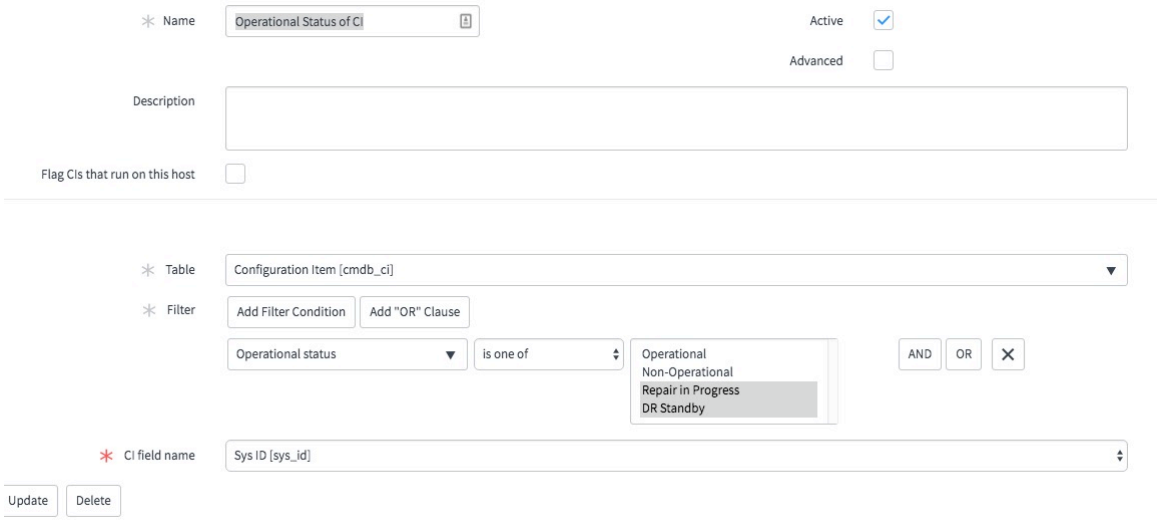

Example of a maintenance rule using a table other than CMDB.

Assume that a company uses Incident records to track maintenance. Any maintenance request is translated to an incident when it has a description that starts with “Performing maintenance on CI”. As long as the status of such an incident is open, this status indicates that the maintenance is in progress. For the **CI field name** field, specify a CI name from the table that was chosen, in this case **Incident**.

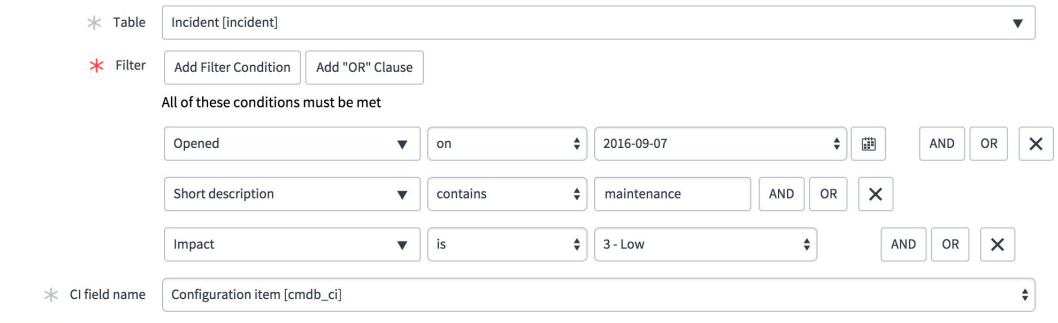

Example of a maintenance rule that uses the advanced script feature.

In the Maintenance Rule page, select the **Advanced** option. Customize the provided example script:

```
var gr = new GlideRecord('cmdb_ci');
gr.addQuery('name', 'your_name_here');
gr.query();
while (gr.next()){
 result.push(gr.sys_id +'');}
```

The advanced script must always return a text string that represents an array of CI IDs, for example, ['sys_id1','sys_id2','sys_id3']. Ensure that every possible execution path returns a value. If the script does not return a value, the maintenance calculator job fails with an exception.

You can use this script as an example to prepare your own customized script.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
(function findCisInMaint(){
var result = [];


//Your code goes here, this is an example of setting CIs in maintenance
var gr = new GlideRecord('cmdb_ci');
gr.addQuery('name', 'your_name_here');
gr.query();
while (gr.next()){
 result.push(gr.sys_id +''); //The array should be a string array, make sure to add strings
}


//Make sure to create a string representation of the array
return JSON.stringify(result);

})();
```

Related topics

- How alerts work with CIs in maintenance
- Place an alert into maintenance

Resolve an incident related to an alert

When you resolve an incident that is associated with an alert, the alert can also close according to the *evt_mgmt.incident_closes_alert* property.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

Procedure

1. Open an incident.
2. Click **Resolve Incident**.

   The corresponding alert is closed, resolved, or left alone as specified by the *evt_mgmt.incident_closes_alert* system property.

   The alert may not close immediately when the incident is closed. Depending on the value of the *evt_mgmt.alert_rule_close_delay* parameter, the alert may close 5 seconds after the last alert update. This value (in seconds) is defined in the parameter. To use the property, you need to create the property with the same name and assign the required value to it. For more information on how to create a property, see Add a system property.

Close an alert

Close an alert by an event or a user action. Closing an alert also closes any related incident that is not already resolved or closed.

Before you begin

Role required: evt_mgmt_admin or evt_mgmt_operator

About this task

An alert can be closed manually, automatically by a close or clear event, or when a related incident is resolved. Whether the alert is closed when the incident is closed is determined by the *evt_mgmt.incident_closes_alert* property. The default value is Yes (True).

When a Clear event is triggered for an open alert, the corresponding alert is set to the "Closed" state. Closing an alert also closes any related incident that is not already resolved or closed. This default behavior can be configured using the *evt_mgmt.alert_closes_incident* property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Event Management** > **All Alerts**.
2. Select the alert you want to close.
3. In the Alert Form, select **Close**.
   The alert closes without confirmation. If the alert has any resolved or closed incidents, a work note is added to the incident indicating that the related alert was closed.

   If an alert has an open incident that is not related to any other open alerts, the incident is either closed, resolved, or left unchanged. This behavior is based on the *evt_mgmt.alert_closes_incident* property. The available values are: Resolve Incident, Close Incident, and Do Nothing.

Reopen an alert

Additional events can cause reopening of alerts, or you can reopen an alert by changing its state. When an alert reopens, any associated incidents can also be updated or reopened according to the incident state and the *evt_mgmt.alert_reopens_incident* property.

Before you begin
Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task
When an alert is reopened, the related incident is processed as follows:

- If the incident is not *Resolved* or *Closed*, a work note is added to indicate that the related alert was reopened.
- If the incident is *Resolved* or *Closed*, the incident is reopened, a new incident is created, or nothing is done, depending on the *evt_mgmt.alert_reopens_incident* property value.
  - If the incident is reopened, work notes are added to the incident.
  - If a new incident is created, any matching alert management rule, alert action rule, and task template applies to the incident. If there is no matching alert action rule or template, fields from the existing incident are copied to a new incident.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Search by alert number.
3. In the **State** field, click **Reopen**.
4. Click **Update**.

Alert insight properties

Use these properties to configure alert insight.

The role required: evt_mgmt_admin

The following alert_insight properties are under sys_properties.

Property	Usage
Time Frame	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Usage
evt_mgmt.alert_insight_alert_history_min	Set the time frame (in minutes) to retrieve repeated and similar alert data. Default `43200` (30 days) **Note:** Alerts are retrieved regardless of their state (open / reopen / flapping / closed).
evt_mgmt.alert_insight_closed_alert_window	Set the time frame (in minutes) to retrieve alerts that were already closed. It is the time after the alert last updated date. Default: `4320` (3 days)
Similarity	
evt_mgmt.alert_insight_alert_same_as_filter	This property is a comma-separated string that defines which of the alert fields is used to consider alerts to be similar. Default: `source,type,resource,metric_name`
Related CIs	
evt_mgmt.alert_insight_related_cis_topology_levels The relationship types are: • CMDB based (metadata rules and suggested relations) • Within application services • Within Alert groups	For 'Within application service' relationship type, this property sets the depth or the maximum level of relationship of retrieved CIs. Default: 3
Score	
evt_mgmt.alert_insight_group_mapping	This property sets the score for within alert group relations. Default: 2
evt_mgmt.alert_insight_level_1_mapping	This property sets the score for level 1 relationship. Default: 3
evt_mgmt.alert_insight_level_2_mapping	This property sets the score for level 2 relationship. Default: 2
evt_mgmt.alert_insight_level_3_mapping	This property sets the score for level 3 relationship. Default: 1
Maximum related tasks	
evt_mgmt.alert_insight_max_tasks	Maximum related tasks to retrieve for alert insight. Default: 10

Metadata rules consideration

The parent-child relationship of CIs is considered. Dependent relationship rules consist of hosting and containment rules, each type modeling the data from a different perspective of the CI.

To manage dependent relationship rules:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To access rules at the class level, use the CI Class Manager. Navigate to **All** > **Configuration** > **CI Class Manager**.
- To access grouped rules, use the Metadata Editor. Navigate to **All** > **Configuration** > **Identification/Reconciliation** > **Metadata Editor**.

Containment rules represent configuration hierarchy of CIs, describing which CI contains which other CIs.

Hosting rules represent placement of CIs in a business definition, describing what CIs run on.

Modify the alert insight properties to configure the way alert information and analysis appears in the Alert Insight pane.

Related CIs configuration

The following properties control which CMDB relationships to consider for related CIs. The CMDB relationships include regular CMDB relation rules, metadata rules (containment rules and hosting rules), and suggested relations.

Property	Usage
evt_mgmt.related_cis_get_all_relation_types	Get all relation types, not including metadata rules. Default: `false`
evt_mgmt.related_cis_use_containment_rules	Use metadata containment rules. Default: `true`
evt_mgmt.related_cis_use_hosting_rules	Use metadata hosting rules. Default: `true`
evt_mgmt.related_cis_use_suggested_relations_rules	Use suggested relations rules. Default: `false`
evt_mgmt.related_cis_validate_relation_rules	This property controls whether to validate relation of CI according to metadata rules. Default: `true`

Score

Scores are configured per relation type or depth. Scores are accumulated. The higher the score, the more relevant is the related CI to the current CI.

Example: For a CI that was found at level `2` in the same application service of the current CI, the score is 2. The same CI is in the same alert group, so there is an extra score of `3`. The accumulated score is therefore `2+3 = 5`.

Affiliation type

The Affiliation Type column in the **Related Incidents**, **Related Change Requests**, and **Related Problems** tabs shows the type of relationship that the CI of the selected alert and the related CI have.

To see affiliation type relationships, navigate to **All** > **Configuration** > **Identification/ Reconciliation** > **Metadata Editor**.

The parent-child relationship between configuration items is considered.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Rotate event and alert table for cleanup

The growth of data tables impedes performance. Preserve instance performance by event table rotation and alert table cleanup for status and alert history retention.

Containing the growth of event tables

By default, the Event Management event [em_event] table uses table rotation. The table rotation schedule ensures that new data is written to one of the tables in the rotation group. According to the value of **Clean base rotation** setting, the base table is deleted and then a new round of writing the tables in a daily rotation commences. By default, this table deletion occurs after seven rotations.

To back up events to a custom table, see Event Management configuration preferences.

Containing the growth of alert tables

By default, a scheduled job cleans out the alert [em_alert] table. Table cleanup is done also for impact status and alert history table.

Related topics

- Purge impact status and alert history
- Alert table clean up
- Table rotation

Modify event table rotation

Table rotation is used by Event Management, by default, to contain the growth of event [em_event] tables within the rotation table group.

Before you begin

Role required: evt_mgmt_admin

About this task

Table rotation ensures that event tables grow to a reasonable size.
Each table rotation has a schedule, according to which new data is written to one of the tables in the rotation group. A table rotation group includes the original table plus several additional tables. After the rotation procedure has written the last table in a rotation, the rotation overwrites the first table in the rotation. By default, the table rotation for Event Management is to rotate the em_table to an offload table and the base table cleanup takes place after seven table rotations.
For information about backing up events to a custom table, see Event Management configuration preferences.

Procedure

1. Navigate to **All** > **System Definition** > **Table Rotations**.
2. Select the [em_event] table rotation group to modify.
3. Edit the required fields.
4. Select **Update**.

Related topics

- Table rotation
- Apply table rotation
- Rotate event and alert table for cleanup

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Purge impact status and alert history

Automatically cleans up outdated impact statuses and alert history from the database to free up space, improve system performance, and ensure only relevant data is retained.

Before you begin

Role required: evt_mgmt_admin

About this task

If records are older than 90 days, by default they are purged. If records are retained for a longer period, larger volumes of data are stored, and performance may be impacted. The data cleanup is done by two jobs: Event Management - Clean Alert History Table and Event Management - Clean Impact Status Table.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Event Management Properties**.
2. To shorten the retention period, use the highlighted properties.

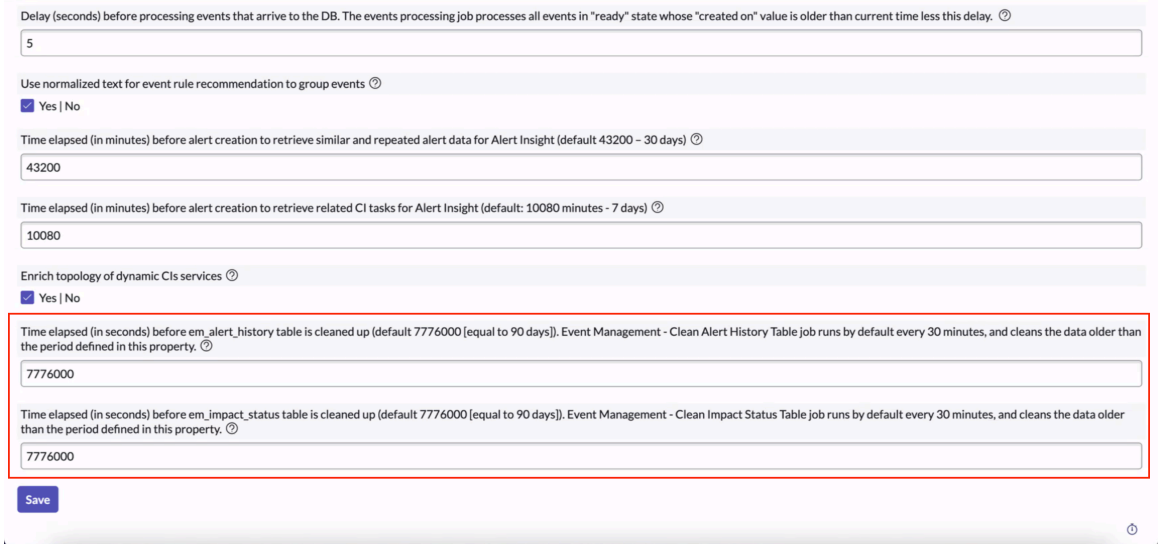

Related topics

- Rotate event and alert table for cleanup
- Modify event table rotation
- Clean alert history and impact status tables
- Alert table clean up

Clean alert history and impact status tables

Schedule jobs to mark and remove old alert records in the Alert History [em_alert_history] and Impact Status [em_impact_status] tables, to prevent the tables from becoming overloaded with data.

Before you begin

Role required: evt_mgmt_admin

About this task

Schedule the following jobs to remove old alert records:

- **Backfill**: Determines when alerts are to be marked for deletion.
- **Clean**: Deletes the alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The records in the [em_alert_history] and [em_impact_status] tables are deleted by default after three months. This is controlled by the properties *evt_mgmt.impact_calculation.cleanup_age_seconds.em_alert_history* and *evt_mgmt.impact_calculation.cleanup_age_seconds.em_impact_status*. The default value is 7,776,000 seconds (equivalent to 90 days).

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
   The **Scheduled Jobs** page appears.
2. To determine how often to mark alerts for deletion on the **Impact Status** table:

   a. Locate and select the **Event Management - Backfill Impact Status Table** job.
   The **Scheduled Script Execution** page appears.

Scheduled Script Execution page

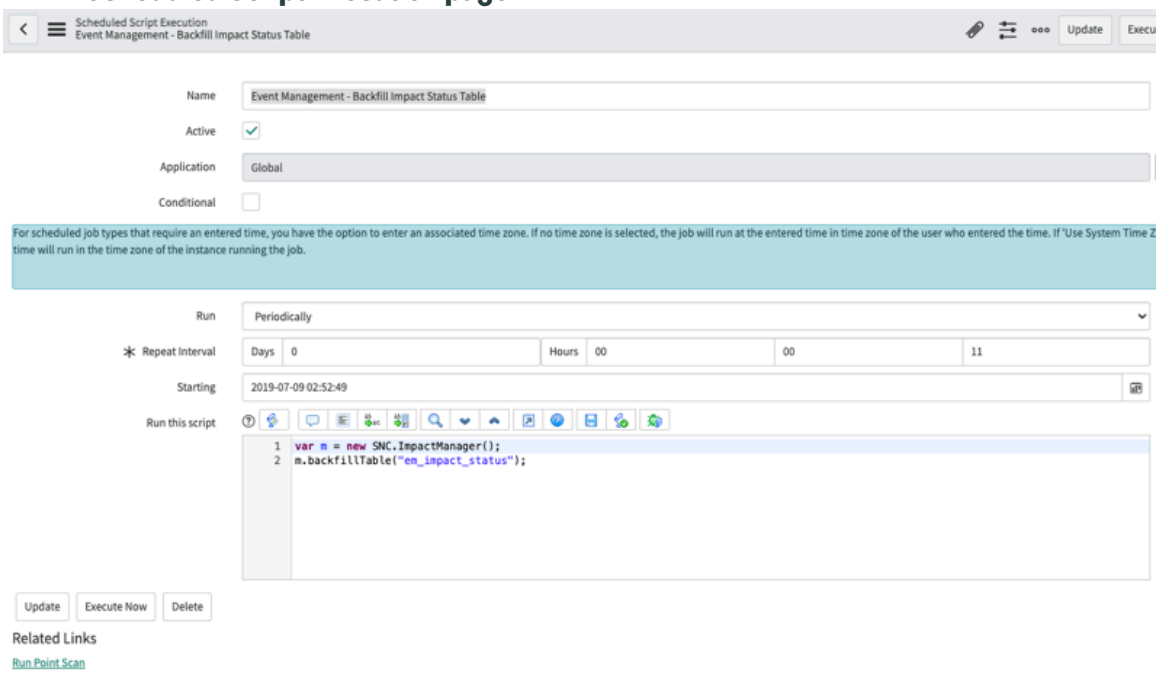

   b. In the **Repeat Interval** field, configure how often you want the job to run.
   To run the job immediately, select **Execute Now**.

   c. Select **Update**.

3. To remove old alerts from the **Impact Status** table:

   a. Locate and select the **Event Management - Clean Impact Status Table** job.
   The **Scheduled Script Execution** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

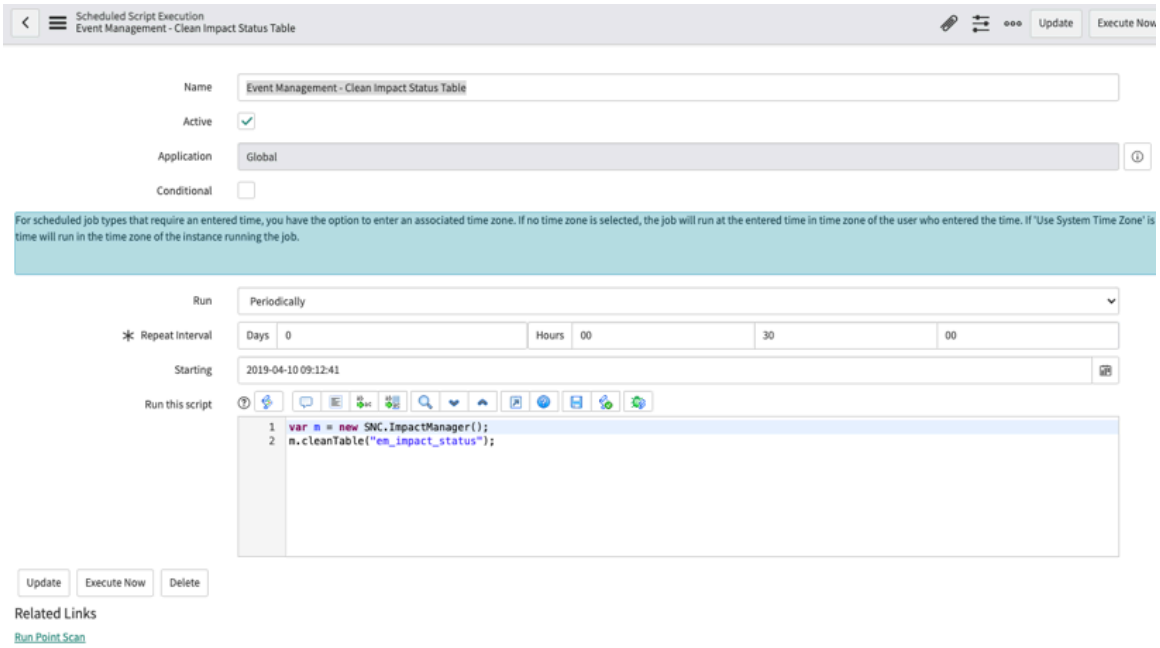

b. In the **Repeat Interval** field, configure how often you want the job to run.
To run the job immediately, select **Execute Now**.

c. Select **Update**.

4. Repeat steps 2 and 3 to perform the Backfill and Clean jobs on the **Alert History Table**.

Alert table clean up

The Scheduled Jobs feature runs a script to automatically close alerts in the Alerts [em_alert] table that meet specific conditions, helping reduce alert noise and keep the system clean and efficient.

Before you begin

Role required: evt_mgmt_admin

About this task

When the Event Management plugin is activated, by default a job is scheduled to run the `Event Management - auto close alerts` script to automatically close alerts in the alerts table that meet the specified conditions. The default retention policy for closed alerts is that the system deletes closed alerts if it is not updated for 90 days. For more information on the scheduled job, see Table cleaner.

Due to the high level of traffic to the table, you cannot audit the Alert table [em_alert] as a whole. Instead, audit selected fields of interest on the table.

Procedure

1. Set the table to `audit = true`.

2. Set `audit = true` on the relevant fields.
To avoid an excess of audit records, audit only the fields that are necessary for auditing.

Result

The `Event Management - auto close alerts` script runs automatically and repeatedly at the scheduled time and closes alerts that meet the specified conditions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Rotate event and alert table for cleanup

Modify event table rotation

Purge impact status and alert history

Alert grouping

Alert grouping is the process of organizing and consolidating related alerts into sets based on common characteristics or criteria. This helps in simplifying alert management by reducing noise, making it easier to prioritize, track, and address issues efficiently. Grouped alerts provide a clearer overview of related incidents, facilitating quicker root cause analysis and remediation.

How alert grouping works

Alert grouping enhances operational efficiency and optimizes alert management by categorizing alerts. Here's how it works:

1. Event consolidation: Event Management collects alerts from multiple sources, providing a centralized view of notifications and alerts. This consolidation enables teams to manage alerts more efficiently and recognize critical issues more easily.
2. Contextual Enrichment: Alerts are enriched with data from the Configuration Management Database (CMDB), which contains detailed information about applications and infrastructure components. This context allows teams to better understand the significance of alerts and prioritize responses based on their impact.
3. Intelligent Correlation: Event Management utilizes different types of machine learning algorithms to automatically group related alerts. This intelligent grouping reduces alert noise by combining similar alerts, allowing teams to concentrate on significant issues rather than being inundated with multiple notifications.

Benefits of alert grouping

- Creating automated alert groups by aggregating alerts based on predefined patterns.
- Correlating alerts using timestamps and CI identification to form automated alert groups.
- Forming CMDB based alert grouping by correlating alerts based on CI relationships in the CMDB.
- Correlating alerts based on text similarity of alerts using NLP (Natural Language Processing).

Related topics

Synchronizing alert response with automated alert grouping

Alert grouping and use cases

Alert grouping methods range from user-defined approaches, like Manual and Rule-based to advanced, fine-tunable algorithms, including Automatic, Mixed, Text-based, Log Analytics, and Network Traffic-based grouping.

Alert grouping types and use cases

Type	Description	Use case
Log Analytics Grouping	Alerts are grouped based on the analysis of log data. This involves correlating log entries to identify related alerts and	An online gaming company enhances server stability by implementing proactive log

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert grouping types and use cases (continued)

Type	Description	Use case
	issues. By leveraging log patterns and sequences, this method can detect complex, multi-step problems across the IT environment.	analytics. They monitor logs from game servers in real-time and use analysis tools to detect patterns of errors that occur before crashes. For instance, the analysis reveals that certain error patterns appear about 30 minutes prior to server crashes. By setting up automated alerts for these patterns, the company can initiate remediation actions, such as restarting services or reallocating resources, before a crash occurs. This proactive approach prevents disruptions, minimizes downtime, and improves the gaming experience by addressing issues before they impact players.
Rule-based Grouping	Alerts are grouped according to predefined rules and criteria set by users. These rules might include specific conditions, such as thresholds or event types. This method is effective for consistent and repeatable patterns but requires maintenance of the rules.	In a data center managing an e-commerce website, rule-based alert grouping helps handle high traffic during events like flash sales. Alerts about server issues, such as high CPU usage, are designated as parent alerts. These parent alerts are linked to child alerts that report related problems, like slow database queries. The rules ensure that server-related alerts are grouped with their symptoms, allowing the IT team to quickly identify and address server overload issues. This approach improves issue resolution efficiency and minimizes downtime.
Automated Grouping	Advanced algorithms automatically identify and group related alerts based on patterns and similarities in the alert data. This method	A large financial institution uses machine learning to manage alerts from numerous servers and applications. The system analyzes historical

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert grouping types and use cases (continued)

Type	Description	Use case
	leverages machine learning and AI to adapt to new and unknown issues, providing proactive alert management. Event Management groups alerts that are similar, but not necessarily identical, based on the proximity in time of the last event generation. Alerts with the same CI and the same pattern identifier are grouped together. Automatic alert grouping consists of the following components. • Alert Aggregation Learner (**Service Analytics Alert Aggregation Learner - Daily**): This offline job runs daily to process past alerts and perform statistical analysis to build alert patterns. For details, see Configure pattern based alert grouping. • Real-time Alert Aggregation job (**Service Analytics group alerts using RCA/ Alert Aggregation**): This job runs every minute to generate alert aggregation groups based on alert patterns, CMDB relationships, text similarity, user-defined alert clustering tags, and network traffic connection between processes.	alert data to recognize patterns, such as database server failures frequently being accompanied by client connection errors. It then automatically groups related alerts together. For instance, when a new database server failure alert is detected, it is grouped with previous connection error alerts. This automated grouping helps the IT and security teams quickly identify and address issues, improving response times and reducing downtime.
Mixed Grouping	Mixed Grouping method combines alerts using multiple grouping strategies, such as CMDB-based grouping and tag-based grouping, into a single, cohesive group. It leverages the strengths of each strategy to reduce alert noise, improve alert correlation, and highlight the true root cause of incidents.	Use case for CMDB-based grouping: A telecommunications company uses CMDB data to manage alerts related to their network infrastructure. Alerts related to a specific network router and its connected devices are grouped together based on their CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert grouping types and use cases (continued)

Type	Description	Use case
	• CMDB-based Grouping: Alerts are grouped based on Configuration Item (CI) relationships and dependencies from the Configuration Management Database (CMDB). This approach ensures that alerts related to specific infrastructure components or services are grouped together, providing context-aware alert management. • Tag Cluster Grouping: Alerts are categorized and grouped using tags or labels that represent common attributes, such as application, server type, or geographic location. This method allows for flexible and dynamic grouping based on evolving tagging strategies.	relationships, enabling the network team to see all related issues and address the root cause efficiently. Use case for tag cluster grouping: An organization without a CMDB manages a Linux server running various services. The IT team uses a **Node** field in each alert to identify the server, and they group all events related to services on the same server based on this node value. For example, they cluster alerts like Service A down and Service B high CPU usage together if they share the same node value. This approach helps the IT team address server-related issues more efficiently. By clustering alerts for the same node, application, or IP address, the team streamlines their response efforts and resolves issues more effectively, even without a CMDB.
Network-traffic based Grouping	Network-traffic-based alert grouping analyzes network connections between processes across hosts to identify related alerts. This method leverages service candidates detected through ML Service Mapping, ensuring that alerts related to network traffic issues are grouped together for better context and faster alert resolution.	A cloud-based e-commerce platform experiences transaction slowdowns, causing delays in payment processing. Traditional alerting generates separate alerts for API timeouts, database lags, and network issues, making it difficult to pinpoint the root cause. With Network-Traffic Based Grouping, alerts are automatically grouped based on process-to-process connections identified through ML Service Mapping. The system detects that payment gateway services,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert grouping types and use cases (continued)

Type	Description	Use case
		fraud detection, and order processing are part of the same service candidate. This reveals that an overloaded fraud detection process is causing transaction bottlenecks. By scaling up the service, the team quickly resolves the issue, minimizing downtime and improving customer experience.
Text-based Grouping	Alerts are grouped by analyzing the text content of alerts to identify similarities and related issues. Natural language processing (NLP) techniques are used to find commonalities in alert description, metric name, and ci class, making this method effective for unstructured data.	In an organization that uses Zoom rooms for virtual meetings, the IT team receives numerous alerts when the Zoom room server experiences an outage. Each alert might indicate a different Zoom room being down, such as Zoom room no 10 is down, Zoom room no 11 is down, and so on, with the only difference being the room number. For organizations with a CMDB, these alerts can be grouped using CMDB relations, as the system can correlate the alerts based on the server's impact on all associated Zoom rooms. However, for organizations without a CMDB, text-based grouping can be used. The system applies natural language processing to group alerts with similar descriptions, helping the IT team quickly identify that multiple Zoom rooms are affected by the same underlying server issue. This approach allows the IT team to efficiently address the root cause of the problem, reducing downtime and improving response times.
Manual Grouping	Users manually select and group related alerts based on their expertise and	A system administrator receives multiple alerts about different services failing on

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert grouping types and use cases (continued)

Type	Description	Use case
	understanding of the system. This approach allows for precise control but can be time-consuming and may miss automated correlations.	a single server. The admin manually groups these alerts, recognizing that they are all related to a single hardware failure on that server, and prioritizes fixing the hardware issue to restore all services.

Manual and rule-based alert grouping differs from algorithm-based grouping mainly in how the parent alert is chosen. In manual, rule-based, or log analytics grouping, one of the real alerts is designated as the parent alert. In Automatic, CMDB, Text-based, Tag Cluster modes, and Network Traffic, a virtual alert, representing the oldest and most severe alert in the group, is created as the parent alert.

By default, the main alert of the group is the oldest and the alert with the highest severity. There is a script include, GroupingProduceGroupMainAlertHook, using which we can change the main alert selection logic. The method that you should customize is getMainAlert.

Note: In domain-separated environments, alert groups are created only for alerts within the same domain.

For information on scheduled jobs and parameters, refer to Scheduled jobs and parameters for alert grouping. For detailed information on different grouping types, see Alert grouping types and creation methods.

Scheduled jobs and parameters for alert grouping

Automate alert organization by configuring jobs to group alerts based on predefined criteria and parameters.

To group alerts in Automated, CMDB, Text-based, Tag Cluster, and Network Traffic Correlation groups, the scheduled job named **Service Analytics group alerts using RCA/Alert Aggregation** is typically run once per minute. This job handles the grouping of alerts based on the specified method. Additionally, you can run multiple scheduled jobs in parallel to manage alert grouping more efficiently. For further details, see Run multiple scheduled jobs for alert grouping.

To define which alerts are grouped, the following parameters are used:

- *sa_analytics.aggregation_enabled*: This parameter enables alert grouping created by the schedule job. Set the property *Enable alert aggregation for Automated, CMDB, and Text-Based groups* to `true` to activate this feature.

  Note: This property also applies to Tag Cluster and Network Traffic Correlation grouping.

- *sa_analytics.agg.query_dynamic_window*: By default, this is set to 10 minutes (600 seconds). It defines the maximum time difference allowed between the last event generation times of two alerts that can be grouped together.
- *sa_analytics.agg.query_max_group_lifetime*: This parameter specifies the maximum time period from the generation of the first alert to the last alert in a group, with a default of 30 minutes (1800 seconds). If events arrive with a delay exceeding this period, the *sa_analytics.agg.group_expiration_time* parameter can be used to extend the grouping time beyond 30 minutes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Some parameters, such as **sa_analytics.agg.query_dynamic_window**, **sa_analytics.agg.query_max_group_lifetime**, and **sa_analytics.agg.group_expiration_time**, are not provided out of the box. To use these properties, you need to create properties with the same names and assign the required values to them. For more information on how to create a property, see Add a system property.

Example: How alerts are grouped

For tag-cluster grouping, alerts are added to a group based on the timeframe parameter defined in the alert tag clustering settings. For automated, CMDB, and text-based, network traffic grouping, alerts are aggregated as follows.

Consider the following alerts with the same CI. (All of them may be added to the same CMDB group).

- Alert1: Initial event generation at 01:00:00 AM
- Alert2: Initial event generation at 01:11:00 AM
- Alert3: Initial event generation at 01:13:00 AM
- Alert4: Initial event generation at 01:16:00 AM
- Alert5: Initial event generation at 01:25:00 AM
- Alert6: Initial event generation at 01:34:00 AM
- Alert7: Initial event generation at 01:43:00 AM

Alert1 and Alert2 are not grouped due to the time gap exceeding 10 minutes. Alert2 and Alert3 create a group at 01:13:00 AM. The 10-minute dynamic window starts at 01:13:00 AM, with the following:

- Alert4 is added to the group at 01:16:00 AM, restarting the 10-minute window.
- Alert5 and Alert6 are added to the group as their event times are within the 10-minute window.
- Alert7 is not added to the group because it arrives 9 minutes after Alert6, exceeding the *sa_analytics.agg.query_max_group_lifetime* limit of 30-minute maximum group lifetime from the initial group creation (01:13:00 AM + 30 minutes = 01:43:00 AM).

Note: After an alert is created, correlation logic is applied only once. Changes to the alert after its creation are not re-evaluated for correlation. If correlation is not established initially, the alert may still be added to a group later—but only if a new incoming alert matches and triggers grouping logic.

Configure scheduled job-based alert grouping

Set up rules and parameters to group related alerts automatically, streamlining alert management and reducing alert noise.

Before you begin
Role required: evt_mgmt_admin

About this task
The scheduled job Service Analytics: Group Alerts Using RCA/Alert Aggregation helps in grouping alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Alert Correlation Properties**.
2. On the Alert Correlation Properties page, enable the relevant properties.
   - Enable CMDB correlation (*sa_analytics.agg.query_cmdb_correlation_enabled*).
   - Enable Network Traffic correlation (*sa_analytics.agg.query_network_traffic_correlation_enabled*).
   - Enable ML based Automation correlation (*sa_analytics.specific_patterns_enabled*).
   - Enable Text based correlation (*sa_analytics.text_based_group_enabled*).
   - Group by property, with comma-separated list of field names. Alerts can be grouped together when the specified field values match each other. The property can contain alert field names (such as assignment_group), CI field names (such as alert_cmdb_ci.location), or alert additional info field names (such as additional_info.state). It can also contain alert tags (such as t_data_center). Property: *sa_analytics.agg.group_alert_with_same_group_by_fields*.
   - Max number of relations between CIs in a topology which form a CMDB group (*sa_analytics.agg.query_cmdb_graph_walk_nodes*).
   - Use all CMDB relations for CMDB group correlation (*evt_mgmt.related_cis_get_all_relation_types*).
3. Select **Save**.

   To group alerts without a CI as Text-based or ML based groups, set *sa_analytics.enable_no_ci_grouping* to `true`. Ensure the Feature Identifier includes both the node and metric name. For details on configuring the feature identifier, see Manage pattern identifier attributes for alert grouping.

4. **Optional:** Set *evt_mgmt.alert_groups_reasoning.enable_worknotes* to `none` to prevent group reasoning work note display.

   (Optional) By default, a work note appears in the **Activities** tab when a group is created or updated, showing the grouping reason.

Related topics

Automated alert grouping

CMDB based alert grouping

Text-based alert grouping

Tag cluster alert grouping

Run multiple scheduled jobs for alert grouping

Run multiple scheduled jobs in parallel to group alerts. This helps prevent overwhelming the system during surges (alert storms).

Before you begin

Role required: admin

About this task

The current scheduled job operates with a single thread. During alert surges (alert storms), processing takes significantly longer. To address this impact, now multiple jobs can process

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

alerts within the scheduled job. This enhances scalability and improves overall performance during high-volume periods.

The number of jobs to be running in the alerts processing is defined in a new property *sa_analytics.agg.alert_grouping.num_of_jobs*. The job number is sent as a parameter to the relevant method.

Note: In this scenario, the procedure for increasing the scale-out from no scale-out (single job) to 2 jobs is shown.

You can efficiently manage and run multiple scheduled jobs in parallel to group alerts using specific properties. These properties allow you to tailor the grouping criteria to meet your needs, ensuring that alerts are processed based on assignment groups, domain separation, or group-by fields. The available properties are:

- *sa_analytics.agg.group_alert_with_same_assignment_group_only*: Groups alerts that share the same assignment group. By default, the value is set to false. If you want to set this property to true, create a property with the same name and set the value to true.
- *sa_analytics.agg.group_alert_with_same_domain_only*: Groups alerts that belong to the same domain. By default, the value is set to true.
- *sa_analytics.agg.group_alert_with_same_group_by_fields*: "Group by" property, with comma-separated list of field names that need to have matching values across alerts to allow alerts to be grouped together. The property can contain alert field names (such as assignment_group), CI field names (such as alert_cmdb_ci.location), or alert additional info field names (such as additional_info.state). It can also contain alert tags (such as t_data_center).

Note: To ensure efficient scale-out, at least one of these properties must be defined to establish some form of separation or grouping logic. This is necessary to ensure that all alerts belonging to the same group are processed correctly by specific jobs, creating a logical division of alerts.

Without setting at least one of these properties to true, the alerts will not be properly segregated. If you do not define any grouping logic, the scale-out will not work. This occurs when both the assignment group and domain separation are set to false. All alerts will be processed by a single scheduled job.

Procedure

1. Navigate to **All** > **System definition** > **Scheduled jobs**.
2. Locate the scheduled job: Service Analytics group alerts using RCA/Alert Aggregation.
3. For all existing scheduled jobs Service Analytics group alerts using RCA/Alert Aggregation, set the value in the **Active** column to **false**.

Deactivate scheduled job

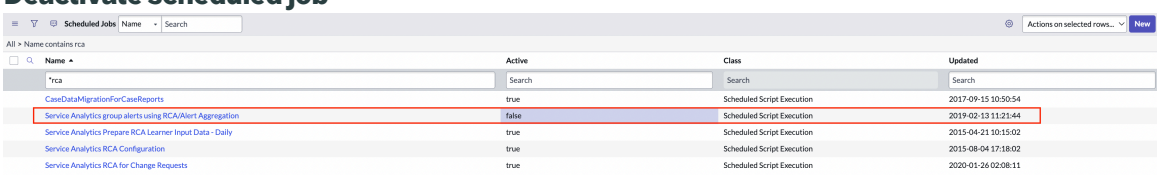

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Create a new property *sa_analytics.agg.alert_grouping.num_of_jobs*, or if it already exists, update the property to the desired number of multiple scheduled jobs for alert grouping such as 2 or 4.
5. Clone the scheduled job definition from the Scheduled jobs [sys_auto] table to the desired number of multiple scheduled jobs for alert grouping.
6. In the new jobs, call the queryS0 function and pass the job number as input.

   In the Run this script section, call queryS0 instead of the old query() function.

   Note: Job numbers must start at 0 and follow consecutively: 0, 1, 2, 3, 4...

7. Fix hashes in the [sa_analytics_status] and [sa_hash] table by selecting the earliest hash with the same prefix and cloning it.

   Add the hash name suffix "_groupingX", where X is the job number.

   In the case of scaling out from no scale-out (single job) to 2 jobs, and for the hash name `last_alert_process_time`, clone it to maintain the same value as in `last_alert_process_time`:

   - `last_alert_process_time_grouping0`
   - `last_alert_process_time_grouping1`

   In the case of scaling out from no scale-out (single job) to 4 jobs, and for the hash name `last_alert_process_time`, clone it to maintain the same value as in `last_alert_process_time`:

   - `last_alert_process_time_grouping0`
   - `last_alert_process_time_grouping1`
   - `last_alert_process_time_grouping2`
   - `last_alert_process_time_grouping3`

Hashes to be cloned

Hash name	Table
query_job_last_run	sa_hash
last_alert_process_time	sa_analytics_status
last_staging_table_update_time	sa_analytics_status
last_staging_table_truncate_time	sa_analytics_status
analytics_trigger_g1	sa_hash

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Activate the new jobs by selecting the **Active** check box for each job.

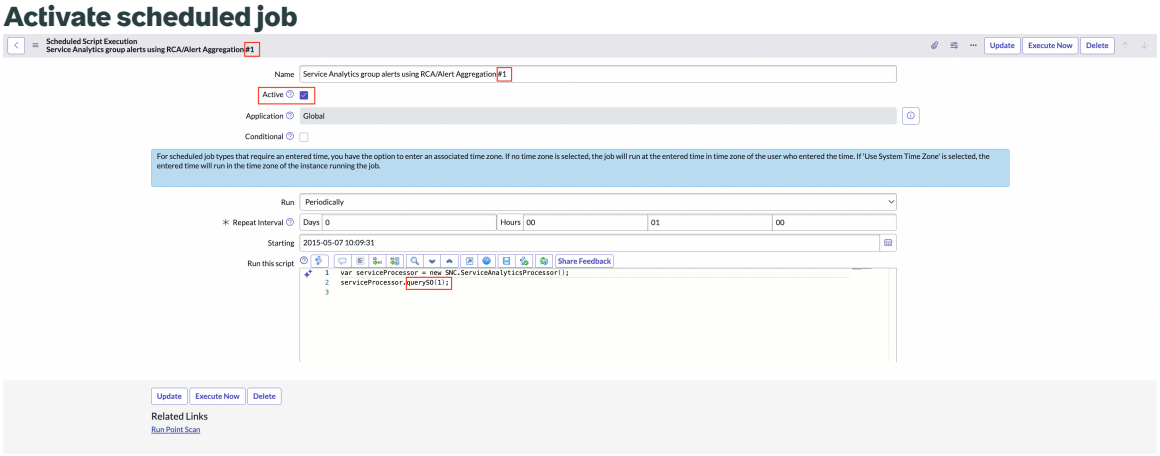

Result

The scheduled job Service Analytics group alerts using RCA/Alert Aggregation is set up to run multiple jobs simultaneously for alert grouping.

Configure alert correlation logic order

Improve alert management by enabling users to customize correlation logic order. This feature empowers you to fine-tune correlation methods to their specific needs, enhancing alert prioritization and response efficiency.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for the property *sa_analytics.agg.query.group_logic_order*.

   Default value is “MIXED,NETWORK_TRAFFIC,PATTERN,GENERALIZED_PATTERNS,TEXTBASE". This is a comma separated list of the grouping types in the order of their execution.

   Note: If one of the grouping types is not specified in the property, it needs to be added manually. Alert correlation rules are trigger-based and applied immediately when an alert is created or updated, before other grouping algorithms.

   - MIXED: Combination of 2 or more grouping criteria
   - NETWORK_TRAFFIC: Network traffic grouping
   - PATTERN: Automated CI-based patterns grouping
   - TEXTBASE: Text-based grouping
3. Use the property *sa_analytics.agg.query.group_logic_order* to define or modify the order of correlation methods based on your preferences.

Related topics

Alert grouping and use cases

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure filters for automatic alert groups

Filter alerts and alert groups to reduce alert noise. Only alerts that match the filter are included in the group of the selected group type (Automated, CMDB, or Text).

Before you begin

Role required: evt_mgmt_admin

About this task

If your filter specifies alerts whose configuration item class is not Linux, then only alerts with non-Linux configuration items are included in the group of the selected group type.

Note: After the filter is created, it will only apply to new alerts.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Automatic Group Filter**.
2. On the Automatic Group Filters page, either select an existing group filter or create a new filter.
3. If you are creating a new filter:

   a. Select **New**.

   b. Enter a name for the filter in the **Name** field.

   c. Select the alert group type:

   - Automated
   - Text
   - CMDB

   You can assign more than one group type with a single filter condition.

   Each group type can have only one active filter.

4. In the **Alert Filter** section, create conditions for the alert to meet to be included in the group.
5. Save the alert group filter.
   - For an existing alert group filter, select **Submit**.
   - For a new alert group filter, select **Update**.

Result

The filter condition appears on the SA Grouped Alert Filters page in the **Alert filter** column.

Apply alert group filters to aggregated groups

Reduce noise by locating only aggregated alert groups that match a configured filter. Aggregated groups are groups created for alerts with identical CIs and pattern identifiers.

Before you begin

Note: This procedure is relevant for advanced configuration only.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Automatic Group Filters**.
2. On the Automatic Group Filters page, add the **Aggregated** column.
   a. Select the Personalize List icon ( ) in the upper right corner.
   b. Move the **Aggregated** entry from the **Available** list to the **Selected** list.
   c. Select the up/down arrows to determine where the **Aggregated** entry is to appear on the form.
   d. Select **OK**.
3. On the Automatic Group Filters page, double-click the value in the **Aggregated** column.

   Aggregated groups apply only to automated alert groups. When the Aggregated option appears, it is automatically activated for an automated alert group.

   Note: If the value of the **Automated** field changes, the value of the **Aggregated** field also changes. You can change the value of the **Aggregated** field only if the value of the **Automated** field is true.

4. Select **true** for the specified filter to apply to groups with identical CIs and pattern identifiers or select **false** to ignore the filter and create aggregated groups for all alerts in the system.

Related topics

Manage pattern identifier attributes for alert grouping

Alert grouping types and creation methods

Explore different alert grouping types, understand their descriptions, and learn about their creation methods to enhance problem identification and streamline alert management.

Viewing and managing alert groups

Navigate to **Event Management** > **All Alerts** to view all alert groups. The icon in the **Group** column denotes the alert group type, while alerts not associated with any group have no entry in the **Group** column. Double-click the **Group** column to open the Grouped Alerts dialog, where you can view all alerts in the group and manually add or remove alerts.

Note: An alert can belong to only one alert group at a time.

Types of alert grouping

Alert grouping types

Type	Icon	Description	Creation method	Additional information
Log Analytics	L	Log Analytics groups are formed when the system identifies multiple related Log Analytics alerts, grouping them based on their significant connections.	Created as part of log analytics event processing.	Kinds of Health Log Analytics alerts

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert grouping types (continued)

Type	Icon	Description	Creation method	Additional information
Rule-based	R	Rule-based groups consist of related alerts that are organized based on compliance with alert correlation rules, which determine how alerts are grouped according to their relationships.	Created via business rule (Calculate correlation rule) on em_alert table when alert is created or updated.	• Rule-based alert grouping • Create an alert correlation rule
Automated	AI	Automated groups are formed by alert aggregation and include a virtual alert as the primary alert of the group. An Aggregated automated group is created when two or more alerts share the same CI type and metric name.	Created via scheduled job.	Automated alert grouping
CMDB	C	CMDB groups are formed based on CI relationships in the CMDB, specifically for CIs that are not included in rule-based or automated groups.	Created via scheduled job.	CMDB based alert grouping
Network traffic based	N	Network traffic alert groups are formed by analyzing network traffic connections between processes across hosts. This method leverages service candidates identified through ML Service Mapping to group alerts related to network traffic issues.	Created via scheduled job.	Network traffic based alert grouping
Text	T	Text groups are formed by grouping alerts based on similar text from frequently used words in following fields. • Description • Metric Name • CI Class	Created via scheduled job.	Text-based alert grouping
Tag Cluster	Tag	Tag Cluster groups are formed by grouping alerts according to user-defined tag-based alert clustering definitions.	Created via scheduled job.	Tag cluster alert grouping
Manual	M	Alerts grouped manually by users to organize related issues.	Created manually by the user.	Manual alert grouping

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information on scheduled jobs and parameters, refer to Scheduled jobs and parameters for alert grouping. For detailed information on configuring alert correlation logic order, see Configure alert correlation logic order.

Log analytics alert grouping

ServiceNow Health Log Analytics identifies patterns in log data and learns pattern behavior. When its artificial intelligence engine detects anomalous behavior, it sends an event to the ServiceNow Event Management application.

Log analytics alert groups are created by the Health Log Analytics application. For more information, see Types of Health Log Analytics alerts.

Rule-based alert grouping

Rule-based alert grouping is created by alert correlation rules. These rules allow you to manually classify alerts as primary or secondary and establish a relationship between them. Use alert correlation rules to group related alerts. The rule runs only for new alerts or alerts whose status changed from close/flapping to open/reopen.

Example: If rule-based alert grouping is applied, secondary alerts indicating that virtual machines or applications on the offline server are also down and are grouped under the primary alert, which is the root alert for the server that is offline. You can view rule-based alert grouping in Express List of Service Operations Workspace (ITOM). For more information, see Viewing links between alerts in rules-based alert groups.

Primary and secondary alerts

- Primary alert: Identifies the root cause of an alert group and represents its secondary alerts.
- Secondary alert: Related to the same issue, they are grouped under the primary alert. The purpose of secondary alerts is to determine which alerts to suppress, reducing alert noise and allowing you to focus on the primary alert.

How primary and secondary alerts interact

In rule-based alert grouping, one of the real alerts is designated as the primary alert. Primary and secondary alert filter conditions specify which alerts are classified as primary and which as secondary. The filter criteria are applied to the Alert [em_alert] table.

In this alert grouping, both primary and secondary alerts are visible but secondary alerts are grouped under the primary alert. The property **Keep the secondary alert open even when primary closed. for manual or rule based** (*evt_mgmt.rule_based_manual_closure*) controls whether secondary alerts remain open or are closed when the primary alert is closed.

If the property is not selected (i.e., set to No), after the primary alert is closed, the parent reference is removed from the secondary alerts (the grouping is dissolved), and the secondary alerts are closed and become stand-alone. If you select the checkbox for the property (i.e., set it to Yes) and the primary alert is closed, the parent reference is removed from the secondary alerts (the grouping is dissolved), causing them to become stand-alone open alerts.

Alert hierarchy

Only one level of secondary alerts is permitted. In situations where a secondary alert has its own secondary alert, the Event Management application flattens the hierarchy to preserve only two levels.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, assume alert A is the primary alert and alert B is the secondary alert. If alert C becomes a secondary alert for alert B, the application flattens the hierarchy so that A remains the primary and B and C become sibling secondary alerts, one level below A.

As another example, assume that there are three correlation rules that produce the following results:

- Rule 1 (with an Order value of 1): B becomes a primary alert for A.
- Rule 2 (with an Order value of 2): A becomes a primary alert for C and D.
- Rule 3 (with an Order value of 3): E becomes a primary alert for A.

When alerts B, C, D, and E are triggered, they all appear in the alert list separately because there are no correlations between them.

When alert A is triggered:

1. Rule 1 makes A a secondary alert under alert B.
2. Rule 2 makes both C and D secondary alerts under alert A, and the hierarchy is flattened so that A, C, and D become secondary alerts to the primary alert B.
3. Rule 3 is not applied because multiple parents are not allowed—A already has B as its primary.

Therefore, if all alerts are triggered, only B appears as the primary alert for A, D, and C, while E remains a standalone alert.

Create an alert correlation rule

Create an alert correlation rule to designate primary and secondary alerts. The primary alert is identified as the root cause of the alert group and the secondary alerts are grouped under the primary alert.

Before you begin

Role required: evt_mgmt_admin

About this task

Grouping automation helps you manage alerts more effectively by collecting similar alerts together. This makes it easier to see patterns, quickly identify issues, and respond efficiently. For more information on how to group alerts easily, see Create Group automation.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Alert Correlation Rules**.
2. Select **New**.
3. On the form, fill in the fields.
   For information on the fields, see Alert correlation rule form.
4. Select **Submit**.

Result

A rule-based alert group is created when a new alert is generated. An alert group is also created when the status of an existing alert changes from Closed or Flapping to Open or Reopened, provided the filter criteria are matched.

Automated alert grouping

Automated alert grouping is a process that uses historical data to automatically organize similar alerts into groups. These alerts could be system issues, like server errors or network outages. By grouping related alerts together, it helps teams quickly identify patterns, manage recurring problems, and reduce the noise from too many individual alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Imagine you're monitoring a city's traffic system. You get a lot of alerts—like reports of accidents, traffic jams, and road closures. Automated alert grouping works like a smart assistant that organizes these alerts based on patterns, so you can see related issues together and respond more efficiently. These automated alert groups are displayed in the Express List within the Service Operations Workspace.

How do you enable this grouping

To enable machine learning-based automation for alert correlation, set the property **Enable ML based Automation correlation** (sa_analytics.specific_patterns_enabled) to true.

If the Domain Support - Domain Extensions Installer is activated, alert aggregation patterns are created based on the domain level defined in the sa_analytics.agg.learner_domain_level property. By default, this domain level is set to two, which corresponds to the second level in the domain hierarchy. For example, in a company, Level 1 might represent the company itself, while Level 2 could represent departments or teams within the company. Alerts are grouped based on this second level, like sorting them by department or team. For more details, Domain separation and Event Management.

How does it work

Automated alert grouping uses machine learning (ML) and historical data to identify patterns among alerts. It looks at specific characteristics, called pattern identifiers—such as the type of issue, the affected system, CI or metric that happened multiple times within a similar time period —to determine if alerts are related. The Alert Aggregation Learner uses algorithms to group alerts based on patterns. Specifically, it uses pattern-based algorithms and probabilistic methods to analyze incoming alerts and identify related ones.

Think of it like noticing that accidents often happen at a particular intersection at rush hour. The system groups similar alerts (like recurring traffic jams at the same spot) together based on certain identifiers (like location or type of problem). The system follows these key steps to group alerts effectively:

- Analyze Historical Data: The system studies past alerts to learn patterns and relationships.
- Apply Machine Learning: ML algorithms analyze historical alert data to identify patterns and relationships among alerts. It enables the system to learn from past incidents and improve its ability to group similar alerts together over time.
- Group Similar Alerts: Alerts with matching patterns are automatically grouped together.

Imagine you're managing a city's traffic system, and you receive multiple alerts:

- 8:00 AM: Accident on Main Street
- 8:05 AM: Traffic jam near Main Street
- 8:10 AM: Road closure on Main Street

Automated alert grouping works like a smart assistant by analyzing these alerts and recognizing a pattern. It groups them together because they all relate to Main Street, likely stemming from the same accident. This helps you see the bigger picture quickly and focus on resolving the root cause (the accident), rather than addressing each alert separately.

Benefits

- Find Recurring Issues: Quickly spot patterns (like a server consistently overheating).
- Save Time: Handle groups of related alerts instead of individual ones.
- Improve Response: Focus on fixing the root cause rather than dealing with scattered issues.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Understanding pattern identifiers

A pattern identifier is a set of criteria or attributes (such as alert type, affected system, etc.) used to group similar alerts. It helps to identify recurring issues, making it easier for teams to respond and address ongoing problems.

How pattern identifiers group alerts

Consider a network monitoring system that generates alerts for various issues, such as high CPU usage, memory leaks, or connection timeouts.

Pattern Identifier: Metric Name and CI

- Alert 1: High CPU usage on Server A at 10:00 AM
- Alert 2: High CPU usage on Server A at 10:05 AM
- Alert 3: Memory leak on Server B at 10:10 AM
- Alert 4: High CPU usage on Server A at 10:15 AM

In this case, the pattern identifier could be set to the Metric Name (e.g., high CPU usage) combined with the Configuration Item (CI) (e.g., Server A). Alerts 1, 2, and 4 would be grouped together because they share the same metric (high CPU usage) and the same CI (Server A), indicating a recurring issue that may need further investigation. Alert 3, however, would not be included in this group because it has a different metric (memory leak) and CI (Server B).

Note: The set of alert fields used for the pattern identifier is also referred to as Feature Identifier Attributes or simply attributes.

How to configure effective pattern identifiers

To configure effective pattern identifiers for alert grouping, follow these three key steps to ensure accurate and meaningful analysis of alerts.

Step	Action	Description
Create an event rule	Define an event rule. To know how to create an event rule, see Create or edit an event rule.	Set up an event rule to populate the relevant alert fields for the pattern identifier.
Manage pattern identifier	Add relevant alert field to the pattern identifier. To know how to add fields to the pattern identifier, see Manage pattern identifier attributes for alert grouping.	After adding the relevant alert fields, select **Deploy** to activate the pattern identifier.
Choose relevant identifiers	Select alert fields that clearly identify the problem. For example, if the issue is that a service is offline or there's no connection to the database, look for specific	• Avoid overly unique fields (e.g., date) that make pattern identification difficult. • Avoid overly common fields that result in too

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Step	Action	Description
	values in the alert that indicate this. Add these types of fields to the pattern identifier. By default, we provide the **Metric Name** field as a pattern identifier.	many alerts being grouped together, making patterns indistinguishable.

Alert grouping and Learned Patterns

Learn how alert patterns are discovered, grouped, and displayed in the system.

Concept	Description
Pattern discovery	When a set of alert fields matches, the alerts are grouped into a "Learned Pattern." For example, alerts with the same Priority Group and Resource are grouped into a pattern.
Pattern reporting	These patterns are displayed on the Learned Patterns report found under **Event Management** > **Administration** > **Learned Patterns**.

Managing Pattern Attributes and Time Frame

Learn the process of managing pattern identifier attributes, deployment of new sets, and how the system identifies issues.

Concept	Description
Active pattern identifier attributes	Only one set of attributes can be active at a time. **Note:** The new set replaces the current one after deployment.
Purpose and time frame	Pattern grouping identifies issues within the last 30 days, controlled by the **sa_analytics.agg.learner_period_days** property.
Issue identification	To identify an issue, the system utilizes a combination of Configuration Items (CIs) and Pattern Identifiers (sometimes referred to as Feature Identifiers). By default, a Pattern Identifier is defined as the Metric Name, but this can be modified. Two alerts are considered similar if they share the same CI and Pattern Identifier, although fields such as Source, Severity, Description, and others may differ.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Concept	Description
	For more information, see Manage pattern identifier attributes for alert grouping. **Note:** The Alert Aggregation Learner also identifies patterns of alerts within manual alert groups. In some cases, you can create patterns from alerts where the CIs share the same value in a specified field. For example, to build patterns from alerts with the same CI **Location** field, enter `location` in the *`sa_analytics.agg.learner_group_by_property`* property. For more information, Configure scheduled job-based alert grouping. When working with CI-based groups, ensure that the pattern identifier includes both the node and the metric name. For details on configuring the Feature Identifier, see Learned patterns report. **Note:** Alerts that lack a CI can still be grouped together as Text-based or CI-based alert groups, treating a node as a CI. To enable this functionality, set the sa_analytics.enable_no_ci_grouping property to true.

Configure pattern based alert grouping

Configure the Alert Aggregation Learner (Service Analytics Alert Aggregation Learner - Daily), an offline job that runs daily to process past alerts. It identifies patterns of related alerts using a combination of pattern-based and probabilistic techniques, enabling quicker detection and resolution of recurring issues.

Before you begin

Role required: evt_mgmt_admin

About this task

The Alert Aggregation Learner tracks manual additions and removals of alerts from automated alert groups. If you undo any previous alert additions or removals, the automatic process adjusts accordingly.

You can review user additions and removals from automated alert groups and undo any action to prevent it from being automatically repeated. The Alert Aggregation Learner also identifies patterns in manual alert groups, enabling automatic formation of new alert groups based on these patterns when new streams of alerts arrive.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Alert Corrlelation Properties**.
2. Enable the following properties.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Enable alert aggregation for Automated, CMDB, and Text-based groups (*sa_analytics.aggregation_enabled*).

  Note: When disabled, disables all other groups.

- Enable ML based Automation correlation (*sa_analytics.specific_patterns_enabled*).

3. **Optional:** To configure the time period for the Alert Aggregation Learner to process alerts, perform the following steps:

   a. Navigate to **All** > **System Properties** > **All Properties**.

   b. On the System Properties page, select the *sa_analytics.agg.learner_period_days* property.

      If the property doesn't exist, you need to define it.

   c. Set the property's **Value** to the number of days by which you want alert aggregation learner job to process.

      Note: Values larger than 30 days increase job processing time. For optimal performance, use values of 30 days or less.

Manage pattern identifier attributes for alert grouping

The Alert Aggregation Learner analyzes alerts and identifies patterns using a defined set of alert and configuration item (CI) attributes. By configuring these attributes as pattern identifiers, you can control which characteristics are used to group alerts. This customization creates meaningful alert groups, improving alert management and response times by reducing noise and enabling focus on critical issues.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Event Management** > **Administration** > **Manage Pattern Identifier**.
2. On the SA Alert Aggregation Pattern Attributes page, select **New** to create a new pattern identifier, or open an existing pattern identifier.
   The SA Alert Aggregation Pattern Attribute page opens.
3. Select the Unlock Feature Identifier Attributes icon ( ) and move attributes from the **Available** list to the **Selected** list.

   Note: Properties from CMDB CI via dot-walking are not populated.

   The **Configuration Item** attribute is automatically considered as an identifier and should not be added explicitly to the Feature Identifier Attributes section. The pattern is defined as a combination of the configuration item and your selected feature identifier attributes.

4. Selecting the Lock Feature Identifier Attributes icon ( ).
5. Select **Submit** for a new pattern identifier or **Update** for an existing pattern identifier.
6. Navigate back to the SA Alert Aggregation Pattern Attribute page by opening the pattern identifier.
7. Activate the pattern identifier attributes by selecting **Deploy**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

- Ensure a corresponding event rule exists: Verify that there is an event rule set up to assign or populate the newly defined pattern identifier attributes to the incoming alerts. Event rules define how attributes are assigned to alerts based on certain conditions.
- Run the Service Analytics Attribute Populator for Historical Alerts job: After the new event rule is in place, this job retroactively populates the pattern identifier attributes for existing alerts that were created before the new attributes were defined. This ensures that past alerts have the necessary attributes for proper grouping.

Learned patterns report

The Learned Patterns report helps assess the efficiency of alert aggregation and identify recurring alert patterns. It enables proactive issue resolution, enhancing overall system performance by providing insights into frequent alerts.

Navigating and sorting the report

To access the Learned Patterns report, navigate to **Event Management** > **Reporting** > **Learned Patterns**.

Sorting Options: You can sort patterns by Pattern Score, Frequency, and Size to prioritize which alerts to address first.

Metrics and data in the report

The report displays key metrics associated with learned patterns, including:

- Frequency: The occurrence count of pattern identifier attributes (CI/MetricName).
- Pattern Identifier Attributes: Describes how alert grouping combines relevant pattern attributes to create a learned pattern. These patterns can reappear across multiple alerts, often indicating the same underlying issue.
- Presentation Format: Metrics are shown in a table format, organized by learned patterns.

This report allows you to:

- Focus on Higher Frequency Alerts: Allocate resources to higher-frequency alerts to expedite the resolution of numerous alerts.
- Target Larger Patterns with Lower Frequency: Allocate resources to larger learned patterns with lower frequency for targeted analysis.

Pattern group details and identifier attributes

- Expand a pattern group to display details about all the identifier attributes that are associated with the pattern.
- Sort the entries within each pattern by Configuration Item and Feature Identifier.

Column	Description
Configuration Item	CI associated with the combined pattern identifier attributes.
Feature Identifier	Feature identifier used for the combined pattern identifier attributes.
Frequency	Number of times the combined pattern identifier attributes appear in the data set.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
Size	Number of occurrences of the combined pattern identifier attributes (like Configuration Item and Feature Identifier) within a learned pattern.
Score	Calculated score that combines the frequency and size of the learned pattern. The formula for this score is: `Score = Frequency × (Size + 1)`. This score helps prioritize patterns based on their recurrence and size, making it easier to focus on high-priority patterns.
CI Link	Link in the **Configuration Item** column to view details of the CI information.

Exclude patterns from learned patterns

Exclude CI-based or CI class-based alerts and patterns when you encounter alerts incorrectly added to a learned pattern by the Learned Patterns job. For example, a pattern might include an alert that occurred at the same time as other alerts but is not actually related to them. This maintains accuracy, ensuring better alert groupings and improved management efficiency.

Before you begin

Role required: evt_mgmt_admin

About this task

You select the incorrect alert in a pattern to exclude the entire pattern to which it belongs.

Note: When you exclude an incorrect alert, any other patterns containing that alert are also excluded.

Procedure

1. Navigate to **All** > **Event Management** > **Reporting** > **Learned Patterns**.

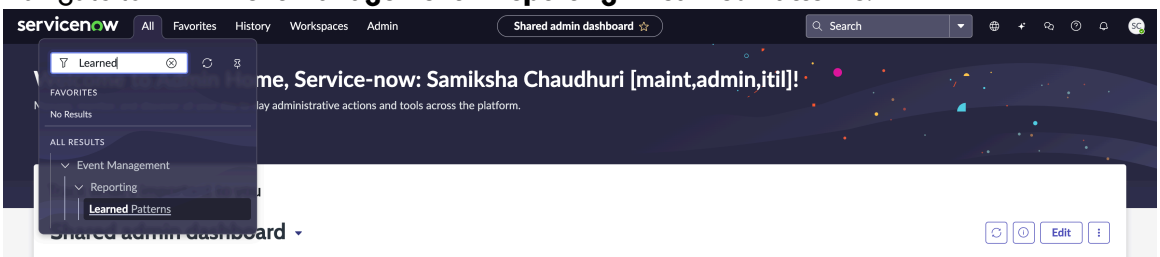

2. On the Learned Patterns page, expand the anomalous pattern.

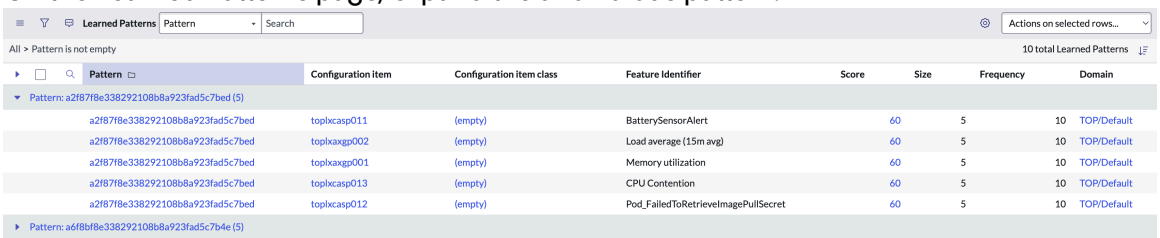

3. Select the pattern name to open it for exclusion.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. On the SA Alert Aggregation Learned Pattern page, select **Exclude**.

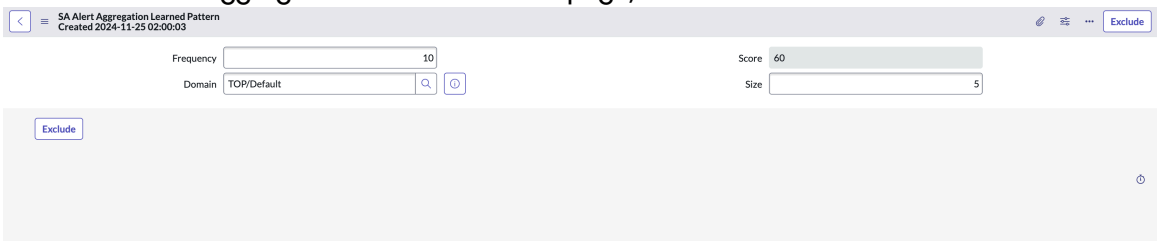

Result

The entire pattern is removed from the Learned Patterns report and listed on the Excluded Patterns page, located at **Event Management** > **Administration** > **Excluded patterns**.

If the pattern includes other alerts, you can restore it by reclaiming those alerts as a learned pattern. For further details, see Restore excluded patterns.

Restore excluded patterns

Restoring excluded patterns to the learned patterns report lets you reintegrate valuable insights lost due to incorrect alerts. This flexibility maintains accurate alert aggregation and enhances monitoring. For example, if you excluded a pattern due to an incorrect alert, you can restore it without that alert, ensuring relevant data remains accessible for analysis and decision-making.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Excluded patterns**.

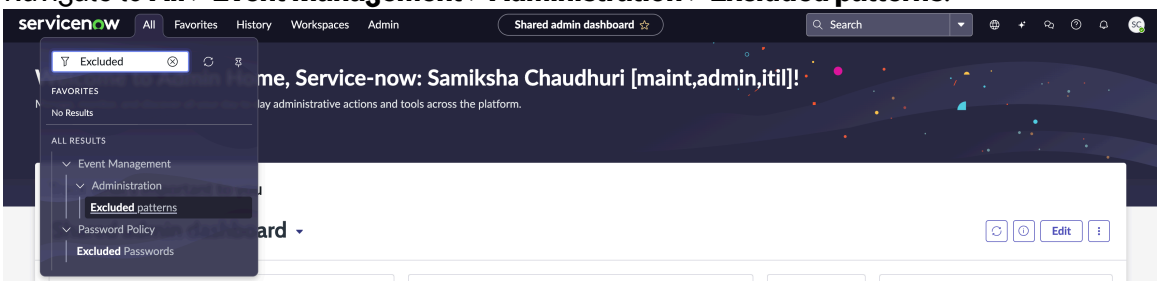

2. On the Excluded Patterns page, expand or select a pattern to view its alerts.

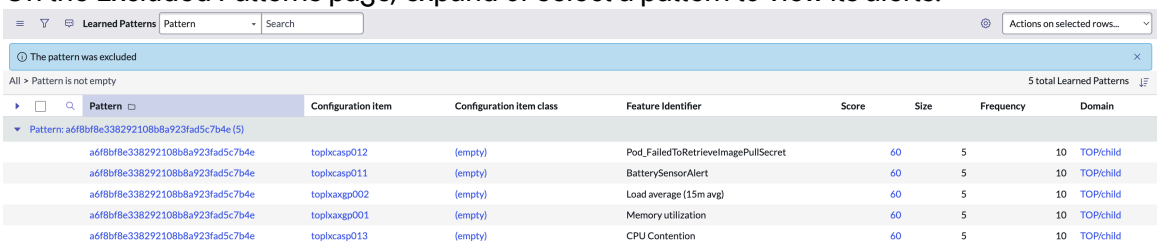

3. Select the patterns you want to restore.
4. In the **Actions on selected rows** drop-down menu, select **Delete**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

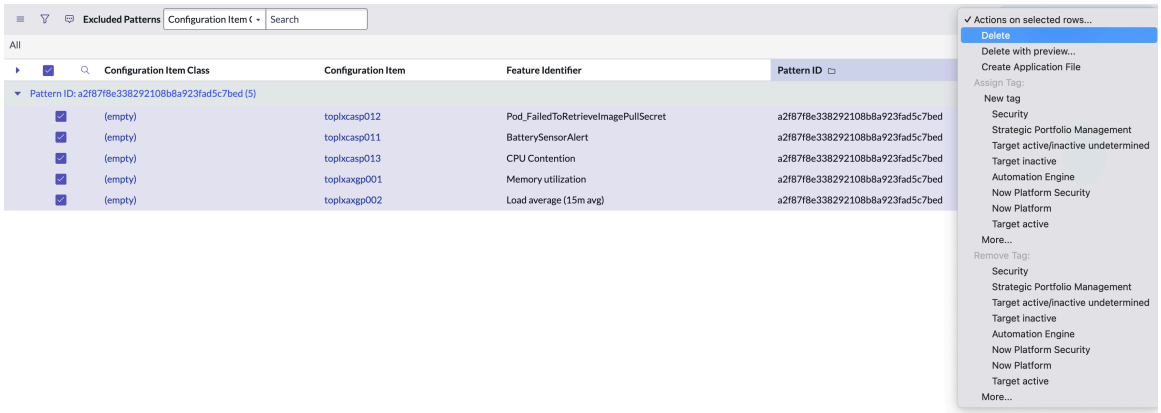

5. In the confirmation window, select **Delete**.
6. Run the Learned Patterns scheduled job.
   a. Select **System Definition** > **Scheduled Job**.
   b. On the Scheduled Jobs page, search for the Service Analytics Alert Aggregation Learner - Daily job.
   c. Open the Service Analytics Alert Aggregation Learner - Daily job.

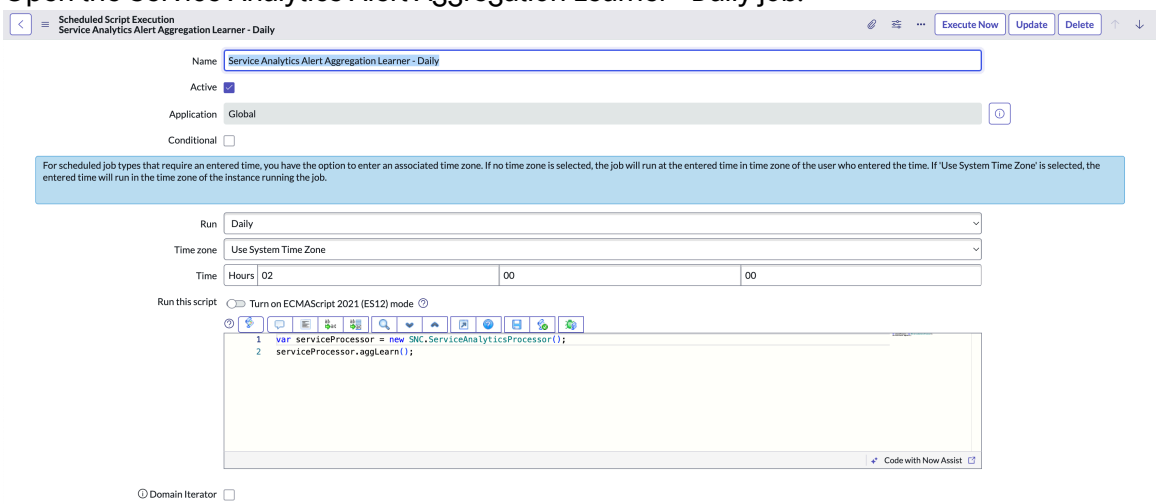

   d. Select **Execute Now**.

Result

The pattern appears with its restored alerts on the Learned Patterns report page. To access this page, navigate to **Event Management** > **Reporting** > **Learned Patterns**.

Network traffic based alert grouping

The Network traffic based alert grouping method groups alerts by analyzing network traffic connections between processes across hosts. It leverages service candidates identified by ML Service Mapping to group alerts related to network traffic issues. This ensures alerts from directly connected processes within the same service candidate are grouped together, offering a more contextual view of network incidents.

Service candidates are potential collections of processes within your IT environment that are identified based on their network connections and interactions. They represent different services or functions that your IT systems provide, even if they are not fully detailed in your

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

system's configuration database. For example, service candidates might group together all the processes involved in email delivery, even if the configuration details are incomplete.

ML Service Mapping uses machine learning to automatically discover and map out these service candidates. It identifies how different processes and components are connected and grouped together based on their network traffic and interactions. This helps in understanding and organizing the IT services and their components, making it easier to manage and troubleshoot issues. For example, ML Service Mapping might automatically identify and map the connections between email servers and mail clients based on their network interactions.

Note: Network traffic based alert grouping is enabled for new customers starting with the Brazil release. Existing customers need to enable the property Enable Network Traffic correlation (*sa_analytics.agg.query_network_traffic_correlation_enabled*) manually.

How it works

- Host identification: Alerts related to network issues are generated from various sources.
- Network context identification: The correlation process uses horizontal discovery results and ML Service Mapping to identify the most relevant service candidates and network connections.

  This process uses the results of the scheduled job **Event Management - Populate Service candidate process to process mapping - Daily**, which runs once a day. The job stores process-to-process connections for Host CIs in the format required by the alert grouping algorithm.

- Alert grouping: Alerts are grouped based on direct process-to-process connections within the context of the same service candidate. Grouping is updated in real-time as new alerts are received.

Benefits

- Improved accuracy: By leveraging network traffic connections and service candidates, this method provides high accuracy in grouping alerts, minimizing false positives.
- Enhanced coverage: It effectively groups alerts even in environments with low CMDB maturity, covering a broader range of alerts and issues.
- Streamlined resolution: IT teams can quickly identify and resolve related issues in bulk, reducing the volume of alerts to manage and improving operational efficiency.

Enable network traffic-based alert grouping

Activate network traffic-based alert grouping to automatically correlate and reduce alert noise by grouping related events based on network traffic patterns. This helps improve efficiency in alert response and streamlines alert management.

Before you begin
Ensure the Service Mapping Plus application (sn_sm_scoped_app) is installed.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Alert Correlation Properties**.
2. Ensure that the property Enable Network Traffic correlation (*sa_analytics.agg.query_network_traffic_correlation_enabled*) is selected.

Disable network traffic-based alert grouping

Disable network traffic-based alert grouping to prevent alerts from being grouped solely by network activity. This reduces noise during traffic spikes and helps critical issues stand out for quicker resolution.

Before you begin
Role required: evt_mgmt_admin

About this task

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Alert Correlation Properties**.
2. Clear that check box for the property Enable Network Traffic correlation (*sa_analytics.agg.query_network_traffic_correlation_enabled*).
3. Set the property *sa_analytics.enable_process_mapping_calculation* to false.

Follow up on work notes and review service candidate

Track and follow up on work notes while reviewing the service candidate to ensure all actions and updates are captured. This helps maintain clarity and continuity in the review process, reducing the chances of missed details.

Before you begin
Role required: evt_mgmt_admin and sm_admin

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.

   Note: Scroll down to an alert with **Network Traffic** as the group, and in the **Group** column, right-click and select **Show Matching**. This displays all alerts related to **Network Traffic**.

2. For the alert whose work notes you want to see, select the alert number.
3. On the Alert page, scroll down and select the **Activities** tab.
   You can view the work notes for the selected alert.
4. In the work notes, select the service candidate to view details.
   The Application Service Candidate page opens.
5. Select **Preview Map Graph** to view the service candidates mapping in the form of a graph.
   Service candidates mapping provides a visual representation of how different service candidates are connected, making it easier to understand relationships, identify potential issues, and analyze dependencies more effectively.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Mixed alert grouping

Mixed alert grouping combines multiple strategies—currently CMDB-based and tag-based grouping—to form a single, cohesive alert group. It reduces noise, bridges gaps when one method isn't enough, and delivers a clearer, end-to-end view of incidents. This approach helps operators identify root causes faster and take informed action, even in complex or incomplete data environments.

To understand how Mixed grouping works in practice, it's important to look at the individual strategies it currently supports. At this stage, Mixed Grouping comprises two core methods: CMDB based alert grouping, which leverages service relationships defined in the CMDB, and Tag cluster alert grouping, which uses shared metadata across alerts. Each plays a distinct role in correlating related alerts effectively.

If you're looking to group alerts, see Create Group automation.

CMDB based alert grouping

CMDB based alert grouping helps organizations manage alerts by organizing them according to their related configuration items (CIs) within the Configuration Management Database (CMDB). This method groups alerts based on CI relations in applications or infrastructure components. Teams can better understand the impact of issues, respond more effectively to alerts, and maintain service availability.

CMDB alert groups are created based on the current relationships between CIs in an organization. This means that when several alerts are triggered on different CIs—for example, if a server is experiencing issues — alerts are grouped based on CI relationships.

By default, all the existing connections between CIs are used to group alerts. However, given the complexity of IT systems, there are many types of relationships, and organizations can filter which ones to consider when aggregating alerts. For instance, you can filter relationship types by suggested relationships that appears in the CMDB Group CI Relations [em_suggested_relation_type] table. You can also filter using the hosting and containment rules (dependent relationship rules).

For example, if a server is hosting multiple applications, and there's a problem, all related alerts can be grouped together. This includes alerts for both the server and the applications it hosts. This helps the IT team quickly identify and address the root cause of the issue rather than dealing with each alert in isolation. Similarly, if different processes are running on a server, grouping alerts based on those processes allows teams to see how problems in one area might affect others. This leads to a more comprehensive understanding of the situation.

The CMDB alert groups are effectively displayed in the Express List within the Service Operations Workspace, which allows teams to easily monitor and manage related alerts. This feature enhances operational efficiency by providing a clear view of alerts that are interconnected based on the current relationships among configuration items (CIs).

Use cases for CMDB based alert grouping

Use cases for CMDB grouping enhance alert management by correlating alerts based on Configuration Item relationships, improving visibility, and facilitating more efficient troubleshooting.

Common CMDB grouping use cases

In the context of CMDB grouping, organizations face several challenges when managing alerts related to Configuration Items (CIs).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Common use cases

Use Case	Challenges	Solutions
Shared Configuration Item (CI) Scenario: An organization monitors a database server experiencing multiple issues, resulting in numerous alerts related to different applications using that database.	• Delayed response: Teams may respond to alerts in isolation, potentially overlooking related alerts, leading to delayed resolutions. • Inefficient resource allocation: Time and resources may be wasted on investigating separate alerts that are actually related. • Lack of context: Alerts related to the same CI can be scattered across different alert groups, making it difficult to see the full picture.	• Aggregate alerts related to the same CI into a single group for a unified view. • Facilitate faster alert resolution by addressing all related alerts together.
Hosting/Containment Relations Scenario: A physical server hosts several virtual machines (VMs), and an alert is generated for a hardware failure on the server. Multiple alerts also arise for the VMs due to their reliance on the server.	• Visibility into dependencies: Teams may lack visibility into how CIs are interconnected, leading to inefficient troubleshooting processes. • Complex alert resolution: Understanding which CIs are affected and how they relate can be complicated, resulting in longer resolution times. • Resource drain: Mismanagement of alerts can lead to duplicated efforts across teams, wasting time and resources.	• Group alerts using CMDB hosting/containment rules to aggregate alerts related to the physical server and its hosted VMs into a single alert group. • Provide a comprehensive view of all alerts tied to the physical server's failure. • Focus remediation efforts on the physical server while monitoring the VMs to ensure all aspects are addressed efficiently.
Applicative Relations Scenario: An enterprise application relies on multiple micro-services, and an issue arises with one of these services, generating alerts across several components, complicating diagnosis.	• Understanding application dependencies: Teams may find it challenging to trace how application components interact, making it difficult to pinpoint issues in complex systems. • Slow incident resolution: Without a clear understanding of applicative flow, alert	• Implement grouping based on applicative flow relations to aggregate alerts related to the affected microservice and its dependent components. • Utilize dependency maps to visualize how different services interact.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Common use cases (continued)

Use Case	Challenges	Solutions
	resolution can be slow and labor-intensive. • Inconsistent monitoring: Alerts related to applications may not be consistently monitored or prioritized, leading to potential oversights.	• Streamline the resolution process by addressing grouped alerts related to the application, improving response times.

CMDB alert grouping — properties and functionality

Learn about the key properties and functionality of CMDB alert grouping, which facilitate efficient alert organization based on relationships and enhance the overall effectiveness of alert management.

Properties for managing CMDB alert group creation

Use the properties listed in the following tables to control CMDB alert group creation. For more information, see Configure scheduled job-based alert grouping.

Property	Setting
Enable CMDB correlation *sa_analytics.agg.query_cmdb_correlation_enabled*	Enable this property to automatically use CI relationships for alert aggregation, correlating alerts and creating CMDB alert groups.
Max number of relations between CIs in a topology that form a CMDB group *sa_analytics.agg.query_cmdb_graph_walk_nodes*	Set the maximum number of enabled relations between CIs in a topology that form a CMDB group. For example, if you set this property to 3, a CI associated with an alert will automatically be grouped with similar CIs within three hops or less. **Note:** This property setting impacts all types of relations used for alert aggregation. Using a value greater than 4 may negatively affect the performance of the alert aggregation job.
Use all CMDB relations for CMDB group correlation *evt_mgmt.related_cis_get_all_relation_types*	Set to true if you want to use all relationship types for CMDB grouping. The default setting for new users is that the option is set to true, meaning all relationship types are used for grouping. Set the property to false to filter the relations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties for managing automatic creation of CMDB alert groups for advanced use cases

In some use cases, you may need to ignore or exclude certain CI relationships from automatic alert group creation. In this case, the *evt_mgmt.related_cis_get_all_relation_types* property must be set to false, and the properties in the below table must also be configured.

Note: If the *evt_mgmt.related_cis_get_all_relation_types* property is set to false, and the relationship isn't specified in the [em_suggested_relation_type] table, the relationship is not considered for alert aggregation.

Properties for managing automatic creation of CMDB alert groups for advanced use cases

Property	Setting
Enable Suggested Relations for CMDB Correlation *sa_analytics.agg.query_cmdb_suggested_relationship_enabled*	Keep this property enabled to use any suggested relationships defined in the CMDB Group CI Relations [em_suggested_relation_type] table when forming CMDB alert groups. To include a new relationship, add it to this table. This property is enabled by default.
Use CMDB relationship based on host-containment Metadata rules *evt_mgmt.related_cis_use_containment_rules*	Containment rules specify which configuration item classes can be contained within another item through a specific relationship when creating a service template. Keep this property enabled when forming CMDB alert groups to use any host-containment rules defined in the Metadata Editor. For more details, see CMDB dependent relationship rules. This property is enabled by default.
Enable filtering by dependent relationship rules *evt_mgmt.related_cis_use_hosting_rules*	Hosting rules define what classes of configuration items are hosted on other classes of configuration items. Keep this property enabled when forming CMDB alert groups to apply any hosting rules, which define which classes of configuration items are hosted on other classes. This property is hidden and enabled by default. **Note:** If you need to change the default value, add the property, as it does not exist in the instance by default. For more information on how to create a property, see Add a system property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties for managing automatic creation of CMDB alert groups for advanced use cases (continued)

Property	Setting
Ignore CMDB alert groups from forming due to CI applicative flow relations *sa_analytics.agg.ignore_cmdb_applicative_flow*	Set this property to true to prevent CMDB groups from forming based on applicative flow relations for CI relations in application services. This property is hidden and enabled by default. **Note:** If you need to change the default value, add the property, as it does not exist in the instance by default. For more information on how to create a property, see Add a system property.
sa_analytics.applicative_flow_relation	This property is hidden and enabled by default.

Event Management stores a map of related CIs in the Event Management CI Graph Reuse [em_ci_graph_reuse] table for each processed CI. You can adjust the storage period by modifying the *sa_analytics.agg.query_cmdb_graph_reuse_period* property, with the default set to one day (86,400 seconds). This process enhances the search performance of relations in the CMDB. To disable the storage process during advanced use-case configurations, set the property value to zero.

Note: If you need to change the default value, add the *sa_analytics.agg.query_cmdb_graph_reuse_period* property, as it does not exist in the instance by default. For more information on how to create a property, see Add a system property.

View the Dependency map for CMDB alerts

The Dependency map illustrates how and why alerts are grouped, simplifying troubleshooting and issue management. It reveals connections between CIs within a CMDB alert group, enhancing visibility into their relationships. Additionally, if a CI is not part of the group but connects to other alert CIs, it remains visible on the map, aiding alert resolution and proactive management.

Before you begin

Role required: evt_mgmt_user, evt_mgmt_operator

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Search for the CMDB group.
3. Open a CMDB group alert.
4. Scroll down and select the **Activities** tab.

   You may see a message similar to the following:

   ```
   Alert2405199 and Alert2405219 were grouped together because of
   CI relationships. Open Dependency Map for CMDB Group Alerts.
   Open Dependency Map for CMDB Group Alerts.
   ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

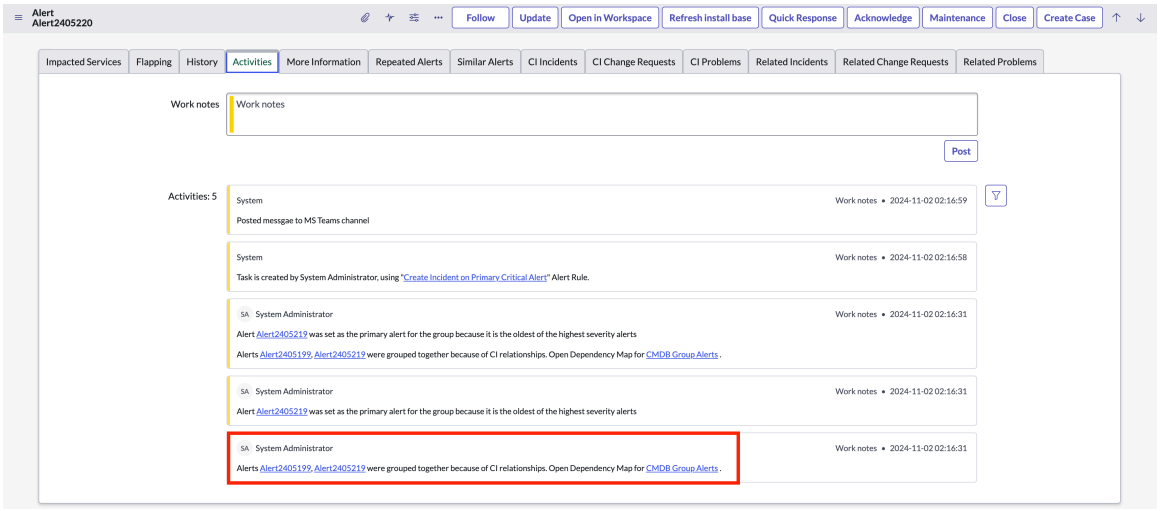

5. Select the link to see the dependency view of the CMDB alerts.

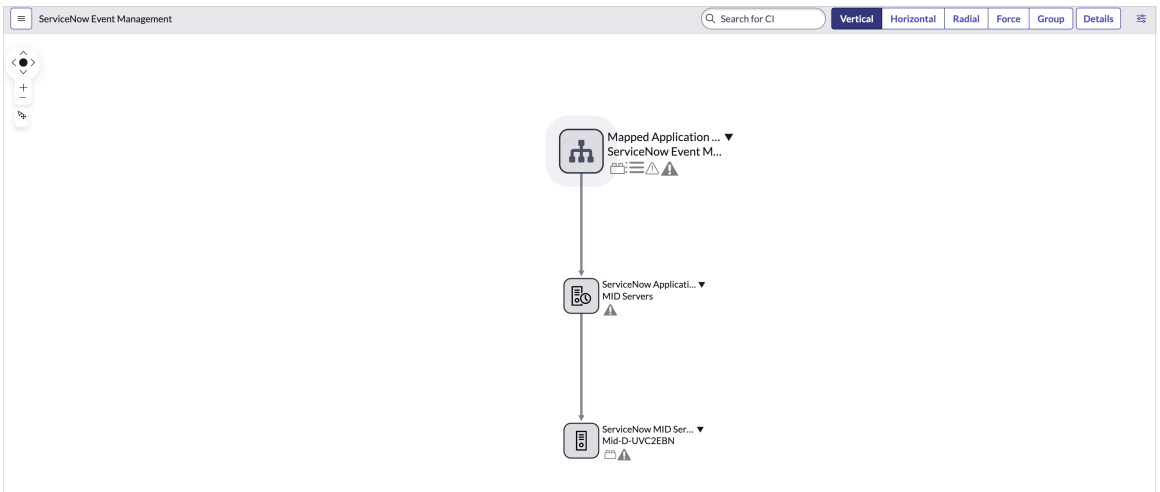

Add a CI relationship to CMDB Group CI Relations

Add a CMDB group CI relationship for a CI class to enable accurate alert group generation.

Before you begin

Role required: evt_mgmt_admin

About this task

Configure the *sa_analytics.agg.query_cmdb_suggested_relationship_enabled* property to further enhance alert grouping by using suggested relationships. For more information on properties, see CMDB alert grouping — properties and functionality.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **CMDB Group CI Relations**.

   a. Click **New**.

   b. Complete the form.

Suggested Relationship fields

Field	Description
Base class	The base class in the relationship serves as either the parent or the child, depending on the type of relationship established.
Relationship	The type of relationship that defines how two or more entities are connected.
Dependent class	The class that relies on another class in the relationship, serving as either the parent or the child depending on the relationship type.

2. Click **Submit**.

Example: CMDB Group CI relation

Base Class	Relationship	Dependent/Target Class
Oracle	Is Hosted On	Linux Server
Oracle	Is Hosted On	Solaris Server

Note: The same parent class and relationship can appear more than once.

Delete a CI relationship from CMDB Group CI Relations

Delete a CMDB group CI relationship to limit available relationships for CMDB alert grouping, helping to reduce noise and improve alert relevance.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **CMDB Group CI Relations**.
2. Select the checkbox next to the CI relationship you want to delete.
3. Select the Actions menu and then select **Delete**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

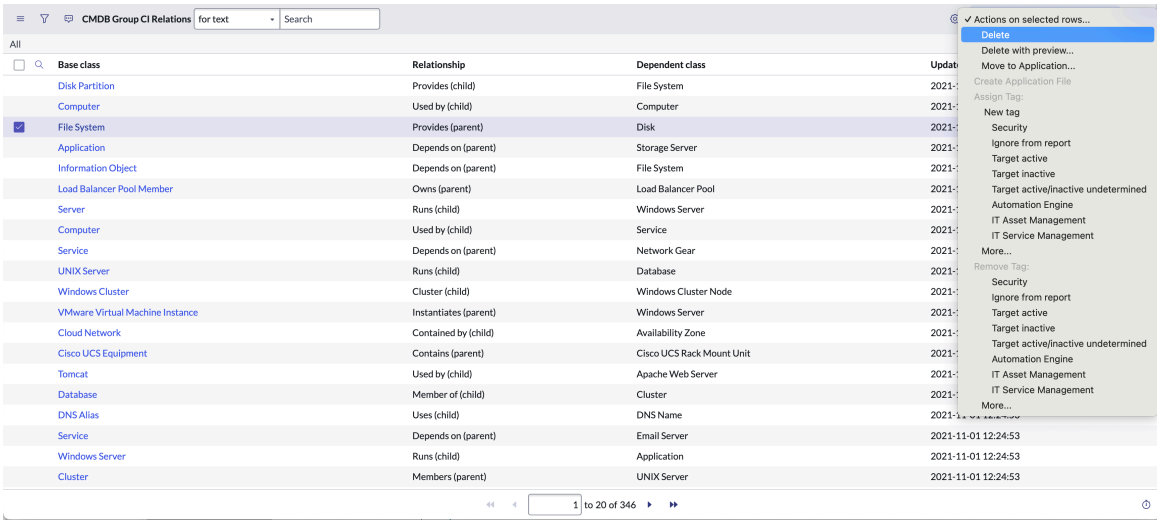

4. On the confirmation window, select **Delete**.
   The CI relationship is deleted.

Tag cluster alert grouping

Tag cluster alert grouping enables you to easily create groups of alerts. It is a non-code method of alert grouping that correlates alerts without having to use CMDB or model training. This simpler way of grouping similar alerts reduces the overall noise of a large quantity of alerts.

Note: To group alerts using tag-based grouping, create a grouping automation in Service Operations Workspace. For more information, see Create Group automation.

Tag cluster alert grouping is enabled immediately after the activation of the Tag-Based Alert Clustering Engine application, available in the ServiceNow Store. This grouping is applied according to the correlation logic order specified in the Configure alert correlation logic order. Alert grouping tags are attached to definitions on a many-to-many (M2M) basis. Multiple tags can be linked to a single definition, and a tag can be part of multiple definitions. Groups formed from tag cluster alert grouping definitions are classified as the Tag Cluster group type.

Tag cluster alert grouping supports domain separation, allowing different domains to have their own distinct alert grouping configurations and logic.

First, create alert grouping tags to define the criteria for grouping alerts. You can set the tags to require an exact match, an approximate ('fuzzy') match, or a character pattern match.

You can also use preconfigured tags to speed up alert clustering. These predefined tags are mapped from alerts and are based on information from sources such as the Alert field, Alert tags, or Alert additional info. If the required data is missing and the selected tag source is Alert CI or Alert CI key, the tag is populated using the Configuration Item (CI) value from the Configuration Management Database (CMDB). Predefined tags are easily identified by their description, which includes out of the box.

You can attach one or more tags to an alert clustering definition, which specifies the conditions for alert correlation. You can either create your own alert clustering definition or use a predefined one provided by the application. Predefined definitions come with associated tags.

Important: Make sure to activate predefined definitions before use. In new systems, several definitions are active by default. The remaining ones must be activated. For more information, see Activate a predefined alert clustering definition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Once one or more alert clustering tags are attached to a definition, the system collects alerts and checks if their tags match all the tag values specified in the definition. Alerts with matching or similar tag values are grouped together. New incoming alerts join an existing group if their tags match the tags in the definition used to create the group.

For tag-cluster grouping, alerts are added to a group based on the timeframe defined in the alert clustering settings. The time between the initial alert (virtual alert) and subsequent alerts is evaluated. If two new alerts are received, and their time difference falls within the defined timeframe, they are added to the group. The initial event's generation time is used to determine the relevance of the timeframe.

Related topics

- Create alert clustering tags
- Create an alert clustering definition
- Activate a predefined alert clustering definition
- Attach a predefined tag to an alert grouping definition
- Configure an event rule to customize alert content
- Alert tags

Create alert clustering tags

Create streamlined alert correlations with alert clustering tags by grouping alerts that share identical or similar tags based on your configured match method. This reduces noise, enhances incident prioritization, and improves operational efficiency, enabling faster issue resolution and response times.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: To group alerts using tag-based grouping, create a grouping automation in Service Operations Workspace. For more information, see Create Group automation.

Procedure

1. Navigate to **All** > **Event Management** > **Tag Based Alert Clustering Engine** > **Alert Clustering Tags**.
2. Click **New**.
3. Fill in the values on the Event Management tag based alert clustering tag form.
4. Select **Submit**.

Result

The tag appears in the Tag Based Alert Clustering Tags table.

Create an alert clustering definition

Define alert clustering conditions to trigger one or more alert clustering tags, which help create alert groups from fewer alerts. Creating alert groups from fewer alerts reduces noise, making it easier to identify critical incidents, prioritize responses, and manage issues effectively.

Before you begin

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

To create a tag definition, users need to use a filter to define which alerts the definition will apply to. They then define a grouping criteria by specifying which field values should be compared to create a group.

Note: To group alerts using tag-based grouping, create a grouping automation in Service Operations Workspace. For more information, see Create Group automation.

Procedure

1. Navigate to **All** > **Event Management** > **Tag Based Alert Clustering Engine** > **Alert Clustering Definitions**.
2. Select **New**.
3. Configure the fields on the Event Management tag based alert grouping definition form.
4. Select **Save**.
   The definition appears in the Tag Based Alert Clustering Definitions table.

Activate a predefined alert clustering definition

Activate the predefined alert clustering definitions provided with the Tag-Based Alert Clustering Engine application before use. Utilizing these preconfigured definitions minimizes setup time and ensures a more efficient configuration process, allowing for quicker implementation and streamlined alert management.

Before you begin

Important: If you are already using the Tag-Based Alert Clustering Engine application, all the predefined definitions are inactive and must be activated. In contrast, new customers will find several definitions active by default, allowing for immediate utilization of the alert clustering capabilities without additional setup.

Role required: evt_mgmt_admin

About this task

Note: To group alerts using tag-based grouping, create a grouping automation in Service Operations Workspace. For more information, see Create Group automation.

Procedure

1. Navigate to **All** > **#Event Management** > **Tag Based Alert Clustering Engine** > **Clustering Definitions**.
2. Select a predefined definition.

   Note: For a complete list of the available predefined definitions, see List of predefined tag-based alert grouping definitions.

3. On the Tag Based Alert Clustering Definition form, select the **Active** check box.
4. Select **Save**.

Result

In the Tag Based Alert Clustering Definitions table, the **Active** field for the activated definition displays the value **true**. You can now associate an alert clustering tag with this definition, enabling effective grouping of alerts based on the defined criteria.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attach a predefined tag to an alert grouping definition

Get started faster with alert clustering by attaching a predefined alert clustering tag to a tag-based alert clustering definition in Event Management. By associating a predefined tag, you ensure that alerts meeting the specified criteria are grouped effectively, facilitating quicker identification and response to related incidents.

Before you begin

Role required: evt_mgmt_admin

About this task

The Tag Based Alert Clustering Engine application comes with a set of predefined alert clustering tags. Attach a tag to a tag-based alert clustering definition to specify the conditions that must be met for alert correlation. This enables effective grouping of alerts based on shared characteristics, improving the accuracy and efficiency of incident response and management.

Note: To group alerts using tag-based grouping, create a grouping automation in Service Operations Workspace. For more information, see Create Group automation.

Procedure

1. Navigate to **All** > **#Event Management** > **Tag Based Alert Clustering Engine** > **Alert Clustering Definitions**.
2. Select a definition.
3. Under the Tag Based Alert Clustering Definitions Tags M2M section, select **Insert a new row**.
4. Use the search field to locate a predefined alert clustering tag, and then select the green check mark to confirm your selection.

   Note: Predefined tags are recognizable by their description, which includes "out of the box." For a complete list of the available predefined tags, see List of predefined alert grouping tags.

5. Select **Save**.

Result

The predefined alert clustering tag is now attached to the tag-based alert clustering definition. This association enables the effective grouping of alerts based on the specified criteria, facilitating improved alert management and response efficiency.

Text-based alert grouping

In text-based alert grouping, alerts are organized and correlated based on specific text patterns or keywords within the alert content. This approach dynamically groups alerts that share similar textual characteristics, such as error messages or event descriptions, allowing for more flexible and adaptive management of alerts.

The EM Alert Clustering Solution is a method used to correlate alerts based on similarities in specific fields and form clusters or groups. In ServiceNow Event Management, it creates clusters based on the Description, Metric name, and Configuration item.Class fields. This solution organizes alerts into text-based groups, and when a new alert arrives, the ML Predictor identifies the appropriate cluster, grouping alerts within the same cluster.

Note: The ML Predictor job is asynchronous and assigns real-time alerts to clusters, which may result in slight delays. This delay can cause text-based groups to be created several minutes later, as the alert grouping job runs once per minute. If prediction results are not available during a run, they are rechecked in the next grouping job.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For text-based logic to execute, you must have the Predictive Intelligence plugin (com.glide.platform_ml) installed and the EM Alert Clustering Solution definition activated.

There are specific settings or limits used to control the behavior of text-based alert grouping. These thresholds define the criteria for how alerts are grouped based on text patterns or attributes. The text-based thresholds are:

- Cluster quality threshold: The Cluster quality threshold (*sa_analytics.alert_grouping_tb_cluster_quality_threshold*) determines the minimum quality required for an alert cluster to be considered valid. This threshold ensures that only clusters with a minimum level of similarity and reliability are used. Clusters that meet this threshold are considered valid, improving the precision of the groupings and reducing noise from irrelevant or low-quality clusters. The range of the threshold is from 1 to 100 and the default value is 70.
- Alert rank threshold: The Alert rank threshold (*sa_analytics.alert_grouping_tb_alert_rank_threshold*) defines the minimum rank required for an alert to be included in a group. This threshold ensures that only alerts with a certain level of similarity are grouped together, filtering out lower-ranked alerts to maintain the quality of the alert group. The default value is 0.3, where smaller values indicate better similarity.

Note: To use these properties, you need to create properties with the same names and assign the required values to them. For more information on how to create a property, see Add a system property.

The EM Alert Clustering Solution definition is located in the [ml_capability_definition_clustering] table. To access it, navigate to **Predictive Intelligence** > **Clustering** > **Solution Definitions**.

To verify if the solution definition is active, see Verify text-based clustering solution. To disable the EM Alert Clustering Solution definition, set the property *sa_analytics.text_based_group_enabled* to `false`. Then clear the **Active** check box in the EM Alert Clustering Solution definition.

Example of text-based alert grouping

Scenario	Example
Network Connectivity Problems: There are widespread network connectivity issues affecting multiple departments.	Alerts from various network monitoring tools might report issues like Network segment down, High packet loss, or Connectivity issues in subnet. Text-based alert grouping uses the EM Alert Clustering Solution and ML Predictor to streamline alert management. The EM Alert Clustering Solution employs Natural Language Processing (NLP) algorithms to analyze and identify common text patterns in alerts. These patterns include issues such as Network segment down or High packet loss. It then clusters these alerts based on their text similarity, grouping related issues together. The ML Predictor further enhances this process by evaluating new alerts in real time and assigning them to the appropriate existing clusters based on their text patterns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scenario	Example
	This dynamic grouping provides a consolidated view of the connectivity problems, allowing network engineers to quickly diagnose and address the root cause of the issues more efficiently.

Verify text-based clustering solution

Event Management uses Natural Language Processing (NLP) algorithms to identify common text patterns in alerts and create alert groups through clustering. This approach not only organizes similar alerts efficiently but also reduces noise and helps focus on significant issues, enhancing overall alert management.

Before you begin

Ensure that the Predictive Intelligence plugin (com.glide.platform_ml) is installed.

Ensure the parameter *sa_analytics.text_based_group_enabled* is set to `true`.

Note: If the Predictive Intelligence plugin is not installed, the text-based logic will not execute, even if *sa_analytics.text_based_group_enabled* is set to true. This is because text-based grouping relies on clustering alerts based on the description, metric name, and CI class. Without the plugin, clustering cannot be performed, and as a result, text-based groups will not be created, regardless of the active status.

Ensure that there are at least 100 alerts in the Alerts table [em_alerts].

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
   The Clustering Definitions page opens.
2. Select EM ALert Clustering Solution.
3. Verify that the **Active** check box is selected.
4. Select **Update & Retrain**.
   For more information on ML solutions, see Reviewing your ML solution training jobs.

Manual alert grouping

Manual alert grouping involves organizing and categorizing alerts based on user-defined criteria and direct intervention.

You can manually select and group alerts to better manage and resolve incidents by tailoring the grouping based on specific needs, relevance, or context. This approach ensures that alerts are categorized effectively, enhancing overall incident management and response efficiency.

Create alert group manually

Manually create an alert group to organize and manage related alerts when not using scheduled jobs. This provides flexibility to group alerts on-demand for effective resolution.

Before you begin

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Select the alert number to open the alert.

   The alert will be used to create an alert group, designated as the primary alert, to which you can add secondary alerts.

   Note: A closed alert cannot be assigned as a parent alert.

3. Scroll down and in the **Secondary Alerts** tab, select **Add To Group**.

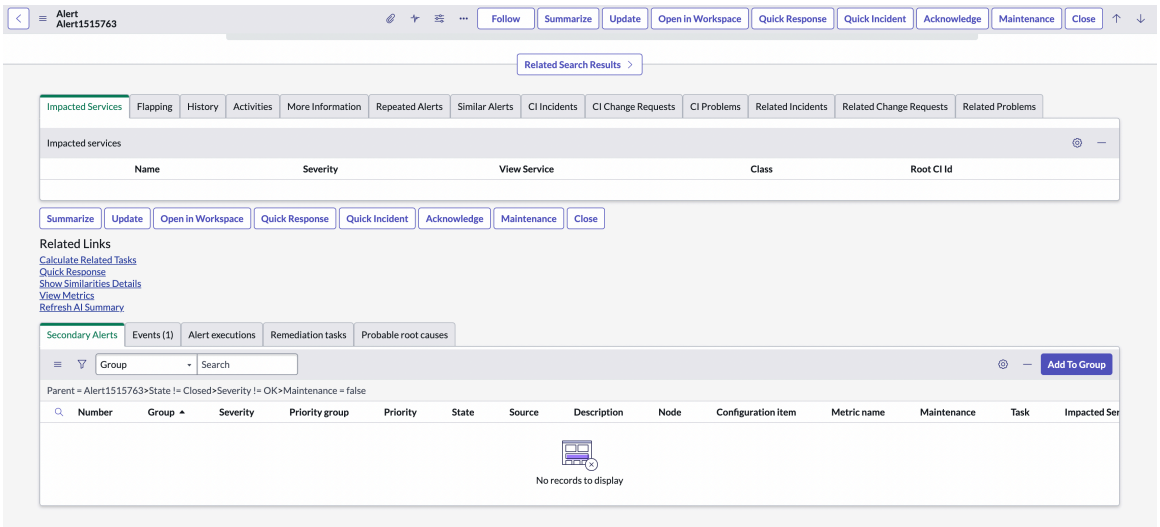

4. Select the check boxes for all alerts you want to add to the group as secondary alerts and select **Add Selected**.

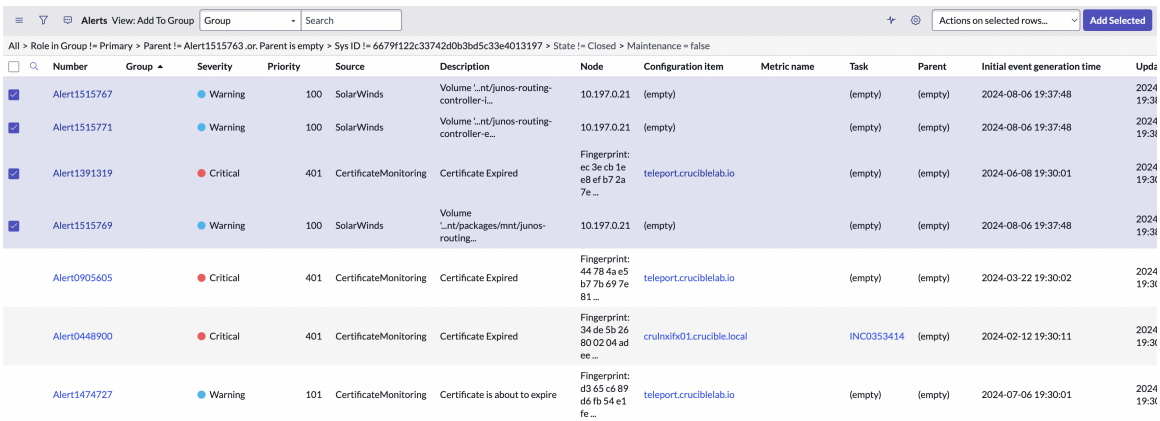

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. In the **Secondary Alerts** tab, view the alerts that have been added as secondary.

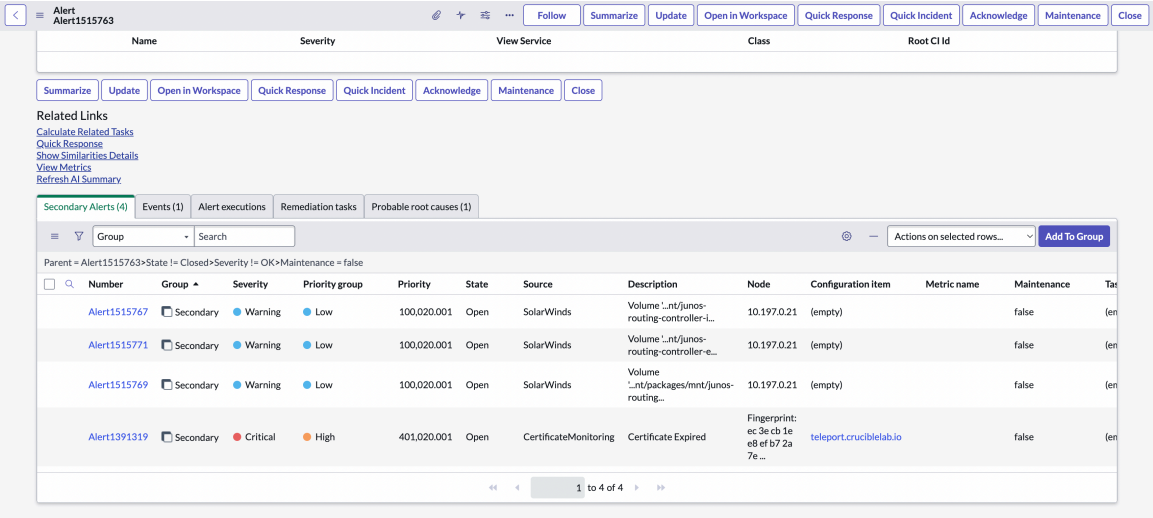

Add secondary alert manually to an existing alert group

Add any relevant alert discovered during the review of an automated alert group as a secondary alert to improve the group's completeness and utility for incident troubleshooting.

Before you begin

Role required: evt_mgmt_admin

About this task

The alert you select to add to a new group may already be part of another group. Adding it to the new group will automatically remove it from its current group, changing its group association. Alternatively, you can update the alert's parent to move it to a different group. By specifying a new parent, the alert is transferred to the new parent's group and removed from the previous one. This can be accomplished by updating the **Parent** field on the alert form or by modifying the **Parent** column in the Alerts list.

Note: When you add an alert to a group, this action is logged as feedback for alert grouping as part of the process. Currently, the feedback option is available but does not trigger further actions after being logged. Future enhancements to this feature are being considered.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Select the alert number to open the alert.
3. In the **Parent** field, search for the alert you want to assign as the parent for the selected alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

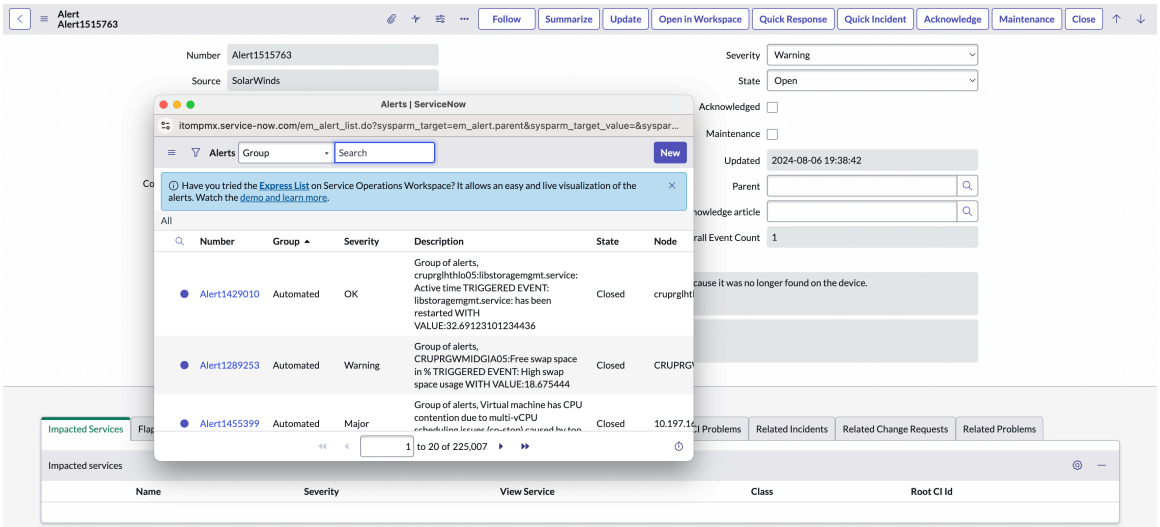

4. Select the parent alert number.
5. Select **Update**.

Result

The alert is added as a secondary alert to the selected parent alert.

Remove an alert from an alert group

Remove an alert if you want to improve the group's accuracy and usefulness in troubleshooting an incident. If this action leaves the group with one or no alerts, the group is deleted and no longer displays as an alert group.

Before you begin

Role required: evt_mgmt_operator

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Select the alert number to open the alert group.
3. Scroll down to view the **Secondary Alerts** tab.
4. Select the check box next to the alert you want to remove from the group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. From the Actions on selected row menu, select **Remove From Group**.

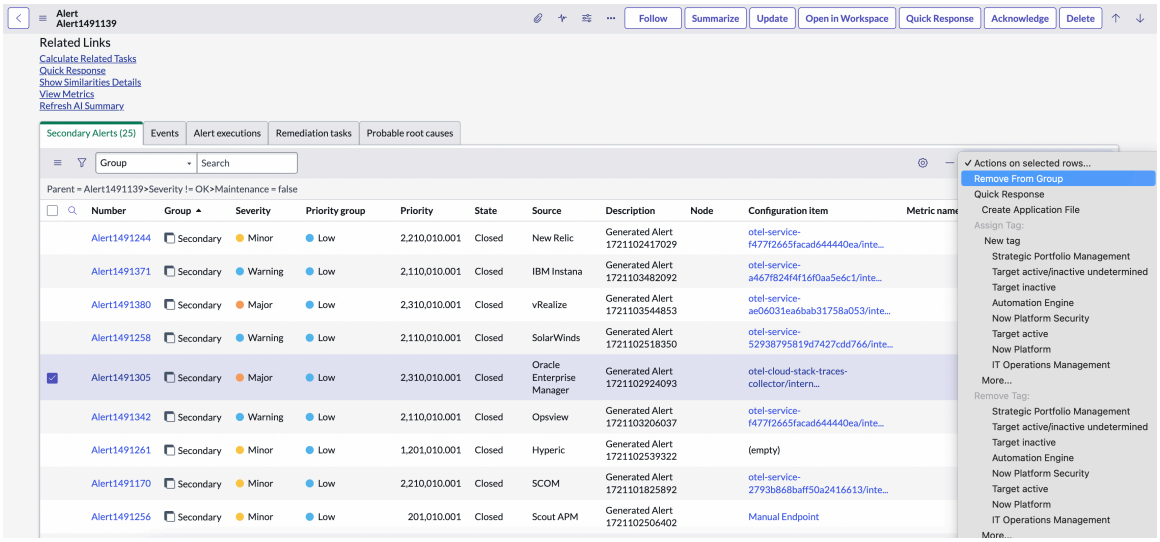

Result

The alert is removed from the group.

Configure work notes for alert grouping justifications

As alerts are added to a group, a message is recorded in the alert's Work notes field to explain why the alert was included in the group. Define alert types for creating these worknotes related to alert group reasoning.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for the property *`evt_mgmt.alert_groups_reasoning.enable_worknotes`*.
3. Open the property.

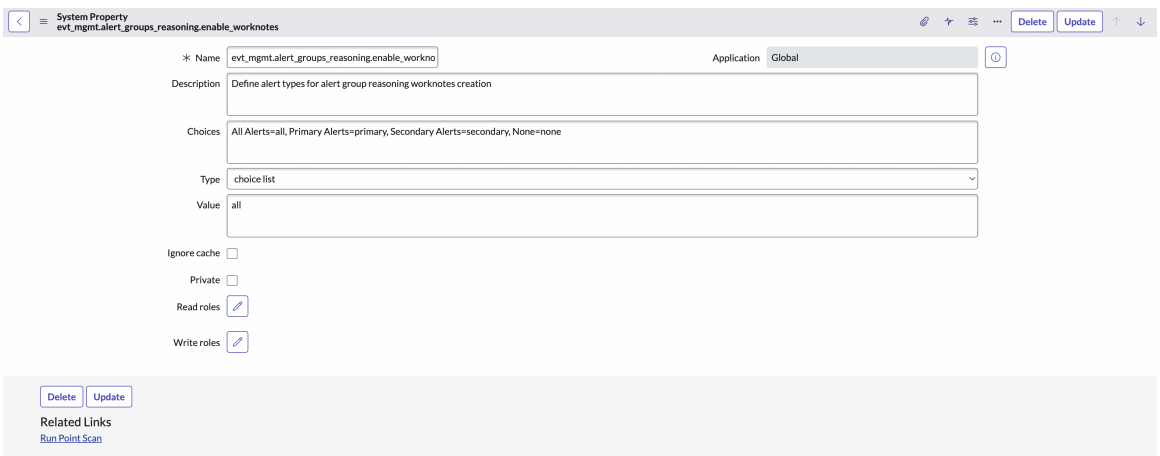

4. In the **Choices** field, define alert types for creating worknotes related to alert group reasoning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property values

Value	Description
All Alerts	Work note information is generated for every alert.
Primary Alerts	Work note information is generated only for primary alerts.
Secondary Alerts	Work note information is generated only for secondary alerts.
None	No work note information is generated.

5. Select **Update**.

Related topics

Alert grouping types and creation methods

Application services in Event Management

An application service is a set of interconnected applications and hosts which are configured to offer a service to the organization.

For general information on application services, see Application services.

Manually create an application service in Event Management

You can manually create an application services. Event Management can use application services to monitor service performance and identify health issues.

Before you begin

Role required: evt_mgmt_admin

About this task

An application services is a set of interconnected applications and hosts which are configured to offer a service to the organization. Application services can be internal, like an organization email system or customer-facing, like an organization website. Use application services to organize, maintain, and monitor services in your organization.

Application services do not support these CI types:

- cmdb_ci_endpoint
- cmdb_ci_translation_rule
- cmdb_ci_config_file
- cmdb_ci_qualifier
- cmdb_ci_application_cluster

Procedure

1. Navigate to **All** > **Event Management** > **Services** > **Application Services**.
2. Click **New**.
3. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a Service form

Field	Description
Name	Enter a unique application service name.
Owned by	Specify an owner who is responsible for this application service in the organization
Add entry points:	
CI Type	Select the CI type (CI class) for the CI you are adding. Every CI belongs to a CI type which contains a set of attributes configured for this kind of CI, for example, cmdb_ci_appl for applications.
CI Name	Select the CI from the list of CIs of the selected CI type.

The CI type list includes only allowed CI types. For example, you cannot add an application cluster.

4. Click **Submit**.

What to do next

For details on grouping application services, see Group application services.

For details on controlling user access to application services, see Control user access to application services.

Convert manual services to application services

You can convert existing manual services to application services. Event Management can use application services to monitor service performance and identify health issues.

Before you begin

Role required: evt_mgmt_admin

About this task

Manual services are no longer supported and this service type, as well as discovered services, have been replaced by application services. Any existing manual services must be converted to application services.
During the conversion process:

- The manual service record is moved from the Manual Services [cmdb_ci_service_manual] table to the Mapped Application Service [cmdb_ci_service_discovered] table by changing the record class.
- All the original manual service attributes, such as, name, owner, and operational status, are retained in the application services.
- The configuration items (CIs) of the manual service are retained in the application services after the conversion.
- The system queries the CMDB for the latest CI changes and applies CI impact rules.

Application service does not support these CI types:

- cmdb_ci_endpoint
- cmdb_ci_translation_rule

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- cmdb_ci_config_file
- cmdb_ci_qualifier
- cmdb_ci_application_cluster

If the preceding CI types are used in the manual service that you want to convert, the conversion fails. CIs of this type should be removed before conversion. They can be replaced by adding new CIs to the application services after conversion.

If your ServiceNow instance uses domain separation, the conversion fails if CIs from multiple domains are in the same service. Replace or remove relevant CIs before conversion. The converted application services contains only CIs assigned to the same domain as the application services.

Procedure

1. Navigate to **All** > **Event Management** > **Services** > **Manual Services**.

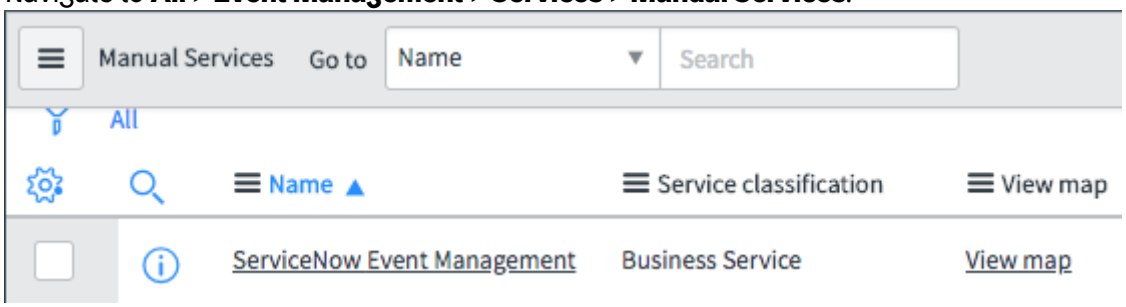

2. Select the manual service that you want to convert to an application services.

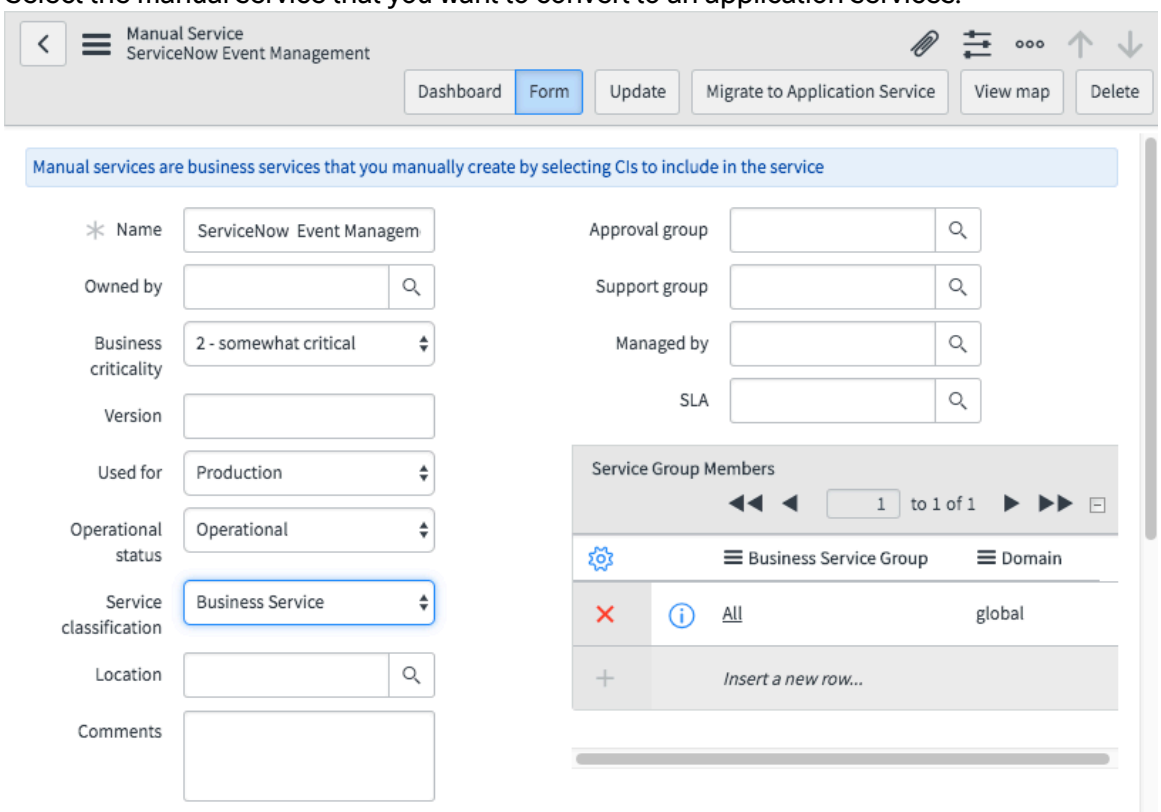

3. Click **Migrate to Application Service**.
4. Click **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

- The CIs from the manual service are added to the converted application services.
- After the conversion process is complete, the application services form opens.

Convert manual services to application services using API

You can use a JavaScript API to convert existing manual services to application services. Event Management can use application services to monitor service performance and identify health issues.

Manual services are no longer supported and this service type, as well as discovered services, have been replaced by application services. Any existing manual services must be converted to application services. Ensure that any automation of manual services (such as business rules or scheduled jobs) have been built and tested before converting to application services.

Note: Manual services are static and retain their original structure after conversion to application services. To update imported manual services with the existing CMDB state, either populate the original manual services to ensure that all CIs are assigned to the correct service, or run the **update with CMDB changes** option after the conversion.

During the conversion process:

- The manual service record is moved from the Manual Services [cmdb_ci_service_manual] table to the Mapped Application Service [cmdb_ci_service_discovered] table by changing the record class.
- All the original manual service attributes, such as, name, owner, and operational status, are retained in the application services.
- The configuration items (CIs) of the manual service are retained in the application services after the conversion.
- The system queries the CMDB for the latest CI changes and applies CI impact rules.

The maximum number of CI connections added to application services during this operation is controlled by the sa.service.max_ci_service_population property. By default, the value is 1,000 (one thousand connections). Increasing the number of CI connections may cause performance issues. To adjust the maximum number of added CI connections, add the sa.service.max_ci_service_population property, as described in Add a system property.

Attempting to process more CIs than the value specified in the sa.service.max_ci_service_population property causes the operation to fail.

Run the conversion process in small chunks of approximately 10 at a time. This makes it easier to pinpoint the problematic service if a conversion problem occurs.

Application services do not support these CI types:

- cmdb_ci_endpoint
- cmdb_ci_translation_rule
- cmdb_ci_config_file
- cmdb_ci_qualifier
- cmdb_ci_application_cluster

If the preceding CI types are used in the manual service that you want to convert, the conversion fails. CIs of this type should be removed before conversion. They can be replaced by adding new CIs to the application services after conversion.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If your ServiceNow instance uses domain separation, the conversion fails if CIs from multiple domains are in the same service. Replace or remove relevant CIs before conversion. The converted application services contains only CIs assigned to the same domain as the application services.

The JavaScript API is:

```
BusinessServiceManager.migrateManualToApplicationService(String
 service_id)
```

Parameters

Name	Type	Description
service_id	String	sys_id of the manual service to be converted

Returns

Type	Description
Boolean	`true` for a successful conversion, `false` for a conversion that failed

Where the `sys_id` of an existing manual service is `451047c6c0a8016400de0ae6df9b9d76`, run the following API to convert it to an application services:

```
var bsManager = new SNC.BusinessServiceManager();
var res =
 bsManager.migrateManualToApplicationService("451047c6c0a8016400
de0ae6df9b9d76");
```

Related topics

BusinessServiceManager - Global

Configuring CIs in a manual service as a manual cluster

Configure or modify a CI as a specific CI or a generic CI class in a manual service (that was not discovered automatically) as a manually defined cluster. A manual cluster delivers redundancy capabilities and provides continued operations if one or more CIs fail.

Manual cluster allows you to group CIs under a logical cluster in a service. In the impact tree you can set the relative impact of a cluster member on the cluster.

Note: Creation of nested manual clusters (cluster inside cluster) is currently not supported.

Once the CI is defined as a manual cluster its impact calculation shows the magnitude of an outage on its children CIs. Impact calculation uses factors such as impact rules and CI relationships to calculate the severity of a generated alert. Impact rules estimate the magnitude or severity of an outage based on affected CIs. They determine how the impact applies to parent or child entities that are part of an application service.

The magnitude of an outage can be 100% down, 60% affected, 40% impaired, or 20% impaired.

Redundancy is defined in impact rules. For example, you can define that the parent CI is critical if at least two child CIs are critical. The impact from child to parent CIs is not always 100%. If the CIs of all three children are affected, the impact is 100% down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In a service map, you can configure or modify a specific CI or a generic CI class in a manual service as a manually defined cluster. When configuring as a generic CI class, all the CIs of the selected class in that service are configured the same.

Note: Visual representation of a manual cluster may differ from the visual representation of a cluster that is discovered automatically.

Related topics

Alert impact calculation

Configure a manual cluster

Provide redundancy capabilities of an entire cluster in case of failure of one or more CIs in that cluster. By viewing the relative impact of a cluster member on the cluster information, you can better understand the source of alerts and take remediation steps.

Before you begin

Role required: evt_mgmt_admin

About this task

In a manual service that was not discovered automatically, you can manually group a specific CI or a generic CI class under a logical cluster. The relative impact of a cluster member on the cluster can be configured in the impact tree.

For more information, see Alert impact calculation.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Manual Clusters**.
2. Select **New**.
3. In the **Select a Service containing this manual cluster** section, select an application service to contain the new manual cluster based on a specific CI or general CI class.
   The cluster number is automatically generated and appears in the **Number** field.

   a. Select the Lookup using list icon ( ) in the **Service** field.

   b. In the CMDB CI Service [cmdb_ci_service_discovered] table, select a manually created service.
   To define a new application service, see Manually create an application service in Event Management.

   c. **Optional:** Preview the record by selecting the Preview this record icon ( ).

4. In the **Select a Specific CI or generic CI class** section, in the **Defined As** field, select how to define the manual cluster.
   Select a specific CI or a generic CI class that applies to all CIs of this class.
   - Specific CI - In the **Configuration item** field, select the Lookup using list icon ( ) and select the configuration item.
   - Generic CI class – In the **Configuration item class** field, select the configuration item class.
5. In the **Provide a Description** section, either provide a description in the **Description** field or leave it blank to contain an automatically generated description when you select **Save** or **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Description** field is automatically updated when you select **Save** or **Submit**. The automatically generated description consists of the configuration item name and the service name. For example: `CI <Alert Processing> in service <linuxApplService> (*** auto generated ***)`.

To prevent an automatic update of the Description field and retain your changes, delete the (*** auto generated ***) suffix that appears as part of the description.

6. Activate the cluster by selecting the **Activate** check box.
7. Select **Submit**.

Result

The manual cluster is saved in the EM Impact Manual Cluster [em_impact_manual_cluster] table. The impact calculation is recalculated automatically.

Related topics

Configuring CIs in a manual service as a manual cluster

Modify a manual cluster

Modify an existing manual cluster by changing the service, CI, CI class, or description.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Manual Clusters**.
2. From the list, select a manual cluster.
3. Change the service containing the manual cluster by selecting a different one in the **Select a Service containing this manual cluster** section.
   The cluster number appears in the **Number** field.
   a. Select the Lookup using list icon ( ) in the **Service** field.
   b. In the CMDB CI Service [cmdb_ci_service_discovered] table, select a manually created service.
   c. **Optional:** Preview the record by selecting the Preview this record icon ( ).
4. In the **Select a Specific CI or generic CI class** section, use the **Defined As** field to modify the manual cluster.
   - Specific CI - In the **Configuration item** field, select the Lookup using list icon ( ) and select the configuration item.
   - Generic CI class – In the **Configuration item class** field, select the configuration item class.
5. **Optional:** In the **Provide a Description** section, modify the description in the **Description** field.

   (Optional) The **Description** field is automatically updated when you select **Save** or **Submit**. The automatically generated description consists of the configuration item name and the service name. For example: `CI <Alert Processing> in service <linuxApplService> (*** auto generated ***)`.

   To prevent an automatic update of the Description field and retain your changes, delete the (*** auto generated ***) suffix that appears as part of the description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Activate or deactivate the cluster by selecting or clearing the **Activate** check box.
7. Select **Update** to save your changes.

Result

The manual cluster is saved in the EM Impact Manual Cluster [em_impact_manual_cluster] table. The impact calculation is recalculated automatically.

Related topics

Configuring CIs in a manual service as a manual cluster

Create an alert query

An alert query is a set of alerts that meet specific criteria for a particular service.

Before you begin

Navigate to **Event Management** > **Administration** > **Event Management Properties** and ensure that the **Enable alert query support** (*evt_mgmt.impact_calulation.alert_group_support*) property is set to **Yes**.

Role required: evt_mgmt_admin or evt_mgmt_operator

About this task

The main motivation to use alert queries is a modeling solution based on data contained in the alert itself, as an alternative to using either discovered, application, or technical services. Create an alert query to combine similar alerts that meet the specific criteria.

You can learn about Event Management basics, including alert queries, from this video:

Procedure

1. Navigate to **Event Management** > **Services** > **Alert Queries**.
2. Select **New**.
3. On the **Alert Query** form, fill in the fields.

   For a description of the field values, see Alert Query form

   Note: In the **Filter** field:

   - When defining an alert query filter, include only fields that appear in the Alert Histories [em_alert_history] table. Impact calculation is based on Alert History data and fields such as Overall Event Count, Priority, and Priority group are not copied to the Alert Histories [em_alert_history] table.
   - Do not specify a dynamic time condition. For example, in the filter, do not specify `Created` condition of `Last 45 minutes` because impact calculation is triggered by a change of alert or alert query. However, for the dynamic time condition, none of these conditions have changed.

     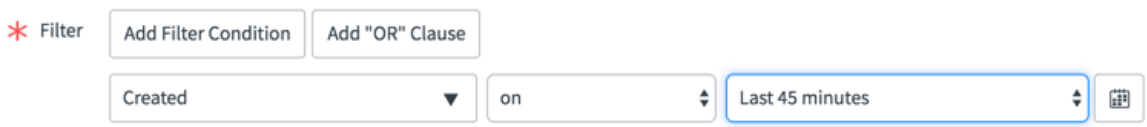

   - Some filters may slow down impact calculation. To solve this problem, adjust your alert query by adding an appropriate index, as described in Index suggestions for slow queries.

4. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Create alert group manually

Create an application service group

Create service groups to combine similar services. Organize services by groups to perform actions simultaneously on multiple services and to control user access to services.

Before you begin

Role required: evt_mgmt_admin

About this task

Service groups are a way to organize and group services based on the application services or technical services they impact. Instead of looking at individual service from different infrastructure components (like servers, databases, or applications), Service groups help consolidate them under the umbrella of a service group. Application service groups enable faster identification of alert impact. Service groups appear in the **Group by** option on the Service Operations Workspace dashboard.

Procedure

1. Navigate to **All** > **Event Management** > **Services** > **Service Groups**.
   The Application Service Group page opens.
2. Select **New**.
3. In the **Name** field, enter the name of the new service group.
4. **Optional:** In the **Parent Group** field, enter the name of another group in which to embed the new group.
   Users with access to a parent service group automatically have access to all its child groups.
5. Select **Submit**.

Related topics

Assign a role to a service group

Group application services

Assign a role to a service group

Assign an Event Management role to the application service group to ensure that group members can manage and act on alerts.

Before you begin

Configure a service group to enable users to manage application services or alert groups.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Event Management** > **Services** > **Service Group Responsibilities**.
2. Click **New**.
3. In the **Application Service Group** field, select the name of the service group.
4. In the **Role** field, select an Event Management role:
   - evt_mgmt_admin
   - evt_mgmt_integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- evt_mgmt_operator
- evt_mgmt_user

5. Click **Submit**.
6. To find users who are assigned to the role, navigate to **User Administration** > **Users** > **Roles** and search for the role.

Related topics

Roles installed with Event Management

Activate and configure Service Mapping for top-down discovery

A top-down discovery provides a list of CIs and their interrelationships. This information is useful for managing software services and hardware issues that are associated with alerts.

Before you begin
Role required: evt_mgmt_admin

About this task
In Event Management, Service Mapping relationships appear on alerts. You can view the related CI and other information in the Event Management dashboard. The mapping relationships appear in a service map with the generated alert, any CIs that are bound to the alert, and all related CIs.

Procedure

Use Service Mapping to get top-down discovery information.
Service Mapping, which is closely related to Event Management, provides tools to discover the relationships between network, hardware, software, storage, and other devices.

Related topics

Service Mapping

Request Service Mapping

Configure email notifications for service severity changes

Configure an email notification to notify users when there is an application service severity change.

Before you begin
Role required: evt_mgmt_admin

About this task
Configure an email notification to notify users when there is an application service severity change. The notification involves specifying when to send it, who receives it, and what it contains.

Procedure

1. Navigate to **All** > **System Properties** > **Email Properties**.
2. In the **Outbound Email Configuration** section, in the **Email sending enabled** option, select **Yes** to enable sending severity information by email.
3. Click **Save**.
4. Navigate to **System Notification** > **Email** > **Notifications**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Click **New**.

   a. In the **Email Notification Actions** form, specify these values.

Field	Value
Name	Type a descriptive name, for example, *Application Service severity change*.
Table	Select the required service type and its corresponding table: Alert Group, Discovered Service, Technical Service, or Application service.
Active	Select the check box to enable email notification.

   b. Click the **Who will receive** tab.
   In the **Users** area, search for and select the required user (the user does not require System Administrator credentials). Repeat this step for as many users that are required to receive this notification.

   c. Click the **When to send** tab and in the **choose field** box, search for and select *Severity*.

   d. Click the **What will it contain** tab and in the **Subject** field, type meaningful text that makes it clear to the receiver of the email notification what the content of the email is.

   e. In the **Message HTML** message area, using the **Select variables** functions, compose the conditions under which the email notification is to be sent.
   For example, configure these variables to change severity:

```
severity changed
```

```
Severity: ${severity}
```

```
on application service
```

```
Name: ${name}
```

6. Click **Submit**.

Configure contextual colors and icons

Use the **Contextual colors and icons** form to configure color, text, and icons to have different default or custom contexts, identified by a context ID.

Before you begin
Role required: evt_mgmt_admin

About this task
You can customize the colors or icons of elements in the **Contextual colors and icons** form to affect the displayed colors or icons of severities in related Event Management screens, either when accessibility mode is off or on. For more information about accessibility, see Accessibility features .

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In each color selection field, you have the option of entering the value of the color as any of the following Cascading Style Sheet (CSS) formats that determine how HTML elements should be displayed:

- Name: predefined color names, for example, red, green, olive, blue, darkblue, or lightblue
- RGB decimal: RGB(102, 153, 204)
- RGB hex: #223344

See the following websites for information about CSS color declarations (including hex or RGB notation):

- HTML Colors (W3CSchools) for more information on hex and RGB notation.
- HTML Color Names (W3CSchools) for valid color names.

The context IDs that are provided with the base instance are:

- information_severity
- warning_severity
- minor_severity
- major_severity
- critical_severity
- topology_connection

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Contextual colors and icons**.
2. Click **New** or modify an existing configuration.
3. Configure the following fields.

Item	Description
Context details	
Context ID	Enter a unique ID for the context, for example, clear_severity.
Context Name	Readable name for the context. This text is available to be read out by screen readers.
Batch	The batch that this context is a member of.
Description	Description of uses for the context.
Colors **Note:** For each entry in the color fields, after you enter a color code, a block appears to the right of the field showing the specified color. Use of invalid color codes triggers an appropriate error notification.	
CSS Color	The corresponding CSS color code for the context color name.
CSS Color on Interaction	CSS color code that appears on hover or focus.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Item	Description
Accessible CSS Color	Accessible CSS color code that appears when accessibility mode is enabled.
Accessible CSS Color on Interaction	Accessible CSS color code that appears on hover or focus.
Icons	
Icon	Icon that is displayed when the context includes an icon.
Accessible Icon	Icon that is displayed when the context includes an icon and when accessibility mode is enabled. To add your own icon, or to see a list of available icons, see Upload one or more images.

4. Click **Submit**.

View Event Management license usage

Event Management is licensed based on the number of CIs bound to alerts during the last year. For alerts not bound to CIs, the system calculates the number of nodes that can send events to the instance during the last year.

Before you begin

Role required: evt_mgmt_admin or evt_mgmt_operator

About this task

The **Event Management - Node Count Store** job calculates the number of unique nodes that send event information, either directly or through third-parties, to Event Management and are eligible for licensing in the last year. The information from this job is stored in the License Usage [em_unique_nodes] table for visibility by all users. CIs bound to alerts are counted towards licensing; if there are no CIs that are bound to alerts, the event node field is used. Node is a reference to a server, not the instances of the software on it. Your node count is based on the amount of infrastructure (servers for example) that will be sending events.

Procedure

Navigate to **All** > **Event Management** > **Reporting** > **License Usage**.

The table is filled automatically based on the alerts in the system. The user has no permission to edit the table. It is used to calculate the usage of the Event Management licensing.

Column	Description
CMDB CI	The CI that is bound to the alert that is gener ated from the event. If there is a value for this field, the **Node** field is empty.
Node	The string value of the **Node** field of the event. If there is a value for this field, the **CMDB CI** field is empty.
Type	One of these values:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
	• Unknown - A CI has not bound to the generated alert. • Server - The source of the event was a server. • PaaS - The source of the event was Platform as a Service.
Is licensable	One of these values: • True - The event listed is counted towards the license usage. • False - The event listed is not counted towards the license usage.

Alert automation in Service Operations Workspace for ITOM

Alert automation is crucial as organizations deal with increasing number of alerts and complex IT infrastructures. Manual alert handling is slow, error-prone and inefficient, underscoring the need for automated systems. Automation can improve the mean time to resolve alerts, improve service reliability and better scale staff resources.

Alert automations also support both centralized administrator and distributed team roles. This enables qualified teams to self-serve and create their own alert automations. For example, you may consider granting access to site reliability engineers (SREs). Members of teams can manage automations for their own team and their own alerts without impacting other teams.

For users familiar with our classic experience, alert automation offers an easier user interface and better team support for event rules, tag-based clustering definitions and alert management rules. Some advanced features are currently only available to admins in the classic experience. These two experiences use the same backend tables. You can use whichever experience is most convenient, and changes in one will also update the other.

https://player.vimeo.com/video/1068873181?h=b03003c9bf&badge=0&autopause=0&player_id=0&app_id=58479

Alert automation types

Currently, Service Operations Workspace ITOM provides the following types of automation.

1. Ignore automation: Reduce irrelevant or false-positive alerts, efficiently manage alert fatigue by filtering out noisy notifications, and allow teams to focus on critical issues.
2. Enrich automation: Enhance raw alerts with contextual information to make them more informative and actionable. In simple terms, this involves taking the raw events generated by monitoring tools and transforming them into a common and standard format to aid automated grouping and response.
3. Group automation: Group multiple related alerts into a single primary alert to reduce alert noise and identify the root cause.
4. Create Respond automation: Respond to alerts automatically by notifying appropriate stakeholders, escalate them as needed or run remediation actions. Determine how and when alerts are escalated based on severity or type. Integrate with third-party systems to create cases, notifications or run remediation actions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert automation process flow

You may start by sending alerts or events from monitoring systems to ServiceNow using the Integrations Launchpad. This is where administrators establish connections between ServiceNow and monitoring tools. These integrations enable the collection of monitored data, generating events from third-party sources.

When alerts are received by ServiceNow, alert automations run in the order shown on the page. First, we ignore alerts to reduce noise. Next, we enrich alerts with extra context, then group the alerts using the added context. Finally, we respond to alerts by escalating or running remediations. There can be several automations for each type. Each automation runs based on specific trigger conditions and executes specific actions. Alerts are only automated when they are received; we do not apply automations to past alerts.

In the alert enrichment phase, administrators add or extract necessary fields from alerts to provide essential information for swift resolution. This ensures that alerts contain all relevant details required for effective incident response. Administrators add context to alerts by modifying and normalizing them. This enhances the correlation of alerts, making it easier to identify patterns and potential threats.

The enriched and composed alerts are then grouped based on predefined criteria, consolidating related alerts. This reduces alert fatigue and facilitates efficient remediation. Finally, escalated alerts trigger notifications to stakeholders through various channels, ensuring timely communication and response to critical alerts.

The following diagram illustrates this process flow.

Alert automation: Reducing noise and improving resolution time

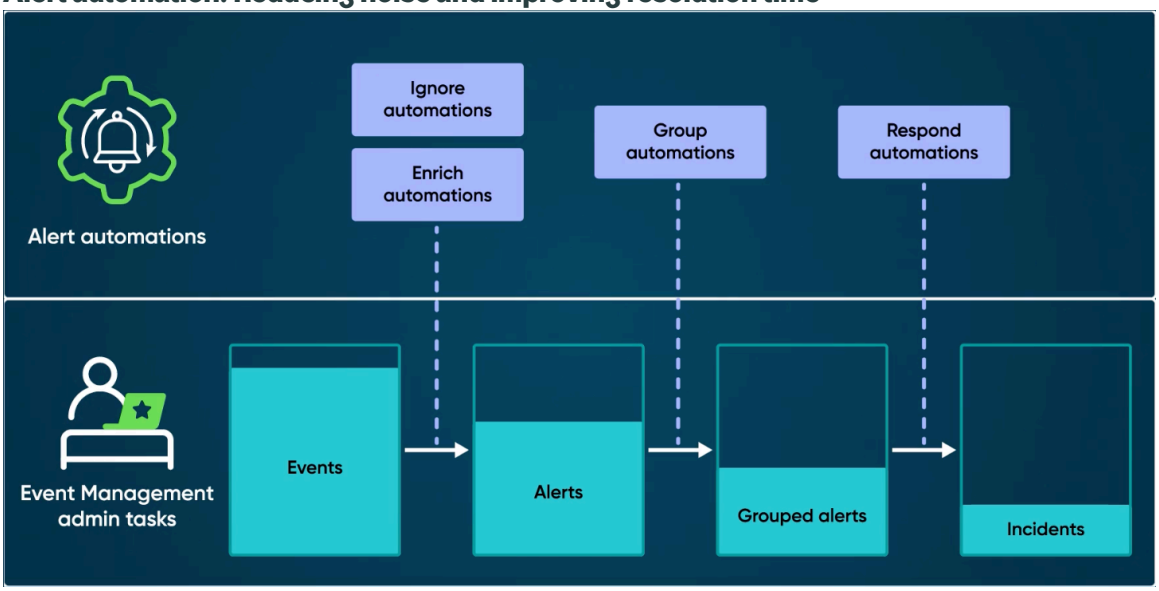

This comprehensive alert automation process can reduce alert noise, improve mean time to resolution (MTTR), enhance service reliability, and boost staff productivity.

Create Ignore automation

Ignore automation streamlines the process of disregarding irrelevant or false-positive alerts from monitoring systems, efficiently managing alert fatigue by filtering out unnecessary notifications. As a result, teams can prioritize and resolve critical issues more effectively.

Before you begin

Role required: evt_mgmt_admin, evt_team_operator, or srm_responder

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Ignore automations filter out alerts from source monitoring systems that match specific conditions. Separately, Event Management ignores alerts that have a duplicated message key field or where the severity is **Clear**.

For users familiar with the classic Event Management experience, ignore automations provide a simpler interface with enhanced team support for the event filter section of event rules.

Note: Ignoring alerts before enriching them improves performance and simplifies the processing flow. The condition filter in ignore automations applies only to the original event, not the enriched one. Enrich automations do not run if **Don't run other enrich alert automations** is set to false, and another enrich automation with a lower order, the same source, and matching filter conditions already applies. To avoid conflicts, consider the enrich automation order, source, and filter conditions when creating enrichments or event rules.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the Optimize section, select **Ignore alerts**.
   The Ignore alerts page is displayed.

Automations page

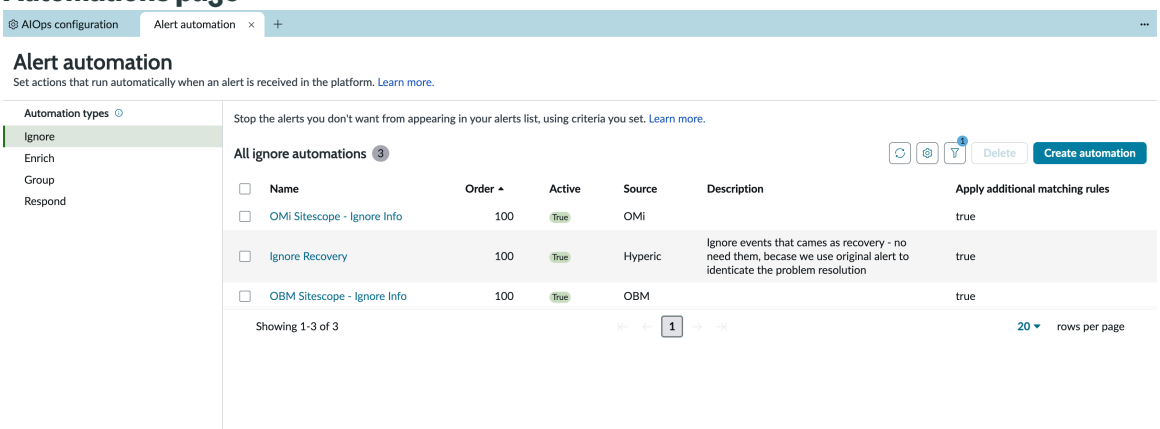

4. Select **Create automation**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

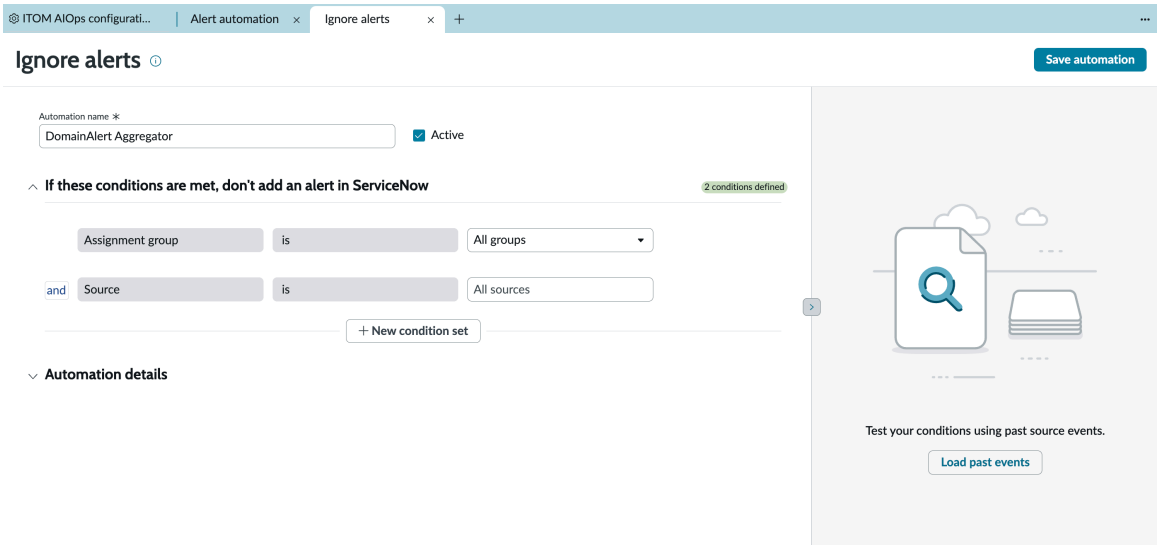

5. In the **Automation name** field, enter the name of the automation for ignoring alerts.
6. Activate the automation by enabling the **Active** toggle switch.
7. In the **If these conditions are met, don't add an alert in ServiceNow** section, set up filter criteria to identify the events you want to capture.
   a. From the **Assignment group** field menu, select the assignment group to determine which team's alerts trigger the automation.

      The **Assignment group** represents a specific team responsible for handling certain alerts. By selecting an assignment group, you ensure that only the alerts assigned to that a particular team trigger the automation. This way, the automation is targeted and only activates for relevant alerts associated with the selected team.

      Note:

      - If you're logged in to the instance with the administrator role (evt_mgmt_admin), all assignment groups are available. Additionally, you can also select **All groups** to generate alerts for any available group.
      - If you're an operator, only the group you're a part of is available.
      - Only members of the selected group or administrators can update or delete the automation.

   b. From the **Source** field menu, select the monitoring tool from where the alert is generated.
   c. Set up the conditions by selecting the field, operator, and field value. Then, add more conditions using OR or AND operators.

      To add another set of conditions, select **+ New condition set**. If you don't see an additional info field, you can manually it in the drop-down list.
8. To verify that you're ignoring the correct events or to see examples of fields that can be used to set the conditions, select **Load past events**.
   You can preview the details of some past events.
9. In the **Automation details** section, provide an order and automation description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

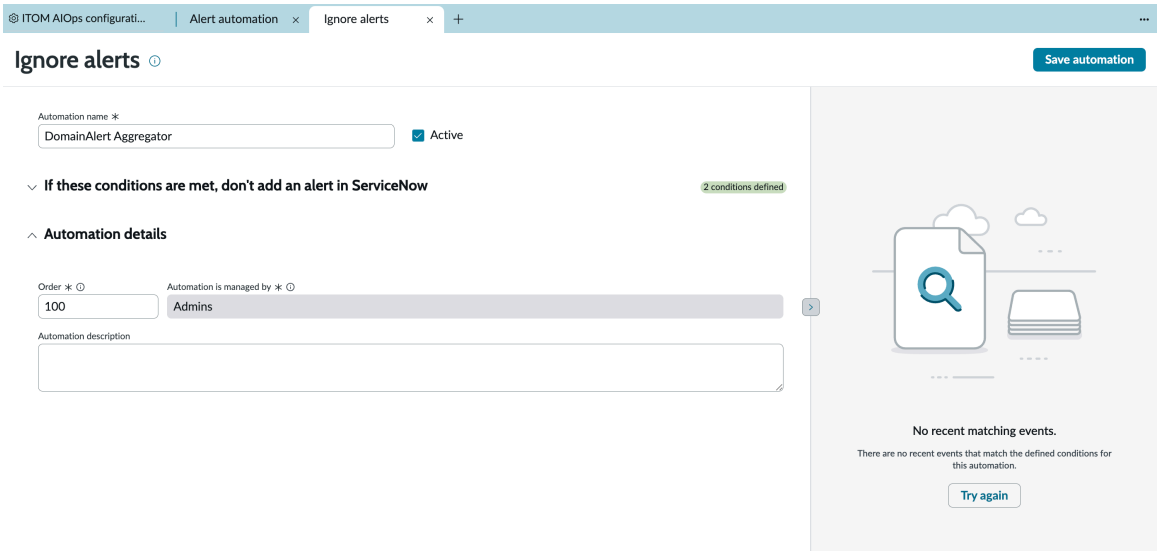

a. In the **Order** field, enter the automation order.

> **Note:** Automations run in order from the lowest to the highest.

The **Automation is managed by** field displays the team or assignment group who owns, edits, and can delete this automation. The assignment group is the same as the one defined in the **If these conditions are met** section.

b. In the **Automation description** field, enter a brief description of the automation.

10. Select **Save automation**.
A notification appears when the automation is successfully saved. Otherwise, an error message is displayed. The ignore automation that you created appears on the **Ignore alerts** page where you can view, edit, or delete the existing automation.

What to do next
You can transform raw events into an understandable format by creating Create Enrich automation.

Create Enrich automation

Alert enrichment involves transforming raw events from monitoring tools into a standard format, aiding automated grouping and response. This includes extracting fields from lengthy alert payloads or composing them into a standardized format. Additionally, you can create tags, which are metadata added to alerts for easier filtering and grouping.

Before you begin
Role required: evt_mgmt_admin, evt_team_operator, or srm_responder

About this task
Extracting takes values from event payload fields and places them in alert output fields, while composing combines multiple alert fields into one. For more information, see Extracting and composing alert fields.

For users familiar with the classic Event Management experience, enrich automations create event rules but with an easier interface and with better teams support. Event rules offer a few advanced features like thresholds and CI identification using IRE that are not yet available in enrich automations. Admins may also enrich alerts with event field mapping rules. Changing alert values creates an event field mapping rule with the mapping type **Map field and transform**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

value (Single field). This rule is linked to the event rule and runs simultaneously, allowing for streamlined mapping and transformation of event data to enrich alerts.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the Optimize section, select **Enrich alerts**.
   The Enrich alerts page is displayed.

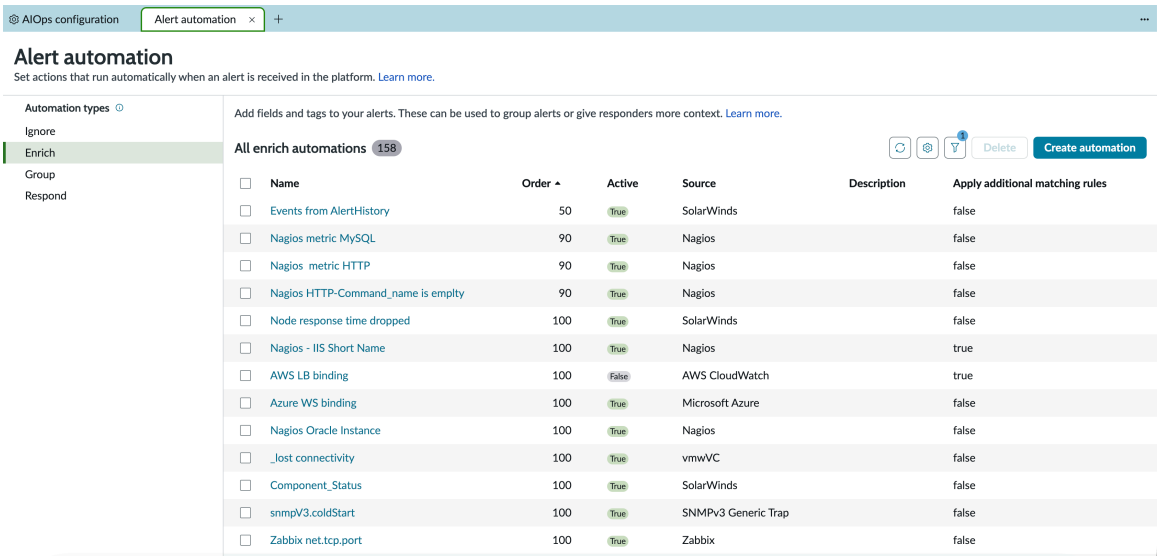

4. Select **Create automation**.
   By default, the **Active** check box is selected.
5. In the **Automation name** field, enter the name of the automation for enriching alerts.
6. In the **If these conditions are met** section, set up filter criteria to identify the alerts you want to enrich.

   The condition is evaluated against the raw event received from the source monitoring system and does not account for enriched fields.

   a. From the **Assignment group** field menu, select the assignment group to determine which team's alerts will trigger the automation.

      The **Assignment group** represents a specific team responsible for handling certain alerts. By selecting an assignment group, you ensure that only the alerts assigned to that particular team will trigger the automation. This way, the automation is targeted and only activates for relevant alerts associated with the selected team.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- If you're logged in to the instance with an administrator role (evt_mgmt_admin), all of the assignment groups are available. Additionally, you can select **All groups** to enable generating alerts for any of the available groups.
- If you're an team operator, only the groups you're a member of are available.
- Only members of the selected group or administrators can update or delete the automation.

b. From the **Source** field menu, select the monitoring tool from where the alert is generated.

c. Set up the conditions by selecting the field, operator, and field value. Then, add more conditions using OR or AND operators.

When using time-based condition filters in Enrich/Event rules (a filter that contains a time-based condition), the value is saved as a static value. The time is evaluated in the system time zone of the instance, defined by the *`glide.sys.default.tz`* system property, which defaults to America/Los_Angeles if left empty.

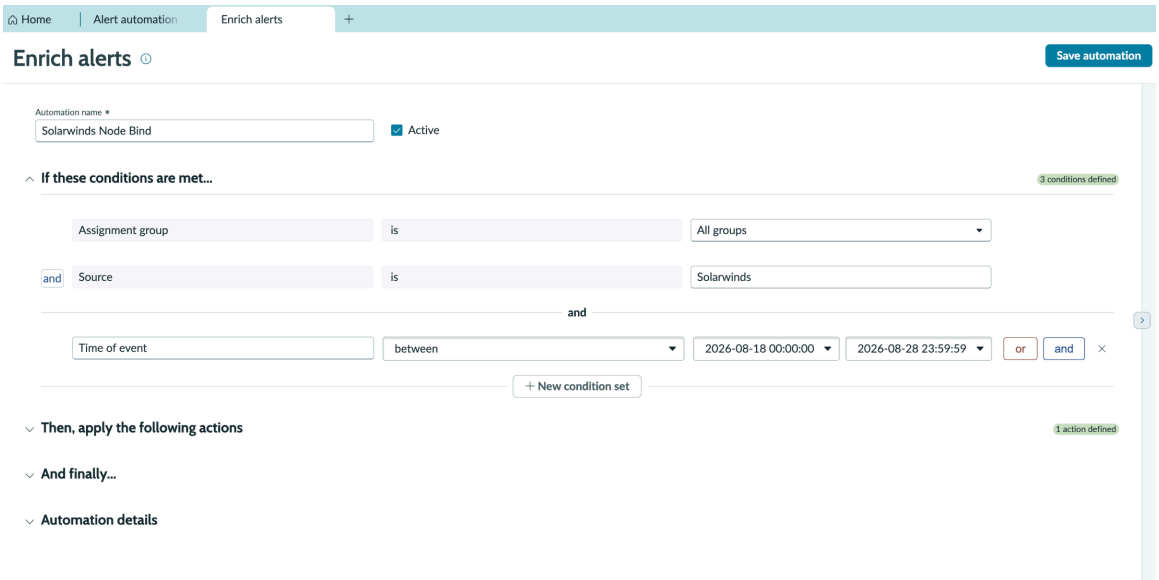

To add another set of conditions, select **+ New condition set**. You can also manually add an additional info field if you don't see it in the drop-down list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Select **Load past events** to view previous events when creating the automation.

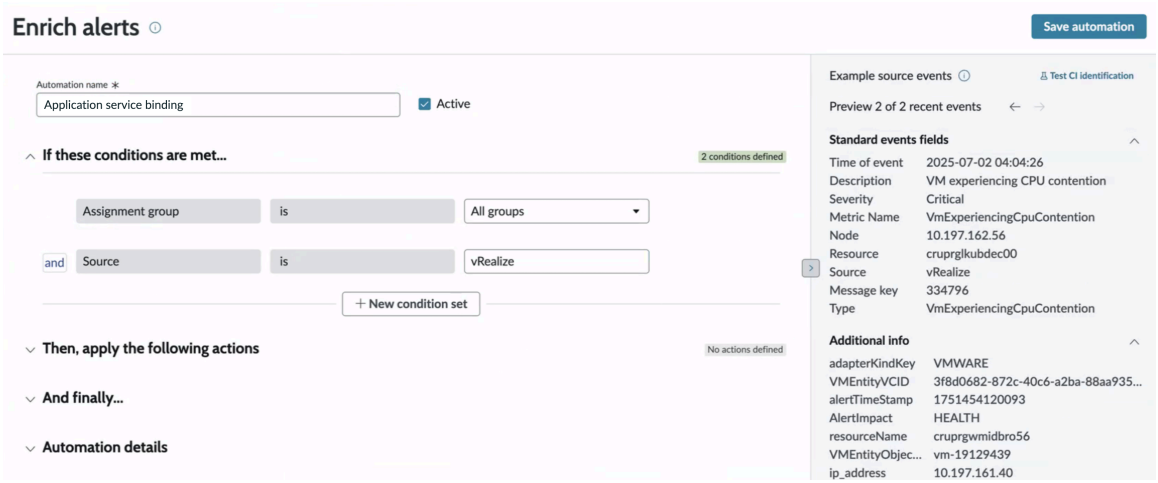

7. In the **Then, apply the following actions** section, select the automation actions that will be triggered by this automation.
   You must select at least one action.
   - **Extract alert fields**: Retrieves alert field values from the event payload and places them in an alert output field.
   - **Copy or compose fields**: Merges various alert fields, tags and text to generate a composed alert output.
   - **Change alert values**: Maps the current value of alert fields to specified new values.
   - **Improve configuration item (CI) identification**: Identifies CIs other than a host for better alert grouping, team assignment, and service health.

Option	Action
Extract alert fields	**a.** Enable the **Extract fields** toggle switch. **b.** From the **Extract from field** menu, select a value. The menu displays the standard event fields, additional info, and tags. The field value is then displayed. You can also manually enter a field name that is not displayed and add your own value. The example source events pane displays a sample of recent events in your system. If no events are displayed, you may cre ate an event, see Create or edit an event rule. **c.** In the **Regular expression** field, create a regular expression to extract the value that you want to extract.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
	Note: You can compose text using regular expression (regex) format conventions. Use one or more capture groups with parentheses to extract parts of the input. Capture groups in the regular expression are assigned to alert outputs based on the order in which they appear. The regex must match the entire input, so consider sur rounding your regex with `.*` on each end. For example, `(\w+).acme.com.*` captures the host name in a fully qualified domain name. The parser for the regex engine is Perl Compatible Regular Expressions (PCRE) compat ible. Nested JSON data is preprocessed to strip quotes and replace colons with equal signs. **d.** In the **Alert output** field, select an alert field or an alert tag. You can also manually enter a new field name. If you want to add an alert tag, select the **Set as a tag** check box. **Tip:** Create alert tags that can be shared across sources for easier filtering and grouping, such as the out-of-the-box tags. 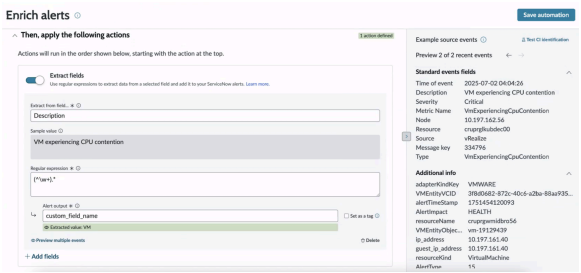 **e.** Select **Preview multiple events** to verify that the regular ex pression (regex) is general-purpose enough to correctly ex tract values across many examples. **Note:** This option is available only when example source events are available and matched with the regex filter. 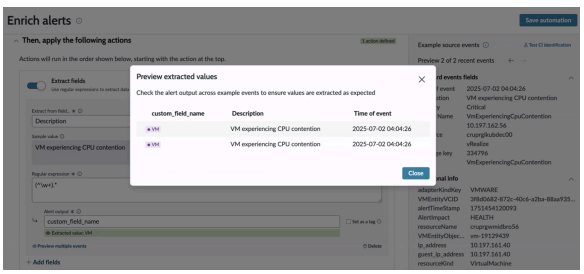 To include additional fields for extraction, select **+ Add fields**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
Copy or compose fields	**a.** Enable the **Copy or compose fields** toggle switch. **b.** From the **Composition** field, select alert fields and/or alert tags, manually enter a field name or even add free text. Alert fields are displayed in the `${field}` syntax format. **c.** In **Output to field**, select an existing alert field or alert tag, or manually enter an alert field. **Output to field** is an enriched alert containing the composed alert data. For easier grouping, you can select a tag from the menu. If you want to use the new field name as a tag for grouping, select the **Set as a tag** check box. **Tip:** Create alert tags that can be shared across sources for easier filtering and grouping, such as the out-of-the-box tags. 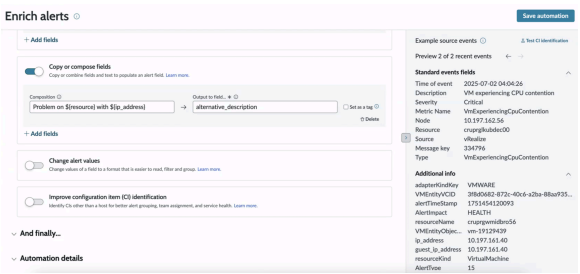 To create additional alert data compositions, select **+ Add fields**.
Change alert values	**a.** Enable the **Change alert values** toggle switch. **b.** In **Alert field to change values for** field, enter the field in the alert that you want to map values for. **c.** In the **When value is** field, enter original value in the alert field that you want to change. **d.** In the **Change it to** field, enter the new value that will replace the original value in the alert field. To add more field values, select **+ Add value** and to add more fields to map, select **+ Add field to map**. 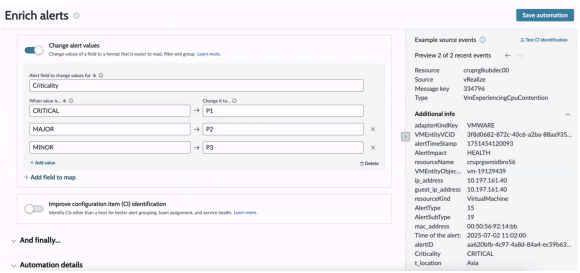
Improve configuration item (CI) identification	This option allows you to change how alerts are bound or linked to a Configuration Item (CI). This ensures alerts are associated with the correct IT components for better visibility and faster is sue resolution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
	The default and most common way to bind alerts to CIs is based on the **Node** field. This works out of the box with no configura tion needed. Use it by populating the **Node** field in your alert with a CI's Name, Fully Qualified Domain Name (FQDN), IP, or MAC ad dress. This supports host CIs including Computers, Operating Systems (OS), Switches, Routers, or any CI type or any class that extends the [cmdb_ci_hardware] table. You may enable this action to improve the CI identification for other types of CIs such as processes or service instances. The system searches for a matching CI in the appropriate CMDB table based on the selected CI type. For example, if you se lect VMware Virtual Machine Instance as the CI class, the sys tem searches for a matching record in the [cmdb_ci_vmware_ instance] table. The search uses details from the event rule record, specifically the **Additional info** fields. 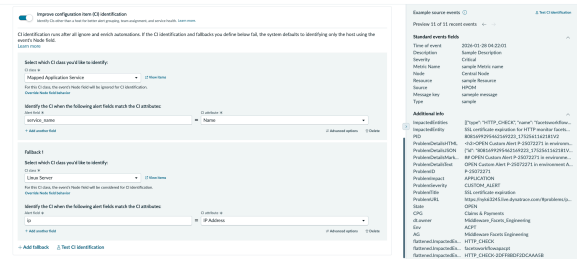 **a.** Enable the **Improve configuration item (CI) identification** tog gle switch. **b. Select which CI class you'd like to identify**: In the **CI class** field, specify the CI class for which the event's node field is evaluated during CI identification. This determines the type of CI the system attempts to identify first. Select **View items** to see the list of CIs of this type and their at tribute values. **c. Identify the CI when the following alert fields match the CI at tributes**: In this section, map alert fields to CI attributes to help the system find the correct CI: ▪ **Alert field**: Select the alert field that contains identifying in formation. ▪ **CI attribute**: Select the CI attribute that the alert field must map to. **d.** Use Advanced options to control how the system treats the alert node field during CI identification. ▪ Node field behavior: Choose one of the following: ▪ Consider node field: Uses the alert's node value during CI identification. ▪ Ignore node field: Ignores the alert's node value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
	Based on the selected CI class, the system recommends whether you should consider or ignore the node field, but you can override this recommendation if needed. ▪ Use all matching additional info fields: Select this check box to use all matching fields from the event's Additional Info for CI identification instead of mapping individual fields. **e. + Add fallback**: Add fallback CI class and field mappings to de fine alternate identification logic. If the primary configuration does not identify a CI, the system evaluates each fallback in or der until it finds a match. Ensure that the **Node** field in the alert is populated correctly to identify a host CI. The CI you are identifying must have a runs on relationship to the host CI or be mapped to the host. **f.** Ensure that at least one CI attribute is present in the **Addition al info** field of the alert. For instructions on how to populate these fields, see Set ad ditional info fields to match CI attribute format. The system at tempts to match values from the **Additional info** field of the alert with the CI table. If a match is found, the alert is bound to the corresponding CI. **g.** Select **Test CI identification** to test the CI identification on sample events. After successful CI identification, the sys tem displays a message with the relevant de tails. 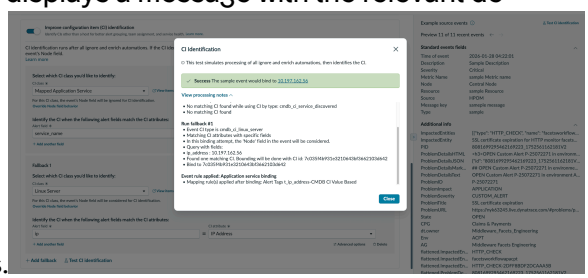 If CI identification fails, the system shows a message explain ing the details. 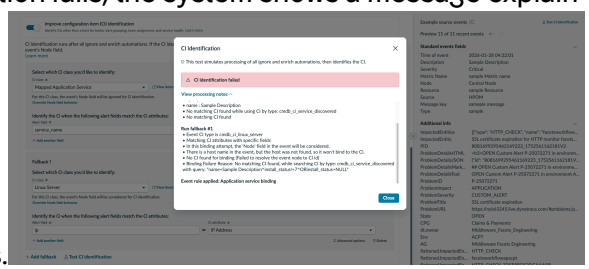 For more information on CI binding, see Binding alerts to CIs.

8. In the **And finally** section, to continue running other enrichment automations with same filter conditions after this automation is executed, select **Run other enrich alert automations**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

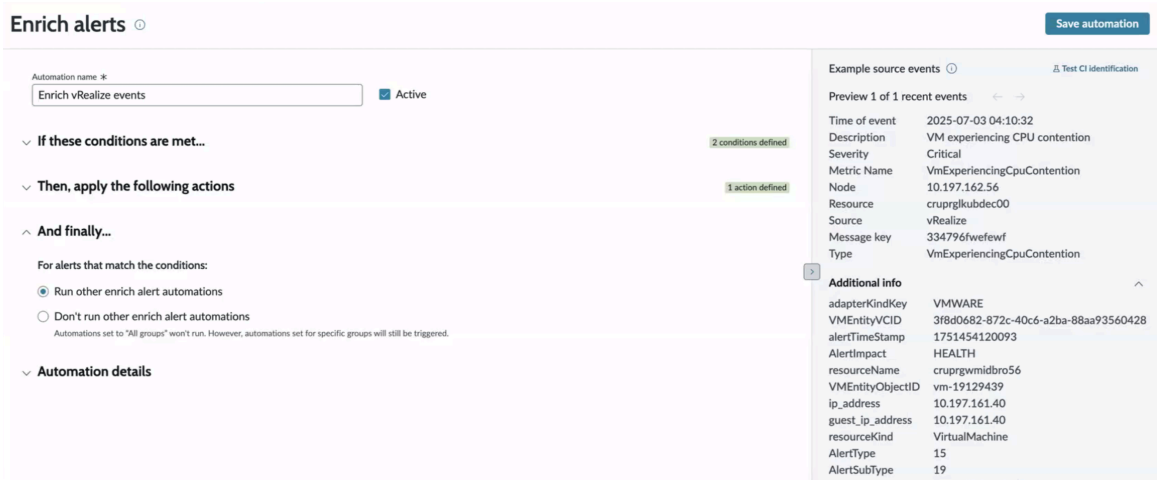

If you select **Don't run other enrich alert automations**, additional automations of this type will stop running after this automation is executed once. If the automation is managed by an administrator, it will stop running administrator-owned automations but will continue to run automations owned by other assignment groups.

9. In the **Automation details** section, provide an order and automation description.

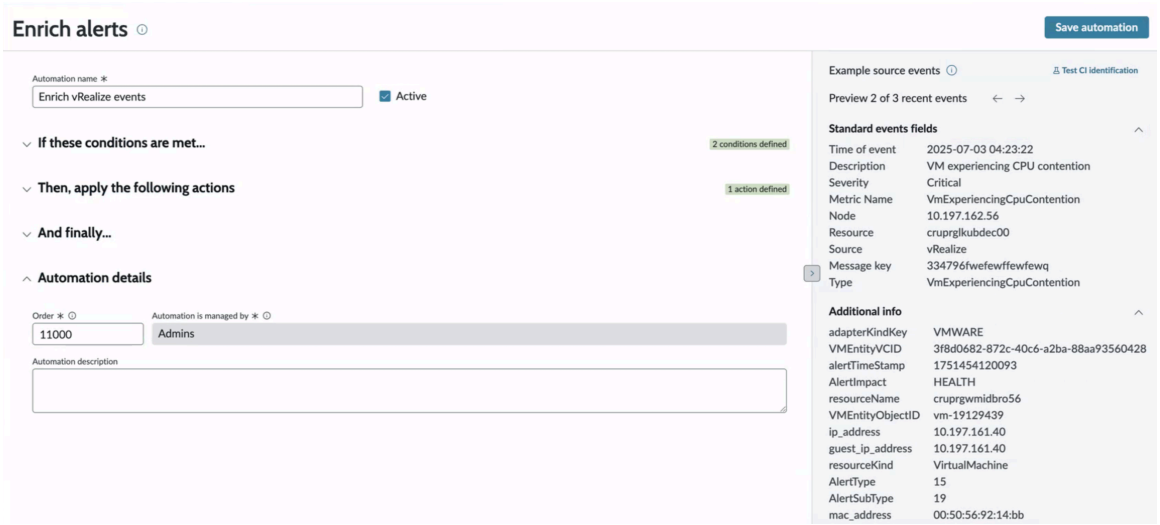

a. In the **Order** field, enter the automation order.

Note: Automations run in order from the lowest to the highest. Ensure there are no enrichment automations with a lower order number that have matching conditions and have **Apply additional automations of this type** set to false. Otherwise, this may prevent subsequent automations from running.

During alert processing, enrichment automations assigned to specific teams are applied first, based on the **Order** field. After these are evaluated, enrichment automations assigned to **Global / All groups** are applied. If multiple automations match, the automation with the highest order value is selected to improve CI identification.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Automation is managed by** field displays the team or assignment group who owns, edits, and can delete this automation. The assignment group is the same as the one defined in the **If these conditions are met** section.

b. In the **Automation description** field, enter a brief description of the automation.

10. Select **Save automation**.
A notification appears when the automation is successfully saved. Otherwise, an error message is displayed. The enrich automation that you created appears on the Enrich alerts page where you can view, edit, or delete the existing automation.

What to do next

You can manage alerts more effectively by grouping similar alerts together with the help of Create Group automation.

Related topics

Extracting and composing alert fields

Extracting and composing alert fields

Extracting and composing are ways to manage what you see in the alert output, making it simpler to filter, group, and read. Alert automation enables you to extract values from event payload's alert field and place it in an alert output field. Composing allows you to merge multiple alert fields into a single output field.

Extracting alert fields

Alert notifications often contain relevant context buried within event payloads. By enriching alert outputs with values from the existing payload, you can better understand the significance of alerts and determine the appropriate steps for resolution. For example, a host name typically includes crucial information such as service, node, cluster, datacenter, and domain. To automatically add the value for a cluster tag based on incoming host data, you can extract just the cluster data.

When extracting alert fields, use regular expressions (regex) to build value formulas. Regex allows you to precisely identify and capture the relevant portions of the payload, enabling the creation of meaningful and actionable alert notifications.

Note: You can compose text using regular expression (regex) format conventions. Use one or more capture groups with parentheses to extract parts of the input. Capture groups in the regular expression are assigned to alert outputs based on the order in which they appear. The regex must match the entire input, so consider surrounding your regex with `.*` on each end. For example, `(\w+).acme.com.*` captures the host name in a fully qualified domain name. The parser for the regex engine is Perl Compatible Regular Expressions (PCRE) compatible.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Extract alert fields

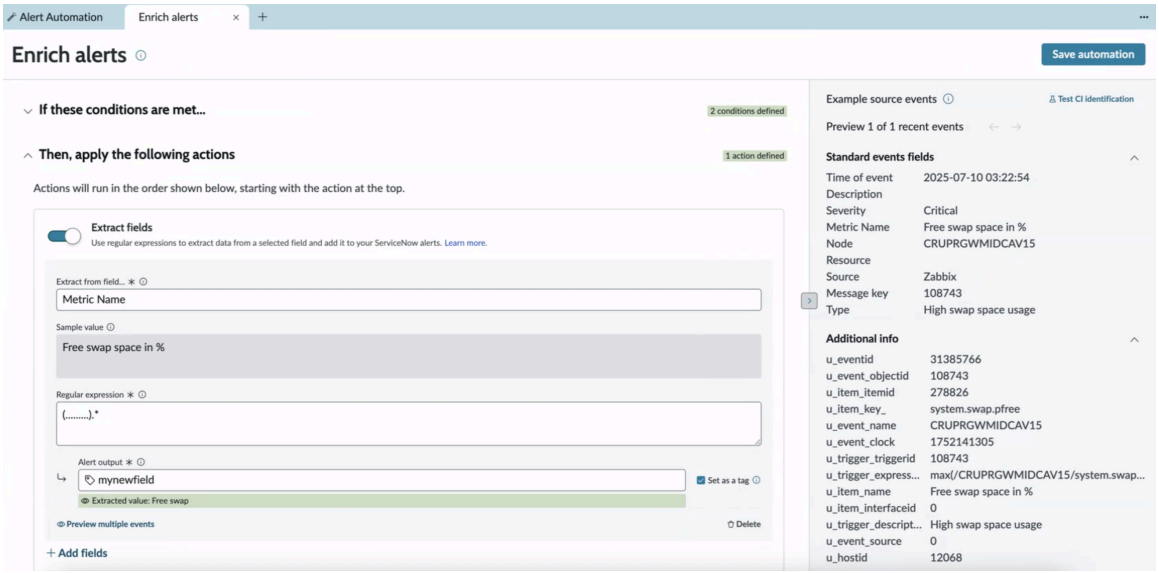

If you use multiple capture groups with parentheses, each extracted value appears in a separate output field.

Extract field for multiple alert outputs

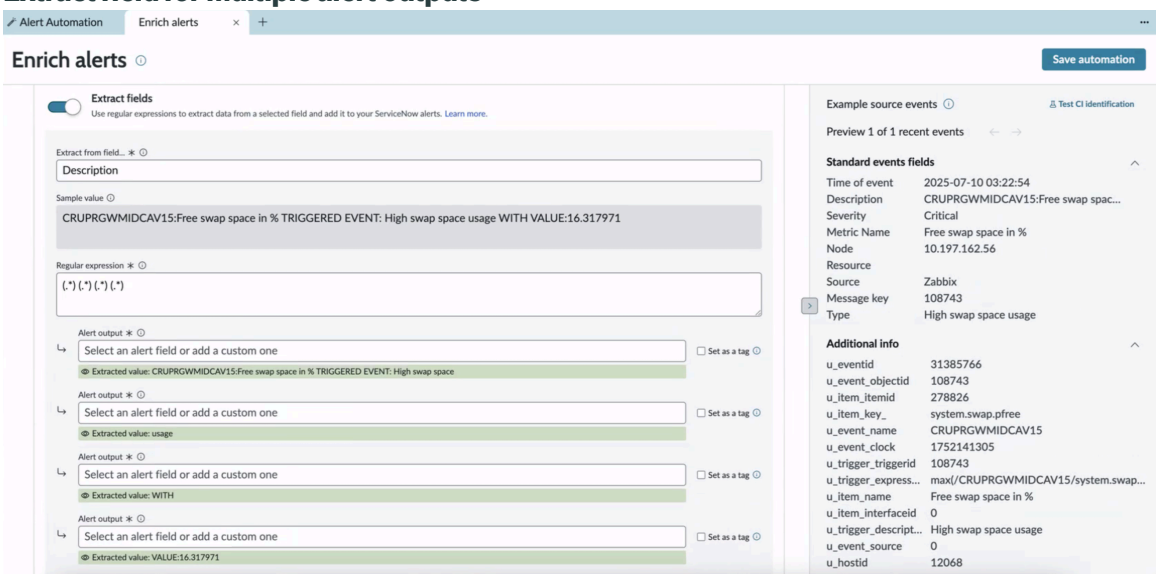

Preview the alert output across example events to verify that the values are extracted as expected by selecting **Preview multiple events**.

Note: This option is available only when example source events are available and matched with the regex filter.

Example: Extracting alert fields

Suppose you need to extract only six letters from a specific field of an event and display them in a new alert field named `mynewfield`. You also want to tag this new alert field for later use in alert grouping. Here's how you can achieve this:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Source input field**: Select the event field from which you want to extract data. In this case, the field is **Metric Name**.
- **Regular expression**: Use a regular expression to extract the specific part you need from the selected field's value. For example, if the **Metric Name** field value contains "Free swap space in %" and you want to extract "Free swap", your regular expression must be (.....).*.
- **Alert output**:
  - Choose an existing alert field, an existing alert tag, or manually enter a new field name. In this case, let's enter a new field name `mynewfield`.
  - Set `mynewfield` as a tag for later use in tag-based grouping. Notice the tag displayed before the field name.

  After applying the regular expression to the selected field's value (in this case the **Metric Name** field value), verify the extracted word displayed below the **Alert output** field. For instance, it should show "Free swap" if the regular expression matches correctly.
- **Preview multiple events**: Previewing multiple events allows you to verify if the regular expression accurately extracts data from a range of example events. This helps determine if any adjustments to the regular expression are needed.

Composing alert fields

When creating an alert output, you can select or manually enter fields, tags, or free text to include. This data can be easily read, filtered, and grouped for better management and understanding of the alerts.

Compose alert fields

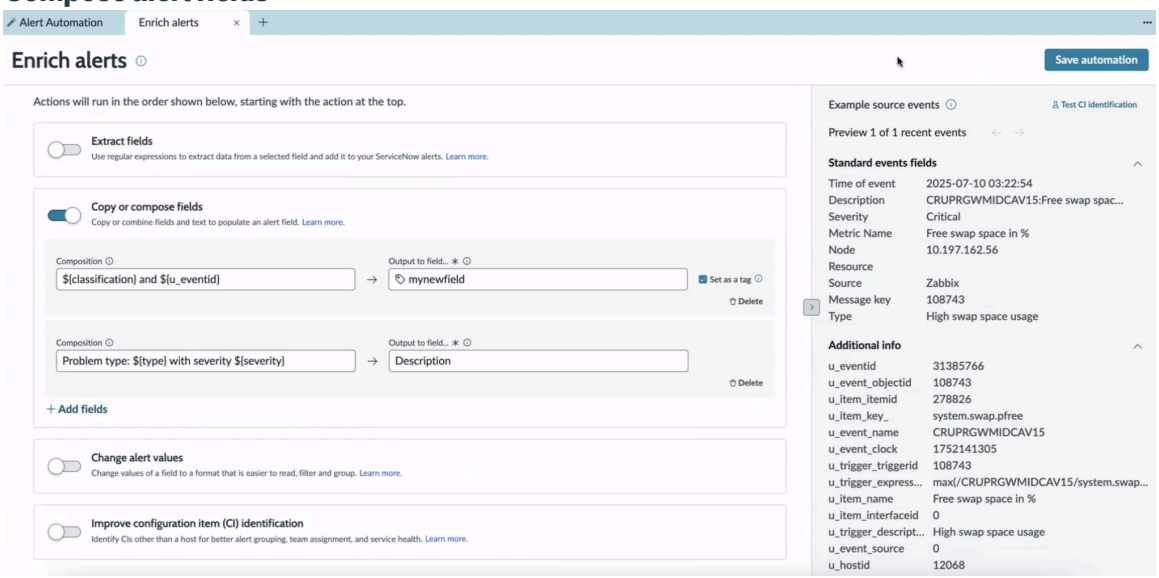

Example: Composing alert fields

Suppose you want to have two input fields that compose data and display it in two different alert output field. In one scenario, you want the source value to come from an existing alert field along with a new field, showing the composed value in a new alert output field. You also plan to tag the new alert output field for later use in tag-based grouping. In the other scenario, you combine existing fields with free text and select the alert output to display in an existing alert field. Here's how you can achieve this:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Scenario 1:
  1. **Source input field**: Select an existing alert field and add the text "and", followed by entering another field **u_eventid**. For example: ${classification} and ${u_eventid}.

     Note that alert fields are displayed in the `${field}` syntax format. You can also select the field name from the drop-down list, and the syntax will be added automatically.

  2. **Alert output**: Enter the name of the new alert field where you want to display the values from the input fields. For example, let's name it `mynewfield`.

     Set `mynewfield` as a tag for later use in tag-based grouping. Notice the tag displayed before the field name.

- Scenario 2:
  1. **Source input field**: Select existing alert fields and include any desired free text for how you want them to appear in the alert output field. For example: Problem type: ${type} with severity ${severity}.

     Alert fields are displayed in the `${field}` syntax format. You can also select the field name from the drop-down box, and the syntax will be added automatically.

  2. **Alert output**: Select an existing alert field where you want to display the values from the input fields. For instance, select **Description**.

Regex preprocessing behavior in Enrich alert automation

Explains how Event Management alert automation preprocesses values before applying regex patterns. Describes why matching behavior differs between pre-populated **Additional Info JSON** fields and free-text sample values.

Preprocessing alert fields

Event Management alert automation preprocesses alert field values before evaluating regex patterns. This preprocessing ensures consistent matching in the backend but can lead to different behavior in the UI depending on how you provide the sample value.

Understanding this behavior helps you build regex patterns that validate correctly and behave as expected at runtime.

How preprocessing works

During regex evaluation, preprocessing occurs only if the value is a JSON value (JSON inside the **Additional Info**) regardless of how you enter the Sample Value. The system automatically preprocesses the Sample Value when you select a pre-populated **Additional Info** field from the **Extract from field** drop-down list.

Before applying regex matching, the backend preprocesses values as follows:

- Removes quote characters (`"`)
- Replaces ": " with "=" (or ":" with "=", if no space is present)
- Converts JSON into `{key=value}`

Free-text Sample Value limitation

When you manually enter a free-text sample value (for example, when the selected field has no data or when testing without matching events):

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The UI does not preprocess the sample value.
- The backend does preprocess the value during regex matching.
- Regex patterns that work in external tools may not appear as matching in the UI.

Example

Current	Preferred
`{"type": "linux_server"}`	`{type=linux_server}`
`{"CI_Type": "server"}`	`{CI_Type=server}`

Impact

A mismatch can occur between the value entered in the **Sample Value** field, and the value the backend evaluates during regex matching.

Workaround

Instead of matching JSON structure, match the preprocessed `{key=value}` format:

- Use `{key=value}` instead of `{"key": "value"}`.
- Avoid matching JSON structure (quotes, colons, whitespace).

Set additional info fields to match CI attribute format

Set additional info fields in alerts to match the field and value format of CI attributes in CI records. This ensures accurate alert-to-CI binding, improving alert tracking and reducing manual effort.

Before you begin

Ensure that you are on the Enrich Alerts page and the **Improve configuration item (CI) identification** toggle switch is enabled.

Role required: evt_mgmt_admin, evt_team_operator, or srm_responder

About this task

The system attempts to match the **Additional Info** fields of the alert with attributes in the CI table. If exactly one matching CI is found, the alert is linked to that CI. For CI identification to be successful, add new fields or ensure that your alert has additional info fields corresponding to column name for CI attributes. For example, use column name **ip_address** instead of the column label **IP Address**. Each matching key name must have matching values. If there are matching fields in your alert that you don't wish to use for binding, copy them to another field then clear the original field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Under **Improve configuration item (CI) identification**, in the **Select which CI class you'd like to identify** field, select the CI class.

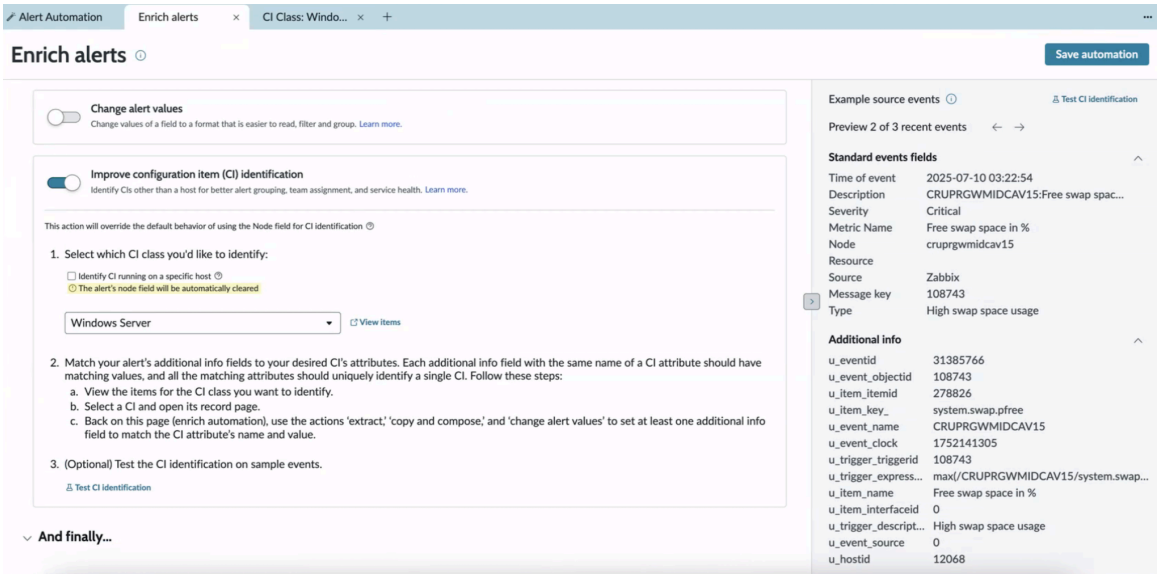

2. Select **View items** next to the **Select which CI class you'd like to identify** field to open the list of CIs of the selected CI class.

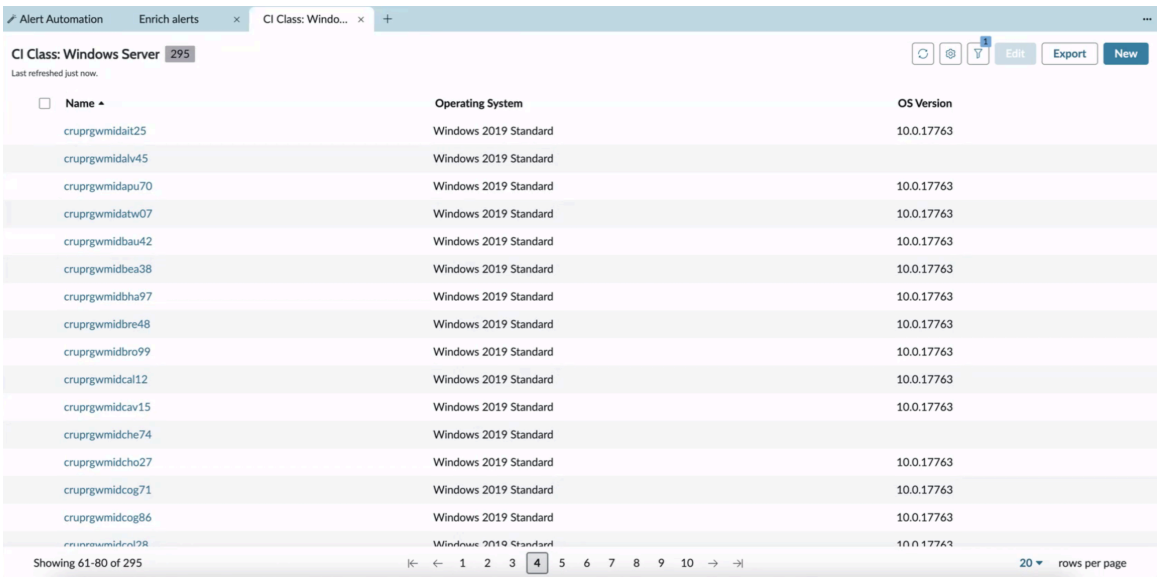

3. Select a CI name to open its details or record page.

   At least one field and its corresponding value from the CI details page must appear in the **Additional info** field of the alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

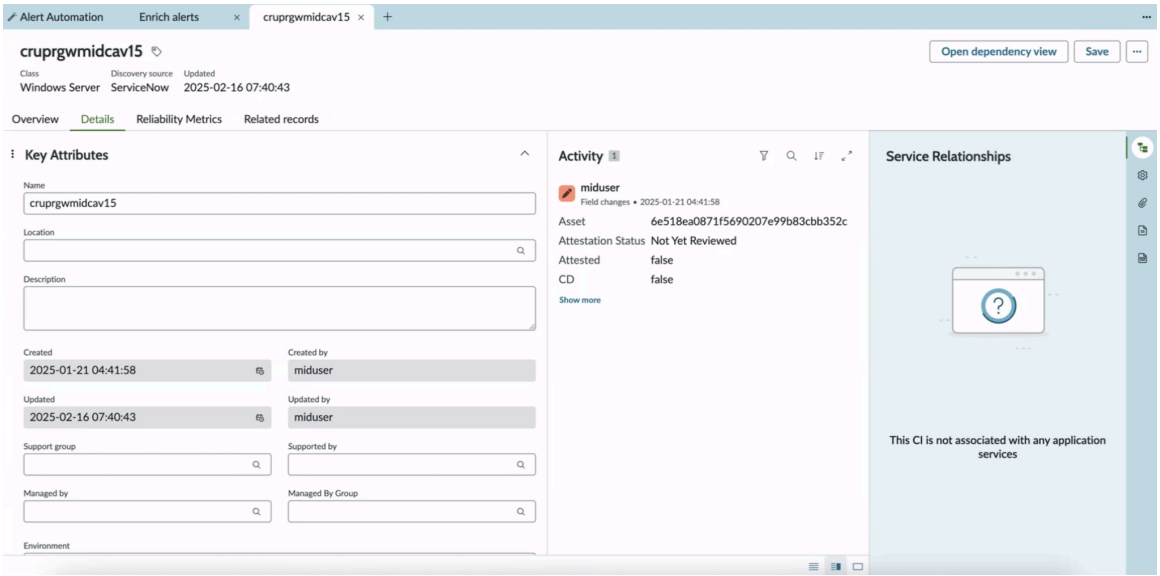

4. Set at least one **Additional info** field in your alert to match the CI attributes in the CI record.

   Use **Extract fields**, **Copy or Compose fields**, and **Change alert values** to precisely match the field and value format.

   Example: Suppose the CI record page has the **name** field with the value **CRUPRGWMIDCAV15**. To include these fields and their values in the **Additional info** field of the alert, set the fields as shown in the image.

   Map the `${u_host_host}` field to a new field called **name**. This field is added to the **Additional info** and is used for CI identification during the binding stage.

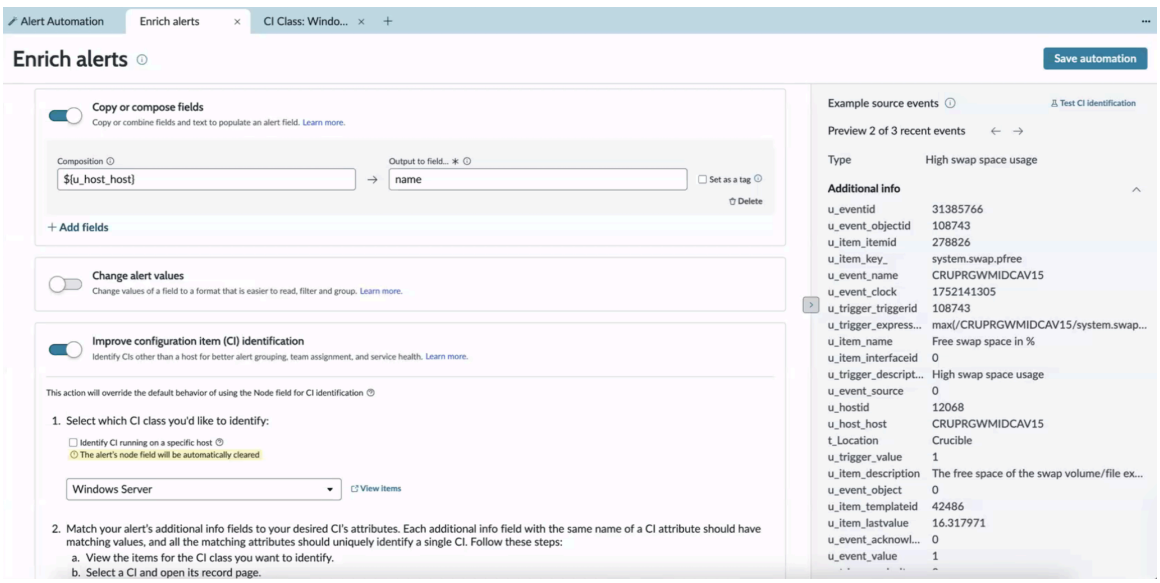

Create Group automation

Grouping automation helps you manage alerts more effectively by collecting similar alerts together. This makes it easier to see patterns, quickly identify issues, and respond efficiently. By organizing alerts in this way, you can reduce alert noise, identify root causes, and assign them to the appropriate teams.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: evt_mgmt_admin, evt_team_operator, or srm_responder

About this task

Grouping of this method is most useful when alerts share common data or tags, such as a node or location. You can use fields or tags populated via an enrich automation. Use of alert tags is the best way to group alerts when your CMDB or service maps are immature. This complements our other grouping algorithms, including alert correlation rules, ML, and text-based grouping. Even if a new alert is matched with multiple groups, it is grouped only with the first match. You can control the priority order of these algorithms via system property. For information on correlation logic order, see Configure alert correlation logic order.

Alert automation also provides a simulation feature allowing you to test how many alert groups would be formed, how many are left ungrouped, and the compression rate. A higher compression rate means your team is more productive and may be able to identify root causes faster. However, consider whether the groups are accurate, operationally correct, and assigned to the right teams. You may adjust the group criteria until you are satisfied with the resulting groups.

For users familiar with the classic Event Management experience, this feature offers an easier interface with improved team support for creating tag-based alert clustering definitions.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the Optimize section, select **Group alerts**

   The Group Alerts page appears. In the Suggested Grouping Automations section, review the suggested automations. If you want to proceed with one, select **Create** on the corresponding automation tile.

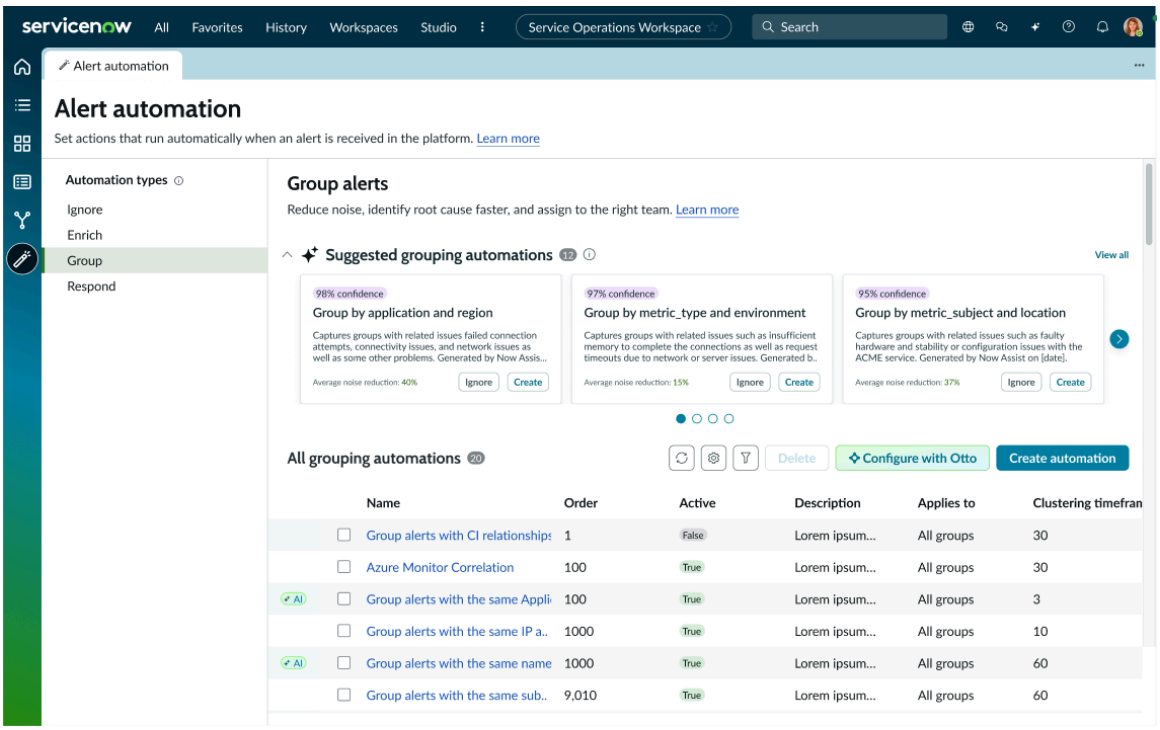

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you want to see all suggested automations—including previously suggested and ignored ones—select **View all** in the Suggested Grouping Automations section.

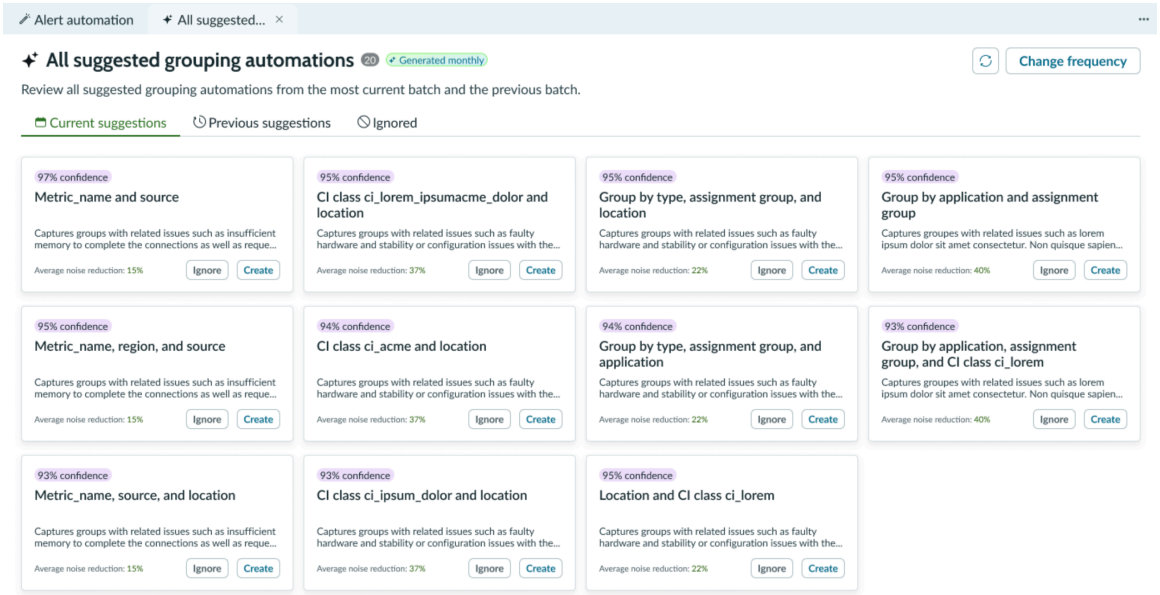

4. Choose how to create the grouping automation.
   - To create the automation using natural language, select **Configure with Otto**. The ServiceNow Otto panel opens with guidance for describing the grouping automation that you want to create. Continue creating the automation in the panel.
   - To create the automation manually, continue to the next step.
5. Select **Create automation**.

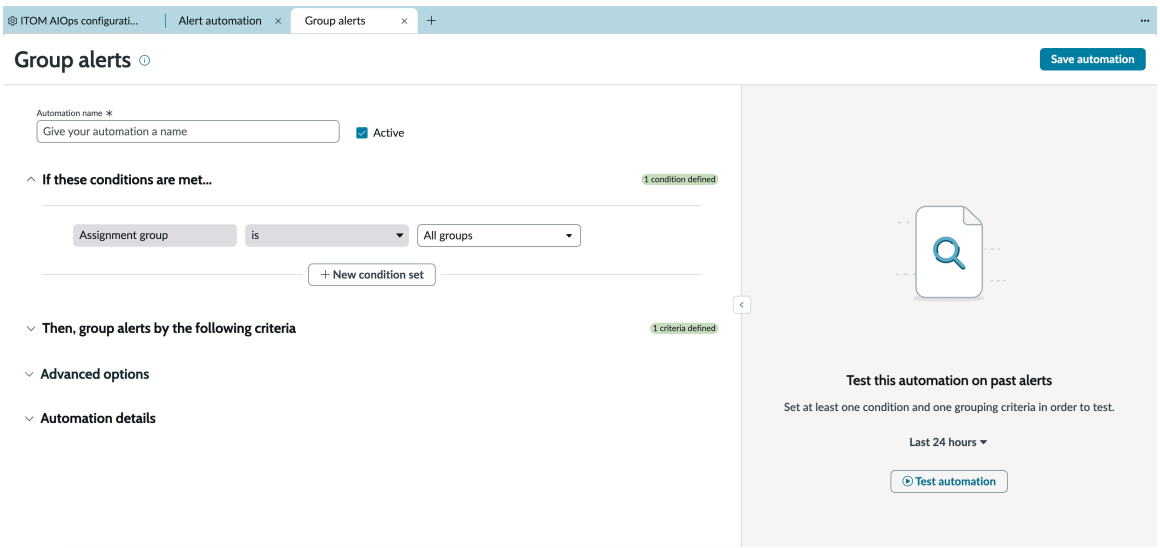

6. In the **Automation name** field, enter the name of the automation for grouping alerts. By default, the **Active** check box is selected.
7. In the **If these conditions are met** section, set up filter criteria to identify the alerts that you want to group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

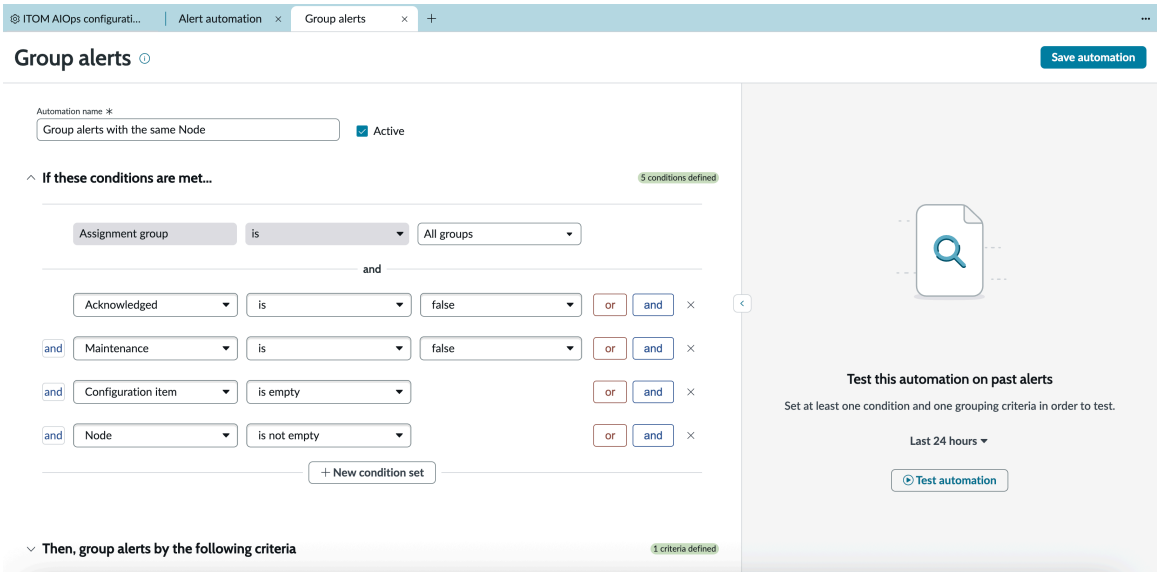

a. From the **Assignment group** field menu, select the assignment group to determine which team's alerts will trigger the automation.

The **Assignment group** represents a specific team responsible for handling certain alerts. By selecting an assignment group, you ensure that only the alerts assigned to that particular team will trigger the automation. This way, the automation is targeted and only activates for relevant alerts associated with the selected team.

Note:

- If you're logged in to the instance with an administrator role (evt_mgmt_admin), all of the assignment groups are available. Additionally, you can select **All groups** to enable generating alerts for any of the available groups.
- If you're an operator, only the group you're a part of is available.
- Only members of the selected group or administrators can update or delete the automation.

b. Set up the conditions by selecting the field, operator, and field value. Then, add more conditions using OR or AND operators. You must add at least one more filter besides the assignment group.

Tip: Select a more specific filter to enhance performance.

To add another set of conditions, select **+ New condition set**. You can also manually add an additional info field if you don't see it in the drop-down list.

8. In the **Then, group alerts by the following criteria** section, perform the following steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

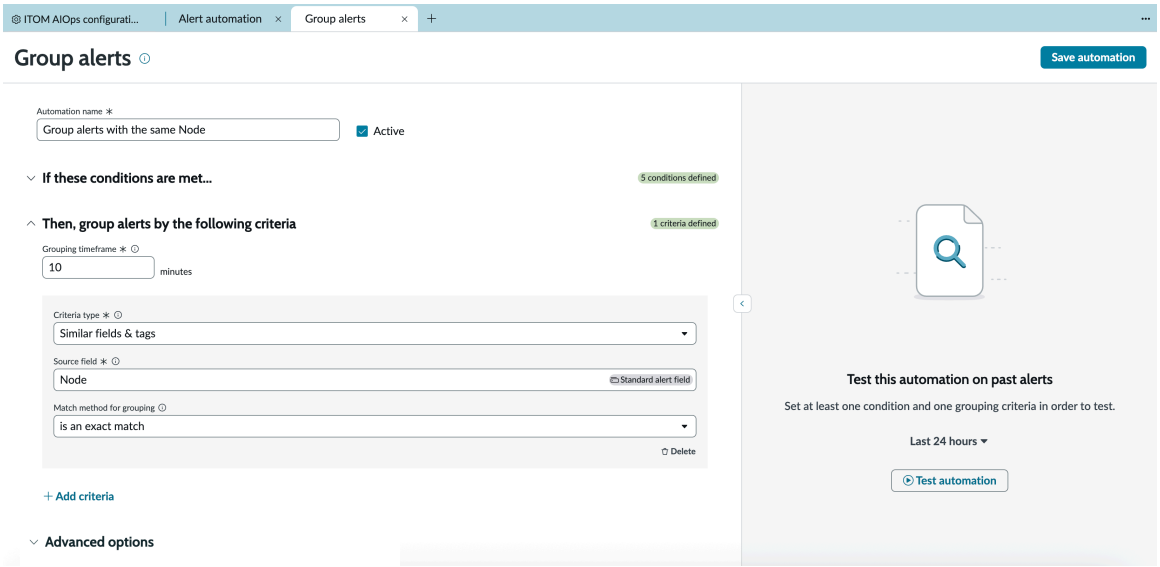

a. In the **Grouping timeframe** field, specify the duration (in minutes) during which alerts are collected and grouped together. This is a fixed window that starts when the first alert arrives.

b. In the **Criteria type** menu, select how you want to group the alerts.

Available options are:

- **Similar fields & tags**: Groups alerts that share common fields or tags.
  - In the **Source field** menu, select the source from which you want the alerts to be grouped.

    Note: You can manually add a field name and select the type of the field. The available options include additional info field, alert tag, and CI tag. This flexibility allows you to customize the information being captured according to your specific needs.

  - In the **Match method for grouping** field, select one of the following options: group alerts based on an exact match, fuzzy match, or pattern match.

    When you select a value for the fuzzy match method in the grouping field, the **Similarity threshold (percentage)** field becomes visible. Alerts are grouped when their similarity is greater than or equal to the specified percentage based on edit distance.

    For example, if you have alerts from USA, CA, and USA, NY, you can group the alerts by country. Set the **Source** field to the location. If the **Match Method for Grouping** is a fuzzy match and the **Similarity threshold (percentage)** is 50%, alerts are grouped when they are at least 50% similar. They share the country "USA" as a common attribute.

  - When you select a value for the pattern match method in the grouping field, the **Pattern matching** field becomes visible. Alerts are grouped when the specified pattern matches. For more information, see Pattern matching.

    Use asterisks (*) in the search string to match any number of characters or a question mark (?) to match any single character. Everything else in the search string matches itself. For example, use "HTTP Error 5??" to match all HTTP 500 errors.

- **Related CIs**: Group alerts according to CI (Configuration Item) relationships. These include parent/child (**VM** > **host**), sibling relationships (for example, multiple VMs on the same

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

host), containment (**process** > **application** > **server**), and applicative flows that describe how components interact within an application. These alerts are connected because the underlying infrastructure components are connected.

- **Related log properties**: Health Log Analytics (HLA) alerts are grouped together because their underlying logs look related based on patterns detected by the HLA algorithms.

  You can create an alert group based on related log properties on HLA alerts combined with alert tags or CMDB logic. When you choose to use related log properties, only log analytics alerts will be considered for grouping. Run the related log properties rule later in the rule order so that other rules—such as CMDB- or service-based grouping—can evaluate and group log analytics alerts first.

- **Impacted service instances**: Alerts are grouped because they affect the same service instance, regardless of topological distance (hops) between the alert CIs.

  You can control which service instance must be considered by specifying relevant impacted service instance. When you select **Impacted service instances**, two options are available: **Any service instance** and **Specific service instance**. If you choose **Specific service instance**, you can select one or more service instances from the list. Alternatively, select the search icon ( ) to open the **Select impacted service instances** pop-up and choose the required service instances.

  By default, the Impacted service instances grouping criterion is based on the impacted service. It can also be configured to use service associations from the [svc_ci_assoc] table. You can control whether to use impacted services or service associations using the *`sa_analytics.use_impacted_services_for_mixed_grouping`* system property. If the property value is `true`, alerts are grouped by impacted services and if the property value is `false`, the alerts are grouped by service associations.

c. To include additional fields for grouping, select **+ Add criteria**.

9. In the **Advanced options** section, you can do the following:

- Enable the toggle switch to set the minimum number of alerts required to form a group. When you enable it, the **Minimum number of alerts in the group** field appears, allowing you to enter the number. The default value is 2, which means grouping occurs only when there are two or more alerts.
- Enable the toggle switch to form group only if at least one alert meets specific conditions you define.

  Note: The threshold and required alert act as prerequisites for forming an alert group. The threshold defines the minimum number of alerts required to create a group. The required alert condition ensures that at least one specific, meaningful alert is present before grouping starts. If that alert is not present, the group is not created — even if there are multiple other alerts. In other words, a group is created only when both conditions are satisfied.

  Example: Imagine a network switch with multiple ports, where alerts can occur either on individual ports or on the switch itself. You want a group to be created only when the switch is affected, not when there is just a port-level issue. In this setup, if alerts occur only on one or more ports, no group is created. A group is created only when an alert occurs on the switch along with one or more alerts on its ports. This ensures that grouping happens only for larger, meaningful issues involving the switch, and not for minor or isolated port problems. This prevents unnecessary alert groups from being created. It helps you focus on more important and meaningful issues and reduces noise caused by minor or isolated alerts.

10. In the **Automation details** section, provide an order and automation description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

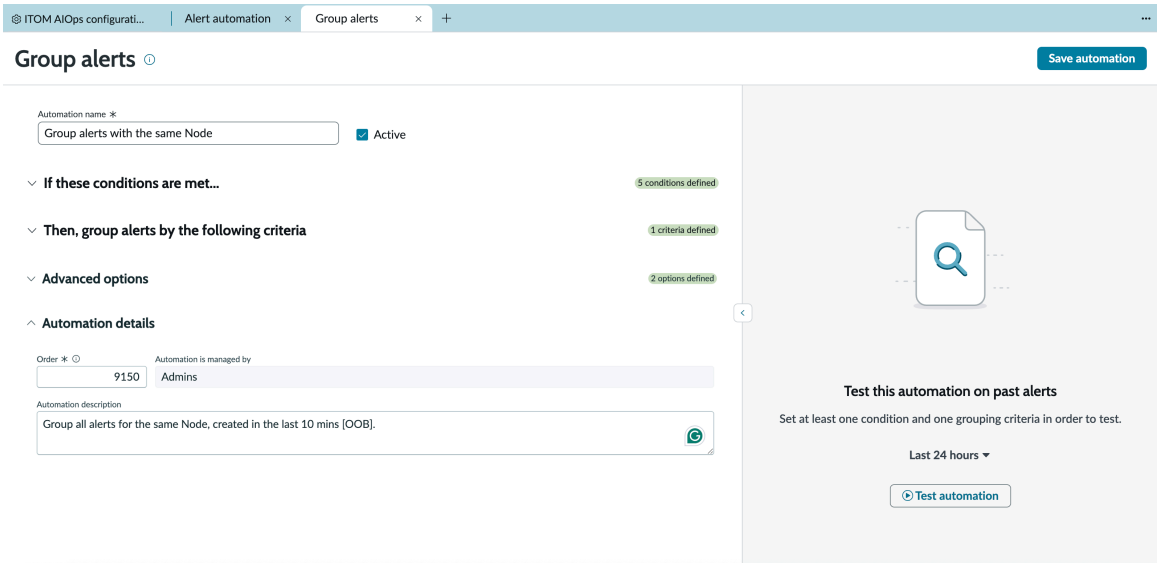

a. In the **Order** field, enter the automation order.

> **Note:** Alerts are grouped based on the first match, executed in order from the lowest to the highest number. The **Automation is managed by** field displays the team or assignment group who owns, edits, and can delete this automation. The assignment group is the same as the one defined in the **If these conditions are met** section.

b. In the **Automation description** field, enter a brief description of the automation.

11. Test the alert grouping by navigating to **Test this automation on past alerts**.

Select the timeframe for the simulation from the drop-down list, select whether you want to consider other grouping types, and then select **Test automation**.

> **Note:** From the **Consider other group types** menu, you can select either **This automation only** or **Consider other group types**. Selecting **This automation only** ignores other alert group types while choosing **Consider other group types** includes all alert grouping automations, such as mixed, automated, and text-based alert grouping.

During the simulation, it shows both the grouped alerts and the ungrouped alerts for the specified timeframe. If any alerts are grouped, you are shown the number of alerts that are grouped. You can select this number to view the grouped alerts. Additionally, selecting an individual alert displays the details of that specific alert. You can also modify any alert grouping conditions or field values and initiate the process again by selecting **Re-run test**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

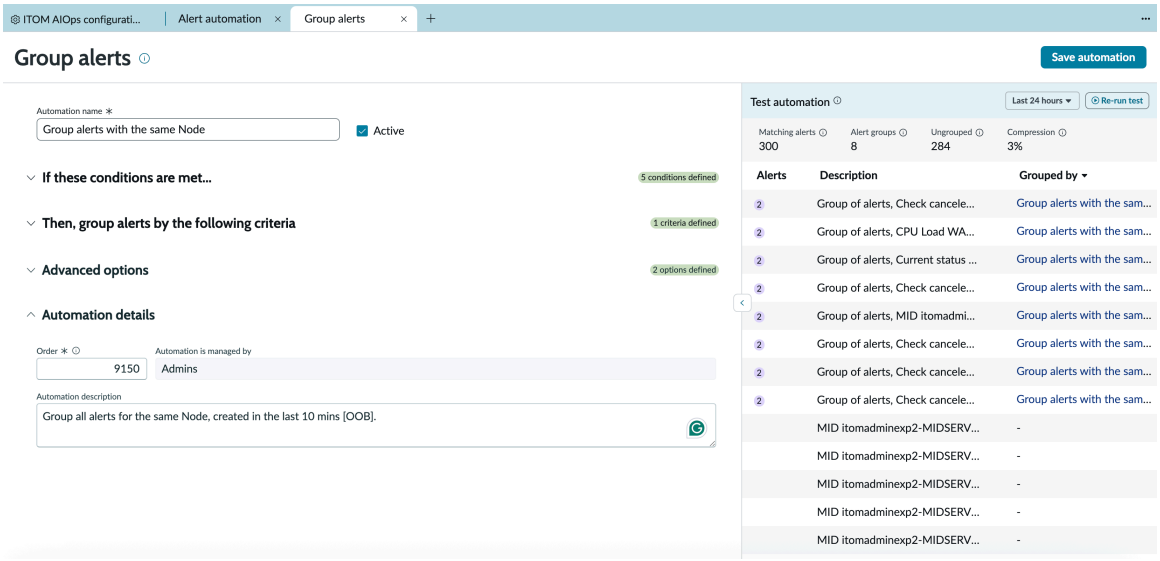

The header of the Test Automation section also displays the following: matching alerts, alert groups, ungrouped alerts, and compression.

- **Matching alerts**: The total number of matching alerts before grouping that match the conditions defined for this automation.
- **Alert groups**: The number of groups containing more than one alerts matching the automation conditions. The smaller number in parentheses represents the number of groups created by this automation.
- **Ungrouped**: The number of alerts matching the automation conditions that remain ungrouped.
- **Compression**: The percentage reduction in the number of total alerts achieved by grouping, calculated as 1 - (Alert groups + Ungrouped) / Total alerts. You can improve the compression rate by grouping your alerts into related problems. The smaller number in parentheses indicates the percentage of matching alerts compressed by this automation.
- When the criteria type is **Related CIs**:
  - The **Open CIs dependency map** link appears in the Test Automation section.
  - The **Configuration Item** field is displayed.

The Test Automation section displays three key columns: **Description**, **Type**, and **Time**. The **Description** column outlines the alert group details. Preceding the **Description** column, a number indicates the total alerts contained within that group. **Type** specifies the category of grouping, such as This grouping automation, Other grouping automation, CMDB group, or Single alert, among others. Selecting **Other grouping automation** directs you to the corresponding automation page that generated the group. Additionally, the **Time** column shows when the test was conducted.

> **Important:** You can run the simulation of alerts on your test as well as production instance. The test automation simulates the automation using up to 200 recent alerts that match the specified conditions. It only considers groups where all alerts meet the conditions for this automation, although additional grouping algorithms may be applied during actual runtime.

12. Select **Save automation**.
    A notification appears when the automation is successfully saved. Otherwise, an error message is displayed. The group automation that you created appears on the Group alerts page where you can view, edit, or delete the existing automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can escalate alerts needing quicker responses from teams or individuals by implementing Create Respond automation.

Create Respond automation

Respond to alerts automatically by notifying appropriate stakeholders, escalating them as needed based on severity and type, or other executing response actions. This process ensures that alerts are managed promptly and effectively.

Before you begin

Role required: evt_mgmt_admin, evt_team_operator, or srm_responder

About this task

For users familiar with the classic Event Management experience, this provides an easy interface with enhanced team support for creating response automations or alert management rules. Alert management rules and response automations are configurable automated actions triggered by alerts matching specific conditions. These actions help IT teams resolve issues faster with minimal manual effort. These advanced features are currently only available to administrators and other personas mentioned in the role required section.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the Optimize section, select **Respond alerts**.
4. Select **Create automation**.
   The Respond page appears.

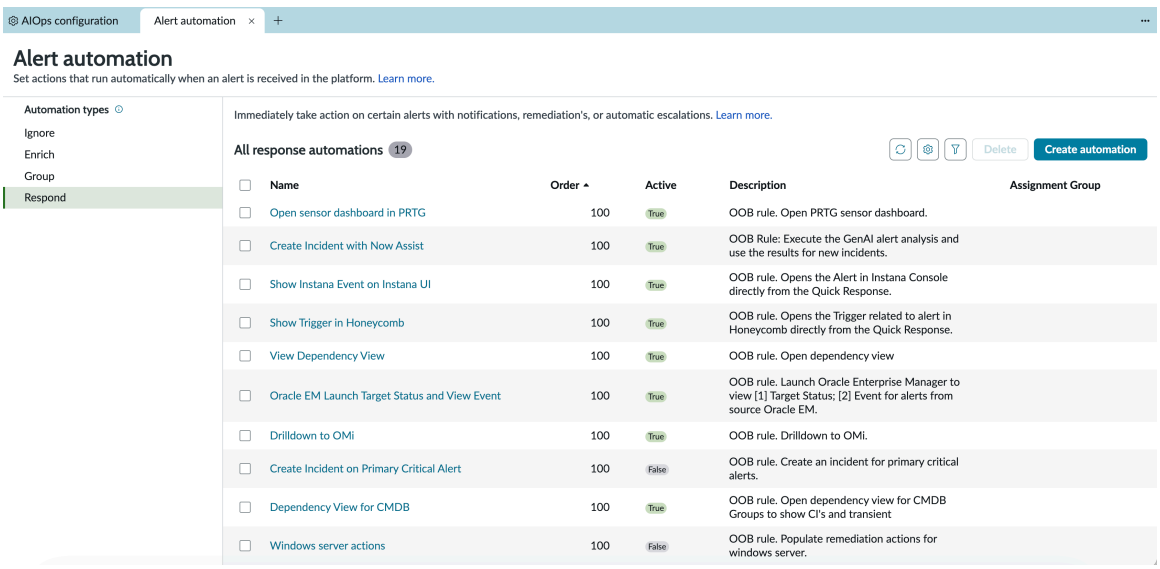

5. In the **Automation name** field, enter the name of the automation.
   By default, the **Active** check box is selected.
6. In the **If these conditions are met** section, create filter criteria to identify the alerts you want to capture.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: You can limit respond automations to primary alerts to prevent secondary alerts from being overwhelmed by unnecessary noise. Ensure these automations are applied only to alerts associated with CIs that aren't currently undergoing maintenance.

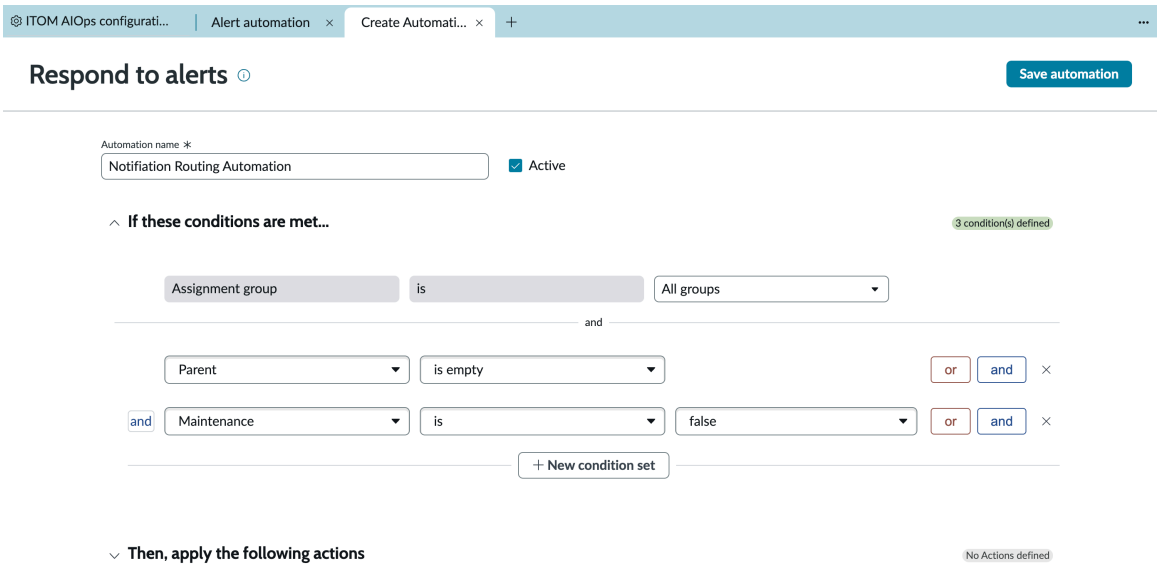

a. From the **Assignment group** field menu, select the assignment group to determine which team's alerts will trigger the automation.

The **Assignment group** represents a specific team responsible for handling certain alerts. By selecting an assignment group, you ensure that only the alerts assigned to that particular team will trigger the automation. This way, the automation is targeted and only activates for relevant alerts associated with the selected team.

Note:

- If you're logged in to the instance with an administrator role (evt_mgmt_admin), all of the assignment groups are available. Additionally, you can select **All groups** to enable generating alerts for any of the available groups.
- If you're an operator, only the group you're a part of is available.
- Only members of the selected group or administrators can update or delete the automation.

b. For the **Parent** field value, select whether the alert for which you want a response is grouped under a parent alert.

c. For the **Maintenance** field value, select whether the alert for which you want a response is under maintenance.

d. To add another set of conditions, select **+ New condition set**. You can also manually add an additional info field if you don't see it in the drop-down list.

Set up the conditions by selecting the field, operator, and field value. Then, add more conditions using OR or AND operators.

7. In the **Then, apply the following actions** section, select one or more of the two automation actions that can be triggered by the automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

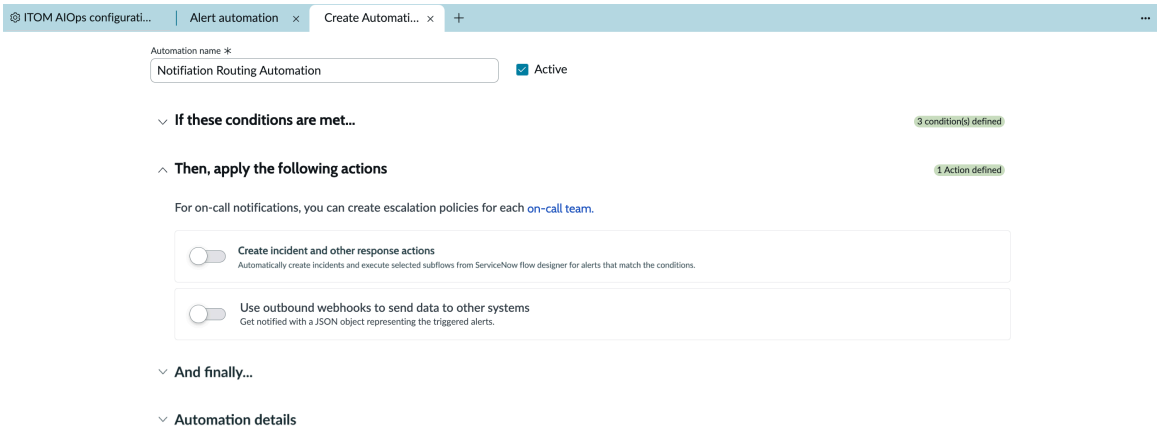

Action	Description
Create incident and other response actions	Creates an incident for the alerts that match the specified filter. For details on mapping the alert fields to the incident fields, see Alert field mapping on the Respond page. **Note:** If you want incident fields to be automatically populated with default values, consider using **Create Incident (advanced)** in the **Response subflows** field. **Create Incident (advanced)** pre-populates incident fields even when you don't explicitly customize them, which helps ensure consistent incident data and reduces configuration effort. The **Create incident and other response actions** option also executes selected subflows from Workflow Studio for alerts that match the specified condition. For example, you may select out-of-the-box subflows to restart services. You can also create your own custom subflows, leveraging hundreds of integrations available in Integration Hub. 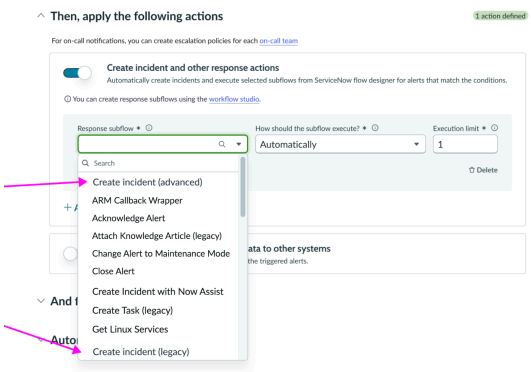

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Description
	a. In the **Response subflow** field, select the subflow. Selecting the Information icon ( ) takes you to the workflow of the subflow. For example, if you select the Information icon ( ) next to Close Alert, the Close Alert workflow opens. **Close Alerts workflow** 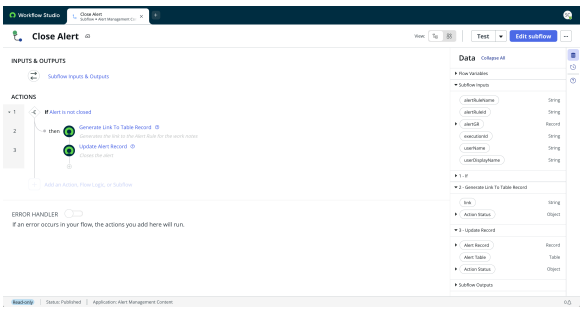 **b.** In the **How should the subflow execute?** field, select **Automatically**, **Manually** or **Automatically and Manually**. **Note:** If you select a subflow with an execution type of **Manually** or **Automatically and Manually**, and then save the automation, it appears in the Express List. 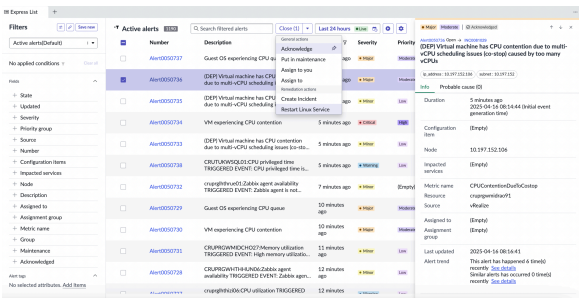 **c.** In the **Execution limit** field, specify the maximum number of times the subflow can run. After reaching the limit, the subflow will not execute again. **d.** In the **Response subflow** field, when you select **Create Incident (advanced)**, two check boxes appear. ▪ **Customize incident fields** ▪ **Add wait time**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Description
	When you select the **Customize incident fields** check box, a mapping section ap pears with the following fields: ▪ **Alert input value**: Generates the value us ing AI or takes it from the alert. ▪ **Incident field name**: Maps the selected alert value to the specified incident field. **Note:** At least one field mapping is required when you select the **Cus tomize incident fields** check box. **Cre ate Incident (advanced)** automatical ly fills incident fields with default val ues, even if you don't customize them. **e.** Select the **Add wait time** check box to de lay incident creation and allow alerts time to auto-resolve or downgrade. You can define conditions that are eval uated after the wait period to determine whether an incident still needs to be cre ated. This helps prevent unnecessary inci dents for alerts that resolve on their own or reduce in severity. To add another response action, select **+ Add response action**.
Use outbound webhooks to send data to oth er systems	Sends notifications to other systems using outbound webhooks. For example, this capa bility can be used to send a chat notification, create a case, or trigger an external runbook to remediate the issue. For details on configuring an outbound web hook, see Outbound webhook parameters. To create an additional property, select **+ Add property**. The property can be any custom key-value pair that you want to include in the webhook payload, such as "Content-Type: ap plication/json" or "Authorization: Bearer <to ken>".

Note: If you don't select at least one action, the automation is deactivated.

8. In the **And finally** section, to continue running other response automations with same filter conditions after this automation is executed, select **Run other response automations**. If you select **Don't run response automations**, additional automations of this type will stop running after this automation is executed once. If the automation is managed by an

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

administrator, it will stop running administrator-owned automations but will continue to run automations owned by other assignment groups.

9. In the **Automation details** section, provide an order and automation description.

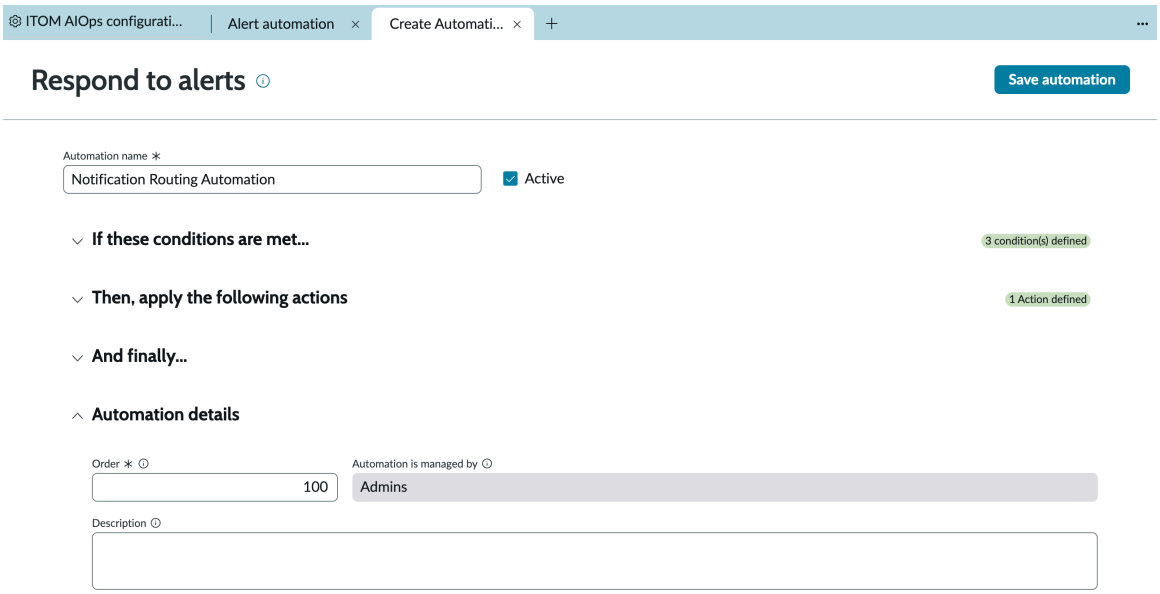

For information on the Automation details fields, see #unique_2282.

10. Select **Save automation**.

A notification appears when the automation is successfully saved. Otherwise, an error message is displayed.

The response automation that you created appears on the All response automations page where you can view, edit, or delete the existing automations.

Delete an automation

Delete an automation to permanently remove an automated alert handling rule from your system.

Before you begin

Role required: evt_mgmt_admin, evt_team_operator, or srm_responder

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the Optimize section, select the automation type that you want to delete, such as Ignore, Enrich, Group, or Respond.

   The page for that automation type opens, displaying all the existing automations.
4. To delete an automation, select the check box next to its name and then select **Delete**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

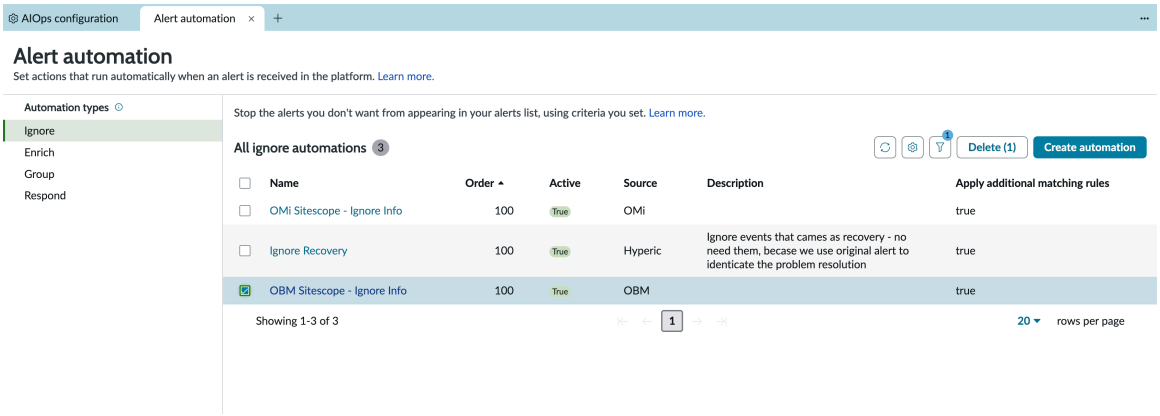

5. In the confirmation window, select **Delete**.
   The alert is permanently deleted.

Express List in the Service Operations Workspace for ITOM

Express List helps you identify health issues across the datacenter on the Service Operations Workspace. It provides quick information on alerts so you can monitor systems and services, resolve alerts, evaluate the alert impact, track issues, and report incidents more efficiently.

https://player.vimeo.com/video/1199445396?h=c7b7151862&badge=0&autopause=0&player_id=0&app_id=58479

The Express List pane sortable alert list reduces the number of clicks necessary to access alert information. Selecting the check box of an alert opens a preview panel where you can view data that helps for prioritization, impact realization, and root cause analysis. You can easily modify the Express List pane to narrow down the display. Use the provided fields list to display additional alert information and to filter out or show matching alerts. You can also modify the displayed time range.

Note:

Exporting alerts is not supported from the Express List. To export alerts, open the list view in Service Operations Workspace.

The following image shows a sample Express List pane.

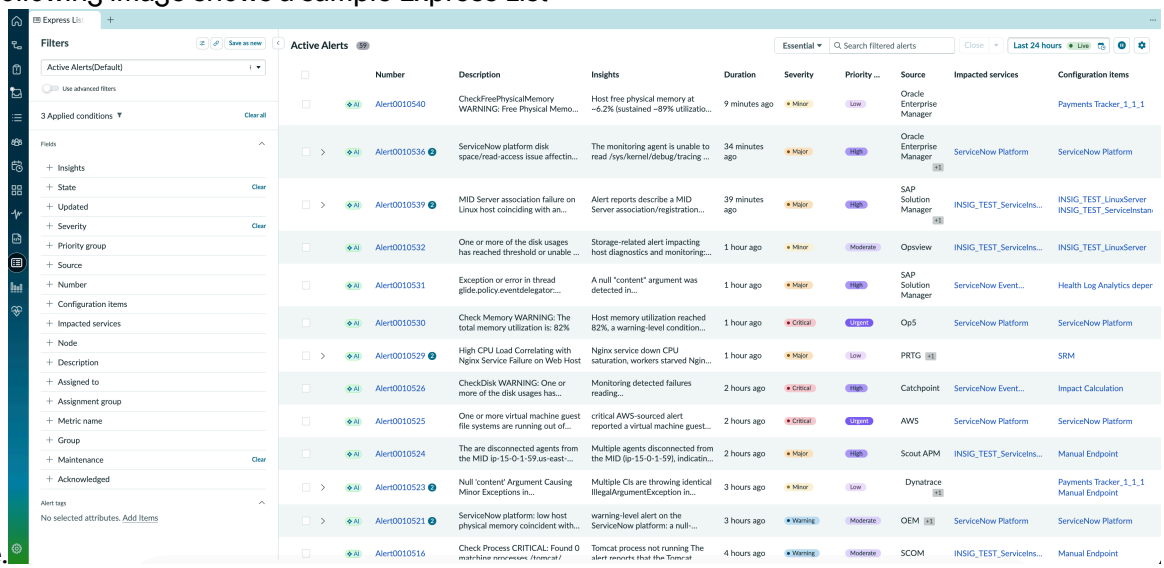

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Express List benefits

The Express List pane enables operators to more efficiently perform the following actions:

- **Monitoring systems or application services**. Monitoring the performance and capability of computer systems to see any discrepancies.
- **Resolving alerts**. Working on discrepancies to troubleshoot issues and promote the efficient functioning of IT infrastructure.
- **Tracking issues**. Tracking and documenting defects and resolutions in detail and reporting incidents.
- **Reporting incidents**. Escalating complex issues to management, other IT resources, third parties or vendors.

Alert information displayed in the Express List preview panel

The Express List preview panel displays additional alert information on the Express List pane. This information enables you to conduct an assessment that helps in resolving incidents without having to open multiple alert forms.

The **Info** tab on the preview panel displays the following information for a selected alert:

- Alert number (link)
- Status
- Duration
- Configuration items (with links)
- Metric name
- Node
- Resource
- Source
- Assigned to
- Assignment group
- Last Updated
- Impacted services

  For more information, see View data on impacted services on the preview panel in Express List.

- Additional info
- Custom field

You can view additional alert information by selecting any of the linked items.

The **Alerts in group** tab on the preview panel displays alert information for a selected alert group. Information for each alert is displayed in a tile with additional options for each alert in the group.

Each tile displays the following alert information:

- Alert number (link)
- Status
- Severity
- Configuration items

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Duration
- Metric name

Viewing alerts: Essential and Extended modes

You can view alerts in two modes:

- **Essential** mode displays primary and single alerts only. Expand an alert group to view its secondary alerts.
- **Extended** mode displays all alert types, including secondary alerts.

Note: You can switch between **Essential** and **Extended** modes without entering search text.

Customizing the alert display time range

You can determine the time range of the displayed alerts. The default time range is the last 24 hours. The list continues to update with new alerts until you select the pause icon ( ). Selecting the current range setting in the upper right displays a dialog box with date and time range options:

- All time - The last 90 days
- Last 24 hours
- Last 2 days
- Last week
- Last 12 hours
- Last hour
- Last 15 minutes
- Custom - Enables you to select a date and time range from the pop-up calendar.

The **Default Time Range** field is introduced on the Express List view in the following versions:

- Service Operations Workspace Express List 26.6.1
- Service Operations Workspace Express List App 26.4.0
- Service Operations Workspace ITOM Apps 26.9.0

Administrators can configure the default time range for a predefined Express List view by using the **Default Time Range** field. The default value is **Last 24 hours**.

The *evt_mgmt.express_list.all_time_days* system property controls the number of days included in the **All time** range. It does not affect the default time range.

Related topics

Roles used by Express List

AI in Express List

ServiceNow Otto for ITOM provides generative AI skills and agentic workflows in Express List that help operators triage alerts, investigate incidents, and respond faster.

Express List includes several generative AI skills and agentic workflows. All features require ServiceNow Otto for IT Operations Management (ITOM) to be installed. For a summary of all AI features across ITOM AIOps, see Install ITOM using ServiceNow Otto for Setup.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

AI features in Express List

AI feature	Description	Use case	Resources
Alert analysis skill	Uses generative AI to generate a human-readable brief of an alert and additional technical information to support investigation. Analyses are provided in English regardless of the language used in the alert description.	An operator opens an alert in Express List and wants a quick summary of what the alert means and what to investigate next.	View an alert analysis by ServiceNow Otto in Express List
Alert group analysis skill	Uses generative AI to generate a simplified, human-readable description of an alert group and technical information to support investigation.	An operator selects an alert group in Express List and wants to understand the group's overall significance before deciding how to respond.	View an alert group analysis by ServiceNow Otto in Express List
Alert group description generation skill	Uses generative AI to generate a comprehensive description of an alert group that encompasses all alerts within the group. The generated description replaces the original alert group description.	An operator wants to replace a system-generated alert group description with a meaningful, AI-generated summary that reflects the full scope of the group.	Generate an alert group description in Express List using ServiceNow Otto
Past incident analysis skill	Queries historical records to find related past incidents and analyzes their frequency, criticality, work notes, and resolution. Presents a summary of the most relevant incidents in the ServiceNow Otto panel, including resolution strategies and contact details for teams who resolved similar incidents.	An operator is investigating an alert and wants to know whether similar issues have occurred before and how they were resolved.	• Accelerate alert resolution with past incident analysis • Generate a ServiceNow Otto summary of past related incidents
AI-generated incident creation skill	Creates an incident with a human-readable, AI-generated description directly from an alert in Express List.	An operator determines that an alert requires an incident and wants the incident description to be generated automatically	Create an incident from an alert with ServiceNow Otto in Express List

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI features in Express List (continued)

AI feature	Description	Use case	Resources
		rather than written manually.	
Analyze alert impact agentic workflow	Uses observability AI agents to investigate an alert and surface information about its severity, impact, ownership, and potential root causes. Integrates with observability tools such as Dynatrace, Kentik, and New Relic.	An operator selects an alert from an observability-connected source and wants to understand its blast radius, identify responsible teams, and find potential root causes before escalating.	• Analyze alert impact agentic workflow • Analyze alert impact in the ServiceNow Otto panel
Analyze potential impact agentic workflow	Analyzes how a change request might affect relevant servers and services. The agent verifies prerequisites, retrieves the change request, selects up to 10 affected servers, identifies matches with suggested services, and displays an impact analysis in the ServiceNow Otto panel. The analysis is also saved to the change request work notes.	An operator or change manager wants to assess the risk of a change request before approving it, without manually reviewing all affected CIs.	• Analyze potential impact agentic workflow • Assess a change request with the Analyze potential impact workflow
Manage alerts autonomously agentic workflow	Provides a unified AI-driven process that automates alert triage, impact analysis, and root cause investigation. Generates reports, summarizes key insights and possible next steps, and stores structured findings in the alert record. Also analyzes Health Log Analytics alerts in an alert group and can classify alerts as proactive or reactive. **Note:** Autonomous workflow is a simplified AI process elaborated in the AI Specialist. You can now activate and manage AI Specialist, controlling the groups it targets and the specific alert type it handles. Activating the AI Specialist automatically disables the Autonomous workflow.	An operations team wants alerts to be triaged, analyzed, and summarized automatically as they arrive, with AI insights available in Express List without manual intervention.	• Manage alerts autonomously agentic workflow • Configure the manage alerts autonomously agentic workflow • Review AI-generated alert insights in Express List

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure ServiceNow Otto for ITOM

Enable Event Management users to view alert analyses that ServiceNow Otto for ITOM creates using generative AI.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

Configure the manage alerts autonomously agentic workflow

Configure an alert management rule to operate the manage alerts autonomously agentic workflow manually or automatically.

Before you begin

- Install ServiceNow Otto for IT Operations Management (ITOM). For more information about installing ServiceNow Otto plugins, see Install plugins for ServiceNow Otto.
- [Optional] Additional configuration is required for some skills used in the agentic workflow:
  - **Observability skills**: Configure the relevant observability skills to surface data from observability tools integrated with Event Management. See Configure observability agents for ServiceNow Otto.
  - **Log analytics skills**: Install and configure Health Log Analytics to enable analysis of log analytics alerts. See Configuring Health Log Analytics.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with your applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

Role required: evt_mgmt_admin, evt_mgmt_operator

About this task

Configure the autonomous workflow alert management rule to operate the manage alerts autonomously agentic workflow manually or automatically. For more information about alert management rules, see Alert management rules for resolving alerts.

When the workflow operates automatically, alerts are addressed as they're created and AI insight information is displayed in Express List.

When the workflow operates manually, users must manually generate AI insights. For more information, see Review AI-generated alert insights in Express List.

For more information about the manage alerts autonomously agentic workflow, see Manage alerts autonomously agentic workflow.

Note:

If you change the alert management rule for the manage alert autonomously workflow, you must update the **sn_aiops_ai_agents.autonomous_alert_rule_sys_id** property, which points to the alert management rule.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.

3. On the ITOM AIOps configuration center page, under the Optimize section, select **Respond to alerts**.
   The Respond page appears.

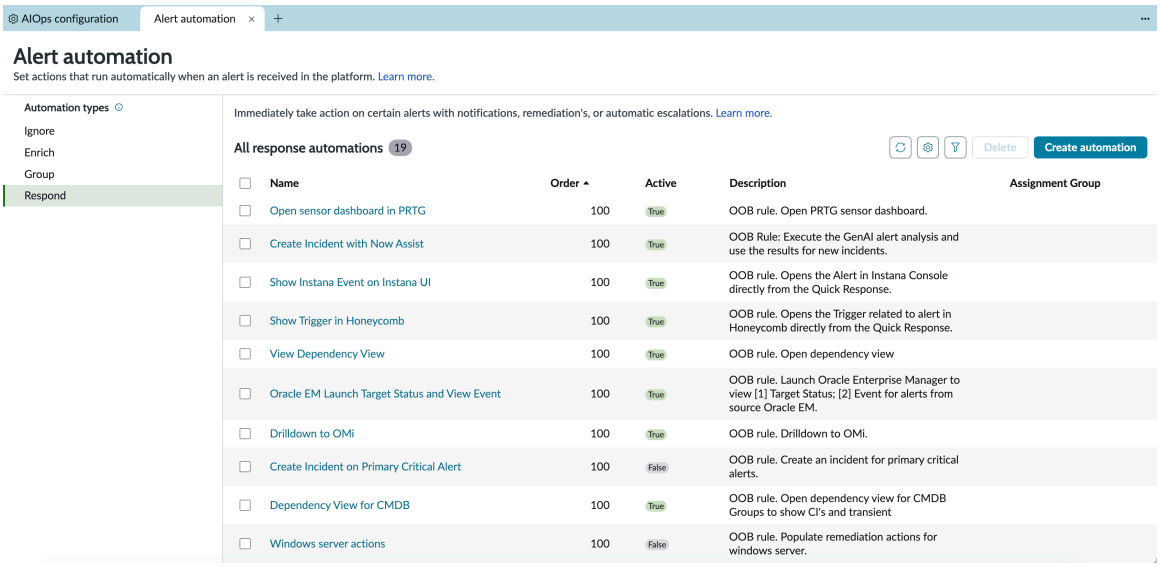

4. In the Name column, select **Autonomous workflow**.
   If the Edit automation pop-up is displayed, select **Open**. The Alert Management Rule form opens to the Alert info tab.

   Note: The **Active** check box is selected when the rule is enabled. To disable the rule and the manage alerts autonomously agentic workflow, select the check box.

5. Configure the execution information for the alert by selecting the **Actions** tab.
6. In the Execution column, double-click the cell for the **Autonomous workflow**.
7. Select the drop-down arrow to choose alert execution, and then select the green check mark to save.

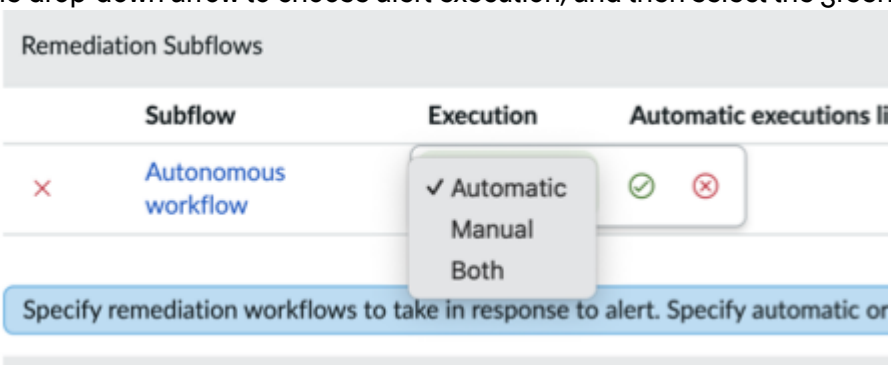

What to do next

To learn more about generating AI insights with the manage alerts autonomously agentic workflow, see Review AI-generated alert insights in Express List.

Set up your workspace

Learn how to set up your individual Express List workspace to suit your needs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Identify an assigned Express List view

Identify Express List views assigned to you by your Event Management admin to make sure that you focus on specific services, priorities, or alerts.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. **Optional:** In the Filters pane, turn on the **Use advanced filters** option to access advanced filters assigned to you.
4. From the drop-down list in the Filters pane, select an Express List view assigned to you. The default view is marked as Default.

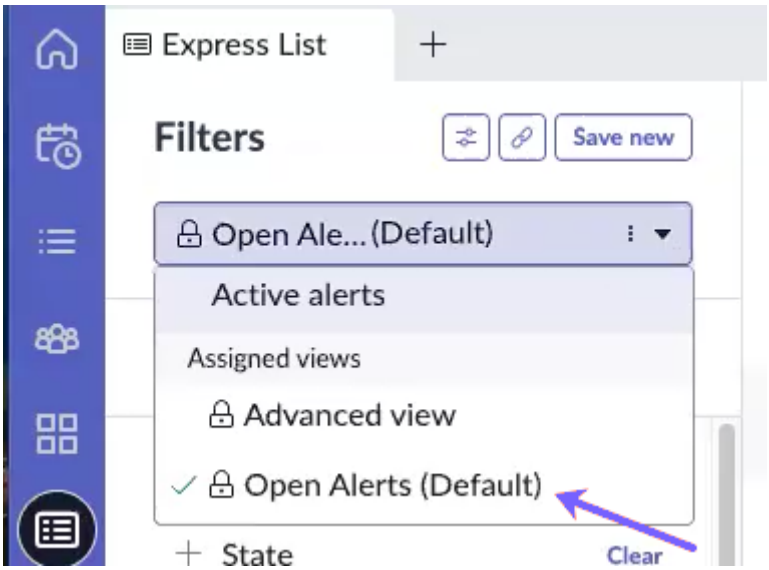

Note: You can't override the filters in views assigned to you. However, you can modify, save, and share these views locally.

The Filters pane lists the predefined filters. The Express List table displays the columns configured for the selected view.

Related topics

Configuring Express List views for users and user groups

Add or modify Express List columns

Add columns to the Event Management Express List display to focus the list of alerts that is displayed or modify the existing columns that are displayed.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Edit the default columns in the pane.
   a. Select the gear icon ( ).
   b. In the **Available columns** list, select the desired columns.
   c. In the **Selected columns** list, drag selected columns up or down to reorder the fields.

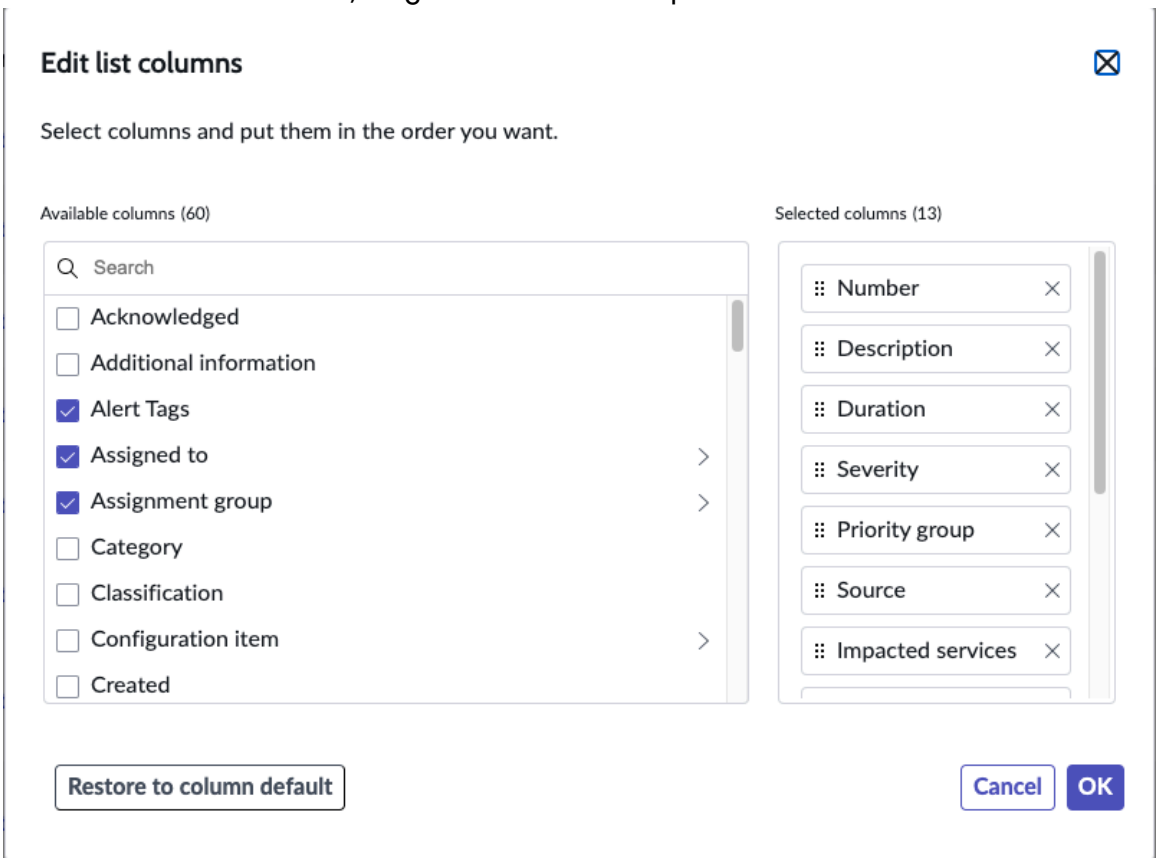

   d. **Optional:** Select **Restore to column default** to return the columns to the default state.
   e. Select **OK**.
4. **Optional:** Modify the column width, order, and sort direction directly in the Express List pane.

Option	Action
Column width	Point to the right side of the column until the bi-directional arrow appears and drag the col umn to the desired width.
Column order	Select the column header and drag it to the desired location.
Column sort	Select the column header and then select the arrow. The arrow direction indicates the sort order. Each selection of the arrow changes the sort direction.

What to do next

- Filter the Express List display by adding or removing fields
- Save a custom filter in Express List

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Filtering the alert display in the Express List pane

Filters help you create different views of categorized alerts, creating a targeted list of alerts to prioritize and focus on, saving time and minimizing distractions.

The fields in the interactive Filters pane in Express List help you create different views of alerts. The default filters are determined by the selected Express List view.

You can interactively view the result of adding or removing any of the fields. The numbers next to each attribute indicate the number of alerts that are returned.

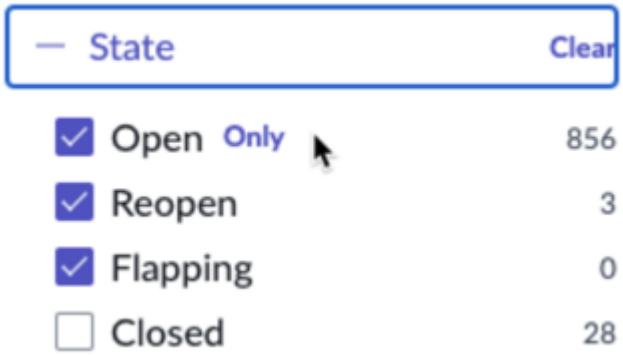

Alternatively, you can turn on the **Use advanced filters** option to set complex filters using a condition builder. The and/or logic provides more flexibility in filtering the alerts.

In the Filters pane, you can perform the following actions:

- Filter by string, number, or date attributes. For more information, see Filter the Express List display using attributes.
- Add filters to or remove filters from the default filter list to focus on the applied filters. The most common filters are State, Priority, CI, and Impacted services. For more information, see Filter the Express List display by adding or removing fields
- Add or remove alert tags. Tags enable you to identify and group alerts easily without the need for discovery or accessing the CMDB. For more information, see Categorize alerts displayed in Express List
- Display or hide alerts that match a selected alert. For more information, see Filter out or show matching alerts.
- Save, delete, assign as the default, or share the filter layout. For more information, see Save a custom filter in Express List.

Filter the Express List display using attributes

Filter the Express List display by text, number, and date attributes using the interactive Express List filter panel. Create a targeted list of alerts to prioritize and focus on.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the primary navigation, select the Express List icon ( ).
3. From the Filters panel on the left side, select whether to filter to text, number, or date attributes depending on your filter type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Filter attribute type	Action
Text	Find elements by name by typing the first few characters to display the elements that start with those characters. Then select the values that you want to view in the Express List pan 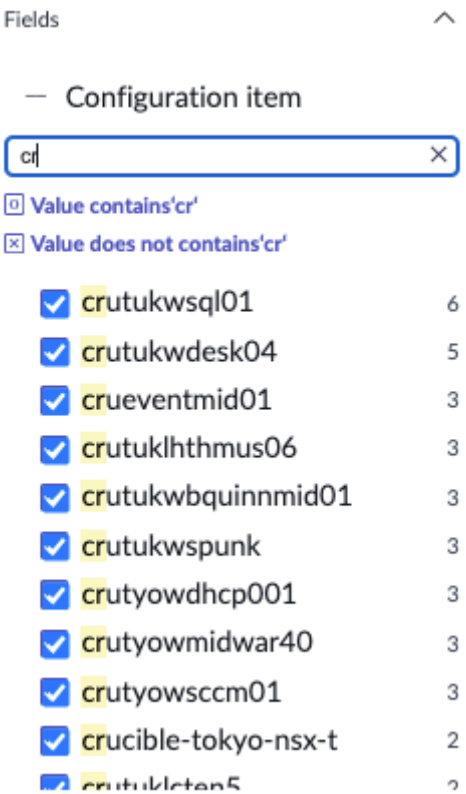 el.
Number of impacted services	Enter a condition to view alerts according to the number of alerts on the selected fil ter. For example, to display alerts that have fewer than 10 impacted services, you would select the **less than** condition and enter 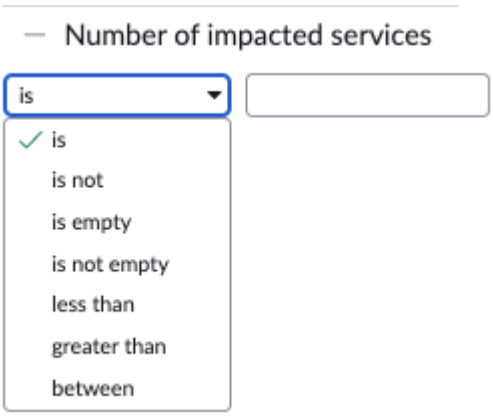 10.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Filter attribute type	Action
Date	View alerts according to the selected date and time range in the pop-up calendar in the Express List panel. 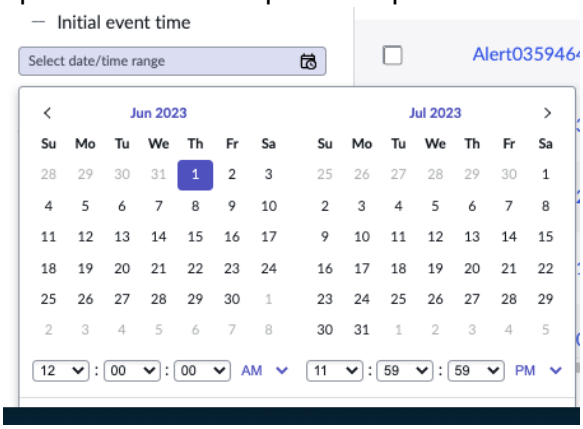

Filter the Express List display by adding or removing fields

Add fields to or remove fields from the Express List pane filter using the interactive filter panel. Use filtering to create different views of categorized alerts.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the primary navigation, select the Express List icon ( ).
3. In the left filter pane, select the Edit filter attributes icon ( ).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the **Edit filter attributes** dialog box, select fields or alert tags to add to your filter.

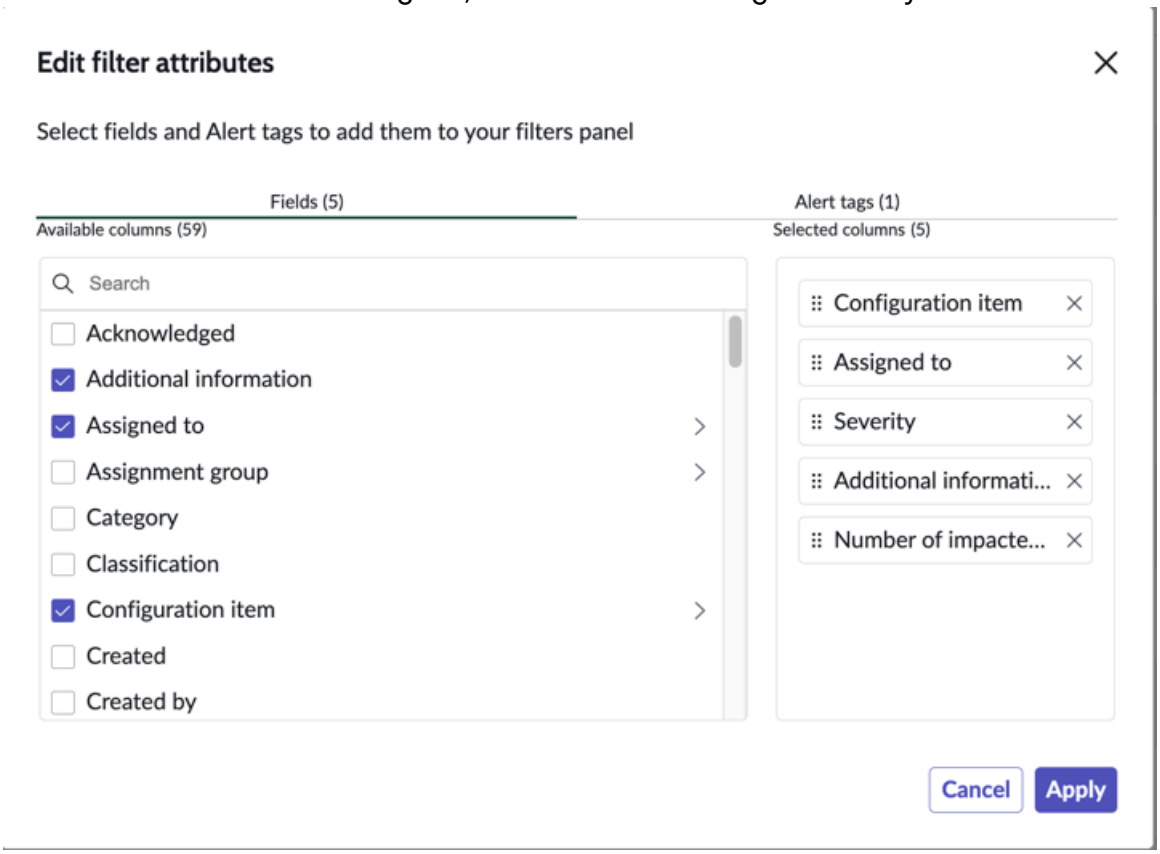

The **Fields** tab is displayed by default and displays the number of selected fields.

The number of available fields is displayed at the top of the column next to **Available columns**.

The selected fields are displayed in the right column. The number of selected columns are displayed at the top of the column next to **Selected columns**.

5. **Optional:** Find a specific field by entering its name in the **Search** field.
6. **Optional:** Rearrange the order of the selected columns by dragging them up or down in the right column.
7. **Optional:** Remove a column by selecting the **X** next to its name.
8. Select **Apply**.
9. **Optional:** Clear a filter by selecting **Clear** next to that filter name.
   Clear appears only if you selected or cleared any of the attribute check boxes of the selected filter.
10. **Optional:** View the applied conditions by selecting the Filter icon ( ) next to **Applied conditions**.

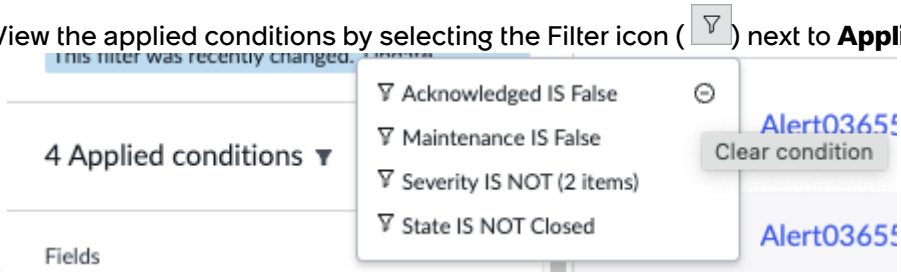

11. **Optional:** Remove a condition by selecting the **Clear condition**( ) icon.

Result

All the selected fields are displayed in the **Filters** panel in the **Fields** column.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Categorize alerts displayed in Express List

Create different views of alerts by adding tags in Express List to enable you to easily identify and group alerts without the need for discovery or accessing the CMDB.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the primary navigation, select the Express List icon ( ).
3. In the left filter pane, select the **Edit filter attributes** icon ( ).
4. In the **Edit filter attributes** dialog box, select the **Alert tags** tab.

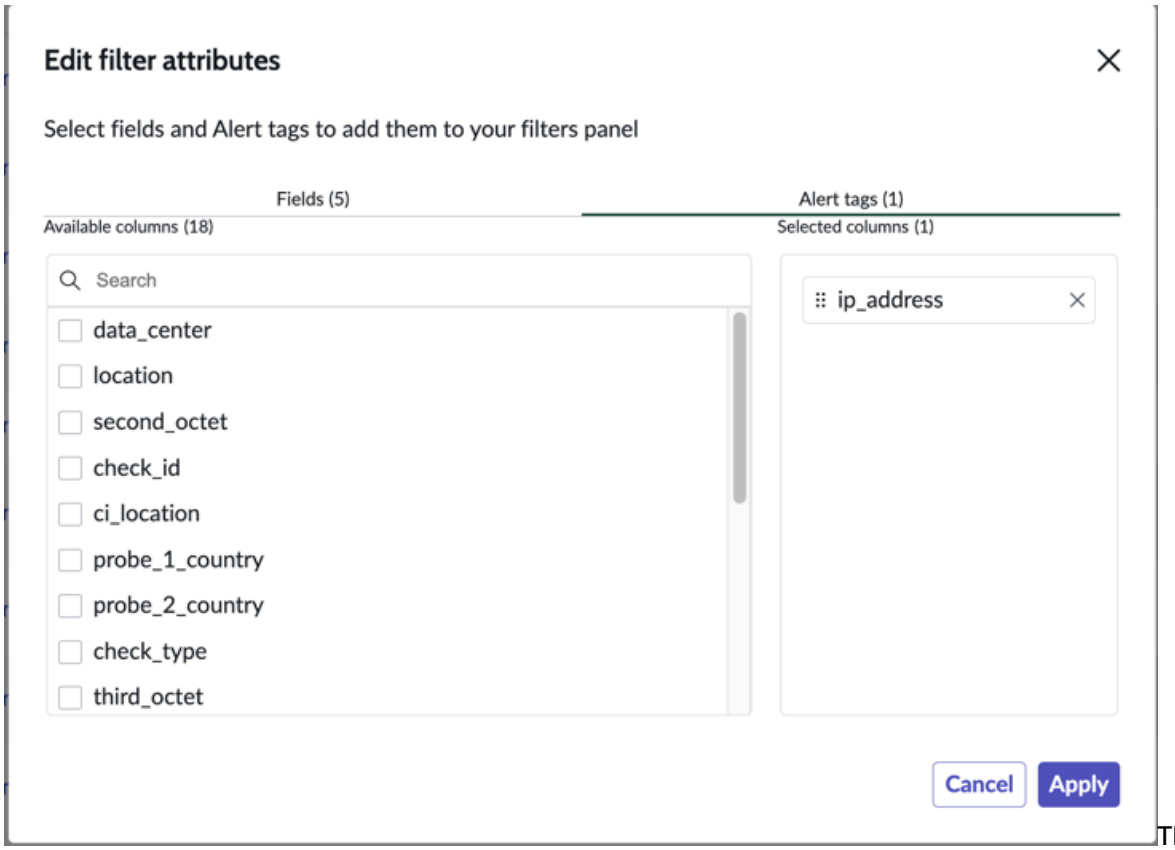

The **Alert tags** tab title displays the number of selected alert tags.

The left column displays the number of available tags in the **Available columns** column title.

The selected tags are displayed in the right column. The number of selected tags is displayed in the **Selected columns** column title.

5. **Optional:** Find a specific tag by entering its name in the **Search** field.
6. **Optional:** Rearrange the selected tags by dragging them up or down in the right column.
7. **Optional:** Deselect a tag by selecting the (**X**) next to it.
8. Select **Apply**.

Result

All the selected tags are displayed in the **Filters** pane under the **Alert tags** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Filter out or show matching alerts

Create different views of categorized alerts in Express List by either displaying or filtering out alerts that match a selected alert.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the primary navigation, select the Express List icon ( ).
3. Right-click an alert row.
4. Match or filter out matching alerts.
   - Display matching alerts by selecting **Show Matching** from the context menu.
   - Filter out matching alerts by selecting **Filter Out** from the context menu.

Save a custom filter in Express List

Save your custom view of categorized alerts in Express List to use again, share, or set as your default view. Saving your custom view saves all filters, tags, and conditions.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the primary navigation, select the Express List icon ( ).
3. Once you have set up a filter that you want to reuse, set as your default view, or share with others, save the filter.
   For more information about filtering the Express List alert display, see the following topics:
   - Filter the Express List display using attributes
   - Filter the Express List display by adding or removing fields
   - Filter out or show matching alerts

   a. In the Express List select **Save new**.

   b. In the **Save new filter** dialog box in the **Filter name** field, enter a name for the filter.

   c. **Optional:** Make this filter your default filter by selecting **Set as default**.

   d. Select **Save** to save the filter.
4. **Optional:** Share a saved filter by copying its link to send to others.

   a. In the **Filters** panel, select the filter from the drop-down list.

   b. Select the Copy link to clipboard icon ( ).
5. **Optional:** Set a saved filter as your default filter.

   a. In the **Filters** panel, select the filter from the drop-down list.

   b. Select the More Actions icon ( ) and select **Set as default**.
6. **Optional:** Delete a saved filter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Filters** panel, select the filter from the drop-down list.

b. Select the More Actions icon ( ) and select **Delete filter**.

Add custom fields to the Express List preview panel

Add custom fields from the Alert [em_alert] table to the Express List Info preview panel to track alert information.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **System Properties** > **All Properties**.
2. In the **System Properties** list, select **evt_mgmt.express_list.custom_fields**.

   If a message appears about the application scope, select **here** to be able to edit the record.
3. In the **Value** field, enter the custom name from the Alert [em_alert] table.
   If you want to provide more than one field, use commas to separate the field names.
4. Select **Update**.
5. **Optional:** View the custom fields on the preview panel:

   a. In the primary navigation, select the Express List icon ( ).

   b. Select a row or the information icon at the beginning of a row.

Result

The Info tab on the preview panel displays the custom fields at the bottom of the fields list.

Assigning and managing alerts

Learn how to assign and manage alerts in Express List.

Monitor incoming alerts

You can monitor incoming alerts in Service Operations Workspace. You can also monitor and manage alerts in the Event Management interface.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator

Note: The evt_mgmt_user role is view-only and does not have permission to perform this action.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. View the alert information.

Alerts column headings

Column heading	Description
Number	Unique ID generated by Event Management to identify the alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column heading	Description
Group	An entry in this column indicates that the associated alert is a member of an alert group. Alerts that do not have an entry in this column are ungrouped alerts. ○ **Automated**: Aggregated automatically by alert aggregation. A virtual alert is added to the group as the primary alert of the group. ○ **Rules-based**: Alert group created as a result of a user-configured correlation rule. ○ **Manual**: This alert is a member of an alert group that is formed when right-clicking an alert and setting it as secondary to the selected primary alert. ○ **CMDB**: CIs without historical data that were aggregated by alert aggregation based on CI relationships in the CMDB. ○ **Secondary**: This alert is a component of an alert group. The alert at the head of the group is known as the primary alert. When **Correlated Alerts** is selected, the secondary alerts that are under the primary alert do not display, making the Alerts list less cluttered and easier to review. ○ **None**: This alert is an ungrouped alert. To make an ungrouped alert become a member of a group, right-click it and select in the topic **Add to Groups**. Select the alert and click **Add Selected**. ○ **Text**: This alert is a member of an alert group that was created based on the similarity of their text or description. ○ **Log Analytics**: This alert is a member of a Log Analytics alert group. ○ **Component-based**: This alert is a member of a Log Analytics group that was created based on the affected IT components or services. ○ **Tag Cluster**: This alert is a member of an alert group that was created based on common tags or labels assigned to them. ○ **Network Traffic**: This alert is a member of an alert group that was created based on network traffic connections between processes of host CIs. ○ **Mixed**: This alert is a member of a grouping method that combines alerts using multiple grouping strategies, such as CMBD-based grouping and tag-based grouping, into a single, cohesive group.
Severity	The severity of the event. The value for this field is copied from the event unless the event closes the alert, in which case the previous severity is retained for reporting. ○ **Critical**: Immediate action is required. The resource is either not functional or critical problems are imminent. ○ **Major**: Major functionality is severely impaired or performance has degraded. ○ **Minor**: Partial, non-critical loss of functionality or performance degradation occurred. ○ **Warning**: Attention is required, even though the resource is still functional. ○ **OK**: An alert is created. The resource is still functional. ○ **Clear**: No action is required. An alert is not created from this event. Existing alerts are closed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column heading	Description
Priority group	Indicates which alerts should be attended to first. Priority is calculated for each open alert and then mapped into one of four priority categories.
Priority	Value providing a guide as to the priority of the alert, based on a number of accumulated categories. For example, alert state and business criticality.
Source	Event monitoring software that generated the event, for example, SolarWinds or SCOM. Optionally, you can enter a description, for example, `Group Alert`. This field has a maximum length of 100.
Description	The alert description.
Node	Node name, fully qualified domain name (FQDN), IP address, or MAC address that is associated with the event, such as IBM-ASSET. This field has a maximum length of 100.
Configuration item	JSON string that represents a configuration item. For example, {"name":"SAP ORA01","type":"Oracle"}. The CI identifier that generated the event appears in the **Additional information** field. This field has a maximum length of 1000. Click to open the alert in dependency view.
Impacted Services	Lists the services affected by this alert group. Select an impacted service to view its record. **Note:** You can access the Service Map in Service Operations Workspace directly from the impacted service form by selecting **Service Map**. The service map shows the impacted path of alerts, enabling you to quickly assess their effect on the service. For more information, see View unified service map and the impact paths in Service Operations Workspace. 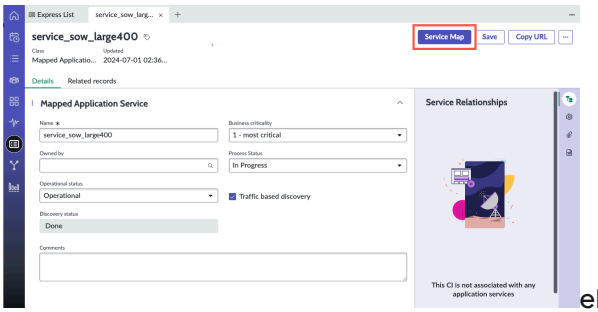 el-imp-service-open-map.png
Metric Name	Unique name that describes which metrics are collected and for which this alert has been created.
Maintenance	Shows whether the resource affected by the alert is in maintenance, Valid values are `true` or `false`.
Task	The corresponding task for the alert, such as an incident, change, or problem.
Parent	Reference to a parent alert.
Initial event generation time	Time that the initial event occurred in the remote system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

If Operational Intelligence is activated, you can right-click an alert and click **View Metrics**. This opens the integrated Insights Explorer and Dependency Views map for the associated CI.

Related topics

- Metric Intelligence
- View alert information
- Alert grouping types and creation methods

Assign alerts in Express List

Assign alerts in Express List to yourself, to someone else, or to an assignment group to balance your organization's work load.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and does not have permission to perform this action.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Select the check box next to the alert number that you want to assign.

Note: You can perform the action on up to 1,000 alerts simultaneously by selecting the **Select All** check box in the Active alerts list.

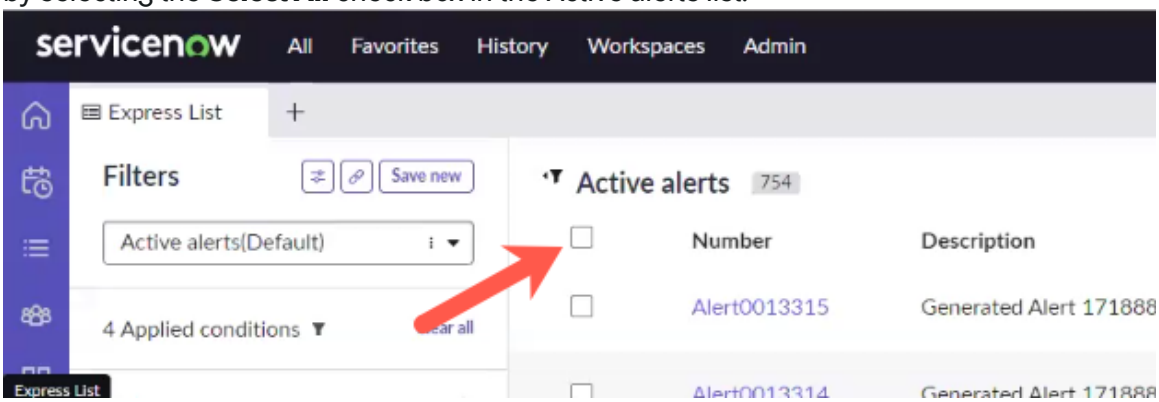

4. Select the **Close** drop-down list.

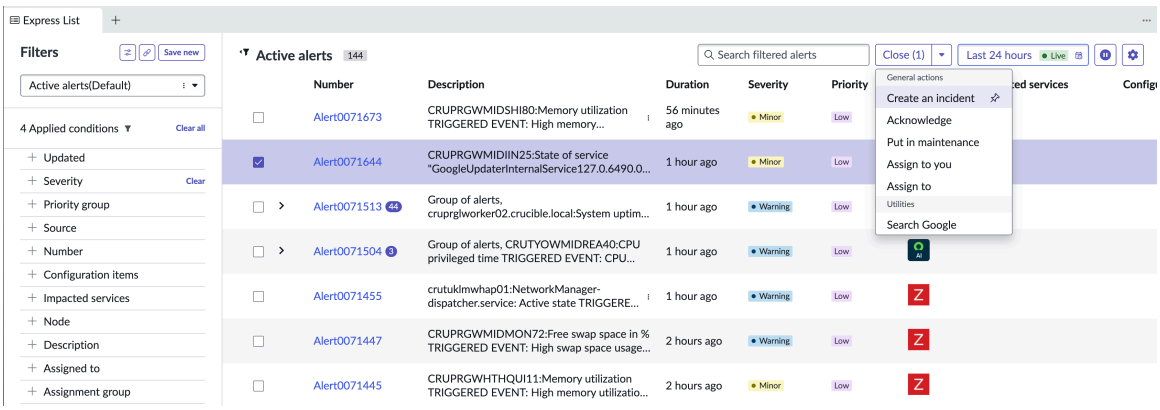

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Perform one of the following actions.

Assignment	Actions
Yourself	Select **Assign to you** and then select **OK**.
Someone else	**a.** Select **Assign to**. **b.** Select the name of the person to whom you want to assign the alert from the **Assigned to** field. **c.** (Optional) Enter information about the as signment or alert in the **Work notes** field. **d.** Select the **Assign <number> alerts** button.
Assignment group	**a.** Select **Assign to**. **b.** Select the name of the group from the **As signment group** field. **c.** Select the name of the person to whom you want to assign the alert from the **Assigned to** field. **d.** (Optional) Enter information about the as signment or alert in the **Work notes** field. **e.** Select the **Assign <number> alerts** button.

Result
The assigned person appears in the **Assigned to** column and the assignment group appears in the **Assignment group** column in the Express List pane.

Find alert records in Express List using text search

Perform a text search for alert records in a filtered list of alerts.

Before you begin
Role required: evt_mgmt_operator, evt_mgmt_admin

About this task
Event Management searches through the alerts that you're currently seeing in the Express List based on the filters and time range you have set. It searches in the following predefined fields even if they aren't visible: **Metric name**, **Node**, **Alert number**, **Alert Tags**, and **Description**. It treats the text that you enter as a single string rather than individual words. The selected alert list mode determines which alert types are displayed and included in the search. **Essential** mode displays primary and single alerts only. **Extended** mode displays all alert types, including secondary alerts.

Procedure
1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. **Optional:** Select **Essential** or **Extended** from the **Switch search scope** menu.
   Essential includes primary and single alerts. **Extended** includes all alert types, including secondary alerts.
4. In the **Search filtered alerts** field, enter the search text.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Ensure that the text string meets or exceeds the minimum character requirement specified in the notification message.
- The maximum character limit is 50.
- The search isn't case-sensitive.

Note: The minimum character requirement is determined by the *evt_mgmt.express_list.min_search_chars* system property setting.

5. Perform the search by selecting the check icon or pressing **Enter**.

Result

The results match both the filters that you have applied and the string you entered in the Search filtered alerts field. In the resulting alert records, the string is highlighted and the filters show the number of records that match the string. Recent searches are saved and accessible from the Search filtered alerts field.

Highlight alerts by condition in Express List

Learn how to apply color-coded highlights to alerts in the Express List based on conditions such as severity, state, and priority, so you can quickly spot critical signals.

Before you begin

Role required: evt_mgmt_admin

About this task

To turn on this feature, enable the system property [sn_sow_exp_app.express_list_conditional_highlighting]. This property is disabled by default.

Procedure

1. Navigate to **All** and search for `sys_highlighted_value_list.do`.
   The Highlighted Values page opens.
2. In the Highlighted Values filter, set **for text** = `em_alert`.
3. Select **New** to create a record where you add the highlight conditions.
4. In the **Table** field, select **Alert[em_alert]**.
5. In the **Field** field, select the field from the drop-down list.
6. Right-click the form header and select **Save**.
7. In the **Highlighted Value Condition** tab, select **New**.
8. In the Highlighted Value Condition form, fill in only these fields.

Field	Description
Conditions	Define the criteria that trigger the highlight. Select **Add Filter Condition** to add a condition, or Add OR Clause to add an alternative set of conditions. Use **AND** or **OR** to combine multiple conditions, and the red **X** to remove one.
Color	Select the highlight color applied to rows that match the condition (for example, yellow).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Application	The application scope in which the condition applies (for example, Global).
Order	The sequence in which the condition is evaluated when multiple conditions exist. Lower numbers are evaluated first.
Show Icon	Select this check box to display an icon with the highlight.
Icon	Icon to display (for example, check-circle-fill).
Variant	Controls the visual style and emphasis of the highlighted text, including the icon (for example, primary, secondary, or tertiary), determining how prominently it appears.
Value Override	Enter a string value to display in place of the field's actual value.
Size	Select the size of the highlight or icon (for example, Small or Medium). Medium is the default. **Note:** This is available only from the Australia release.

9. Select **Submit**.
   You return to the Highlighted Values page.
10. In the **UX Highlighted Values Configurations** tab, select **Edit**.
11. In the Collection list, select **SOW Highlighted Value Config** and use the arrow button to move it to the **UX Highlighted Values Configurations** List.
12. Select **Save**.
    The highlight conditions appear on the alerts in the Express List.

Manage alert states in Express List

Learn how to change alert states.

Acknowledge viewed alerts in Express List

Acknowledge alerts that you've viewed to remove clutter in the Express List pane.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and does not have permission to perform this action.

About this task

Use Acknowledge to denote that the alert is known, and can temporarily be ignored. Do not use Acknowledge to mark an alert as needing attention.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Acknowledge selected alerts.
   You can acknowledge alerts regardless of whether you are working in the main pane or in the preview panel. The preview panel opens only when a single alert or alert group is selected.

Location	Action
The Express List pane	**a.** In the Express List pane, select the alert to acknowledge. You can select more than one alert row. **Note:** You can perform the action on up to 1,000 alerts simultaneously by selecting the **Select All** check box in the Active alerts list. 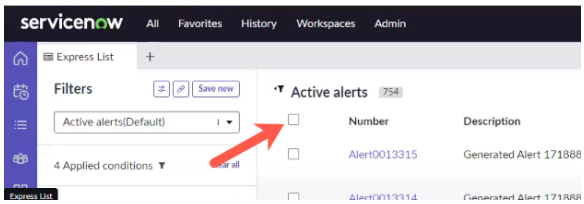 To display the individual alerts inside a group, select the chevron icon ( ) at the beginning of the alert group row. **b.** From the **Close** drop-down list at the top right of the alert list, select **Acknowledge**.
The Express List preview panel for group alerts	**a.** Select an alert group row check box. **b.** In the **Alerts in group** tab, select an alert tile. **c.** Select the more actions icon ( ). **d.** Select **Acknowledge**.

Result

The alert is acknowledged and its status is displayed at the top of the preview panel.

Place an alert in the Acknowledged state

Placing an alert into the Acknowledged state indicates that the alert has issues that need to be addressed.

Before you begin

Role required: evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and does not have permission to perform this action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Filter with **Active alerts**.
4. Select an alert.

5. On the **Details** subtab, select the **Acknowledged** check box to place the alert into the Acknowledged state.
6. Click **Save**.

Place an alert in the Maintenance state

Placing an alert into the Maintenance state indicates that there are issues with the alert's CIs that are being addressed.

Before you begin

Role required: evt_mgmt_admin or evt_mgmt_operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Filter with **Active alerts**.
4. Select an alert.

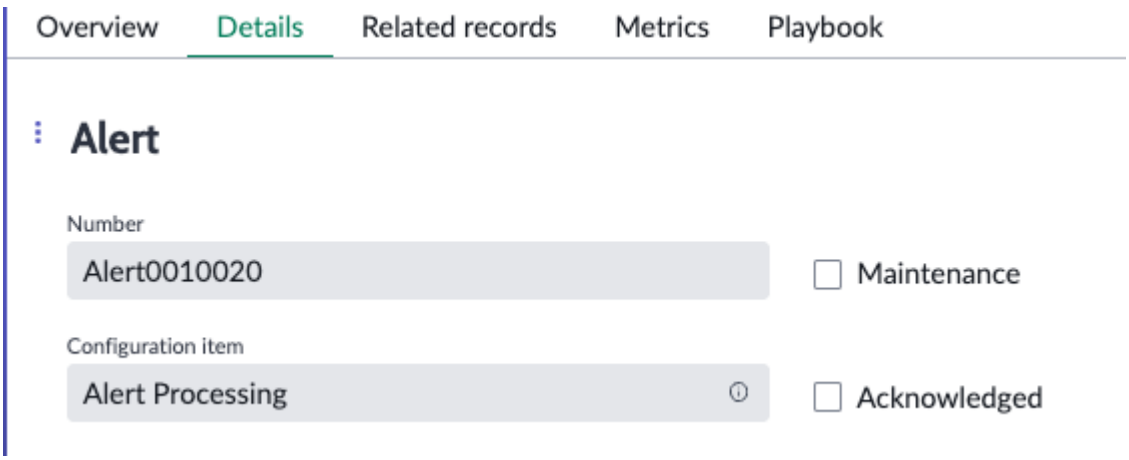

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. On the **Details** subtab, select the **Maintenance** check box to place the alert into the Maintenance state.
6. Click **Save**.

Mute a CI's alerts in Express List

Mute a CI's alerts in Express List to avoid alerts being issued when replacing or upgrading the CI.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and does not have permission to perform this action.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Select the alert row check box.

Note: You can perform the action on up to 1,000 alerts simultaneously by selecting the **Select All** check box in the Active alerts list.

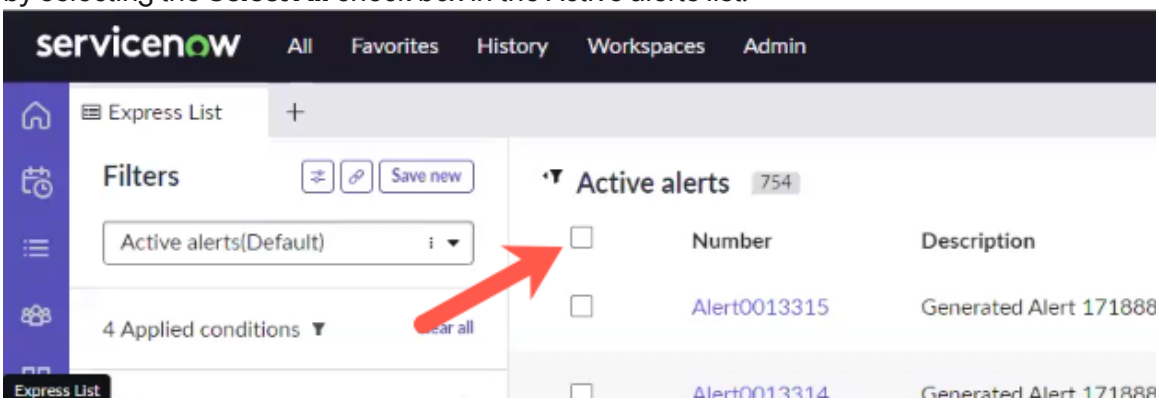

4. From the **Close** drop-down list, select **Put in Maintenance**.

Result

Alerts are muted for the configuration items of that alert.

Note: The maintenance timeframe isn't provided in the **Put in Maintenance** action. Alert maintenance status is calculated based on whether the related CI is currently in maintenance, such as during an active change request or through a maintenance rule. The **Maintenance Calculation** scheduled job updates associated alerts when the CI enters or exits maintenance.

Close an alert in Express List

Manually close alerts from the Express List pane or preview panel if you know the root cause or have fixed the problem.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and does not have permission to perform this action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Close selected alerts.
   You can close alerts regardless of whether you are working in the main pane or in the preview panel. The preview panel opens only when a single alert or alert group is selected.

Location	Action
The Express List pane	**a.** In the Express List pane, select the alert. You can select more than one alert row. **Note:** You can perform the ac tion on up to 1,000 alerts simulta neously by selecting the **Select All** check box in the Active alerts list. 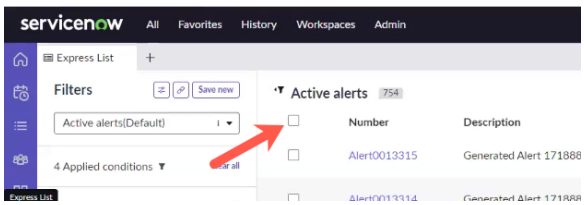 To display the individual alerts inside a group, select the chevron icon ( ) at the beginning of the alert group row. **b.** From the **Close** drop-down list at the top right of the alert list, select **Close alert**.
The Express List preview panel for group alerts	**a.** Select an alert group row. **b.** In the **Alerts in group** tab, select an alert tile. **c.** Select the more actions icon ( ). **d.** Select **Close alert**.

Result
The alert is closed and its status is displayed at the top of the preview panel.

Reopen alert closed by AI Specialist

Reopen an alert that the AIOps AI Specialist auto-closed as noise so it bypasses the noise classifier and runs the full investigation pipeline. Your override is honored for the alert's entire lifecycle, preserved in the audit trail, and fed back to improve future classification accuracy.

Before you begin
Role required: evt_mgmt_operator, evt_mgmt_admin

About this task
When the AI Specialist classifies an alert as noise, it closes the alert automatically. If you determine the alert is significant, reopen it. Reopening tells the system the classification was wrong, so the alert is not classified as noise again. Instead, it runs through the full investigation as if it were new and significant, and your decision helps improve future classification.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Open an alert that was closed by AI and you want to reopen.
   When an alert is auto-closed by the AIOps AI Specialist, its reasoning appears under the **Overview** tab in the Summary section. If you determine the alert is significant, reopen it. The alert can be reopened by selecting the **Reopen** link provided following the reasoning or you can reopen the alert from the **Details** tab as shown in the procedure.
4. Select the **Details** tab.

   In the Activity section, the work note states the alert is closed by an AI Specialist. The **State** field shows **Closed**.
5. Change the **State** field value to **Reopen**.

   Note: Select **Reopen**, not **Open**. Reopen is the state transition the system watches for to trigger the classifier bypass and full investigation pipeline. Setting the state to **Open** does not register as an operator override of the AI close, so the alert would not be guaranteed the bypass-and-reinvestigate treatment.

Hide closed alerts in Express List

Choose whether to display the closed alerts in an alert group as well as the open alerts on the preview panel and in Link View.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. In the Active alerts list, select an alert group.

   The preview panel shows the default display:

   - The Alerts tab and Timeline view show only the open alerts. The closed alerts are hidden.
   - Link View shows the attributes of both open and closed alerts.
4. On the relevant tab, change the display.

   On the preview panel, Alerts tab:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Show both the open and the closed alerts by switching the **Hide closed alerts** toggle to the closed position.

b. (Optional) Sort the displayed alerts by time or by severity by selecting the sort icon ( ) and then choosing the relevant menu item.

Note: If you switch the toggle or sort the alerts in the Timeline view, the Alerts tab adjusts to reflect your selection.

In Link View: Show only the attributes of open alerts while hiding those of closed alerts by switching the **Hide attributes of closed alerts** toggle to the open position.

Responding to and working with alerts in Express List

Learn how to respond to an alert and manage its workflow.

Manage alerts autonomously agentic workflow

Enhance IT operations with AI-driven, autonomous alert management using the manage alerts autonomously agentic workflow.

Manage alerts autonomously agentic workflow overview

The manage alerts autonomously workflow introduces a unified AI-driven process that significantly streamlines alert management and reduces resolution times. This workflow supports alert management in the following ways:

- Automates the triage
- Impact analysis
- Root cause investigation of IT alerts
- Generates reports, summarizes key insights and possible next steps

For information on how to review key insights and data derived from the workflow in Express List, see Review AI-generated alert insights in Express List.

For information about configuring this workflow, see Configure the manage alerts autonomously agentic workflow.

Use the information on this page to learn about the actions related to the manage alerts autonomously agentic workflow. To modify the workflow, you must duplicate it and adjust the settings according to your requirements. For more information, see Duplicate an agentic workflow .

Important: When you modify an agentic workflow, AI agent, or tool, make sure that you update all instructions accordingly.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Manage alerts autonomously agentic workflow

The manage alerts autonomously agentic workflow uses the manage alerts AI agent to perform alert management and resolution tasks.

AI agent used in the manage alerts autonomously agentic workflow

AI agent	AI agent role
Manage alerts AI agent	Investigates alerts, summarizes alert-related reports, and stores structured insights with key findings.

The manage alerts autonomously agentic workflow performs several actions in the course of the workflow. These actions may include the following:

- Triages alerts
  - Evaluates and categorizes alert
  - Analyzes alert history to identify noise patterns
  - Updates alert group description based on analysis
  - Performs related incidents analysis to detect focal points and common closure patterns
- Determines alert impact
  - Evaluates impact on services
  - Determines user impact by finding recent incidents or cases
  - Uses observability skills for deeper service state validation
- Investigates relevant information
  - Retrieves and summarizes similar closed alerts
  - Analyzes recent changes for causal relationships
  - Summarizes related KB articles for relevant information
  - Identifies trends or anomalies in related metrics
  - Uncovers errors, exceptions, or warnings in related logs
  - Analyzes non-primary log analytics alerts in an alert group to support investigation, including the following:
    - Classifies alerts as proactive or reactive, indicating emerging risk or an active issue
    - Assess if triage is required
    - Provides contextual insights, including potential failure scenarios and what may occur if the alert isn't addressed
- Summarizes and stores information
  - Consolidates findings
  - Generates a final summary
  - Saves the summaries in the alert record, such as

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- reasoning, evidence, and confidence level behind AI-driven investigation of alert significance and automatic closure of insignificant alerts, and
- summaries of Service Observability dashboards to provide health analysis of each service associated with an alert. For more information see, Analyze service health in Service Observability.

- Provides clear, actionable insights

Important: This agentic workflow is turned on by default. For more information, see AI agents, skills, and agentic workflows on by default.

Review AI-generated alert insights in Express List

Access alert information in Express List that is consolidated autonomously by AI skills and agents. Use the AI insights badge, column, and filter to monitor alert statuses and review of AI-generated insights.

Before you begin

For this feature, you must have ServiceNow Otto for IT Operations Management (ITOM) installed on your instance. For more information about installing ServiceNow Otto plugins, see Install plugins for ServiceNow Otto.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with your applications have specific roles that come included with the application. If you select *Users with specific roles* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

Role required: None

About this task

The manage alerts autonomously workflow investigates alerts, summarizes alert-related reports, and stores structured insights with key findings for use in Express List. For more information about the manage alerts autonomously agentic workflow, see Manage alerts autonomously agentic workflow.

Autonomous workflow is a simplified AI process elaborated in the AI Specialist. You can now activate and manage AI Specialist, controlling the groups it targets and the specific alert type it handles. Activating the AI Specialist automatically disables the Autonomous workflow.

Note: Currently, ServiceNow Otto for ITOM only supports tag-based, CMDB, Log Analytics, Mixed, Automated, and Network Traffic-based alert groups. For all other alert group types, it only analyzes the parent alert.

There are several ways to explore AI insights in Express List.

- Quickly review if alerts have been processed by checking for the AI insights badge AI.
- Search for specific information and key words in the Insights column. To add the column, see Add or modify Express List columns for more information.
- Filter for information by using the **Insights** filter attribute in the filter panel.
- Review **AI insights** in the preview panel or the alert record.

These options present differently, depending on whether the workflow is configured for automatic or manual execution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- When the workflow operates automatically, alerts are addressed as they're created and AI insight information is displayed in Express List.
- When the workflow operates manually, you must manually generate AI insights. For more information, see details in the following procedure.

For information about configuring this workflow, see Configure the manage alerts autonomously agentic workflow.

Procedure

1. Navigate to **Workspace Experience** > **Workspaces** > **Service Operations Workspace**.
2. Select the Express List icon .
3. Review the AI insights through the following options.

Review AI insights	Procedure
Check for the AI Insights badge for alert status	○ If the AI insights icon AI is visible next to the check box, insights are available for that alert. ○ If insights aren't available for an alert, you can initiate the process manually. Details for generating insights are in the following options.
Search for alerts with AI Insights information and key words	Search for content with the free text search. For more information, see Find alert records in Express List using text search.
Filter using AI Insights filter attribute	Filter using the **Insights** attribute with a minimum string of two characters. For more information, see Filtering the alert display in the Express List pane.
Review AI insights in the preview panel	**a.** In the alerts list, select an alert by selecting the check box or the information icon ⓘ next to the alert. **b.** ▪ If data isn't available for this alert, you can initiate the process by selecting **Generate**. ▪ In the **Insights** tab, review **AI insights**. ▪ If the AI Specialist is in the active state, you can see the processing steps of the agentic workflow. After processing completes, the **View AI activity** link appears. Select the link to go to the **AI activity** tab on the alert details page. In the **AI activity** tab, view the detailed report of the alert along with the workflow steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Review AI insights	Procedure
Review AI Insights in the alert record overview	**a.** In the alerts list, select an alert number to open the alert record. The **Overview** tab is selected by default. **b.** ▪ If data isn't available for this alert, you can initiate the process by selecting **Gen erate**. ▪ If data is available, you can review the **AI insights** summary in the Summary sec tion. ▪ If the AI Specialist is in the active state, you can see the processing steps of the agentic workflow. After processing com pletes, the **View AI activity** link appears. Select the link to go to the **AI activity** tab on the alert details page. **c.** Select the **AI activity** tab to view the de tailed report of the alert along with the work flow steps.

Analyze alert impact agentic workflow

Use the analyze alert impact agentic workflow to investigate an alert and get the context that you need to respond efficiently.

Analyze alert impact agentic workflow overview

The analyze alert impact agentic workflow uses AI agents to interact with observability tools, such as Dynatrace, Kentik, and New Relic. The agents surface information to help you do the following:

- Understand the severity and impact of alerts.
- Find potential root causes using insights, such as recent deployments.
- Identify ownership by surfacing relevant services and responsible teams.
- Support stakeholder communication with clear, actionable information.

Use the information on this page to learn about the agents related to the analyze alert impact agentic workflow. To modify the analyze alert impact agentic workflow, duplicate it and adjust the settings.

Important: When you modify an agentic workflow, AI agent, or tool, make sure that you update all instructions accordingly.

Analyze alert impact agentic workflow

Investigate an alert and get insights into its severity, impact, ownership, and root causes.

To access the agentic workflow in AI Agent Studio:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Analyze alert impact**.

The Analyze alert impact page lets you manage the agentic workflow, including defining key requirements, defining security controls, and adding triggers.

Important: This agentic workflow is turned on by default. For more information, see AI agents, skills, and agentic workflows on by default.

AI agents used in the analyze alert impact agentic workflow

The analyze alert impact agentic workflow uses observability AI agents to gather information from alerts and request insights. The observability AI agents require additional configuration. For configuration procedures and detailed information about the data returned by these agents, see Configure observability agents for ServiceNow Otto.

AI agents and their roles

AI agent	What it does
SRE Investigate AI Agent	Coordinates observability and network agents to collect and correlate investigation findings.
Observability agents	Retrieves data from the observability vendor associated with the alert. For details, see Configure observability agents for ServiceNow Otto

Generating the alert analysis

In the agentic workflow record:

1. Review the information in the Define key requirements, Define security controls, and Add triggers screens, make the necessary updates, and then select **Continue**.
2. In the Select channels and status screen:
   a. Turn on the ServiceNow Otto panel display.
   b. Select **Save and test**.
3. Test the agentic workflow by asking a question, for example, `What is the impact of Alert0010056NR?`
4. Select **Continue to Test Chat Response**.

   The page shows the chat responses, visualizes the AI agents involved, and lists the AI agent decision logs.

In *AI Agent Studio*, you get notified when the analysis is generated. You can then act on the information or ask more questions about the alert. For more information about using the agentic workflow in the ServiceNow Otto panel, see Analyze alert impact in the ServiceNow Otto panel.

Analyze alert impact in the ServiceNow Otto panel

Learn how to use the analyze alert impact agentic workflow in the ServiceNow Otto panel. The agentic workflow helps you investigate an alert and get the context that you need to respond efficiently.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Make sure ServiceNow Otto for ITOM is installed and observability agents are configured for third-party vendors. For more information, see Install plugins for ServiceNow Otto and Configure observability agents for ServiceNow Otto.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with your applications have specific roles that come included with the application. If you select *Users with specific roles* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon ( ).
3. Select an alert from a source associated with an observability agent.
4. Open the panel by selecting the ServiceNow Otto icon.
5. Ask a tool-specific question about the alert.
   For example, for New Relic alerts, you can ask `Have there been recent deployments for this alert?`. To better understand the type of data you can query, see the **Overview of data returned** row in the corresponding vendor's table in Configure observability agents for ServiceNow Otto.

   AI agents analyze your question and respond with a summary or insight. The first response for an alert might take several seconds to appear.

What to do next

Use the information that you received to understand the alert impact and respond accordingly. You can also ask more questions about the alert.

Note: Your conversation in the ServiceNow Otto panel is specific to the alert that you selected in Step 3. To ask about a different alert, open that alert and select the ServiceNow Otto icon to start a new conversation.

Triage and analyze alerts agentic workflow

Use the Triage and analyze alerts agentic workflow to complete preliminary alert tasks and analysis for alerts.

Triage and analyze alerts agentic workflow overview

The Triage and analyze alerts agentic workflow uses AI agents to support alert triage, investigation, and analysis. AI agents can perform the following functions:

- Acknowledge alerts.
- Assign alerts to individual users and assignment groups.
- Summarize alert and alert group data to create a human-readable description and add technical analysis.
- Investigate relevant past incidents to analyze the significance of the alert and present options for resolution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: These functions can also be used for the origin alert of an incident when the workflow is accessed from the incident form or Express List using the ServiceNow Otto panel.

Use the information on this page to learn about the agents related to the Triage and analyze alerts agentic workflow. To modify the Triage and analyze alerts agentic workflow, you must duplicate it and adjust the settings according to your requirements. For more information, see Duplicate an agentic workflow.

Important: When you modify an agentic workflow, AI agent, or tool, make sure that you update all instructions accordingly.

Triage and analyze alerts agentic workflow

Acknowledge, assign, and investigate current and past alerts to determine significance and possible resolutions.

To access the agentic workflow, use the ServiceNow Otto panel. For more information about using the agentic workflow in the Now Assist panel, see ServiceNow Otto panel.

AI agents used in the Triage and analyze alerts agentic workflow

The Triage and analyze alerts agentic workflow uses a team of AI agents to perform preliminary tasks and support alert resolution.

AI agents used in the Triage and analyze alerts agentic workflow

AI agent	AI agent role
Alert handling AI agent	Assign, acknowledge, and maintain an up-to-date alert record.
Alert analysis AI agent	Perform alert analysis, and update alert description, when applicable.
Alert history analysis AI agent	Analyze past occurrences, assess its significance, and close the alert when applicable.
Related incidents analysis AI agent	Analyze past incidents and provide insights on common assignments and summarized resolution notes.
Alert verification AI agent	Assess data completeness and determine whether the alert is suitable for analysis.

View an alert analysis by ServiceNow Otto in Express List

View an alert analysis created using generative AI. Alert analysis include a human-readable brief of the alert and technical information to help you investigate the alert more effectively.

Before you begin

Install the ServiceNow Otto for ITOM. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

Note: Currently, ServiceNow Otto for ITOM only supports tag-based, CMDB, Log Analytics, Mixed, Automated, and Network Traffic-based alert groups. For all other alert group types, it only analyzes the parent alert.

Role required: evt_mgmt_operator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. In the Active alerts list, select the information icon ( ) next to the alert.
4. On the preview panel **Info** tab, select **Analyze** in Alert analysis by ServiceNow Otto.
5. View the information provided in the Alert analysis.
6. **Optional:** Use the Alert analysis icons to perform related tasks.

Alert analysis icons

Icon	Description
	Copy the content of the alert analysis to the clipboard.
	Refresh the alert analysis. **Note:** Refreshing regenerates the results. Past results are deleted.

View an alert analysis by ServiceNow Otto in Express List

View an alert analysis created using generative AI. Alert analyses include a human-readable brief of the alert and technical information to help you investigate the alert more effectively.

Before you begin

Install the ServiceNow Otto for ITOM plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

Note: Currently, ServiceNow Otto for ITOM only supports tag-based, CMDB, Log Analytics, Mixed, Automated, and Network Traffic-based alert groups. For all other alert group types, it only analyzes the parent alert.

Important: This generative AI skill is turned on by default. The skill will be automatically available to appropriate role users for the application. For more information, see AI agents, skills, and agentic workflows on by default .

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the navigation bar, select the Express List icon ( ).
3. In the Active alerts list, select the information icon ( ) next to the alert.
4. On the preview panel **Info** tab, select **Analyze** in Alert analysis by ServiceNow Otto.
5. View the information provided in the Alert analysis.
6. **Optional:** Use the Alert analysis icons to perform related tasks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert analysis icons

Icon	Description
	Copy the content of the alert analysis to the clipboard.
	Refresh the alert analysis. **Note:** Refreshing regenerates the results. Past results are deleted.

View an alert analysis by ServiceNow Otto in Service Operations Workspace

View an alert analysis by ServiceNow Otto for an alert on the alert's Overview tab in the Service Operations Workspace. Alert analyses include a human-readable brief of the alert and technical information to help you investigate the alert more effectively.

Before you begin

Install the ServiceNow Otto for ITOM plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

Note: Currently, ServiceNow Otto for ITOM only supports tag-based, CMDB, Log Analytics, Mixed, Automated, and Network Traffic-based alert groups. For all other alert group types, it only analyzes the parent alert.

Important: This generative AI skill is turned on by default. The skill will be automatically available to appropriate role users for the application. For more information, see AI agents, skills, and agentic workflows on by default.

Role required: evt_mgmt_user, evt_mgmt_operator

Procedure

1. Navigate to **Workspace Experience** > **Workspaces** > **Service Operations Workspace**.
2. Select the List icon ( ).
3. In the Alerts list, select **All Alerts**.
4. Select an alert.
   The alert opens in a separate tab, displaying the **Overview** subtab.
5. In the **Analysis** section, select **Analyze alert**.
6. Review the information provided in the alert analysis.

Launch an alert analysis from the ServiceNow Otto panel

Initiate an alert analysis from the ServiceNow Otto panel. The alert analysis displays directly in the panel for convenient review.

Before you begin

Install the ServiceNow Otto for ITOM plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

For comprehensive information about the ServiceNow Otto panel, see ServiceNow Otto panel.

Role required: evt_mgmt_operator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the List icon ( ) in the navigation bar.
3. Select **Lists** > **Alerts** > **All Alerts**.
4. Initiate an alert analysis from the ServiceNow Otto panel using one of the available methods.

Launch an alert analysis	Procedure
From the list	**a.** Select the ServiceNow Otto panel icon ( ). The panel displays. **b.** Enter an alert analysis request, including the alert number, in the field at the bottom of the panel. **Note:** If you provide incomplete information, ServiceNow Otto prompts you for the missing details. The alert analysis displays in the panel.
From an alert form	**a.** Open an alert. **b.** Select the ServiceNow Otto panel icon ( ). The panel displays. **c.** Select **Perform alert analysis**. The alert analysis displays in the panel.

5. **Optional:** Enter follow-up questions or requests.
   For example: `Summarize the analysis` or `Show alert severity`.

Accelerate alert resolution with past incident analysis

Use ServiceNow Otto to investigate relevant past incidents and present effective strategies used to resolve them.

ServiceNow Otto queries historical records to find past incidents and analyzes their frequency, criticality, work notes, and resolution. It presents a summary of the most relevant related incidents in the ServiceNow Otto panel. The summary includes resolution strategies used, as well as contact details for individuals or teams who resolved the incidents and could assist with the current alert. The information provided helps you better understand the context and significance of the current alert and work more efficiently toward resolving it.

Generate a ServiceNow Otto summary of past related incidents

View an AI summary of past incidents related to the current alert, including strategies used to resolve them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Install ServiceNow Otto for ITOM. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.
- Ensure that the Alert investigation skill is active. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

For comprehensive information about the ServiceNow Otto panel, see ServiceNow Otto panel.

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the List icon ( ) in the navigation bar.
3. Select **Lists** > **Alerts** > **All Alerts**.
4. Open an alert.
5. Select the ServiceNow Otto panel icon ( ).
   The panel displays.
6. Request a past incidents investigation by selecting **Analyze related incidents**.
   A summary of the most relevant related incidents displays in the panel.
7. **Optional:** Inspect the details of a related incident by selecting the relevant link under Related Incidents.
   The incident page opens, showing the Overview tab.

Respond to an alert in Express List

Run a response action to remediate an alert issue manually from the Express List pane.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and doesn't have permission to perform this action.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Select the alert number to open it.
4. Select the response action for the alert from the **Close** drop-down list.

Note: If the list is empty, no response actions are available.

Run response actions on multiple alerts in Express List

Run response actions to efficiently remediate and respond to issues on multiple alerts directly from the Express List pane.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. In the **Active Alerts** list, select the check boxes for desired alerts.
4. Select the **Alert actions** drop-down arrow to view **Response actions**.

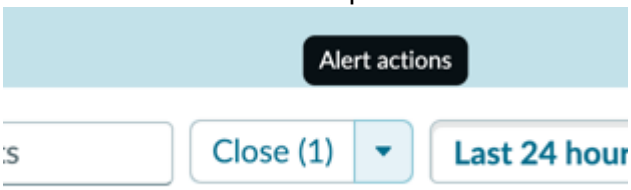

5. Select a response action.

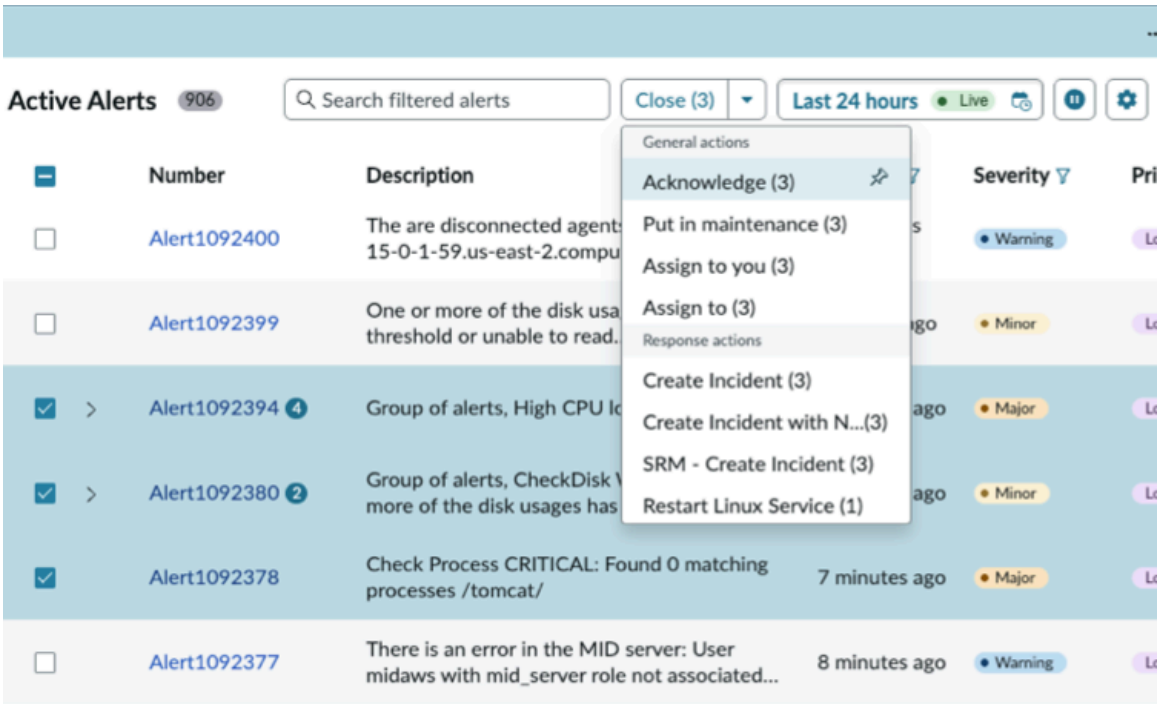

Note: Numbers next to actions represent the number of alerts that can be addressed with that action. If the list is empty, no response actions are available.

The **Run response action** modal appears. The modal describes the number of alerts addressed with the selected action, and lists the selected alerts that aren't addressed.

View details about an alert in Express List

View information about the impact and cause of an alert on the alert's Overview tab in the Express List.

Before you begin

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select the alert number to open it.
   The alert opens in a separate tab, displaying the **Overview** tab.
4. View alert information on the **Overview** tab.

Overview tab

Section	Description
Summary	Contains a description of the selected alert, describing why it is considered an issue.
Impact	Lists the Configuration item and any services impacted by the alert. The **Impacted Services** tile includes the following subtabs: ○ **Application services:** The application services affected by the alert. For details on application services, see Application services. ○ **Related service offerings:** The service offerings connected to the application services affected by the alert. Application services and their related service offerings display on a many-to-many basis. That is, several service offerings can relate to one or more application services, and several application services can relate to one or more service offerings.
Cause	Lists the log properties with information contributing the alert issue, as well as probable root causes for the alert.

View data on impacted services on the preview panel in Express List

View extra information on the preview panel about impacted services that are bound to alerts.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon.
3. In the Active alerts list, select an alert check-box.
4. On the preview panel, open the Info tab.
5. Hover over the Impacted services section to display information icons.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Select the relevant icon to display extra information on an impacted service.

Impacted service information icons

Icon	Name	Description
ⓘ	Info	Displays the severity, the business criticality, and a short description of the impacted service. From this pop-up, you can open the service map in Service Operations Workspace by selecting **View service map**. 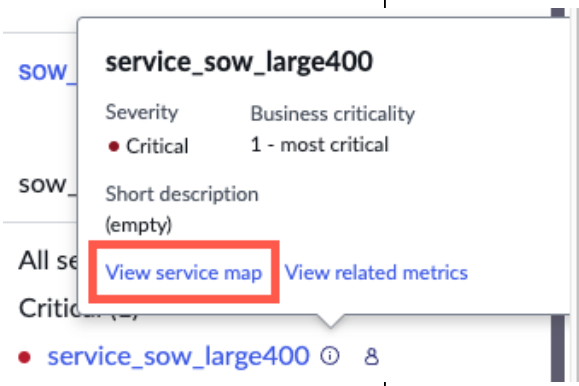 The service map shows the impacted path of alerts, enabling you to quickly assess their effect on the service. For more information, see View unified service map and the impact paths in Service Operations Workspace. If metrics data exists for this service, you can open the Metric Explorer from the pop-up by selecting **View related metrics**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Icon	Name	Description
	Contact person/group	Displays the name of the owner of the impacted service and the support group assigned to it.

View metrics for an alert

Viewing metrics enables you to understand historical data related to an alert's CI.

Before you begin

- Ensure that you install the Operational Intelligence (com.snc.sa.metric) plugin.
- Configure the metrics to be displayed for the alert, as described in Configure alert metrics.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Select an alert to open it.
4. Click the **Metrics** subtab.
   The configured metrics appear on the page.

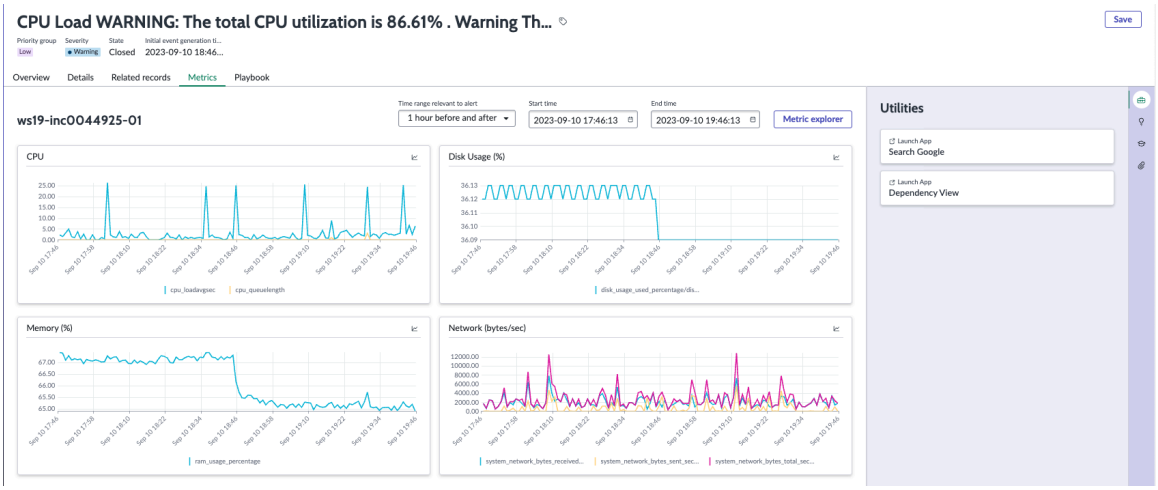

   To customize the metrics display, configure the following properties on the **System Properties** > **All Properties** page:

   - *ws_metric.widgets_number_limit*: Specify the number of widgets/records that can be displayed on the metrics tab.
   - *ws_metric.datasets_per_widget_limit*: Specify the number of charts/datasets that can be displayed on a chart.

5. In the **Time range relative to alert** list, select the time range for which you want the metrics to be displayed for the alert:
   - 1 hour before and after (default setting)
   - 3 hours before and after

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 5 hours before and after
- Custom range (to configure a customized time range for the metrics to measure)

6. Click a parameter below a chart to hide the specified parameter's data on the chart.
7. **Optional:** Click the **Insights Explorer** button to view metrics by dragging and dropping them onto a timeline.

Related topics

View metric values in the Insights Explorer

View anomaly alert metric data on the preview panel in Express List

View visualizations for anomaly alerts to investigate anomalies using metric data. You can view visualizations from the Express List preview panel or the alert record.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. In the Active alerts list, locate an anomaly alert.
   Information in the **Description** column can help you determine which alerts are anomaly alerts. For example, an alert description might warn about values nearing or exceeding configured thresholds, or metrics above or below a defined boundary.
4. Choose where to view the anomaly chart.

Option	Procedure
To view the chart in the preview panel	**a.** Select the check box for the anomaly alert. **b.** In the preview panel, select the **Info** tab. A chart with a visual representation of the anomaly appears.
To view the chart in the alert record	Select the number of the anomaly alert to open the alert record. The **Overview** tab opens by default. The anomaly chart is displayed.

Depending on how the metric anomaly is configured, one of two charts is displayed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- When the Metric Intelligence statistical model is used to detect anomalies, the anomaly alert chart is displayed with upper and lower bounds based on machine learning models.

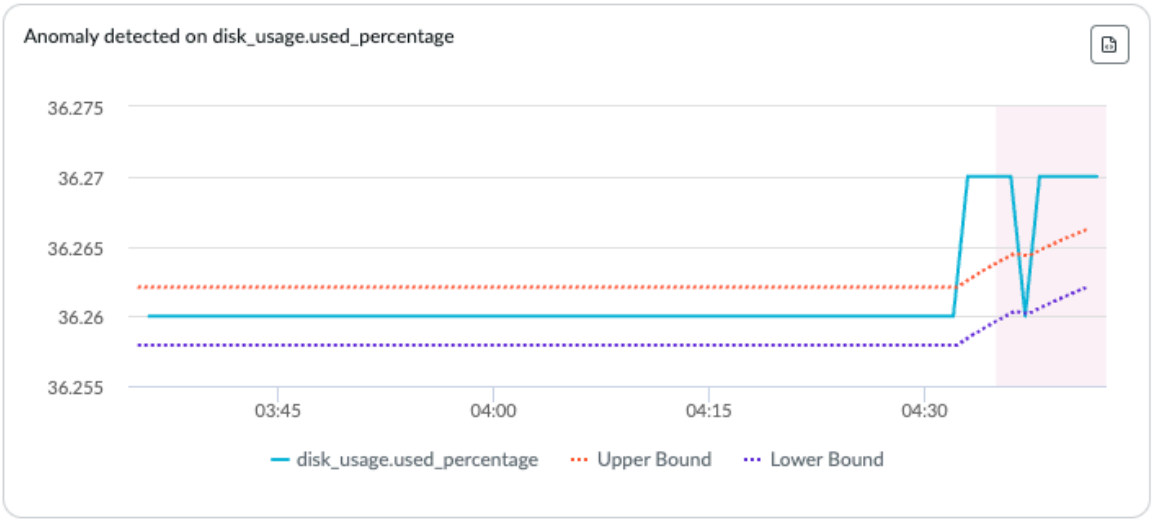

  For more information, see Understanding Metric Intelligence.

- When thresholds are configured by administrators, the static threshold metric anomaly alert chart shows anomalies defined using the configured thresholds.

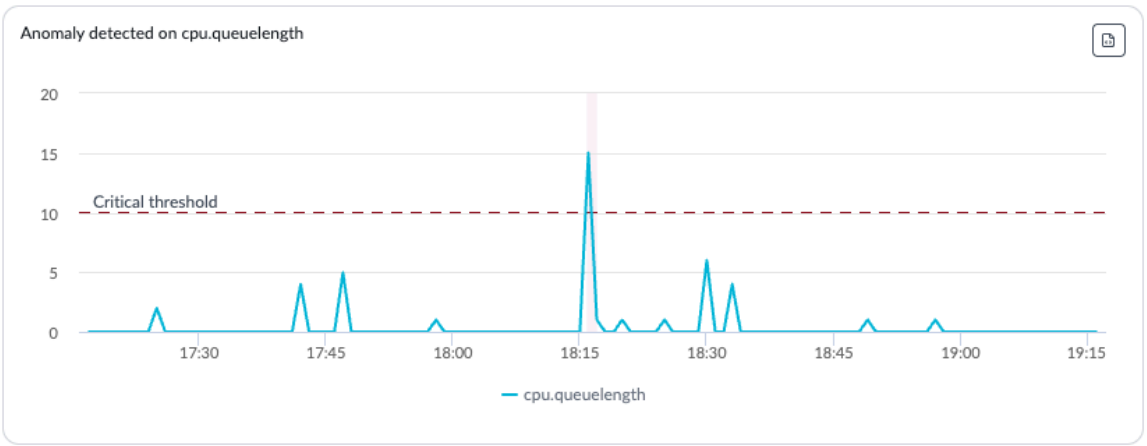

Note:

The raw data used for the metric chart in the preview panel is available only for seven days. If older alerts are selected and the raw data is no longer available, a chart isn't shown.

For more information, see Create metric rules.

5. **Optional:** Review surrounding metric data by selecting the **Open in Metric explorer** icon ( ) in the information panel.

   (Optional) The **Metric Explorer** tab displays the time frame of the anomaly. For an open alert, the chart shows one hour before and after the last time of event generation. For a closed alert, the chart shows one hour before and after the first event that created an alert. If there's no data, the chart isn't displayed.

   For more information, see Metric Explorer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View Log Analytics anomaly alert charts on the preview panel in Express List

View visualizations for Health Log Analytics anomaly alerts to identify periods of behavior that deviate from expected ranges. You can view visualizations from the Express List preview panel or the alert record.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. In the Active Alerts list, locate a Log Analytics based alert.
   This option is available only for individual alerts or the underlying individual alerts within a group.
4. Choose where to view the anomaly chart.

Option	Procedure
View the chart in the preview panel	**a.** Select the check box for a Log Analytics based alert. **b.** In the preview panel, select the **Info** tab. A chart with a visual representation of the anomaly appears.
View the chart in the alert record	Select the number of the anomaly alert to open the alert record. The **Overview** tab opens by default. The anomaly chart is displayed.

5. **Optional:** Select the parameter to visualize anomaly data within that time period:
   - **Events per minute** - The default setting, which shows the detected anomaly.
   - **Day ago**
   - **Week ago**
6. **Optional:** Review relevant logs in context by selecting the Open in Surrounding logs icon: .
   The **Surrounding logs** tab displays the list of log lines that were generated one minute before and one second after the Log Analytics alert. For more information, see How Health Log Analytics generates alerts.

Run actions to resolve alert issues in Service Operations Workspace for ITOM

Resolve the underlying issue of an alert by running predefined remediation tasks, ensuring faster issue resolution and ensuring consistent resolution.

Before you begin

Role required: evt_mgmt_operator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Select the alert number to open it.
   The alert opens in a separate tab, displaying the **Overview** tab.
4. Select **Launch playbook**.
5. Select the **Playbook** tab.
   The **Run remediation** card displays the available remediation actions.

Playbook tab

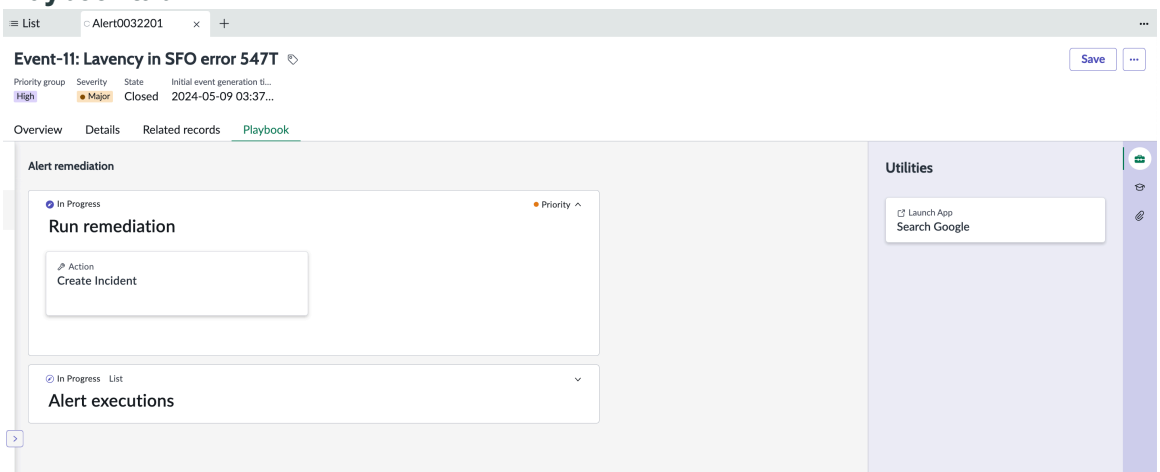

6. Select the tile corresponding to the action you want to run.
   The selected action appears on the **Alert executions** card.

Result

When the action is complete, it appears in the **Completed** card underneath the Alert remediation section.

Launch a web application for an alert in Service Operations Workspace for ITOM

Launch a web application that was configured in an alert management rule to quickly access necessary tools, improving response times and efficiency.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Select the alert number to open it.

   The alert opens in a separate tab, displaying the **Overview** subtab. In the right panel, the **Utilities** tab is selected by default.

4. In the Utilities section, select the web application to launch the app.

> **Note:** For information on how to add a web application that can be launched, see Create an alert management rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Working with alert groups in Express List

Learn how to manage alert groups.

Viewing an alert group analysis by ServiceNow Otto in Express List

View an analysis of alert groups in Express List, generated using AI. The analysis helps you better understand the nature of the alert group, why the alerts in the group were correlated, and how to proceed in the remediation process.

The AI-driven alert group analysis offers a simplified, human-readable description of the group and technical information to help you investigate it more efficiently. The information provided is based on descriptions and Configuration Item (CI) details relating to the individual alerts in the group. Alert group analysis is supported for the following alert group types:

Tag-based alert groups

Tags help categorize alerts based on common attributes, such as impacted systems or services. Alerts in tag-based alert groups share certain tags, indicating similarities in the issues. The alert analysis for these groups presents the shared tags that were used to correlate the alerts in the group. It also provides insights into why these alerts were grouped. In addition, the analysis offers suggestions for a course of action based on the similarities between the tags.

CMDB alert groups

When an alert is created on a CI in the Configuration Management Database (CMDB), Event Management looks for alerts on other CIs that are closely related to it in the CMDB topology. A close topological relationship between CIs suggests an interdependence between components in the IT infrastructure and contributes to alert correlation. When a relationship between CIs is found, a CMDB alert group is formed. The alert analysis for CMDB alert groups explains the nature of the group and why it was formed, and the relationships between the CIs. It also provides technical information to help you decide on which alerts to concentrate your investigation.

Log Analytics alert groups

When Event Management identifies multiple Log Analytics alerts that are related in important ways, it groups them into a Log Analytics group. The system generates a Log Analytics group when the Log Analytics alerts share one or more relationships related to time, metadata, message text, and trend. The alert analysis for Log Analytics alert groups is based on anomaly information provided by Health Log Analytics. It also considers the descriptions, CIs, and tags of the alerts in the group.

Network Traffic-based alert groups

The Network Traffic-based alert grouping method groups alerts by analyzing network traffic connections between processes across hosts. It leverages service candidates identified by ML Service Mapping to group the alerts. Alert analysis for network traffic-based alert groups gives a description for all the alerts in the group and the connections between them.

Mixed alert groups

The Mixed grouping method combines alerts using multiple grouping strategies, such as those used in CMDB alert groups and tag-based groups, in a single, cohesive group. This method leverages the strengths of each strategy to reduce alert noise, improve alert correlation, and highlight the true root cause of incidents.

Automated alert groups

Automated alert grouping uses aggregation algorithms to identify and group related alerts. The algorithms group alerts that share the same alert identifier, such as the CI and metric identifier, and that occurred multiple times within a similar time frame.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The algorithms rely on historical data to detect recurring patterns and organize similar alerts. Alert analysis for automated alert groups provides details about the CI or time-based patterns that contributed to the grouping. The analysis also includes a summary of historical incidents related to the group, including frequency, criticality, and resolution strategies.

View an alert group analysis by ServiceNow Otto in Express List

View an alert group analysis created using generative AI. The analysis offers a simplified, human-readable description of the alert group and technical information to help you investigate it more efficiently.

Before you begin

- Install ServiceNow Otto for ITOM. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.
- View important information about the ServiceNow Otto for IT Operations Management (ITOM) application in AI in IT Operations Management.

Note: Currently, ServiceNow Otto for ITOM only supports tag-based, CMDB, Log Analytics, Mixed, Automated, and Network Traffic-based alert groups. For all other alert group types, it only analyzes the parent alert.

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the navigation bar, select the Express List icon ( ).
3. In the Active alerts list, select the information icon ( ) next to an alert group to display the preview panel.
4. On the preview panel **Info** tab, select **Analyze** in Alert analysis by ServiceNow Otto.
5. View the information provided in the Alert analysis.
6. **Optional:** Use the Alert analysis icons to perform related tasks.

Alert analysis icons

Icon	Description
	Copy the content of the alert analysis to the clipboard.
	Refresh the alert analysis. **Note:** Refreshing regenerates the results. Past results are deleted.

Related topics

Viewing an alert group analysis by ServiceNow Otto in Express List

Update an alert description in Express List based on a ServiceNow Otto for ITOM analysis summary

Update alert descriptions in alert groups or single alerts in Express List, to a human-readable summary generated by AI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Updating the alert description is available for the alert types:

- Tag-based alert group.
- CMDB correlation alert group.
- Network Traffic correlation alert group.
- Log Analytics single alerts and alert group.
- Mixed alert group.
- Automated alert group.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the navigation bar, select the Express List icon ( ).
3. In the Active alerts list, select the information icon ( ) of an alert.
4. On the preview panel **Info** tab, select **Analyze** in Alert analysis by ServiceNow Otto.
5. Review the human-readable, AI-generated summary in the Alert analysis.
6. On the preview panel **Info** tab, select **Update description** to update the description based on the summary content of the alert group.
   The AI-generated summary content replaces the original description of the alert group.

View an alert group analysis by ServiceNow Otto in Express List

View an alert group analysis created using generative AI. The analysis offers a simplified, human-readable description of the alert group and technical information to help you investigate it more efficiently.

Before you begin

Install the ServiceNow Otto for ITOM plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the Express List icon ( ) in the left navigation bar.
3. In the Active alerts list, select the information icon ( ) to the left of an alert group to display the preview panel.
4. On the preview panel Info tab, select **Analyze** in Alert analysis by ServiceNow Otto.
5. View the information provided in the Alert analysis.
6. **Optional:** Use the Alert analysis icons to perform related tasks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert analysis icons

Icon	Description
	Copy the content of the alert analysis to the clipboard.
	Refresh the alert analysis. **Note:** Refreshing regenerates the results. Past results are deleted.

Related topics

Viewing an alert group analysis by ServiceNow Otto in Express List

Generate an alert group description in Express List using ServiceNow Otto

Use AI to generate a meaningful description of an alert group in Express List that encompasses all the alerts within the group. The generated description replaces the original description of the group.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the Express List icon ( ) in the navigation bar.
3. In the Active alerts list, select the information icon ( ) for an alert.
4. Below the displayed alert group description, select **Update description**.
   ServiceNow Otto generates a comprehensive new description of the alert group, which encompasses all the alerts within the group. If you select **Update**, the generated description replaces the original description of the alert group taken from the earliest, most severe secondary alert in the group.
5. **Optional:** Restore the original alert group description by selecting the revert icon ( ) next to **Update description** and then choosing **Restore**.

Save the ServiceNow Otto alert analysis as an alert group description

Save the human-readable summary generated by ServiceNow Otto in the Alert analysis as the description of an alert group.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the Express List icon ( ) in the left navigation bar.
3. In the Active alerts list, select the information icon ( ) to the left of an alert.
4. On the preview panel Info tab, select **Analyze** in Alert analysis by ServiceNow Otto.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Review the human-readable, AI-generated Alert analysis summary.
6. Select **Save description** to save the summary content as the description of the alert group. The AI-generated summary content replaces the original description of the alert group.

Viewing links between alerts in alert groups in Express List

Gain a better understanding of the relationships between alerts in alert groups in the Express List by using Link View. Link View offers a visual representation of the relationships between the alerts in a group.

When Event Management generates an alert group, Link View shows how the attributes of the alerts in the group are linked. The colored tags represent configuration items (CIs) and other environment items in relation to the alerts.

The information shown in Link View is available without the need for a populated Configuration Management Database (CMDB). However, when the CMDB is populated, Link View offers additional value by providing the probable cause of the alerts and the service that the alert group impacts.

Sample alert group in Link View

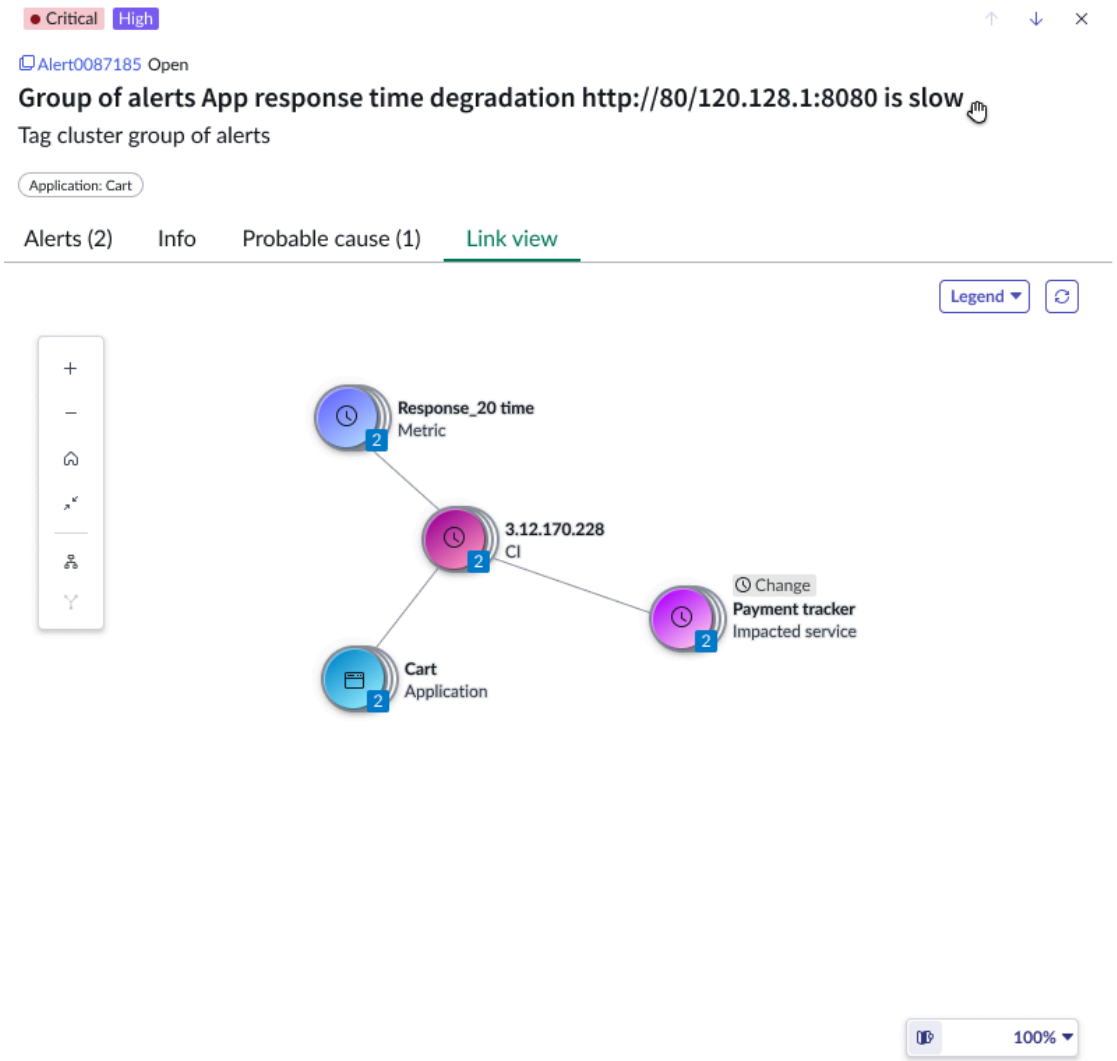

You can focus on your areas of interest by dragging the nodes in Link View to different positions. When you refresh an alert group, rearranged nodes appear in their original position again. Therefore, Link View is not refreshed automatically, but waits for you to do so manually.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If an alert on a CI impacts a service in the Configuration Management Database (CMDB), Link View shows the impacted service. You can view it at a glance for quick triage.

A stacked node indicates that multiple nodes were mapped for the same tag. When the same key-value pair appears in more than one alert, the corresponding node is shown with a badge. For example, when the same key-value pair appears in two alerts, the badge on the node shows the number 2. See the Payment tracker node in the sample alert group figure. When a node has no badge, the key-value pair appears in only one alert. An active change request, a probable cause of the alert, is marked by a Change badge.

Node with a change badge

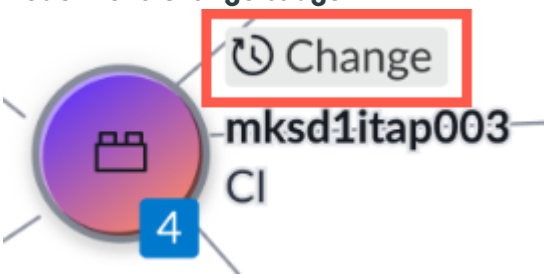

The Link View legend lists the meaning of the symbols and colors used and enables you to toggle between hiding and showing types of tags to reduce noise. In addition, the legend describes the meaning of the various lines linking the alert attributes. Attributes linked by a solid line share one or more alerts, whereas attributes linked by a dotted line are correlated by grouping criteria. For a description of each tag, see Attributes in Express List Link View. Hovering over a node displays a tooltip that includes the tag name, class, and severity. It also shows the number of alerts in which the tag appeared. If applicable, the tooltip indicates whether the alert is primary or secondary or the probable cause.

Node tooltip

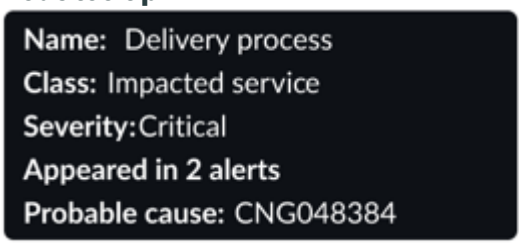

Currently, Link View is supported for several alert groups. For more information, see the following topics:

- Viewing links between alerts in tag-based alert groups
- Viewing links between alerts in rules-based alert groups
- Viewing links between alerts in CMDB-based alert groups
- Alert links in network traffic-based alert groups
- Alert links in log analytics-based alert groups
- Viewing links between alerts in mixed alert groups

Related topics

View links between alerts in a group in Express List

Viewing links between alerts in tag-based alert groups

View the connections between alerts in a tag-based alert group in Express List by using Link View. Link View shows how the attributes of the alerts in the group are linked with each other.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When you've set alert tags and Event Management has generated an alert group based on tag-based rules, Link View offers a visual representation of the relationships between the alerts. The colored tags represent Configuration Items (CIs) and other environment items in relation to the alerts.

Sample tag-based alert group in Link View

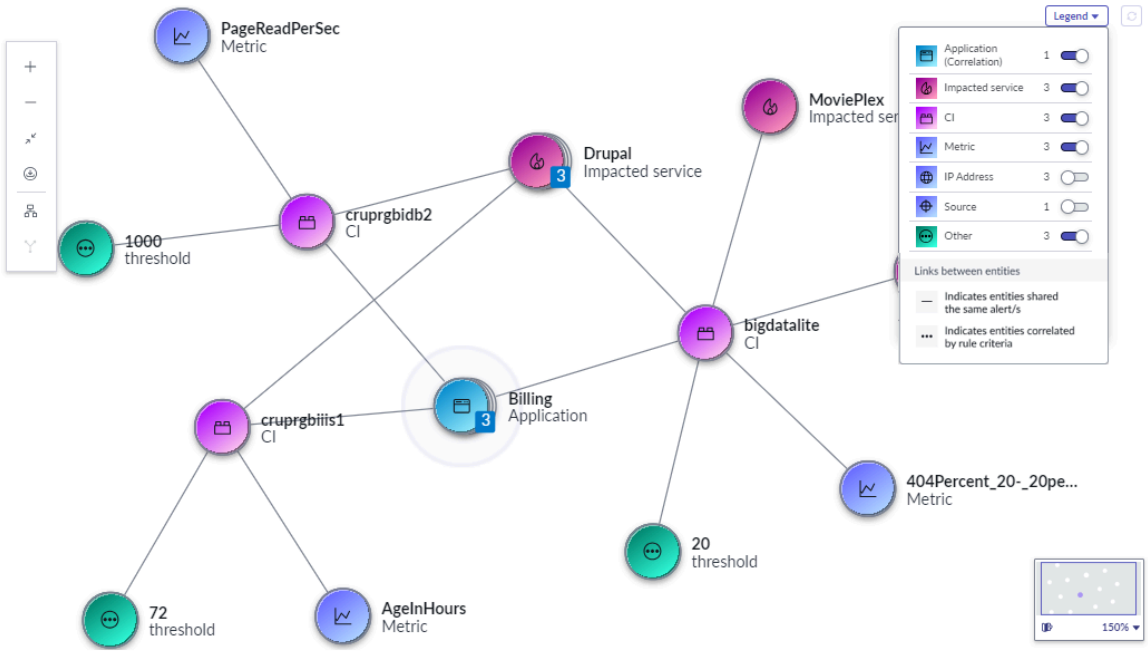

In this example, the defining tag for the correlation, Application, is linked to the tags of the relevant CIs. However, a correlation can be defined by any tag or combination of two tags, such as Resource and Metric name. When the defining tag is missing from Link View, the correlation was based on information that is not represented in the view, such as the description. Each CI node is further linked to the tags that appeared on the same alert as that CI.

In an alert group that was generated from tag rules, nodes connected by a dashed line signify an approximate, "fuzzy" match. This means that the nodes were correlated when the match on the tag rule was not exact, but close enough to be significant. For example, this could happen when a location, email address, or IP pattern was similar enough to qualify as a match.

The Link View legend lists the meaning of the symbols and colors used to represent the tags and their number of unique nodes. In the legend, defining tags are marked as Correlation. You can toggle between hiding and showing tag types to reduce noise. For a description of each tag, see Attributes in Express List Link View.

Related topics

- Viewing links between alerts in alert groups in Express List
- View links between alerts in a group in Express List

Viewing links between alerts in rules-based alert groups

View the connections between alerts in a rules-based alert group in Express List by using Link View. Link View shows how the attributes of the alerts in the group are linked with each other.

When you've set alert rules and Event Management has generated an alert group based on these rules, Link View offers a visual representation of the relationships between the alerts. The colored tags represent Configuration Items (CIs) and other environment items in relation to the alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample rules-based alert group in Link View

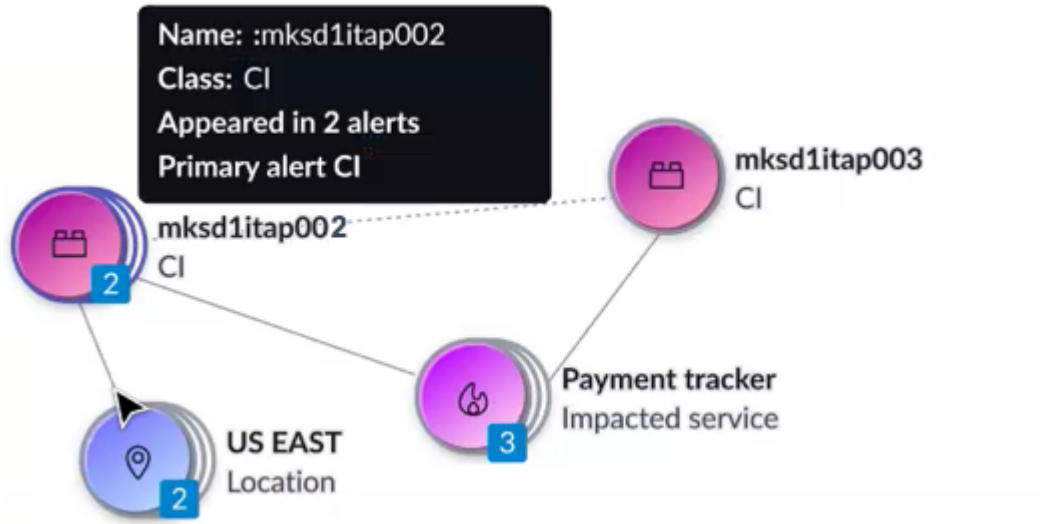

Rules-based alert groups are comprised of a primary alert and one or more secondary alerts. The primary alert is generated directly and is represented in Link View as the Home node. Secondary alert nodes are connected to the Home node by a dashed line, indicating that they don't share the same alert, but are correlated using rule criteria. A solid line between nodes indicates that the connected nodes share the same alert or multiple alerts.

If an alert on a CI impacts a service in the Configuration Management Database (CMDB), Link View shows the impacted service. You can view it at a glance for quick triage. Hovering over a node displays a tooltip. The tooltip includes the tag name, type, number of alerts in which it appeared, and whether it is a primary or secondary alert.

Related topics

- Viewing links between alerts in alert groups in Express List
- View links between alerts in a group in Express List

Viewing links between alerts in CMDB-based alert groups

View the connections between alerts in alert groups in Express List that were created based on the proximity of Configuration Items (CIs) in the Configuration Management Database (CMDB). Link View shows how the attributes of the alerts in the group are linked with each other.

When alerts on CIs that are located near each other in the CMDB topology occur within a certain time-frame, Event Management generates a CMDB-based alert group. Event Management correlates CIs in a service map when they are separated by a maximum of three other CIs. The colored tags in Link View offer a visual representation of the CIs and other environment items in relation to the alerts in the group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample CMDB-based alert group in Link View

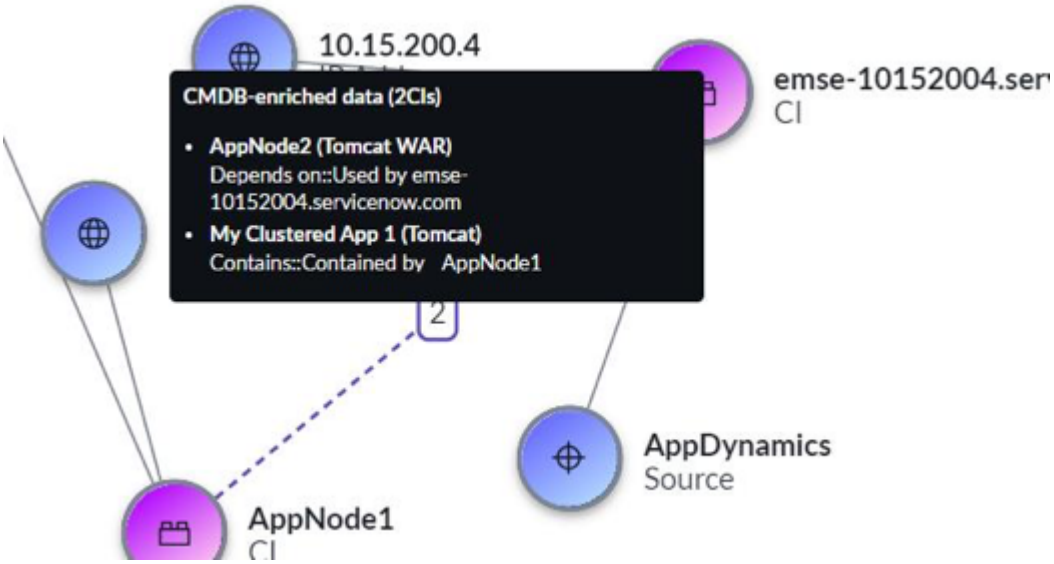

In this sample Link View of a CMDB-based alert group, the dashed line with the number 2 between the CIs indicates that the CIs are not connected directly. They are separated by two CIs in the service map in the CMDB topology. The sample CMDB service map also shows correlated CIs separated by two CIs.

Sample service map with correlated CIs

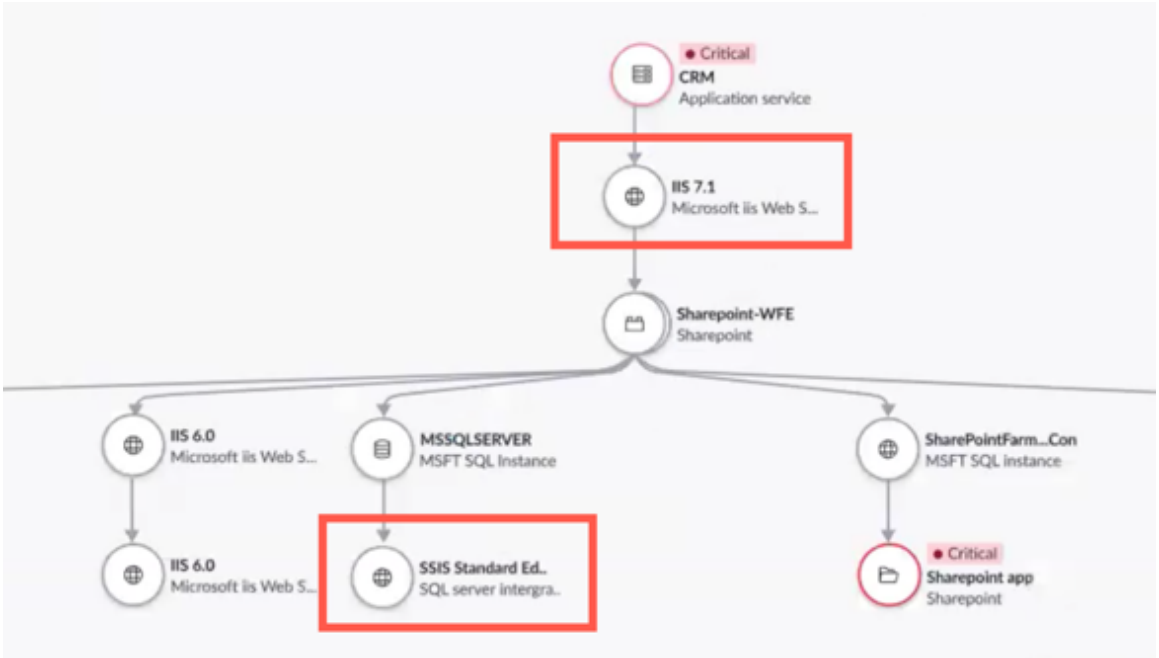

Related topics

- Viewing links between alerts in alert groups in Express List
- View links between alerts in a group in Express List

Alert links in network traffic-based alert groups

View the connections between alerts in network traffic-based alert groups in Express List by using Link View. Network traffic-based alert groups are created by analyzing network traffic connections between processes across hosts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service candidates identified by ML Service Mapping, group alerts based on network traffic connections between processes. Link View shows a contextual view of network incidents based on these groups of directly connected processes. Service candidates are potential collections of processes within your IT environment that are identified based on their network connections and interactions.

Sample network traffic-based alert group in Link View

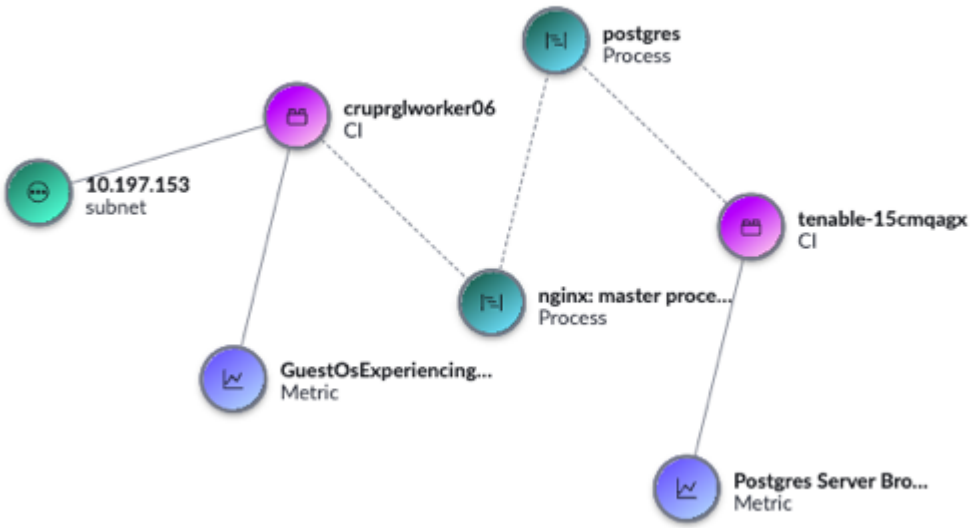

In this sample Link View of a network traffic-based alert group, the dashed lines indicate that there are connected processes running on the CIs. Attributes linked by a solid line share one or more alerts.

Related topics

- Viewing links between alerts in alert groups in Express List
- View links between alerts in a group in Express List

Alert links in log analytics-based alert groups

Use Link View in Express List to see why alerts in a log analytics-based alert group are correlated together through the visualization of shared attributes.

Link View for log analytics-based alert groups shows the connections between component-based alerts. Each component-based alert may involve multiple host configuration items (CIs). A component represents a logical part of a service instance where the alert was detected. Components can consist of multiple CIs that perform the same function, such as multiple redundant hosts. For more information, see Types of Health Log Analytics alerts.

The colored tags in the link view represent attributes of the alerts such as hosts, users, or correlation indicators extracted from the logs. The lines connect attributes that share the same alert. This visualization helps you understand the relationships between different alerts in the group using their shared attributes.

Log analytics-based groups often contain multiple correlation attributes within a single group. Link View enables you to explore why anomalies detected on different components were correlated into one group. By manipulating the attribute display, you can better understand the relationships between alerts that share common attributes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample log analytics-based alert group in Link View

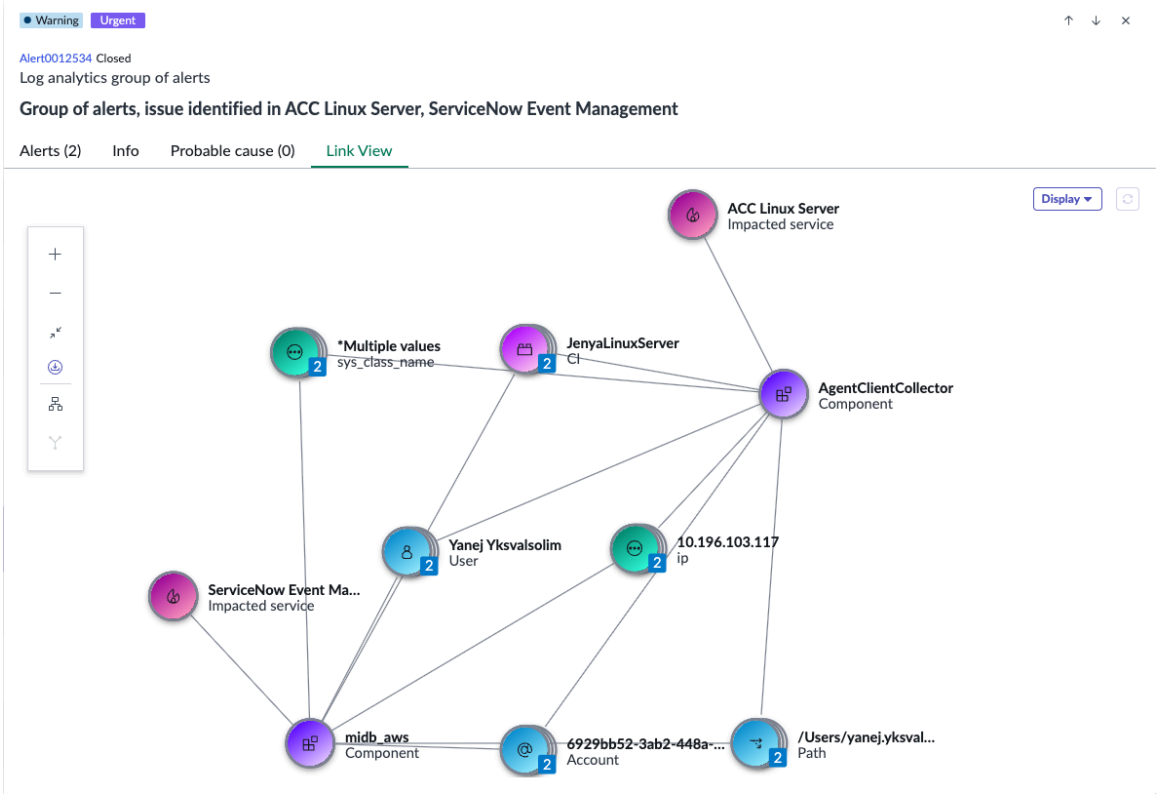

In this sample Link View, a component attribute is shown for each child alert in the log analytics alert group. Other attributes show connections between the component alerts.

The Link View legend lists the meaning of the symbols used to represent the attributes. You can toggle between hiding and showing attributes types to reduce noise. For a description of each attribute, see Attributes in Express List Link View.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

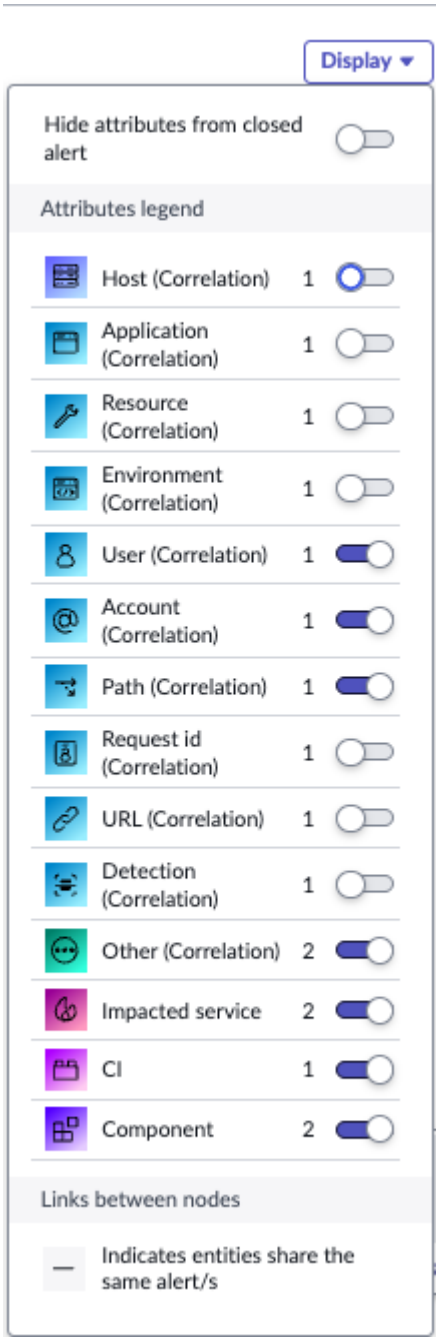

Viewing links between alerts in mixed alert groups

Link View displays the connections between alerts in mixed alert groups in Express List, including how alert attributes relate to each other within the group.

Link View for mixed alert groups shows the connections between alerts in groups that combine different alert group strategies to form a single cohesive group.

To view how different alert groups are visualized in Link View, see the description of Link view for the relevant alert group types. Currently, mixed groups are composed of Tag cluster alert groups, CMDB-based alert groups, related log entities and shared impacted services.

For more information about alerts in the mixed alert group, see Mixed alert grouping.

For more information about Link View for related groups, see Viewing links between alerts in tag-based alert groups, and Viewing links between alerts in CMDB-based alert groups.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample mixed alert group in Link View

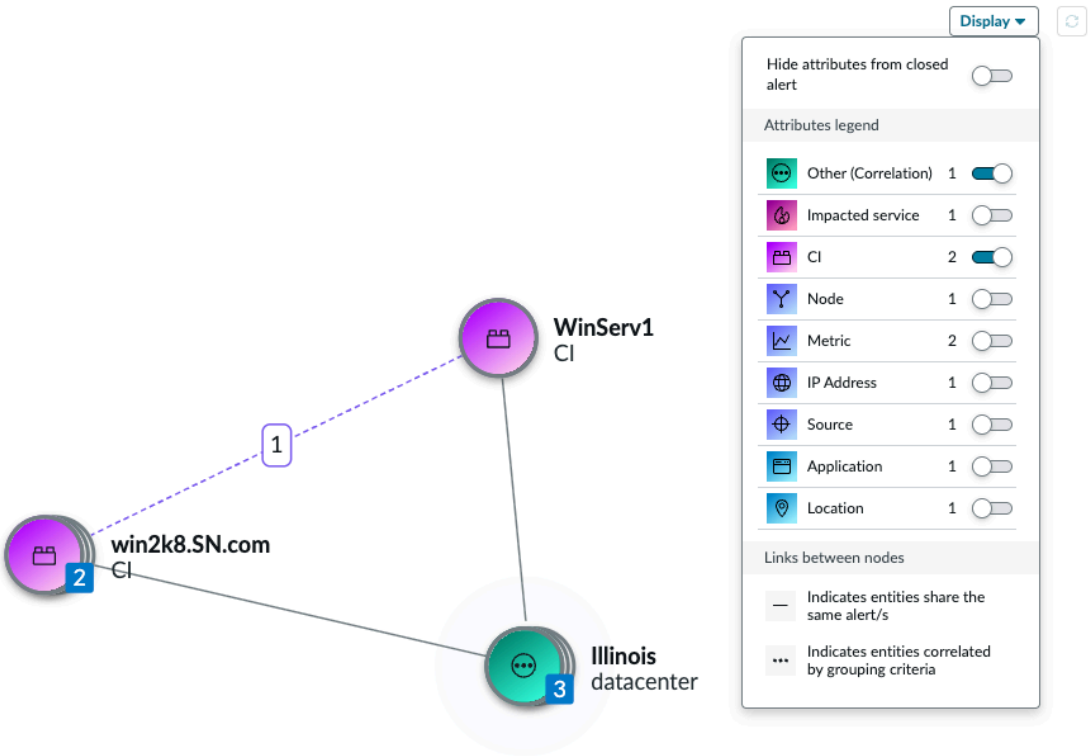

The preceding figure shows the Link View of a mixed alert group. The dashed line with the number 1 between the CIs indicates that the CIs aren't connected directly. They are separated by a CI in the service map in the CMDB topology. The sample also shows the defining attribute for the correlation, datacenter, linked to the relevant CIs.

Sample mixed alert group with related log entities and shared impacted services in Link View

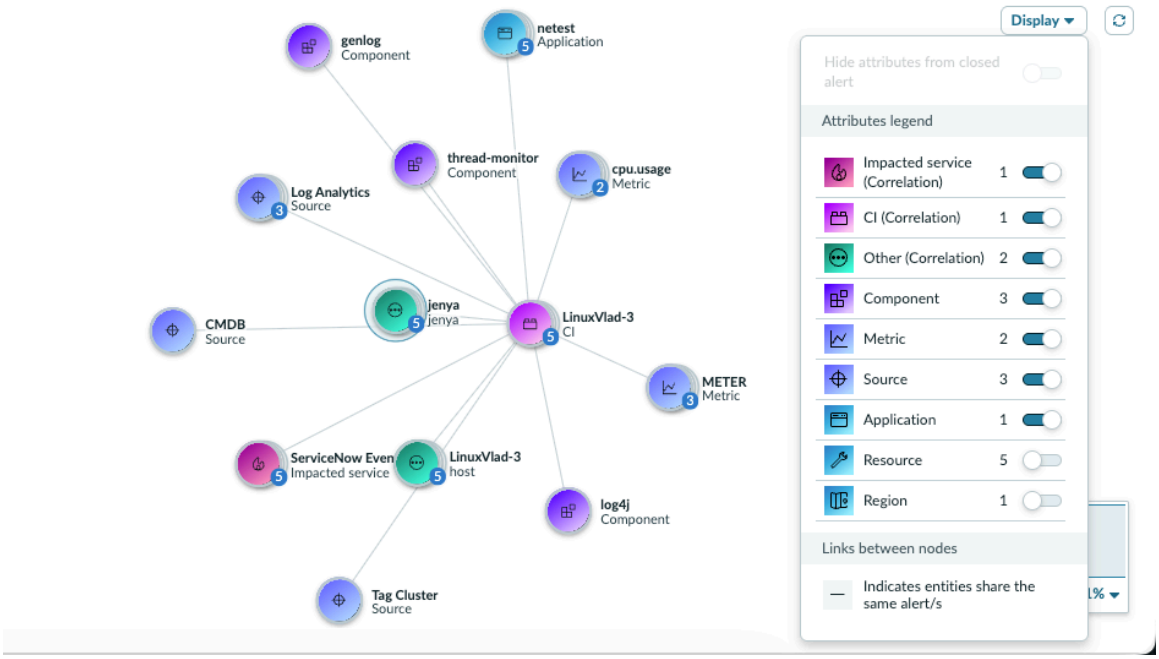

The Link View for a mixed alert group that includes related log entities and shared impacted services displays connected nodes. Each node represents a correlation source and attribute,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

including log analytics data, CMDB information, and impacted services. This information gives operators a picture of the scope and nature of the grouping at a glance.

The Link View **Display** legend lists the meaning of the symbols used to represent the tags and their number of unique nodes. In the legend, defining tags are marked as **Correlation**. You can toggle between hiding and showing tag types to reduce noise. For a description of each tag, see Attributes in Express List Link View.

Related topics

Viewing links between alerts in alert groups in Express List

View links between alerts in a group in Express List

View links between alerts in a group in Express List

Use Link View in Express List to view a visual representation of the relationships between alerts in a group.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task

For an overview of Link View in Express List, see Viewing links between alerts in alert groups in Express List.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the primary navigation, select the Express List icon ( ).
3. In the fields in the interactive filter panel, select **Group** and then the check box for the alert groups to display.

   Note: Link View is supported for the following alert group types:
   - Network traffic-based
   - Rules-based
   - Mixed alert groups, which include:
     - CMDB-based alert groups
     - Tag cluster-based alert groups
     - Related log entities
     - Shared impacted services

4. In the **Active alerts** list, select the description of an alert group.
   The preview panel opens to the **Alerts in group** tab, which lists all the correlated alerts in the selected group.
5. Navigate to **Link View**.
   Link View opens, displaying the relationships between the alerts in the selected group.
6. **Optional:** Customize the Link View display.

Task	Action
Focus on an area of interest	Select one or more nodes and rearrange them in Link View by dragging them to a new location.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Task	Action
Refresh the alert group	Select **Refresh**. After refreshing the alert group, rearranged nodes appear in their original position again. Newly added nodes are marked as New. **Note:** The **Refresh** button is enabled when new data for the alert group is available. Link View doesn't refresh auto matically.
View the meaning of icons and colors	Select the Link View legend. (Optional) The legend also indicates the num ber of unique nodes displayed per tag. For a description of each tag, see Attributes in Ex press List Link View.
Reduce noise	In the Link View legend, toggle between hid ing and showing a tag type.
View information about an alert	Hover over a node to display text with informa tion about the alert.

Chronological alert data in an alert group

Visualize the chronological sequence of events within an alert group in Express List using the Timeline view. This feature provides a comprehensive overview of when the alerts occurred, their severity changes, and other pertinent data for efficient triage and Mean Time to Resolution (MTTR).

In Timeline view, each alert is represented by a bar, enabling you to quickly discern the order of occurrence, severity transitions, and closure status. The timeline begins with the creation of the first alert and extends to the present moment.

Sample timeline view

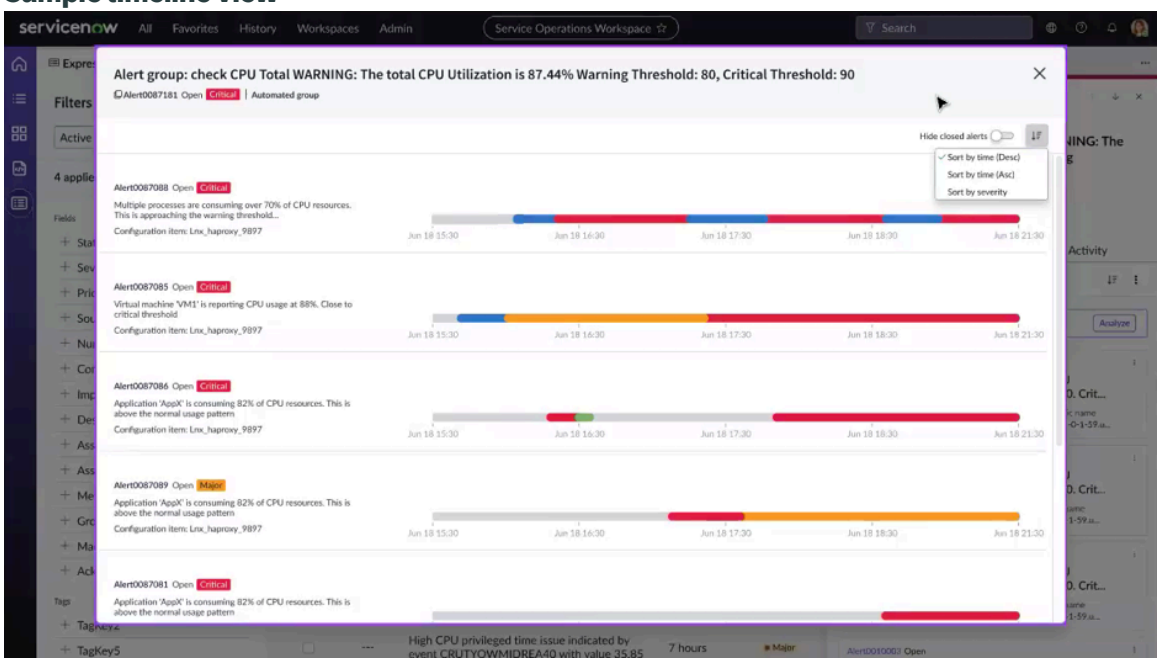

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Standard severity colors are used in the bars, with gray segments indicating periods before alert creation or after its closure. Hovering over a severity bar reveals a tooltip with key alert details.

Severity range by color in Timeline view

Color	Severity
	Critical
	Major
	Minor
	Warning
	OK

The title of the Timeline view for an alert group summarizes the group's description, number, state, severity, and grouping type, derived from the primary alert. Detailed information including alert number, state, severity, a brief description, and associated Configuration Item (CI) or node, when available, is displayed for each alert within the group.

Related topics

View a timeline of the alerts in an alert group

View a timeline of the alerts in an alert group

Gain insight into the sequence of events relating to an issue by viewing chronological information of the alerts in an alert group in Express List.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

About this task

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the primary navigation, select the Express List icon ( ).
3. In the Active alerts list, select the information icon (for an alert.
4. On the preview panel Alerts tab, select **Timeline view**.
5. Analyze the information provided in the Timeline view.
6. **Optional:** Sort the displayed alerts by ascending or descending time or by severity.
   a. Select the sorting icon ( ).
   b. Select the preferred sorting option from the pop-up menu.
7. **Optional:** Show or hide closed alerts by toggling the **Hide closed alerts** slider.

Related topics

Chronological alert data in an alert group

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remove alerts from an alert group in Express List

Remove secondary alerts from an alert group directly from the Express List pane or from the preview panel.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and does not have permission to perform this action.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Remove a secondary alert from an alert group.
   You can remove alerts from a group regardless of whether you are working in the main pane or in the preview panel. The preview panel opens only when a single alert or alert group is selected.

Location	Action
The Express List pane	**a.** In the Active alerts list, select the chevron icon ( ) at the beginning of one or more alert group rows to display the alerts con tained in the group. **b.** Select one or more secondary alerts. **c.** From the **Close** drop-down list at the top right of the list, select **Remove from group**. **Note:** When you are trying to remove the only secondary alert from an alert group, thereby ungrouping the group, a message displays that enables you to undo the action.
The Express List preview panel for a group of alerts	**a.** In the **Alerts in group** tab, select an alert tile. **b.** Select the more actions icon ( ). **c.** Select **Remove from group**.

Promoting alerts in Express List

Learn how to promote an alert to an incident or generate a case.

Create an incident with ServiceNow Otto in Express List

Create an incident with a human-readable, AI-generated description.

Before you begin

Currently, this feature is available for CMDB, Log Analytics, tag-based, and component-based alerts.

Role required: evt_mgmt_operator, evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The evt_mgmt_user role is view-only and doesn't have permission to perform this action.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Create an incident from a selected alert.
   a. In the Express List pane, select the check box for the desired alert.

      Note: To display the individual alerts inside a group, select the chevron icon (›) at the beginning of the alert group row.

   b. Select the **Alert actions** drop-down list.

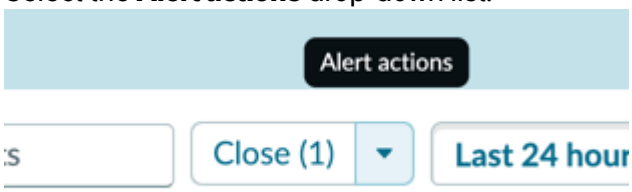

   c. Under **Response actions**, select **Create Incident with Otto**.

   An incident with a human-readable, AI-generated description is created from the selected alert and a confirmation message is displayed.

Create an incident from an alert with ServiceNow Otto in Express List

Create an incident with a human-readable, AI-generated description from the Express List pane by using AI.

Before you begin

Note: Currently, ServiceNow Otto for ITOM only supports tag-based, CMDB, Log Analytics, Mixed, Automated, and Network Traffic-based alert groups. For all other alert group types, it only analyzes the parent alert.

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Create an incident from a selected alert.
   a. In the Express List pane, select the alert.

      Note: To display the individual alerts inside a group, select the chevron icon (›) at the beginning of the alert group row.

   b. Select the **Alert actions** drop-down list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

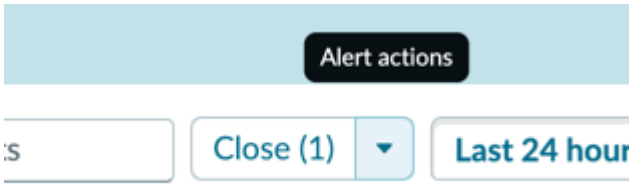

c. Under **Response actions**, select **Create Incident with ServiceNow Otto**.

An incident with a human-readable, AI-generated description is created from the selected alert and a confirmation message is displayed.

Create an incident from an alert in Express List

Create an incident from an alert directly from the Express List pane, saving you the time and effort of navigating to the incident record.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Note: The evt_mgmt_user role is view-only and doesn't have permission to perform this action.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation bar, select the Express list icon .
3. Create an incident from a selected alert.

a. In the Express List pane, select the check box for the desired alert.

Note: To display the individual alerts inside a group, select the chevron icon ( ) at the beginning of the alert group row.

b. Select the **Alert actions** drop-down list.

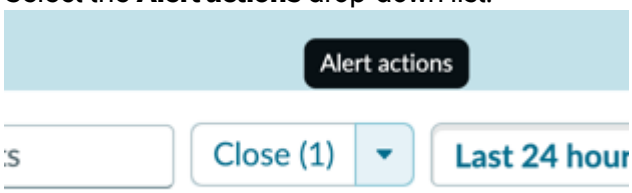

c. Under **Response actions**, select **Create Incident**.

An incident is created from the selected alert and a confirmation message is displayed.

Generating case from an alert

Enable swift resolution by generating cases from alerts, either manually or automatically.

Customer Service Management provides integration with ITOM Event Management. This integration enables you to improve the customer experience by breaking down the silos that exist between the processes and systems used by front line customer service and back-office

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

operations teams. By monitoring issues and creating cases proactively, you can better correlate customer issues and provide faster responses.

If you are using the integration between the Customer Service Management and Event Management applications, you can create a case for the customer install base affected by an alert. For more details, see Manually create a proactive case from an alert. You can also activate an alert rule to automatically create a proactive case from an alert when an incident is created. For more details, see Activate an alert rule to automatically create a proactive case from an alert.

Express List reference

Reference topics provide additional information for configuring and using Express List.

Roles used by Express List

Detailed information on all roles that can access the Express List feature on the Service Operations Workspace.

Role title	Description	Contains roles
Event Management Administrator [evt_mgmt_admin]	Has read and write access to all Event Management features to configure Event Management. **Note:** Be careful with the evt_mgmt_admin role because it can be elevated to the admin role.	• evt_mgmt_user • template_editor_global
Event Management Operator [evt_mgmt_operator]	In combination with the evt_mgmt_user permissions, can also activate operations on alerts such as acknowledge, close, open incident, and run remediations.	evt_mgmt_user
Event Team Operator [evt_team_operator]	Manages Event Management operations within a specific team as defined in the **Assignment Group** field. This role allows the operator to read and write alerts exclusively for their assigned team. Additionally, the operator can make configuration changes specific to their team, including updates to Alert Automation and Integrations Launchpad.	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title	Description	Contains roles
	Note: The evt_team_operator role must be assigned to an assignment group to view and manage alerts for that group. If the role is created but not associated with any groups that have alerts assigned, the operator cannot see any alerts.	
Event Management User [evt_mgmt_user]	Has read-only access to all Event Management features. The contained itil role grants access to manage incidents that are created from alerts. Can only view Express List.	itil
Event Management Integrator [evt_mgmt_integration]	Has create access to the Event [em_event] and Registered Nodes [em_registered_nodes] tables to integrate with external event sources.	

Alerts in Service Operations Workspace

The Service Operations Workspace interface displays an alerts list and details on specific alerts.

When you click an alert in the alerts list, the **Details** tab of the selected alert appears. The issue that caused the alert appears in the alert title. Only the sub tabs relevant to the alert appear on the resulting page. For example, the **Alerts in Group** option appears in the **Related Records** tab only for alert groups.

View the alerts that Azure grouped into the issue in the alert list. The **Source** column displays A for these alerts.

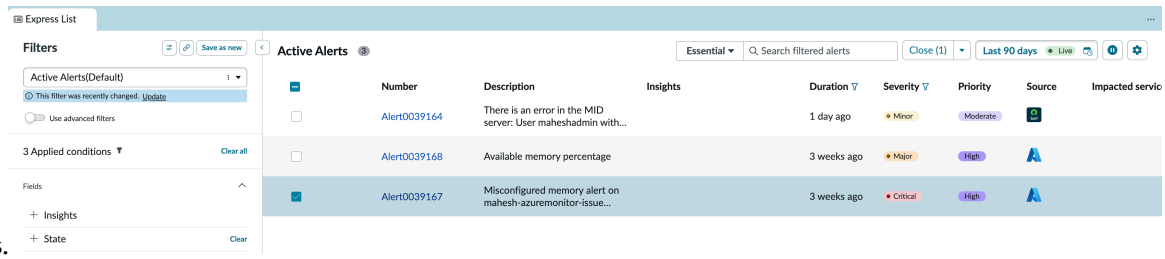

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Details tab

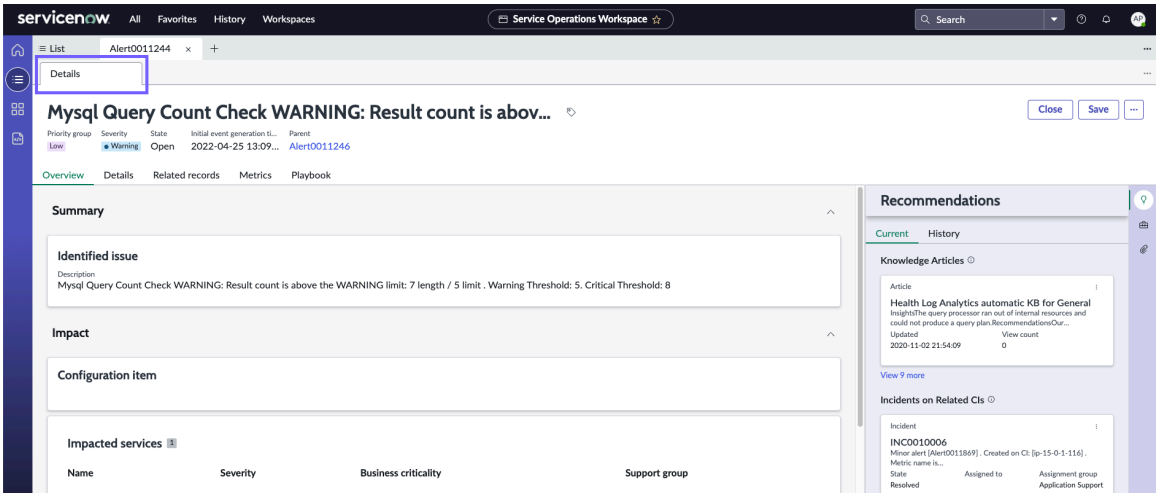

The following table describes areas on the alert form.

Alert form areas

Feature	Description
Header	The header includes these details: • Description: The text from the **Description** field of the alert is displayed. • Priority group: Priority group to which the alert belongs. • Severity: Severity of the alert. • State: The state of the alert. • Initial event generation time: Date and time when the initial event was generated. • Parent: The primary alert in the group (relevant only for child alerts in a group). • Group: Group type to which the alert belongs (relevant only for alerts in a group). • Task: Number of the incident with which the alert is associated. The displayed information varies according to the alert type.
Overview tab	When viewing an alert with an assigned CI, this tab opens when selecting the alert. The information displayed varies depending on the type of alert (grouped, secondary, or primary). If you have installed the Health Log Analytics app, alert types include Health Log Analytics alerts and component-based alerts.
Details tab	Shows the alert's core attributes, including its number, configuration item, class, resource, type, severity, state, and full description of the identified issue. Alongside are assignment fields (assigned to, assignment group, task, parent) plus the work notes and a chronological Activity log tracking state changes and system events.
Related Records tab	Displays lists of records linked to the alert, grouped into categories: events, impacted services, service offerings, changes, configuration items, remediation tasks, incidents, and similar alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert form areas (continued)

Feature	Description
	Note: ServiceNow represents an Azure Issue as a single alert. You can view the alerts that Azure grouped into the Issue in this tab. 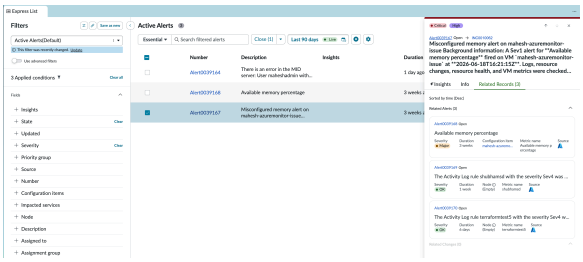

Express List keyboard shortcuts

Use keyboard shortcuts to quickly perform common actions in the Express List.

Navigation shortcuts

Action	Keyboard shortcut
Go to the previous/next row	Arrow up/down
Move focus forward between clickable elements	Tab
Move focus backward between clickable elements	Shift + Tab

Action shortcuts

Action	Keyboard shortcut
Open an alert record	Option/Alt + R
Acknowledge an alert	Option/Alt + K
Assign an alert to yourself	Option/Alt + Y
Close an alert	Option/Alt + X
Create an incident	Option/Alt + I
Search free text	Option/Alt + S
Pause/Restart live updates	Option/Alt + L
Open/close the preview panel	Enter

Selection shortcuts

Action	Keyboard shortcut
Expand/collapse a selected group of alerts	Space
Select multiple neighboring alert rows	Shift + arrow up/down

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Selection shortcuts (continued)

Action	Keyboard shortcut
Select multiple separate alert rows	Command/Control + arrow up/down
Activate a selected link or button	Enter

Note: The information presented in this section is also available as in-product help in the Help Center panel for the Express List.

Attributes in Express List Link View

The table lists the node attributes available in Link View with their icon and description.

Note: "Attribute" is a generic term that represents fields and tags found on the alerts or configuration items (CI).

Attributes in Link View

Attribute name	Icon	Description
Impacted service		Service in the Configuration Management Database (CMDB) impacted by an alert on a CI.
Resource		Node resource that is relevant to the event. For example, Disk C, CPU-1, the name of a process, or service. Maximum length: 100 characters.
Source		Event monitoring software that generated the event, such as SolarWinds or SCOM. Maximum length: 100 characters.
Region		Predefined attribute extracted from connector payload.
Location		Predefined attribute extracted from connector payload.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attributes in Link View (continued)

Attribute name	Icon	Description
Assignment group		Alert assignment group.
Environment		Predefined attribute extracted from connector payload.
Metric		Unique name that describes which metric data is collected.
IP address		Predefined attribute extracted from connector payload.
CI		JSON string that represents a configuration item. For example, `{"name":"SAP ORA01","type":"Oracle"}`. The CI identifier that generated the event appears in the Additional information field. Maximum length: 1000 characters.
Node		Node name, fully qualified domain name (FQDN), IP address, or MAC address that is associated with the event, such as IBM-ASSET. Maximum length: 100 characters.
Application		Predefined attribute extracted from connector payload.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attributes in Link View (continued)

Attribute name	Icon	Description
Process		A process or a cluster of processes running on the host that are identified by the horizontal discovery.
Component	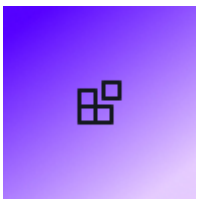	A logical part of a service instance where the alert was detected, with one component shown for each child alert.
Service		Attribute extracted from the log file. The impact of this alert on the service might not be clear or defined.
Host		Attribute extracted from the log file. It might not be identified in the CMBD.
User		Attribute extracted from the log file.
Account		Attribute extracted from the log file.
Path		Path or route extracted from the log file.
Request id		The request ID extracted from the log file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attributes in Link View (continued)

Attribute name	Icon	Description
URL		Web address or URL extracted from the log file.
Detection		An identified anomaly on a metric measured by Health Log Analytics.
Other **Note:** This attribute is user-defined. All the other attributes are predefined in the system.		User-defined attribute. **Note:** All the other attributes are predefined in the system.

Lines linking attributes

Link	Visual representation	Description
Solid line		Solid line linking attributes in Link View, indicating that the attributes share one or more alerts.
Dotted line		Dotted line linking attributes in Link View, indicating that the attributes are correlated by grouping criteria.

Alert field mapping on the Respond page

The fields to be configured on the Respond page, when selecting an action to be triggered by an alert automation.

Select automation actions fields

Field	Description
Incident field name	Name of the field that will appear in the incident.
Map field	Option that indicates how the incident field is to be mapped:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select automation actions fields (continued)

Field	Description
	• **Same as**: The selected incident field will have the same value as the indicated alert field. • **To**: The selected incident field will be set to a specified value.
Alert field value	Value that will be set in the alert field: • If you select **Same as** in the Map field column, the alert field value that you want to appear in the incident field also appears in the Incident field name column. • If you select **To** in the Map field column, enter the alert field value that you want to appear in the specified incident field.

To return to the procedure, see Create Respond automation.

Using Event Management

As an Event Management operator, your role is to find alerts, analyze them, and take action to help resolve the underlying issue.

This guide can assist you on a regular basis.

View alert information

View a list of all alerts for application services s, and then manage individual alerts as necessary.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

Multiple related events may correlate into a single alert. Event Management only creates alerts when one or more events meet the conditions defined in event rules, alert action rules, and alert configuration settings.

When updating an alert, empty field values do not override values inherited from the alert's event.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.

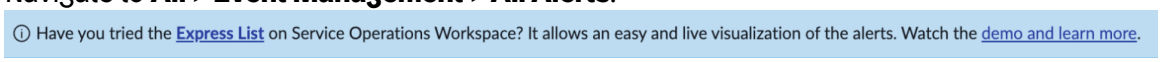

Note: To stop the message from appearing, set the *show_sow_enhanced_experience_message* property in the [sys_properties] table to false.

2. To view or manage an alert, click the alert number.
3. Review the information on the Alert form.
   You can click tabs on the form for further information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To view flapping information, click the **Flapping** tab.
- To view alert history, click the **History** tab.

Alert form

Field	Description
Number	If an alert was created as a result of the event, this field contains the unique ID that Event Management generates to identify the alert.
Source	Event monitoring software that generated the event, such as SolarWinds or SCOM. Maximum length: 100 characters.
Node	Node name, fully qualified domain name (FQDN), IP address, or MAC address that is associated with the event, such as IBM-ASSET. Maximum length: 100 characters.
Type	The metric type to which the event is related, such as Disk or CPU, which is used to identify an event record from which alerts are created. Maximum length: 100 characters.
Resource	Node resource that is relevant to the event. For example, Disk C, CPU-1, the name of a process, or service. Maximum length: 100 characters.
Configuration item	JSON string that represents a configuration item. For example, `{"name":"SAP ORA01","type":"Oracle"}`. The CI identifier that generated the event appears in the **Additional information** field. Maximum length: 1000 characters. Note: Reference pop-ups and click-throughs are hidden by default for read-only fields. For **Configuration item** and other read only fields, you can optionally change the read-only setting. For more information, see Configure pop-ups on read-only fields . Click to open the alert in dependency view.
Task	The corresponding task for the alert, such as an incident, change, or problem.
Metric Name	Unique name that describes which metric data is collected.
Description	A description of the alert.
Message key	Unique event identifier to identify multiple events that relate to the same alert. If this value is empty, it is generated from the **Source**, **Node**, **Type**, **Resource**, and **Metric Name** field values. Maximum length: 1024 characters.
Severity	The severity of the event. The value for this field is copied from the event unless the event closes the alert, in which case the previous severity is retained for reporting. ○ **Critical**: Immediate action is required. The resource is either not functional or critical problems are imminent. ○ **Major**: Major functionality is severely impaired or performance has degraded. ○ **Minor**: Partial, non-critical loss of functionality or performance degradation occurred. ○ **Warning**: Attention is required, even though the resource is still functional.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ **OK**: An alert is created. The resource is still functional. ○ **Clear**: No action is required. An alert is not created from this event. Existing alerts are closed.
State	The state of the alert. ○ **Open**: The alert requires user action. ○ **Reopen**: The previously closed alert requires additional user action. ○ **Flapping**: After the alert has been closed, it receives a high frequency of identical events from the same source that causes many alert reopenings. User action is required. ○ **Closed**: The alert is closed and no further user action is required.
Acknowledged	Select **Acknowledged** to indicate that a user has acknowledged the alert.
Maintenance	A check box that shows whether the resource that is affected by the alert is in maintenance.
Updated	The most recent time that the alert information was updated.
Parent	○ If this field is blank, you can specify an alert number. The current alert becomes a secondary alert to this parent alert. ○ If an alert number is displayed, then the current alert is already a secondary alert to this current alert, which is the parent alert.
Feedback	This field appears in alert forms only if the alert being modified is a parent alert. Your response provides feedback about the accuracy and usefulness of the group. Possible values are: ○ None - No value was selected. ○ No Feedback - There is no feedback for this alert. ○ Yes - The alert group is accurate and useful. ○ No - The alert group is not accurate and not useful. **Note:** Currently, the feedback option is available but does not trigger further actions after being logged.
Knowledge article	The knowledge article associated with the alert, if any.
Overall Event Count	The counter is increased every time an event is bound to the alert. The count value is kept even after the actual event has been purged. Primary alerts (virtual alerts) are updated based on their secondary alerts. The actual number of events that affected the alert is displayed. **Note:** Do not create an Alert Group based on Event Count. Where an Alert Group is based on Event Count, the impact of the group does not calculate properly because the Event Count field is not copied to the em_alert_history table for impact calculation.
Source instance	The name of the machine or software that generated the event. For example, SolarWinds on 10.22.33.44.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
User name and role	The user and role of the person who made the most recent alert updates.
Acknowledged	The **Acknowledged** check box value after the most alert recent update. ○ *True*: The **Acknowledged** is selected. ○ *False*: The **Acknowledged** check box is cleared.
Correlated Alerts section	The secondary alerts that are correlated with this alert, where this alert is the primary alert. For more information, see Rule-based alert grouping.
Primary Alerts section	The primary alert that is correlated with this alert, where this alert is a secondary alert. For more information, see Rule-based alert grouping.

Additional tabs.

Field	Description
Impacted Services tab	
Impacted Services	The impacted services that are related to this alert are listed. For more information, see View the impact tree.
Flapping tab	
Flap count	The number of times the alert has flapped—that is, has fluctuated between a closed and a non-closed state—within the flap interval since the start time in the **Flap start window**.
Flap start window	The initial start time to measure the flapping occurrences.
Flap last update time	The last time flapping occurred. This time is the platform processing time, not the source system time.
Flap last state	The state before the alert entered the flapping state.
History tab	
Initial event generation time	The time when the event that generated the alert first occurred. This time is the platform processing time, not the source system time.
Last event generation time	The last time the event that is linked to the alert occurred. This time is the platform processing time, not the source system time.
Created	The time that the alert was created.
Work notes	The additional notes about the alert.
Activities tab	
Activity	A record of the work that was performed. The text has a date and time stamp.
More Information tab	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Priority Breakdown	Displays the computation of the alert priority score.
Additional information	A JSON string that gives more information about the event. The JSON data is supported for String values only, other value types are not supported. You must convert numbers to String values by enclosing them in double quotes. For example, this value is not supported: {"CPU":100 } while this value is supported: {"CPU":"100"}. Another example of a valid JSON string is: {"evtComponent":"Microsoft-Windows-WindowsUpdateClient","evtMessage":"Installation Failure: Windows failed. Error 0x80070490"}. This information can be used for third-party integration or other post-alert processing. Values in the **Additional information** field of an Event that are not in JSON key/value format are normalized to JSON key/value format when the event is processed. For example, assume that the following plain text is in the **Additional information** field "Connection instance is successful". When the event is processed, all this plain text becomes one JSON string and might not be useful within an alert. In the resultant alert, this string is in the **Additional information** field in JSON key/value format, containing the data: {"additional_content": "Connection instance is successful"}.

What to do next

You can respond to the alert in the following ways:

Alert response options

Option	Description
Submit	Save the modifications that were made to the form and return to the **Alerts** list.
Acknowledge the alert.	Click **Acknowledge**. If the alert is reopened, this button reappears so you can reacknowledge the alert.
Create an incident.	Click **Create incident**.
If Security Incident Response is activated, create a security incident response.	Click **Create Security Incident**.
Designate that the alert is in maintenance.	Select the **Maintenance** check box. For more information, see View all alerts by the maintenance status.
Close the alert.	Click **Close**. For more information, see Close an alert.

Related topics

Security incidents created from events and alerts

Priority group

For better triage and focus, alerts that have a higher priority are brought to the top of the alert list. This placement brings to your attention those alerts that require you to handle them at a higher priority than other alerts.

The priority group indicates which alerts should be attended to first. Priority is calculated for each open alert and then mapped into one of four priority categories.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Thresholds

Thresholds on the calculated priority are used to determine to which category the alert is mapped to. Alerts with a priority above the value of the *evt_mgmt.top_priority_group_threshold* property are categorized as **Urgent**, which is the highest priority. The default value of this threshold is 1M.

Priority group categories

The priority groups into which alerts fall are listed in the following table.

Alert priority groups

Priority group	Icon
Urgent	
High	
Moderate	
Low	

The value of the alert priority group is more meaningful than the alert severity on its own. For example, an alert with a severity of Major, may not be in need of resolution ahead of an alert whose priority is Critical. The combined group priority value takes the weight of each factor into account, giving you a better indication of which alert to attend to first.

The alert priority group is available in alert lists in the Dashboard, Alert Intelligence, and All Alerts page. For example, in the **All Alerts** list, the alert priority group is displayed as depicted in the following image.

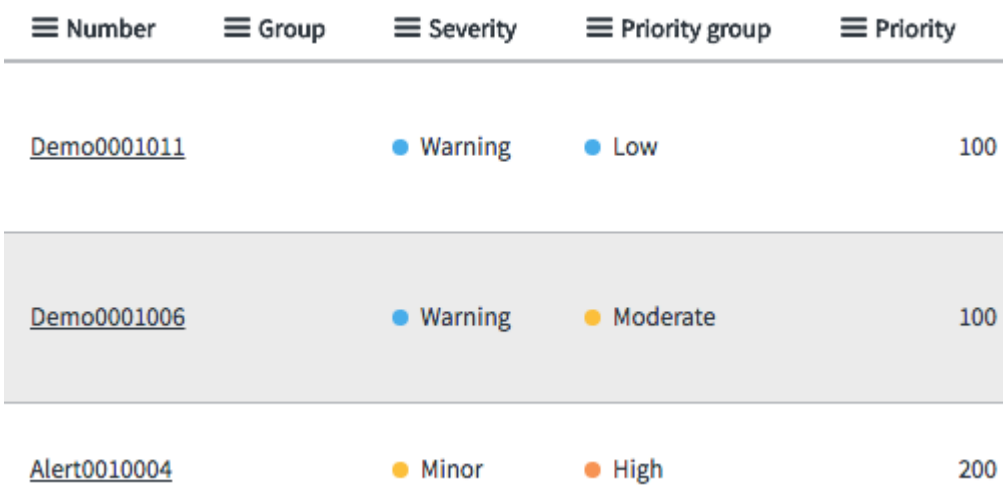

Triggers that cause recalculation

Changes to open alerts trigger the recalculation of the priority and then the alert is again categorized into the relevant priority group. Closed alerts are not considered for alert priority group calculation.

Related topics

Alert priority

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert execution information

Alert execution information provides a reference to the actions that have been performed concerning the alert. Among the information presented is which alert management rules ran on the alert, incidents that were opened, and which remediation workflows ran.

Alert Executions

In an alert form, the Alert Execution tab shows key related records for the alert. For example, it shows which actions were performed and which alert management rules were run. If an incident was opened, the related tasks are shown.

Alert Executions

Fields in the Alert Executions section

Field	Description
Alert	Reference to the alert that ran under this rule.
Action name	Indicates how the alert was resolved.
Link to execution	Link action that was taken to resolve the alert.
Related task	Link to the incident that was created when the alert was resolved.
Log	Link to the log that details the actions taken to resolve the alert
Automatic run	Indicates if the rule was automatic (true) or manual (false).

Action name

This indicates what occurred, for example, if an incident was created.

Log

The log reveals which alert management rule initiated the action. It shows, for example, automatic execution 1 out 3, or manual execution, 2 out of 4.

View alerts in the flapping state

You can view alerts that are specifically in the flapping state.

Before you begin

Before starting this procedure, ask your administrator to configure alert flapping properties.
Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

Flapping occurs when the event source continues to generate events even after its associated alert has been closed. Flapping causes the resource's status to repeatedly fluctuate between the *OK* severity and a severity requiring attention, for example, *Critical*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The frequency of events from an identical source within a given time interval determines whether an alert is in a flapping state or a new issue has occurred. Based on the *evt_mgmt.flap_frequency* and *evt_mgmt.flap_interval* property values:

- If the same issue is recurring, Event Management associates the new event with the existing alert and the state of the alert is set to **Flapping**.
- If the issue occurs after the time interval expires, Event Management creates an alert.

For example, you can respond to an alert by rebooting a problematic server. After no events are generated for several minutes, it is assumed that the issue is fixed and the alert is closed. If the reboot did not actually fix the issue, this server can generate more events later. Then additional alerts are generated for the same issue.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Click the number of an alert that is in the **Flapping** state.
3. On the alert, click the **Flapping** tab.

Flapping tab on Alert form

Field	Description
Flapping tab	
Flap count	The number of times the alert has flapped—that is, has fluctuated between a closed and a non-closed state—within the flap interval since the start time in the **Flap start window**.
Flap start window	The initial start time to measure the flapping occurrences.
Flap last update time	The last time flapping occurred. This time is the platform processing time, not the source system time.
Flap last state	The state before the alert entered the flapping state.

4. If the **Parent** field is empty, address this alert as a new issue.

Related topics

Configure alert flapping

View all alerts by the maintenance status

The **Maintenance** status indicates that the CI is under maintenance. For example, there is a software upgrade, and the issues can result from that activity, therefore all maintenance alerts are discarded.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

The alert maintenance status can change due to any of these reasons:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- A user selected the **Maintenance** check box on the alert, or selected the alert from the list view, and then clicked the **Maintenance** UI action.
- A new alert that had been bound to a CI whose status is in maintenance status on the CMDB table.
- If all previous events associated with the alert are in maintenance, the alert **Maintenance** field can automatically be set to *true* showing maintenance is in progress. However, if an alert **Maintenance** field is *false* but the related CI status is changed to *In Maintenance* after the alert was created, any subsequent events do not put the alert into maintenance. For example, if an alert is generated for a router and later the router is taken down for maintenance, it is still assumed that there was a problem before maintenance started and the issue must not be ignored.

Note:

- If a maintenance flag was set on an alert in Geneva, after upgrading to Helsinki it affects the impact as a regular alert.
- If an alert is bound to a CI by a user action (such as an alert management rule) and the CI is in the Maintenance state, you must manually bind the CI to the alert and mark it with the **In Maintenance** status.

Procedure

1. Navigate to **All** > **Event Management** > **All alerts**.
2. Review the **Maintenance** column for each alert.
   A value of *true* indicates that the CI is under maintenance.

Note: If the **Maintenance** field is not visible on the Alert form, personalize the list to add the field. For more information, see Personal lists.

Collaborate from within an alert

You can collaborate with colleagues and write work notes while working in an alert.

Before you begin
Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task
You can select to work either in a compact view of Connect that overlays the standard user interface or in a full-screen workspace. The conversations are not copied into the alert.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Click the required alert.
3. In the alert screen, click **Follow**.
   The label of the button changes to **Following**, showing the current connection to collaboration status.
4. To view the Connect sidebar in a full window, click the down arrow in the **Follow** button and select **Open Connect Full**.
   You can write text in the **Worknote** field. This text is not copied into the alert.
5. To collaborate with colleagues or with support, click the **Add User** icon.
6. To conclude the Connect collaboration session, click **Following**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Place an alert into maintenance

You can manually place any alert into maintenance to hide it from the Alerts list and Agent Workspace.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator

About this task

Placing the alert in maintenance does not place the CI into maintenance. The **Status** field on the CI record remains the same.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Perform one of the following actions.

Option	Description
To put multiple alerts into maintenance	Select the check boxes next to each alert, and then click the **Maintenance** UI action at the top of the list.
To put one alert into maintenance	Open the alert, and either select the **Maintenance** check box and click **Update**, or click the **Maintenance** UI action at the top of the list.

Related topics

Create maintenance rules

Alert insight information

Alert insight aids faster alert triage, enabling a quicker way to find a solution and expose the probable root cause of the selected alert.

Use alert insight to find a solution and help find the root cause of the alert

Results from data analysis provide further information about the alert, beyond that of the alert form. This aids faster alert triage, enabling a quicker way to find a solution and expose the probable root cause of the alert.

Information shown is from repeated alerts, similar alerts, incidents, problems, and change requests on the same CI, similar CIs, and other CIs that are related by CMDB relations or grouping.

Apply a quick response in an alert

In an alert, use the *Quick Response* feature to apply remediation to the alert or to launch a web application.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

About this task

Quick Response is a feature that enables you to perform a quick action on the selected alert. The action is to apply remediation or to launch an application. If you have upgraded from Istanbul or earlier, this feature is a combination of the **Remediation** and **Launch an application** features.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you select a remediation link when invoking a Quick Response, a workflow is executed according to the criteria specified in the alert management rule. Click a launch application link to launch the web application that was configured in the alert management rule. For further information, see Migrate an alert action rule to an alert management rule.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. In the Alerts list, right-click the required alert and select **Quick Response**.

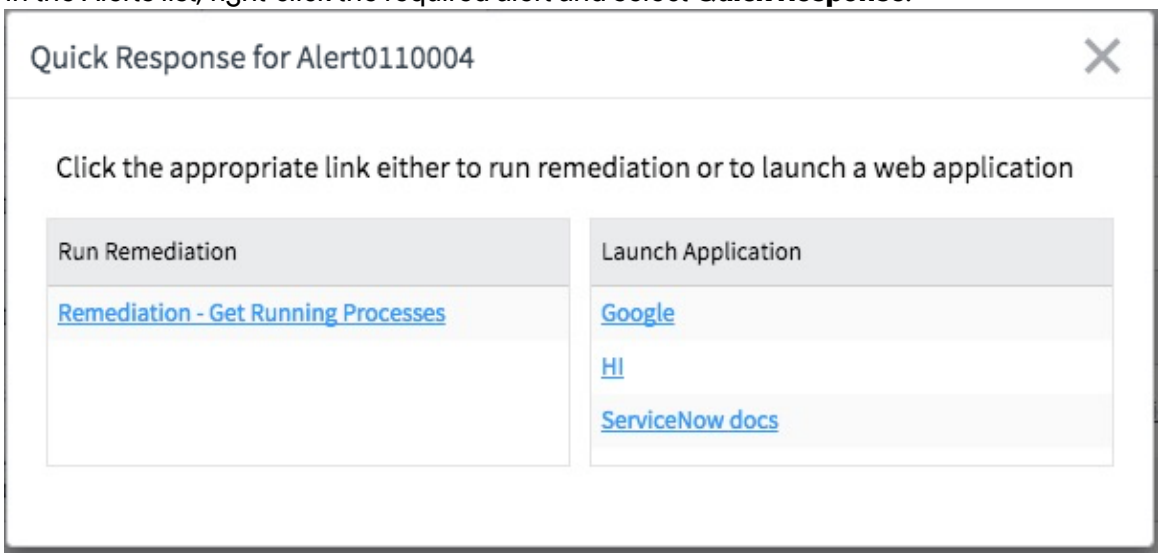

3. In the Quick Response page, click the required link to perform the required action.
   If the link that you require is not available, perform one of these actions.

Option	Description
For Remediation, define a remediation task using an Alert rule based on the selected alert	**a.** Navigate to **Alert Rules**. **b.** Select the required alert management rule to edit. **c.** Click **Remediation**. **d.** Click **Enable Remediation**. **e.** In the **Execution** field, select either Auto matic or Manual. **f.** In the **Orchestration** workflow field, specify the remediation workflow. ▪ For **Orchestration**, in the workflow set tings, select **Remediation Task [em_re mediation_task]** in the **Table** field. ▪ After you finish configuring the orchestra tion workflow, publish it. **g.** Click **Update**. For further information, see Migrate an alert action rule to an alert management rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
For Launch Application, create a web application link using an Alert rule based on the selected alert	**a.** Navigate to **Alert Rules**. **b.** Select the required alert management rule to edit. **c.** Click **Launcher**. **d.** Click **Enable**. **e.** In the **Display Name** field, specify a name for the link. **f.** In the **URL** field, compose the URL using data from the alert in the format: `${source}.com:${port}/${cmdb_ci.name}` **g.** Click **Update**. For further information, see Migrate an alert action rule to an alert management rule.

What to do next

To create a remediation task or web application link, see Migrate an alert action rule to an alert management rule.

Related topics

Launch web application from alert

Monitor service health

On the Operator Workspace, you can view alerts by application services, technical service, and alert group. For services, you can also open a service map to view relationships between CIs in the service.

Before you begin

Ask an administrator to create at least one application services, technical service, or alert group if none of these items exist.

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

The Operator Workspace tiles represent a service, alert group, or CI. The tile color represents the overall alert severity.

Note: Event Management Dashboard is not supported for new instances in the Paris release. However, instances upgraded from a release prior to Orlando that use Event Management Dashboard can continue to do so.

These tile icons are available:

Dashboard tile icons

Icon	Meaning
	Represents the status of a mapped application services.
	Represents the status of an Alert group with a set of corresponding open alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dashboard tile icons (continued)

Icon	Meaning
	Represents the status of a dynamic CI group.

Each tile represents the highest severity of an alert for the service, alert group, or CI. For example, a tile appears in red if an alert is critical. A green tile indicates either information alerts or no currently active alerts. As the alert severities change, the tile colors update accordingly. The following severities are available:

- **Red**: critical severity (highest severity). Immediate action is required. The resource is either not functional or critical problems are imminent.
- **Orange**: major severity. Major functionality is severely impaired or performance has degraded.
- **Yellow**: minor severity. Partial, non-critical loss of functionality or performance degradation occurred.
- **Blue**: warning severity (lowest severity). Attention is required, even though the resource is still functional.
- **Green**: informational. No severity. An alert is created. The resource is still functional.

Procedure

To integrate with Operational Intelligence, activate the Operational Intelligence (com.snc.sa.metric) plugin.
Follow the steps in the Get started with Metric Intelligence topic. Metrics information appears on the **Metrics** tab in the Service Operations Workspace.

Related topics

View metric values in the Insights Explorer

View discovered service history

The discovered service history shows the frequency of discovered services for a particular time period.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

About this task

The discovered services history appears only when there are multiple discovered services that occur over a period. A discovered service is a service that is discovered by Service Mapping. The discovered services history is not available for application services. If no discovered services are generated for a particular service, the discovered services history is hidden.

Procedure

1. Navigate to **All** > **Event Management** > **Service Operations Workspace**.
2. Click a service tile and click **Service Map**.
3. Click the **History map** button.

   The history timeline appears, with changes to the service indicated by a clock icon ( ).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. To view a map showing changes in the service's severity, click the More icon ( ) and select **Advanced map**.

The displayed color corresponds to the discovered services severity. The length of the bar in each color corresponds to how long the discovered services stayed at that severity.

The following table explains the severity colors:

Service severity colors

Severity	Icon color
Critical	Red
Major	Orange
Minor	Yellow
Warning	Blue

5. To change the information that appears on the discovered services history Advanced map, click a history icon.

- **Calendar icon**: Show information for currently active discovered services.
- **Hours**: Show information for discovered services that occurred in the past hour.
- **Days**: Show information for discovered services that occurred in the past 24 hours.
- **Weeks**: Show information for discovered services that occurred in the past week.
- **Months**: Show information for discovered services that occurred in the past month.

Monitor alerts for an application services

To view information for application services only, navigate to the application services list. From this list, you can open service maps to view and manage alerts for the CIs in each service.

Before you begin

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Event Management** > **Services** > **Application Services**.

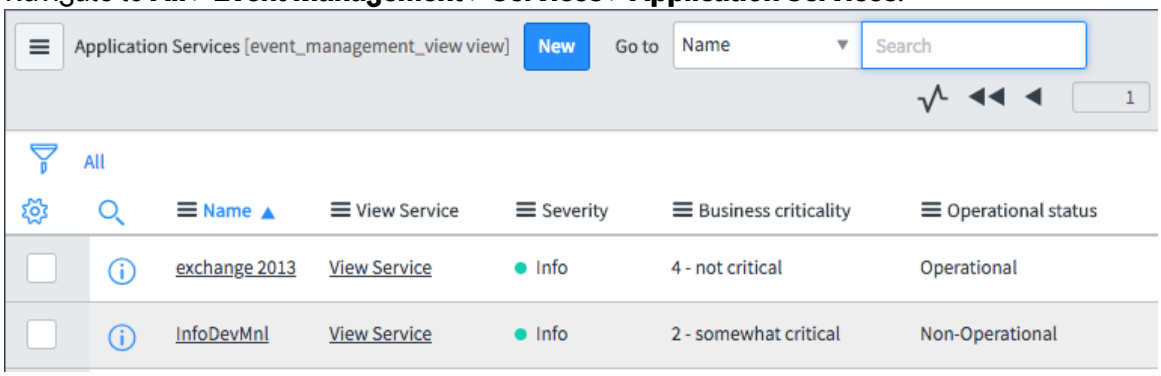

2. In the row of the required application services, click **View Map**.
3. Do one or more of the following.

Option	Action
View alert details for a CI	In the service map: **a.** Click a CI tile. **b.** Below the map, click the **Alerts** tab and re view the listed alerts.
View impact on the CI parent	In the service map: **a.** Click a CI tile. **b.** Below the service map, click the **Impact** tab and review the listed impact rules. **c.** Adjust the impact rules as necessary.
Change the map display, map layout, or map indicators	In the service map header: **a.** Click the menu icon. **b.** Configure the appropriate settings.
Navigate to another application services	In the service map header: **a.** Click the down arrow next to the service name or the folder icon for all services. **b.** Search for and select another application services ( ). The icon color represents the highest impact or severity for active alerts on the application services. ○ **Critical**: Red (highest severity). ○ **Major**: Orange. ○ **Minor**: Yellow. ○ **Warning**: Blue (lowest severity).
View properties for a CI	In the service map:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Action
	a. Click a CI tile. **b.** In the side pane, click the **Properties** tab and review information about the CI. **c.** If you want to view more detailed informa tion, scroll to the end of the pane and click **Detailed Properties**.

View monitored services

View all services that Event Management supports, such as, alert groups, discovered services, application services, and technical services. According to the type of service, you can view service definition details or drill down into the service.

Before you begin

Role required: evt_mgmt_user

Note: Application services that were created before activating the Event Management plugin are not listed in the in the **View Service** column of the Monitored Services list.

Procedure

1. Navigate to **All** > **Event Management** > **Monitored services**.
2. Click **View Service** to open the view that is relevant to that service.
   When you drill down to the contents of the service, either a map, alert list, or CI list appears according to the type of service.

   For example, if you click **View Service** for SNMP1000, you can drill down to the contents of the service and the related topology map appears.

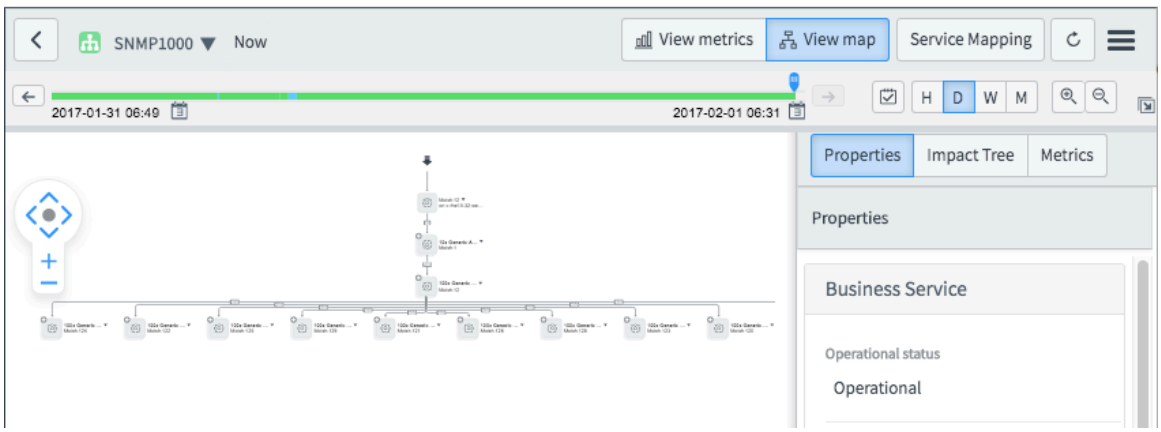

3. To view service definition details of a service, click the name of the service on the **Monitored services** page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, in the **Name** column, click SNMP1000 (SNMP1000).

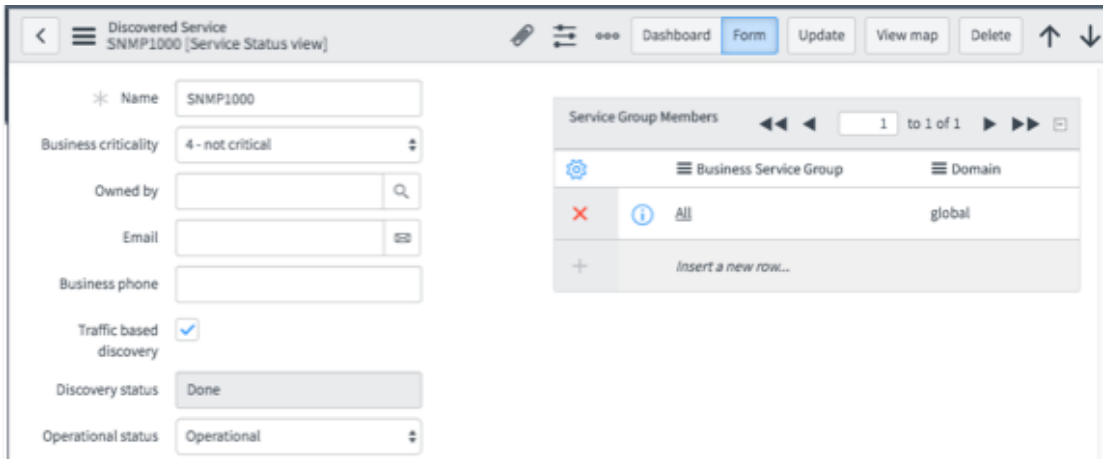

Related topics

Monitor service health

Use the Event Management overview dashboard

The Event Management overview module uses Performance Analytics to present data from your instance for you to better visualize and understand your processes and drive continual improvement.

About this task

From the Xanadu release onward, the Event Management overview module has been deprecated and replaced with the AIOps dashboards. The dashboards deprecated include the following:

- Event Management insights
- Event Management scorecards
- Event Management overview

The AIOps dashboards offer comprehensive visualizations of critical operational data. You can monitor key performance indicators (KPIs), track service health, and gain valuable context on the overall IT environment. For more information, see AIOps Dashboards in Service Operations Workspace for ITOM.

Related topics

Performance Analytics and Reporting

Launch web application from alert

You can launch a web application from an alert that matches the conditions set in an alert action rule.

Before you begin

Create an alert action rule that is configured to prepare a URL when the conditions that you stipulate are met. For example, based on the results when the alert parameters are resolved, this feature can open a Knowledge Base article or construct a URL.

Role required: evt_mgmt_admin, evt_mgmt_operator, or evt_mgmt_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Right-click an alert in All Alerts or the Alert list in the Event Management Dashboard (for instances upgraded from a release before Orlando), and choose **Quick Response**. Click the required link to launch the web application that was configured in the alert action rule. For further information, see Migrate an alert action rule to an alert management rule.

Procedure

1. Navigate to **All** > **Event Management** > **All Alerts**.
2. Right-click the required alert and select **Quick Response**.
   A window appears with a list of available web applications that were configured in the alert action rule and the filter that matches the alert. To see the target URL, point to the cursor over the application name.

   Note: If one or more of the alert action rule parameters cannot be resolved, the related application name appears in black. When the alert parameters are resolved, then the application name becomes a link.

3. Click the name of an application to open the URL in another tab or window in your browser.

Example:

A common use case is launch in context to the source management system.
Other examples include to search in knowledge bases, not only within ServiceNow, but externally as well. Any URL-based action can utilize the alert parameters and the URLs can refer to wikis, messaging services, REST APIs, and so on.

Related topics

Migrate an alert action rule to an alert management rule

Apply a quick response in an alert

Event Management Operator Tutorial

As an Event Management operator, your role is to find alerts, analyze them, and take action to help resolve the underlying issue.

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Take the tutorial

Before you can work as an operator, take a tutorial to become familiar with several Event Management concepts. The tutorial helps you understand what operators typically do on a regular basis.

Start the tutorial by clicking the first lesson:

Lesson 1	An overview of events and alerts
Lesson 2	An overview of application services

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Lesson 3	Event Management operator workspaces
Lesson 4	What operators do

An overview of alerts for Event Management operators

As an Event Management operator, understand how alerts are generated from events, what to look for in alerts, and how alerts are grouped.

This is the first lesson in the Event Management tutorial.

Lesson 1		**An overview of events and alerts**
Lesson 2		An overview of application services
Lesson 3		Event Management operator workspaces
Lesson 4		What operators do

Your organization already has an event monitoring tool in place, such as Microsoft System Center Operations Manager (SCOM), Nagios, SolarWinds, and so on. When an issue occurs on your network, such as a computer going down or a database failure, the event monitoring tools send events to your ServiceNow instance. The Event Management application processes the events according to the settings that your administrator configured, and then generates alerts. An alert is an indicator that the issue requires some type of action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert generation

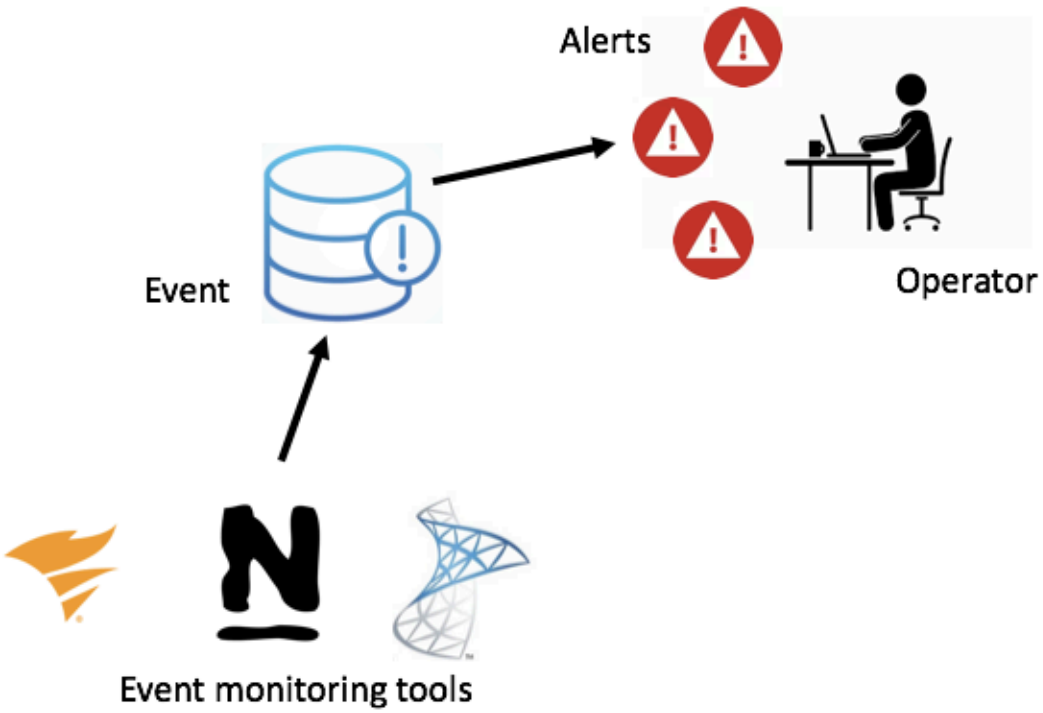

As an Event Management operator, your role is to view alerts. Depending on how Event Management is implemented in your organization, take an action to help resolve the underlying issue or notify someone who can. Later in this tutorial, you will see the phases of a typical alert management process.

Alert priority and severity

The two most common characteristics of an alert are the priority and the severity.

- The priority of an alert is a score that helps you determine how important the impact is to application services. Multiple factors determine the alert priority score. Your Event Management administrator can configure the algorithm that the Event Management application uses to calculate priority.
- The severity of an alert is an indicator of how serious the underlying issue is. The event monitoring tool in your organization usually sends severity values with the event, which then gets carried over in the alert. These are the default severity types that you will see in this tutorial:

Severity	Description
Critical	The resource is either not functional or critical problems are imminent.
Major	Major functionality is severely impaired or performance has degraded.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Severity	Description
Minor	Partial, non-critical loss of functionality or performance degradation occurred.
Warning	Attention is required, even though the resource is still functional.
OK	No severity. An alert is created. The resource is still functional.
Clear	The alert no longer needs action.

Correlated alerts

Some alerts are related to each other. For example, if a router goes down, several separate alerts could be generated, one for each server connected to the router. All of these alerts are related, or correlated. To help you manage correlated alerts, Event Management can automatically group them and establish a two-level hierarchy. One root alert, called the primary alert, appears at the top, and other related alerts, called secondary alerts, appear under the primary alert. When you view alerts, primary alerts stand out by default so you know which alert to focus on without being distracted by the secondary alerts.

In our example, if a router goes down on your network, network communication is also affected for connected servers, assuming they cannot reach any other routers. The router outage becomes the primary alert and the alerts generated on the server are secondary alerts that are correlated under the router alert.

Secondary alert generation

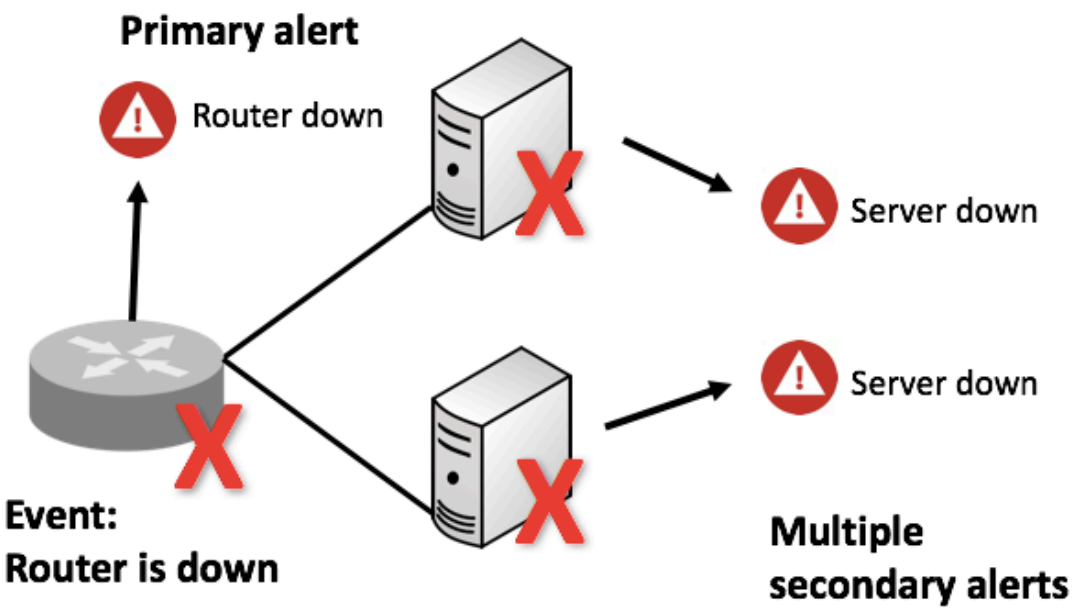

Depending on your organization's Event Management implementation, alerts might be grouped automatically based on correlation rules that your administrator sets up. Your instance can also learn how to improve the way it correlates alerts based on these rules. As an operator, you should still verify the accuracy of the correlation and, if necessary, manually correlate additional alerts with the primary alert. Later in the tutorial, you will learn how to do this.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In this tutorial, you will learn how to manually correlate alerts.

Alert flapping

An alert can flap, meaning that it gets multiple open-close events in rapid succession. Flapping indicates that Event Management does not know whether the underlying events are genuine or not. The events could indicate small issues with the way CIs are configured, or larger issues, like network outages.

Alert flapping

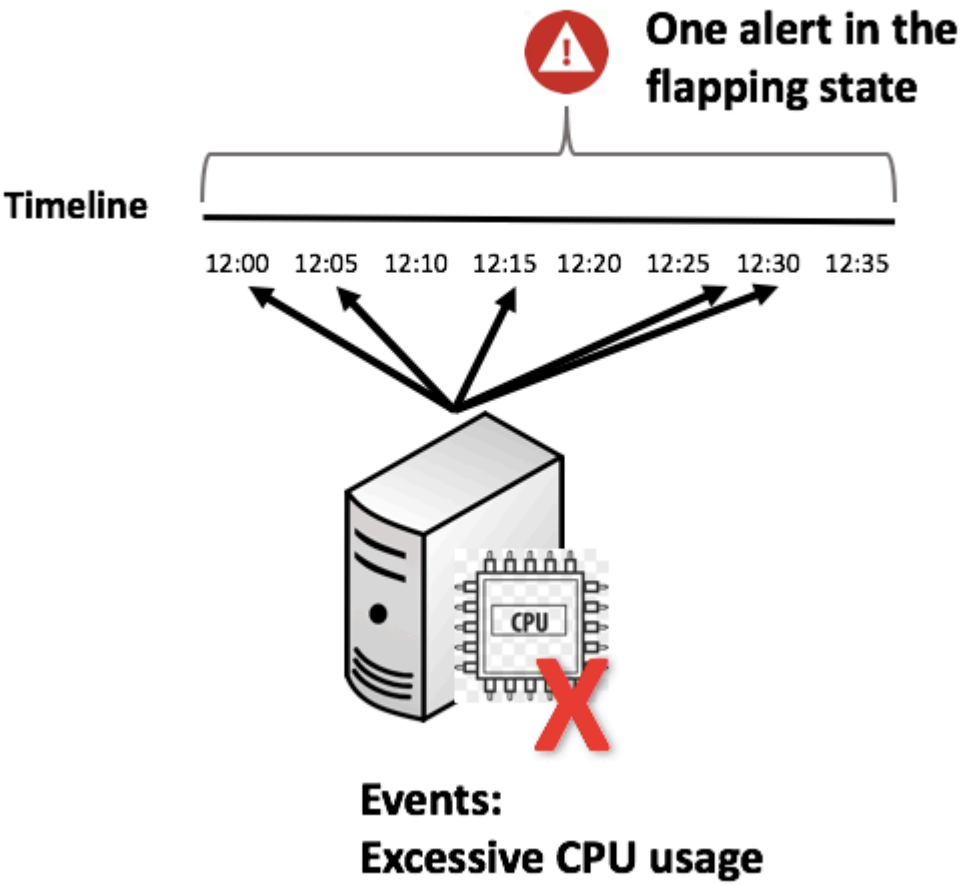

For example, if a server that hosts a web service has too many active processes, it might trigger an event about excessive CPU usage. Since CPU usage can fluctuate rapidly depending on web service requests, several events might be triggered, leading to the alert being put in the flapping state. As an operator, you might need to create an incident to have the server restarted. Someone might have to reconfigure the CPU or make a hardware change on the device.

As another example, consider a loose network cable that causes momentary, repeated network outages. The thresholds that your administrator configures might not be optimal for this kind of alert and Event Management considers it a flapping alert.

Continue the tutorial

Proceed to the next lesson: Application services for Event Management operators.

Application services for Event Management operators

As an Event Management operator, you need to understand what application services are.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This is the second lesson in the Event Management tutorial.

Lesson 1	✓	An overview of events and alerts
Lesson 2	○	**An overview of application services**
Lesson 3	○	Event Management operator workspaces
Lesson 4	○	What operators do

An application service is a collection of components, such as network devices, computers, and applications that offer a service to your organization. The services can be something like an email system or a website that tracks orders or requests in a database. Your administrator should have already specified the application services in your ServiceNow instance.

Each component that makes up an application service is referred to as a configuration item or CI. This service map provides a visual representation of an application service:

Service map

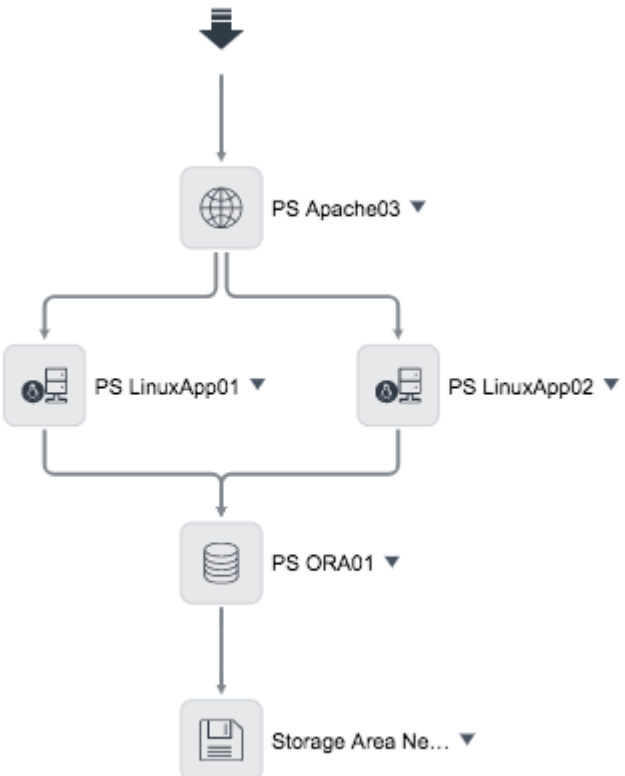

In this example, you can see a Web Portal application service with these CIs:

- **PS Apache03**: An Apache web server that hosts a company website.
- **PS LinuxApp01** and **PS LinuxApp02**: Two Linux servers that share the workload from the web server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **PS ORA01**: A database server that both Linux servers need to access.
- **Storage Area Network 001**: A mass storage device on which the other CIs depend.

You can see service maps like this on the Service Operations Workspace dashboard. Later in this tutorial, you will learn about the dashboard and what an application service looks like when an alert is associated with a CI.

Types of application services

Event Management classifies application services into the following types:

Technical services

A technical service is a dynamic grouping of CIs based on some common criteria. For example, a technical service could be comprised of all web servers or all Oracle databases for a specific location, like North America.

Application services

An application service can consist of discovered services, manual services, or both. A discovered service is an application service that the Service Mapping application finds (if your organization uses Service Mapping). A manual service is an application service that your administrator configures by selecting and adding each CI and specifying the relationships between CIs.

Alert groups

Alerts that are grouped together, either manually or automatically.

Alert impact

Application services are critical to the operations of your organization. If an issue occurs on one CI, the entire application service can be affected. Part of your role as an operator is to analyze alerts on CIs and see how they impact the application service as a whole. You can then take an action to help remediate or solve the underlying issue. Your administrator can configure impact rules that go into calculations for the severity of an alert.

Later, when you learn how to use the Service Operations Workspace dashboard, you will learn how to view an impact tree for an application service. This helps you understand the relationship between the severity of an alert and the overall application service.

In this example, you can see how a **Major** alert on an Oracle database (**PS ORA01**) also causes a major alert on its parent CIs and on the Web portal application service itself.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Impact tree

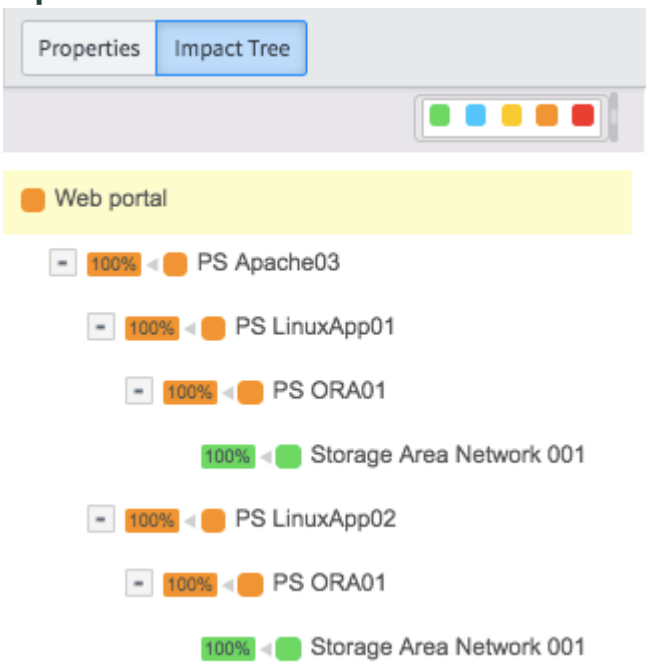

Continue the tutorial

Proceed to the next lesson: Event Management operator environment.

Event Management operator environment

As an Event Management operator, your primary work environment is the Service Operations Workspace dashboard.

In this, the third lesson in the Event Management tutorial, you get an overview of your workspaces so you know how to find the information you need.

Tutorial lessons		
Lesson 1	✓	An overview of events and alerts
Lesson 2	✓	An overview of application services
Lesson 3	○	Event Management operator environment
Lesson 4	○	What operators do

The Service Operations Workspace dashboard

Your main area of work is the Service Operations Workspace dashboard, which provides a view that focuses on how alerts relate to application services. From here, you can drill into each application services to see the affected CIs and get an understanding of the overall impact of whatever caused the alert. To open the dashboard, navigate to **Event Management** > **Service Operations Workspace**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Operations Workspace

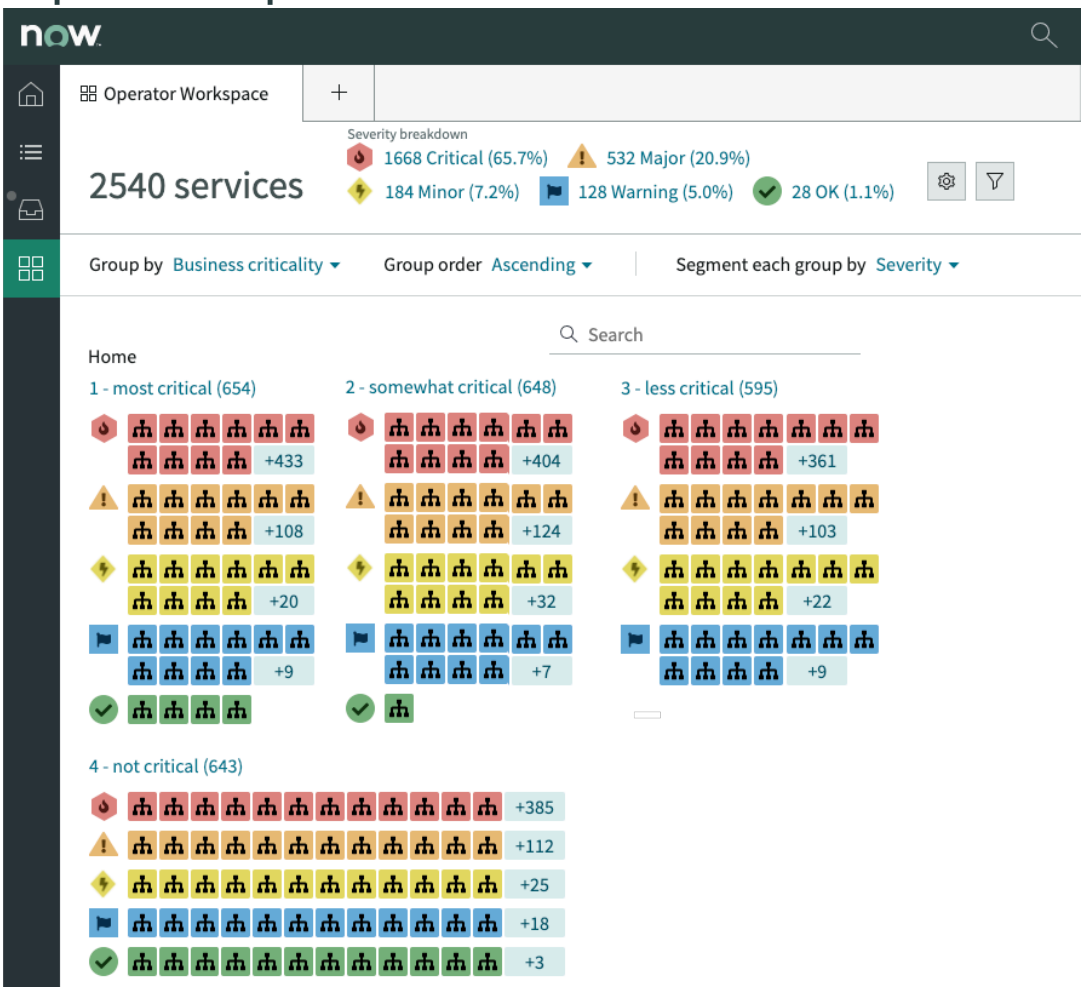

The main sections of the dashboard are:

Banner

The banner contains controls that show or hide application services based on the criteria you select:

Click a severity level to show or hide application services based on the alerts associated with them.

Use the **Group** and **Segment** controls to organize the view.

Note: Your administrator assigns an application the criticality and cost value to your application services.

Group and Segment filters

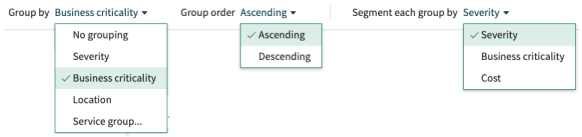

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application services

The Service Operations Workspace dashboard displays tiles that represent application services. The color of a tile represents the severity of the alerts that are associated with the application service.

Click a tile to show a summary of alerts associated with the application service. Click to view details or the service map.

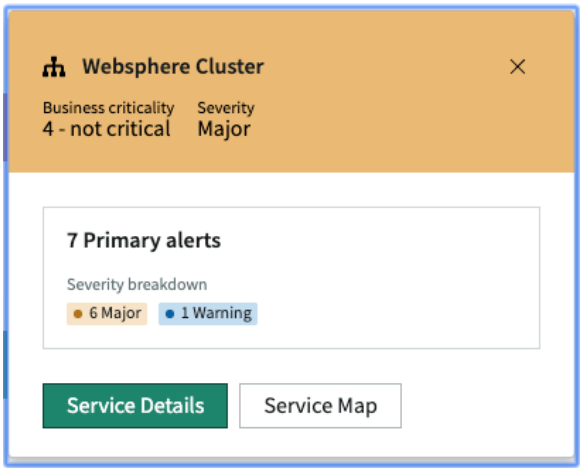

Alerts

Click the List icon ( ). On the Lists tab, click the kind of alert to view.

You can filter or sort the list to find an alert. Sort by any alert details, such as the **Priority**, which considers multiple factors for how serious the alert is. You can also sort by the **Severity**, which is a value provided by the event monitoring tool.

Open any alert by clicking the number.

You will learn about what each of the columns means for an alert later on when you analyze an alert.

Application service map views on the Service Operations Workspace dashboard

Double-click the name of an application service tile to open one of the application service map views. The view that you see depends on the type of application service.

For a manual or standard application service, this view appears:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

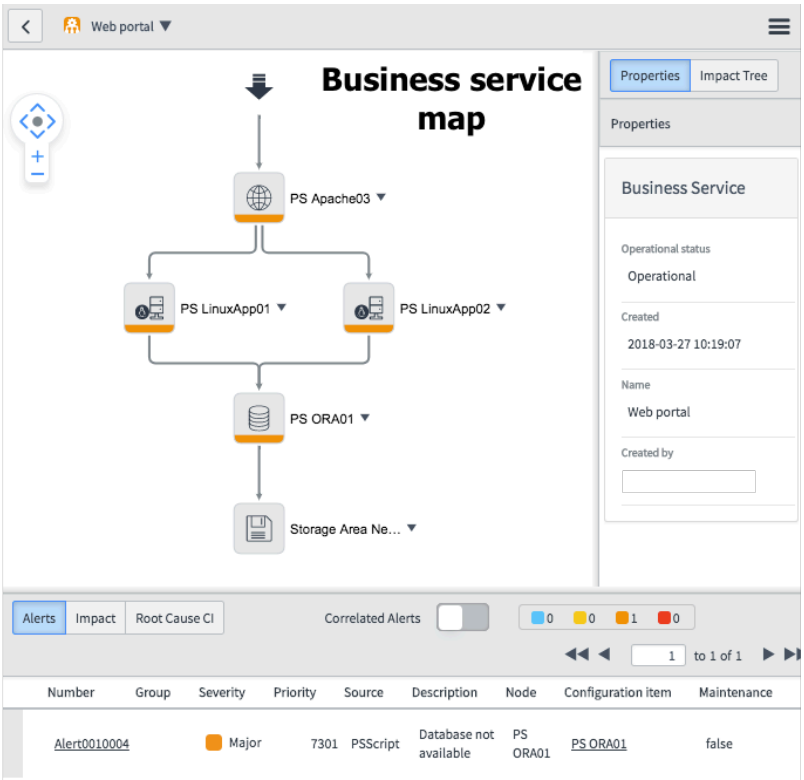

The application service map that you were introduced to in a previous lesson appears in the main panel. This is what you can do from this view:

Click any of the CIs to see the alerts only for that CI and to display the details about that CI in the **Properties** pane.

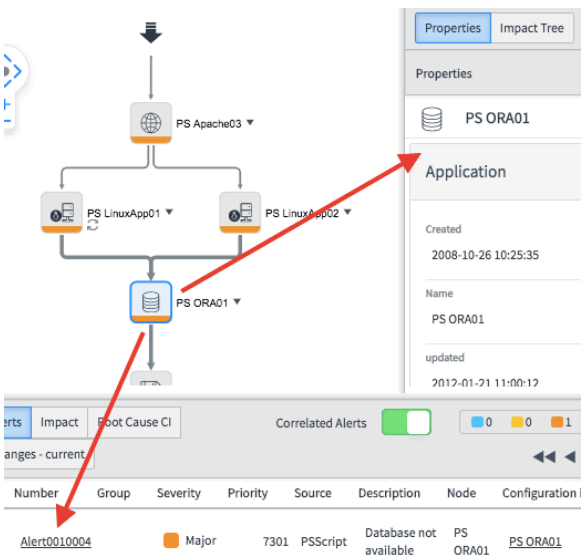

Click **Impact Tree** to see the state of all the CIs and how they affect each other when receiving an alert. You will learn more about the impact tree later on in the tutorial.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

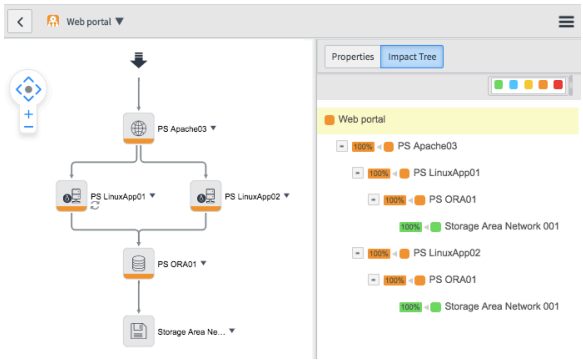

Open any alert from the **Alerts** list at the bottom by clicking the number.

For a technical application service, this view appears:

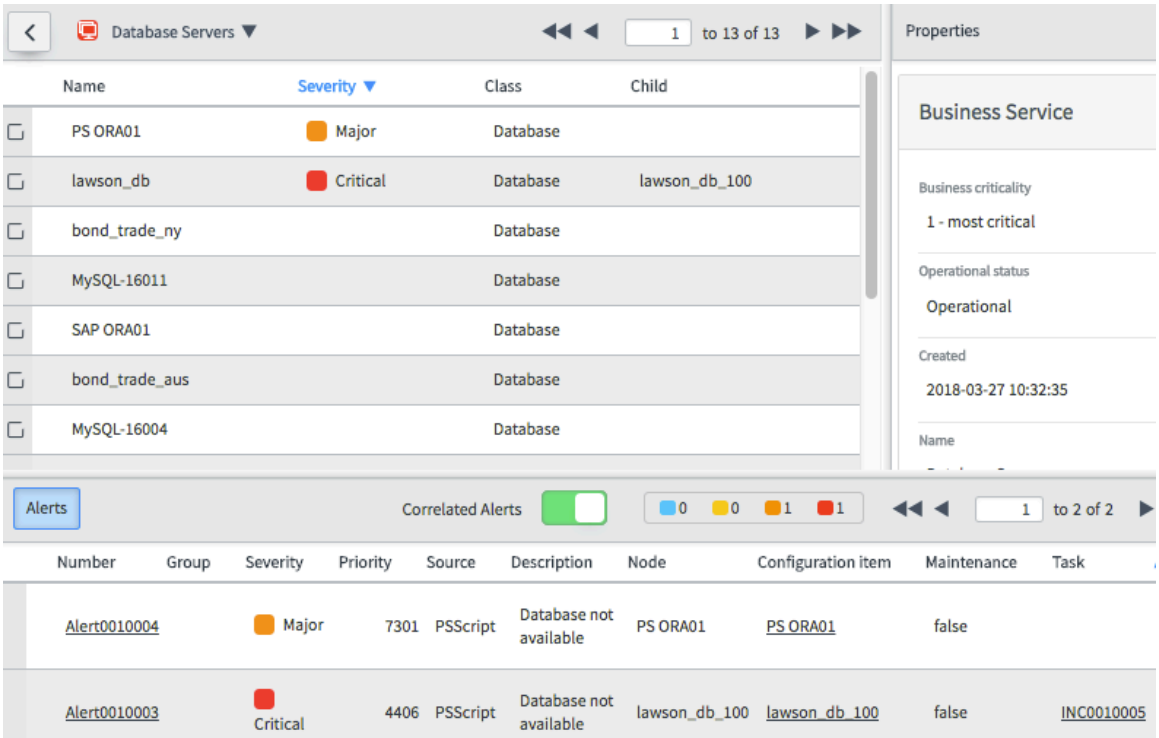

This is what you can do from this view:

Click any CI to view details about it.

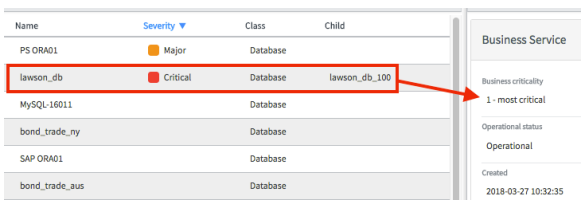

Open any alert from the **Alerts** list at the bottom by clicking the number.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Continue the tutorial

Proceed to the next lesson: What Event Management operators do.

What Event Management operators do

As an Event Management operator, your typical workflow involves three phases: analyzing an alert and its effect on application services, taking some type of action, and making sure the alert is finally closed.

This is the last lesson in the Event Management tutorial. In this lesson, you will run through the three phases of a typical operator day-to-day workflow.

Lesson 1	✓	An overview of events and alerts
Lesson 2	✓	An overview of application services
Lesson 3	✓	Event Management Service Operations Workspace
Lesson 4	○	**What operators do**

Your organization will have specific policies and procedures that may differ from the phases outlined in this topic. In addition, your administrator might have customized the Event Management application so that some of these phases are automated and do not actually require any action on your part.

For the purposes of this tutorial, we will walk you through each phase and allow you to perform the tasks manually.

Phase 1	Analyze and acknowledge an alert	In this phase, you will find an alert to work on, analyze the details, and acknowledge it so that other operators know it is a legitimate alert.
Phase 2	Triage an alert	In this phase, you take an action to help remediate the issue that caused the alert. The most common action to take is to create an incident and assign it to someone who can solve the underlying issue.
Phase 3	Close an alert	In this phase, you will verify that the alert is resolved, and then close the alert.

Start here

Start out by learning how to Analyze and acknowledge an alert.

Operator phase 1: Analyze and acknowledge an alert

As an Event Management operator, the first thing you should do is access alerts and then find the ones you want to focus on. Open the Alert form to analyze the details. Acknowledge the alert to let other operators know that the issue should be addressed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Phase 1	○	**Analyze and acknowledge an alert**
Phase 2	○	Triage alerts
Phase 3	○	Close an alert

You must understand both of your operator environments: the Service Operations Workspace dashboard and Service Operations Workspace. Review Event Management operator environment if you have not already done so. This tutorial uses the Service Operations Workspace dashboard.

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Event Management** > **Workspaces** > **Service Operations Workspace**.

   In the list, click **Alerts** > **All Alerts**. By default, all alerts appear in the list sorted by **Group**. Click any heading to sort by that value. To view current filter settings, click the filter icon

   . To modify filter settings, click **Advanced view**.

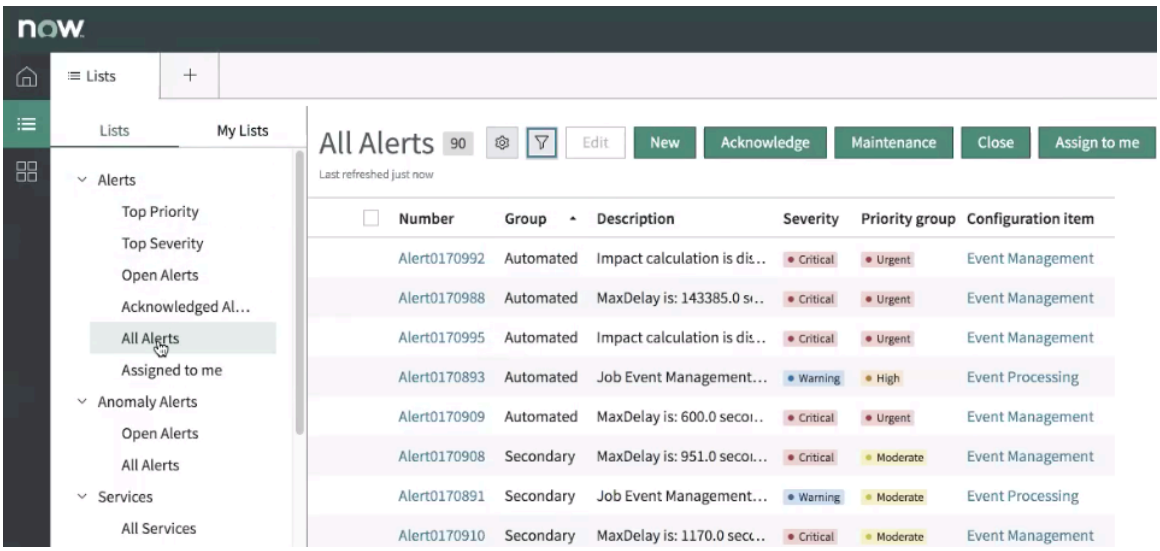

2. Find the alerts to focus on by clicking **Application Services** > **Services**, or by using any of the filtering and sorting options.
3. Click an alert name to open the **Details** tab for the alert.

   Example
   The important fields on the form and related items are described after the screenshot.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

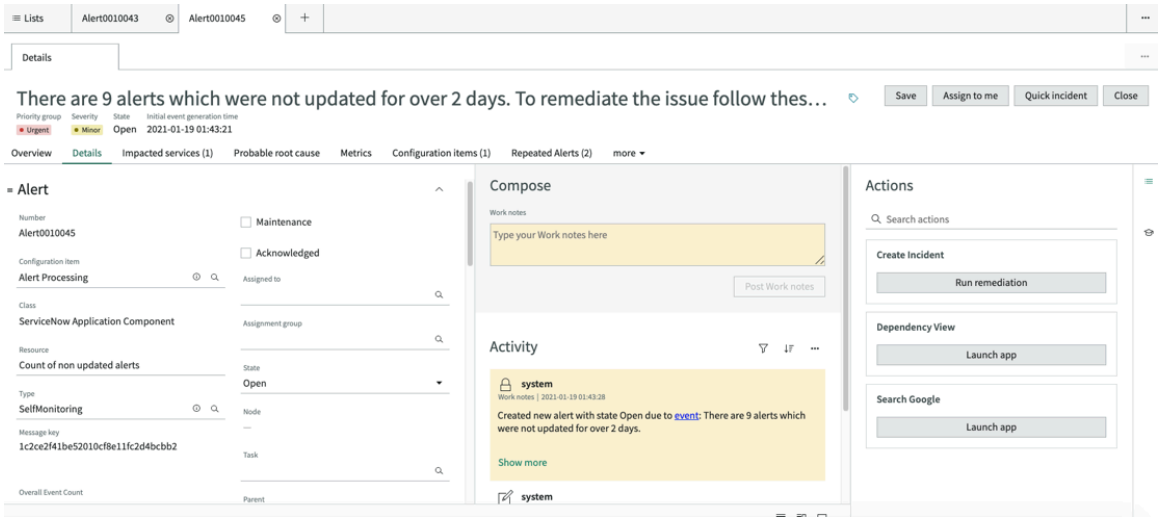

Fields on the Alert Details tab for an alert that is associated with a single CI

Field	Description
Number	The automatically generated alert number that you saw in the list of alerts on the Service Operations Workspace dashboard.
Node	The name of the CI with which the alert is associated. The event that generated the alert occurred on this CI. You might see a computer name, or an IP address or MAC address of a device.
Resource	The object on the node that caused the event. You might see values like the name of a hard disk, a CPU, or the name of a process or a service. You can add information to this field if necessary.
Configuration item	The CI in the CMDB to which this alert applies. If the **Node** field is filled in but this field is empty, and if you know which CI the node represents, you can enter the CI in this field to manually associate the alert with it. **Note:** Your administrator can configure rules that automatically bind CIs to an alert, usually based on what is in the **Node** field. However, it is possible that the system is unable to find the correct CI to bind the alert to. In this case, you might need to manually select a CI. For the purposes of this tutorial, we assume that the correct CI is bound to the alert. When an alert contains a CI, the **Overview** tab opens when selecting the alert.
Task	Any existing task that is associated with the alert, such as an incident, change request, or problem, is shown in this field. When you first look at an alert, this field is blank. You can associate a task when you triage the alert. You associate an incident to an alert and assign the incident to someone at a later stage in this tutorial.
Description	A description of the alert that came from the alert monitoring software in your organization. You can add information if necessary.
Severity	Severity value for the alert:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ Critical: Immediate action is required. Either the resource is not functional or critical problems are imminent. ○ Major: Major functionality is severely impaired or performance has degraded. ○ Minor: There is a partial, non-critical loss of functionality, or performance has degraded. ○ Warning: Attention is required even though the resource is still functional. ○ Info: An informational message. An alert is created, but the resource is still functional. ○ Clear or Resolved: No action is required. An alert is not created from this event. Existing alerts are closed.
State	The processing state that the alert is in. A newly generated alert is in the **Open** state. Other states are: ○ **Reopen** indicates that a previously closed alert is open again, and it requires your attention. ○ **Flapping** indicates that the alert is receiving a high frequency of identical events from the same source, which can cause many alerts to re-open even after it is closed. ○ **Closed** indicates that the alert is closed and does not require any further action. You will close an alert as part of the last stage in this tutorial.
Maintenance	Indicates that the alert does not need to be worked on and is hidden from the list of alerts in the Service Operations Workspace dashboard. In this example, we do not want to put the alert in maintenance.
Acknowledged	Indicates that an operator verified that the alert is genuine and is triaging it. You acknowledge the alert at a later stage in this topic. Your Event Management administrator can also have alerts automatically acknowledged immediately after they are generated. If this is the case, the **Acknowledge** check box on the Alert form might already be selected.
Parent	Displays a value only for secondary alerts. The value indicates the alert number of the parent alert. The field is empty for primary alerts (alerts with no parent alert).
Knowledge article	Any existing knowledge base (KB) article that is associated with the alert. When you first open an alert, this field is blank. You can associate a KB article when you triage the alert.

4. Click the **Impacted Services** tab and then click the service name.
   Before you acknowledge the alert, you might want to learn more about the services that are affected due to the service that is the source of the alert.

   The service opens and displays the **Details** tab. Click **Service Map** to view the map of related application services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

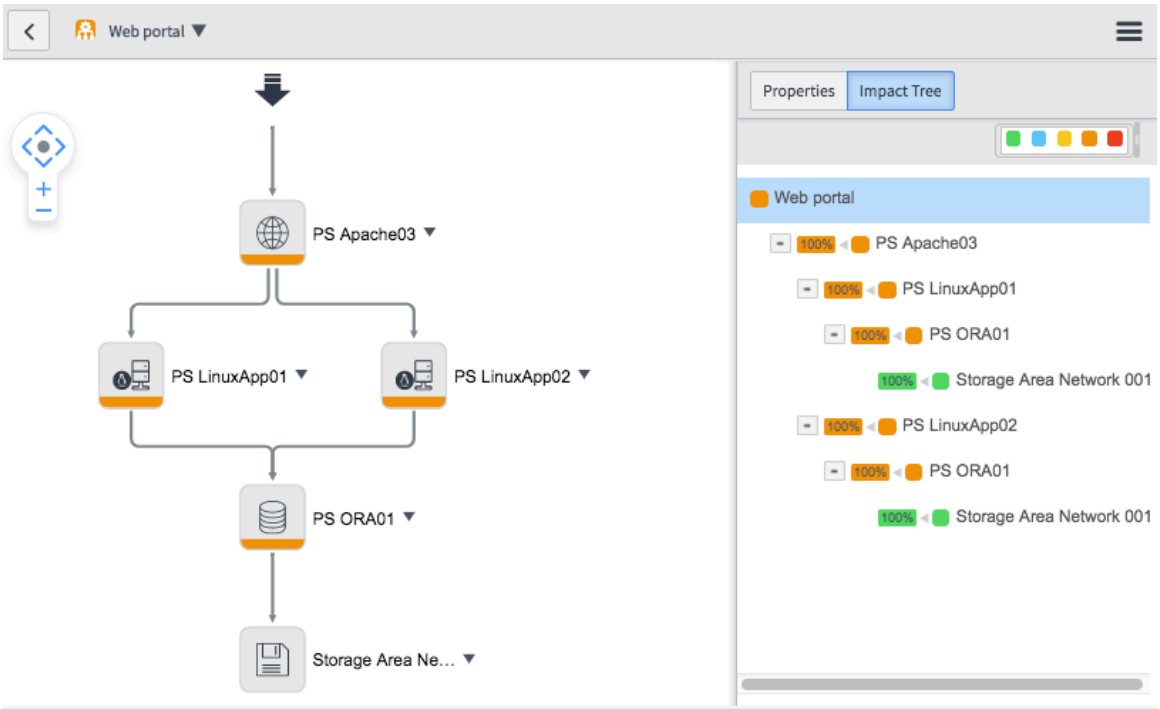

Notice that the major severity indicator in orange is displayed on several CIs. It starts with the Oracle database (**PS ORA01**) and propagates up to all CIs that depend on the database. From this, you can surmise that the alert with the highest level of severity on the Oracle database indicates the root cause of the issue.

This does not necessarily mean that the system generated separate alerts for each CI, but it is possible. When you triage an alert, you can determine whether alerts on other CIs should be correlated under a primary alert.

5. Click the **Impact Tree** tab and review the impacts on each CI.

   The percentages next to each CI in the impact tree show you the effect on each CI. In this example, each CI is impacted 100% by the alert. This means the underlying issue is serious enough to impact the functionality of the application service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

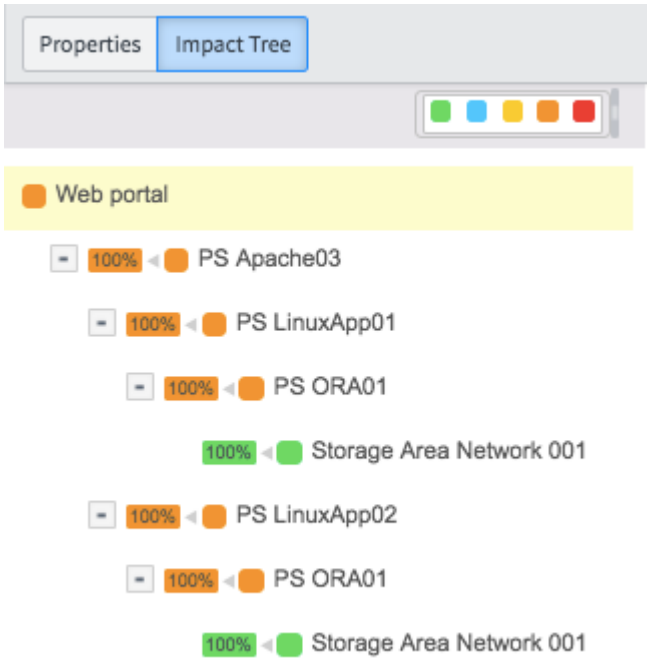

6. Acknowledge the alert.

   Acknowledging an alert does not assign it to you, nor does it create a task like an incident or change request. It simply lets other operators know that you are aware of the issue. After you acknowledge an alert, you will take further action during the triage stage.

Option	Description
Acknowledge a single alert	With this process, you can acknowledge only the alert that you are viewing. **a.** Open the Alert Details form as described earlier. **b.** Select the **Acknowledged** check box.
Acknowledge one or more alerts	With this process, you can acknowledge one alert or multiple related alerts. **a.** Open the list of alerts as described earlier. **b.** Select each alert to acknowledge. **c.** Click **Acknowledge**.

7. Navigate back and look for the alert that you acknowledged.
   You can see that the **Acknowledged** value is **true**.

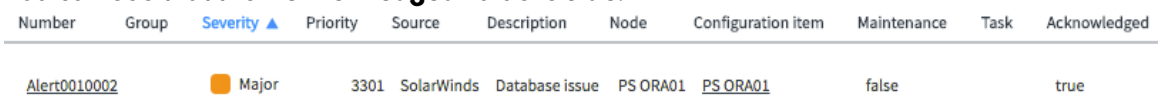

Number	Group	Severity ▲	Priority	Source	Description	Node	Configuration item	Maintenance	Task	Acknowledged
Alert0010002		Major	3301	SolarWinds	Database issue	PS ORA01	PS ORA01	false		true

What to do next

Go to Operator phase 2: Triage an alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Operator phase 2: Triage an alert

After you analyze and acknowledge an alert, you must triage it. The triage phase involves verifying alert correlation and taking an action to help resolve the issue that caused the alert. This topic covers the most common triage task: creating an incident from an alert.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Phase 1		Analyze and acknowledge an alert
Phase 2		**Triage alerts**
Phase 3		Close an alert

Role required: evt_mgmt_operator

Procedure

1. Before you start to work on the alert, check whether other alerts should be correlated with the one you just reviewed:

   a. Navigate to **Event Management** > **Service Operations Workspace**.

   b. Click the same application service tile and review the alerts in the **Alerts** list.

   c. If multiple alerts are present, open one to determine if it is related to the other alert, or if the underlying issue is the same in both alerts.
   For example, an alert might be associated with a web server while another alert exists for a database issue. Correlate the web server alert as a secondary alert with the database alert as the primary alert.

   d. Open the secondary alert, click the lookup icon ( ) in the **Parent** field, and then select the primary alert.
   This creates a primary (parent) / secondary (child) relationship between the two alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

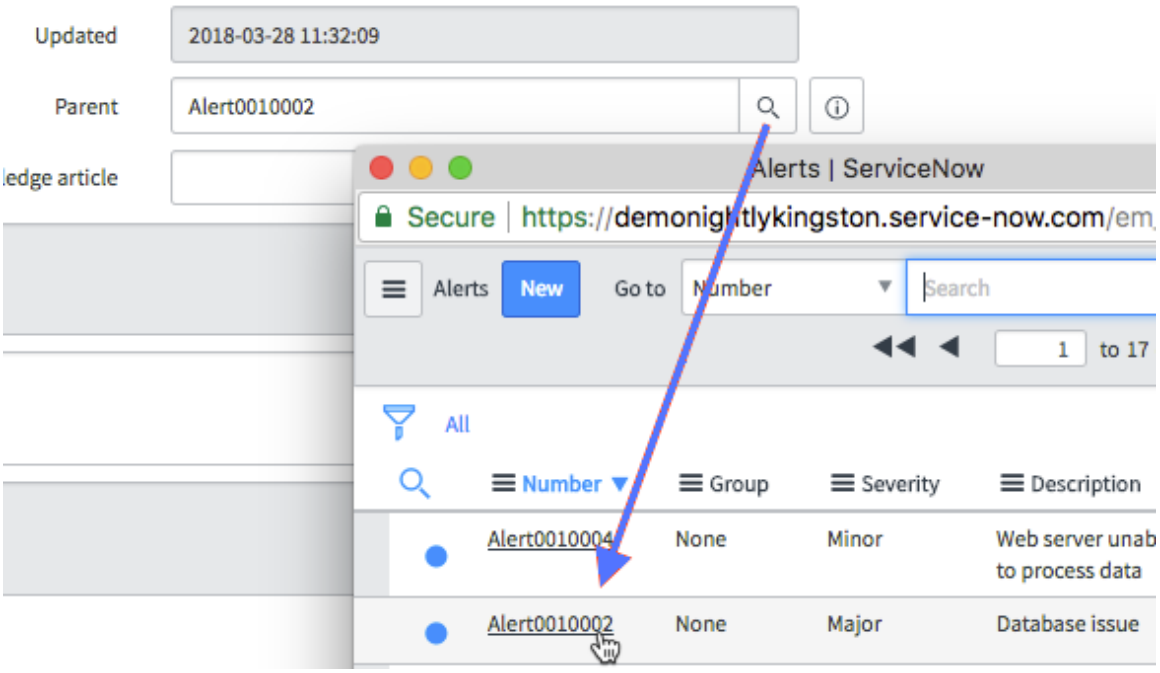

e. Repeat this process to correlate as many secondary alerts as necessary.

f. Navigate back to the Service Operations Workspace dashboard and verify that the primary alert displays an icon in the **Group** column.

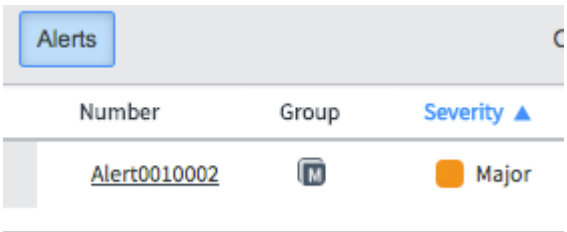

Note: Your administrator can set up rules that let the system automatically correlate alerts so that you do not need to do so manually. In that case, a **Feedback** field appears on the Alert form. Select **Yes** if the system correlated the alert correlated correctly, or **No** if not. Currently, the feedback option is available but does not trigger further actions after being logged.

2. Open the primary alert.
3. On the Alert form, click **Actions** > **Create Incident**.
   If your organization uses Security Incident Response, the button is **Create Security incident**.

   The Flow Designer opens.

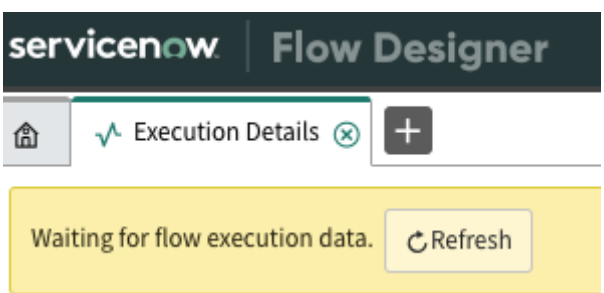

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Click **Refresh**.

   The **Execution Details** page opens.

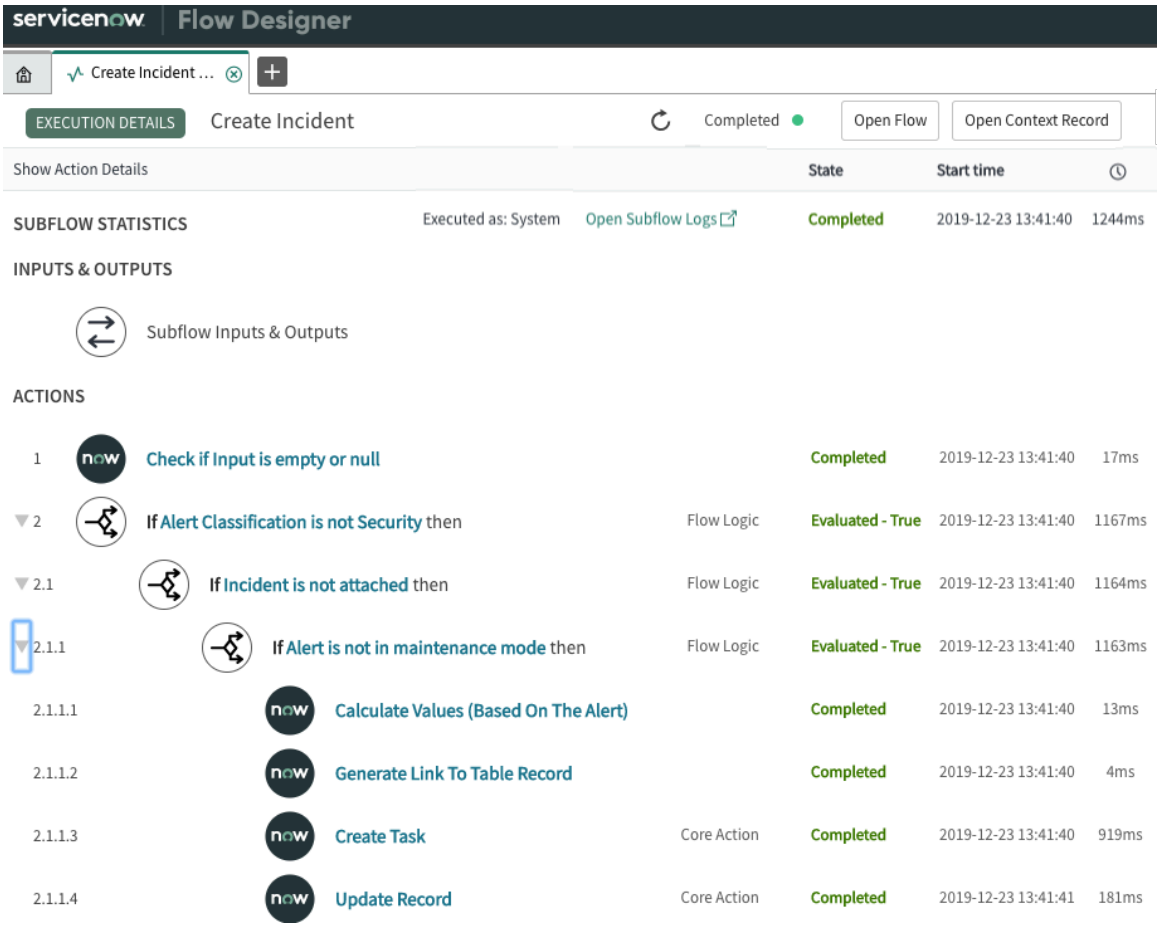

5. Click **Open Context Record** and then fill in the Incident form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

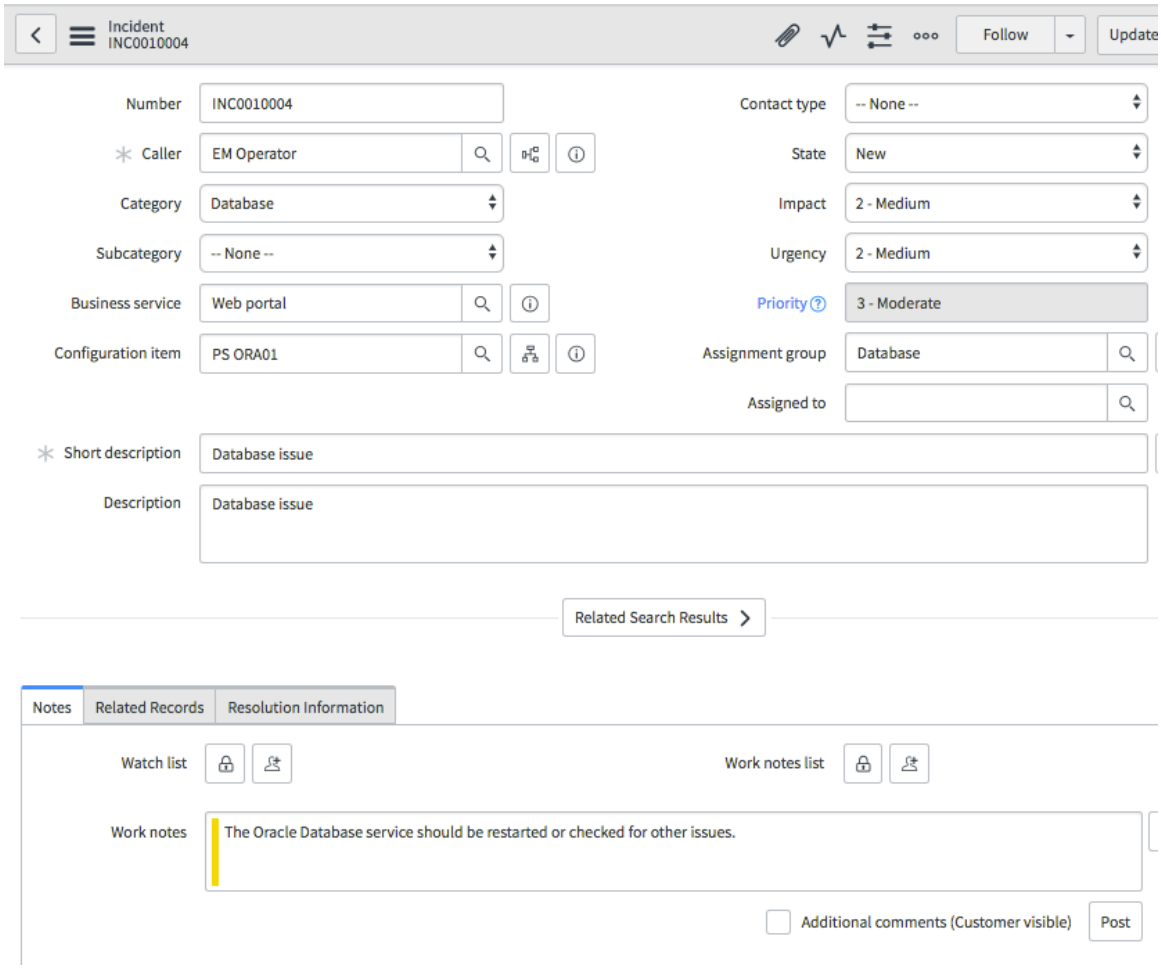

Field	Description
Caller	Click the lookup icon ( ) and then select your name. The caller is the person who discovered the issue that led to the incident.
Category and Subcategory	Select a category that best describes this alert. In this example, **Database** is the best choice.
Application service	Select the application service to which the CI belongs. In this example, the application service is **Web portal**.
Configuration item	If it is not already populated, select the CI. In this example, the Oracle database (**PS ORA01**) is automatically populated into the field.
Impact and Urgency	Select the impact and urgency levels that you think appropriate.
Assignment group Assigned to	Click the lookup icons ( ) for both or either of these fields, and then select the group or the individual that can handle the issue.

6. Click **Submit** to create the incident.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Alert you are working on reappears. On the Alert form, the incident is populated in the **Task** field. You can also see the incident number in the Task column on the Service Operations Workspace dashboard.

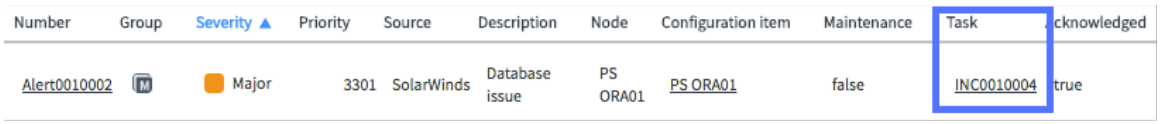

Number	Group	Severity ▲	Priority	Source	Description	Node	Configuration item	Maintenance	Task	cknowledged
Alert0010002		Major	3301	SolarWinds	Database issue	PS ORA01	PS ORA01	false	INC0010004	true

What to do next

There are other tasks you perform as part of the triage stage:

- Run a remediation workflow on an alert if your Event Management administrator already set up a workflow in your ServiceNow instance and your policies allow you to trigger it from the alert.
- Launch a web application from an alert to open a website or an event monitoring tool that provides more information about the alert.
- Associate a knowledge base article with an alert if there is existing information about the alert that might help resolve the underlying issue.
- Put an alert into maintenance to temporarily hide it from the Service Operations Workspace dashboard if the alert does not require action at this time.

If you do not need to perform any other triage actions, proceed to Phase 3: Close an alert.

Run a remediation workflow on an alert

As an Event Management operator, you can also run a workflow on your ServiceNow instance that helps remediate the alert. For example, you might run a workflow that automatically restarts a server on your network, which might resolve an alert about CPU usage.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Phase 1	✓	Analyze and acknowledge an alert
Phase 2	○	**Triage alerts**
Phase 3	○	Close an alert

Note: You can run a remediation workflow if your administrator already set up workflows for you to choose from. You should be familiar with your organization's policies regarding triaging of alerts.

Role required: evt_mgmt_operator

Procedure

1. From the Service Operations Workspace dashboard, open the alert that you acknowledged in Phase 1: Analyze and acknowledge an alert.
2. On the Alert form, click **Quick Response**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

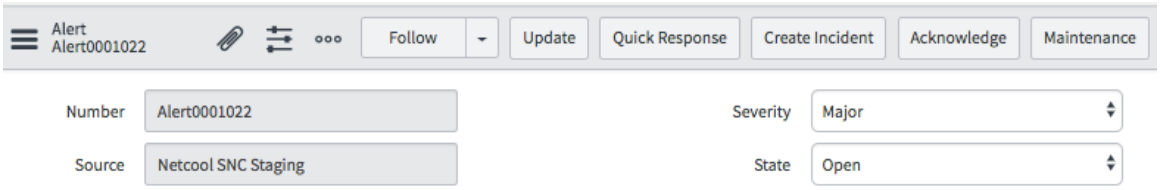

3. In the Quick Response window, click the name of the remediation under **Run Remediation**.

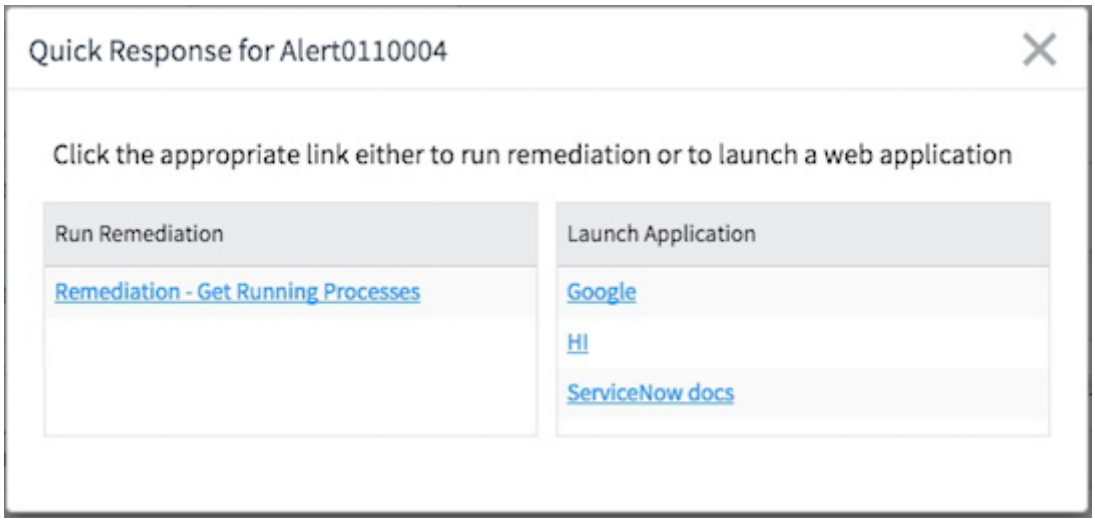

What to do next

There are also other tasks you can take as part of the triage stage:

- Launch a web application from an alert to open a website or an event monitoring tool that provides more information about the alert.
- Put an alert into maintenance to temporarily hide it from the Service Operations Workspace dashboard if the alert does not require action at this time.
- Associate a knowledge base article with an alert if there is existing information about the alert that might help resolve the underlying issue.

If you do not need to perform any other triage actions, proceed to Phase 3: Close an alert.

Launch a web application from an alert

As an Event Management operator, you can also launch a web application from an alert. The web application might be a console for your event monitoring tool or any external website with additional alert information.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Phase 1	✓	Analyze and acknowledge an alert
Phase 2	○	**Triage alerts**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Phase 3	Close an alert

Role required: evt_mgmt_operator

You can launch a web application if your administrator has set up a link to the application. You should be familiar with your organization's policies on triaging alerts.

About this task

Procedure

1. From the Service Operations Workspace dashboard, open the alert that you acknowledged in Phase 1: Analyze and acknowledge an alert.
2. On the Alert form, click **Quick Response**.

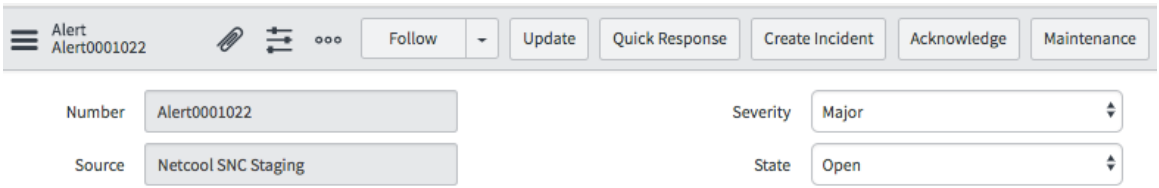

3. In the Quick Response window, click the name of the application under **Launch Application**. Your administrator has specified which applications should appear in the list.

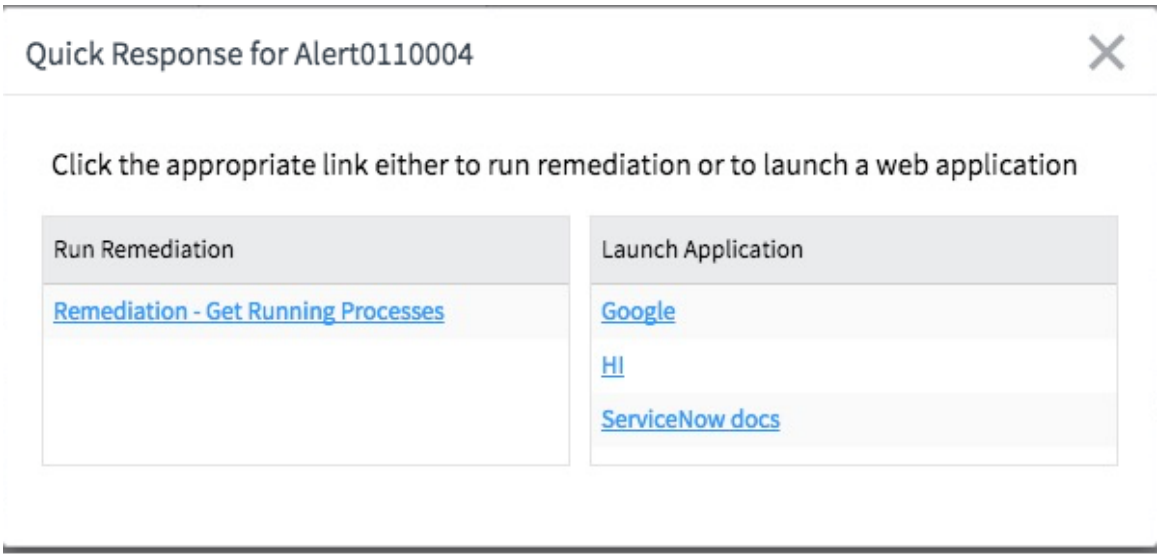

What to do next

There are also other tasks you can take as part of the triage stage:

- Run a remediation workflow on an alert if your Event Management administrator already set up a workflow in your ServiceNow instance and your policies allow you to trigger it from the alert.
- Put an alert into maintenance to temporarily hide it from the Service Operations Workspace dashboard if the alert does not require action at this time.
- Associate a knowledge base article with an alert if there is existing information about the alert that might help resolve the underlying issue.

If you do not need to perform any other triage actions, proceed to Phase 3: Close an alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Associate a knowledge base article with an alert

As an Event Management operator, you can associate a knowledge base (KB) article with the alert to capture additional information about the alert. This might include a procedure that someone has to follow to resolve the underlying issue on your network, or a best practice to prevent the issue from reoccurring.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Phase 1	✓	Analyze and acknowledge an alert
Phase 2	○	**Triage alerts**
Phase 3	○	Close an alert

This task assumes that your organization uses the Knowledge Base application in your ServiceNow instance.

Role required: evt_mgmt_operator

Procedure

1. From the Service Operations Workspace dashboard, open the alert that you acknowledged in Phase 1: Analyze and acknowledge an alert.
2. On the Alert form, click the lookup icon (🔍) next to the **Knowledge article** field.
3. Filter the list of existing KB articles by first selecting a field, such as **Short Description**, and then entering related text into the search text field.
   You can use the `contains` (*) operator to search for articles that contain keywords. For example, entering `*oracle` in the short description filters the KB articles that contain the word `oracle` somewhere in the short description.

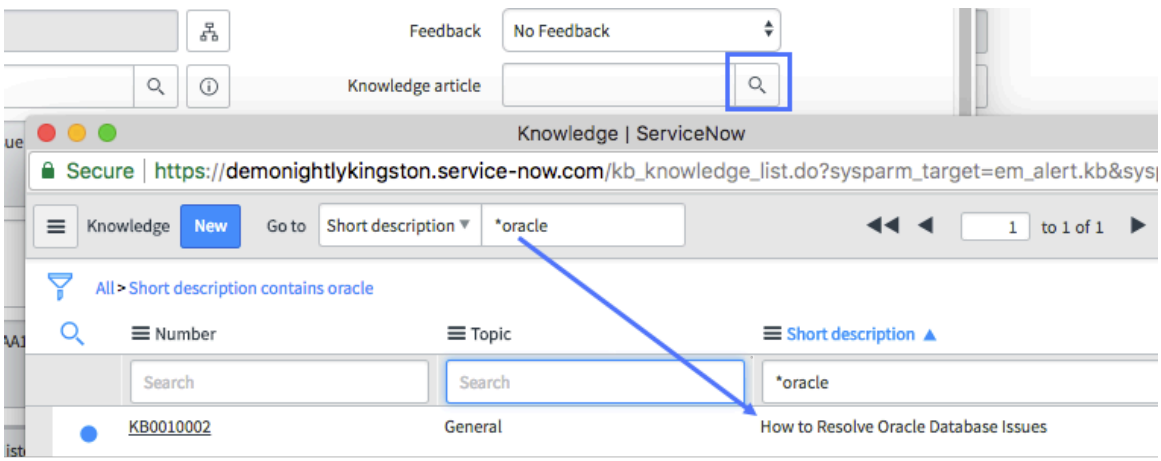

4. If you cannot find any related KB articles, you can click **New**, create a new one, and then click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The KB article number appears in the **Knowledge article** field on the Alert form.

5. Click **Update** on the Alert form to save the information.

What to do next

There are also other tasks you can take as part of the triage stage:

- Run a remediation workflow on an alert if your Event Management administrator already set up a workflow in your ServiceNow instance and your policies allow you to trigger it from the alert.
- Launch a web application from an alert to open a website or an event monitoring tool that provides more information about the alert.
- Put an alert into maintenance to temporarily hide it from the Service Operations Workspace dashboard if the alert does not require action at this time.

If you do not need to perform any other triage actions, proceed to Phase 3: Close an alert.

Put an alert into maintenance

As an Event Management operator, you can put an alert into maintenance if the alert does not require further action but you still want to keep it active. Putting the alert into maintenance hides it from the Service Operations Workspace dashboard so that other operators do not need to access it, but it does not close the alert.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Phase 1	✓	Analyze and acknowledge an alert
Phase 2	○	**Triage alerts**
Phase 3	○	Close an alert

Role required: evt_mgmt_operator

Procedure

1. From the Service Operations Workspace dashboard, open the alert that you acknowledged in Phase 1: Analyze and acknowledge an alert.
2. Select the **Maintenance** check box and click **Update**, or click the **Maintenance** button at the top of the list.
3. Click **Update**.
4. Navigate to **Event Management** > **Operators Workspace**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Notice that alerts in maintenance are not visible in the list by default. The filter excludes alerts where `Maintenance | = | true`.

What to do next

If the issue that led to the alert still needs further attention later, navigate to **Event Management** > **All Alerts**, open the alert, and then clear the **Maintenance** check box and click **Update**. From there you can perform other triage actions.

There are other tasks you can take as part of the triage stage:

- Run a remediation workflow on an alert if your Event Management administrator already set up a workflow in your ServiceNow instance and your policies allow you to trigger it from the alert.
- Launch a web application from an alert to open a website or an event monitoring tool that provides more information about the alert.
- Associate a knowledge base article with an alert if there is existing information about the alert that might help resolve the underlying issue.

If you do not need to perform any other triage actions, proceed to Phase 3: Close an alert.

Operator phase 3: Close an alert

After you take action on an alert, you can verify several items on the alert and then close it.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Phase 1	✓	Analyze and acknowledge an alert
Phase 2	✓	Triage alerts
Phase 3	○	**Close an alert**

Role required: evt_mgmt_operator

About this task

Procedure

1. Navigate to **Workspace Experience** > **Workspaces** > **Service Operations Workspace**.
2. Find the alert you want to close.
   If the alert does not appear in the list, for example because it is in maintenance, change the filter options as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click the alert number to open the Alert Details tab and review key information about the alert.

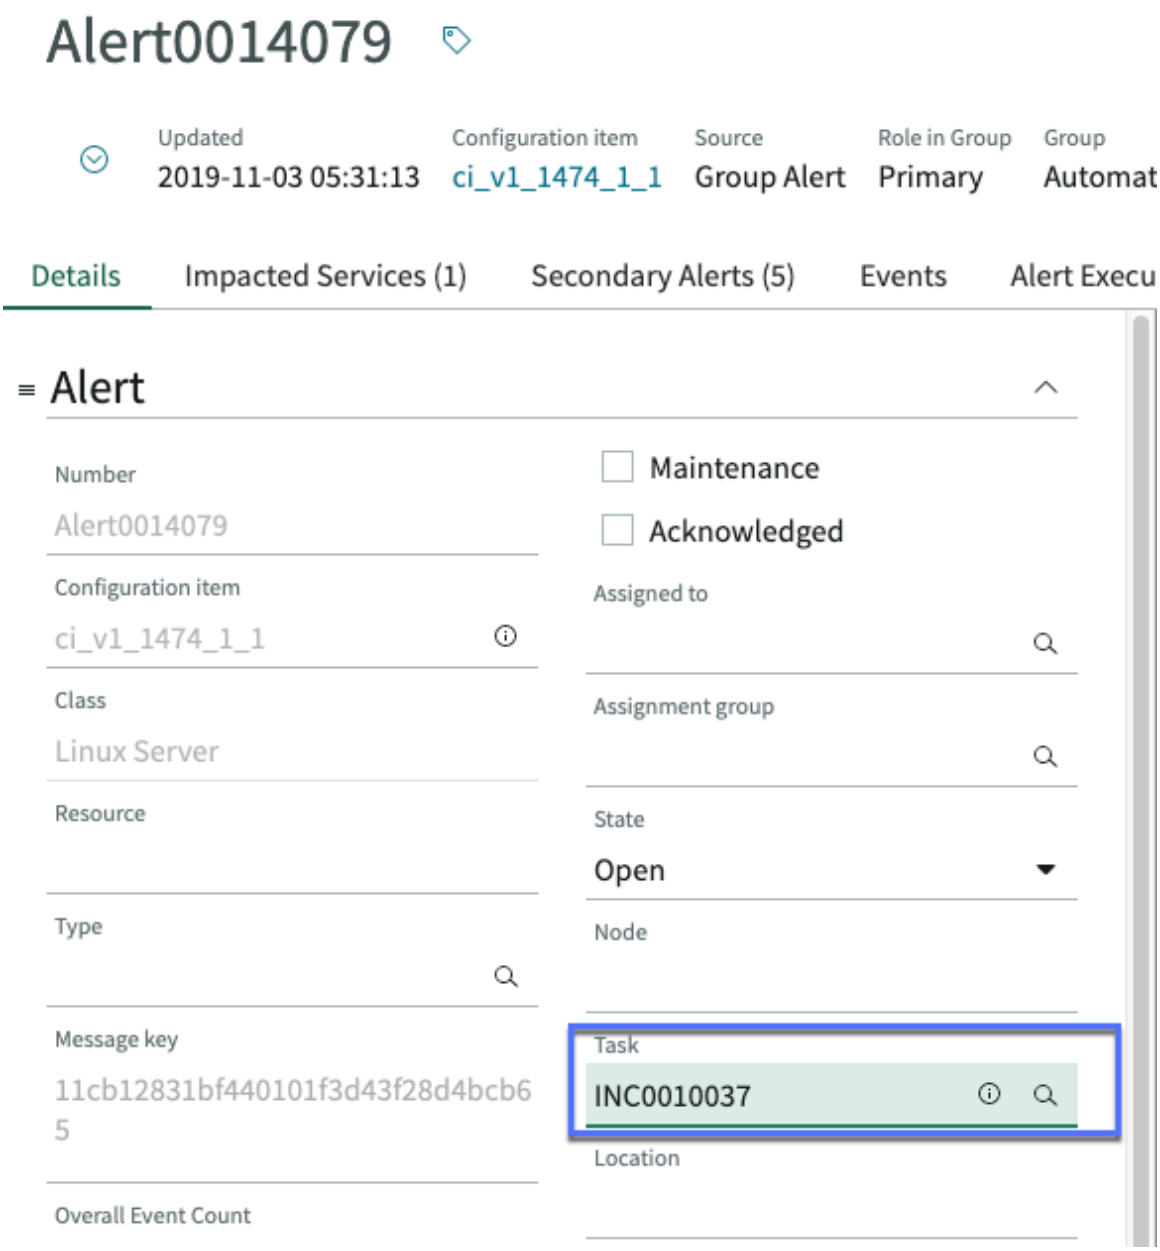

4. Review information about application services on the **Impacted Services** tab.

In this example, the service_1474 application service appears in the list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

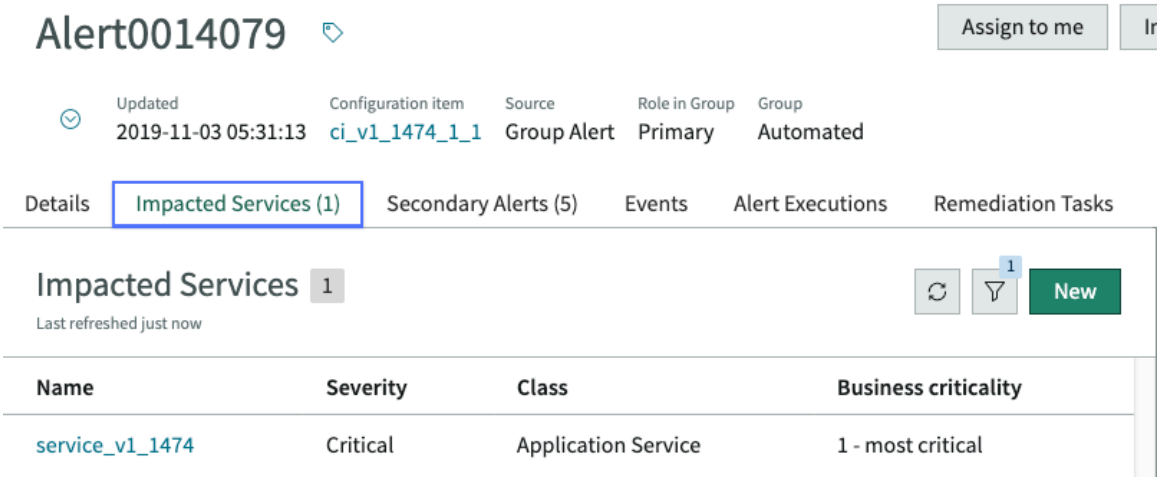

If you need more information about an application service and want to see the service map, click the service name and then click **Service Map**.

5. Look at the alerts related to this application service, and modify the correlation if necessary. The process for a primary alert differs from the process for a secondary alert.
   - On a primary alert, the secondary alerts that are correlated with the primary appear on the **Secondary Alerts** tab.

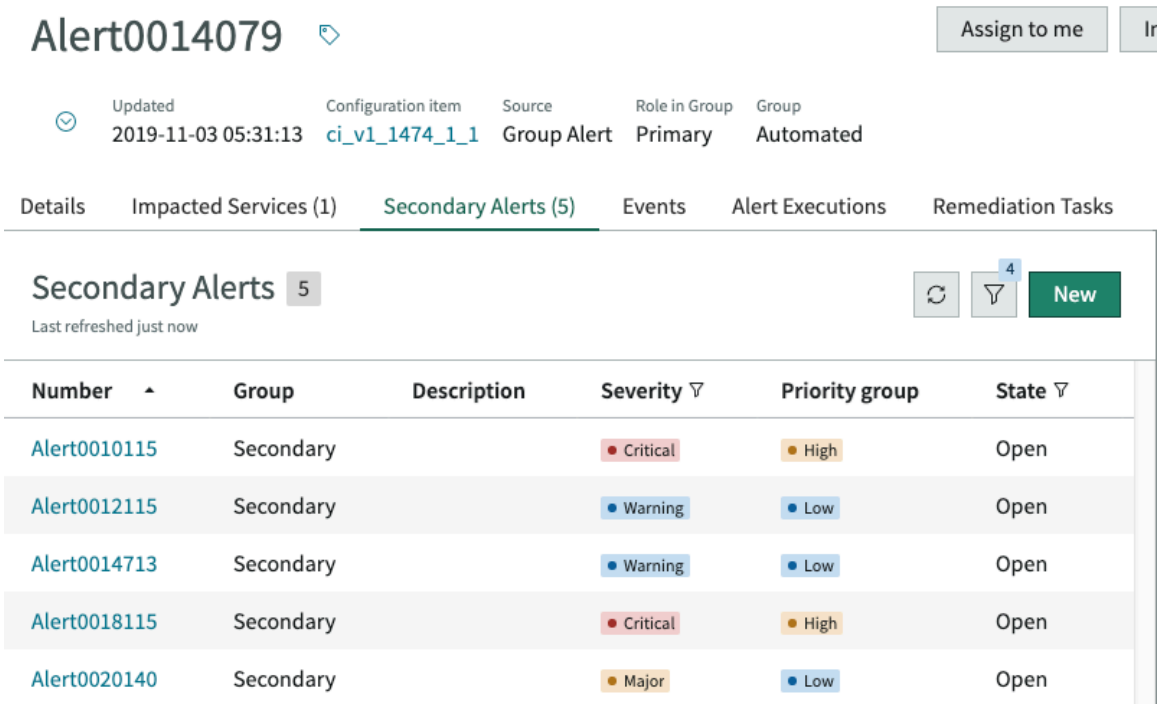

Typically, you do not have to do anything with the secondary alerts unless you want to add more secondary alerts or remove the correlation altogether.

Note: When you close a primary alert, the secondary alerts are also closed by default.

   - On a secondary alert:

     The alert number of the primary alert appears in the **Parent** field. You do not have to do anything with the parent alert, unless you want to re-correlate this secondary alert with another primary alert, or remove the correlation altogether by clearing the field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: When you close a secondary alert, the primary alert is not closed. You must close the primary alert separately.

6. When you are sure the underlying issue causing the event and corresponding alert has been rectified, click **Close**.
   If you think that the alert needs attention, do not close it. Instead, put it into maintenance by selecting the **Maintenance** checkbox.

What to do next

You are finished with the Event Management Operator tutorial, which covers a basic workflow of analyzing alerts, triaging them, and closing them out.

After you close an alert, you might need to follow up on these items:

Reopen an alert

If you think an alert requires further action, you can manually change its state to **Reopen**. When an alert reopens, any associated incidents can also be updated or reopened according to the incident state and the rules that your administrator configured.

Note: Any time you reopen an alert, always go through the triage and close phases again.

Work with alerts in the flapping state

In some cases, the same event that created the alert you just closed could be triggered again, even repeatedly, in a short amount of time. If this occurs, the **State** of the alert that you just closed changes to **Reopen**, and eventually to **Flapping** depending on settings that your administrator configures.

See Handle alerts in the flapping state for more information on what to look for.

For a summary of other tasks that you might need to perform at some time, see Advanced tasks for the Event Management operator.

Advanced tasks for the Event Management operator

As an Event Management operator, you might need to perform additional tasks that are outside of your typical workflow, or tasks that you need to perform only once.

Advanced task	Description
Customize your alert list view	Get to the alerts you need to analyze as soon as possible.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced task	Description
Work with flapping alerts	If an alert is in the flapping state, it means that the underlying issue keep reoccurring. You should triage the alert again and look for another solution.
Handle alerts while CIs are in maintenance	Learn how to handle alerts on a CI that is in the maintenance state, which means it is undergoing some type of change.
Adjust alert impact while triaging an alert	You can change impact rules, which affects how alert impact is calculated and alert severity.

Customize your alert list view

You can create one or more customized alert list views that show only the information pertinent to you. For example, you might want to focus only on the alerts for application services in a specific location, or only the alerts for application services of high criticality.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Role required: evt_mgmt_operator

Procedure

1. Navigate to **Event Management** > **Operators Workspace**.
2. On the **Lists** tab, click **My Lists**.

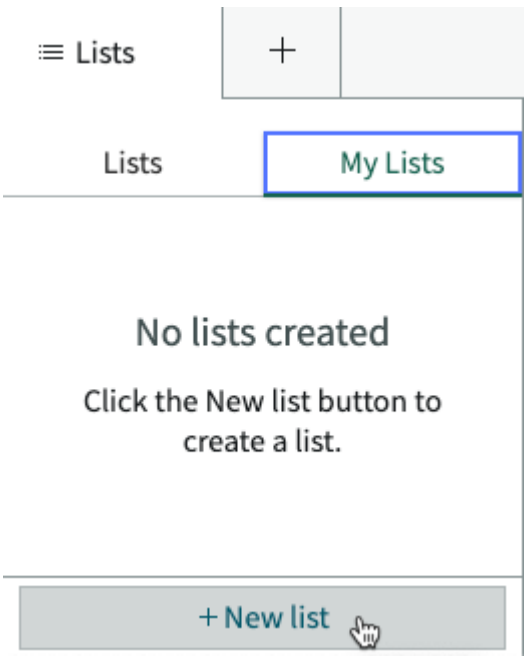

3. Click **+ New list**, fill in the form with name and filter conditions, and then click **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

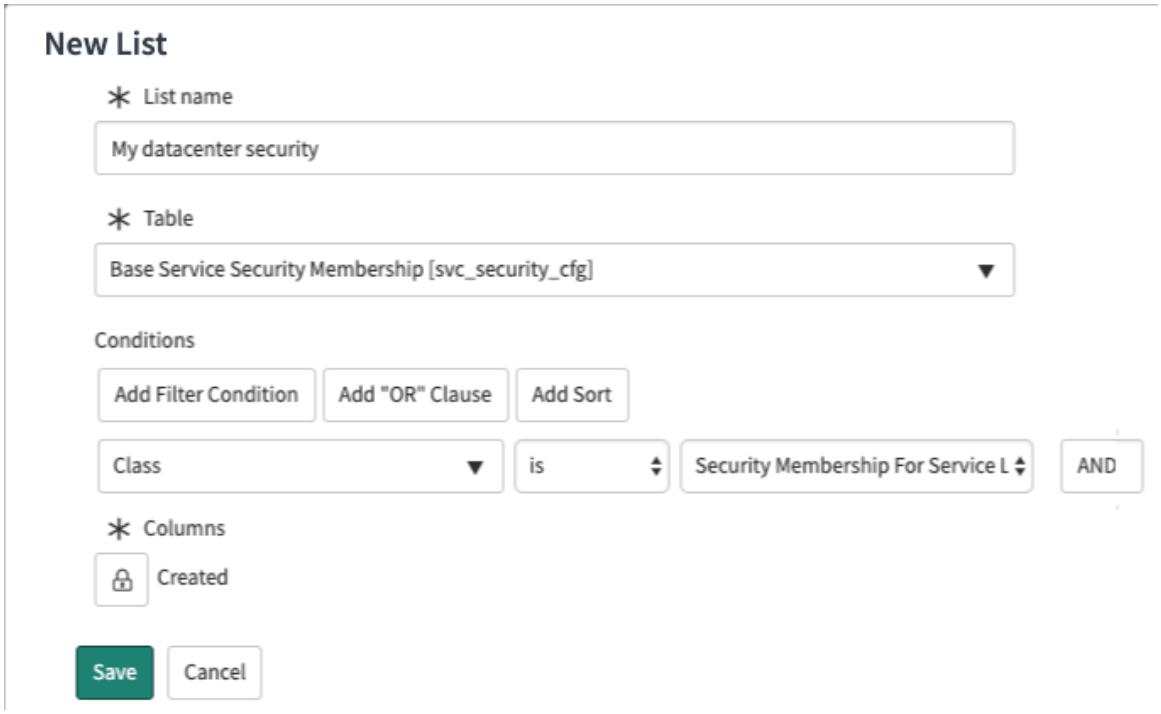

The list appears on the **My Lists** tab.

Note: The list applies only to application services, and not to application service groups. The groups appear whether or not they match the filter, and whether or not they contain any services that match the filter.

Work with flapping alerts

If an alert is in the flapping state, you might need to triage the alert again.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Role required: evt_mgmt_operator

About this task

If an event repeatedly triggers the same alert in a short amount of time, the alert is put into the flapping state, even if an operator closed it previously. The event can also be fluctuating between severity levels, such as Info and Critical. You should triage the alert again, and take action to try to prevent the event from reoccurring.

Note: Your Event Management administrator can configure several settings that determine when to mark an alert as flapping, and other settings that determine what action you can take. This topic covers a generic example.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Find alerts in the flapping state:
   a. Navigate to **Event Management** > **Operators Workspace**.
   b. Modify the list of alerts by clicking the filter icon ( ) and specifying criteria to find flapping alerts.
   Click **Advanced view** to modify filter settings. For example, you can specify:
   `State | is | Flapping`.
2. Click the alert number to open it.
3. On the alert, click the **Flapping** tab and review the information:

Field	Description
Flap count	The number of times the alert started flapping since the time in the **Flap start window** field.
Flap start window	The time that flapping started.
Flap last update time	The last time flapping occurred. This time is the instance processing time, not the source system time.
Flap last state	The state of the alert before it entered the flapping state.

4. Decide which action to take based on how often flapping has occurred.
   You can do any of the following, depending on what your organization's standard operating procedures are and what your Event Management administrator configured:
   - Run a remediation workflow that might do something like reboot a server or reset a router.
   - Access a web application, like the application for your event monitoring software on your network.
   - Modify an associated incident or a KB article to include additional information or steps.
5. After you are certain that the underlying issue is addressed, close the alert.

Handle alerts while CIs are in maintenance

When an alert occurs on a CI that is in the maintenance state, the alert state is also changed to maintenance. You should find and monitor the states of these alerts. Later, when changes on the CI are complete, finish triaging the alert and close it.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before completing this task, familiarize yourself with basic form configuration in an instance. You might need to add fields and related lists to forms.

The example in this topic assumes that a network administrator opened a change request to reconfigure a problematic web server.

Role required: evt_mgmt_operator

Procedure

1. Find alerts in maintenance:
   a. Navigate to **Event Management** > **Operators Workspace**.
   b. Click the filter icon ( ) and in the Filter box, click **Advanced view** and specify:

      `State|is|Maintenance`.
2. In the list of alerts, find an alert with a value in the **Configuration Item** column. This is the CI on which the alert occurred.
3. Click the CI name to open the form that displays information about the CI.

   In this example, the status of the web server is **In Maintenance** and there is a change request scheduled so that someone can reconfigure it.
4. Click the CI name to open the form that displays information about the CI.

   In this example, the status of the server (on the **Details** tab) is **In Maintenance**. You can view the scheduled change request on the **Change Requests** tab (someone will reconfigure the CI).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

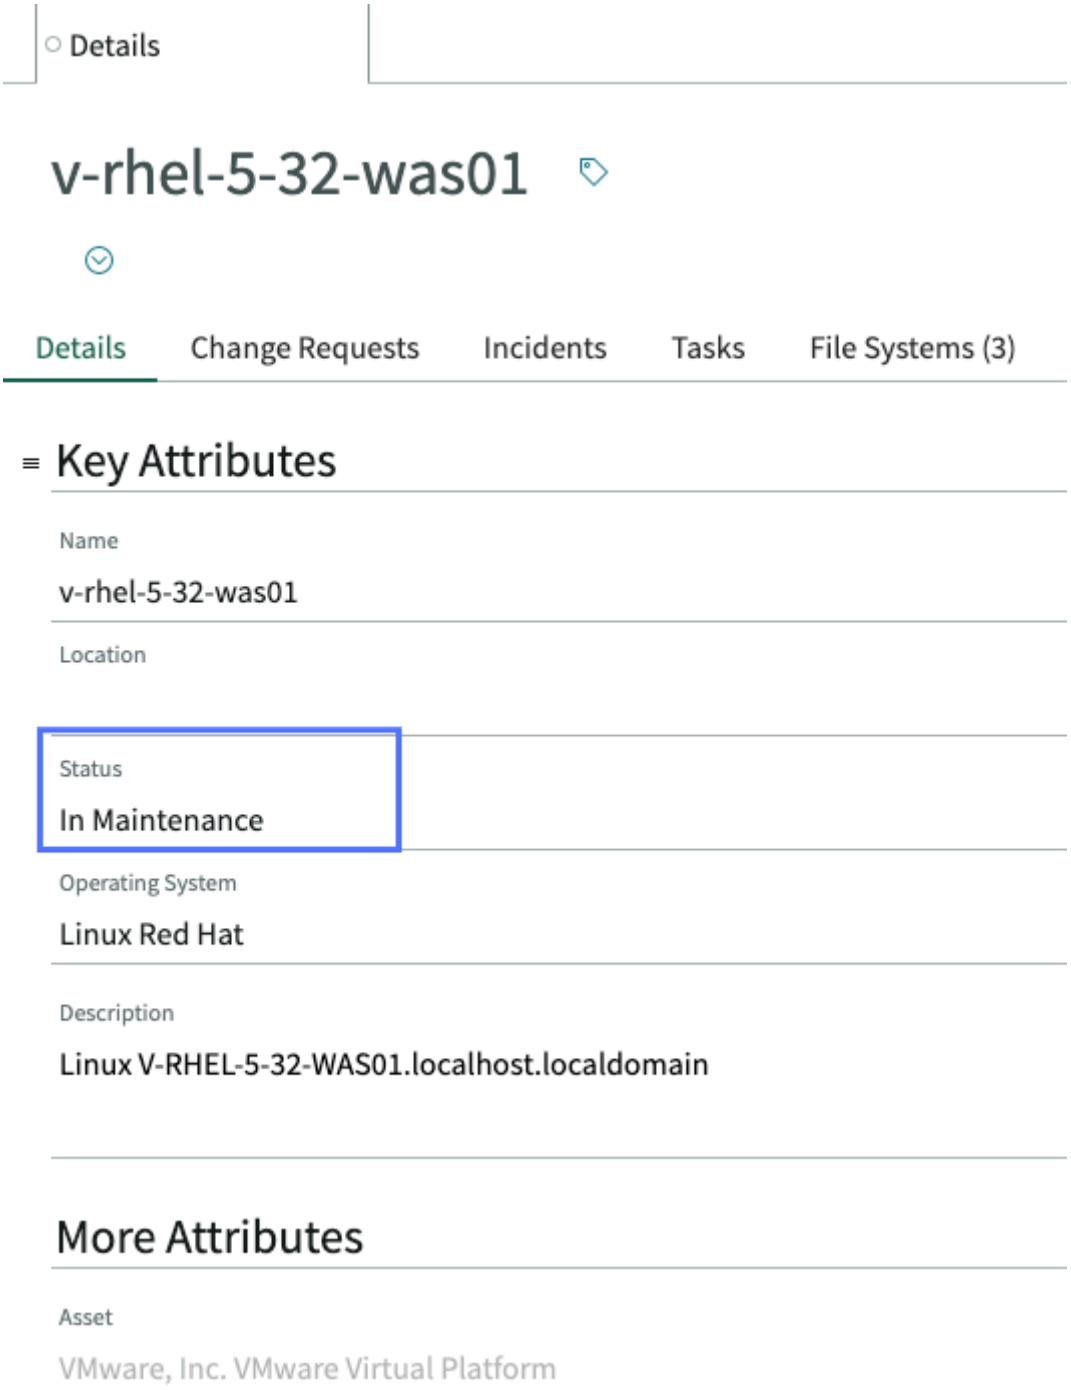

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

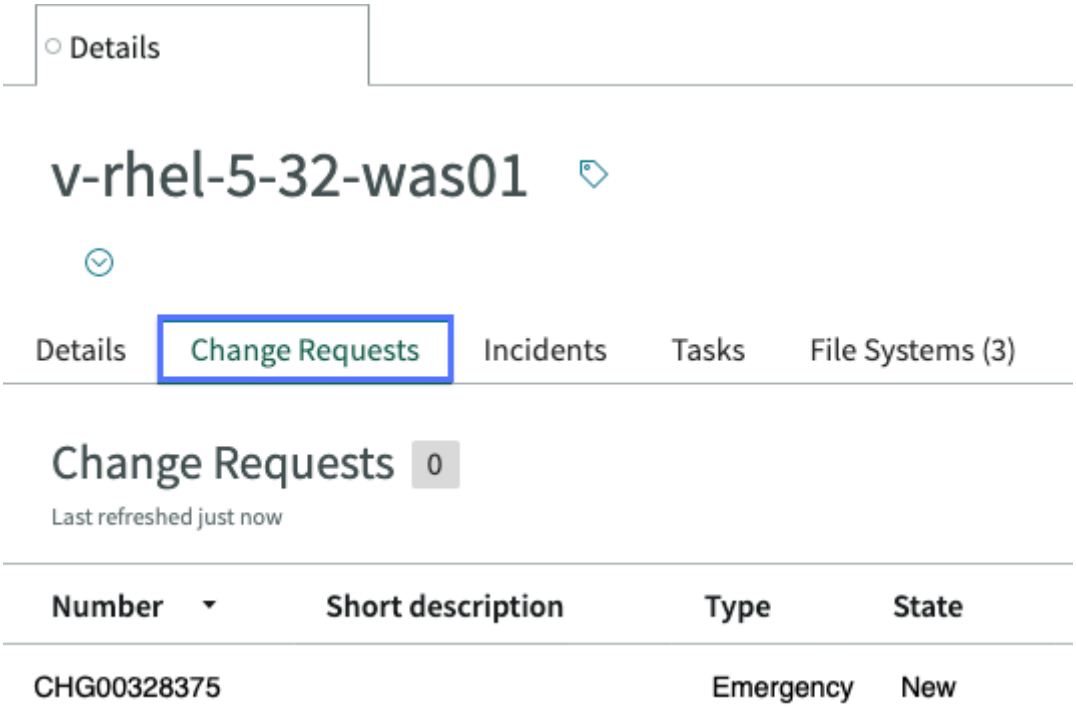

5. At this point, you should not modify the alert associated with this CI.
   You can open the change request to see if a schedule is attached. Coordinate with your team to find out when the CI will no longer be in maintenance.
6. After you verify that the **Status** field on the CI is not **In Maintenance**, open the alert again and clear the **Maintenance** check box.
7. Continue to triage the alert.

Adjust alert impact while triaging an alert

As an Event Management operator, you might need to modify the impact that an alert has on an application service and on the CIs in a service. Adjust the impact when it does not accurately represent your network environment.

Before you begin

Note: The Operator Workspace interface is available only to customers who have upgraded from a release prior to the Utah release. New customers as of the Utah release can use the Service Operations Workspace for ITOM, which offers an enhanced UI for managing alerts.

Role required: evt_mgmt_operator

About this task

Before you make any modifications to impact calculations, make sure that you understand how the application services and the CIs actually function. You might need to perform tests on network infrastructure or contact someone who can.

Procedure

1. Navigate to **Event Management** > **Operators Workspace**.
2. Double click the name of the application service that has the alerts you want to modify.
3. On the Details tab, click **Service Map**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. On the **Service Map** tab, click **Full Map** Tree.
5. On the map, click **Impact Tree**.
6. Analyze the application service to see whether the impact tree accurately reflects the impact of the alerts based on what you know about the CIs involved.

   In this example, a critical alert indicates a serious issue with a Linux server that is part of a Web portal application service.

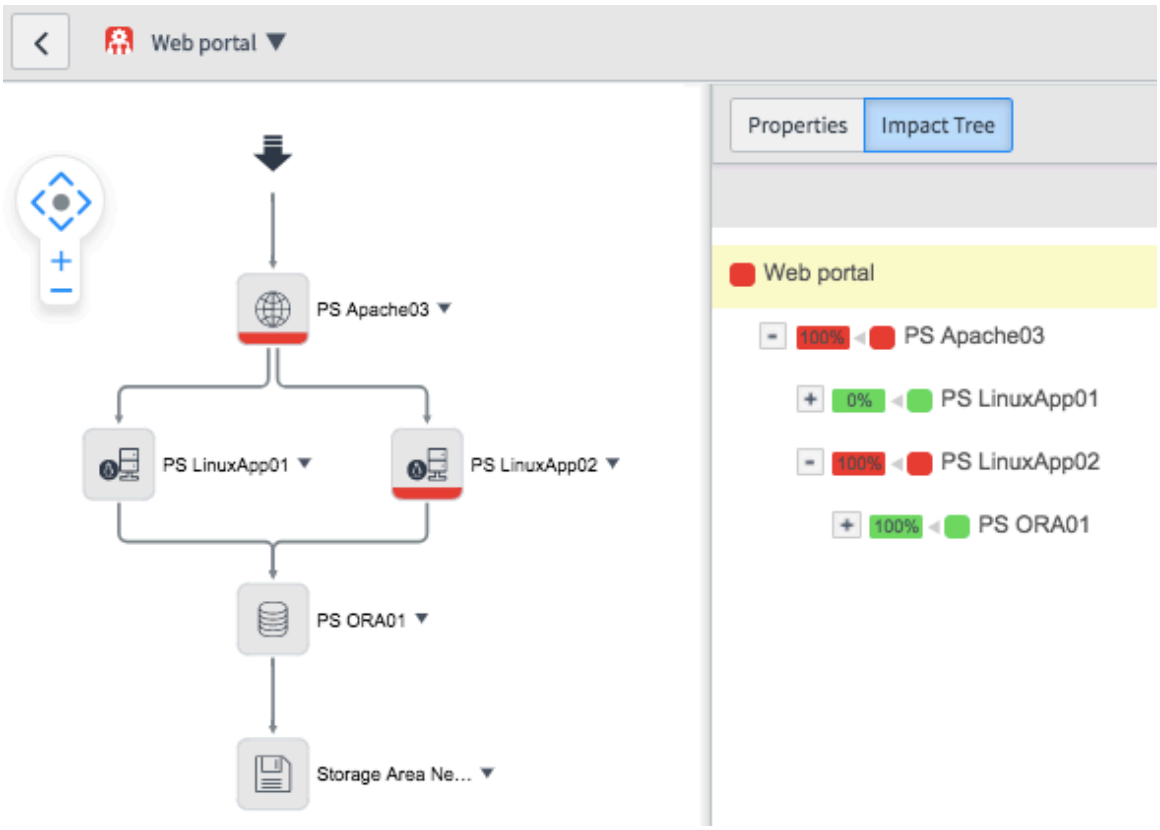

   Notice that the impact on the Web portal itself is red, meaning critical. This is true even though only one of the Linux servers that it relies upon actually has a critical alert. The other Linux Server is not impacted by the alert. Assume for this example that the Linux servers perform different roles for the web service. Even though **PS LinuxApp02** has a critical alert, the **PS LinuxApp01** server can adequately provide all the necessary services to keep the application service operational.

7. Click the **PS Apache03** CI, and then click **Impact** to display the impact rules for the CI. By default, the CI impact rule appears, showing you this information:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The **Impact On** and **Influence** columns show you that the CI (**PS Apache03**) should inherit 100% of the impact of all child CI impacts. In this case, the child CI impacts include the critical impact on the **PS LinuxApp02** server. This setting is why the **PS Apache03** web server, and the application service as a whole, inherits the impact 100%.
- The **Critical** column shows you that a critical alert (red) is inherited as critical alert (red).

8. Change the **Influence** value to 0 since the web server is still operational.
9. Change the **Critical** column to indicate that a critical impact status (red) should be inherited as a non-impact (green).
10. Wait for the dashboard to update automatically, or click **Event Management** > **Service Operations Workspace** and double-click the same application service to get an updated application service map.

    Notice that the influence on the Apache web server is 0% and the impact status is green. Also notice that the critical alert status (red) remains unchanged, even though the impact changed to green.

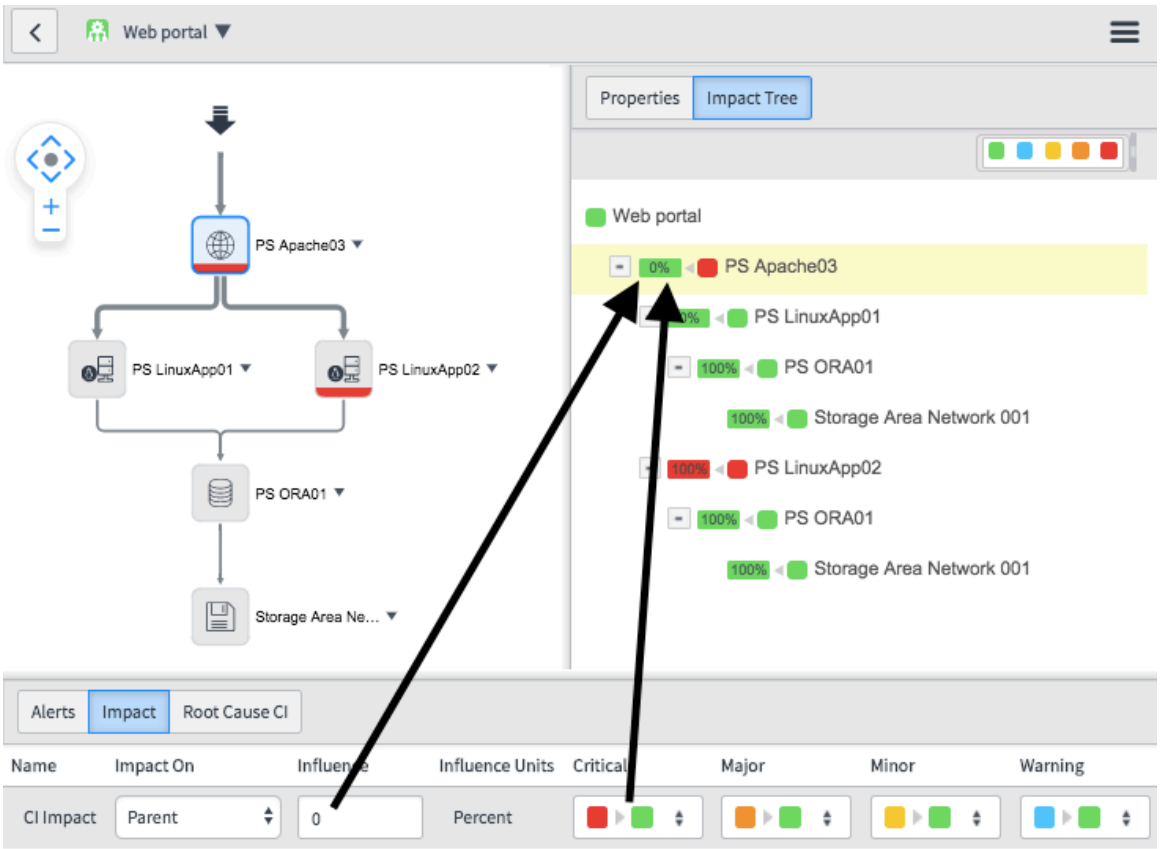

11. Navigate back to the Service Operations Workspace dashboard.
    Notice that the application service is now green, indicating that the overall application service is operational. The critical alert still displays at the bottom in the Alerts list, indicating that there is still an issue that you or other operators can address.

What to do next

Your administrator might configure additional impact rules that display when you click a CI in the application service map. You can make adjustments to the **Influence** fields on any impact rule to get the most accurate impact tree.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event Management reference

Reference topics provide additional information about mapping and fine-tuning application services using Event Management lists and forms.

Components installed with Event Management

Activating the Event Management (com.glideapp.itom.snac) plugin adds several roles, scheduled jobs, and tables.

Note: The Application Files table lists the components that are installed with this application. For instructions on how to access this table, see Find components installed with an application.

Roles installed with Event Management

Roles used by the Event Management application.

Event Management adds these roles.

Role title [name]	Description	Contains roles
Event Management Administrator [evt_mgmt_admin]	Has read and write access to all Event Management features to configure Event Management. **Note:** Exercise caution with the evt_mgmt_admin role, as it can be elevated to the admin role. A user with the evt_mgmt_admin role has the ability to add and modify scripts that run on a global scope. Ensure proper access control. With this role, the user can create and/or update the following scripts: • Alert correlation rules • Alert management rules • Maintenance rules • Advanced scripts • Event field mapping • Pre- and post-binding scripts	• evt_mgmt_user • template_editor_global
Event Management Operator [evt_mgmt_operator]	In addition to the evt_mgmt_user permissions, can also activate operations on alerts such as acknowledge, close, open incident, and run remediations.	evt_mgmt_user
Event Management User [evt_mgmt_user]	Has read access to all Event Management features. Has the itil role to be able to manage incidents that are created from alerts.	itil
Event Management Integrator	Has create access to the Event [em_event] and Registered Nodes	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description	Contains roles
[evt_mgmt_integration]	[em_registered_nodes] tables to integrate with external event sources.	
Event Management AIOps Manager [evt_aiops_manager]	Responsible for overseeing NOC teams performance and can onboard AI Specialists as part of the agentic workforce features in AIOps.	evt_team_operator
Event Management Team Operator [evt_team_operator]	Can create and manage their own integrations, set up their own alert automations, and enhance control over alert management for their teams.	

Scheduled jobs installed with Event Management

List of scheduled jobs that are provided with Event Management.

To review the list of scheduled jobs, navigate to **System Scheduler** > **Scheduled Jobs**. Event Management adds the following scheduled jobs.

Scheduled job	Description
Event Management — Connector execution job	Compares current time with time when active connector instances were last run and sets relevant connectors to execute. Runs every 10 seconds.
Event Management — Handle Impact Stuck Service	Releases stuck services and marks them as requiring rebuilding in the Impact Changes table to rebuild the impact tree. Runs every 2 minutes, 31 seconds.
Event Management — Impact Calculator Trigger	Trigger the impact calculation. The Event Management dashboard and impact tree are refreshed using the calculated figures. Runs every 6 seconds.
Event Management — Impact Topology Consumer	Consumes topology changes and marks the related services as 'require rebuilding' in the Impact Changes table to rebuild their impact trees. Runs every 19 seconds.
Event Management — Update stuck connectors	Release connector instances that are stuck. Runs every 2 minutes.
Event Management — Alert Priority Queue	Calculate alert priority. Two Alert Priority Queue jobs are active and available and can be run multi-thread. Runs every minute.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scheduled job	Description
Event Management — Auto Close Alerts	Alerts that are idle longer than 7 days (default time period) are closed. Modify the default using the `evt_mgmt.alert_auto_close_interval` property. Runs every 10 minutes.
Event Management — Calculate Alert Priority Grouping	Runs and calculates the priority groups: urgent, high, moderate and low for the alerts according to the highest and lowest priority score in the system. Runs every 30 minutes.
Event Management — Close Flapping Alerts	Close flapping alerts. Runs every 5 minutes.
Event Management — Close Threshold Alerts	Close threshold alerts. Runs every 2 minutes.
Event Management — Evaluate Scoped Alert Rules Management0	Execute alert management rules. Two separate jobs run to ease the load on the system. Runs every 11 seconds.
Event Management — Evaluate Scoped Alert Rules Management1	
Event Management — Filter Services for Impact - Queue Handler	Process changes from em_impact_inclusion_class and em_impact_filter_service. Runs every minute.
Event Management — Impact Tree Builder	Handles all services with changes from the em_impact_changes table and rebuilds their impact trees. Runs every 11 seconds.
Event Management — Insert Health Monitor	Job to produce the **ServiceNow Event Management** application services. Runs once every hour.
Event Management — Maintenance Calculator	Calculate the maintenance for CIs. Runs every minute.
Event Management — Node Count	Calculate license usage. Runs once every hour.
Event Management — Populate Impacted Services Table	Scan cmdb_ci_service_auto and copy all services to em_impacted_service, along with all related classes to em_impact_inclusion_class.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scheduled job	Description
	Runs once after upgrade.
Event Management — Process Events	Job that runs and processes each Ready event (apply event rule, mapping rule, and create or update alert) Runs every 5 seconds.
Event Management — Process Metric Binding	Process metric binding. Runs every 5 seconds.
Event Management — Process Records in em_extra_data_json	Creates events for log analytics anomalies. Runs every 13 seconds.
Event Management — Queue Connector Processor	Bi-directional functionality. Processes all pending alerts in the Update Queue and sends them to the MID Server. By default, this dequeue process is performed in batches of 1,000 alerts. Runs every 30 seconds.
Event Management — Recalculate Impact for Groups	By default, this job is not active. Can be run on demand to correct the impact on service groups. Runs on demand.
Event Management — Recover Stuck Events	Handle all events that are in queued state and switch back to Ready to handle events from the beginning. Runs at system startup.
Event Management — Update Health Monitor	Update the **ServiceNow Event Management** application services. Runs once every hour.
Event Management — Update SLA Configuration Result	Synchronizes the CIs that match the SLA configuration filter with the Event Management SLA [em_ci_severity_task] table. Runs every 10 minutes.
Event Management — Update SLA Severity	Updates Event Management SLA [em_ci_severity_task] table with the new severity. Runs once every minute.
Event Management — Convert IT Service	Run this property on demand to convert manual services to application services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scheduled job	Description
	Runs every 30 minutes.
Event Management — Collect xmlstats	Collect event processing statistics. Runs once every minute.
Event Management — Impact Calculator for Alert Groups and SLA	Calculates the effect of alerts on alert group services.
Event Management — Manual Impact Calculator Trigger	Manually trigger of Impact Calculation on all operational Services.
4x Event Management — Impact Calculator for BS_0	Calculates the impact of alerts on application services, and builds the impact tree on the Event Management dashboard. Four separate jobs run to ease the load on the system.
4x Event Management — Impact Calculator for BS_1	
4x Event Management — Impact Calculator for BS_2	
4x Event Management — Impact Calculator for BS_3	
Event Management — Backfill Alert History table	Locates redundant records in the em_alert_history table and cancels them.
Event Management — Backfill Impact Status table	Locates redundant records in the em_impact_status table and cancels them.
Event Management — Impact for Groups	Calculates the impact of alerts on service groups.
Event Management — Clean Alert History Table	Cleans the Alert History (em_alert_history) table by removing records more than 90 days old. You can customize the amount of time after which alerts are removed by configuring *evt_mgmt.impact_calculation.cleanup_age_sec* in the sys_properties.list table.
Event Management — Clean Impact Status Table	Cleans the Impact Status (em_impact_status) table by removing records more than 90 days old. You can customize the amount of time after which alerts are removed by configuring *evt_mgmt.impact_calculation.cleanup_age_sec* in the sys_properties.list table.

Event Management adds the following scheduled jobs to support alert aggregation and RCA.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
Service Analytics Purge Old Observation Data — Daily	Cleans the staging data.
Service Analytics Prepare RCA Learner Input Data — Daily	Prepares RCA input data. Stores and probes MID Server to learn statistical information about alerts.
Service Analytics group alerts using RCA/Alert Aggregation	Applies RCA and alert aggregation to open alerts and prepares automated alert groups.
Service Analytics Alert Aggregation Learner — Daily	Learns information about existing alerts and groups new open alerts.
Service Analytics RCA Configuration	Configures root cause analysis.
Service Analytics Check File System Space on Analytics MID — Daily	Checks disk usage on the dedicated MID Server, and generates an event if it exceeds the threshold set in the *sa_analytics.rca.mid_max_allowed_space* property.
Service Analytics Gather Value Report Data — Daily	Gathers data for the Value Report.
Service Analytics — Update virtual alerts for aggregation groups	Update the virtual alerts that were created to represent alert aggregation groups, with any changes to alerts belonging to that group. Runs every minute.
Service Analytics Attribute Populator for Historical Alerts	Populate attributes used in feature identifier for historical alert data using event rules. Runs on demand.
Event Management - Analytics Alert Syncer	Gathers alert data that is used by grouping job.

Domain separation properties for Event Management connectors

Properties provide the metadata to identify the domain.

You can change the values if you want to use any other table or fields for domain identification but make sure that the table is domain separated.

Note: When personalizing domain separation for event management connectors, all the properties can be overridden at the connector level as well.

There are three caches maintained for personalizing domain separation (main, user_access, missing_domain). You can create a few system properties to control them.

Default System Properties

Property	Description
evt_mgmt. connector_domain_info_table_name	The table where you store domain information for personalizing domain separation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Default System Properties (continued)

Property	Description
	• **Type**: string • **Default value**: core_company
evt_mgmt. connector_domain_info_column_name	The field name in the table to identify the provided domain for personalizing domain separation. • **Type**: string • **Default value**: name
evt_mgmt. connector_domain_id_column_name	The field to get domain ID from for personalizing domain separation. • **Type**: string • **Default value**: sys_domain
evt_mgmt. connector_domain_path_column_name	The field to get the domain path from for personalizing domain separation. • **Type**: string • **Default value**: sys_domain_path

Property	Description
evt_mgmt.connector_custom_domain_sep_cache_expire_in_seconds	For personalizing domain separation, if you don't want your main cache to expire every week.
evt_mgmt.connector_custom_domain_sep_user_access_on_domain_cache_size	For personalizing domain separation, if you want to increase the size of the user access cache.
evt_mgmt.connector_custom_domain_sep_missing_domain_cache_expire_in_seconds	For personalizing domain separation, if you want to increase the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	expire time for cache storing information regarding the missing domain.
evt_mgmt.connector_custom_domain_sep_cache_size	For personalizing domain separation, if you need more main cache size for storing domain information.

Tables installed with Event Management

Tables that are provided when Event Management is activated.

Event Management adds these tables.

Table	Description
Alert [em_alert]	Alerts that Event Management manage.
Alert Correlation Rule [em_alert_correlation_rule]	Rules specifying primary and secondary correlated alerts.
Alert Aggregation Group Alerts [em_agg_group_alert]	Stores alerts associated with aggregated alert groups.
Alert Aggregation Group [em_agg_group]	Relationships between aggregated groups and primary alerts.
Alerts History [em_alert_history]	History of alerts. Used for impact calculation.
Alert Rule [em_alert_rule]	Mappings of alert fields to the Incident [incident] table.
Alert Template [em_alert_template]	Alert templates. This table extends the Template [sys_template] table.
Event Management SLA [em_ci_severity_task]	Event Management SLA tasks for CIs and services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Connector Definition [em_connector_definition]	Settings for gathering events from external event sources.
Connector Instance [em_connector_instance]	Connection details for external event sources.
MID Server to Connector Instance [em_connector_instance_to_mid]	Mappings of MID Servers to connector instances.
Event Management License Usage [em_unique_nodes]	When events are received by ITOM AIOps, an entry is added or updated in this table based on the monitored target specified in the received message. The entry links to its corresponding CMDB CI. If none is found, the entry is assigned *Type = Unknown*.
Event [em_event]	Events received by Event Management.
Event Filter [em_event_filter]	Storage for defined event filters.
Event Match Rule [em_match_rule]	Updated events for alert processing. Used by event rules.
Event Match Field [em_match_field]	Mappings of event fields to alert fields. Simple mapping. Used by Event Rules.
Event Compose Field [em_compose_field]	Mappings of event fields to alert fields. Composite mapping. Used by Event Rules.
Event Mapping Rule [em_mapping_rule]	Updated event fields for alert processing.
Event Processing Statistics [em_event_stats]	Statistics on Event Management performance.
Event Type [em_event_type]	Event types.
Task Template [em_incident_template]	Templates that define how to populate new tasks. For example, how fields of an incident that is being created from an alert, must be populated. This table extends the Template [sys_template] table.
Registered Nodes	Registered nodes data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[em_registered_nodes]	
Threshold Rule [em_threshold_rule]	Alert threshold rules.
Binding Device Map [Em_binding_device_map]	Event binding to network paths and storage paths.
Process to CI Type Mappings [Em_binding_process_map]	Event binding to specific processes.
CI Remediation [em_ci_remediation]	Remediation rule definitions.
Impact Graph [em_impact_graph]	Impact tree of CIs containing CI hierarchy and impact rules to be used for impact calculation.
Impact Graph History [em_impact_graph_history]	History of changes in impact tree.
Impact Rule Definitions [em_impact_rule_definition]	Definition of rules used for impact calculation.
Impact Rule instance [em_impact_rule]	Rules based on impact rule definitions.
Impact Inclusion Classes [em_impact_inclusion_class]	Specifies the classes to be used for impact calculation.
Impact Filter Services [em_impact_filter_service]	Specifies services to be included or excluded from impact calculation.
Infrastructure Relations [em_impact_infra_rel_def]	Child-parent pairs or CI types. CIs matching these definitions are added to impact trees.
Impact Maintenance CIs [em_impact_maint_ci]	CIs that are in maintenance and therefore are excluded from impact calculation.
Impact Status [em_impact_status]	Calculated status of CIs and services to be displayed in the dashboard and service maps for dynamic CI groups.
SLA Configuration [em_sla_configuration]	SLA configuration records that identify the CIs that SLAs can run on.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Service Analytics Metric Type Registration [sa_metric_registration]	Source registration details for processing raw data.
Health monitor scripts [em_monitor_scripts]	These scripts determine how to monitor or check, for example, when using the Connectors Monitor script. You can create customized script to monitor a device or an entity. The scripts provided with the base instance are: • Check delay in event processing • Connectors Monitor • Get Event Processing state • MID Server Threshold Alerts
Monitoring configuration [em_monitor_conf]	Use this table to configure what to monitor according to the scripts that are listed. Configure how often to run each script. If a script has a threshold, it determines what alert severity to display. Threshold values are in units of minutes and specify the delay time. Navigate to **Event Management** > **Settings** > **Self-Health configuration** to see the list of Monitoring Configurations or to create a new one. Use this script to test Data Center Monitoring. The scripts provided with the base instance are: • Connector's idle state monitoring-monitor to verify whether any of the connectors was in idle state that surpassed the threshold [in minutes] that was configured. • Connectors Status- monitor to track the active status of the connectors. • Delay in event processing-monitor to track the duration [in minutes] of events that remained in 'ready' state and were not processed. • Event Processing job-monitors the state of the event processing jobs. • MID Server Threshold Alert-monitors MID Server health.
Monitoring Event Management Jobs [em_monitor_jobs_state]	Monitor Event Management jobs by adding the relevant jobs to the table. **Note:** The following jobs are not monitored by this table: • Event Management - Process Events • Event Management - Impact Tree Builder • Event Management - Recover Stuck Events • Event Processing

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Monitoring state [em_monitor_state]	Use this table to set the threshold for each connector. When there is a value above the threshold, an alert is generated.
EM XMLStats Data [em_xmlstats_data]	Self-health statistics and diagnostic details for Metric Intelligence and Event Management, which are used to produce the XMLStats page.

Event Management adds the following tables to support alert aggregation and RCA.

Table	Description
SA RCA Status [sa_rca_status]	Information (such as IDs) for the latest messages that were sent to the ECC Queue for a service during RCA.
SA RCA Output [sa_rca_output]	RCA learner output data.
SA RCA Group [sa_rca_group]	Automated alert groups for the RCA query.
SA Analytics Alert Staging [sa_analytics_alert]	Staging table for alerts used for analytics.
SA RCA Input [sa_rca_input]	Input data for the RCA learner.
SA Analytics Status [sa_analytics_status]	Last run information to be used for alert aggregation and RCA.
SA RCA Group Alert [sa_rca_group_alert]	Alerts associated with automated alert groups.
SA RCA Service Configuration Item Association [sa_rca_svc_ci_assoc]	Associations between CIs and services.
SA Alert Aggregation Learned Pattern [sa_agg_pattern]	Learned patterns from alert aggregation.
SA Alert Aggregation Learned Pattern Elements [sa_agg_pattern_element]	CI/Metric Name pairs associated with learned patterns.
SA Alert Aggregation Query Group Patterns	Relationships between groups discovered in alert aggregation queries and patterns found in learning.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sa_agg_group_pattern]	
SA Alert Aggregation Query -- Staged (Recent) Alerts [sa_agg_group_alert_staging]	A staging table for alerts that have not yet been associated with any aggregated alert group.
SA Agg Pattern Attribute [sa_agg_pattern_attribute table]	CI/alert attributes to be used for finding patterns for alert aggregation.
SA Alert Attribute Populator Status [sa_alert_attribute_populator_status table]	State and statistics for attribute populator job.
SA Alert Aggregation Learned Pattern Elements Pair wise Mutual Information and Joint Probability [sa_agg_pattern_element_pair]	Pairwise probabilities for pattern elements.
EM Agg Group Prediction [em_agg_group_prediction]	Alert predictions for alert groups.

Event Management Platform Analytics Solutions

Platform Analytics Solutions contain preconfigured dashboards. These dashboards contain actionable data visualizations that help you improve your business processes and practices.

Platform Analytics Solutions

Platform Analytics data visualizations use Performance Analytics indicator data to show you data over time, helping you analyze your business processes and identify areas of improvement. With Platform Analytics Solutions, you can get value from Performance Analytics for your application with minimal setup. You can always create your own objects as well.

To enable the solution for Event Management, an admin can navigate to **Performance Analytics** > **Guided Setup**. Click **Get Started** then scroll to the section for Event Management. The guided setup takes you through the entire setup and configuration process.

Data visualization in ITOM now includes a funnel component to help you track metric progression across stages, making it easier to identify drop-offs or bottlenecks in processes. For more details, see Data visualization in ITOM.

Related topics

Activate your Performance Analytics subscription

Customize Platform Analytics elements

Customize the appearance of widgets when viewing Platform Analytics visualizations.

Before you begin

Role required: user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Performance Analytics** > **Dashboards**.
2. On the **Dashboards** page, select a dashboard to view.
3. Select the **Edit** button.
4. Select the relevant element to open the Configuration panel.
5. In the Configuration panel, edit the element's options.
6. Select **Save**.

What to do next

Users with the admin role can customize colors and fonts. To customize fonts, navigate to **sys_ux_theme_property.list** and edit the applicable fonts.

Machine learning solutions for Event Management

Build solutions for Event Management with Predictive Intelligence.

Solution definition

Solution Definition for Event Management

Solution Definition	Description
Closed alert similarity	Closed alert similarity is a similarity solution definition that is used to recommend similar Alert records based on these fields: Description, Node, Resource, CI Type,Source, and Metric Name from the Alert [em_alert] table. Results are visible to users on the Alert form and on the Similar Alerts tab in Workspace.
EM Alert Clustering Solution	EM Alert Clustering Solution is a clustering solution definition used in Event Management. Allows building clustering models enabling identification of common text patterns in alerts to create alert groups. Text clusters are based on these fields: Description, Metric Name, Configuration Item Class from the Alert [em_alert] table.

Performance Analytics Edit Widget dialog box

Fields in the dialog box for editing a Performance Analytics dashboard widget.

Section	Description
Show real-time updates	Select for updates to appear in the widget in real-time.
Show	Select or clear the boxes to show or hide the following on the widget: • Border • Header • Title
Title alignment	Select the alignment of the widget title:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Section	Description
	• Left • Center • Right
Title color	Select to configure the color of the widget title.
Header color	Select to configure the color of the widget header.
Interactivity	Select from the displayed options. • Follow interactive filter • Show when following filter

Event Management MID Web Server extension form

Fields in the form for creating or modifying a MID Web Server extension.

MID Web Server Context form

Field	Description
Name	Unique name for the extension for easy identification.
Short description	Brief, meaningful description of the extension.
Status	This field is auto-populated with the status of the extension. The field is blank until the extension is started. After issuing a command to the extension, one of the following values is displayed: • Started: The extension is running. • Stopped: The extension is not running. • Offline: The MID Server is down. • Error: The extension failed with an error. The error message is displayed in the **Error Message** field. • Warning: A run-time exception has occurred. The extension continues to work.
HTTP/HTTPS Port	Port number the extension listens on for incoming requests.
Authentication Type	The available authentication types are: • **API Key** (default): Any client connecting to the extension must enter an API key into the request's authorization header. For details on viewing existing API keys, see Configure MID Web Server API key authentication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MID Web Server Context form (continued)

Field	Description
	• **Keybased** (deprecated and kept only for compatibility with releases prior to the Brazil release): The value in the **Secret Key** field is used to generate an authentication token sent in the **Authorization** header of incoming client requests. For details, see Configure key-based MID Web Server authentication. • **Basic**: Define a username and password in the **Basic Auth User** and **Basic Auth User Password** fields, which appear when selecting this option. • **mTLS**: A certificate-based authentication that uses the TLS protocol. For details, see MID Web Server and agent mTLS Authentication.
Secret Key	The value generated when Keybased authentication is selected for the value in the Authentication Type field.
Basic Auth User	Username to be used when selecting Basic authentication.
Basic Auth User Password	Password to be used when selecting Basic authentication.
Secure Connection	Option to indicate that the extension is to use the TLS protocol to encrypt incoming connections. A certificate and private key combination must also be provided to the extension. When selecting this option, the **MID Server Unified Keystore** check box appears.
Use MID Unified Keystore	Option for using the MID Server Unified Keystore option for a secure connection. Clear this option to use **Web Server Keystore**. For details, see Configure a secure MID Web Server extension.
Keystore Certificate Alias	The alias name of the keypair inside the selected keystore that the extension uses to establish a secure connection. Leave this field empty to use the default alias when using the MID Unified Keystore or when using the Web Server Keystore with a single alias.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MID Web Server Context form (continued)

Field	Description
Keystore Password	The password for the Web Server Keystore when it is used to establish a secure connection.
Keystore Key Password	The password of the private key entry mapped to the specified alias of the Web Server Keystore. This field can be left empty if the password matches the entry in the Keystore Password field.
Truststore Password	The password to the truststore that holds the Certificate Authority (CA) certificates used to validate the signatures of the client certificates. Appears when the **mTLS** authentication type is selected. Leave this field empty if the extension is configured to use the cacerts truststore of the JRE running the MID Server.
Error Message	Message describing any error that causes a command, such as Start or Stop, to fail. This field appears only when the value in the Status field is Error.
Execute on	Location for running this extension. The available options are **Specific MID Server** or **Specific MID Server Cluster**.
MID Server	Depending on your selection in the **Execute on** field, the name of the designated MID Server or MID Server cluster. • If you selected **Specific MID Server**, the name of the designated MID Server. • If you selected **Specific MID Server Cluster**, the name of the designated MID Server cluster.
Executing on	Displays the name of the MID Server on which the extension is running even if the MID Server is down. If the user stops the extension, this field is empty.

Event Management tag based alert clustering tag form

The form for creating or modifying a tag based alert clustering tag displays detailed information about the tag.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag based alert clustering tag form

Field	Description
Name	Name of the clustering tag. Defaults to a description of the configured tag (such as, **Exact match on Alert Field "metric name"**). The default name is visible only after saving the tag. Tag names must be unique.
Customized name	Select the check box to customize the value in the **Name** field.
Domain	The domain in which the current record was created. Read-only.
Description	Enter an optional description of the tag.
Source	Select the source from which to choose the field to be matched. • **Alert field** • **Alert additional info** • **Alert CI** • **Alert CI key** • **Alert tags** If you select **Alert additional info**, choose an **Additional Info Key** instead of a field. **Alert additional info** is a field on alert containing the key value pair in Jason format and we can use a specific key. **Alert CI key**: Create CMDB key values to cluster alerts based on CMDB Key Values [cmdb_key_value] table. Key values provide an additional method for determining commonalities between alert CIs.
Selected field	Indicate the field that has to match between alerts for the alerts to be included in a group. Appears when you select **Alert Field** or **Alert CI** in the **Source** field. If you select **Alert CI** in the **Source** field, you need to select a CMDB key. A CMDB key is assigned to the CI of the alert.
Key	Indicate the Key that has to match between alerts for the alerts to be included in a group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag based alert clustering tag form (continued)

Field	Description
	Appears when you select **Alert additional info** in the **Source** field, or appears when you select **Alert tags** in the **Source** field. When **Alert additional info** is selected, enter the name of one of the fields in the additional info field. When **Alert tags** is selected, enter the name of the relevant alert tag.
CMDB key	Indicate the CMDB key to match for the alerts to be included in a group. Appears when you select **Alert CI Key** in the **Source** field.
Match method	Select the type of match required for the alerts to be included in a group. • **Exact**: Indicates that the field value must be an exact match for the alert to be included in a group. For example, you can configure an alert clustering tag indicating that the alert's **Metric name** field must be an exact match to form a group. When invoking that tag, all alerts with identical values in the **Metric name** field are included in the same group. • **Fuzzy**: Indicates that the field value needs to be an approximate match (depending on the value configured in the **Similarity** field) for the alert to be included in a group. • **Pattern**: Indicates that the field value needs to follow the pattern in the **Pattern** field. For correct syntax and usage examples, see Pattern matching.
Similarity	Specify the similarity percentage that must be met by the alerts to be included in a group. For example, entering 50 indicates that at least 50 percent of the indicated value must appear in the alert for the alert to be included in the group. Appears only when **Fuzzy** is selected as the **Match Method** value. Default value = 90

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

List of predefined alert grouping tags

A list of the predefined alert clustering tags provided with the Tag Based Alert Clustering Engine # application.

Predefined alert clustering tags

Name	Source	Match method
CI	em_alert.cmdb_ci	Exact match
Description	em_alert.description	Exact match
Metric Name	em_alert.metric_name	Exact match
Node	em_alert.node	Exact match
Resource	em_alert.resource	Exact match
Source	em_alert.source	Exact match
Source instance	em_alert.event_class	Exact match
Type	em_alert.type	Exact match
Assignment group	em_alert.assignment_group	Exact match
CI Class	cmdb_ci.sys_class_name	Exact match
Dns domain	cmdb_ci.dns_domain	Exact match
ip address	cmdb_ci.ip_address	Exact match
Location	cmdb_ci.location	Exact match
Environment	cmdb_ci.environment	Exact match
CI name	cmdb_ci.name	Exact match
Namespace	cmdb_ci.namespace	Exact match

Predefined tags extracted from connector payload

Name	Description	Default mapping from CMDB
t_ip_address	IP address	cmdb_ci.ip_address
t_application	Application	-
t_region	Region	-
t_location	Location	cmdb_ci.location
t_environment	Environment	cmdb_ci.environment
t_namespace	Namespace	cmdb_ci.namespace

Event Management tag based alert grouping definition form

The form for creating or modifying a tag based alert clustering definition displays detailed information about the definition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag based alert clustering definition form

Field	Description
Name	Name of the alert clustering definition. Definition names must be unique.
Active	Select to activate the definition. This option is selected by default.
Order	The order by which definitions are tested for incoming alerts. Those with lower Order values are tested first. When an alert matches one of the definitions' filters, it continues searching for more definitions. Default value = 1000
Domain	The domain in which the current record was created. Read-only.
Assignment group	Assignment group that works on the alert. If no assignment group is defined in the alert rule, then this alert rule is considered as a global rule. When the rules are running – first the global rules run and then the rules that belong to the assignment group of the alert.
Description	Enter an optional description of the alert clustering definition.
Filter	Set conditions that incoming alerts must meet to be measured by the alert clustering definition's tags. If the tags correspond to alerts that exist in the system and are within the **Clustering timeframe (minutes)** value, the incoming alerts join with the existing alerts to form an alert group. After configuring the filter, you can click the **Preview** button to view how many existing alerts match the filter's condition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag based alert clustering definition form (continued)

Field	Description

	Note:
	• Matching alerts are not automatically included together in an alert group. Alerts are grouped only if they have corresponding alert clustering tags.
	• Filter parameters are case sensitive by default. To disable case sensitivity, set the *sa_analytics.correlation_case_sensitive* parameter to **false**.
	• You can also configure alert fields to be excluded from the search, using the *sn_em_tbac.tag_excluded_alert_fields* property. By default, the following are excluded by this property:
	  ○ type
	  ○ event_class

Override group description	Default group descriptions begin with a "Group of alerts" prefix, followed by the description of the primary alert in the group. You may override this group description by selecting the **Override group description** check box. Then, in the **Custom description** field, type a description. This description is used as the description of the groups that are created by this alert clustering definition. **Note:** You cannot save the form if you left the Custom description field blank or with the default 'Group of alerts' text.
Clustering timeframe	The maximum time, in minutes, allowed between the first alert and subsequent alerts for them to be grouped together. For example, a value of 60 indicates that any alert generated within 60 minutes of the first alert in the group (the oldest alert) is included in the alert group. Any alert generated after this 60-minute window from the initial alert is not included in the group. Permitted values = 0-1440
Tag Based Alert Clustering Definitions Tags M2M	Select the alert clustering tags to be assigned to the alert clustering definition. Alerts that meet the criteria specified in the selected tags are included in the alert group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tag based alert clustering definition form (continued)

Field	Description
	The available options are the tags created on the **Tag Based Alert Clustering Tags** page.

List of predefined tag-based alert grouping definitions

A list of the predefined alert clustering definitions provided with the Tag Based Alert Clustering Engine#application.

Predefined alert clustering definitions

Name	Description	Order
Group alerts from the same Application	Group all alerts from the same application, created in the last 10 minutes. In new systems, this definition is activated by default.	9010
Group all alerts from the same IP address	Group all alerts from the same IP address, created in the last 10 minutes.	9020
Group all alerts from the same Namespace	Group all alerts from the same namespace, created in the last 10 minutes. In new systems, this definition is activated by default.	9030
Group all alerts from the same Subnet	Group all alerts from the same subnet, created in the last 10 minutes. In new systems, this definition is activated by default.	9040
Group alerts from the same CI class and Location	Group all alerts from the same CI class and location, created in the last 10 minutes.	9050
Group alerts from the same Application and Environment	Group all alerts from the same application and environment, created in the last 10 minutes.	9060
Group all alerts from a similar Node	Group all alerts from a similar node name, created in the last 10 minutes.	9070
Group alerts from the same Location and Assignment group	Group all alerts from the same location and assignment group, created in the last 10 minutes.	9080
Group alerts from the same Region and Metric	Group all alerts from the same region and metric, created in the last 10 minutes.	9090
Group alerts from the same CI class and Metric	Group all alerts from the same CI class and metric, created in the last 10 minutes.	9100
Group alerts from the same Node and Metric	Group all alerts from the same node and metric, created in the last 10 minutes.	9110
Group alerts from the same Assignment group and Class	Group all alerts from the same assignment group and class, created in the last 10 minutes.	9120
Group alerts from the same Type, Metric and Source	Group all alerts from the same type, metric, and source instance, created in the last 10 minutes.	9130
Group alerts from the same CI	Group all alerts from the same CI, created in the last 10 minutes.	9140

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predefined alert clustering definitions (continued)

Name	Description	Order
	Important: When this rule is active, CMDB grouping must be disabled.	
Group alerts from the same Node	Group all alerts from the same node, created in the last 10 minutes. In new systems, this rule is activated by default.	9150

Domain properties installed with Event Management

Use the domain properties installed with Event Management to provide metadata that identifies the appropriate domain table for event creation.

Default domain metadata properties

Property
evt_mgmt. connector_domain_info_table_name
evt_mgmt. connector_domain_info_column_name
evt_mgmt. connector_domain_id_column_name
evt_mgmt. connector_domain_path_column_name

Connector domain metadata

Override push connector default domain metadata values installed with Event Management with the values in the event HTTP request, including the URL, headers, request body (payload), and connector parameters. Override pull connector domain metadata values with connector parameters only.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Domain metadata parameters

Name	Description	Overrides System Property
connectorDomainInfoTableName	Name of the table containing the records with the domain information.	evt_mgmt.connector_domain_info_table_name • **Type**: string • **Default value**: core_company
connectorDomainInfoColumnName	Name of the column that has a value that should be matched with the domain identifier. For example, you might want to identify a record in the Domain Info table that has a matching value in the **company_name** field.	evt_mgmt.connector_domain_info_column_nam • **Type**: string • **Default value**: name
connectorDomainIdColumnName	Name of the table column that stores the domain ID.	evt_mgmt.connector_domain_id_column_name • **Type**: string • **Default value**: sys_domain
connectorDomainPathColumnName	Name of the table column that stores the domain path.	evt_mgmt.connector_domain_path_column_na • **Type**: string • **Default value**: sys_domain_path

Order of preference for accessing domain metadata

You must provide domain metadata to configure personalized push connector domains. The metadata in preferred accessing order are:

1. Push connector parameters
2. URL query parameters
3. HTTP headers
4. Request body (payload)
5. Default system properties

URL query parameters

Push connectors have the following URL syntax:

- Instance Type: `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=<connector-source>`
- MID Server type: `http://<MID_Server_IP>:<MID_Web_Server_Port>/api/mid/em/inbound_event?Transform=<connector-transform-script-name>`

You can include domain metadata as additional query parameters in the ServiceNow push connector webhook URL. For example, if the domain information source is the sys_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

table, the URL could be `https://<instance-name>.service-now.com/api/sn_em_connector/em/inbound_event?source=<connector-source>&connectorDomainInfoTableName=sys_user`

Request headers

Any of the domain metadata properties can be passed as part of HTTP request headers. Check the documentation from your third-party monitoring tool for information about how request headers are set.

Request body (payload)

You can reference domain metadata from the request body (JSON payload) by setting the metadata parameters as root-level keys in the payload. The following example shows how the *connectorDomainInfoTableName* and *connectorDomainInfoColumnName* parameters might be passed as part of the request body:

```
{
    “source”: “Azure Monitor”,
    “host”: “127.0.0.1”,
    “company”: “ServiceNow”,
    “connectorDomainInfoTableName”: “sys_user”,
    “connectorDomainInfoColumnName”: “first_name”
    }
```

Note: For information on connector parameters, see Configure domain separation for push connector events.

Dynatrace connector instance form

The Dynatrace connector instance form displays the fields you must fill in when creating a Dynatrace connector instance.

Connector instance page

Field	Description
Name	Unique name for the Dynatrace metric connector instance.
Description	Description for the use of connector instance.
Connector Definition	The connector definition to be used by the connector. Select **Dynatrace Metrics**.
Host IP	Hostname/IP address of the Dynatrace server.
Credential	Select the search icon and select the existing API key credential. If this credential does not exist, create a new credential, as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance page (continued)

Field	Description
	1. Ensure that you have an access token with Read Metric (metrics.read) permissions, created on the Dynatrace UI. 2. Select **New** after selecting the search icon. 3. Select **API Key Credentials**. 4. In the **API Key** field, enter the access token from the Dynatrace UI. 5. In the **Name** field, enter a descriptive name to uniquely identify the credential. 6. Fill in any other relevant fields, and select **Submit**.
Metric collection	Selected by default. Read-only.
Metric collection schedule (seconds)	Time interval (in seconds) at which the metric connector pulls metrics from the Dynatrace server. Default: 60.
Active	Select to activate the connector instance.

Dynatrace connector instance value parameters

The following table displays the Dynatrace connector instance value parameters that you can fill in, as needed, when creating a Dynatrace connector instance.

Connector instance value parameters

Parameter	Description
BaseApiPath	Custom base path for Dynatrace API calls. If your Dynatrace instance uses a different path than the default `/api/v2/metrics`, enter the custom path to ensure successful collection of metrics.
connectionAliasSysID	Connection credential alias sys ID of the Service graph connector you have configured. By default, the classic sys ID is populated. You can provide the sys ID of the Grail connection, if configured.
debug	Set to **true** to enable debug logs. Default: false
initialSyncInMins	Number of minutes before the current time that the initial metric collection runs. For example, if the value is 15, the connector fetches metrics from the past 15 minutes. Default: 15

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance value parameters (continued)

Parameter	Description
isGrailApiEnabled	Set to **true** to enable Grail connection. Grail API related metrics are displayed in the **Configure metrics to collect** list. If this value is set **false**, the classic connection is enabled and classic API related metrics are displayed. Default: false
logPayload	Determines whether to print the raw log payload. Set to **true** only for debugging purposes, as printing the raw payload quickly fills up the MID Server logs. Default: false
maxMetricsPerQuery	Maximum number of metric ids to query per API call to Dynatrace. The value cannot exceed 32, the maximum limit allowed by Dynatrace. Default: 32
pageSize	Maximum number of records to return in an API response from Dynatrace. Default: 500
port	Port number for the Dynatrace server to connect to. empty
protocol	Protocol to use when connecting to the Dynatrace server. Default: https
throttleApiCalls	Determines whether throttling is enabled. When enabled (**true**), a limit is set on the number of API calls to Dynatrace per minute. Default: false
requestPerMinute	The limit on number of API calls to Dynatrace per minute. Relevant only when *`throttleApiCalls`* = **true**. Default: 1000

Push connector instance form

Push Connector Instance form displays the fields that you must fill when you create or modify a connector.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	A unique name for the connector instance record, such as the name of an event source host.
Description	Any optional information that the administrator wants to use to identify this record.
Push Connector definition	Select the definition for the external system from which you want to receive events. You can select any existing definition, for example: • Thousandeyes • Dynatrace • Catchpoints For a list of the connector definitions available in the base system, see Configure Event Management connectors. **Note:** If there are available configuration parameters for the push connector type you select, they appear under Push Connector Instance Configuration Parameters.
MID Server	It shows the list of MID Servers that has at least one MID Server Event Listener context up and running. The MID Server section appears only when the MID type of definition is selected.
URL	This field is auto-populated only after saving the details of the connector instance. Instance: The URL has to be configured in an external system with the ServiceNow credential having a evt_mgmt_integration role to receive the events. MID push connector instance: The URL has to be configured from an external system with Configure the MID Web Server extension credential to receive the events. This field appears only after the form has been saved.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Active	Select this check box to enable receiving the events from this external event source. This field appears only after the form has been saved.

Here are some options that depend on the type of **Push Connector Definition** you selected:

- Datadog: For case or work items, while creating push connectors select the **Configure Work Item Alerts** button. You are redirected to the Datadog work Items pull connector instance form, where you configure related-alert fetch and bi-directional synchronization. For more information, see Configure a Datadog work item connector instance.
- Dynatrace: Specify the connection credential alias sys ID of the Service graph connector as the parameter value of the **connection_sys_id** parameter. For details on configuring this the Service graph connector, see Configure Service Graph Connector for Observability - Dynatrace using guided setup.

Related topics

Configure a push connector

Nagios connector instance value parameters

The following table displays the Nagios connector instance value parameters that you can fill in, as needed, when creating a Nagios connector instance.

Connector instance value parameters

Parameter	Description
debug	Set to **true** to enable debug logs. Default: false
max_fetch_interval_min	Limits metric history retrieval to prevent system overload. Default: 180 minutes
offset_min	Adds a small time buffer to ensure no metrics are missed. Default: 5 minutes

Nagios connector instance form

The Nagios connector instance form displays the fields you must fill in when creating a Nagios connector instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance page

Field	Description
Name	Unique name for the Nagios metric connector instance.
Description	Brief explanation of the purpose or use of the connector instance.
Connector Definition	Connector definition to be used by the connector. Select **Nagios Metrics**.
Metrics database host	Host IP address of the host where the Nagios server collects metrics.
Metrics database credentials	Credentials required to access the Nagios database.
Metric collection	Option is pre-selected and read-only.
Metric collection schedule (seconds)	Time interval (in seconds) at which the connector retrieves metrics from the Nagios server. Default: 60 seconds.
Active	Option to activate the connector instance.

Zabbix connector instance form

The Zabbix connector instance form displays the fields you must fill in when creating a Zabbix connector instance.

Connector instance page

Field	Description
Name	Unique name for the Zabbix metric connector instance.
Description	Brief explanation of the purpose or use of the connector instance.
Connector Definition	Connector definition to be used by the connector. Select **Zabbix Metrics**.
Host IP	Host IP address of the host where the Zabbix server collects metrics.
Credential	Credentials required to access the Zabbix database. Select the search icon and select the Basic Auth credential.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance page (continued)

Field	Description
	• Name: Provide the unique name for Zabbix cred. • UserName: Provide Zabbix server username • Password: Provide Zabbix server password • Select **Submit**.
Metric collection	Option is pre-selected and read-only.
Metric collection schedule (seconds)	Time interval (in seconds) at which the connector retrieves metrics from the Zabbix server. Default: 60 seconds.
Active	Option to activate the connector instance.

Zabbix connector instance value parameters

The following table displays the Zabbix connector instance value parameters that you can fill in, as needed, when creating a Zabbix connector instance.

Connector instance value parameters

Parameter	Description
api_endpoint_suffix	Set this value based on your Zabbix server web interface installation: https://www.zabbix.com/documentation/current/en/manual/installation/frontend#welcome-screen Default value: /zabbix/api_jsonrpc.php [For Apache Server]
debug	Set to **true** to enable debug logs. Default value: false
enable_batch_processing	If batch processing is enabled, metrics details are fetched in batch with 1000 (can be altered by setting appropriate value to - 'max_hosts_per_batch' parameter) hosts per batch. Default value: false
groupids	Provide the host group IDs. These IDs will be used to fetch metrics from the hosts attached to the specified groups. This applies only if `enable_batch_processing` is set to true.
logPayloadForDebug	Determines whether to print the raw log payload. Set to **true** only for debugging purposes, as printing the raw payload quickly fills up the MID Server logs. Default value: false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance value parameters (continued)

Parameter	Description
max_fetch_interval_min	Configured from the 'max_fetch_interval_min' parameter: If the time exceeds 'max_fetch_interval_min', fetching starts from 'max_fetch_interval_min' before the current time; for an INITIAL run without 'lastMetricSignature', the time is set to 'max_fetch_interval_min' before the current time. Default value: 180 minutes
max_hosts_per_batch	If batch processing is enabled, metrics details are fetched in batch with 1000 (can be altered by setting appropriate value to - 'max_hosts_per_batch' parameter) hosts per batch. Default value: 1000
offset_min	Specifies the time offset in minutes for adjusting the fetch start time relative to the current time. Default: 5 minutes
port	Port number for the Zabbix server to connect to. Default value: empty
protocol	Protocol to use when connecting to the Zabbix server. Default value: http
with_triggers	It will fetch the trigger information. Default value: false

Note: When monitoring a large number of hosts in Zabbix, fetching massive metric data can impact performance and cause API timeouts; enabling batch processing (enable_batch_processing = true) mitigates this by fetching metrics in batches (default: 1000 hosts, configurable via *max_hosts_per_batch*), though it increases network calls and must be used when monitoring 10,000+ hosts or facing performance issues.

SCOM connector instance form

The SCOM connector instance form displays the fields you must fill in when creating a SCOM connector instance.

Connector instance page

Field	Description
Name	Unique name for the SCOM connector instance.
Description	Description for the use of the SCOM event collection instance.
Active	Select to activate the connector instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance page (continued)

Field	Description
	Appears only when selecting a connector instance that has already been created.
Connector definition	The vendor and protocol used to gather events from the external event source. For a regular SCOM connector, select **SCOM**. For a connector instance that you want to monitor metrics, select **SCOM Metrics**. When selecting this option, the fields in the **Connector instance - metrics fields** table appear.
Assignment group	Select the group to whom the connector instance is assigned. For example, this may be a group of administrators or a support group.
Host IP	Specify the SCOM IP address. This field does not appear when **SCOM metrics** is selected as the Connector definition.
Credential	Select Windows credentials. **Note:** When configuring Windows credentials, if the credential applies to specific MID servers, ensure that the MID Servers you specify on the credential record match the MID Servers specified in the **MID Servers for Connectors** section on the Connector Instance form. This field does not appear when **SCOM metrics** is selected as the Connector definition.
Event collection last run time	The last event collection run time value updates automatically. This field does not appear when **SCOM metrics** is selected as the Connector definition.
Last event collection status	The last event collection run time status updates automatically. This field does not appear when **SCOM metrics** is selected as the Connector definition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance page (continued)

Field	Description
Event collection schedule (seconds)	The frequency in seconds that the system checks for new events from SCOM Operations. This field does not appear when **SCOM metrics** is selected as the Connector definition.
Bi-directional	Select to invoke the bi-directional option. This option enables the bi-directional exchange of values to-and-from the external event source. There is default implementation for SCOM. The **Last bi-directional status** option displays only when this option is selected. This field does not appear when **SCOM metrics** is selected as the Connector definition.
Last bi-directional status	The value of this field populates automatically. This field does not appear when **SCOM metrics** is selected as the Connector definition.
Last error message	The last error message field is automatically updated. This message is the last error message received by the connector. If the test connector fails, an error message is displayed in this field. This field does not appear when **SCOM metrics** is selected as the Connector definition.
MID Servers (MID Server for Connectors section)	Specify one or more MID Servers that are up and valid. MID Servers sort according to the order that their details were entered into the MID Server for Connectors section. The port requirement from the MID Server to the SCOM server is `1433`. If the specified MID Server is in a cluster, it may not be selected and another available MID Server is selected in its place. Ensure that the specified MID Servers match the MID Servers specified on the Windows Credentials page, as the connector uses Windows credentials to access MID Servers. If no MID Server is specified, an available MID Server that has a matching IP range is used.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance - metrics fields (displayed when Connector definition = SCOM Metrics)

Field	Description
Metrics collection	Selected by default. Read-only.
Metrics collection last run time	The last run time of the Metrics Collection scheduled job. This field is read-only.
Last metrics collection status	Status of the Metric Intelligence collection activity. The value of this field populates automatically. This field is read-only.
Metrics collection schedule (seconds)	The time, in seconds, to repeat the Metric Intelligence Collection scheduled job.
Metrics database host	The IP address or the host name of the metrics database host.
Connect using a named instance	When selected, the connection is made using the specified named instance. Otherwise, the connection is made using the specified port.
Metrics database port	The port used by the metrics database. The connection is made using JDBC. Default port number is 1433. If **Connect using a named instance** is selected, this option does not appear.
Metrics database named instance	The metrics database instance name.
Database login with Windows authentication	Perform database login with the credentials of the logged-in user defined on the MID Server service.
Metrics database credential	Credentials for the metric database. Use JDBC credentials for the local database user. If **Database login with Windows authentication** is selected, this option does not appear.

Solarwinds connector instance form

The Solarwinds connector instance form displays the fields you must fill in when creating a Solarwinds connector instance.

Connector instance page

Field	Description
Name	Unique name for the Solarwinds metric connector instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance page (continued)

Field	Description
Description	Brief explanation of the purpose or use of the connector instance.
Connector Definition	Connector definition to be used by the connector. Select **Solarwinds Metrics**.
Host IP	Hostname/IP address of the Solarwinds server.
Credential	Credentials for a SolarWinds user with Basic Authentication and read access to the SWIS REST API.
Metric collection	Option is pre-selected and read-only.
Metric collection schedule (seconds)	Time interval (in seconds) at which the connector retrieves metrics from the Solarwinds server. Default: 60 seconds.
Active	Option to activate the connector instance.

Solarwinds connector instance value parameters

The following table displays the Solarwinds connector instance value parameters that you can fill in, as needed, when creating a Solarwinds connector instance.

Connector instance value parameters

Parameter	Description
debug	Set to **true** to enable debug logs. Default value: false
max_fetch_interval_in_min	Number of minutes before the current time when the initial metric collection runs. For example, if the value is 180, the connector fetches metrics from the past 180 minutes. Default value: 180 minutes
logPayloadForDebug	Determines whether to print the raw log payload. Set to **true** only for debugging purposes, as printing the raw payload quickly fills up the MID Server logs. Default value: false
offset	Specifies the time offset in minutes for adjusting the fetch start time relative to the current time. Default: 0 minutes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connector instance value parameters (continued)

Parameter	Description
Port	Port number for the Solarwinds server to connect to. Default value: 17774
protocol	Protocol to use when connecting to the Solarwinds server. Default value: https

Event forwarding properties

Several system properties enable you to customize an event forwarding job.

Note: To access the System Properties [sys_properties] table, enter `sys_properties.list` in the navigation filter.

Property	Description
evt_mgmt.event_forwarding.sources_to_exclude	List of event sources from which events should not be forwarded. By default, all events are forwarded. Separate multiple sources by commas. • **Type**: string • **Default value**: Empty
evt_mgmt.event_forwarding.event_fields_to_include	List of event fields to include when sending events to the target instance besides the required default fields: resolution_state, event_class, additional_info, node, resource, cmdb_ci,severity, time_of_event, message_key, sys_domain, metric_name, ci_type, ci_identifier, type, source, description, and sys_domain_path. Separate multiple fields by commas. • **Type**: string • **Default value**: Empty
evt_mgmt.event_forwarding.event_batch_size	Maximum number of events to batch while sending to the target instance. • **Type**: integer • **Default value**: 500
evt_mgmt.event_forwarding.max_payload_size_in_mb	Maximum size of the event payload, in MB, sent to the target instance at a single time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	• **Type**: integer • **Default value**: 5
evt_mgmt.event_forwarding.max_events_to_forward_per_day	Maximum number of events that can be forwarded from the instance per day to minimize the performance impact on the source instance. • **Type**: integer • **Default value**: 100000
evt_mgmt.event_forwarding.log_debug	Enables or disables debug logging for event forwarding. Set this property to true to log debugging information. • **Type**: Boolean • **Default value**: False

Team-based integration properties

Team-based integration system properties enable you to customize the assignment group functionality for existing and new customers for event rules and event field mapping.

Team-based integration system properties

Property	Description
evt_mgmt.alert_auto_assignment_field	The property determines automatic alert assignm list of comma-separated fields. The default value of this property is cmdb_ci.sup IntegrationGroup. The assignment group value is **group** field (cmdb_ci.support_group) of the alert the value is taken from the IntegrationGroup valu **group** field in the Connector instances table [em For upgraded instances from previous family rel to avoid any impact on production instances. • **Type**: String • **Default value**: ○ Empty - for upgraded customers ○ cmdb_ci.support_group, IntegrationGroup - **Note:** When you use an alert field in this p alert field in the **Additional Info** field, an aler the value of the **Assignment group** field mu For example, `{"Alert field name":"` `desired sys_user_group record>`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Team-based integration system properties (continued)

Property	Description
evt_mgmt.enable_group_level_event_rules	Activates team-based functionality for executing field-mapping rules. Global rules take precedence based rules. • **Type**: Boolean • **Default value**: ○ False - for upgraded customers ○ True - for new customers

Alert Query form

You can combine similar alerts that meet specific criteria for a particular service by creating an alert query.

Alert Query form

Field	Description
Name	The alert query name.
Business criticality	The importance of this alert query to the business. This field can determine disaster recovery strategies for this alert query. • 1: most critical • 2: somewhat critical • 3: less critical • 4: not critical
Used for	The purpose of the alert query. Select from: • Production (Default) • Staging • QA • Test • Development • Demonstration • Training • Disaster recovery
Owned by	The party that is responsible for the alert query.
Email	The email alias for the alert query.
Business phone	The phone number for the business.
Operational status	The status of the alert query.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert Query form (continued)

Field	Description
	• **Operational**: The alert query is running. Impact calculation runs only on alert queries in the Operational status. Select this option to use this service. If you customized preconfigured options for this category, select the option whose value is set to 1. Event Management calculates impact only on services in use. If you do not want to use this service now, select one of these options: • **Non-Operational**: The alert query is unavailable, disabled, or down. • **Repair in progress**: The alert query is receiving maintenance. • **DR Standby**: The alert query is available as a disaster recovery (DR) standby. • **Ready**: The alert query is prepared for use. • **Retired**: The alert query is withdrawn from active use.
Comments	Any additional notes and information.
Service Group Members	The service group which contains this service.
Filter	The conditions for sending a notification to the alert query. The conditions must meet the alert values. **Note:** • When defining an alert query filter, include only fields that appear in the Alert Histories [em_alert_history] table. Impact calculation is based on Alert History data and fields such as Overall Event Count, Priority, and Priority group are not copied to the Alert Histories [em_alert_history] table. • Do not specify a dynamic time condition. For example, in the filter, do not specify `Created` condition of `Last 45 minutes` because impact calculation is triggered by a change of alert or alert query. However, for the dynamic time condition, none of these conditions have changed. • Some filters may slow down impact calculation. To solve this problem, adjust your alert query by adding an appropriate index, as described in Index suggestions for slow queries.

Alert correlation rule form

Manage the fields that define how alerts are correlated and grouped.

Alert correlation rule form

Field	Description
Name	Name of the correlation rule.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert correlation rule form (continued)

Field	Description
Order	The evaluation priority for the rule. Rules with lower numerical values are given higher priority. An alert is evaluated against each alert action rule until a match is found. For example, if you have two alert correlation rules with priorities 10 and 20, respectively, the rule with priority 10 will be evaluated first. If an alert matches the criteria of the rule with priority 10, no further rules will be checked. If it doesn't match, the alert will then be evaluated against the rule with priority 20.
Active	Option to activate or deactivate the rule.
Advanced	Option to switch to advanced mode, which lets you use custom scripts to define your own logic. The sample correlation rule, **Alert correlation rule SAMPLE**, is provided out-of-the-box for reference. You can use the available script as a guide. **Note:** The **Filter** condition specifies which alerts the rule will apply to. Ensure that the same condition is used in the advanced script to identify alerts to be included in the group.
Description	Description of the rule.
Primary Alert	The filter condition to identify the alert that is the primary alert, or most important alert, in a set of related alerts. This field does not appear when **Advanced** is selected.
Secondary Alert	The filter condition to identify the alert that is related to the primary alert, however it is of lesser importance. This field does not appear when **Advanced** is selected.
Filter	The filter condition to identify the alert on which the script is run. **Filter** is available only when **Advanced** is selected. Filter parameters are case sensitive by default. To disable case sensitivity, set the *sa_analytics.correlation_case_sensitive* parameter to **false**.
Relationship Type	Specify the type of relationship between the primary and secondary alert: • **No Relationship**: Ignore the relationship when looking for a match. • **Same CI or Node**: Relate both alerts with the same CI. If the CI field is blank, the alerts must have the same Node value. • **Primary is Parent**: Relate the alerts where the primary alert is the parent in a parent-child relationship, as described in the CI Relationship Types table [cmdb_rel_ci]. • **Primary is Child**: Relate the alerts where the primary alert is the child in a child-parent relationship, as described in the CI Relationships table [cmdb_rel_ci].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert correlation rule form (continued)

Field	Description
	This field does not appear when the **Advanced** check box is selected.
Time Difference in Minutes	The minutes between which the primary and secondary event must occur to match this rule. The default value is 60 minutes. **Note:** The value for this entry cannot exceed 1440 minutes (one day). This field does not appear when **Advanced** is selected.
Script	Custom script that you can modify to return a JSON string that specifies the primary and secondary alerts. Select **Advanced** to display the script field. `(/* The function needs to return a JSON-{correlationType:[correlatedAlerts]} for example: if your filter matches the alert, set the alert as the primary alert and set alerts 1, 2 and 3 each as secondary alerts.` `You can use both multiple primary alerts and multiple secondary alerts. The correlationType can be PRIMARY or SECONDARY, and the alerts ID must be in an array. CurrentAlert is the GlideRecord of the currentAlert on which that rule runs. The system supports only one primary per alert, so: Do not correlate more than one alert under the PRIMARY array. Do not correlate alerts that already have a primary under the SECONDARY array. The system supports open alerts only, so do not correlate alerts that have been closed under either one of the arrays. */`  `(function findCorrelatedAlerts(currentAlert){`  `var result = {}; //Insert your code here result = {'SECONDARY':['alertID1','alertID2','alertID3']};` `return JSON.stringify(result);`  `})(currentAlert);`
Relationship	Description of the CI relationship between primary and secondary, for example, *Allocated from::Allocated to* or *Allocated to::Allocated from*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert correlation rule form (continued)

Field	Description
	This field displays only if either **Primary is Parent** or **Primary is Child** is selected for the **Relationship Type**. 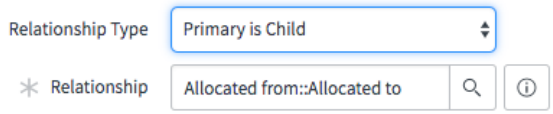 This field does not appear when **Advanced** is selected.

Health Log Analytics

The ServiceNow® Health Log Analytics application helps prevent IT issues before your users are affected. It helps you identify the root cause of an issue by enabling you to triage related logs and analyze the raw data.

Health Log Analytics collects raw log data streaming into your instance and structures textual data into key-value pairs. It detects meaningful patterns in the data, and when it discovers an anomaly, it sends an event to the Event Management application.

Health Log Analytics is available in the ServiceNow Store.

Request apps on the Store

Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Get started

Explore

Explore Health Log Analytics workflow, architecture, and high-level tasks.

Configure

Configure Health Log Analytics setup and configuration tasks.

Remediate

Investigate log data and take action to resolve the underlying issue.

Control

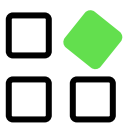

Influence how HLA generates and prioritizes alerts and detects anomalies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

System features and health alerts

View system features and core components with related system health alerts.

Analytics and reports

Gain insights into the number of issues predicted before users were affected and the money that was saved.

Reference

Find reference information about Health Log Analytics.

Health Log Analytics licensing

The ServiceNow AI Platform® uses a licensing method where your organization is billed for using Health Log Analytics, which is available as a separate license from ITOM AIOps. The ServiceNow Product Documentation doesn't provide information on prices, packaging, or other details determined by your organization customer contract. For general information about licensing and subscriptions, see ITOM/OT SU Licensing and subscriptions.

Dependencies

For a list of Health Log Analytics plugin dependencies, see Plugins or applications installed with ITOM AIOps.

Troubleshoot and get help

- Ask questions and explore other resources for in the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Health Log Analytics

ServiceNow Health Log Analytics (HLA) predicts IT issues before they affect your users by collecting, analyzing, and correlating machine-generated log data in real time. It discovers anomalies and alerts you to potential issues.

Health Log Analytics overview

Health Log Analytics typically receives and processes log data and sends events to ServiceNow Event Management. The application discovers anomalies as they happen. It helps you identify the root cause of an issue by enabling you to triage related logs and analyze the raw data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics can handle any kind of machine-generated textual log data. It can process application, infrastructure, and network logs, as well as other types of textual log data. Although a Configuration Management Database (CMDB) can be helpful to generate high-quality events and alerts, it is not necessary.

Note:

- Health Log Analytics supports only UTF-8 logs. It does not support binary logs.
- If you are sending logs in a language other than English, additional configuration may be required.

For a brief explanation of key terms and concepts used in HLA, see the Health Log Analytics terminology.

Health Log Analytics users

Health Log Analytics Users

User	Description	Role
Administrator	Configures the Health Log Analytics application to make it ready for use by Operators. Performs administration tasks to keep the system running efficiently.	evt_mgmt_admin, admin
Operator	Analyzes Log Analytics alerts and takes action to help resolve the underlying issue.	evt_mgmt_operator

Health Log Analytics workflow

Health Log Analytics collects and processes log data automatically. It structures the data logically for operators to analyze, and generates meaningful alerts and suggestions that display in Event Management.

The diagram shows the Health Log Analytics workflow from collecting the data through sending an event or alert to Event Management.

Health Log Analytics workflow

Ingestion

This layer connects your environment to Health Log Analytics. You can stream your logs directly from servers and endpoints or from log repositories. The optional guided setup helps you create data input connectors for common data sources, such as:

- Rsyslog
- Beats
- Splunk

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Elasticsearch
- MID Server
- TCP

Structuring

This layer deals with structuring log data and auto-mapping it to logical silos, called Components. Data structuring can be done automatically or manually.

The system auto-structures log data by extracting the following properties from incoming log messages: Message, Timestamp, Host, Severity, and External-IDs. It extracts explicit values, like "property-name" and "value is IP." and semantic ones such as length, number of English words, and variance.

Auto-mapping assigns log samples and metadata to the appropriate tags automatically. The system tries to map log lines by analyzing the source that streams the data. The mapping is based on agent hints and common transport header fields.

Enrichment

This layer handles identifying the variable parts of a log message.

Health Log Analytics workflow - Enrichment

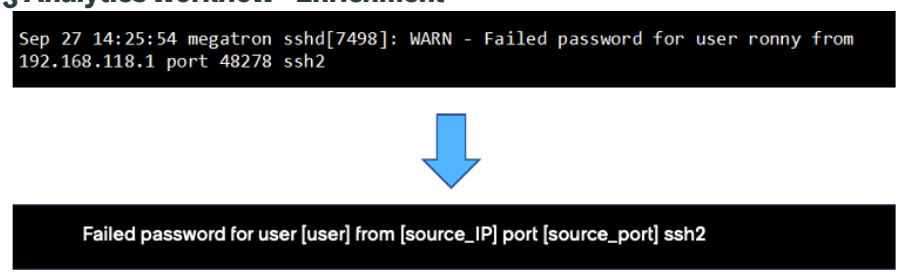

It also identifies keywords and contextual properties. In the image, "WARN" and "Failed" are the keywords to track. "User," "source IP," and "port" are the contextual properties.

Analysis

In this layer, each log line is indexed. Health Log Analytics extracts properties from the inner log message that contribute to models of behavior that the system learns to expect. Anomalous behavior departs from this expected behavior. You can search for an event and its most significant properties for manual triaging.

Machine Learning (ML) and Artificial Intelligence (AI)

Health Log Analytics uses advanced unsupervised machine-learning algorithms to discover patterns within logs and learn their unique data behavior. It then sets dynamic thresholds based on the data signature in real time to detect issues when they first occur. When the system detects a deviation from the typical pattern, it sends an event to Event Management.

Alert in Event Management

Health Log Analytics sends events to Event Management. In Event Management, Health Log Analytics alerts appear in the **All alerts** list. This list enables operators to see alerts from the event and the Health Log Analytics alert type in a single location.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics benefits

Health Log Analytics benefits

Benefit	Feature	User
Use predictiveLog Analytics alerts to handle emerging IT issues before they escalate and impact users.	Analyzing and resolving Log Analytics alerts	Operator
Set up log data connector integrations quickly and conveniently from the Integrations Launchpad.	Log data connector integrations	Administrator
Shorten onboarding time by installing content packs.	Content packs	Administrator
Save time and reduce errors by migrating data input configurations between instances.	Data input migration	Administrator
Identify the root cause of an alert by analyzing the logs that surround the anomaly.	Surrounding logs	Operator
Visualize anomalous log data.	Log viewer	Operator
Detect relationships in log data.	Log correlators	Operator
Assign higher or lower significance to alerts.	Mute alert metrics	Operator
Reduce noise by creating log filters.	Log alert filters	Operator
Influence how Health Log Analytics finds anomalies by managing keywords it looks for in the log data.	Lexical keywords	Operator
Create alerts for specified metrics by adding, changing, or deleting rules.	Custom alert rules	Operator

What to explore next

- Health Log Analytics architecture
- How Health Log Analytics generates alerts
- Types of anomalous behavior in Health Log Analytics
- Types of Health Log Analytics alerts

Health Log Analytics architecture

Health Log Analytics collects logs streaming into your ServiceNow instance from endpoints or data lakes, such as Splunk and Elasticsearch.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The ServiceNow instance receives the logs either via a MID Server connector instance or via MID-less Ingest, which supports integration with services such as Amazon Data Firehose. Health Log Analytics identifies and triages anomalies in your log data using unsupervised machine-learning (ML) models. It then groups the anomalies together and applies further algorithms to help identify the root cause of the issue.

Health Log Analytics scaled architecture

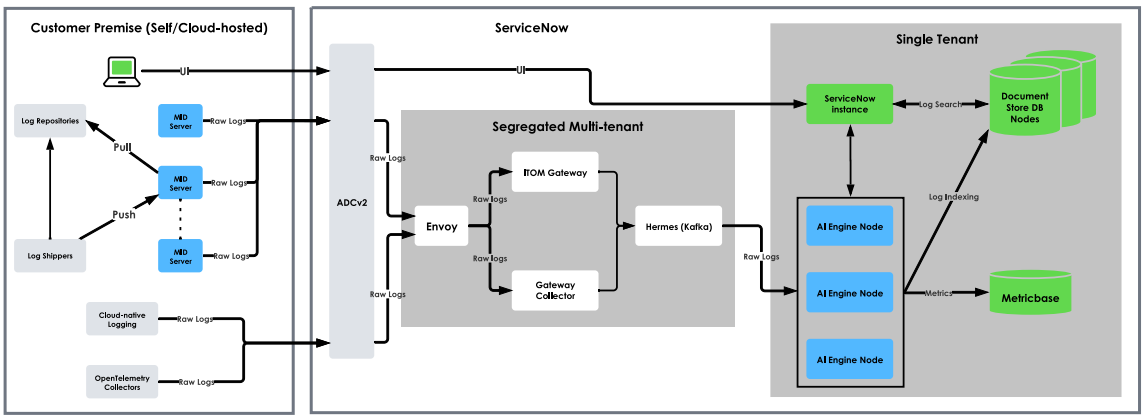

Health Log Analytics terminology

Before getting started with Health Log Analytics, it's important to familiarize yourself with some key concepts used in the application.

Health Log Analytics (HLA) terminology

Term	Description
Alert	A notification that HLA generates when it finds a statistically significant anomaly in metrics based on your system's log data, indicating a potential IT issue. HLA sends these alerts to Event Management, where you can see them in the **All Alerts** list.
Anomaly	Unusual or unexpected behavior in log data that differs from the baseline of normal behavior that HLA has learned from historical patterns.
Baseline	The statistical model of normal behavior of your system that HLA has learned from historical log data patterns. This baseline helps HLA identify anomalies.
Classification	A data type that determines how the HLA engine analyzes and processes a parsed log field. Available classifications are: Meter, Gauge, Histogram, Automatic Root Cause (ARC Only), Timeless Gauge, and Invalid.
Component	Representation of the smallest part of a service instance, typically a single micro service, module, or daemon. For example, a "checkout flow" service instance may include components such as cart-service, payment-service, and inventory-service. Every log or metric is assigned to a single component. This ensures that anomalies or incidents are attributed to that specific component.
Configuration Item (CI)	An individual item in your IT environment, such as a server, database, or application, that is tracked and managed in the Configuration Management Database (CMDB).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics (HLA) terminology (continued)

Term	Description
Configuration Management Database (CMDB)	The central database in ServiceNow that stores and manages information about all CIs and their relationships. This database provides HLA with the necessary context to correlate logs, anomalies, and alerts to specific services.
Correlation	The process of connecting related log events, anomalies, and alerts. By using factors like patterns and relationships defined in the CMDB service relationships, HLA identifies a shared root cause and reduces "noise."
Data input	A configurable connector that enables HLA to collect, transform, and ingest log data from external sources. **Note:** The term data input is used in HLA's classic interface (UI16) and back-end records. See also: Integration.
Data input mapping	The process of mapping raw log data to its specific log source, enabling HLA to connect logs to the corresponding service instance and component for analysis with full context. It also involves mapping your raw log data to source types so HLA can understand the format and structure of your logs. **Note:** The term data input is used in HLA's classic interface (UI16) and back-end records. See also: Integration. For a description of the simplified mapping interface in the new HLA UI, see: Log context mapping.
Detection algorithm	The statistical or machine-learning logic that HLA uses to analyze metrics and patterns to identify anomalous behavior. Examples include sigma-based detection, trend analysis, and comparing against a baseline.
Enrichment	The process of adding extra, contextual information to raw log data to make it more meaningful and actionable for analysis. Enrichment helps HLA connect technical events to their potential impact on your business.
Grouped alerts	A collection of related alerts that are combined into a single group to reduce noise and simplify triage.
Incident	A record created in ServiceNow when correlated alerts point to a significant disruption in a service, requiring investigation and resolution.
Integration	A configurable connector that establishes data pipelines so HLA can collect, transform, and ingest log data from external sources. Setup of integrations is done through the Integrations Launchpad, which significantly reduces implementation time compared to manual data input setup. **Note:** The term integration is used in HLA's new UI experience and front-end records. See also: Data input.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics (HLA) terminology (continued)

Term	Description
Integrations Launchpad	A framework that provides integration setup workflows for connecting external log data sources to HLA. Setting up integrations through the Integrations Launchpad significantly reduces implementation time compared to manual data input setup.
Label	A semantic identifier assigned to common log properties, such as MESSAGE, HOST, TIMESTAMP, SEVERITY, and EVENT-ID. Labels tell HLA what role a parsed log field plays in the log structure.
Lexical keywords	Specific words found in log data, such as "crashed" or "failed," that can point to important issues. HLA tracks lexical keywords to detect widespread issues and long-forming trends. For more information, see: Add, edit, or delete lexical keywords in Health Log Analytics.
Log context mapping	The process of mapping your raw log data to the appropriate log source, enabling HLA to connect logs to their corresponding service instance and component for contextualized analysis. The user-friendly Log context mapping interface in the new HLA UI focuses only on the service context. It's simpler to use than Data input mapping in HLA's classic interface (UI16), which also includes source type configuration. See also: Data input mapping.
Log ingestion	The process of streaming of logs from servers and endpoints or log repositories to HLA, using data input connectors or integrations.
Log property	A structured data element, such as timestamp, severity, or error code, that is extracted from a log entry.
Log source	A logical representation of the source of the log data that HLA ingests. Each log source is defined by a service instance-component pair. HLA performs anomaly detection and generates alerts on issues within the scope of each individual log source.
Log Viewer	An interface within HLA that allows you to search, filter, and examine raw or parsed log data from multiple sources. The Log Viewer helps you investigate events and better understand your system's behavior.
Metric	A quantifiable measurement extracted from log data in HLA. HLA supports various generic metric types, such as METER (event count), GAUGE (numeric value), and HISTOGRAM (distribution bucket). It uses the extracted log data to automatically create message-pattern-based metrics (METER), severity metrics (METER), keyword metrics (METER), and raw log metrics (METER, GAUGE, and HISTOGRAM). Each metric is associated with a specific log source and monitored for anomalous behavior.
MID Server	An intermediary agent that facilitates secure communication between on-premise or private cloud log sources and the ServiceNow instance. It enables HLA to collect log data securely from environments that the instance can't access directly.
Parsing	The process of extracting meaningful, structured data fields from raw log entries. HLA uses its AutoParser and AutoExtraction capabilities to

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics (HLA) terminology (continued)

Term	Description
	parse the logs, which then enables effective analysis, correlation, and visualization of the data.
Pattern	A recurring format found in log messages (free text fields within log records). HLA groups similar log messages that share this format into a pattern. It then monitors how often that pattern occurs to establish normal system behavior and detect anomalies. **Note:** HLA only performs pattern monitoring and anomaly detection on log properties labeled as MESSAGE.
Raw log data	The original, unprocessed log entries collected from external sources, which HLA ingests before parsing and analysis.
Root Cause Analysis (RCA)	An automated analysis that identifies meaningful factors, entities, and highlights that contributed to an alert or incident, helping to identify the underlying cause.
Service instance (formerly application service)	A logical grouping of Configuration Items (CIs), such as servers, databases, and applications, that collectively deliver a complete business service. In HLA, log data is correlated to service instances to provide a comprehensive, real-time view of their operational health and performance. **Note:** A single service instance can consist of multiple CIs, whereas the same CIs might be part of multiple service instances.
Severity	A classification in HLA that indicates the level of impact or urgency of an alert. This helps operators prioritize which issues to investigate and resolve first.
Source type	A configuration profile for a log data type that determines how HLA interprets, parses, and extracts fields from logs of that format and prepares the data for processing. Think of it as telling HLA: "Here's a specific type of log, and here's how you should process it."
Source type structure	A template within a Source Type that defines the fields and their arrangement in logs of that specific type. The Source type structure ensures that HLA can consistently extract and understand the incoming data. Think of it as telling HLA: "This part of this log is the timestamp." "This part is the severity." "This part is the error code," and so on.
Tag	Metadata applied to logs, metrics, or alerts. Tags are used for categorization, filtering, and correlation.
Time series	A sequence of metric values recorded in chronological order over a period of time. This data is used to identify trends, seasonality, and anomalies.
Variable extraction	The process of identifying and extracting dynamic, user-defined fields from log messages to enable custom metrics and analytics in HLA.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How Health Log Analytics generates alerts

Health Log Analytics identifies patterns in your log data and learns pattern behavior. When HLA's AI engine detects anomalous behavior, it sends an event to the ServiceNow Event Management application. As an operator, you can use these predictive alerts to handle emerging IT issues before they impact users.

Log anomaly detection

Anomalies are abnormal or unexpected behavior that occur when activities deviate from established baselines. There are many kinds of anomalies. In this example, the system tracks the baseline rate (the average number of events per minute) for a specific log pattern. When this typically inactive log generates a spike in events, the system detects the deviation from the baseline and generates an alert.

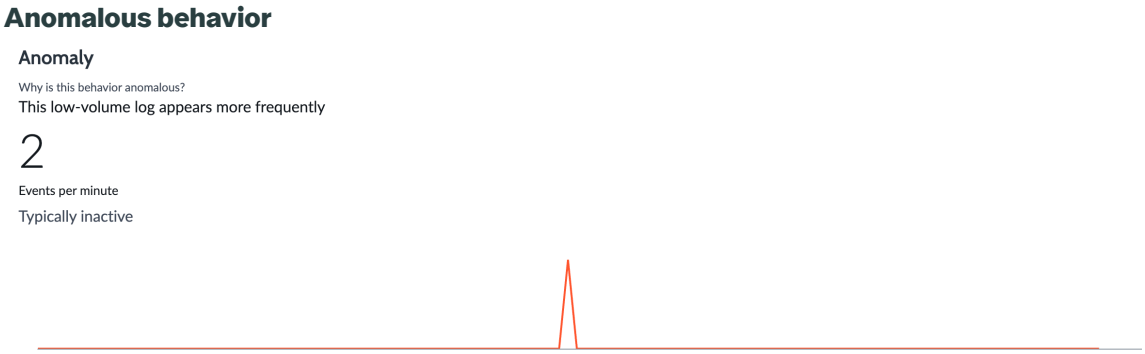

Health Log Analytics uses various methods to detect anomalies and generate alerts.

Alert metrics

Health Log Analytics monitors multiple metrics in the log stream to detect anomalous behavior. Each metric is associated with a unique source: the combination of service instance and component. When the system identifies an anomalous pattern for a metric, it generates an alert.

As an operator, you can provide feedback about the generated alerts. Your feedback "teaches" Health Log Analytics that a specific alert is significant or irrelevant to you. The application then either raises the priority of the alert metric or mutes it to reduce noise.

- A significant alert is more likely to be included in a Log Analytics group when the associated metric behaves anomalously. For more information, see Mark an alert as significant in Health Log Analytics
- Mute an alert for a specified source to eliminate distracting new alerts for unimportant issues. When a metric is muted, Health Log Analytics removes the current alert and any other alerts based on that metric from the feed. It also stops generating new alerts from that metric. For more information, see Mute an unimportant alert in Health Log Analytics.
- When the situation changes, you can return a significant metric to its default significance. You can also reactivate a muted metric to cause the system to start generating alerts again. For more information, see Restore normal importance to an alert metric in Health Log Analytics.

Lexical keywords

Health Log Analytics scans your logs for words that can indicate important issues. Lexical keywords such as "crashed" or "failed" signal a condition that can merit attention.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The system sets a threshold for each lexical keyword that is based on what it considers the normal occurrence pattern and frequency of that keyword in your logs. When it scans your logs, it finds all occurrences of the keyword. If the number exceeds the threshold, it generates an alert. For more information, see View the lexical keywords that generate alerts in Health Log Analytics.

For information about managing global keywords, see Add, edit, or delete lexical keywords in Health Log Analytics. To create or delete keywords for a specific source type, see Configure source type capabilities in Health Log Analytics.

Correlations

Log correlators are keys or values in log data that detect correlations between alerts. For example, a log correlator could detect when the ID of a particular network device occurs simultaneously in multiple warnings across different service instances. For more information, see Log correlators for identifying related alerts.

Advanced alert filtering

Add advanced log alert filters to scan alerts for conditions that you specify. The filters reduce noise by dropping alerts that do not indicate a significant issue. While developing a filter, you can test, update, publish, or activate the filter at any time. For more information, see Create advanced log alert filters.

Custom alert rules

Define a Log Analytics alert rule when you encounter log data that should generate an alert. The alert rule generates an alert for a specified metric with a threshold that you specify and sets the properties of the generated alert. For more information, see Alert rules in Health Log Analytics.

Types of anomalous behavior in Health Log Analytics

Anomalous behavior in a CI or a service can indicate an important issue. For example, a spike in the frequency or number of messages of a particular type can indicate a problem.

To build models of expected behavior, Health Log Analytics monitors the log stream to learn baselines for patterns, metrics, and gauges over various time periods. Time periods can be hourly, daily, weekly, or unlimited. Behavior that departs from the learned models is considered anomalous behavior.

Types of log property

Pattern

A pattern is a value or rate that repeats, whether in text, time, or relationships.

Meter

A meter property is a numeric or text value. For example, a status code, a response code, an action, or a pattern.

Gauge

A gauge property has a numerical value that is reported continuously. Gauge properties represent operations that consume resources. For example, CPU usage, memory usage, or response time.

How anomalies appear in Health Log Analytics

The Anomaly card illustrates the anomalous activity that led to the alert. The chart shows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Recent anomalous activity
- Expected behavior (the learned baseline)
- Baseline values from one day earlier
- Baseline values from the previous week

In this example, the system tracks the baseline rate (the average number of events per minute) for a specific log pattern. When this typically inactive log generates a spike in events, the system detects the deviation from the baseline and generates an alert.

Anomaly card

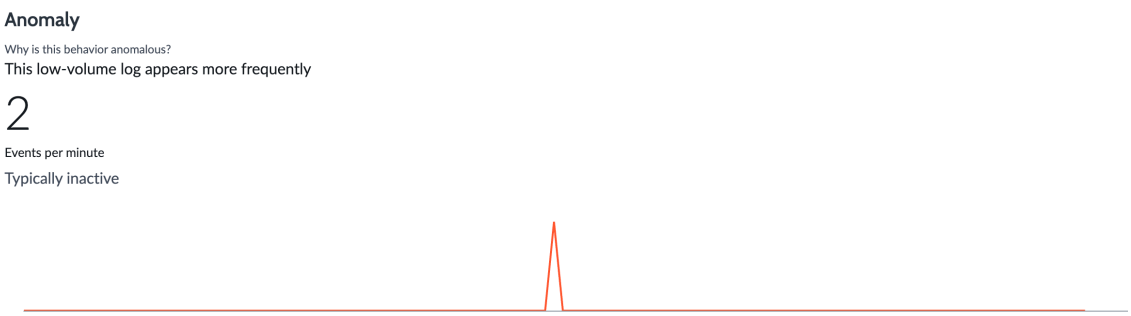

Kinds of anomalies

Some of the kinds of anomalies

Behavior	Description
New behavior	A pattern that has not ever been seen. The New Behavior alert type does not display a chart.
Signal dead/Stopped appearing	All pattern or log data from a source has stopped. There has been no signal for at least five minutes.
Signal alive/Appearing again	A pattern or log data from a "dead" source is appearing again. For a baseline of one hour, a pattern is "dead" if it appears less than once per minute.
Anomaly above average or below average	Activity that deviates from expected baseline behavior for pattern or meter or gauge metrics, such as keywords metrics or severity metrics.
Baseline reference increase or decrease	An increase or decrease in the value or volume of a log property as compared to the one-hour or one-week baseline.
Correlation of severity and keyword alerts	An increase in the volume of a severity level or keyword.

Types of Health Log Analytics alerts

Health Log Analytics generates several types of alerts.

In the list of alerts, all alerts that are generated by Health Log Analytics have the value **Log Analytics** in the **Source** column.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The value in the **Group** column identifies the type of alert as one of the following:

Component-based alert

Component-based alerts involve more than one configuration item (CI). A component is a logical component of a service instance that caused the alert. A component can be multiple CIs that perform the same function, such as multiple redundant hosts.

Important: Each Component-based alert is the parent of a number of read-only alerts. You do not work directly on read-only alerts. You work only on the parent Component-based alert.

In this example service instance, the identical Java apps X, Y, and Z make up a single component: Component B. Tomcat servers Q, R, and S and their hosts make up a different single component: Component C.

Example service instance

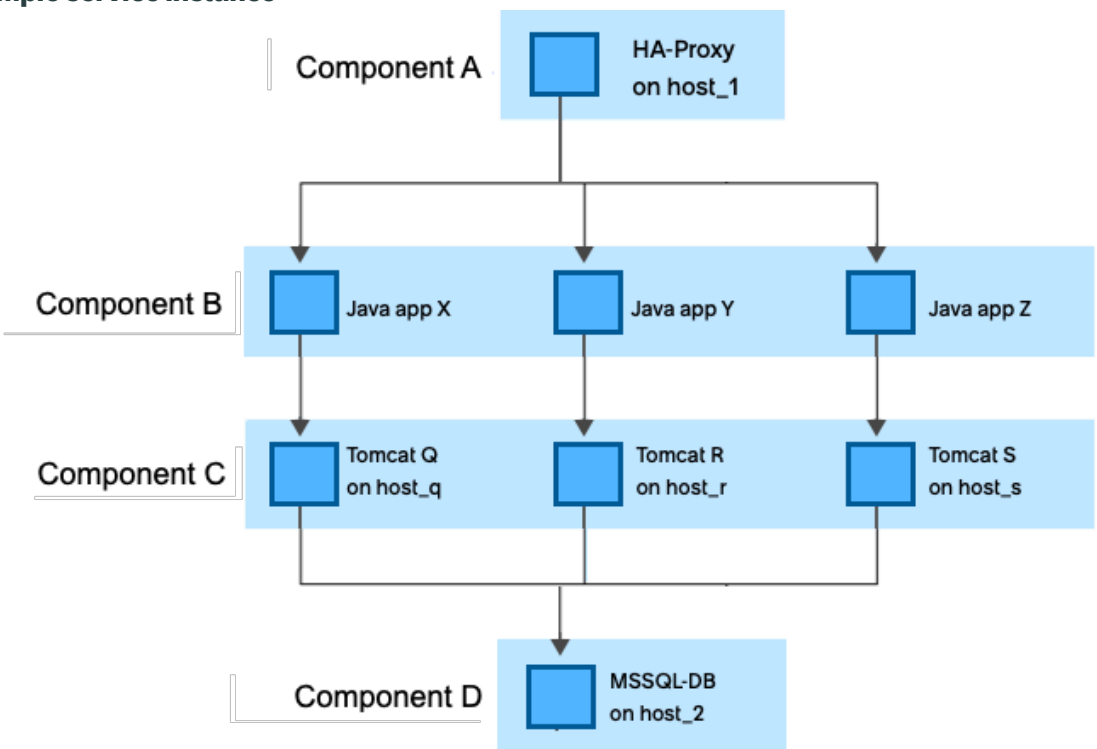

Log Analytics alert

A Log Analytics alert identifies an anomaly that involves a single CI. A Log Analytics alert has the value **None** in the **Group** column. The anomaly that leads to the alert can be an unexpected number of log entries or an unexpected value of a metric.

Log Analytics group

When the system identifies multiple Log Analytics alerts that are related in important ways, it groups them into a Log Analytics group. A Log Analytics group can group up to four alerts. The system generates a Log Analytics group when the Log Analytics alerts share one or more of the following relationships:

- Time: The events all occurred within a configured time interval.
- Metadata: The alerts have matching values in log-line metadata. For example, all alerts involve the same host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Message text: The message text in the log data is similar or identical between alerts.
- Trend: The alerts show a similar tendency in values or rates. For example, a particular metric value is increasing in all alerts.

Note: You can mark an alert as significant. A significant alert is more likely to be included in a Log Analytics group when the associated metric behaves anomalously. For more information, see Mark an alert as significant in Health Log Analytics.

Configuring Health Log Analytics

As an Administrator, set up and configure Health Log Analytics for Operators to use, and carry out administration tasks to ensure that the system runs efficiently.

Configuration overview

Complete the following tasks for a successful setup and configuration of Health Log Analytics.

Important: Regardless of how you implement Health Log Analytics, for log data ingestion via a MID Server you must first configure MID Server system requirements and ensure that the MID Server log ingestion capability is enabled.

For MID Server proxy requirements, see MID Server proxy preconditions for streaming logs to Health Log Analytics.

Setup and configuration tasks:

1. Install Health Log Analytics (HLA).

   Install the Health Log Analytics application.

2. Set up log data inputs for streaming raw log messages to your ServiceNow instance.

   Use any of the following methods:

   - Set up integrations for Health Log Analytics from the Integrations Launchpad.

     You're encouraged to set up data input integrations from the Integrations Launchpad, which provides a unified interface for convenient integration.

   - Set up data inputs in Health Log Analytics using guided setup.
   - Set up data inputs in Health Log Analytics manually.

3. Identify and resolve a log streaming issue in Health Log Analytics.

   Identify and address log streaming issues to ensure that the data inputs you have configured are streaming data properly to your instance.

4. Edit raw log data before processing in Health Log Analytics.

   Modify raw log messages before they are processed in the MID Server, and therefore before Health Log Analytics maps and structures it.

5. Map your logs to service instances and components for log context.

   Map your logs to service instances and components so that Health Log Analytics can generate alerts in context.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you turn off automatic mapping, you must define a mapping script manually in the **Data Input Mapping** page. For more information, see Map raw log data manually.

6. Refine the source type structure in Health Log Analytics.

   Fine-tune how Health Log Analytics reads your inner log messages and detects anomalies by customizing the extracted properties in the source type structure.

   Health Log Analytics enables you to reclassify auto-classified properties and change auto-mapped labels. These adjustments help the system's machine learning better understand your priorities.

7. Perform additional data input setup tasks.

   After performing the mandatory data input setup and configuration, you can continue with optional setup tasks.

8. Perform advanced data input configuration tasks.

For a brief explanation of key terms and concepts used in HLA, see the Health Log Analytics terminology.

Install Health Log Analytics (HLA)

Install Health Log Analytics by requesting ServiceNow HLA installation from ServiceNow Customer Support.

Before you begin

Verify that the Health Log Analytics application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. To obtain the entitlements, contact your ServiceNow Solution Consultant.

The minimum version of the ServiceNow AI Platform required for Health Log Analytics installation is Brazil.

Important: Health Log Analytics uses the default user sn_occ_occultus for internal communication. Make sure that this user is active and do not deactivate it.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. On the ServiceNow Store website, opt into the ITOM AIOps product.

   a. Go to the ServiceNow Store website and search for the Health Log Analytics app.

Health Log Analytics app in the ServiceNow Store

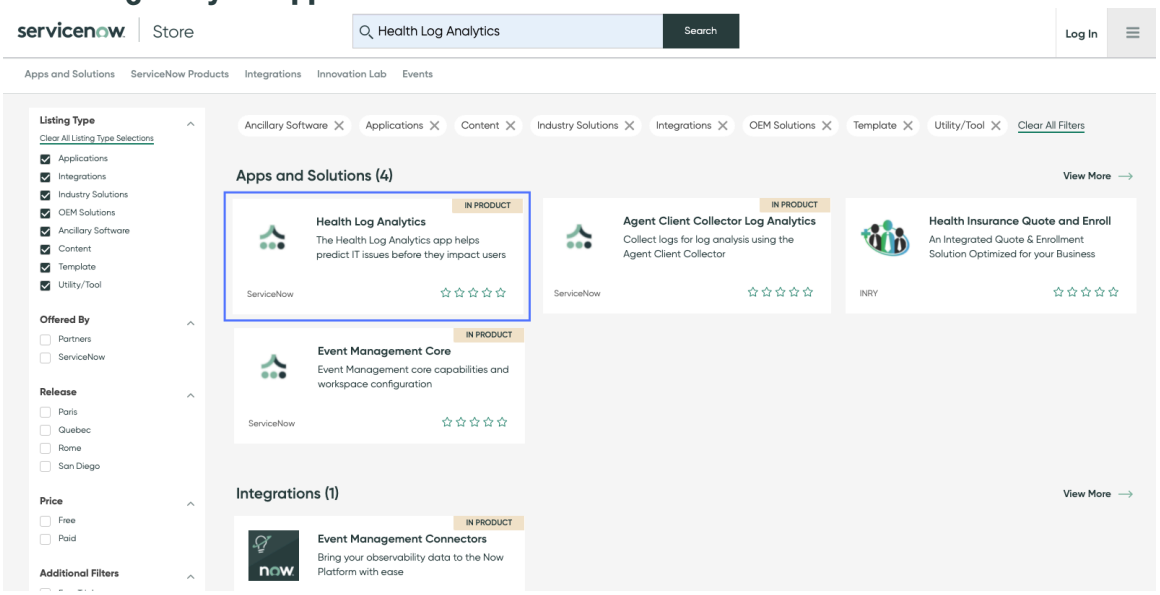

   b. Select the Health Log Analytics tile.
   The Health Log Analytics product page displays.

Health Log Analytics product page in the ServiceNow Store

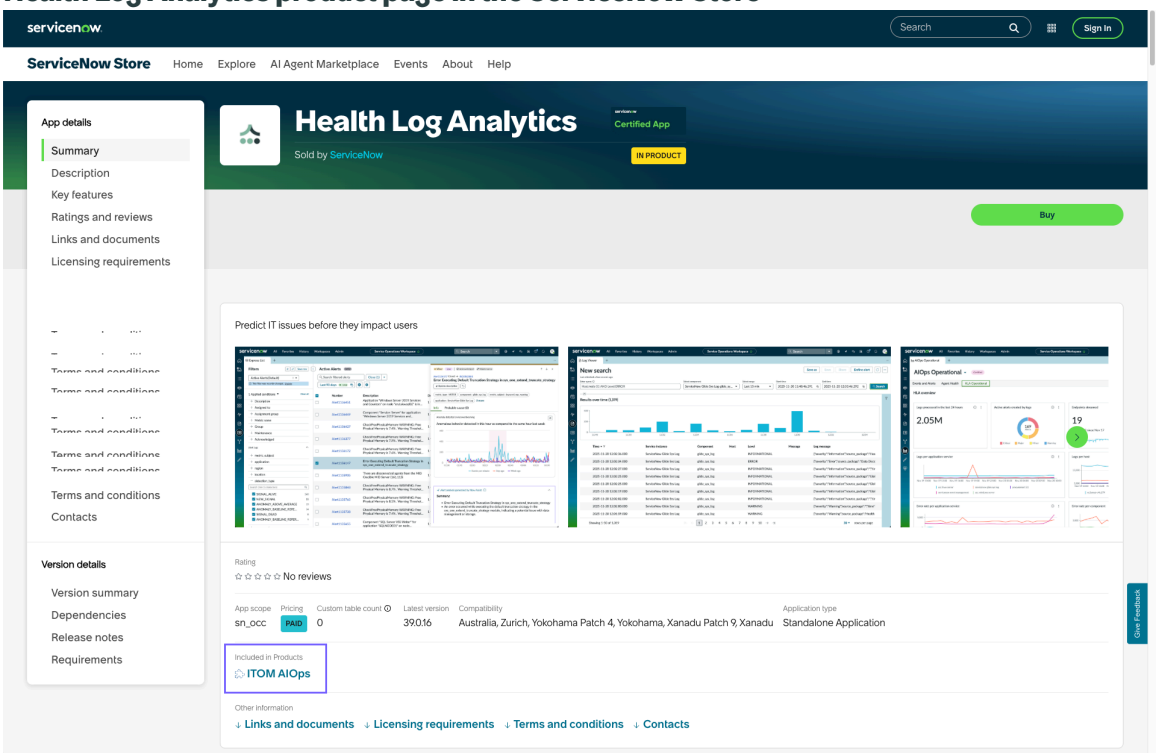

   c. Under **Included in Products**, select ITOM AIOps.

   d. Select **Opt-in** to opt-in to ITOM AIOps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Install the Health Log Analytics application on your ServiceNow instance.

   a. On your instance, navigate to **System Applications** > **All Available Applications** > **All**.

   b. Find the Health Log Analytics application using the search bar.

   c. On the Health Log Analytics tile, select **Install**.
   When the application is successfully installed on the instance, an automatic change request (CHG) is triggered to request provisioning for the necessary endpoints. The instance admin receives the request and any relevant notifications.

   > **Note:** Even when the Health Log Analytics application is installed successfully, it will not work until the endpoints have been provisioned. Provisioning endpoints configures your source systems to collect and send their log data to your ServiceNow instance. Health Log Analytics depends on this ingested data to detect anomalies, correlate events, and generate alerts.

   d. On the **Installed** tab, verify that the application and its dependencies were installed successfully.

Health Log Analytics is installed on the instance

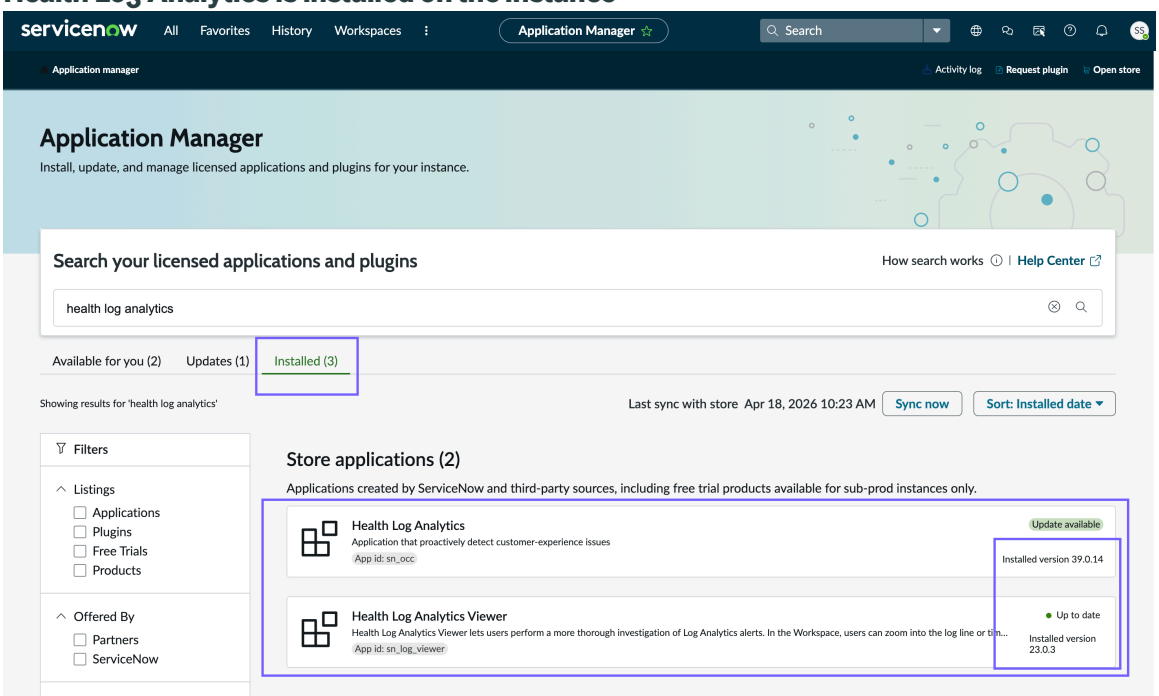

3. **Optional:** Observe the status of the provisioning process by reviewing the CHG.

   > **Note:** Some steps in the provisioning process might require your explicit approval. In such cases, ServiceNow Customer Support will reach out.

4. When provisioning is completed, verify that the Health Log Analytics application is up and running on your ServiceNow instance.

   a. On the instance, enter `sn_occ_stats.do` in the navigator and press the **Enter** keyboard key.
   The Health Log Analytics Package Dependencies and Versions table displays.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics Package Dependencies and Versions table

Health Log Analytics Package Dependencies and Versions

Dependency	Version
Health Log Analytics	39.0.14
Service Operations Workspace Log Analytics	26.5.0
Alert Intelligence	23.13.0
---	---
Occultus Version	1.36.0.27
MetricBase Version	25.2.0.6
ElasticSearch Version	8.17.3
---	---

Occultus

Node: svc241172.aus101.service-now.com ,Status: Active ,Version: 1.36.0.27

The provisioned endpoints are listed along with their version number:

- Occultus
- MetricBase
- ElasticSearch

The Health Log Analytics application is up and running on your instance.

Result
The Health Log Analytics application is installed and provisioned on your instance.

What to do next

- Set up the data input process for Health Log Analytics
- (Optional) Install the ServiceNow® Agent Client Collector Log Analytics (ACC-L) plugin.

  Agent Client Collector Log Analytics enables you to stream log data from Linux and Windows hosts to a ServiceNow instance using the Agent Client Collector. For more information, see the Agent Client Collector Log Analytics documentation.

Set up Health Log Analytics on your ServiceNow instance

Implement Health Log Analytics on your ServiceNow instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

- Install Health Log Analytics on your instance before implementation. For more information, see Install Health Log Analytics (HLA).
- Regardless of how you implement Health Log Analytics, you must configure MID Server system requirements and ensure that the MID Server log ingestion capability is enabled.

  For MID Server proxy requirements, see MID Server proxy preconditions for streaming logs to Health Log Analytics.

Setting up data inputs

You're encouraged to integrate data inputs with Health Log Analytics from the Integrations Launchpad. This tool provides a unified interface for convenient integration with connectors that feed raw log messages from external sources into your instance for processing and analysis. For more information, see Set up integrations for Health Log Analytics from the Integrations Launchpad.

Alternatively, you can set up data inputs in one of the following ways:

- Using guided setup. The guided setup provides a sequence of tasks to help you configure data inputs on your ServiceNow instance. Using the guided setup ensures that you have the minimum required setup for the data input process. For more information, see Set up data inputs in Health Log Analytics using guided setup.
- Manually. For more information, see Set up data inputs in Health Log Analytics manually.

For supported data inputs, see Supported data inputs for Health Log Analytics.

Set up integrations for Health Log Analytics from the Integrations Launchpad

Set up integrations from the Event Management Integrations Launchpad in Service Operations Workspace for ITOM.

Integrations Launchpad

The Integrations Launchpad tool provides a unified interface for convenient integration with connectors that feed raw log messages from external sources into your ServiceNow instance for processing and analysis. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Integrations for Health Log Analytics

The Integrations Launchpad enables the following integrations for Health Log Analytics:

Pull integrations

These integrations pull log data from external data sources and stream the data to your instance, typically via a MID Server. Select an integration in the table to open a page with the setup procedure.

Pull integrations

Integration	Description
Amazon CloudWatch	Streams log data from Amazon CloudWatch to your instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pull integrations (continued)

Integration	Description
Amazon S3	Streams log data from Amazon S3 (Simple Storage Service) buckets to your instance.
Apache Kafka	Streams log data from Apache Kafka to your instance.
Elasticsearch	Streams log data from Elasticsearch indices to your instance.
Microsoft Azure Event Hubs	Streams events from Microsoft Azure Event Hubs to your instance.
Microsoft Azure Log Analytics	Streams log data from Microsoft Azure Log Analytics to your instance. The connector points the Health Log Analytics AI engine to a data source in your Microsoft Azure Log Analytics account.
MID Server	Collects log messages from the MID Server and streams them to your instance.
ServiceNow System Logs Retriever	Sends log data from the ServiceNow System Log table to the Health Log Analytics AI engine. This integration doesn't run on a MID Server. **Note:** Only a single ServiceNow System Logs Retriever data input can exist in the system, and only users with the admin role can create and configure it.
Splunk Poller	Pulls log data from Splunk to your ServiceNow instance periodically by query.

Push integrations

These integrations connect to external data sources that push log data to your instance, typically via a MID Server. Select an integration in the table to open a page with the setup procedure.

Push integrations

Integration	Description
ACC Log Analytics	The ACC-L agent is installed on the monitored host and sends raw log data to your instance via the MID Server.
Cribl	Enables Health Log Analytics to process Cribl log messages streaming into the ServiceNow instance.
Edge Delta REST	Enables Health Log Analytics to process logs it receives from Edge Delta in a distinct format. These logs stream into the ServiceNow instance via REST.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Push integrations (continued)

Integration	Description
Edge Delta TCP	Enables Health Log Analytics to process logs it receives from Edge Delta in a distinct format. These logs stream into the ServiceNow instance over the TCP transport protocol.
GCP PubSub	Receives log messages that were published to a Google Cloud Pub/Sub topic and streams them to your instance.
REST API	Streams log data to your instance in JSON format.
Splunk TCP	Streams log messages to your ServiceNow instance over the TCP transport protocol using a Splunk heavy forwarder.
Splunk UDP	Streams log messages to your ServiceNow instance over the UDP transport protocol using a Splunk heavy forwarder.
TCP	Sends raw log messages to your instance directly over a TCP/SSL socket.
UDP	Sends raw log messages to your instance directly over a UDP socket.
Vector Agent	Enables Health Log Analytics to process log messages that are streaming into the ServiceNow instance via a Vector Agent.

MID-less integrations

These integrations stream log data directly to your ServiceNow instance, without a MID Server.

MID-less integrations

Integration	Description
AWS Firehose	Streams log messages from Amazon Data Firehose directly to the collector service in ITOM Gateway, where it's queued for Health Log Analytics processing.
Cribl Stream	Streams Cribl log data directly to your ServiceNow instance.
Microsoft Azure Event Hubs (MID-less)	Streams events from Microsoft Azure Event Hubs to your instance.
OpenTelemetry Collector	Streams log data directly to your ServiceNow instance using the OpenTelemetry (OTLP) protocol.
Splunk OpenTelemetry Collector	Streams Splunk log data directly to your ServiceNow instance using the OpenTelemetry (OTLP) protocol.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MID-less integrations for Health Log Analytics

Health Log Analytics (HLA) supports integrations that stream log data directly to your ServiceNow instance without a MID Server. Use these integrations to simplify your deployment and reduce infrastructure overhead.

Role of the MID Server in Health Log Analytics

A Management, Instrumentation, and Discovery (MID) Server is a Java application that runs on a server in your network. In a standard HLA deployment, the MID Server acts as a relay between your log data sources and your ServiceNow instance. It collects, filters, and forwards log data on behalf of the instance.

While MID Servers provide flexibility and support a wide range of integrations, they introduce additional infrastructure to deploy, maintain, and monitor. Organizations with high log volumes may also have to scale MID Server capacity over time.

How MID-less integrations work

MID-less integrations bypass the MID Server entirely. Log data streams from your third-party source directly to ServiceNow over HTTPS, where it is queued and processed by HLA.

This architecture relies on two components:

ITOM Gateway

A lightweight, cloud-native component that receives incoming log data and routes it to your ServiceNow instance. You deploy ITOM Gateway based on your expected log volume. For more information, including deployment options and prerequisites, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

JSON Web Token (JWT)

A secure, time-limited token used to authenticate the data stream between your log source and the ServiceNow datacenter. JWT authentication removes the requirement to store credentials such as AWS keys directly on your ServiceNow instance. For more information, see Configure a JWT provider and token.

Supported MID-less integrations

The following integrations support MID-less log streaming to Health Log Analytics.

- AWS Firehose
- Microsoft Azure Event Hubs (MID-less)

MID-less integrations using the OpenTelemetry (OTLP) protocol:

- Cribl Stream
- OpenTelemetry Collector
- Splunk OpenTelemetry Collector

Note:

- OTLP uses a standardized, vendor-neutral protocol to send application and system logs to HLA. It enables logs to be easily collected, processed, and correlated regardless of their origin or destination.
- The MID-less OpenTelemetry integrations all use the same transport model and share a common setup procedure in the Integrations Launchpad.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When to use MID-less integrations

Consider MID-less integrations when your organization wants to:

- Reduce infrastructure by eliminating the requirement to deploy and maintain a MID Server
- Simplify onboarding with a guided setup in the Integrations Launchpad
- Stream log data from cloud-native sources, such as Amazon Data Firehose or Microsoft Azure Event Hubs, where a MID Server intermediary adds unnecessary complexity
- Avoid storing third-party credentials on the ServiceNow instance

Before choosing a MID-less integration, review the following:

- MID-less integrations require ITOM Gateway. You must set up and configure ITOM Gateway before activating any MID-less integration. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.
- If your deployment requires on-premises data collection, bidirectional communication, or integration types not yet supported in MID-less mode, a MID Server based integration may be more appropriate.

MID-less integration setup from an MCP Client

Set up MID-less integrations for HLA from an AI-enabled MCP Client without opening the ServiceNow UI. This capability requires the ITOM MCP Server Console to be active on your instance. In addition, your MCP Client must be connected to the ITOM MCP Server Console with valid OAuth credentials. For more information, see MID-less log ingestion for HLA from an MCP Client.

MID-less log streaming via ITOM Gateway in Health Log Analytics

Health Log Analytics (HLA) can receive log data from external sources directly through the ITOM Gateway, without routing data through a MID Server. This architecture supports cloud-native log sources such as Amazon Data Firehose, Cribl, and OpenTelemetry, and is required for high-volume HLA deployments.

How log streaming via ITOM Gateway works

External log sources send data to the ITOM Gateway over gRPC. The ITOM Gateway forwards the data to the#Hermes#Messaging Service, which delivers it to the AI Engine, the HLA processing back-end. This design separates log streaming from ingestion, enabling higher throughput than direct MID Server streaming.

Deployment scenarios

The setup path depends on your expected log volume.

Log streaming via ITOM Gateway deployment scenarios

Deployment type	Setup
Standard	For typical log volumes, enable ITOM Gateway and the#Hermes#Messaging Service on your instance, configure a JSON Web Token (JWT) provider and token, and set up your log source from Integrations Launchpad.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Log streaming via ITOM Gateway deployment scenarios (continued)

Deployment type	Setup
	For more information, see Configure a JWT provider and token and Set up log streaming via ITOM Gateway for Health Log Analytics.
High-volume	For deployments requiring 30,000 or more log events per second, you must scale the HLA infrastructure before enabling ITOM Gateway. This process involves resizing the AI Engine and Elasticsearch nodes and coordinating a cross-team migration. Contact ServiceNow Support to request infrastructure scaling. For details and the required information to provide, see Set up log streaming via ITOM Gateway for Health Log Analytics **Note:** Infrastructure scaling requires a 6-hour change window and involves expected downtime of 2–6 hours for HLA functions only.

Supported log sources

Currently, the following log sources can stream data to HLA via ITOM Gateway:

- AWS Firehose
- Cribl Stream
- OpenTelemetry Collector

Each source has a dedicated tile in Integrations Launchpad for configuration.

Authentication

Log sources authenticate to HLA using a JWT token. You must configure a JWT provider and generate a token on the ServiceNow instance before activating an ITOM Gateway integration.

Key components

Key components

Component	Description
ITOM Gateway	Receives incoming log data from external sources over gRPC and forwards it to the #Hermes#Messaging Service.
Hermes#Messaging Service	Message broker that routes data from ITOM Gateway to the HLA back-end.
AI Engine	HLA back-end component that processes and analyzes ingested log data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Key components (continued)

Component	Description
Integrations Launchpad	ServiceNow interface for configuring log source integrations. Each supported log source has a dedicated tile.

Related topics

Configure a JWT provider and token

Set up log streaming via ITOM Gateway for Health Log Analytics

Configure a JWT provider and token

Configure a JSON Web Token (JWT) provider and token to authenticate log streaming integrations sending data to Health Log Analytics (HLA) via ITOM Gateway. This configuration is required before you can activate any ITOM Gateway integration in Integrations Launchpad.

Before you begin

- Review the MID-less log streaming via ITOM Gateway in Health Log Analytics documentation.
- Verify that Health Log Analytics version 36.0.19 or higher is installed on your instance.
- Verify that the user creating the KeyStore has appropriate access permissions for the `sys_certificate`, `jwt_key_store`, `jwt_provider`, and `sn_ics_jwt_config` tables and for the relevant cryptographic modules. For ACL recommendations, see the How to resolve Key Management Framework access denied errors for Password2 decryption [KB1112530] article in the Now Support Knowledge Base.

Role required: sys_admin

About this task

This procedure creates a JWT token for ITOM Cloud Services, which enables HLA to authenticate incoming log data from external sources via ITOM Gateway. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

Procedure

1. In a terminal or PowerShell session, generate a private key, certificate, and keystore. Use the following commands:

```
# Generate the private key
openssl genrsa -out private.key 2048

# Create the certificate (metadata prompts are optional)
openssl req -new -x509 -key private.key -out public_key.cer
 -days 9125 -sha256

# Create the keystore
openssl pkcs12 -export -in public_key.cer -inkey private.key
 -out my_keystore.p12 -name my_keystore
```

   These commands produce two output files used in later steps: `my_keystore.p12` and `public_key.cer`.

2. On the ServiceNow instance, create a certificate record that stores the keystore.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **All** > **System Definition** > **Certificates**.

b. Select **New** to create a record.

c. Fill in the required fields.

d. Set **Type** to Java Key Store.

e. In the **Key store password** field, enter the keystore password generated in step 1.

f. Attach the `my_keystore.p12` file created in step 1.

g. Select **Submit**.

Note:

- If you encounter keystore password errors, re-enter the password carefully and use the Validate Keystore option on the certificate record to confirm it is correct.
- If files are not attaching, confirm that both `my_keystore.p12` and `public_key.cer` are attached to their respective records.

3. Create a JWT Key that links the signing credentials to the certificate.

a. Navigate to **All** > **System OAuth** > **JWT Keys**.

b. Select **New** to create a record.

c. Fill in the **Name** field.

d. Set **Signing Keystore** to the certificate record created in step 2.

e. In the **Signing Key** field, enter the keystore password generated in step 1.

f. Select **Submit**.

4. Create a JWT Provider that defines how tokens are generated.

a. Navigate to **All** > **System OAuth** > **JWT Providers**.

b. Select **New** to create a record.

c. Fill in the **Name** field.

d. Set **Signing Configuration** to the JWT Key created in step 3.

e. Set **Expiry Interval (sec)** to `3600`.

f. Select **Submit**.

5. Create an ICS JWT Config record that completes the connection between the JWT Provider and HLA.

a. Navigate to **All** > **ITOM Cloud Services** > **ICS JWT Config**.

b. Select **New** to create a record.

c. Fill in the **Name** field.

d. Set **JWT Provider** to the JWT Provider created in step 4.

e. Attach the `my_keystore.p12` file created in step 1.

Note: This record requires its own copy of the `my_keystore.p12` file.

f. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The JWT provider and token are configured, enabling Health Log Analytics to authenticate incoming log data from external sources via ITOM Gateway.

What to do next
Set up log streaming via ITOM Gateway for Health Log Analytics.

Set up log streaming via ITOM Gateway for Health Log Analytics

Set up log streaming via ITOM Gateway to enable Health Log Analytics (HLA) to receive log data directly from external sources without a MID Server.

Before you begin
- Review the MID-less log streaming via ITOM Gateway in Health Log Analytics documentation.
- Verify that Health Log Analytics version 36.0.19 or higher is installed on your instance.
- Verify that HLA provisioning is complete. Confirm that the AI Engine and Elasticsearch show green status at: `https://<instance>.service-now.com/xmlstats.do?include=services_status`.
- Verify that ITOM Cloud Services Core is installed on your instance.
- Verify that a JWT provider and token are configured on your instance. For more information, see Configure a JWT provider and token.
- (High-volume deployments only) If you expect 30,000 or more log events per second, contact ServiceNow Support to request infrastructure scaling. Proceed to the log streaming via ITOM Gateway setup procedure only after ServiceNow confirms that scaling is complete.

  Provide the following information:

Required information for infrastructure scaling

Information	Example
Expected maximum log events per second	100,000 log events per second from all sources
Environment URL or URLs	<customer>prod.service-now.com
Preferred migration date and time	Sundays, overnight US time preferred

Note: Scaling requires a 6-hour change window and expected HLA downtime of 2–6 hours.

Role required: evt_mgmt_admin

About this task
Log streaming via ITOM Gateway removes the MID Server from the log ingestion path. External log sources send data directly to the ITOM Gateway, which routes it through the#Hermes #Messaging Service to the HLA AI Engine. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

Setup involves two stages: enabling ITOM Gateway and Hermes on your instance, and creating an integration from Integrations Launchpad.

Supported log sources include AWS Firehose, Cribl Stream, and OpenTelemetry Collector.

For more information about connecting a Cribl Stream instance to the ServiceNow ITOM Gateway, see the Set up Cribl to send data to HLA without a MID server [KB2271238] article in the Now Support knowledge base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Enable ITOM Gateway and the#Hermes#Messaging Service.
   a. Navigate to **All**.
   b. In the navigation filter, enter **sys_service.list** and press Enter.
   c. In the **Services** table, verify that an entry for `itom-hla-gw` exists with a corresponding `sys_service_endpoint` entry.
   d. Navigate to **All** > **MID Server** > **Properties**.
   e. Search for the property `mid.hla.itom_gateway_streaming.enabled` and set it to true.
   All MID Servers automatically switch to ITOM Gateway streaming mode.

   Note: MID Servers re-establish gRPC connections. Data inputs are not restarted.

2. Set up the integration from Integrations Launchpad.
   For more information, see Set up integrations for Health Log Analytics from the Integrations Launchpad.
3. On the **Overview** tab of the integration, confirm that the ITOM Gateway URL, Integration ID, and Token values are populated.
   If these values appear, log streaming via ITOM Gateway setup is complete.

Set up an Amazon Data Firehose integration for Health Log Analytics

Set up an integration to stream log data from Amazon Data Firehose directly to the ServiceNow datacenter, where it's queued for Health Log Analytics processing. There's no need to store AWS keys on your ServiceNow instance.

Before you begin

Set up MID-less log streaming via ITOM Gateway. Choose the deployment method based on your expected log volume. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

As an alternative to the provided procedure, set up a MID-less integration for HLA from an AI-enabled MCP Client. For more information, see MID-less log ingestion for HLA from an MCP Client.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Firehose` in the search field.
5. Select the Amazon Data Firehose (MID-less) integration tile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see Amazon Data Firehose integration configuration fields.
7. Select **Next**.
8. Follow one of the procedures provided on the **Set-up instruction** screen to install the integration in AWS.
   You can perform the installation procedure using either AWS CloudFormation or the AWS Management Console. Both methods require the same input. The appropriate choice depends on your requirements and the size of your deployment:
   - Use AWS CloudFormation when you need automation, repeatability, and infrastructure-as-code benefits. It's best for deploying standardized configurations across multiple environments. AWS CloudFormation prompts for input automatically.
   - Use the AWS Management Console when you need manual setups or quick configurations without automation.

   Before you perform either of these procedures, copy the access token, the ServiceNow endpoint URL, and the integration ID to the clipboard. You need these values during the installation process.
9. Do one of the following.
   - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
     - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

       When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
     - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

     On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:

- Log data auto-mapping and mapping
- Source type structure adjustment
- Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Cribl Stream (MID-less) integration for Health Log Analytics

Set up a Cribl Stream integration to stream Cribl log data directly to your ServiceNow instance, without a MID Server.

Before you begin

Set up MID-less log streaming via ITOM Gateway. Choose the deployment method based on your expected log volume. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

As an alternative to the provided procedure, set up a MID-less integration for HLA from an AI-enabled MCP Client. For more information, see MID-less log ingestion for HLA from an MCP Client.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, search for the Cribl Stream integration and select its tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

5. On the **Provide details** form, fill in the fields and then select **Next**.
6. On the **Set-up instruction** screen, follow the instructions to enable sending Cribl log data to ServiceNow using OpenTelemetry.

   a. Copy the provided access token, endpoint, and integration ID to the clipboard.
   These credentials enable Cribl to securely send logs to the ServiceNow ITOM Gateway.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Credentials

Credential	Description
Access Token	The JWT credential passed with every request (servicenow-access-token).
Endpoint	The Gateway gRPC endpoint, the address to which Cribl Stream will send the log data.
Integration ID	The ID by which the integration identifies to ServiceNow (servicenow-integration-id).

b. In the Cribl UI or configuration files, depending on your deployment, set up a Cribl destination that uses the OpenTelemetry Protocol (OTLP) to forward data to ServiceNow.

Note: Cribl integrates directly with ServiceNow via API. No collector or agent is needed.

c. In Cribl, paste the copied credentials into the destination configuration.

For more detailed instructions, see the Set up Cribl to send data to HLA without a MID Server [KB2271238] article in the Now Support knowledge base.

7. Do one of the following.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu ( ) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:

- Log data auto-mapping and mapping
- Source type structure adjustment
- Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Microsoft Azure Event Hubs (MID-less) integration for Health Log Analytics

Set up an integration for streaming events from Microsoft Azure Event Hubs to your ServiceNow instance without a MID Server.

Before you begin

Set up MID-less log streaming via ITOM Gateway. Choose the deployment method based on your expected log volume. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

As an alternative to the provided procedure, set up a MID-less integration for HLA from an AI-enabled MCP Client. For more information, see MID-less log ingestion for HLA from an MCP Client.

Role required: evt_mgmt_admin

About this task

This Microsoft Azure Event Hubs integration eliminates the need for a MID Server, which simplifies setup and enables more efficient log data ingestion into your instance. By removing the MID Server requirement, log data collection and processing are streamlined to promote faster analysis and root cause identification.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Azure Event Hubs (MID-less)` in the search field.
5. Select the Azure Event Hubs (MID-less) integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields and then select **Next**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Event Hubs MID-less integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Description	Option to add a brief description of the integration to help identify it.
Transport	Protocol used to stream the log data. This field is read-only. This integration uses Microsoft Azure Event Hubs to stream log data to your instance.

7. On the **Set-up instruction** screen, perform the appropriate procedure to install the integration in your third-party console.
   - Install the integration with Bicep, a domain-specific language (DSL) that uses declarative syntax to deploy Azure resources.

     Using Bicep enables automation and consistency across environments through Infrastructure as Code (IaC), allowing you to manage your Azure environment more efficiently. It also supports version control and integration with CI/CD pipelines.

   - Install the integration in the Azure portal.

     If you're doing a one-time setup or prefer a graphical interface, installing in the Azure portal is a good choice. It provides a convenient way to configure resources, making it suitable for straightforward deployments.

   The screen provides a link to detailed configuration instructions for each of these two installation options.
8. Do one of the following.
   - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
     - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

       When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

     - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

     On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

   - If you didn't complete all the prior requirements, select **Save draft**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu ( ) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up an OpenTelemetry Collector integration for Health Log Analytics

Set up an OpenTelemetry Collector integration to stream log data directly to your ServiceNow instance using the OpenTelemetry (OTLP) protocol, without a MID Server.

Before you begin

Set up MID-less log streaming via ITOM Gateway. Choose the deployment method based on your expected log volume. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

As an alternative to the provided procedure, set up a MID-less integration for HLA from an AI-enabled MCP Client. For more information, see MID-less log ingestion for HLA from an MCP Client.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, search for the OpenTelemetry (MID-less) integration and select its tile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

5. On the **Provide details** form, fill in the fields and then select **Next**.
6. On the **Set-up instruction** screen, follow the provided procedure to enable sending log data to ServiceNow using OpenTelemetry.

   a. Copy the provided access token, endpoint, and integration ID to the clipboard.
   These credentials enable the collector to receive telemetry data and forwarding it to the ServiceNow ITOM Gateway.

   Credentials

Credential	Description
Access Token	The JWT credential passed with every request.
Endpoint	The address to which the data is sent: the ITOM Gateway.
Integration ID	The ID with which the integration identifies to ServiceNow.

   b. In your OpenTelemetry Collector configuration file, exporters section, add an OTLP exporter to forward telemetry data to ServiceNow.
   Typically, the configuration file is `otel-collector-config.yaml`.

```
exporters:
  otlp/servicenow:
    endpoint:
itomgw-prod-gateway-phxbwi.sncapps.service-now.com:443
    headers:
      # Reads the access token from the
SERVICENOW_ACCESS_TOKEN env variable
      servicenow-access-token: ${SERVICENOW_ACCESS_TOKEN}
      servicenow-integration-id:
f7452b23877c4b10ebdd4335dabb3553
```

   This step tells the collector where to send the data and how to authenticate with your instance.

   For more detailed instructions, see the Configure an OpenTelemetry Collector or 3rd Party agent to send data to HLA [KB2117238] article in the Now Support knowledge base. For further assistance, contact your OpenTelemetry Collector administrator.

7. Do one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Splunk OpenTelemetry Collector (MID-less) integration for Health Log Analytics

Set up a Splunk OpenTelemetry Collector integration to stream Splunk log data directly to your ServiceNow instance using the OpenTelemetry (OTLP) protocol, without a MID Server.

Before you begin

Set up MID-less log streaming via ITOM Gateway. Choose the deployment method based on your expected log volume. For more information, see MID-less log streaming via ITOM Gateway in Health Log Analytics.

As an alternative to the provided procedure, set up a MID-less integration for HLA from an AI-enabled MCP Client. For more information, see MID-less log ingestion for HLA from an MCP Client.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, search for the Splunk OpenTelemetry (MID-less) integration and select its tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

5. On the **Provide details** form, fill in the fields and then select **Next**.
6. On the **Set-up instruction** screen, follow the provided procedure to enable sending Splunk log data to ServiceNow using OpenTelemetry.

   a. Copy the provided access token, endpoint, and integration ID to the clipboard.
   These credentials enable Cribl to securely send logs to the ServiceNow ITOM Gateway.

   Credentials

Credential	Description
Access Token	The JWT credential passed with every request.
Endpoint	The address to which the data is sent: the ITOM Gateway.
Integration ID	The ID with which the integration identifies to ServiceNow.

   b. In your OpenTelemetry Collector configuration file, exporters section, add an OTLP exporter to forward telemetry data to ServiceNow.
   Typically, the configuration file is `otel-collector-config.yaml`.

```
exporters:
  otlp/servicenow:
    endpoint:
 itomgw-prod-gateway-phxbwi.sncapps.service-now.com:443
    headers:
      # Reads the access token from the
 SERVICENOW_ACCESS_TOKEN env variable
      servicenow-access-token: ${SERVICENOW_ACCESS_TOKEN}
      servicenow-integration-id:
 7b36df1387f00f10d696c845dabb351f
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This step tells the collector where to send the data and how to authenticate with your instance.

For more detailed instructions, see the Configure an OpenTelemetry Collector or 3rd Party agent to send data to HLA [KB2117238] article in the Now Support knowledge base. For further assistance, contact your OpenTelemetry Collector administrator.

7. Do one of the following.
   - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
     - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

       When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

     - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

     On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up an ACC Log Analytics integration for Health Log Analytics

Set up an ACC Log Analytics (ACC-L) integration for streaming log messages to your ServiceNow instance. This integration can be used when agent-less log collection is impractical, for example because of security requirements or host accessibility.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify you're using the latest release version of the Health Log Analytics plugin to enable ACC Log Analytics integration setupfrom the Integrations Launchpad.
- Verify that a MID Server is installed and configured. For more information, see MID Server system requirements.
- The MID Server must support the Agent Client Collector Listener. The ACC Listener can be enabled during ACC Log Analytics integration setup from the Integrations Launchpad.
- The Agent Client Collector Log Analytics application has a dependency on the following ServiceNow applications, available from the ServiceNow Store:
  - Health Log Analytics, Version 22.0.12 - December 2021 and later releases.
  - Agent Client Collector - Framework, Version 2.7.0 - December 2021 and later releases.

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, search for `Agent Log Analytics`.
5. Select the Agent Log Analytics integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, complete the fields.
   For a description of the fields, see the **Provide details** table in ACC Log Analytics (ACC-L) integration configuration fields.
7. **Optional:** Select **Advanced settings** and complete the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in ACC Log Analytics (ACC-L) integration configuration fields.
8. Select **Next** to save your settings and proceed to the **Setup instructions** tab.
9. Enable Agent Client Collector Log Analytics to stream logs from your IT infrastructure to Health Log Analytics by creating log collection policies for your configuration items (CIs).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Create an Agent Client Collector log policy.

10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

    - If you didn't complete all the prior requirements, select **Save draft**.

      The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up an Amazon CloudWatch integration for Health Log Analytics

Set up an integration for streaming log data from Amazon CloudWatch to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

> **Important:** Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, search for `CloudWatch`.
5. Select the Amazon CloudWatch integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Amazon CloudWatch integration configuration fields.
7. Select **Next** to save your settings and proceed to the next configuration form.
8. On the **Set data retrieval method** form, fill in the fields.
   For a description of the fields, see the **Set data retrieval method** table in Amazon CloudWatch integration configuration fields.
9. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Amazon CloudWatch integration configuration fields.
10. Do one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up an Amazon S3 integration for Health Log Analytics

Set up an integration for streaming log data from Amazon S3 (Simple Storage Service) buckets to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

> **Important:** Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, search for S3.
5. Select the Amazon S3 integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Amazon S3 integration configuration fields.
7. Select **Next** to save your settings and proceed to the next configuration form.
8. On the **Set data retrieval method** form, fill in the fields.
   For a description of the fields, see the **Set data retrieval method** table in Amazon S3 integration configuration fields.
9. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Amazon S3 integration configuration fields.
10. Do one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up an Apache Kafka integration for Health Log Analytics

Set up an integration for streaming log data from Apache Kafka to your ServiceNow instance for processing by Health Log Analytics.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

> **Important:** Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

- Consider whether to pull the log data to a single MID Server or to a MID Server cluster.

  In MID Server clusters, multiple MID Servers are grouped together for failover protection. The active integration runs on a single MID Server in the cluster. If that MID Server fails, the system moves its tasks to the next available MID Server in the cluster in a configured order.

  If you decide to use a MID Server cluster, verify the following:

  - The MID Server cluster includes only MID Servers that support basic authentication.

    Note: mTLS is not supported for log ingestion.

  - Log ingestion is enabled for each MID Server in the cluster.

    Note: If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically.

  - No more than the default maximum of 10 integrations will pull logs to a single MID Server.

    A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 integrations running on it, even when that MID Server is down.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, search for `Kafka`.
5. Select the Kafka integration tile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Apache Kafka integration configuration fields.
7. Select **Next**.
8. On the **Set data retrieval method** form, fill in the fields.
   For a description of the fields, see the **Set data retrieval method** table in Apache Kafka integration configuration fields.
9. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Apache Kafka integration configuration fields.
10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
    - If you didn't complete all the prior requirements, select **Save draft**.

      The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Cribl integration for Health Log Analytics

Set up an integration to enable Health Log Analytics to process Cribl log messages streaming into your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

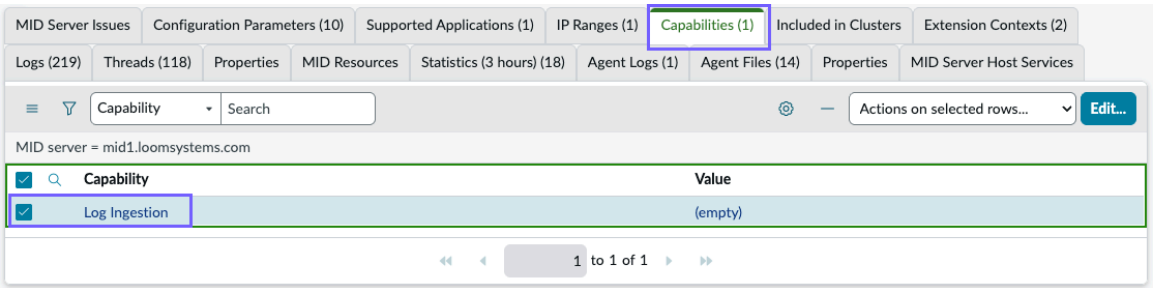

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

If your organization uses Cribl for filtering and routing large volumes of log data from various sources, the log format received by Health Log Analytics is distinct from other types. The Cribl integration enables HLA to detect and separate transport headers from inner log messages in this format. It then forwards only the inner message to the source type structure for processing.

You set up integrations through the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.

3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Cribl` in the search field.
5. Select the Cribl integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see Cribl integration configuration fields.
7. Select **Next**.
8. On the **Set up connection** form, fill in the fields.
   For a description of the fields, see Cribl integration configuration fields.
9. Do one of the following.
   - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
     - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

       When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

     - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

     On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Log data auto-mapping and mapping
- Source type structure adjustment
- Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up Cribl integration authentication credentials

Configure credentials in Health Log Analytics for the authentication method used by the Cribl integration.

Before you begin

Decide on the appropriate authentication method for your instance:

- Basic authentication is preferred for self-hosted (On-Prem) Cribl instances.
- Token-based authentication is required for Cribl Cloud instances.

> **Note:** Token-based authentication is also supported for self-hosted (On-Prem) Cribl instances.

For token-based authentication, create a client ID and a client secret on your Cribl instance and use them to generate a token. For example:

- Client ID: 1234567890abcdef
- Client secret: abcdef1234567890

For detailed information on the Cribl API, refer to the Cribl API documentation.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
2. Select the default Cribl credential alias: Cribl_Credential_Alias.
3. Select **New** to create a credential record.
4. Select the appropriate credentials type for your Cribl instance:
   - For basic authentication, select **Basic Auth Credentials**.
   - For token-based authentication, select **Custom Header API Key Credential**.
5. On the displayed form, fill in the fields.
   - For basic authentication:

Basic Auth credentials fields

Field	Description
Name	A name that describes the credentials.
Username	Your Cribl server username. For example: admin.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Password	Your Cribl server password. For example: admin123.

- For token-based authentication:

Custom Header API Key credential fields

Name	A name that describes the credential.
API Key	The authentication token you generated on the Cribl instance.

6. Select **Save** to save the credentials configuration.

What to do next

- Verify that the provided credentials have appropriate permissions to access your Cribl instance.
- Verify that only one active credential record of each credentials type exists under the credential alias.

Set up an Edge Delta REST integration for Health Log Analytics

Set up an Edge Delta REST integration to enable Health Log Analytics to process logs it receives from Edge Delta in a distinct format. These logs stream into your ServiceNow instance via REST.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

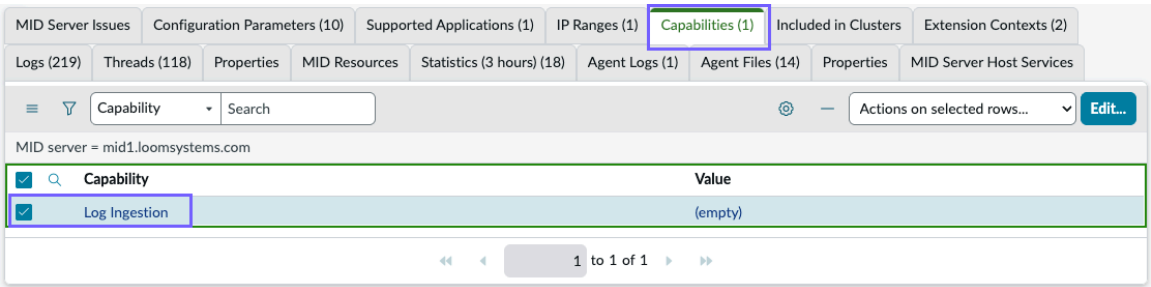

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

If your organization uses Edge Delta to handle large volumes of log data from multiple sources before sending it to Health Log Analytics, the log format HLA receives from Edge Delta is different from other types. The Edge Delta integration enables HLA to detect and separate transport headers from inner log messages in this format. It forwards only the inner message to the source type structure for processing.

You set up integrations through the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Edge Delta` in the search field.
5. Select the Edge Delta REST integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see Edge Delta REST integration configuration fields.
7. Select **Next**.
8. Follow the instructions in the **Set-up instruction** page to install the integration in the third-party console.
   a. Copy the displayed URL.
   b. In Edge Delta, create an HTTP destination.
   c. Paste the copied URL into the HTTP destination.
   d. Add the following headers:
      - Authorization

        Use either of the following:
        - Key <API_KEY>

          Find the API_KEY by navigating to **All** > **MID Web Server API Key** > **View API Key**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Basic <Base64 (username:password)>

  Note: Verify that you use a username and password for a MID web server with Basic Auth enabled.

- Content-Type: application/json.

9. Do one of the following.
   - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
     - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

       When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

     - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

     On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu ( ) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up an Edge Delta TCP integration for Health Log Analytics

Set up an Edge Delta TCP integration to enable Health Log Analytics to process logs it receives from Edge Delta in a distinct format. These logs stream into your ServiceNow instance over the TCP transport protocol.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

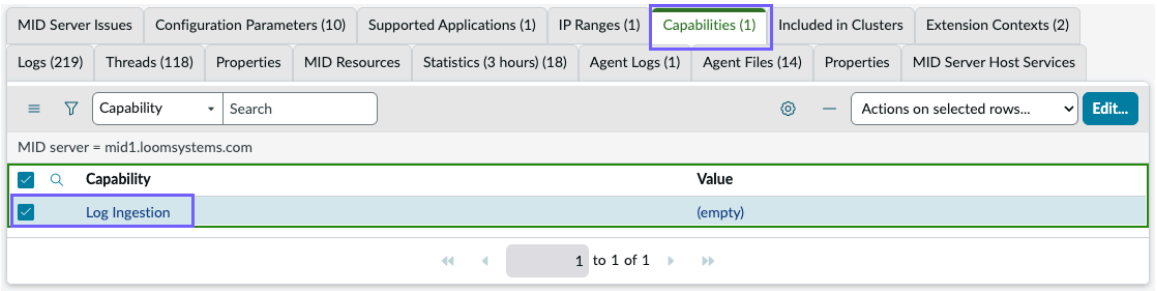

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

If your organization uses Edge Delta to handle large volumes of log data from multiple sources before sending it to Health Log Analytics, the log format HLA receives is different from other types. The Edge Delta integration enables HLA to detect and separate transport headers from inner log messages in this format. It forwards only the inner message to the source type structure for processing.

You set up integrations through the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Edge Delta` in the search field.
5. Select the Edge Delta TCP integration tile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

> **Note:** If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Edge Delta TCP integration configuration fields.
7. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Edge Delta TCP integration configuration fields.
8. Select **Next**.
9. In the **Set-up instruction** screen, copy the displayed Host IP and Port values and paste them in your Edge Delta pipeline under TCP destination.

   > **Note:** The Host IP and Port values are based on your configuration on the **Provide details** form.

10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
    - If you didn't complete all the prior requirements, select **Save draft**.

      The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Log data auto-mapping and mapping
- Source type structure adjustment
- Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up an Elasticsearch integration for Health Log Analytics

Set up an integration to stream log data seamlessly from Elasticsearch indices to your instance for Health Log Analytics processing.

Before you begin

Note: Health Log Analytics supports Elasticsearch versions above 7.10.2 and below 8.18.2. For advanced Elasticsearch log streaming guidance, see the Stream logs using Elasticsearch data input - Advanced guide [KB1080162] article in the Now Support Knowledge Base.

- Verify that the Health Log Analytics application is installed and provisioned on your instance. For more information, see Install Health Log Analytics (HLA).
- Verify that a service instance is available.
- Verify that the Health Log Analytics AI Engine is up and running.

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

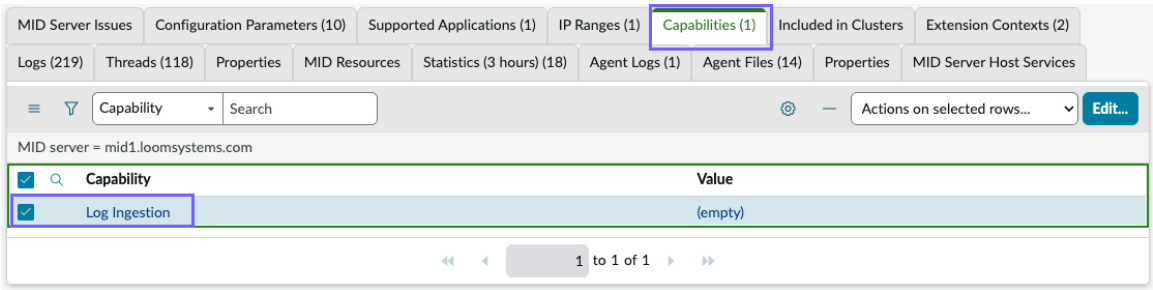

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named `mid.public_ip` with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter *`Elasticsearch`* in the search field.
5. Select the Elasticsearch integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Elasticsearch integration configuration fields.
7. Select **Next**.
8. On the **Set data retrieval method** form, fill in the fields.
   For a description of the fields, see Elasticsearch integration configuration fields.
9. **Optional:** Select **Advanced settings** to set advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Elasticsearch integration configuration fields.
10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

    - If you didn't complete all the prior requirements, select **Save draft**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a GCP PubSub integration for Health Log Analytics

Set up an integration for receiving log messages that were published to a Google Cloud Platform (GCP) Pub/Sub topic and streaming them to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

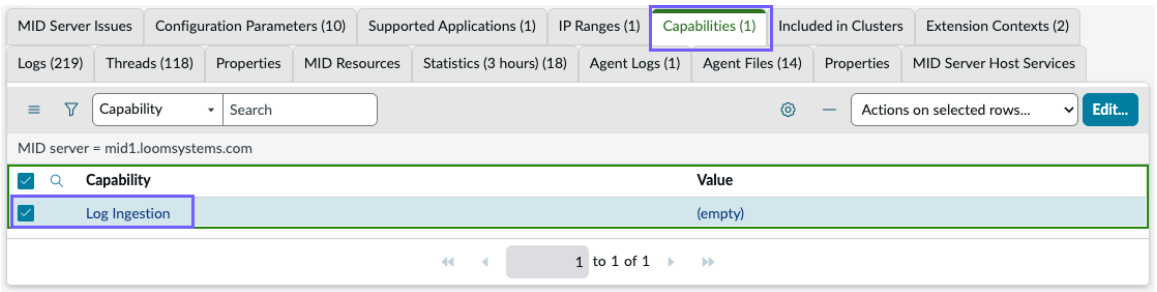

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Google Cloud` in the search field.
5. Select the Google Cloud Platform integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in GCP PubSub integration configuration fields.
7. Select **Next**.
8. On the **Set up connection** form, fill in the fields to connect to GCP PubSub.
   For a description of the fields, see the **Set up connection** table in GCP PubSub integration configuration fields.
9. **Optional:** Select **Advanced Settings** and fill in the fields.
   For a description of the fields, see the **Advanced Settings** table in GCP PubSub integration configuration fields.
10. Do one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Microsoft Azure Event Hubs integration for Health Log Analytics

Set up an integration for streaming events from Microsoft Azure Event Hubs to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

> **Important:** Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Event Hubs` in the search field.
5. Select the relevant Azure Event Hubs integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields and then select **Next**.
   For a description of the fields, see the **Provide details** table in Microsoft Azure Event Hubs integration configuration fields.
7. On the **Set data retrieval** form, fill in the fields.
   For a description of the fields, see the **Set data retrieval** table in Microsoft Azure Event Hubs integration configuration fields.
8. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Microsoft Azure Event Hubs integration configuration fields.
9. Do one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Microsoft Azure Log Analytics integration for Health Log Analytics

Set up an integration for streaming log data from Microsoft Azure Log Analytics to your ServiceNow instance. The integration points the Health Log Analytics AI engine to a data source in your Microsoft Azure Log Analytics account.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

> **Important:** Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Azure Log` in the search field.
5. Select the Azure Log Analytics integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields and then select **Next**.
   For a description of the fields, see the **Provide details** table in Microsoft Azure Log Analytics integration configuration fields.
7. On the **Set data retrieval method** form, fill in the fields.
   For a description of the fields, see the **Set data retrieval method** table in Microsoft Azure Log Analytics integration configuration fields.
8. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Microsoft Azure Log Analytics integration configuration fields.
9. Do one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a MID Server integration for Health Log Analytics

Set up an integration for collecting and streaming MID Server log messages to your ServiceNow instance for processing by Health Log Analytics.

Before you begin

Important: You can create only one MID Server integration per MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

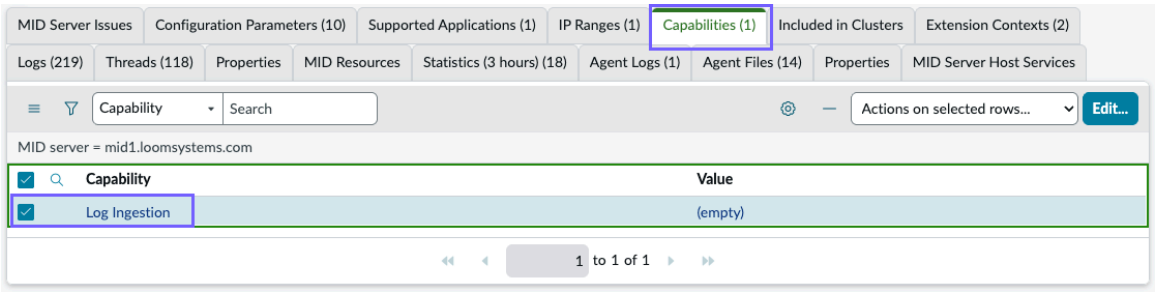

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `mid` in the search field.
5. Select the MID Server integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a description of the fields, see the **Provide details** table in MID Server integration configuration fields.

7. Select **Advanced Settings** and fill in the fields.
   For a description of the fields, see the **Advanced Settings** table in MID Server integration configuration fields.
8. Do one of the following:
   - If you completed all the prerequisites before starting the configuration, select **Activate**.

     The integration is activated and the **Overview** tab is displayed. On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. You can complete the prerequisites and activate the installation later. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

Set up a REST API integration for Health Log Analytics

Set up an integration for streaming log data to your ServiceNow instance for processing by Health Log Analytics.

Before you begin

- Verify that the Health Log Analytics application is installed and provisioned on your instance. For more information, see Install Health Log Analytics (HLA).
- Verify that a service instance is available.
- Verify that the Health Log Analytics AI Engine is up and running.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

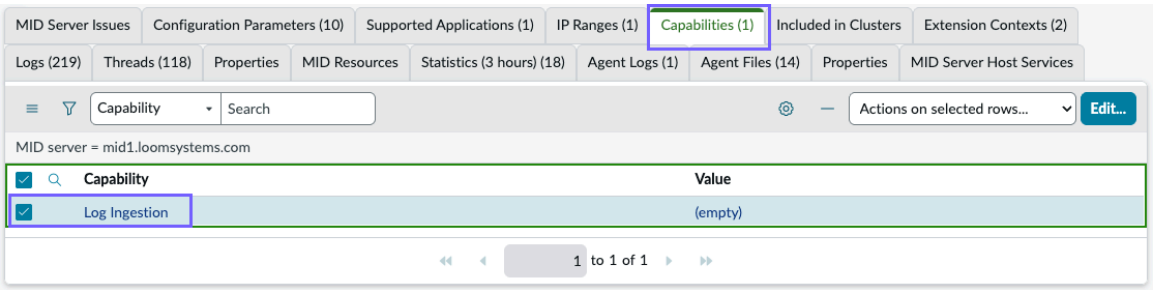

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

On the MID Web Server Context form, in the **Execute on** field, select *Specific MID Server* and not the cluster option. In the **MID Server** field, select the specific MID Server to which the log data is pulled.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named `mid.public_ip` with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Note: REST API integrations support only UTF-8 encoding for incoming data.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter REST in the search field.
5. Select the REST API integration tile.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

> **Note:** If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see REST API integration configuration fields.
7. Select **Next**.
8. In the **Set-up instruction** screen, copy the provided URL to the clipboard.

   The format of the URL is: `http://(MID_SERVER_IP):(MID_WEB_SERVER_PORT)/api1/mid/hla/raw`, where `/aps/mid/hla/raw` is the endpoint. HLA has replaced `(MID_SERVER_IP)` and `(MID_WEB_SERVER_PORT)` with the values you configured on the **Provide details** form.

9. Do one of the following.
   - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
     - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

       When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

     - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

     On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

Instruct your API client to send log messages to the MID Server.

1. In the API client, configure a request with the following parameters:
   - HTTP Method: POST.
   - URL: The URL you copied in the **Set-up instruction** screen.
   - Authentication: The required authentication credentials. For example, Basic authentication (username and password).

     > **Note:** The authentication method must be the same as the method configured on the MID Web Server.

   - Headers: Set the `Content-Type` header to match your payload format.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For JSON: `application/json`.
- For raw text: `text/plain`.

- Body: The logs payload in JSON or raw text.

2. Send the request.

   A successful response indicates that the MID Server has received the payload. When the MID Server has processed the log messages, the data input streams them to your ServiceNow instance using the REST protocol.

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a ServiceNow System Logs Retriever integration for Health Log Analytics

Set up an integration for streaming log data from your ServiceNow platform's System Log table to the HLA engine.

Before you begin

- Verify that the Health Log Analytics application is installed and provisioned on your instance. For more information, see Install Health Log Analytics (HLA).
- Verify that a service instance is available.
- Verify that the HLA Engine is up and running.

Note:

- This integration doesn't run on a MID Server.
- Only a single ServiceNow System Logs Retriever integration or data input can exist in the system.
- The HLA store app installation automatically creates the ServiceNow System Logs Retriever integration. Use this procedure if you need to recreate it manually.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `System Logs` in the search field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Select the ServiceNow System logs tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. The **Provide details** form is displayed.
   These fields are auto-populated, because only one ServiceNow System Logs Retriever integration or data input can exist in the system. For a description of the fields, see the **Provide details** table in ServiceNow System Logs Retriever integration configuration fields.
7. Select **Next**.
8. **Optional:** On the **Set data retrieval method** page, build conditions for filtering log data from the System Log table before it is streamed to the HLA engine.

   (Optional) This page helps you focus on the specific logs you want to monitor. For example, you might want to look at Discovery logs or exclude email logs. By default, all logs are included.

9. Do one of the following:
   - If you completed all the prerequisites before starting the configuration, select **Activate**.

     The integration is activated and the **Overview** tab is displayed. On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. You can complete the prerequisites and activate the installation later. For more information, see Activate a draft integration in Health Log Analytics.

Result

Log data starts streaming from the ServiceNow System Log table to the Health Log Analytics AI engine, based on the configured filters. The dedicated HLA engine for your instance starts analyzing your system logs in real-time.

Note: Starting with Health Log Analytics version 37.0.15, you can use the ServiceNow System Logs Retriever integration to automatically detect errors and anomalies in your ServiceNow instance's system log. By combining HLA anomaly detection with ServiceNow Otto for ITOM, you can identify and resolve platform issues more effectively for a better user experience. For more information, see Use Case: Proactive monitoring of your ServiceNow instance in Health Log Analytics.

What to do next

Use the information in the **Overview** tab to refine how HLA reads the log data by adjusting your integration configuration. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.

Note:

- The MID Server streaming status and the **Log streaming sources** table do not appear on the **Overview** tab for this integration, because it doesn't run on a MID Server.
- When the integration is first configured, it may take up to 24 hours for the first alerts to appear. Once the data is populated, you can refresh the screen to see real-time metrics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the More menu to go directly from the **Overview** tab to the **Data Input Mapping**, **Source Type Structures**, and **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More menu also provides direct access to the **Log Viewer**, where you can see the raw logs that the integration ingests.

Set up a Splunk Poller integration for Health Log Analytics

Set up an integration that periodically pulls log data from Splunk to your ServiceNow instance for processing by Health Log Analytics.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

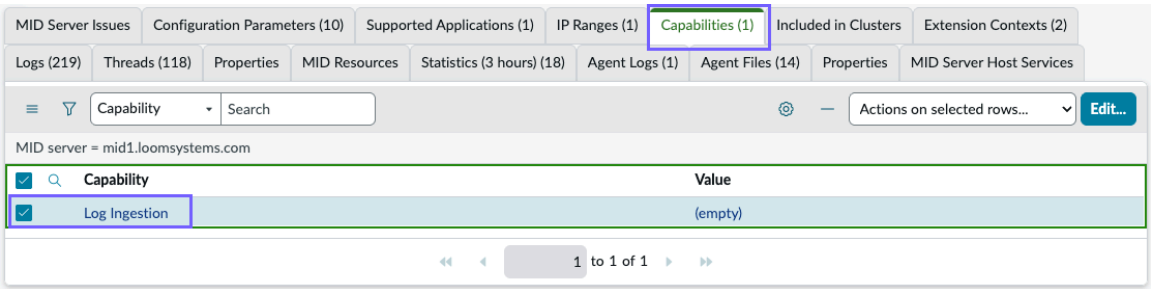

> **Important:** Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.
- The MID Server must support basic authentication.

  > **Note:** mTLS is not supported for log ingestion.

- No more than the default maximum of 10 integrations will stream logs to a single MID Server. You can modify the maximum number by adding the property `sn.occ.log_ingestion.max_datainputs_per_mid` to the MID Server and then changing the default value.

  To find out how many data inputs are streaming logs to the same MID Server, navigate to the **Streaming Sources** table. Then count the data inputs that stream to a specific MID Server.

- Consider whether to pull the log data to a single MID Server or to a MID Server cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In MID Server clusters, multiple MID Servers are grouped together for failover protection. The active integration runs on a single MID Server in the cluster. If that MID Server fails, the system moves its tasks to the next available MID Server in the cluster in a configured order.

If you want to use a MID Server cluster, verify the following:

- The MID Server cluster includes only MID Servers that support basic authentication.
- Log ingestion is enabled for each MID Server in the cluster.

  Note: If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically.

- No more than the default maximum of 10 integrations can stream logs to a single MID Server.

  A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 integrations running on it, even when that MID Server is down.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Splunk Poller` in the search field.
5. Select the Splunk Poller integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Splunk Poller integration configuration fields.
7. Select **Next**.
8. On the **Set data retrieval method** form, fill in the fields.
   For a description of the fields, see the **Set data retrieval method** table in Splunk Poller integration configuration fields.
9. Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see Splunk Poller integration configuration fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

    - If you didn't complete all the prior requirements, select **Save draft**.

      The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu ( ) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Splunk TCP integration for Health Log Analytics

Set up an integration to stream log messages to your ServiceNow instance over the TCP transport protocol using a Splunk heavy forwarder. Health Log Analytics processes the ingested log data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

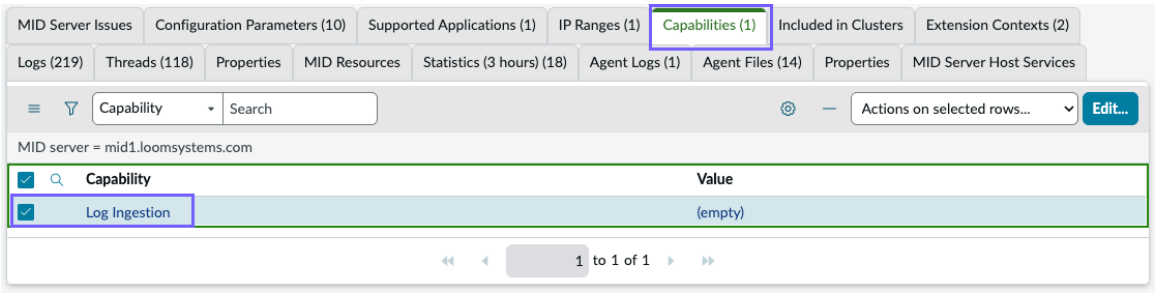

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.
- The MID Server must support basic authentication.

  Note: mTLS is not supported for log ingestion.

- No more than the default maximum of 10 integrations will stream logs to a single MID Server. You can modify the maximum number by adding the property `sn.occ.log_ingestion.max_datainputs_per_mid` to the MID Server and then changing the default value.

  To find out how many data inputs are streaming logs to the same MID Server, navigate to the **Streaming Sources** table. Then count the data inputs that stream to a specific MID Server.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Splunk` in the search field.
5. Select the Splunk TCP integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Splunk TCP integration configuration fields.
7. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Splunk TCP integration configuration fields.
8. Select **Next**.
9. Follow the procedure on the **Set-up instruction** screen to install the integration in the third-party console.

   Note: The procedure varies based on your configurations.

10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
    - If you didn't complete all the prior requirements, select **Save draft**.

      The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a Splunk UDP integration for Health Log Analytics

Set up an integration to stream log messages to your ServiceNow instance over the UDP transport protocol using a Splunk heavy forwarder. Health Log Analytics processes the ingested log data.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

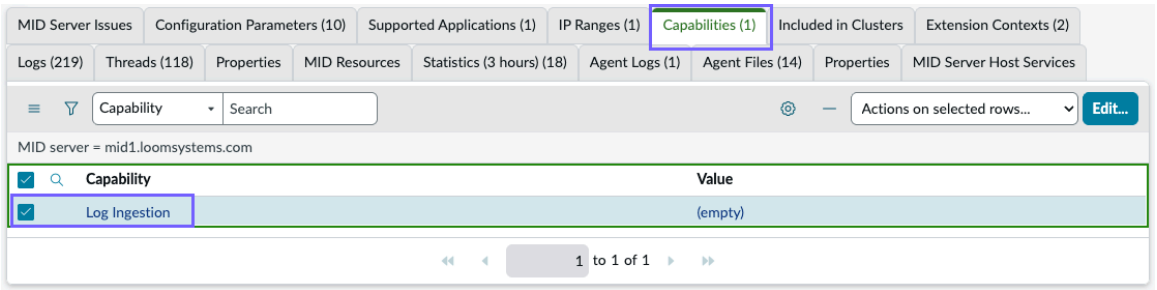

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Splunk` in the search field.
5. Select the Splunk UDP integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see **Provide details** table in Splunk UDP integration configuration fields.
7. **Optional:** Select **Advanced settings** and fill in the advanced configuration fields.
   For a description of the fields, see the **Advanced settings** table in Splunk UDP integration configuration fields.
8. Select **Next**.
9. Follow the procedure on the **Set-up instruction** screen to install the integration in the third-party console.

   Note: The procedure varies based on your configurations.

10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

    - If you didn't complete all the prior requirements, select **Save draft**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a TCP integration for Health Log Analytics

Set up an integration for sending log data to your ServiceNow instance directly over a TCP/SSL socket.

Before you begin

- Verify that the Health Log Analytics application is installed and provisioned on your instance. For more information, see Install Health Log Analytics (HLA).
- Verify that a service instance is available.
- Verify that the HLA Engine is up and running.
- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

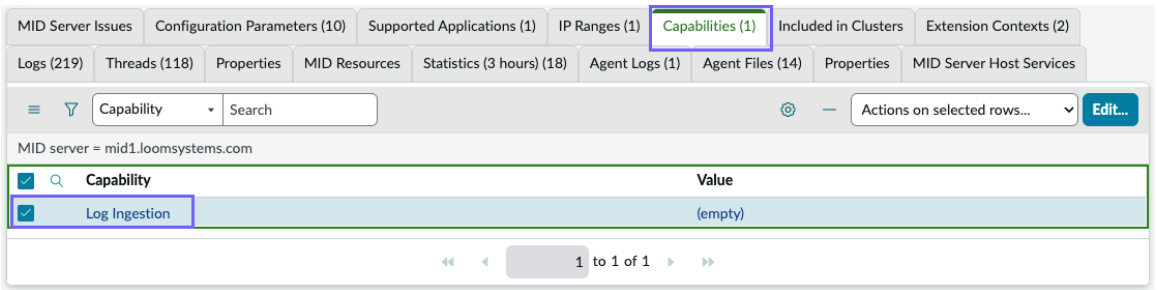

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Note: For information about streaming data from Logstash to your instance, see the Streaming data from Logstash to your HLA instance [KB0994735] article in the Now Support knowledge base.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter TCP in the search field.
5. Select the TCP integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in TCP integration configuration fields.
7. Select **Next**.
8. In the **Set-up instruction** screen, follow the instructions on the screen to install the integration in the third-party console.
9. Do one of the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
  - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

    When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
  - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

  On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
- If you didn't complete all the prior requirements, select **Save draft**.

  The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Set up a UDP integration for Health Log Analytics

Set up an integration for sending log data to your ServiceNow instance directly over a UDP socket.

Before you begin

- Verify that the Health Log Analytics application is installed and provisioned on your instance. For more information, see Install Health Log Analytics (HLA).
- Verify that a service instance is available.
- Verify that the HLA Engine is up and running.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

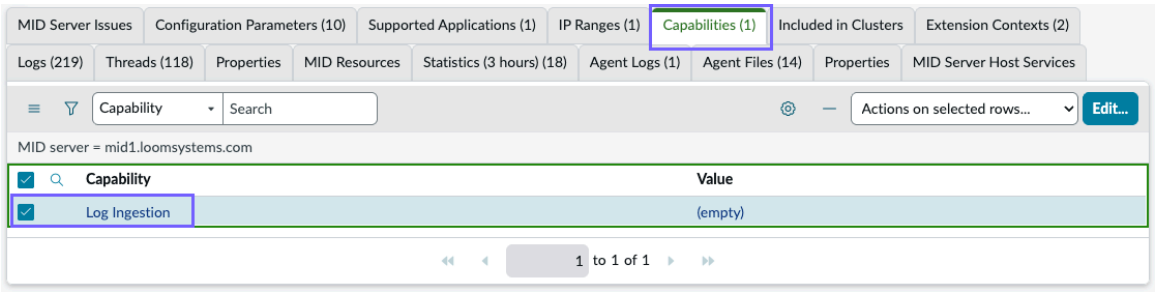

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

Set up an integration from the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter UDP in the search field.
5. Select the UDP integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a description of the fields, see the **Provide details** table in UDP integration configuration fields.

7. **Optional:** Select **Advanced Settings** to set advanced configuration fields.
   For a description of the fields, see the **Advanced Settings** table in UDP integration configuration fields.
8. Select **Next**.
9. In the **Setup instruction** screen, follow the instructions on the screen to install the integration in the third-party console.
10. Do one of the following.
    - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
      - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

        When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.

      - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

      On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.

    - If you didn't complete all the prior requirements, select **Save draft**.

      The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:
  - Log data auto-mapping and mapping
  - Source type structure adjustment
  - Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up a Vector Agent integration for Health Log Analytics

Set up a Vector Agent integration to enable Health Log Analytics to process log messages that are streaming into your ServiceNow instance via a Vector Agent.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

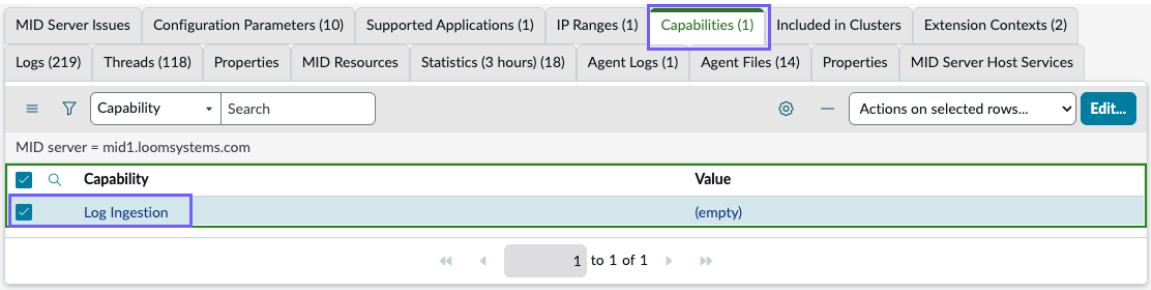

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

- Verify that you have a `vector.yaml` file on your machine, or create one.

Role required: evt_mgmt_admin

About this task

If your organization uses Vector by Datadog to handle log data from multiple sources before sending it to Health Log Analytics, the log format HLA receives is distinct from other types. The Vector Agent integration enables HLA to detect and separate transport headers from inner log messages in this format. Only the inner messages are forwarded to the source type structure for processing.

You set up integrations through the Integrations Launchpad in Service Operations Workspace, which you access from the ITOM AIOps configuration center. The AIOps configuration center is a centralized workspace for configuring and managing AIOps features from a single place. The integrations setup process reduces implementation time compared to manual data input setup in the classic interface in Health Log Analytics. For more information, see Integrations Launchpad in Service Operations Workspace for ITOM.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. From the Integrate section, under Integrations, select **Add integration**.
   The Integrations Launchpad appears.
4. In the **Browse integrations** tab, enter `Vector` in the search field.
5. Select the `Vector` integration tile.

   Note: If you start an integration setup before meeting all prerequisites, a message appears. You can cancel the setup and complete the prior requirements first. Alternatively, you can continue in draft mode and complete the requirements later. Note that you can't activate the integration until you have completed all the prerequisites.

6. On the **Provide details** form, fill in the fields.
   For a description of the fields, see the **Provide details** table in Vector Agent integration configuration fields.
7. Select **Next**.
8. In the **Set-up instruction** screen, follow the instructions on the screen to install the integration in your third-party console.
   Update the `vector.yaml` file as explained in the steps.
9. Do one of the following.
   - If you completed all the prerequisites before starting the configuration, activate the integration either with or without AI-powered capabilities. Select the appropriate option.
     - **Activate with AI** enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts.

       When the integration is activated successfully, the **Overview** tab is displayed. ServiceNow Otto for ITOM collects and analyzes log data. An AI icon indicates that AI is active.
     - **Activate** activates the integration without AI-powered mapping. The integration is activated and the **Overview** tab is displayed.

     On the Integrations Launchpad, the integration tile is available in the **Installed integrations** tab.
   - If you didn't complete all the prior requirements, select **Save draft**.

     The system saves the integration as a draft in the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**. At a later time, complete the prerequisites and activate the installation. For more information, see Activate a draft integration in Health Log Analytics.

What to do next

On the **Overview** tab, do the following:

- Use the displayed information to refine how Health Log Analytics reads the log data. For more information, see Review log streaming data and adjust integration settings in Health Log Analytics.
- Use the More options menu (⋮) to open the **Data Input Mapping**, **Source Type Structures**, or **Log Sources** pages with context from the integration. If your log data is not properly mapped, structured, or sourced, go back and adjust the configuration. If the Service Operations Workspace Log Analytics application is installed, the More options menu also provides direct access to the **Log Viewer**. Use the **Log Viewer** to review raw log messages ingested by the integration. For more information, see:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Log data auto-mapping and mapping
- Source type structure adjustment
- Review alert-related logs on the Log Viewer

If you activated the integration with AI, verify that AI correctly auto-mapped log data to service instances and components. To do this, select **View mapping** under **Log context mapping**. You can override the AI mapping by selecting a different log field from each list. For more information, see Map logs to service instances, components, source types.

Monitor log data flow and optimize integration settings

The Overview screen in Health Log Analytics provides a comprehensive view of the components in the log-processing pipeline of a specific active integration. From this screen, you can troubleshoot any streaming issues for this integration and adjust its settings if needed.

The **Overview** screen shows the log data streaming status and streaming sources of an active integration. It provides direct access to the **Data Input Mapping**, **Source Type Structures**, and **Log Sources** pages, as well as the **Log Viewer**, all with context from the current integration.

Overview screen and View menu items

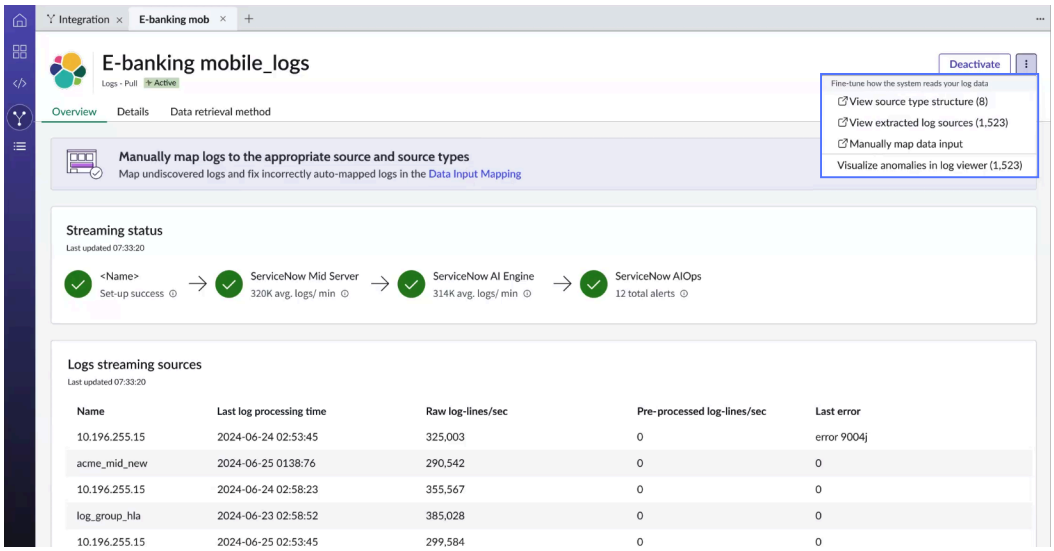

Streaming status

The **Streaming status** section on the **Overview** screen shows the status of data streaming between the log source and the MID Server and between the MID Server and the HLA engine. In addition, it displays the number of logs per minute passing through the MID Server and through the HLA engine.

Note: For the ServiceNow System Logs Retriever integration, the **Overview** screen doesn't display the MID Server streaming status because that integration doesn't run on a MID Server.

The ServiceNow AIOps component shows the total number of alerts that the HLA engine has created. These statistics are updated when the **Overview** screen loads and are automatically refreshed every minute to show live data. You can change the default auto-refresh time interval through the system property *sn_itom_integ_app.overview_page_data_input_stats_auto_refresh_interval_secon*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If data streaming fails, the integration is automatically deactivated and the **Streaming status** marks the component where the failure occurred. In addition, a banner explains the failure and either proposes steps to take to fix it or refers to ServiceNow support.

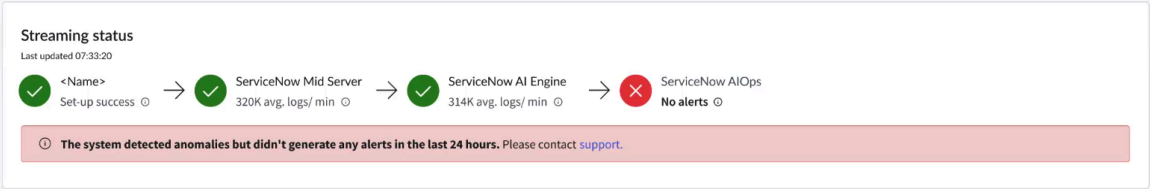

For MID-less or OpenTelemetry Protocol (OTLP) integrations, such as Amazon Data Firehose, the **Overview** screen displays the ITOM Gateway as a component in the log-processing pipeline. The MID Server component is not shown in the pipeline, because log data is sent directly from the source to the ITOM Gateway. The logs are then processed by the HLA engine to find anomalies.

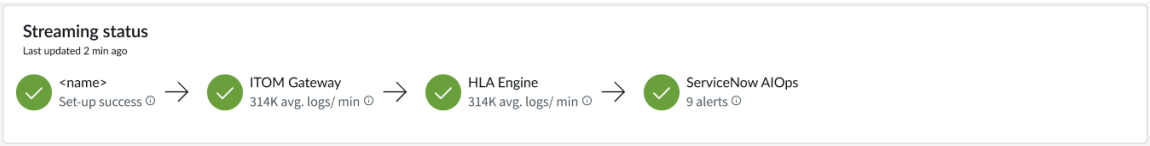

For these integrations, the **Overview** screen shows the average rate of logs per minute over the last 15 minutes passing through the ITOM Gateway and the HLA engine, similar to the metrics shown for MID-based ingestion.

Note: The ITOM Gateway component is shown only if the MID Server property *mid.hla.itom_gateway_streaming.enabled* is set to true. For more information, see MID Server properties.

When the integration isn't streaming live data from the source, a warning message is displayed. There is a 4-hour interval between the last log source time and the current time before this message is generated. You can change the default time interval through the system property *sn_itom_integ_app.overview_page_log_source_time_threshold_hours*.

Note: The warning message is displayed only when the HLA engine is up and running.

Log streaming sources

The **Log streaming sources** section displays all sources that are currently streaming log data to HLA through this integration. Depending on your integration, the **Log streaming sources** table displays the following information about the integration's log streaming sources.

Note: This table isn't available for the ServiceNow System Logs Retriever integration.

Log streaming sources

Column	Description
Name	The name of the log data source.
Status	The streaming status: Active or Not active.
Data input	The integration streaming the data to your ServiceNow instance.
MID Server	The MID Server to which the log data is streaming.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Log streaming sources (continued)

Column	Description
Last event time	The time when the integration received the latest event.
Last log processing time	The time when the last log was received or processed.
Raw log lines/sec	The average number of raw log lines that streamed to the MID Server per second in the last one-minute interval. **Note:** This value represents the number of raw log lines before preprocessing.
Pre-processed log lines/sec	The average number of preprocessed log lines that streamed to the MID Server per second in the last one-minute interval. **Note:** This value can differ from the number of raw log lines per second. For example, the difference can be a result of logs having been dropped during preprocessing.

Example

As an admin, you use the integration **Overview** page to verify continuous data flow into the ServiceNow instance. If log streaming fails, HLA doesn't receive the real-time data needed to detect anomalies or generate alerts.

For example, the **Streaming Status** for an active Elasticsearch integration may show a red circle with a white X for the MID Server component in the log-processing pipeline. This indicates that logs are not reaching the MID Server.

The **Log Streaming Sources** table shows a state of Authentication Failed or Connection Error, with an error message indicating invalid credentials. This points to an issue with the Elasticsearch credentials. Fixing the credentials and either restarting the log service or waiting for the next polling cycle, restores the connection.

The **Log Streaming Sources** table now shows a connection state of Connected or Active with successful authentication. The **Streaming Status** displays a green circle with a white check mark for the MID Server component. With log streaming restored, HLA can resume processing data and generating anomaly alerts.

Related topics

Review log streaming data and adjust integration settings in Health Log Analytics

Review log streaming data and adjust integration settings in Health Log Analytics

Review the log data streaming status and streaming sources of an active integration for Health Log Analytics on the integration's Overview screen. From this tab, you can investigate streaming issues and refine the integration settings. Leverage the displayed data to refine how HLA reads the log data by adjusting the integration's configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

For an overview of this feature, see Monitor log data flow and optimize integration settings.

Confirm that the integration has been activated on the Integrations Launchpad. For more information, see Set up integrations for Health Log Analytics from the Integrations Launchpad.

For MID-less or OpenTelemetry Protocol (OTLP) integrations, verify that the *mid.hla.itom_gateway_streaming.enabled* property is set to true in the MID Server properties. This enables displaying the ITOM Gateway component on the Overview screen.

To access the Log Viewer from the Overview tab with context from the integration, you must have the ServiceNow Service Operations Workspace Log Analytics application installed.

Role required: evt_mgmt_admin

About this task

For more information about interpreting and using the data shown in the **Overview** screen, see Monitor log data flow and optimize integration settings.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Manage installed integrations**.
   The Integrations Launchpad appears.
4. In the **Installed integrations** tab, search for an active integration and open it.
   The **Overview** screen displays, showing the integration's data streaming status and its log data sources data.

   Note: The ServiceNow System Logs Retriever integration doesn't run on a MID Server. Therefore, for this integration the **Overview** screen doesn't display the MID Server streaming status.

5. Analyze the provided data and investigate any data streaming issues.
6. **Optional:** Adjust the configuration of the integration, if needed.

   Use the options on the View menu ( ) to navigate to the relevant table or the Log Viewer.

   Note: The View menu content changes dynamically depending on whether the referenced tables contain data.

View menu options

Option	Description
Data Input Mapping	Opens the **Data Input Mapping** page. On this page, you can map sources manually if HLA didn't discover properties automatically. For more information, see Map raw log data manually.
Source Type Structure	Opens the **Source Type Structures** page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	In the Source Type Structure, log data is classified and organized based on its origin or type. You can adjust the structure and verify that the HLA AI engine extracts properties properly and classifies them correctly. For more information, see Source type structure adjustment in Health Log Analytics.
Log Sources	Opens the **Log Sources** table. This table enables you to verify that HLA has created all your log sources. If a log source is missing, you can add it manually. For more information, see Verify your log sources in Health Log Analytics.
Log Viewer	Opens the **Log Viewer**, which shows the frequency of anomalies in the log data during a specific time period. For more information, see Review alert-related logs on the Log Viewer. **Note:** The Log Viewer option appears only if the Service Operations Workspace Log Analytics application is installed.

7. When you have resolved any data streaming issues and fixed the configuration, select **Test and save**.
   The system saves the updated configuration to the database and tests the configured port, returning either success or an error. If an error is returned, fix the configuration according to the suggestions on the screen and then select **Test and save** again. Once the test is successful, continue with the next step.
8. Reactivate the integration.

Activate a draft integration in Health Log Analytics

Activate an integration that was configured and saved as a draft in Health Log Analytics (HLA).

Before you begin

Configure an integration in draft mode through the Integrations Launchpad and save it.

Role required: evt_mgmt_admin

Procedure

1. Complete all the prior requirements for the integration.
2. In the Integrations Launchpad **Installed integrations** tab, under **Waiting for your action**, locate and select the integration tile.
3. On the configuration page, select **Activate**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For Push integrations: The integration is activated and the **Overview** tab is displayed. The integration tile is available in the **Installed integrations** tab on the Integrations Launchpad.
- For all other integrations: Continue with the following steps.

4. In the pop-up window, select **Test & Save** to save the integration to the database and test connectivity.
   If an error is returned, adjust the configuration as suggested in the error message and then try to activate the integration again.
5. When the test is successful, select the appropriate activation option.
   - **Activate**

     The integration is activated and the **Overview** tab is displayed.

     Note: If your integration supports activation with AI capabilities, you can select this option to activate the integration without enabling these capabilities.

   - **Activate with AI**

     Note: This option is only available if your integration supports activation with AI capabilities.

     The integration is activated with AI capabilities enabled and the **Overview** tab is displayed. ServiceNow Otto collects and analyzes log data, and an icon indicates that it auto-maps log data to service instances and components for contextual alert generation.

   The integration tile is available in the **Installed integrations** tab on the Integrations Launchpad

Edit an installed integration in Health Log Analytics

Edit an installed integration for streaming log data to Health Log Analytics on the Integrations Launchpad. For example, you can switch to a different service instance.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Manage installed integrations**.
   The Integrations Launchpad appears.
4. In the **Installed integrations** tab, search for and select the tile for the integration you want to edit.
5. Select a tab that you want to edit.
6. Select **Deactivate** to deactivate the integration.

   Note: You must deactivate the integration before you can edit it.

7. Select **Edit**.
8. Edit the integration per your requirements.
9. Select **Test and save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics tests the MID Server connectivity, returning either success or an error. If an error is returned, make adjustments to your configuration as suggested on the screen, and then select **Test and save** again. When the test is successful, you can reactivate the integration.

10. Select **Activate**.

Deactivate a log data integration in Health Log Analytics

Deactivate an integration for streaming log data to Health Log Analytics on the Integrations Launchpad. A deactivated integration stops streaming logs to your ServiceNow instance, but is still available.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. From the Integrate section, under Integrations, select **Manage installed integrations**.
   The Integrations Launchpad appears.
4. In the **Installed integrations** tab, search for and select the integration tile for the integration you want to deactivate.
5. In the **Overview** tab, select **Deactivate**.

Result

The integration is deactivated and stops streaming log data to the instance. The deactivated integration is still available on the Integrations Launchpad in the **Installed integrations** tab, under **Waiting for your action**.

Set up data inputs in Health Log Analytics using guided setup

The Health Log Analytics guided setup provides a sequence of tasks to help you create data inputs on your ServiceNow instance. Data input configuration is an essential step in setting up the Health Log Analytics application. Using guided setup ensures that you have the minimum required setup for the data input process.

Before you begin

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

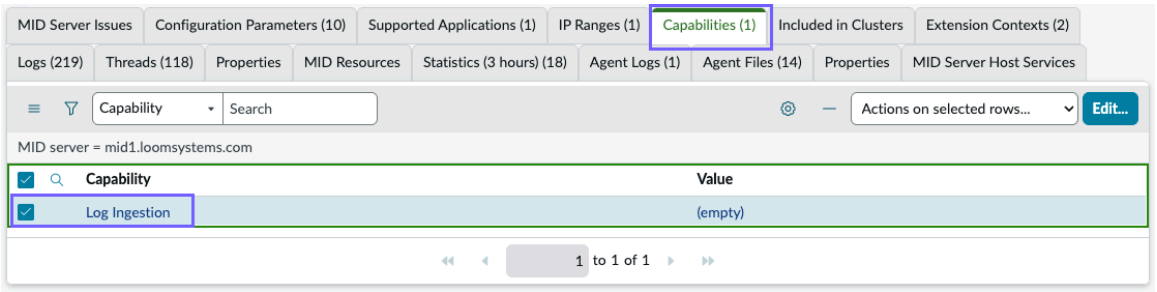

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.
- For MID Server proxy requirements, see MID Server proxy preconditions for streaming logs to Health Log Analytics.

Role required: evt_mgmt_admin

About this task

For information about using the guided setup interface, see Using guided setup.

Procedure

1. Navigate to **All** > **Guided Setup** > **ITOM Guided Setup** > **Health Log Analytics**.
2. Follow the instructions to select and set up a data input.

What to do next

Once you've created the data input connector, proceed to map your raw log data.

Related topics

Supported data inputs for Health Log Analytics

Set up data inputs in Health Log Analytics manually

Set up your Health Log Analytics data inputs for Health Log Analytics manually. Data input configuration is an essential step in setting up the Health Log Analytics application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Note: Consider using the Health Log Analytics data input guided setup, which ensures that you have the minimum required setup for the data input process. For more information, see Set up data inputs in Health Log Analytics using guided setup.

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

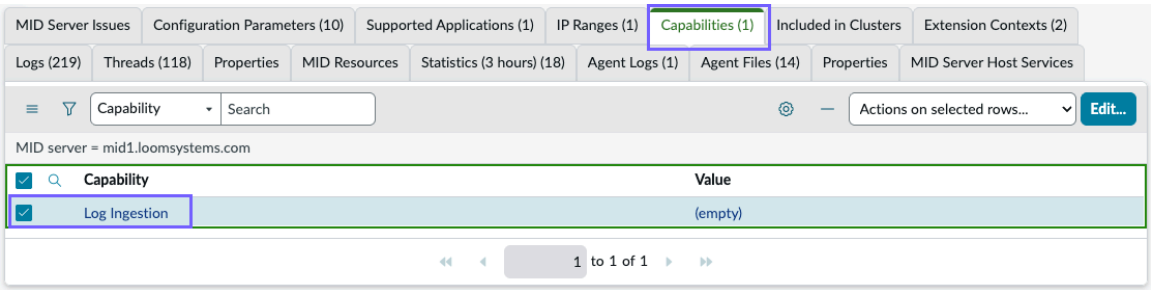

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.
- For MID Server proxy requirements, see MID Server proxy preconditions for streaming logs to Health Log Analytics.

Role required: evt_mgmt_admin. For the ServiceNow System Logs data input: admin.

Procedure

1. Setup a data input manually by performing the relevant procedure described in the product documentation.

Data Inputs

Data Input	Description
Agent Client Collector	The data input streams log messages to your ServiceNow instance using the ServiceNow Agent Client Collector. This data input is supported for use with the Agent Client Collector Log Analytics

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data Input	Description
	application, available from the ServiceNow Store.
Amazon CloudWatch	The data input streams log data from Amazon CloudWatch to your ServiceNow instance.
Amazon S3 Bucket	The data input streams log data from Amazon S3 (Simple Storage Service) buckets to your ServiceNow instance.
Apache Kafka	The data input streams log data from Apache Kafka to your ServiceNow instance.
Cribl	The data input to enables Health Log Analytics to process Cribl log messages streaming into your ServiceNow instance.
Edge Delta	The data input enables Health Log Analytics to process Edge Delta log messages streaming into your ServiceNow instance.
Elasticsearch	The data input pulls log data from Elasticsearch indexes into your instance.
GCP PubSub	The data input receives log messages that are published to a Google Cloud Pub/Sub topic and streams them to your ServiceNow instance.
Microsoft Azure Event Hubs	The data input streams events from Microsoft Azure Event Hubs to your ServiceNow instance.
Microsoft Azure Log Analytics	The data input streams log data from Microsoft Azure Log Analytics to your ServiceNow instance.
MID Server	The data input collects MID Server log files and streams them to your instance.
REST API	The data input streams log data to your ServiceNow instance in JSON format.
Rsyslog or Beats	The data input streams log data into your instance using Rsyslog or Beats.
ServiceNow System Logs Retriever	The data input streams log data from the ServiceNow System Log table to the Health Log Analytics AI engine. **Note:** Only a single ServiceNow System Logs Retriever data input can exist in the system, and only users with the admin role can create and configure it. This data input doesn't run on a MID Server.
Splunk	The data input streams log data into your instance using Splunk.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data Input	Description
Splunk Polling	The data input periodically pulls log data from Splunk by using a query.
TCP	The data input sends raw log messages to your instance directly over a TCP/SSL socket.
UDP	The data input streams raw log messages to your ServiceNow instance directly over a UDP socket.
Vector Agent	The data input enables Health Log Analytics to process log messages that are streaming into your ServiceNow instance via a Vector Agent.

Note: Testing the connection to the MID Server at the end of the data input setup procedure ensures that your data input is configured correctly. You can only publish a data input configuration when the connection between the MID Server and the data repository has been established.

2. Identify and address streaming issues to confirm that the data input is streaming log data to the MID Server from all sources.
   For more information, see Identify and resolve a log streaming issue in Health Log Analytics.
3. **Optional:** Edit raw log data before Health Log Analytics maps and structures it.
   For more information, see Edit your raw log data before processing.
4. Determine how Health Log Analytics handles raw log data that is streaming into your instance. By default, every incoming log line is auto-mapped to the correct tag. If properties aren't discovered automatically, map the data input sources manually by defining a JavaScript function. For more information, see Map the raw data.
5. **Optional:** Tweak the source type structure to make sure that Health Log Analytics extracts and classifies all properties correctly.
   For more information, see Refine the source type structure in Health Log Analytics.
6. **Optional:** Perform additional data input setup tasks.
   For more information, see Additional data input setup tasks in Health Log Analytics.

Related topics

Supported data inputs for Health Log Analytics

Configuring data inputs for Health Log Analytics manually

Configure the data input process manually in Health Log Analytics. Data input configuration is an essential step in setting up the Health Log Analytics application.

Configure an Amazon CloudWatch data input in Health Log Analytics manually

Set up a data input for streaming log data from Amazon CloudWatch to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled.
  For more information, see MID Server system requirements.

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Amazon CloudWatch data input.
4. On the form, fill in the fields.
   For a description of the fields, see Amazon CloudWatch data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   For a description of the fields, see Amazon CloudWatch data input configuration fields. For information about configuring the advanced settings later, see Configure advanced settings for Amazon CloudWatch data inputs in Health Log Analytics manually.
6. On the **Query settings** tab, fill in the fields.
   For a description of the fields, see Amazon CloudWatch data input configuration fields.
7. On the **Transport** tab, fill in the fields.
   For a description of the fields, see Amazon CloudWatch data input configuration fields. These fields are required.
8. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
9. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

10. Select **Publish** to publish the data input to the MID Server.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data from Amazon CloudWatch to your ServiceNow instance.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Amazon CloudWatch data inputs in Health Log Analytics manually

Configure an Amazon S3 data input in Health Log Analytics manually

Set up a data input for streaming log data from Amazon S3 (Simple Storage Service) buckets to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

  Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Amazon S3 data input.
4. On the form, fill in the fields.
   For a description of the fields, see Amazon S3 data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a description of the fields, see Amazon S3 data input configuration fields. For information about configuring the advanced settings later, see Configure advanced settings for Amazon S3 data inputs in Health Log Analytics manually.

6. On the **Query settings** tab, fill in the fields.
7. On the **Transport** tab, fill in the field.
   For a description of the fields, see Amazon S3 data input configuration fields.
8. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
9. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

10. Select **Publish** to publish the data input to the MID Server.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data from Amazon S3 buckets to your ServiceNow instance.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Amazon S3 data inputs in Health Log Analytics manually

Configure an Apache Kafka data input in Health Log Analytics manually

Set up a data input for streaming log data from Apache Kafka to your ServiceNow instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

  Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Apache Kafka data input.
4. On the form, fill in the fields.
   For a description of the fields, see Apache Kafka data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   For a description of the fields, see Apache Kafka data input configuration fields. For information about configuring the advanced settings later, see Configure advanced settings for Apache Kafka data inputs in Health Log Analytics manually.
6. On the **Query Settings** tab, fill in the field.
   For a description of the fields, see Apache Kafka data input configuration fields.
7. On the **Transport** tab, fill in the fields.
   For a description of the fields, see Apache Kafka data input configuration fields.
8. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
9. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
- If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

Note: You can only publish the data input configuration when the connection is created successfully.

Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

10. Select **Publish** to publish the data input to the MID Server.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data from Apache Kafka to your ServiceNow instance.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Apache Kafka data inputs in Health Log Analytics manually

Configure a Cribl data input in Health Log Analytics

Configure a dedicated Cribl data input to enable Health Log Analytics to process Cribl log messages streaming into your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

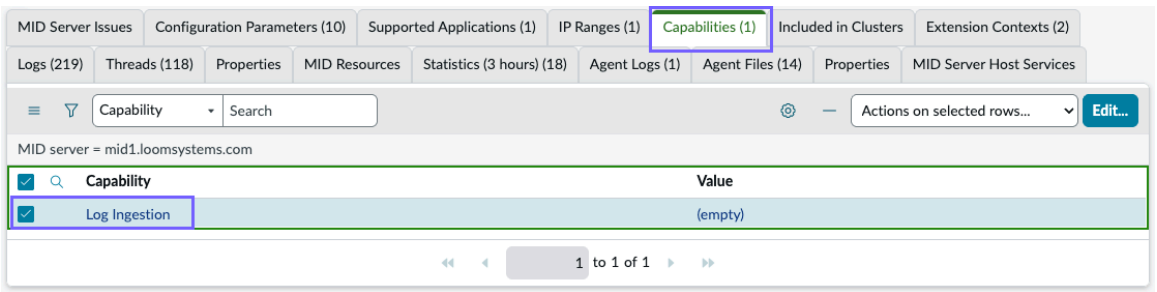

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

If your organization uses Cribl for filtering and routing large volumes of log data from various sources, the log format Health Log Analytics receives is distinct from other types. The Cribl data input enables HLA to detect and separate transport headers from inner log messages in this format. It forwards only the inner message to the source type structure for processing.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Cribl data input.
4. On the **Getting Started** tab, fill in the form fields and then select **Next**.
   For a description of the fields, see Cribl data input configuration fields.
5. On the **Config** tab, fill in the form fields and then select **Next**.
   For a description of the fields, see Cribl data input configuration fields.
6. On the **Process** tab, select the required route and then select **Next**.

   For more information, see Cribl data input configuration fields.

   Note: For troubleshooting issues related to this step, see the Cribl - Setting up Cribl data input | Troubleshooting setup & data ingestion issues [KB0558611] article in the Now Support Knowledge Base.

   The **Finish** tab opens.
7. On the Cribl instance, verify that the configuration of the data input is correct.
   You can use the Cribl commands displayed on the **Finish** tab as needed.
8. When the configuration is correct, return to the data input configuration **Finish** tab and select **Publish** to publish the data input.

Result

The data input configuration process is complete. The data input starts streaming log data to your ServiceNow instance.

For resolving issues related to the ingestion of log data from Cribl, see the Cribl - Setting up Cribl data input | Troubleshooting setup & data ingestion issues [KB0558611] article in the Now Support Knowledge Base.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure an Edge Delta data input in Health Log Analytics manually

Set up an Edge Delta data input to enable Health Log Analytics to process Edge Delta log messages streaming into your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

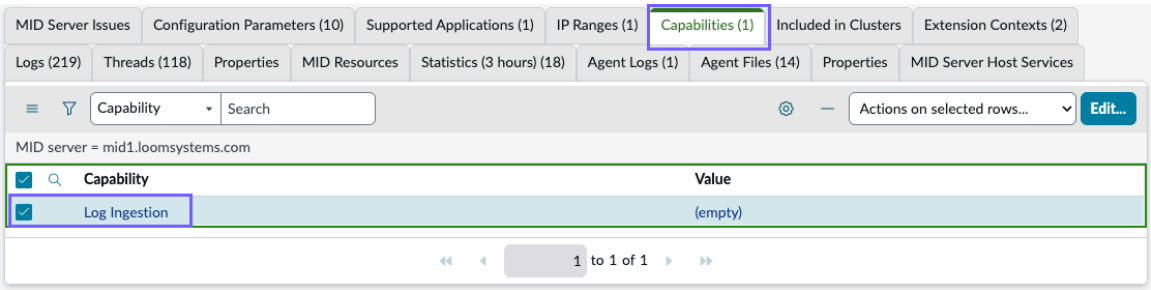

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

About this task

If your organization uses Edge Delta to handle large volumes of log data from multiple sources, the log format is distinct from other types. Edge Delta sends data to Health Log Analytics. The Edge Delta data input enables HLA to detect and separate transport headers from inner log messages. It forwards only the inner message to the source type structure for processing.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Edge Delta data input.
4. On the form, fill in the fields.
   For a detailed description of the fields, see Edge Delta data input configuration fields.
5. Select **Save** to save your configuration.
6. Select **Publish** to publish the data input.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table. The data input starts streaming log data to your ServiceNow instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Configure an Elasticsearch data input in Health Log Analytics manually

Set up a data input for streaming log data from Elasticsearch indices to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

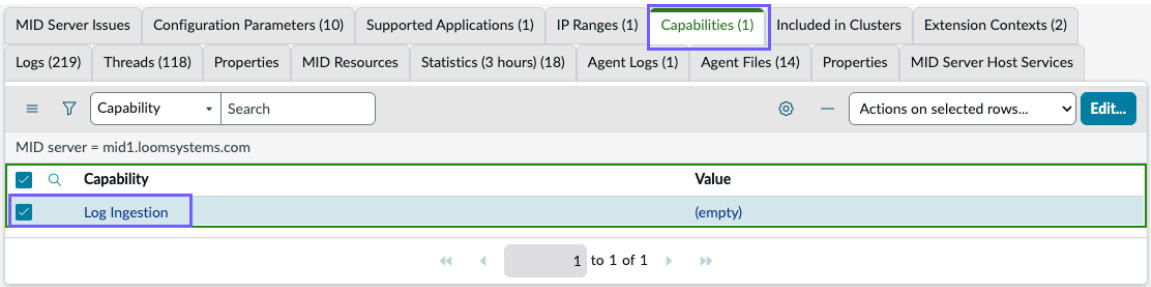

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Health Log Analytics supports Elasticsearch versions above 7.10.2 and below 8.18.2.

Note: For advanced information about streaming log data from Elasticsearch indices to your instance, see the Stream logs using Elasticsearch data input - Advanced guide [KB1080162] article in the Now Support knowledge base.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Elasticsearch data input.
4. On the form, fill in the fields.
   For a description of the fields, see Elasticsearch data input configuration fields.
5. **Optional:** Select **Advanced** to set advanced configuration fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a description of the fields, see Elasticsearch data input configuration fields. For information about configuring the advanced settings later, see Configure advanced settings for Elasticsearch data inputs in Health Log Analytics manually.

6. On the **Transport** tab, fill in the fields.
   For a description of the fields, see Elasticsearch data input configuration fields.
7. On the **Query Settings** tab, fill in the fields.
   For a description of the fields, see Elasticsearch data input configuration fields.
8. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
9. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. This field displays only when a streaming error has occurred.

     Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

10. Select **Publish** to publish the data input to the MID Server.

Result
The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data from Elasticsearch indices to your ServiceNow instance.

For more information about streaming logs using the Elasticsearch data input, see the Stream logs using Elasticsearch data input - Advanced guide [KB1080162] article in the Now Support Knowledge Base.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next
Make sure that the data input is streaming data.

Note: If you experience permissions-related issues with streaming log data from Elasticsearch, refer to the Granting privileges for data streams from Elasticsearch [KB0967366] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Configure advanced settings for Elasticsearch data inputs in Health Log Analytics manually

Configure a GCP PubSub data input in Health Log Analytics manually

Set up a data input for receiving log messages that were published to a Google Cloud Platform (GCP) Pub/Sub topic and streaming them to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

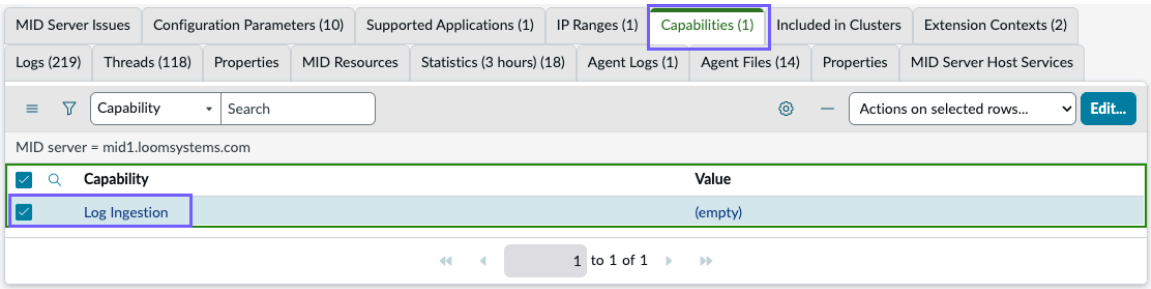

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the GCP PubSub data input type.

   Note: The selected data input type complements the passive data input (listener). For more information, see Supported data inputs for Health Log Analytics.

4. On the form, fill in the fields.
   For a description of the fields, see GCP PubSub data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   For a description of the fields, see GCP PubSub data input configuration fields.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you changed the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Result
The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts to receive log messages from the Google Cloud Pub/Sub topic and streams them to your ServiceNow instance via the MID Server.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

Configure a Microsoft Azure Event Hubs data input in Health Log Analytics manually

Set up a data input for streaming events from Microsoft Azure Event Hubs to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

  Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Microsoft Azure Event Hubs data input.
4. On the form, fill in the fields.
   For a description of the fields, see Microsoft Azure Event Hubs data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   For a description of the fields, see Microsoft Azure Event Hubs data input configuration fields.
   For information about configuring the advanced settings later, see Configure advanced settings for Microsoft Azure Event Hubs data inputs in Health Log Analytics manually.
6. On the **Query settings** tab, fill in the fields.
   For a description of the fields, see Microsoft Azure Event Hubs data input configuration fields.
7. On the **Transport** tab, fill in the fields.
   For a description of the fields, see Microsoft Azure Event Hubs data input configuration fields.
8. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
9. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

10. Select **Publish** to publish the data input to the MID Server.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data from Microsoft Azure Event Hubs to your ServiceNow instance.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Microsoft Azure Event Hubs data inputs in Health Log Analytics manually

Configure a Microsoft Azure Log Analytics data input in Health Log Analytics manually

Set up a data input for streaming log data from Microsoft Azure Log Analytics to your ServiceNow instance. The data input points the Health Log Analytics AI engine to a data source in your Microsoft Azure Log Analytics account.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

> **Important:** Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Microsoft Azure Log Analytics data input.
4. On the form, fill in the fields.
   For a description of the fields, see Microsoft Azure Log Analytics data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   For a description of the fields, see Microsoft Azure Log Analytics data input configuration fields. For information about configuring the advanced settings later, see Configure advanced settings for Microsoft Azure Log Analytics data inputs in Health Log Analytics manually.
6. On the **Query settings** tab, fill in the fields.

   Health Log Analytics uses the values set in these fields to generate the query for receiving log data from Microsoft Azure Log Analytics.

   For example:

```
{"query": "ContainerLog | where TimeGenerated > %s | take 500",
 "project": ["TimeGenerated", "LogEntry", "LogEntrySource"]}
```

   where:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- `ContainerLog`: The data source name.
- `TimeGenerated`: The event time property name.
- `500`: The maximum number of rows retrieved in each query.
- `["TimeGenerated", "LogEntry", "LogEntrySource"]`: The event time and the comma-separated list of column names to return.

Alternatively, you can configure a custom query in the **Log query** field. The value you set in this field overrides the values in all other fields on the **Query settings** tab except the **From** field. If the **Log query** field is empty, Health Log Analytics generates the query using the values set in the other fields on the tab.

For a description of the **Query settings** fields, see Microsoft Azure Log Analytics data input configuration fields.

7. On the **Transport** tab, fill in the fields.
   For a description of the fields, see Microsoft Azure Log Analytics data input configuration fields.
8. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
9. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

10. Select **Publish** to publish the data input to the MID Server.

Result
The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data from Microsoft Azure Log Analytics to your ServiceNow instance.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Microsoft Azure Log Analytics data inputs in Health Log Analytics manually

Configure a MID Server data input in Health Log Analytics manually

Set up a data input for collecting and streaming MID Server log messages to your ServiceNow instance.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

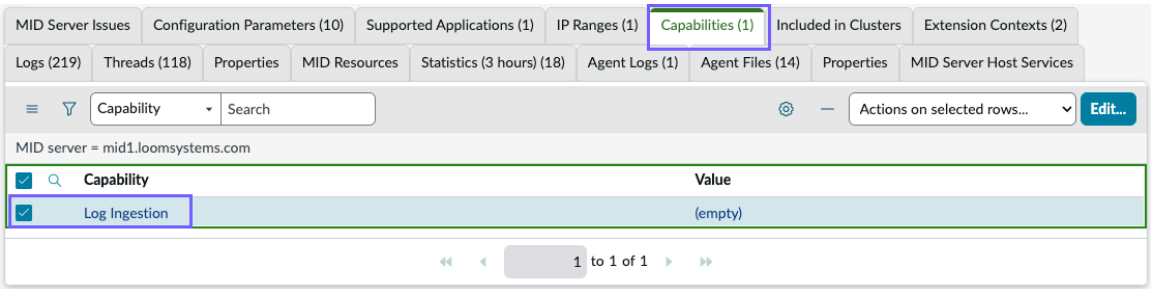

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Role required: evt_mgmt_admin

About this task

This procedure describes how to create a MID Server data input from the Data Inputs page. To configure the MID Server data input automatically, click the related link on the MID Server form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

- When you create a MID Server data input, all logs are automatically mapped to the predefined mid_server_log source type. This source type includes preconfigured mapping settings to maximize the value of monitoring your MID Server logs. To review the mapping script, navigate directly from the data input configuration form to the mapping form by selecting **Go to Mapping** under Related Links. Review the source type by navigating to Source Type Structures and select the mid_server_log source type.
- For optimal performance, debug MID Server log messages are not processed by default. If you want to send debug messages to Health Log Analytics for processing, add the following to your MID Server properties: mid.hla.mid_data_input.ignore_debug_level_logs = false.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the MID Server data input.
4. On the form, provide a description of the data input and the name of the MID Server from which the logs are streamed.
   For a description of the fields, see MID Server data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   On the form, fill in the fields. For a description of the fields, see MID Server data input configuration fields. For information about configuring the advanced settings later, see Configure advanced settings for MID Server data inputs in Health Log Analytics manually.
6. Select **Submit**.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts collecting and streaming MID Server log messages to your ServiceNow instance.

For a description of how this data input can help you detect and resolve emerging issues in your organization's ServiceNow instance, see Use Case: Proactive monitoring of your ServiceNow instance in Health Log Analytics.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Configure a REST API data input in Health Log Analytics manually

Set up a REST API data input for streaming log data to your ServiceNow instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

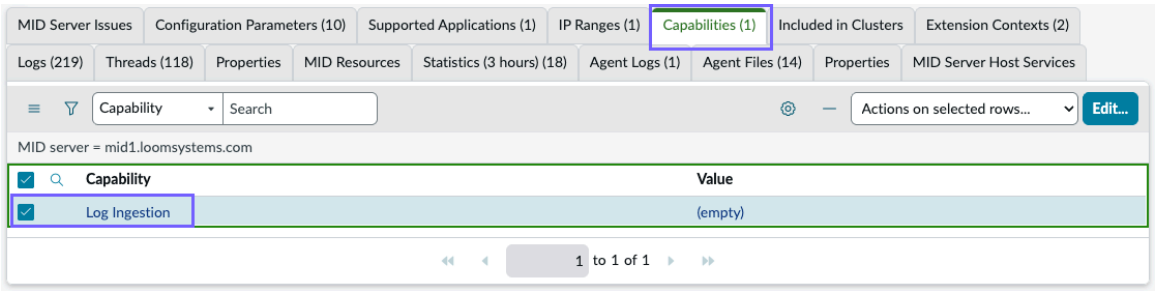

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

On the MID Web Server Context form, in the **Execute on** field, select *Specific MID Server* and not the cluster option. In the **MID Server** field, select the specific MID Server to which the log data is pulled.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the REST API data input.
4. On the form, fill in the fields.
   For a description of the fields, see REST API data input configuration fields.
5. Select **Save**.
   - Health Log Analytics adds the data input record to the **Data Inputs** table.
   - The system automatically generates the REST API data input URL for sending logs to the MID Server. The generated URL will be in this format: `http://(MID_SERVER_IP):(MID_WEB_SERVER_PORT)/ap1/mid/hla/raw`. HLA replaces `(MID_SERVER_IP)` and `(MID_WEB_SERVER_PORT)` with the values you configured.
6. Copy the generated URL.
7. Instruct your API client to send log messages to the MID Server.

   a. In the API client, configure a request with the following parameters:
      - HTTP Method: POST.
      - URL: The URL you copied in the previous step.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Authentication: The required authentication credentials. For example, Basic authentication (username and password).

  Note: The authentication method must be the same as the method configured on the MID Web Server.

- Headers: Set the `Content-Type` header to match your payload format.
  - For JSON: `application/json`.
  - For raw text: `text/plain`.
- Body: The logs payload in JSON or raw text.

b. Send the request.
A successful response indicates that the MID Server has received the payload. When the MID Server has processed the log messages, the data input streams them to your ServiceNow instance using the REST protocol.

Configure a Rsyslog, Filebeat, or Winlogbeat data input in Health Log Analytics manually

Set up a data input for streaming log messages to your ServiceNow instance using an Rsyslog, Filebeat, or Winlogbeat agent.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

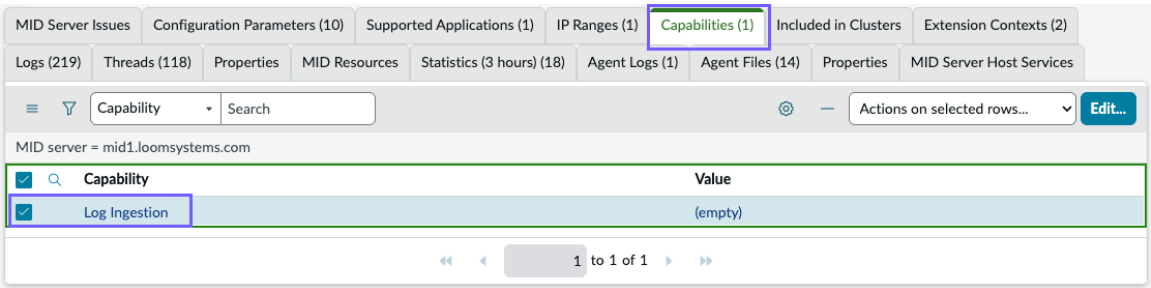

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the data input type to create from the available data input types described in the table.

   Note: The selected data input type complements the passive data input (listener). For more information, see Supported data inputs.

   Data input types

Type	Description
Rsyslog	Streams log messages from UNIX-based servers to the ServiceNow AI engine using the Rsyslog agent.
Linux using Filebeat	Streams system log messages and local files from Linux servers to the instance using the Filebeat agent.
Windows Application Logs using Filebeat	Streams local files from Microsoft Windows devices to the ServiceNow instance using the Filebeat agent.
Windows OS using Winlogbeat	Streams Windows event logs to the ServiceNow instance using the Winlogbeat agent.

4. On the **Getting Started** tab, fill in the form.

   For a description of the fields, see Rsyslog, Filebeat, or Winlogbeat data input configuration fields.

   Note: When creating a data input for Linux using Filebeat, you can select a content pack from the **Content pack** drop-down. The content pack contains default source types and mapping script templates that save you the time it takes to create them from scratch. For more information, see Health Log Analytics content packs for quicker time to value.

5. If the Rsyslog, Filebeat, or Winlogbeat agent has not already been installed, download and install it from the **Installation** tab.
   For installation instructions, refer to the product documentation of the respective agent.

   Note: Make sure that you're running the latest version of the agent. Earlier versions will work, but with limited functionality.

6. On the **Tagging and Binding** tab, assign logs to an service instance in the Configuration Management Database (CMDB) to enable the service to correlate the log data and enable the system to do root cause analysis.

   a. For each source, configure the path and service instance for the logs to be streamed.

   Note: By default, only the required fields **Path** and **Service instance** appear.

   For a description of the fields, see Rsyslog, Filebeat, or Winlogbeat data input configuration fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. If you want to ship multiline logs using Filebeat, configure the properties that control how Filebeat handles messages that span multiple lines of text.

Field	Description
Match	Specifies how Filebeat combines matching lines into an event.
Negate	Defines whether the pattern identified in the log lines is negated.
Regex	Specifies the regular expression to match.

Note: Health Log Analytics currently doesn't support multiline properties for Rsyslog.

c. Optional: Define additional log paths to enable the data input to stream log types from multiple paths.
Do the following for each additional log path:

i. Insert a new row.

ii. Configure the log path.

iii. Choose a service instance.

iv. (Optional) Choose a component and a source type.

Note: This option is not available or necessary when using Winlogbeat, because Health Log Analytics streams the Windows event logs.

7. On the **Finish** tab, complete the configuration for your data input type.

- Rsyslog:

  a. Download the configuration file and install it on the endpoint device, in the `/etc/rsyslog.d/rsyslog.conf` directory.

  Note: If you are using the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store , do the following instead:

  i. On the endpoint device, install the configuration file in the `/etc/rsyslog.d/` directory.

  ii. Create a spool directory by running the `sudo mkdir -p/var/spool/rsyslog` command.

  b. Validate the configuration by running the `rsyslogd -N1` command and verify the output.

  c. Restart Rsyslog by running the `sudo systemctl restart rsyslog` command.

  d. Verify the output. If it contains errors, check the `/var/log/messages` system log file for error messages and fix the errors.

- Linux using Filebeat:

  a. Download the configuration file and install it on the endpoint device, in the `/etc/filebeat/` directory.

  b. Start the agent service by running the `sudo service filebeat start` command.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The generated configuration ignores files that were last changed more than six hours ago. If needed, you can change this setting in the configuration.

c. Restart the agent service by running the appropriate command.

- Windows using Beats (Filebeat or Winlogbeat):

a. Download the configuration file and install it on the endpoint device, in the `C:\Program Files\` directory.

b. Start the agent service by running the appropriate command in PowerShell.

- Filebeat: `PS > Start-Service filebeat`
- Winlogbeat: `PS > Start-Service winlogbeat`

Note: The generated configuration ignores files that were last changed more than six hours ago. If needed, you can change this setting in the configuration.

c. Restart the agent service by running the appropriate command.

8. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
9. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. This field displays only when a streaming error has occurred.

     Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

10. Select **Publish** to publish the data input to the MID Server.

Result
The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data to your ServiceNow instance.

Note: The configuration file Health Log Analytics generates for new Filebeat data inputs is compatible with Filebeat version 7.14 and above. The minimum Filebeat version that Elasticsearch currently supports is 7.17, which is nearing its end of life (EOL).

Existing Filebeat data input configuration files in HLA are compatible with Filebeat versions up to 8.x. For Filebeat version 9.0 and above, either migrate from log input to filestream input or generate a new configuration file. For more information on the migration process, refer to the Filebeat documentation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Rsyslog, Splunk, TCP data inputs

Configure advanced settings for Beats data inputs in Health Log Analytics manually

Configure a ServiceNow Log Export data input in Health Log Analytics manually

Set up a data input for monitoring ServiceNow instance node logs from both Java code and JavaScript in Health Log Analytics (HLA). This data input provides a complete view of your instance status, including server context.

Before you begin

Note: Only a single ServiceNow Log Export data input can exist in the system. This data input doesn't run on a MID Server.

- Verify that the Log Export Service (LES) is installed on your ServiceNow instance. LES is a ServiceNow Platform tool under a different license than Health Log Analytics that enables you to export your system logs seamlessly into HLA. For more information, see Log Export Service (LES).
- Verify that the Hermes Messaging Service is up and running:
  1. Navigate to **All** > **Hermes Messaging Service** > **Diagnostics**.
  2. Ensure that the status icon is green, signifying that the service is active.

     If it isn't, contact ServiceNow support.
- Create an instance-signed PKI Certificate:
  1. Navigate to **All** > **Certificate Generator** > **Instance PKI Certificate Generator**.
  2. In the **Certificate Password** field, set a password to use for your certificate.
  3. Select **Generate**.

     The system generates the certificate.
  4. Download the `keystore` and `truststore` files to your machine.
- Create a Basic Auth credential and attach it to a credential alias.
  1. Navigate to **All** > **IntegrationHub** > **IntegrationHub Usage** > **Connections & Credentials** > **Credentials**
  2. Select **New**.
  3. Select **Basic Auth Credentials**.
  4. In the **Password** field, enter the certificate password.
  5. In the **Credential alias** field, attach the new Basic Auth credential to a new credential alias.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The ServiceNow Log Export data input streams system node logs from Java code and JavaScript into HLA. This provides a complete view of your ServiceNow instance status with server context. This enhanced visibility facilitates improved monitoring and generates more meaningful alerts to help you identify instance issues, detect security threats, optimize performance, and maintain regulatory compliance.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose **ServiceNow Log Export Hermes Consumer**.
4. Select the **Manage Attachments** icon ( ) and attach the `keystore` and `truststore` files to the form.
5. On the form, fill in the fields.
   For a description of the fields, see ServiceNow Log Export data input configuration fields.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record.

Alerts generated by this data input have the `glide_node_log alerts` component alert tag in the **Express List Active alerts** page.

Configure a ServiceNow System Logs Retriever data input in Health Log Analytics manually

Set up a data input for streaming log data from the ServiceNow System Log table to the HLA engine (aka Occultus).

Before you begin

Note: Only a single ServiceNow System Logs Retriever data input can exist in the system. This data input doesn't run on a MID Server.

Role required: admin

About this task

The HLA store app installation automatically creates the ServiceNow System Logs Retriever data input. Use this procedure if you need to recreate it manually.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the ServiceNow System Logs Retriever data input.
4. On the form, fill in the fields.
   For a description of the fields, see ServiceNow System Logs Retriever data input configuration fields.
5. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record.

The data input starts streaming ServiceNow log data from the System Log table to the Health Log Analytics AI engine, based on the configured filters. Admin users can set filters to query the System Log table. Operators can monitor the logs and view the alerts that Health Log Analytics generates from them.

For a description of how this data input can help you detect and resolve emerging issues in your organization's ServiceNow instance, see Use Case: Proactive monitoring of your ServiceNow instance in Health Log Analytics.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Configure a Splunk data input in Health Log Analytics manually

Set up a data input for streaming log messages to your ServiceNow instance using a Splunk heavy forwarder.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

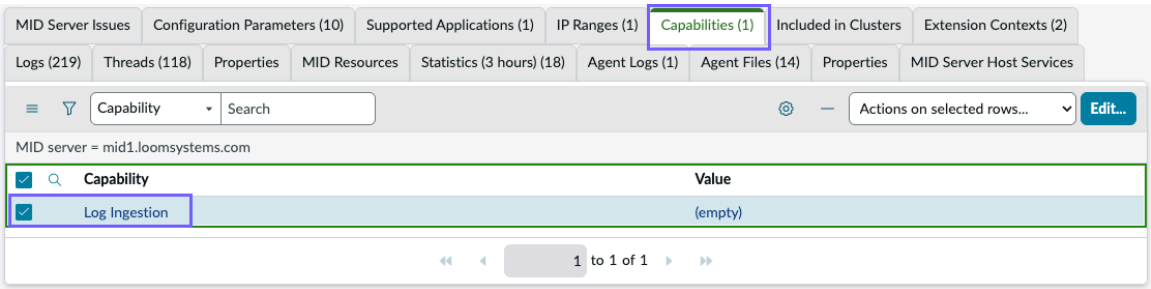

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Configure Splunk to forward logs to your ServiceNow instance using Syslog.
- The configuration of this data input assumes the existence of an environment variable named $SPLUNK_HOME. In Unix-like environments, this variable typically points to `/opt/splunk`.

> **Note:** The Windows environment uses the same directory structure, but with backslashes (\).

Role required: evt_mgmt_admin

About this task

This setup procedure is for streaming logs to your instance using a Splunk heavy forwarder. If you can't use a heavy forwarder, you can use a Splunk Universal Forwarder instead. For more information, see the Splunk Universal Forwarder as a Shipping method [KB0961378] article in the Now Support Knowledge Base.

Starting in the Yokohama family release, you can use new Splunk data inputs to ingest data in the preprocessed ("cooked") log-forwarding format that Splunk uses by default. In cooked mode, the Splunk forwarder embeds configuration details, such as host, source type, source, and other settings into the log data. Ingesting the data into HLA in this format verifies that each log line retains all the relevant contextual information. If you are using the cooked data option in HLA, there is no need to edit the `props.conf` and `transforms.conf` files during Splunk data input configuration.

> **Note:** All Splunk configuration files are located in the $SPLUNK_HOME/`etc/system/local/` folder. If a configuration file that you need to modify doesn't exist, create it and save it to this folder.

> **Note:** A MID Server that is down can cause a blockage in your Splunk pipeline. A full processing queue does not affect the pipeline.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Splunk data input for streaming logs via a Splunk heavy forwarder or universal forwarder.
4. On the **Getting Started** tab, fill in the form fields.
   For a description of the fields, see Splunk data input configuration fields.
5. On the **Outputs.conf** tab, add the following stanzas to the `outputs.conf` file to make the shipper forward log data over the selected transport protocol on the selected port.

   Select **Next**.

   > **Note:** If you have already configured outputs, merge these lines with your existing configuration.

   - Forwarding over TCP:

     > **Note:** Use the first stanza only if you have not already configured a tcpout stanza. The second stanza is required for forwarding to Health Log Analytics over TCP.

     ```
     [tcpout]
     indexAndForward = 1
     ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
defaultGroup = nothing

[tcpout:health_log_analytics]
server = ip-172-31-17-121.us-east-2.compute.internal:5000
sendCookedData = false
compressed = false
```

- Forwarding over UDP:

  Note: Use the first stanza only if you have not already configured a syslog stanza. The second stanza is required for forwarding to Health Log Analytics over UDP.

```
[syslog]defaultGroup = nothing

[syslog:health_log_analytics]
server = ip-172-31-17-121.us-east-2.compute.internal:5000
```

6. On the **Props.conf** tab, edit the `props.conf` file, and then select **Next**.

   Note: If you selected the **Use Cooked Data** option on the **Getting Started** tab, there is no need to edit the `props.conf` file.

   a. Modify existing stanzas or add stanzas to mark source types, service instances, and hosts for forwarding to Health Log Analytics.

      Note: For best results, mark only source types for forwarding.

      When adding stanzas, use the following name formats:

      - Source types: [<source type>]. For example: [syslog]
      - Sources (not recommended): [source::<source>]. For example, [source::myApp]
      - Hosts (not recommended): [host::<host>]. For example, [host::10.9.8.7]

   b. Add the following line to the end of each stanza that you want to forward to Health Log Analytics over either TCP or UDP.

      - Forwarding over TCP:

```
TRANSFORMS-cloned_data = clone_for_hla_store_sourcetype,
 clone_for_hla
```

      - Forwarding over UDP:

```
TRANSFORMS-cloned_data = clone_for_hla_store_sourcetype,
 clone_for_hla_udp
```

      This line applies the CLONE_SOURCETYPE transform on the data to prevent manipulation needed for Health Log Analytics processing from affecting your existing data pipeline. For example, to send all logs from the source type "syslog" to Health Log Analytics:

```
[syslog]
                                    #existing configuration
 goes here
                                    TRANSFORMS-cloned_data =
 clone_for_hla_store_sourcetype, clone_for_hla_udp
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Add the following stanza to apply all the relevant transforms needed for Health Log Analytics processing.

Note: Splunk enables you to anonymize sensitive data on the cloned source type for the selected protocol. For more information, see the "Anonymize data" section in the Splunk documentation.

- Forwarding over TCP:

```
[send_to_hla_tcp]
TRANSFORMS-health_log_analytics =
 health_log_analytics_add_sourcetype,
 health_log_analytics_add_host,
 health_log_analytics_add_source,
 health_log_analytics_add_syslog5424,
 health_log_analytics_add_index,
 health_log_analytics_add_syslogHost,
 health_log_analytics_add_time
```

- Forwarding over UDP:

```
[send_to_hla_udp]
TRANSFORMS-health_log_analytics =
 health_log_analytics_add_sourcetype,
 health_log_analytics_add_host,
 health_log_analytics_add_source,
 health_log_analytics_add_syslog5424,
 health_log_analytics_add_index,
 health_log_analytics_add_syslogHost,
 health_log_analytics_add_time
```

7. On the **Transforms.conf** tab, add the following stanzas to the `transforms.conf` file, and then select **Next**.

Note: If you selected the **Use Cooked Data** option on the **Getting Started** tab, there is no need to edit the `transforms.conf` file.

The third stanza clones the logs for further manipulation without affecting your existing indexing. The remaining stanzas add the information needed to enable correct Health Log Analytics processing.

Note: You can obfuscate sensitive data by adding a transform here and then modifying the stanza of the cloned source type in the `props.conf` file.

```
[accepted_keys]
#Custom field for preserving sourcetype
hla_sourcetype_preservation=_hla_sourcetype

#Store sourcetype in a custom field, since CLONE_SOURCETYPE
 overwrites it
[clone_for_hla_store_sourcetype]
SOURCE_KEY = MetaData:Sourcetype
REGEX = ^sourcetype::(.+)$
FORMAT = hla_sourcetype::$1
DEST_KEY = _hla_sourcetype
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
[clone_for_hla]
REGEX=.
DEST_KEY = _TCP_ROUTING
FORMAT = health_log_analytics
CLONE_SOURCETYPE=send_to_hla_tcp

#Only used in case of UDP forwarding
[clone_for_hla_udp]
REGEX=.
DEST_KEY = _SYSLOG_ROUTING
FORMAT = health_log_analytics
CLONE_SOURCETYPE=send_to_hla_udp

#Add metadata to the log message, since metadata is lost when
 forwarding externally
[health_log_analytics_add_sourcetype]
SOURCE_KEY = _hla_sourcetype
REGEX = ^hla_sourcetype::(.+)$
FORMAT = sourcetype="$1"] $0
DEST_KEY = _raw

[health_log_analytics_add_host]
SOURCE_KEY = MetaData:Host
REGEX = ^host::(.+)$
FORMAT = host="$1" $0
DEST_KEY = _raw

[health_log_analytics_add_source]
SOURCE_KEY = MetaData:Source
REGEX = ^source::(.+)$
FORMAT = source="$1" $0
DEST_KEY = _raw

[health_log_analytics_add_syslog5424]
REGEX=.
FORMAT = - - [sdid@1234 $0
DEST_KEY = _raw

[health_log_analytics_add_index]
SOURCE_KEY = _MetaData:Index
REGEX = ^(.+)$
FORMAT = $1 $0
DEST_KEY = _raw

[health_log_analytics_add_syslogHost]
SOURCE_KEY = MetaData:Host
REGEX = ^host::(.+)$
FORMAT = $1 $0
DEST_KEY = _raw

[health_log_analytics_add_time]
SOURCE_KEY = _time
REGEX = ^(.+)$
FORMAT = <182>1 $1 $0
DEST_KEY = _raw
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. On the **Finish.conf** tab, restart Splunk by running the `$SPLUNK_HOME/bin/splunk restart splunkd` command.
9. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
10. Verify that the data input is configured correctly by selecting **Test connection**.

    Health Log Analytics tries to connect the MID Server to the data repository.

    - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
    - If the connection failed, the reason for the failure displays in the **Error message** field. This field displays only when a streaming error has occurred.

      Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

      Note: You can only publish the data input configuration when the connection is created successfully.

    Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

11. Select **Publish** to publish the data input to the MID Server.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts streaming log data to your ServiceNow instance using a Splunk shipper.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Rsyslog, Splunk, TCP data inputs

Configure a Splunk Polling data input in Health Log Analytics manually

Set up a data input that periodically pulls log data from Splunk by using a query.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

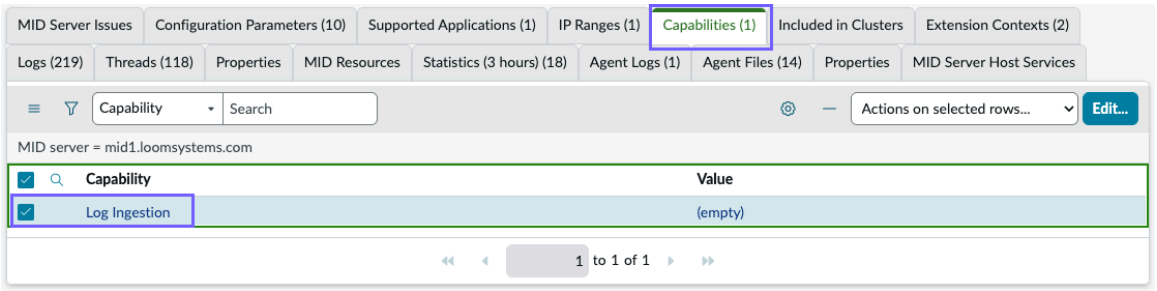

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the Splunk Polling data input.
4. On the form, fill in the fields.
   For a description of the fields, see Splunk Polling data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   On the **Transport tab** and **Advanced tab**, fill in the fields. For a description of the fields, see Splunk Polling data input configuration fields.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Verify that the data input is configured correctly by selecting **Test connection**.
   Health Log Analytics tries to connect the MID Server to the data repository. If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected.

   Note: You can only publish the data input configuration when the connection is created successfully.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
- If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Configure a TCP data input in Health Log Analytics manually

Set up a data input for sending raw log messages to your ServiceNow instance directly over a TCP/SSL socket.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

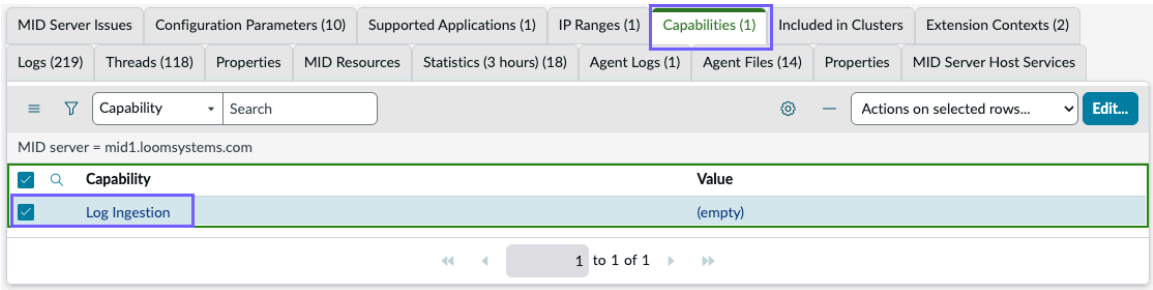

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Note: For information about streaming data from Logstash to your instance, see the Streaming data from Logstash to your HLA instance [KB0994735] article in the Now Support knowledge base.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the TCP data input type.

   Note: The selected data input type complements the passive data input (listener). For more information, see Supported data inputs for Health Log Analytics.

4. On the form, fill in the fields.
   For a description of the fields, see TCP data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   For a description of the fields, see TCP data input configuration fields. For information about configuring the advanced settings later, see Configure advanced settings for Rsyslog, Splunk, TCP data inputs.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts sending raw log messages to your ServiceNow instance directly over a TCP/SSL socket.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Rsyslog, Splunk, TCP data inputs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure a UDP data input in Health Log Analytics manually

Set up a data input for sending raw log messages to your ServiceNow instance directly over a UDP socket.

Before you begin

This feature is supported in the Health Log Analytics application, Version 25.0.17 - November 2022 and later, available from the ServiceNow Store.

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

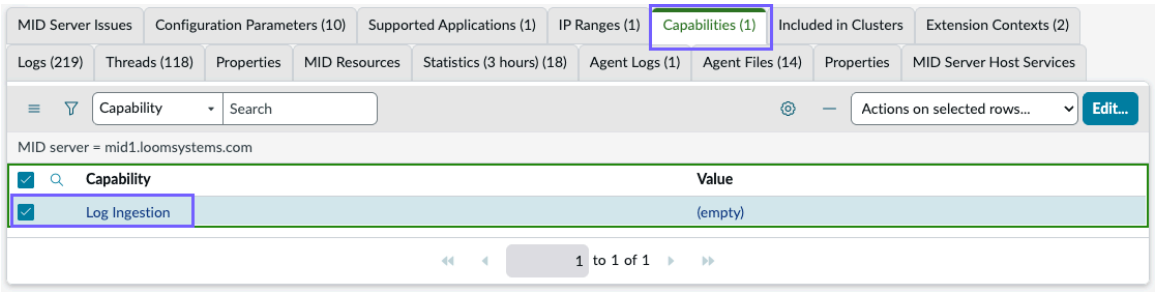

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *mid.public_ip* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.
- For shipping your logs encrypted using SSL TLS, see the Streaming Data With Rsyslog & Filebeat Using SSL [KB0866319] article in the Now Support Knowledge Base.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose the UDP data input type.

   Note: The selected data input type complements the passive data input (listener). For more information, see Supported data inputs for Health Log Analytics.

4. On the form, fill in the fields.
   For a description of the fields, see UDP data input configuration fields.
5. **Optional:** Select **Advanced** to set additional configuration fields.
   For a description of the fields, see UDP data input configuration fields.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Result

The data input configuration process is complete. Health Log Analytics adds the data input record to the **Data Inputs** table and attaches the configuration file to the data input record. The data input starts sending raw log messages to your ServiceNow instance directly over a UDP socket.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

What to do next

Make sure that the data input is streaming data.

Related topics

Configure advanced settings for Rsyslog, Splunk, TCP data inputs

Configure a Vector Agent data input in Health Log Analytics manually

Set up a Vector Agent data input to enable Health Log Analytics to process log messages that are streaming into your ServiceNow instance via a Vector Agent.

Before you begin

- Verify that a MID Server is installed and configured with the Log Ingestion capability enabled. For more information, see MID Server system requirements.

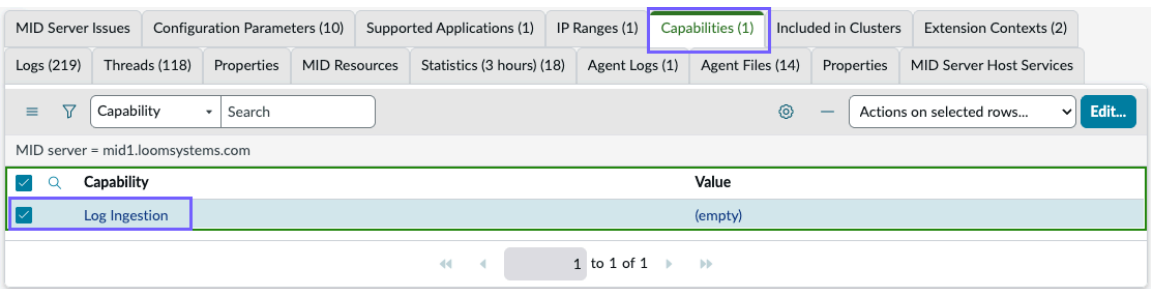

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Unless the MID Server and external clients are on the same network, the MID Server must have a public IP address. This is required when its IP is exposed through network address translation (NAT), a load balancer, or a similar device. The public IP address enables external clients, such as Filebeat agents located outside its network, to reach the MID Server. Private IP addresses aren't routable over the internet. Without a public IP, external clients can't connect to the MID Server even if they are configured with its address. In the MID Server properties, add a property named *`mid.public_ip`* with the public IP address as the value. For more information, see Create a MID Server property. If the MID Server and external clients are on the same network, connections can be made using the private IP address.

- Verify that you have a `vector.yaml` file on your system, or create one.

Role required: evt_mgmt_admin

About this task

If your organization uses Vector by Datadog to handle log data from multiple sources, the log format is distinct from other types. The Vector Agent data input enables HLA to detect and separate transport headers from inner log messages. It forwards only the inner message to the source type structure for processing.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. On the Data Inputs page, select **New**.
3. Choose **Vector - Stream logs from Vector Agent**.
4. On the **Getting Started** tab, fill in the form and then select **Next**.
   For a detailed description of the fields, see Vector Agent data input configuration fields.
5. On the **vector.yaml** tab, update the code with the relevant values as needed.
   For example, you can specify the sources or source file from which the logs will be streamed.

   Note: Code that you should not modify is preceded by a commented note.

   For more information, consult the Vector documentation.
6. Copy the updated code and paste it into the `vector.yaml` file on your system.
7. Select **Next**.
   The **Finish** tab displays a message confirming that the data input configuration process is complete.
8. Start streaming log data by running Vector with the configuration file.
   Run the following command:`vector --config /path/to/your/vector.yaml`, where `/path/to/your/` is the actual path to your YAML file.
9. **Optional:** Use the other commands displayed on the **Finish** tab as needed.
10. Select **Save** to save the data input.

Result

Health Log Analytics adds the data input record to the **Data Inputs** table. The data input starts streaming log data to your ServiceNow instance.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Make sure that the data input is streaming data.

Identify and resolve a log streaming issue in Health Log Analytics

Find and address log streaming issues to verify that your data inputs are streaming log data to your instance properly.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Streaming Sources**.
   The Streaming Sources page shows all data inputs and the MID Servers that are receiving logs from them.

   Note:

   - When **Look up hostnames** is selected in the advanced data input configuration, the Streaming Sources page shows the hostname of devices that use an Rsyslog or a Filebeat shipper. For Elasticsearch indices, it displays the index name.
   - Streaming Sources is also available as a related list on the data input form. The related list displays only the endpoint devices that are relevant to that data input.
   - If the HLA engine is down and data has stopped streaming, a notification appears at the top of the Streaming Sources page. When this happens, contact ServiceNow support.

2. Select a data input record to view the streaming data of its sources and identify streaming issues and their possible cause.
   For example, if the last recorded event time for a data input's endpoint server is yesterday, that server might be down or configured incorrectly. A streaming issue might also be caused by the data input configuration file not being installed on the endpoint.

Filter	Description
Status	The status of the source. A red bullet indicates that this source has not streamed data in the last hour.
Last event time	The last recorded time an event arrived at the MID Server in the last one-minute interval. Health Log Analytics continuously updates the last event time. If the last event time is not up to date, data is not streaming.
Raw log lines/sec	The average number of raw log lines that streamed to the MID Server per second in the last one-minute interval. **Note:** This value represents the number of raw log lines before preprocessing.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Filter	Description
Preprocessed log lines/sec	The average number of preprocessed log lines that streamed to the MID Server per second in the last one-minute interval. **Note:** This value can differ from the number of raw log lines per second. For example, the difference can be a result of logs having been dropped during preprocessing.

3. Investigate and resolve any data streaming issues.

Note: If you experience permissions-related issues with streaming log data from Elasticsearch, refer to the Granting privileges for data streams from Elasticsearch [KB0967366] article in the Now Support Knowledge Base.

What to do next

When the logs are streaming properly, proceed to map your raw log data.

Note: You can choose to edit incoming raw log data before Health Log Analytics processes it. For example, preprocessing enables you to discard log portions or remove sensitive data from your logs. This task is optional.

Additional data input setup tasks in Health Log Analytics

After performing the initial data input setup and configuration in Health Log Analytics, continue with the remaining data input setup tasks.

Add a timestamp format in Health Log Analytics

Define any timestamp format that does not appear in the list of defined formats. Health Log Analytics must be able to read timestamps of any format contained in your log files.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Timestamp Formats**.
2. To determine whether your timestamp format matches a predefined format in the system, do the following:
   a. Select **Find format**.
   b. Enter a sample of your timestamp.
   c. Select **Find format**.
   If the sample matches an existing format, the system provides its ID.
3. If your timestamp sample doesn't match an existing format, then manually define your custom format and add it to the available timestamp formats.
   a. On the Timestamp Formats page, select **New**.
   b. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Timestamp format form

Field	Description
ID	ID of the timestamp format. This value is selected automatically.
Format	Valid timestamp format. For example, dd-MMM-yyyy, HH:mm. **Note:** Health Log Analytics uses the Java 8 API for timestamp syntax.
Sample	Example of the timestamp format. For example, 03-May-2018, 10:20.
Description	Description of the timestamp format.

4. Select **Submit**.

Result
The system validates the timestamp format and adds it to the list of available timestamps.

Add a source type manually in Health Log Analytics

Create a source type manually before you configure a data input if you want to stream log data to a specific source type. By default, Health Log Analytics automatically extracts the source type during the mapping process.

Before you begin
This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.

Role required: evt_mgmt_admin

Procedure
1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Source Types**.
2. On the Source Types page, select **New**.
3. On the form, fill in the fields.

Source type form

Field	Description
Name	The name of the source type. The name must be between 3-50 characters long. Acceptable characters are letters, numeric values, dashes, and underscores.
Examples	One or more raw input samples for this source type. Enter each sample on a new line. **Note:** This field must contain data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The fields in the following table are read-only.

Field	Description
Created manually	Option for indicating that this source type was created manually. Default: true.
Auto-extraction enabled	Auto-extraction for this source type is disabled by default.
Learning mode	Option for indicating that Learning mode for the source type is ongoing. The AI engine must learn the format of log messages in each new log stream. In Learning mode, it learns the format of messages for this source type and extracts properties based on that knowledge.

4. Right-click the form header and select **Save**.
5. Select **Load examples** to load the raw input samples for this source type that you provided in the **Examples** field.
   When loading the log samples is completed successfully, the **Examples** field and the **Load examples** buttons no longer appear on the form.
6. **Optional:** Check how the AI engine has parsed the raw input samples by selecting the **Go to Source Type Structure** related link and modify the extracted properties if necessary.
7. Click **Submit**.

What to do next

Configure a data input and then map it to the new source type. For more information, see Set up data inputs in Health Log Analytics manually and Map raw log data manually.

Related topics

Refine the source type structure in Health Log Analytics

Configure source type capabilities in Health Log Analytics

Health Log Analytics extracts source types automatically in the mapping process. You can add timestamp formats and specify, delete, or exclude keywords for individual source types.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Source Types**.
   The Source Types table lists all source types that were extracted automatically during raw data mapping.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Source Types

Field	Description
Auto extraction enabled	Auto-extraction for this source type is either enabled or disabled. This value is set on the Source Type Structure form. Default: true.
Learning mode	Learning mode for the source type is either ongoing or completed. The AI engine must learn the format of log messages in each new log stream. In Learning mode, it learns the format of messages for this source type and extracts properties based on that knowledge.
Stack trace detection enabled	Detection of active stack frames for this source type is either enabled or disabled. Typically, a stack trace enables tracking an error to the log in which it originated.

2. Open a record.
3. **Optional:** Add a timestamp format for this source type.
   a. On the **Source Type Timestamps** tab, select **New**.
   b. In the **Timestamp format** field, select a format from the list of available timestamp formats.
   c. Select **Submit**.
4. **Optional:** Specify a keyword for this source type.
   a. On the **Lexical Keywords** tab, select **New**.

      Note: The **Lexical Keywords** tab contains both global and specified keywords. Selecting **New** creates a specified keyword for this source type.

   b. On the form, fill in the fields.

Field	Description
Name	Unique and descriptive name for the keyword.
Regular expression	The regular expression ("regex") that defines matches.
Exact match	Boolean value. If selected, Health Log Analytics matches the exact regex. For example, 'NullPointerException' in a message would not match the regex 'exception'. Default: True
Case-sensitive	Boolean value. If selected, Health Log Analytics looks for a case-sensitive match of the regex.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Default: False

c. Select **Submit**.

5. **Optional:** Delete a keyword specified for this source type.
   a. Select the row of the specified keyword that you want to delete.
   b. From the **Actions on selected rows** list at the bottom of the page, select **Exclude keyword**.
   c. Select **Continue** to delete the keyword.
6. **Optional:** Exclude a global keyword for this source type
   When you exclude a keyword, Health Log Analytics no longer looks for it in the log data of this source type.
   a. Select the row of the global keyword that you want to exclude.
   b. From the **Actions on selected rows** list at the bottom of the page, select **Exclude keyword**.

   Note: You can delete global keywords from the Lexical Keywords page.
7. Select **Update**.

Verify your log sources in Health Log Analytics

Verify that all your log sources are present and active. Health Log Analytics tagging assigns logs to service instances and components, and automatically maps logs to sources.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Log Sources**.
   The Log Sources table provides information about the sources created when you mapped the streaming data. The table lists these sources by service instance.

Column	Description
Service instance	The service instance. Service instances represent large or complete business processes. A single service instance can consist of multiple components, whereas different service instances can include the same components.
Component	The components comprising the service instance. Components represent small technical parts or business processes. Typically, they represent Configuration Items (CIs). Several components are often clustered together in a single service instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
Status	The status of the log source: Active or Inactive. The Active status indicates that logs streamed into the system from this source in the last hour.
Created	The creation date and time of this log source.
Last event time	The last time logs streamed into the system from this source.

2. Make sure that the system has created all your log sources.
3. Verify the creation, status, and log samples of a specific source in the source record.

   Note: Because auto-mapping can map log data to multiple sources, be sure to verify this information per source.

4. If a log source is missing, add it.

   a. Navigate to **Health Log Analytics** > **Data Input** > **Data Input Mapping**.

   b. Use the JavaScript function to manually map a log sample to the correct service instance, component, and source type.
   For more information, see Map raw log data manually.

   c. Navigate back to the Log Sources page and verify that the status of the new log source is Active.

   The system creates a log source.
5. Make sure that the status of all your log sources is Active.
   If the status of a source is Inactive, double-click it and change it to Active manually.
6. Select **Update**.

Related topics

Delete a log source in Health Log Analytics

Stop extraction of unneeded log data in Health Log Analytics

Extract specific log data in Health Log Analytics

Delete a log source in Health Log Analytics

Delete a log source with or without its associated log data in Health Log Analytics.

Before you begin

Role required: evt_mgmt_admin

About this task

When you delete a log source, Health Log Analytics automatically deletes its related log data. You can also delete a source while keeping the log data. The associated data will then be deleted according to the log retention policy. For more information, see Storage space for log retention in Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Log Sources**.
2. Open a source record.
3. Delete the log source in one of the following ways:
   - Delete the source and its related data by selecting **Delete**.
   - Delete the source without its related data by selecting the **Delete without logs** related link.

Review properties extracted from a source type in Health Log Analytics

Inspect the properties that were extracted from all the source types in a source type structure in a single table to identify any setup issues.

Before you begin

Role required: evt_mgmt_admin

About this task

The Source Type Properties Stats table provides a comprehensive overview of the properties that were extracted from the source types. By reviewing the statistical distribution of their values, you can assess the classification of these properties and adjust your settings if needed.

The procedure in this section explains how to access the Source Type Properties Stats table via the Source Type Structure. From a source type record, you can navigate to the Source Type Structure by selecting the **Go to Source Type Structure** related link.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Source Type Structures**.
2. Open a source type structure record.
3. Open the **Properties Statistics** related links tab.

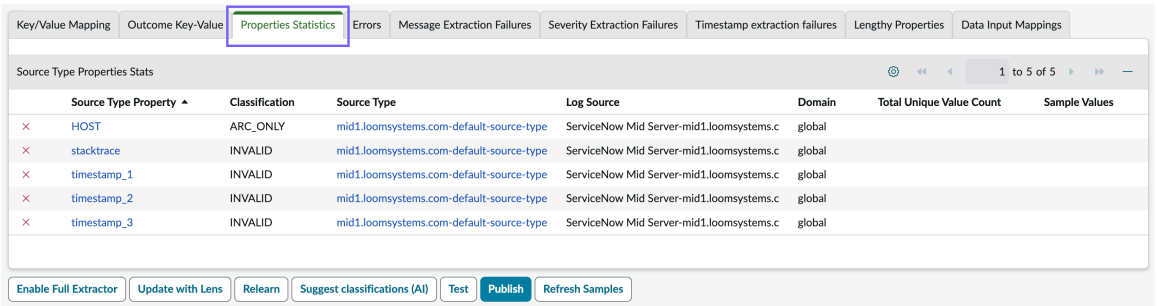

The **Source Type Properties Stats** table shows the properties of all the source types that are part of this source type structure.

Source Type Properties Stats

Column	Description
Source Type Property	The property extracted from the source type.
Classification	The classification of the property. The available types are: METER, GAUGE, ARC, and INVALID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
	Note: A property classified as INVALID (the default classification) isn't tracked for anomaly detection or shown in the Meaningful Properties analysis.
Source Type	The source type from which the property was extracted.
Log Source	The list of Log Sources in which this property appeared.
Domain	The domain to which the log source belongs.
Total Unique Value Count	Only for properties classified as METER: The total number of distinct values contained in it. This field is inactive when the property is classified as ARC or GAUGE. **Note:** The Total Unique Value Count is updated once a day. Records older than 24 hours are deleted automatically.
Sample Values	Only for properties classified as METER: Sample unique values tracked by the system. Default is 30 samples. This field is inactive when the property is classified as ARC or GAUGE.

4. Review the properties and identify possible classification issues.
5. **Optional:** View statistics for individual source type properties classified as ARC_ONLY or GAUGE.
   a. Open a source type property classified as ARC_ONLY or GAUGE from the **Properties Statistics** related links tab.
   b. On the **Source Type Property** form, select **Source Type Properties Stats**.
   The following information displays in a pop-up:
   - For an ARC property: The percentage distribution of terms, based on the processed logs. For example, the percentage of Bad Request vs. Success.
   - For a GAUGE property: The average, minimum, and standard deviation from the expected numerical value.
6. **Optional:** Change classification settings on the **Key/Value Mapping** tab.
   For more information, see Refine the source type structure in Health Log Analytics.

Review patterns extracted from a source type in Health Log Analytics

Inspect all learned patterns extracted from a source type in a source type structure, together with the log sources in which these patterns appeared. Reviewing these patterns can provide valuable insights into the log message patterns that Health Log Analytics tracks for each source type and log source.

Before you begin

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Source Type Structures**.
2. Select a source type in the **Source type name** column.
3. On the Source Type form, select the **Aggregated Log Patterns** related list.
   The tab shows the following data for each pattern that was extracted from this source type. The displayed data is updated every hour.

Aggregated Log Patterns

Column	Description
Pattern text	The content of the aggregated pattern.
Sources	The sources in which the pattern appeared.
Domain	The domain to which the pattern belongs.

Modify a data input configuration in Health Log Analytics

Change the configuration of a data input for Health Log Analytics by adding a new path to an existing data input configuration or modifying the data input's MID Server destination and port.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open a record from the Data Inputs table.
3. Modify the data input configuration.

Column	Description
Name	Name of the data input.
Description	Description of the data input.
Port	The port on the MID Server. **Note:** The port must not be occupied by another process. Make sure that your organization's security team opens the selected port.
MID	The MID Server to which the logs stream.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
	Note: ○ You can select only MID Servers with log ingestion capability that support basic authentication. MID Servers that support mTLS are not listed. ○ The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties.

Settings

Column	Description
Path	The full path from which to stream logs. You can use a wildcard. **Note:** This column is not available on Windows systems using Winlogbeat.
Service instance	The service instance to which to bind the log data. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.
Component	The device type or stack layer as context for the logs that is used for anomaly detection and correlation. For example, Tomcat. Components typically represent CIs in the CMDB. Several components are often clustered together in a single service instance.
Source Type	The source type that defines how Health Log Analytics handles a specific application and parses the log data. For example: Tomcat Catalina. Each data input can have multiple source types depending on the diversity of its log formats. Service instances and components can have any number of source types.
For handling multiline messages on Linux / Windows systems using Filebeat only:	
Match	Specifies how Filebeat combines matching lines into an event, either **after** or **before**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
Negate	Boolean that defines whether the pattern identified in the log lines is negated. The default is **false**.
Regex	The regular expression to match.

Note: You can modify the Rsyslog configuration file to make the agent ship system logs in addition to application logs. For more information, see the Shipping system logs using Rsyslog [KB0954507] article in the Now Support Knowledge Base.

4. Select **Update**.
5. For data inputs that use Rsyslog or Beats agents only, rebuild the server-side configuration file and install it on the endpoint device.
   a. Select **Rebuild configuration file**.

      Health Log Analytics rebuilds the file and saves it in the Manage Attachments section. Depending on the agent used, the rebuilt file is saved as either `rsyslog.yml`, `filebeat.yml`, or `winlogbeat.yml`.

      The system automatically renames the previous configuration file by adding a suffix with the date and time the file was rebuilt to the file name.

   b. Install the rebuilt configuration file on the endpoint according to your data input type.

Data input type	Action
Rsyslog	i. Download the file and install it on the endpoint device in the `/etc/rsyslog.d/rsyslog.conf` directory. ii. Validate the configuration by running the `rsyslogd -N1` command. iii. Verify the output. If it contains errors, check the `/var/log/messages` system log file for error messages and fix the errors. iv. Restart Rsyslog by running the `sudo systemctl restart rsyslog` command.
Linux	i. Download the file and install it on the endpoint device in the `/etc/filebeat/` directory. ii. Restart the agent service by running the `sudo service filebeat restart` command. **Note:** The generated configuration ignores files that were last changed more than six hours ago. If needed, you can change this setting in the configuration file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data input type	Action
Windows using Beats (Filebeat or Winlogbeat):	i. Download the file and install it on the endpoint device in the C : \Program Files\ directory. ii. Restart the agent service by running the appropriate command in PowerShell: ▪ Filebeat: `PS > Restart-Service filebeat` ▪ Winlogbeat: `PS > Restart-Service winlogbeat` **Note:** The generated configuration ignores files that were last changed more than six hours ago. If needed, you can change this setting in the configuration file.

Related topics

Configure advanced settings for Rsyslog, Splunk, TCP data inputs

Configure advanced settings for Beats data inputs in Health Log Analytics manually

Advanced data input configuration in Health Log Analytics

When you have configured a data input successfully, Health Log Analytics adds a record to the Data Inputs table and attaches the configuration file to it. You can configure advanced settings for your data input. Configuring advanced settings is optional.

Configure advanced settings for Amazon CloudWatch data inputs in Health Log Analytics manually

Configure advanced settings for data inputs used for streaming log data from Amazon CloudWatch to your instance.

Before you begin

Role required: evt_mgmt_admin

About this task

You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open an Amazon CloudWatch data input record from the Data Inputs table.

   Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced** and then select the **Advanced** tab.
4. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a description of the fields, see Amazon CloudWatch data input configuration fields.

The **Character encoding** field is read-only. By default, the character encoding used for this data input is UTF-8.

5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Related topics

Configure an Amazon CloudWatch data input in Health Log Analytics manually

Configure advanced settings for Amazon S3 data inputs in Health Log Analytics manually

Configure advanced settings for data inputs used for streaming log data from Amazon S3 buckets to your instance.

Before you begin
Role required: evt_mgmt_admin

About this task
You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open an Amazon S3 data input record from the Data Inputs table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced** and then select the **Advanced** tab.
4. On the form, fill in the fields.
   For a description of the fields, see Amazon S3 data input configuration fields.
5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Related topics

Configure an Amazon S3 data input in Health Log Analytics manually

Configure advanced settings for Apache Kafka data inputs in Health Log Analytics manually

Configure advanced settings for data inputs used for streaming log data from Apache Kafka to your instance.

Before you begin

Role required: evt_mgmt_admin

About this task

You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open an Apache Kafka data input record from the Data Inputs table.

   Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced** and then select the **Advanced** tab.
4. On the form, fill in the fields.
   For a description of the fields, see Apache Kafka data input configuration fields.
5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Related topics

Configure an Apache Kafka data input in Health Log Analytics manually

Configure advanced settings for Beats data inputs in Health Log Analytics manually

Configure advanced settings for data inputs that use Beats agents.

Before you begin

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

For information about how to change settings that were configured when the data input was created, such as adding a new path or changing the data input's MID Server destination or port, see Modify a data input configuration in Health Log Analytics.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open a Beats data input record from the Data Inputs table.
   The data input configuration displays.

   Note: The number of log sources that the data input has created is shown in the **Sources count** field. For more information about data input sources, see Log data auto-mapping and mapping in Health Log Analytics.

   Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced**.
4. On the form, fill in the fields.
   For a description of the fields, see Rsyslog, Filebeat, or Winlogbeat data input configuration fields.
5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. This field displays only when a streaming error has occurred.

     Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Configure a Rsyslog, Filebeat, or Winlogbeat data input in Health Log Analytics manually

Configure advanced settings for Elasticsearch data inputs in Health Log Analytics manually

Configure advanced settings for data inputs used for streaming log data from Elasticsearch indices to your instance.

Before you begin

Role required: evt_mgmt_admin

About this task

You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

For additional information about streaming logs using the Elasticsearch data input, see the Stream logs using Elasticsearch data input - Advanced guide [KB1080162] article in the Now Support Knowledge Base.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open an Elasticsearch data input record from the Data Inputs table.
   The data input configuration displays.

   Note: The number of log sources that the data input has created is shown in the **Sources count** field. For more information about data input sources, see Log data auto-mapping and mapping in Health Log Analytics.

   Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced** and then select the **Advanced** tab.
4. On the form, fill in the fields.
   For a description of the fields, see Elasticsearch data input configuration fields.
5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.

   Note: If you experience permissions-related issues with streaming log data from Elasticsearch, refer to the Granting privileges for data streams from Elasticsearch [KB0967366] article in the Now Support Knowledge Base.

6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

- If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
- If the connection failed, the reason for the failure displays in the **Error message** field. This field displays only when a streaming error has occurred.

  Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

  Note: You can only publish the data input configuration when the connection is created successfully.

Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Related topics

Configure an Elasticsearch data input in Health Log Analytics manually

Configure advanced settings for Microsoft Azure Event Hubs data inputs in Health Log Analytics manually

Configure advanced settings for data inputs used for streaming log data from Microsoft Azure Event Hubs to your instance.

Before you begin

Role required: evt_mgmt_admin

About this task

You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open a Microsoft Azure Event Hubs data input record from the Data Inputs table.

   Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced** and then select the **Advanced** tab.
4. On the form, fill in the fields.
   For a description of the fields, see Microsoft Azure Event Hubs data input configuration fields.
5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics adds the data input record to the Data Inputs table.

7. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Related topics

Configure a Microsoft Azure Event Hubs data input in Health Log Analytics manually

Configure advanced settings for Microsoft Azure Log Analytics data inputs in Health Log Analytics manually

Configure advanced settings for data inputs used for streaming log data from Microsoft Azure Log Analytics to your instance.

Before you begin

Role required: evt_mgmt_admin

About this task

You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. If no advanced settings are configured, the system uses the default values.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open a Microsoft Azure Log Analytics data input record from the Data Inputs table.

   Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced** and then select the **Advanced** tab.
4. On the form, fill in the fields.
   For a description of the fields, see Microsoft Azure Log Analytics data input configuration fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   If the data input is configured to run on a MID Server cluster, the system tries to connect all the MID Servers contained in the cluster to the repository. The cluster passes the test if at least one of its MID Servers gets connected. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Related topics

Configure a Microsoft Azure Log Analytics data input in Health Log Analytics manually

Configure advanced settings for MID Server data inputs in Health Log Analytics manually

Configure advanced settings for data inputs that are used for collecting and streaming MID Server log messages.

Before you begin
Role required: evt_mgmt_admin

About this task
You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open a MID Server data input record from the Data Inputs table.
   The data input configuration displays.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The number of log sources that the data input has created is shown in the **Sources count** field. For more information about data input sources, see Log data auto-mapping and mapping in Health Log Analytics.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

3. Select **Advanced**.
4. On the form, fill in the fields.
   For a description of the fields, see MID Server data input configuration fields.
5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Update**.

Result
Health Log Analytics updates the system parameters.

Related topics

Configure a MID Server data input in Health Log Analytics manually

Configure advanced settings for Rsyslog, Splunk, TCP data inputs

Manually configure advanced settings for data inputs that use Rsyslog, Splunk, or TCP agents in Health Log Analytics.

Before you begin
Role required: evt_mgmt_admin

About this task
You can set system parameters for reading log data that determine the actions that the system performs on log data arriving on the MID Server. For example, you can set the time zone to use if a log lacks a timestamp. If no advanced settings are configured, the system uses the default values.

For information about how to change settings that were configured when the data input was created, see Modify a data input configuration in Health Log Analytics.

Procedure
1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. Open a Rsyslog, Splunk, or TCP data input record from the Data Inputs table.

   The data input configuration displays.

Note: The number of log sources that the data input has created is shown in the **Sources count** field. For more information about data input sources, see Log data auto-mapping and mapping in Health Log Analytics.

Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the data input configuration page. When this happens, contact ServiceNow support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Advanced**.
4. On the form, fill in the fields.
   For a description of the fields, see Rsyslog, Filebeat, or Winlogbeat data input configuration fields.
5. **Optional:** In the Streaming Sources related list, verify that this data input is streaming log data from all relevant endpoint devices.
   For more information about streaming sources, see Identify and resolve a log streaming issue in Health Log Analytics.
6. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
7. Verify that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. This field displays only when a streaming error has occurred.

     Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

8. Select **Publish** to publish the data input to the MID Server.

Stop or restart a data input in Health Log Analytics

You can stop using a data input for streaming log messages to your ServiceNow instance. Restart the data input if you want it to resume streaming data.

Before you begin

This functionality is available only when the data input configuration has been published.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Log Data Inputs**.
2. Open a data input record from the Data Inputs table.
3. Select **Stop Data Input**.
   The name of the button changes to **Start Data Input**.
4. To start using the data input again, select **Start Data Input**.

Edit raw log data before processing in Health Log Analytics

Use the Data Input Preprocessor to filter, split, or sanitize raw log data before it is treated in the MID Server and mapped and structured by Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: evt_mgmt_admin

About this task
Define a JavaScript function to preprocess your raw log data before it reaches Health Log Analytics. For example, you can drop irrelevant or redundant messages to reduce log volume. You can also split multi-event log lines into discrete messages for accurate mapping or sanitize sensitive or personally identifiable information so that it never enters the ServiceNow instance.

Procedure
1. Navigate to **All** > **Health Log Analytics** > **Data Input Preprocessor**.
2. Open a record.

   Note: The first time the preprocessing form is displayed, Health Log Analytics fetches log samples automatically. In subsequent sessions, fetch fresh samples manually by selecting **Refresh Samples**.

3. **Optional:** Prevent showing log samples that are taken from your organization's raw logs by selecting **Disable raw log samples**.
   Select this option if you don't want to expose your logs to Health Log Analytics.
4. **Optional:** View how the current JavaScript function affects the log lines.
   a. Add a sample message in the **Test manual sample** field.
   b. Select **Go**.
   c. Note how the JavaScript function affects the log lines.
5. In the **Raw input samples** field, choose a log sample that will show the effect of the new JavaScript function you're going to define on log lines when you test it.
6. Define a JavaScript function that modifies your raw log data before Health Log Analytics maps and structures it.

   Note:

   (ACC data inputs only) When the Agent Client Collector switches to a different MID Server to provide failover protection, it must switch to a different ACC data input. Therefore, all ACC data inputs must have the same JavaScript function. Health Log Analytics provides the latest published JavaScript function to all existing and future ACC data inputs, replacing the previous script. This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. For information about failover protection in Agent Client Collector Log Analytics (ACC-L), see Agent Client Collector Log Analytics.

   a. In the JavaScript console, either change the provided default JavaScript function, modify an existing custom JavaScript function, or define a new one.

      Note: In addition to the default JavaScript function for data preprocessing, Health Log Analytics provides a JavaScript function template named `Scrubbing`. This JavaScript function removes personally identifiable information such as email addresses, social security numbers, and passwords, from the raw log data. The template can act as a starting point for your custom script code. This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.

      Note: (ACC data inputs only) Make sure that your JavaScript function can be used to handle data streamed by all ACC data inputs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

JavaScript functions for preprocessing raw log data use the following objects:

- Signature: function process(sample, metadata)

Object	Description
sample	The current log sample.
metadata	Object that contains the event stream. To access the event stream, call the *`metadata.eventStream`* function. Because event stream is an object, you can access each of its properties by calling the *`metadata.eventStream.origin`* or *`metadata.eventStream.LocalPort`* functions. The list of available event stream properties is available in the **Event Stream: access via metadata** tab.

- Return type and structure

  Note: The JavaScript function returns a map of two entries. Do not change this return structure.

Object	Description
modifiedInput	The current log message after the JavaScript function has modified the original message. If null, the original log message is used.
splitEvents	Array of log messages after the JavaScript function has divided the original message.

- To discard a log message, call `return drop()`.

b. Test your JavaScript function by selecting **Test** and viewing the outcome on the log sample in related lists.

Tab	Description
Outcome	The result of running the JavaScript function on your log sample.
Event Stream: access via metadata	The key-value pairs that present the event stream metadata.
Failures	The log samples on which the JavaScript function failed to run successfully.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If your new JavaScript function is not behaving as expected, you can revert to the last published one by selecting the **Revert JS Function** related link.

c. **Optional:** Make any necessary adjustments and then test the JavaScript function again.

7. Select the **Save template** option to save the JavaScript function.
   You can either save the JavaScript function as a new template or overwrite the currently selected template.
   - To save the JavaScript function as a new template, enter a new name in the **Template name** field.
   - To overwrite the template currently selected in the **JS function templates** field, leave the **Template name** field empty.
8. Select **Publish** to save the JavaScript function to the database.

Result

When the JavaScript function is published, Health Log Analytics uses it to preprocess your raw log data before mapping and structuring it.

The new script is automatically added to the list of JS function templates from which you can choose. This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.

(ACC data inputs only) Health Log Analytics provides the published JavaScript function to all existing and future ACC data inputs. The new JavaScript function replaces the previous script.

What to do next

Proceed to the **Data Input Mapping** page by selecting the **Go to Mapping** related link. Map raw data that streams into your instance to determine how Health Log Analytics processes it.

Log data auto-mapping and mapping in Health Log Analytics

By default, the HLA Engine tries to auto-map every incoming log line to the correct tags. You can change automatic mapping results manually by defining a JavaScript function.

Auto-mapping incoming log lines

Health Log Analytics auto-mapping assigns log samples and metadata to three tags: service instance, component, and source type. The service instance assignment is based on the service instance specified in the data input setup. The remaining tags are assigned automatically.

For example, in the following example log line, Health Log Analytics uses the "source" field to find the component and source type.

```
{"beat":{"version":"6.8","name":"abc3.prd.acme.com","hostname":"
abc3.prd.acme.com"},"@timestamp":"2020-08-27T10:12:24.792Z","pro
spector":{"type":"log"},"message":"****
 User null is requesting the following page
 http://www.acme.com PROPS:{"subcategory1":"home
 pages","httpStatus":"200","loginLevel":"Anonymous","userAgent":
"Mozilla5.0",
 ("pageUrl":\"http://www.acme.com","host":"abc3.prd.acme.com","@
version":"1","source":"/opt/oracle/weblogic/abc/online_store3/lo
gs/online_store3.out","offset":3951550786}
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In the example, Health Log Analytics extracts the string "online_store". It analyzes the following fields if they exist in the log line: source, path, channel, namespace_name, name, pod_name, source_name, and aws_lambda_name. When data is sent over Syslog, it also analyzes the syslog tag.

Stop extraction of unneeded data

If an extracted string is not descriptive enough or contains redundant text or information, you can stop extracting such expendable data. For more information, see Stop extraction of unneeded log data in Health Log Analytics.

Ensuring extraction of specific data

You can make sure that Health Log Analytics extracts specific desired terms. For more information, see Extract specific log data in Health Log Analytics.

AI-assisted log mapping with ServiceNow Otto for ITOM

AI-assisted log mapping using ServiceNow Otto for ITOM recommends optimal mapping configurations based on your log patterns. This approach reduces manual setup effort and improves accuracy without requiring custom JavaScript functions.

For this capability to be enabled, the integration must be activated by selecting the Activate with AI option. This option enables AI-powered automatic parsing and mapping of log data to service instances and components for contextual alerts. When the integration is activated, ServiceNow Otto for ITOM collects and analyzes log data and suggests the best log field for mapping. You can override the suggested mapping by selecting a different log field from each list, if needed. For more information, see Map logs to service instances, components, source types.

Mapping data input sources manually by defining a JavaScript function

You can change automatic mapping results manually by defining a JavaScript function. Data input mapping enables you to organize your log data by service instance and by availability zone. A single service instance can include multiple components, and a component can receive logs from many different source types. An service instance-component pair, however, is unique. Source types are based on a specific log structure and format. Service instances and components are defined more broadly and are therefore used mainly for logical mapping.

Activating **Test mode** avoids blowing up Elasticsearch storage with sample data that is used only for perfecting the log data mapping. When the data input is in Test mode, Health Log Analytics doesn't create the source types, sources, or any other objects it creates in the standard flow. It saves the streamed data in dedicated temporary Elasticsearch indices that appear as components in the Log viewer. When you publish the script and exit Test mode, these temporary indices are deleted to minimize storage space consumption.

When you're defining a JavaScript function, select **Test** to view the outcome of the script as it is currently specified. This functionality enables you to preview the created source types and sources. You can then refine the script until it achieves the desired outcome. For example, it can be useful to compare the test outcome of several versions of the JavaScript function.

Note: By default, the test processes 100 log data samples. You can customize this number in the system properties. For more information, see Configure global Health Log Analytics system properties.

During the data input setup, the system might create an excessive total number of data input sources. For example, this can be due to a faulty mapping script. You can configure limits for the number of sources created per data input in the system properties:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

System property	Description	Default
log_source.sources_warning_limit	The warning limit for the number of sources created per data input.	500
log_source.sources_critical_limit	The critical limit for the number of sources created per data input.	600

The number of log sources that a specific data input has created displays in the **Sources count** field for that data input. When the total number of sources created during the data input setup reaches the warning limit, the system sends a warning notification by email. It also displays a message on the **Data Input mapping**, **Log sources**, and **Data input** pages. The notification and the message include the total number of sources created so far and the three data inputs that contributed the most sources to this total. If no action is taken, the system continues to create sources until the total number reaches the critical limit. When this happens, the data input setup and streaming from all data inputs stops automatically. You can't start data inputs again manually until the issue has been resolved. You can resolve this state by following the instructions in the How to handle too many sources in data inputs [KB0963067] article in the Now Support Knowledge Base.

Binding log data

Binding log data to Configuration Items (CIs) in the Configuration Management Database (CMDB) enables you to search the CMDB for endpoints that match a log. When you configure a data input, you bind log entries to a service instance that is bound to a CI in the CMDB. Binding log entries, service instances, and CIs enables the HLA Engine to correlate them for use in root cause analysis (RCA). For more information, see Configure a Rsyslog, Filebeat, or Winlogbeat data input in Health Log Analytics manually or Configure an Elasticsearch data input in Health Log Analytics manually.

Map raw log data manually

Mapping raw log data that streams into your ServiceNow instance determines how Health Log Analytics processes the data. If HLA doesn't discover properties automatically, you can map data input sources manually.

Before you begin

For an overview of mapping, see Log data auto-mapping and mapping in Health Log Analytics.

Role required: evt_mgmt_admin

About this task

Health Log Analytics structures logs, creates metrics for anomaly detection, and presents alerts based on how your data is tagged. By default, HLA tries to auto-map every incoming log line to the correct tag. To manually map data input sources when the system doesn't automatically discover properties, create a JavaScript function.

In the JavaScript function, you're required to map only the service instance (here called application service). Mapping the component and the source type is optional: Health Log Analytics tries to extract their values from the log data automatically. If the attempt fails, it assigns the default values. If you map the component but not the source type or vice versa, the system tries to extract the missing value from the log data. If it fails, it assigns the component value to the source type or vice versa, depending on which one you have mapped. This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

(ACC data inputs only) When the Agent Client Collector switches to a different MID Server to provide failover protection, it must switch to a different ACC data input. Therefore, all ACC data inputs must have the same JavaScript function. Health Log Analytics provides the latest published JavaScript function to all existing and future ACC data inputs, replacing the previous script. This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. For information about failover protection in Agent Client Collector Log Analytics (ACC-L), see Agent Client Collector Log Analytics.

Note: You can manipulate raw log data before Health Log Analytics maps and structures it. For more information, see Edit raw log data before processing.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Mapping** > **Data Input Mapping**.
2. Open a record.

   Note: The first time the Data Input Mapping form is displayed, Health Log Analytics fetches log samples automatically. In subsequent sessions, fetch fresh samples by selecting **Refresh Samples**.

   Note: If the HLA engine is down and data has stopped streaming, a notification appears at the top of the Data Input Mapping page. When this happens, contact ServiceNow support.

3. **Optional:** To forward the complete raw log message, disable automatic header properties detection for this data input by selecting **Disable header detection**.
   For more information, see Header properties detection in Health Log Analytics.
4. **Optional:** Activate Test mode by setting the **Test mode** value to ON.
   In Test mode, Health Log Analytics doesn't create the source types, sources, or any other objects it creates in the standard flow. This is to avoid blowing up Elasticsearch storage with sample data. For more information, see Log data auto-mapping and mapping in Health Log Analytics.
5. **Optional:** View how the current JavaScript function affects log lines.
   a. Add a sample message in the **Test manual sample** field.
   b. Select **Go**.
   c. Note how the JavaScript function affects the log lines.
6. In the **Raw input sample** field, choose a log sample that will show the effect of your new JavaScript function on log lines when you test it.
7. **Optional:** Enable HLA to identify service instances and components from log fields and automatically map log data to them by selecting **Enable auto map from log field**.

   Note:

   - Automatic log context mapping can only be configured in the Service Operations Workspace. For more information about this feature, see Map logs to service instances, components, source types.
   - If a manually-defined mapping script exists, automatic log context mapping overwrites it.
   - To be able to define a JavaScript function manually in the current form, the **Enable auto map from log field** check box must be clear.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Define a JavaScript function that maps your data input sources to the correct service instance (here called application service), component, and source type.

   Note: (ACC data inputs only) Make sure that your JavaScript function can be used to handle data streamed by all ACC data inputs.

   a. In the JavaScript console, either modify the provided default JavaScript function or a custom JavaScript function selected from the **JS functions templates** drop-down list, or define a new JavaScript function.
   The provided templates can act as a starting point for your custom script code. Available templates are:

   - Default Mapping Script
   - Glide Node Log Mapping Script
   - Glide Sys Log Mapping Script
   - Linux OS logs streamed with Syslog

     Note: Logs must have a syslog wrapper for this JavaScript function to work properly.

   - Mid Server Log Mapping Script
   - Windows - Winlogbeat logs streamed with Winlogbeat

   The JavaScript function for mapping raw log data uses the following objects:

   - Signature: function map(sample, metadata)

Object	Description
sample	The current log sample after preprocessing.
metadata	Object that contains: ▪ Event Stream: Access via: <metadata value>. For example: `metadata.eventStream.origin` ▪ Transform header 1: Access via: <metadata value>. For example: `metadata.headers.i1.type` ▪ Transform header 2: Access via: <metadata value>. For example: metadata.headers.i2.type ▪ type Transform header 3: Access via: <metadata value>. For example: `metadata.headers.i3.type` ▪ Fallback Assignment: Access via: <metadata value>. For example: `metadata.fallBacks.host`

   Note: You can create and associate multiple services per source type.

   - Return type and structure

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The JavaScript function returns a map of two entries. Don't change this return structure.

Object	Description
applicationService	The existing application service (that is, the service instance) to which this sample will be assigned. **Note:** The service instance must be bound to anomaly alerts related to the log.
component	The component to which this sample will be assigned.
sourceType	The source type to which this sample will be assigned.

- To discard a log message, call `return drop()`.

b. Test the JavaScript function by selecting **Test**.

Testing the JavaScript function enables you to view the outcome of the script on the log sample. For a description of the displayed fields, see JavaScript function test outcome fields.

Note: If your new JavaScript function is not behaving as expected, you can revert to the last published one by selecting the **Revert JS Function** related link.

c. Optional: Make any necessary adjustments and then test the JavaScript function again.

d. Optional: Compare the outcome of multiple tests.
Comparing the test results of several versions of the JavaScript function can help you refine the script until it achieves the desired outcome.

9. When you have finalized the JavaScript function, select **Save template** to save it.
You can either save the JavaScript function as a new template or overwrite the currently selected template.

- To save the JavaScript function as a new template, enter a new name in the **Template name** field.
- To overwrite the template currently selected in the **JS function templates** field, leave the **Template name** field empty.

10. Select **Publish** to save the JavaScript function to the database.

Result

When the JavaScript function is published, Health Log Analytics uses it to map data input sources.

The new script is automatically added to the list of JS function templates from which you can choose. This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

(ACC data inputs only) Health Log Analytics provides the published JavaScript function to all existing and future ACC data inputs. The new JavaScript function replaces the previous script.

What to do next

- (Optional) Edit your raw log data before Health Log Analytics maps and structures it. If you want to perform this task immediately, select the **Go to Preprocessor** related link to proceed to the **Data Input Preprocessor** page.
- (Optional) Fine-tune how the system reads your log data by refining the Source Type structure. This step enables you to reclassify auto-classified properties and change auto-mapped labels.

Related topics

Log data auto-mapping and mapping in Health Log Analytics

View source type and log source relationships

Header properties detection in Health Log Analytics

Stop extraction of unneeded log data in Health Log Analytics

Extract specific log data in Health Log Analytics

Header properties detection in Health Log Analytics

Automatic header properties detection in Health Log Analytics separates the transport header from the inner log message and forwards only the inner log message to the source type structure. The inner message contains the actual log data, not the shipping information.

The Health Log Analytics application supports header properties detection for Fluentd, Beats, and Syslog (RFC 3164, RFC 5424).

Starting with Version 33.0.27 - August 2024, the system also supports header properties detection for logs that follow the OpenTelemetry logs data model and semantic conventions. For more information, see the OpenTelemetry Logs Data Model documentation.

You can disable header properties detection for a data input. In this case, the Health Log Analytics AI Engine stops extracting properties from the header. Forwarding the complete raw message to the source type structure can be useful in the following situations:

- The inner log message lacks information for parsing, such as timestamp and severity.
- The header contains information needed for structuring the log data.
- The data input forwards the logs fully parsed.

For the procedure to disable header properties detection, see Map raw log data manually.

Extract specific log data in Health Log Analytics

Set Health Log Analytics to extract specified terms from logs and map them to specific components.

Before you begin

For an overview of mapping, see Log data auto-mapping and mapping in Health Log Analytics.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Mapping** > **Extracted values**.
2. Select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the **Value to extract** field, enter the term that you want Health Log Analytics to extract from your logs.
4. In the **Map to component** field, enter the component that you want the value mapped to.
5. In the **Property** field, choose a property from the list.
6. **Optional:** Further specify the value to extract by selecting additional options:
   - Exact match
   - Case sensitive
7. **Optional:** In the **Description** field, add a description of the value.
8. Select **Submit**.

Stop extraction of unneeded log data in Health Log Analytics

If an extracted string of data is not descriptive enough or contains redundant text or information, you can set Health Log Analytics to stop extracting such data from your logs.

Before you begin

For an overview of mapping, see Log data auto-mapping and mapping in Health Log Analytics.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Mapping** > **Excluded values**.
2. Select **New**.
3. In the **Value to exclude** field, enter the string that you want Health Log Analytics to stop extracting.
4. **Optional:** Further specify the values to stop extracting by selecting additional options:
   - Starts with
   - Ends with
   - Contains
   - Case sensitive
5. **Optional:** In the **Description** field, add a description of the value.
6. Select **Submit**.

Mapping logs for contextual alerts in Health Log Analytics

Map your incoming log data to service instances, components, and source types so that Health Log Analytics can generate alerts in context.

> **Important:** To enable AI-powered automatic parsing and mapping of log data, the integration must be activated with AI.

Mapping your log data to the correct context is especially important when the integration processes logs from multiple service instances and components. ServiceNow Otto suggests the best log field for mapping to service instances and components. When you use the AI-suggested field, or when that field is the default, an AI sparkle icon (✨) appears. You can select a different field if needed. If the AI agent can't find a suitable match, HLA uses the system default. The system default also applies if the selected field is not present in the sample log.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For a walkthrough of how to set up and review AI-suggested mappings, see AI-assisted log mapping in Health Log Analytics.

Example

A large financial institution might face performance issues with its e-banking application, which relies on various components like web, application, and database servers. Without log context mapping, logs from these components appear isolated, making it difficult to correlate issues. An anomaly in a Tomcat server log might be detected, but without proper context, the operator struggles to assess its impact. Log context mapping enables you to define rules to map logs to the e-banking application service instance and the Tomcat server component. This mapping provides a contextualized view for root cause analysis and resolution.

Map logs to service instances, components, source types

Map your logs to service instances, components, and source types so that Health Log Analytics (HLA) can generate contextual alerts.

Before you begin

For an overview of this feature, see Mapping logs for contextual alerts in Health Log Analytics.

Configure and activate an integration with AI before you perform this procedure. For more information, see Set up integrations for Health Log Analytics from the Integrations Launchpad.

Role required: evt_mgmt_admin

About this task

HLA uses the settings that you define in this procedure to identify service instances, components, and source types in log examples and map log data to them automatically. HLA links logs to the corresponding service instances, components, and source types for contextualized analysis. By default, automatic log context mapping is enabled.

AI suggests the optimal log field for mapping to service instances and components. When you use the AI-suggested field or if it's the default field, an AI sparkle icon (✦) appears. You can select a different field if needed. If AI can't find an optimal match or if your selected field is missing from the sample log, HLA uses the system default.

Important: If a manually defined mapping script exists in the **Data Input Mapping** page, automatic mapping overwrites it.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. Under Integrations, select **Manage installed integrations**.
   The Integrations Launchpad appears.
4. On the **Installed integrations** tab, search for an integration and open it.
   The **Overview** screen displays.
5. Under **Next steps**, on the **Log context mapping** card, select **View mapping**.
   The **Log context mapping** table displays.
6. Select the relevant integration.
   The integration's **Log context mapping** page displays.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

> **Note:** If the selected integration has features that are unsupported in Service Operations Workspace, you're referred to the **Data Input Mapping** page. For more information, see Map raw log data manually.

7. **Optional:** Turn off automatic service instance and component mapping by using the toggle.
   If you turn off automatic mapping, you must define a mapping script manually in the **Data Input Mapping** page. For more information, see Map raw log data manually.

   The remaining steps in this procedure are only relevant when automatic mapping is enabled.

   > **Note:** When you open the **Log context mapping** form for the first time, the system default values display in all fields.

8. Configure the log field by which HLA identifies a service instance from the **Raw log field** drop-down list.
   A service instance represents a complete business process.
9. Configure the log field by which HLA identifies a component from the **Raw log field** drop-down list.
   A component represents the smallest part of a service instance. A single service instance typically comprises multiple components.
10. **Optional:** Select **Preview** to view the service instances and components the system has identified based on your settings in the previous steps.
    The service instances and components are grouped by service.
11. In **Advanced settings**, review the fallback service and set the log field by which the system identifies a source type.

    a. In the **Selected fallback service** field, review the fallback service the system uses if it can't identify the service instance in the log data.
    The displayed fallback service is the default service instance that you selected when you configured the integration.

    > **Note:** For some integrations, such as Filebeat, this field shows `N/A` instead of a value. In these integrations, if HLA can't identify the service instance, it uses the system default value rather than the fallback service.

    b. Configure the log field by which HLA identifies a source type from the **Raw log field** drop-down list.

    c. **Optional:** Select **Preview** to view the log sources the system has identified based on your source type settings.

12. **Optional:** In the Simulator, review a log processing simulation based on your current mapping.
    The provided simulation helps you understand how the current mapping impacts HLA log processing. Use this information to fine-tune your settings:
    - The **Log fields** tab shows the default service, component, and source type and the raw log.
    - The **Extracted values** tab shows the service instance, component, and source type that the system has extracted from the raw log based on your current settings.
13. Select **Save mapping** to save your settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View source type and log source relationships

Explore the many-to-many relationships between source types and log sources to help you optimize data input mapping in Health Log Analytics.

Before you begin

For an overview of mapping, see Log data auto-mapping and mapping in Health Log Analytics.

Role required: evt_mgmt_admin

About this task

Health Log Analytics uses source types to parse logs and utilizes log sources as the framework for detecting anomalies. Each source type provides logs that are analyzed by multiple log sources, and each log source receives logs that are parsed by multiple source types. Reviewing the associations between source types and log sources can help you improve your data input mapping for better performance and outcomes.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Mapping** > **Data Input Mapping**.
2. Verify that log data is streaming into your instance, and then open a data input mapping record.
3. View the associations between the log sources and the source types.
   - View the log sources that are associated with a source type:
     a. Open the **Source Types** tab to view the generated source type structures.
     b. Open a source type structure record.
     c. Select a source type.
     d. Under **Sources**, view all the log sources that receive data from this source type.
   - View the source types that are associated with a log source:
     a. Open the **Sources** tab to view the generated log sources.
     b. Select a log source record.
     c. On the **Source Types** tab, view all the source types that feed data into this log source.

Note: The displayed results are updated once a day. Results that are older than 24 hours are deleted automatically.

Source type structure adjustment in Health Log Analytics

Health Log Analytics (HLA) enables you to reclassify auto-classified log properties and change auto-mapped labels. These adjustments help HLA machine learning analyze your data accurately.

What HLA does automatically

HLA automatically separates the transport header from the inner log message and sends only the inner log message to the source type structure. The system extracts properties from incoming log messages and auto-maps labels to source type fields.

- Classification determines how the HLA engine analyzes and processes each field. Available classifications are: METER, GAUGE, ARC_ONLY, HISTOGRAM, INVALID.
- Labels tell HLA what role a field plays in the log structure. Available labels are: MESSAGE, HOST, TIMESTAMP, SEVERITY, EVENT-ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Reviewing and, if necessary, modifying automatically mapped labels and auto-classified properties in the source type structure verifies the HLA engine interprets log fields correctly.

Note: Because a single data input can contain multiple source types, HLA structures log data by source type rather than by data input.

AI-assisted log parsing with ServiceNow Otto

AI-assisted log parsing using ServiceNow Otto provides recommendations for optimal classifications and labels for parsed log properties based on your log patterns. This approach reduces manual effort and improves accuracy without requiring custom JavaScript functions.

To enable this capability, activate the integration with AI-powered capabilities by selecting the **Activate with AI** option. This option enables AI-powered automatic parsing and mapping of log data. When the integration is activated, ServiceNow Otto collects and analyzes log data and displays suggested classifications and labels for your source type properties. An AI icon (✨) indicates that AI is active. Verify that ServiceNow Otto correctly classified and labeled your log properties. You can override the suggestions by selecting a different classification or label from each list, if needed.

While ServiceNow Otto generates a parser automatically for a source type, a notification banner on the **Source Type Structure** page shows the current status of the generation process.

Example

The following example shows how to modify values of extracted properties in the source type structure with JavaScript.

Consider the following log:

```
{
  "TenantId": "abc-01-02-03-04-05050708091011121314",
  "@timestamp": "2020-08-28T08:29:23.967Z",
  "Computer": "john Doe_computer",
  "EventType_s": "LogMessage",
  "Job_s": "johnDoe_cell",
  "IP_s": "1.00.00.00",
  "message": "This is the extracted message. This part of the
 message includes superfluous content and values",
  "MessageType_s": "OUT",
  "Timestamp_d": 1598603359017850000,
  "Type": "my_LogMessage_is",
  "_ResourceId": ""
}
```

The sample code contains `"key":"value"` pairs: The key is the property name. The value is the property value.

The key "message" in the sample has the following value: `"This is the extracted message. This part of the message includes superfluous content and values"`.

If you wanted your logs to contain only the meaningful part of that message, you would add JavaScript code instructing the system to extract only that part:

```
//Added JavaScript to extract only the first sentence in the
 message! if (output['message'] != null){ output['message'] =
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
output['message'].slice(0, output['message'].indexOf("\.")); }
(edited)
```

You could use the same logic to reclassify a value. For example, if the key "Computer" is insignificant, you could set its value to "Invalid."

For more information about the source type structure, see the Source Type Structure – Labels & Classifications [KB0863562] article in the Now Support Knowledge Base.

Note: You can reactivate the learning mode for a source type structure to use the HLA engine's learning functionality again after the initial learning period becomes outdated. On the **Source Type Structure** form, select **Relearn** to reset parsing, extraction, classification, and labeling for the source type structure and restart the learning process.

Refine the source type structure in Health Log Analytics

Fine-tune how Health Log Analytics reads your inner log messages and detects anomalies by customizing the extracted properties in the source type structure.

Before you begin

For an overview, see Source type structure adjustment in Health Log Analytics.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Inputs** > **Source Type Structure**.
2. Open a record.

   Note:

   - Health Log Analytics automatically separates the transport header from the inner log message when **Auto extraction** is selected (default).
   - Health Log Analytics fetches log samples automatically the first time the **Source Type Structure** form displays. In subsequent sessions, fetch the latest samples by selecting **Refresh Samples**.

3. **Optional:** Get AI-powered classification and labeling suggestions.

   AI can classify parsed log properties and suggest appropriate labels.

   a. Select **Suggest classifications (AI)** from the **Source Type Structure** page.

      On the **Key/Value Mapping** tab, AI-suggested classifications and labels appear in the **AI-suggested classification** and **AI-suggested label** columns.

   b. Review the suggestions.

   c. In the **Classification** and **Labels to assign** columns, select your preferred classifications and labels from the lists.

4. **Optional:** View how the current JavaScript function affects the log lines.

   a. Add a sample message in the **Test manual sample** field.

   b. Select **Go**.

   c. In the **Key/Value Mapping** tab, note how the JavaScript function affects the log lines.

5. In the **Raw input sample** field, choose a log message.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When you test your JavaScript function, Health Log Analytics uses this sample message to show the joint effect of automatic extraction and the JavaScript function on the log lines.

The following fields are read only:

Field	Description
Duration (ms)	The processing time of all samples, in milliseconds.
Dropped	The total number of logs dropped in all the samples.
Errors	The total number of errors that occurred in all the samples.
Timestamp extraction failures	The number of timestamp extraction failures that occurred in all the samples.
Severity extraction failures	The total number of severity extraction failures that occurred in all the samples.
Message extraction failures	The total number of message extraction failures that occurred in all the samples.
Lengthy properties	The total number of lengthy properties in all the samples. Lengthy properties are properties with more than 256 characters. **Note:** Because Health Log Analytics doesn't extract such properties, they aren't indexed as keywords in Elasticsearch.

6. Define a JavaScript function that customizes automatically extracted properties or adds properties to the source type structure.

   a. In the JavaScript console, either change the provided default JavaScript function, modify an existing custom JavaScript function, or define a new one.

   Note: In addition to the default JavaScript function, Health Log Analytics provides several JS function templates for refining the source type structure. The templates can act as a starting point for your custom script code. This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.

JS function templates	Description
Extract Key_Values Using Regex	Script used for parsing logs into a series of key:value pairs by using regular expressions to identify specific regex patterns. Depending on the log data being sent to the source type, it parses either the inner message or the full log line. You can view

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

JS function templates	Description
	how the data appears in the samples that are loaded into the source type structure. This process is iterative. It ’s repeated until all key:value pairs are found.
Extract Key_Values Using split-regex	Script used for parsing logs into a series of key:value pairs by using regular expressions to identify specific regex patterns. This process captures the values in a LIST form first, then uses the Java split function to create the key:value pairs. Depending on how the data appears in the raw input, this function might be more efficient than `Extract Key_Values Using Regex`. The raw input is either the inner message that header detection passes on, or the full raw log. If header detection is turned off or doesn't work on that particular log, the full raw log is used.
JSON Parse - flatten	Script used for extracting JSON information that is part of another text string from the raw input. For example, the information could be a JSON request written as part of a longer inner message. Inner JSON sections are difficult to split. In such cases, this script can be used to "flatten" or parse out the key:value pairs.
Extract a new field from raw input	Script that extracts a new field from the raw input, using a regular expression with capture groups to identify the new field's pattern.
Parse XML format	Script used for extracting key:value pairs from the XML by using a regular expression to identify the XML format. This process is iterative. It's repeated until all key:value pairs are found.
Set numerical severity level as textual values	Script used for converting numerical severity values to their corresponding textual severity values. **Note:** To enable the system to identify the severity level of a log correctly, severity must be given in text format. No numerical severity values should remain.
Trim function	Script used for removing the double quotes that surround the string from VALUE.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

JS function templates	Description
	You can adapt this function to remove any other characters that surround the output VALUE of a key:value pair.

The JavaScript function for refining the source type structure uses the following objects:

Signature - function construct (sample, output)

Object	Description
sample	The inner message extracted from the sample message.
output	Object that contains the key-value pair map.

b. Test the JavaScript function by selecting **Test**.

c. View the outcome of the JavaScript function in the related lists and modify if needed.

- The **Key/Value Mapping** tab shows the effect of your JavaScript function combined with the system's automatic extraction on your raw input sample.

  Modify keys where appropriate.

  - The **Classification** field enables you to reset the classification of a property. The available types are:

Classification

Type	Description	Example
Meter	A property with this classification detects anomalies in the number of times the property appears in each log message. It presents changes in the quantity of this value as part of Automatic Root Cause analysis. **Note:** Properties classified as Meter consume resources.	Status codes, response codes, actions, or patterns
Gauge	A property with this classification detects anomalies in a numerical value that is reported continuously.	CPU, memory, or response time

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Description	Example
	Note: Properties with a classification of Gauge consume resources.	
Timeless Gauge	A property with this classification detects anomalies in a numerical value that is not reported continuously. The system notifies about an anomaly in this value regardless when the anomaly occurred. **Note:** This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.	
Automatic Root Cause (ARC Only)	A property with this classification reports the property as part of the Automatic Root Cause analysis of another anomaly and not as an anomaly by itself.	User name, IP addresses, application components, or datacenter
Invalid	A property with this classification isn't calculated or shown in the Automatic Root Cause analysis. **Note:** Properties classified as Invalid save resources.	

- The **Labels to assign** field enables you to set a property on a label.

Labels to assign

Label	Description
Timestamp	The property that holds the timestamp of the event.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Label	Description
	Note: If automatic header properties detection is enabled and the source type lacks a timestamp, the system extracts the timestamp from the transport header. If header detection is turned off or doesn't work for the relevant log lines in the data input, all log lines must have proper timestamps.
Severity	The property that represents the severity level of the log.
Message	The log message. The system uses this property to identify textual patterns in the data.
Host	The property that represents the host from which the event was sent. **Note:** If automatic header properties detection is enabled and the source type lacks a host, the system extracts the host from the transport header. If header detection is turned off or doesn't work for the relevant log lines in the data input, all log lines must have a host.
External ID	The property that serves as a unique identifier for this event type. For example, Event ID in Windows event log.

    - The **Rename key** field enables you to rename the key.
- The **Outcome Key-Value** tab shows how your JavaScript function processed the data.
- The **Data Input Mappings** tab shows the data input mappings for the current source type.

  The **Last event time** column on this tab shows the last date and time that the AI engine processed a log for the data input mapping. The system refreshes this information every five minutes. Knowing when the last log was processed for this data input mapping makes it easier to verify that log data is streaming.

- The remaining tabs show errors, message extraction failures, severity extraction failures, timestamp extraction failures, and lengthy properties.

Note: If your new JavaScript function isn't behaving as expected, you can revert to the last published one by selecting the **Revert JS Function** related link.

d. Optional: Make any necessary adjustments and then test the JavaScript function again.

7. Select the **Save template** option to save the JavaScript function.
   You can either save the JavaScript function as a new template or overwrite the currently selected template.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To save the JavaScript function as a new template, enter a new name in the **Template name** field.
- To overwrite the template currently selected in the **JS function templates** field, leave the **Template name** field empty.

8. Select **Publish** to save the JavaScript function to the database.

Result
When the JavaScript function is published, Health Log Analytics uses it to fine-tune how it reads your inner log messages and detects anomalies.

The new script is automatically added to the list of JS function templates from which you can choose. This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.

What to do next
Continue with the remaining data input setup tasks: Verify that all your log sources are present and active, and add timestamp formats if necessary.

Related topics
Source type structure adjustment in Health Log Analytics

Health Log Analytics content packs for quicker time to value
Shorten onboarding time for the Health Log Analytics application by installing content packs. The packs contain default source types and mapping script templates that save you the time it takes to create them from scratch.

After installing a content pack, Health Log Analytics begins to stream and parse your log data immediately. The application recognizes and analyzes the data, generates more relevant alerts, and reduces noise from the moment of the integration. You can still customize the mapping configuration later on, if needed.

You can integrate content packs in two different ways:

- Select a content pack when you're configuring a data input. Health Log Analytics activates the selected pack automatically and uses its mapping script for mapping the data input sources.
- Activate a content pack before you configure a data input. When you have configured the data input and you're mapping the raw log data, select the content pack's mapping script from the **Data Input Mapping** form.

Currently, a content pack is available for streaming logs from Linux using Filebeat.

Note: All configurations in content packs exist in the global domain. If you're using domain separation, the content pack content is only available for logs in the global domain.

This feature is supported in the Health Log Analytics application, Version 25.0.17 - November 2022 and later, available from the ServiceNow Store.

Activate a content pack in Health Log Analytics
You can activate a content pack before configuring a data input.

Before you begin
For an overview of this feature, see Health Log Analytics content packs for quicker time to value.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Content Packs**.
2. Open the appropriate content pack for your data input.

   Note: Currently, a content pack is available for streaming logs from Linux using Filebeat.

3. On the **Content Pack** form, select **Activate**.
   The **Active** check-box appears selected. The activated content pack displays the script templates and source types that are associated with it.

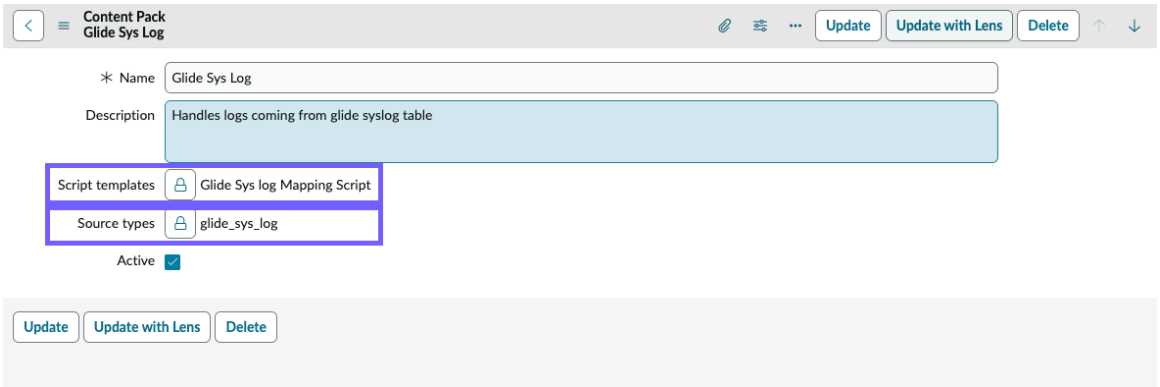

Note: The source types included in the content pack may already exist on your ServiceNow instance. In this case, Health Log Analytics automatically renames the content pack source types by adding a content pack suffix to their name. If you want to use the content pack source types and their mapping, delete or rename the original source types and then remove the suffix from the content pack source types. Otherwise, Health Log Analytics uses the original source types.

On the **Data Input Mapping** form, the mapping script is listed as a JavaScript functions template option.

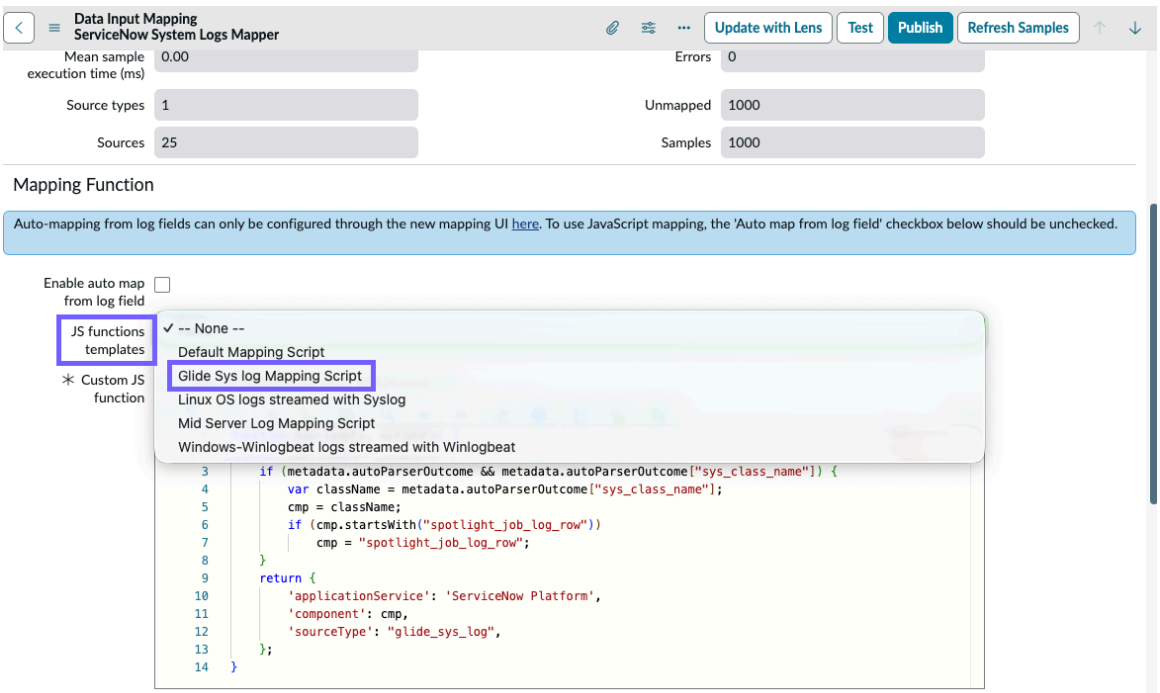

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

When you have configured the data input and you're mapping your raw log data, you can select the content pack's mapping script from the **JS functions templates** drop-down.

Migrating a Health Log Analytics data input configuration between instances

Export a Health Log Analytics data input and source types configuration as an update set and import it to a different ServiceNow instance. In the target environment, you can use the migrated data input for streaming and processing log data. This functionality saves time and reduces possible errors by avoiding the need to configure the settings again on the target instance.

The migrated configuration includes mapping and all other relevant settings, except credentials. Depending on the data input that you're exporting, you may need to configure credentials in the target environment. For example, when exporting Amazon CloudWatch or Amazon S3 data inputs, you must configure Amazon Web Services (AWS) credentials.

This feature is supported in the Health Log Analytics application, Version 25.0.17 - November 2022 and later, available from the ServiceNow Store.

Starting with Health Log Analytics Version 38.0.17 - December 2025, you can also export all source types related to one or more selected log sources to the update set together. For more information, see Export source types by log source.

For general information about update sets in the ServiceNow AI Platform, see System update sets.

Export a Health Log Analytics data input configuration

Export the configuration of a Health Log Analytics data input with or without the related source types to an update set. You can then import the update set to the target environment.

Before you begin

For an overview of this feature, see Migrating a Health Log Analytics data input configuration between instances.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Data Inputs**.
2. In the Data Inputs list, select the data input that you want to export.
3. Export the data input to an update set.

> **Note:** The exported configuration doesn't include the service instance configured for the data input.

   a. From the **Actions on selected rows** drop-down menu, choose **Export Data Input**.

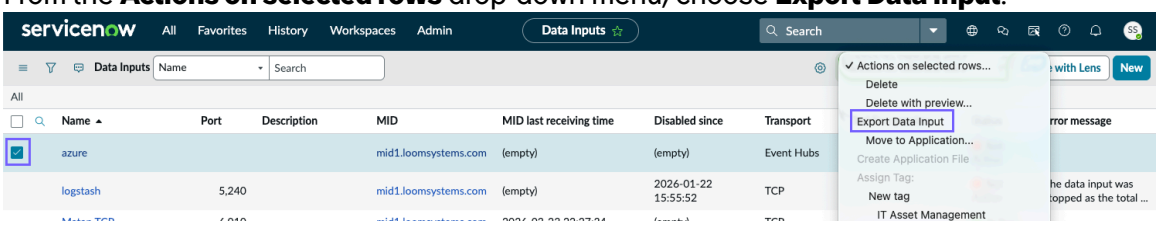

   b. In the **Export Data Input** dialog, do one of the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

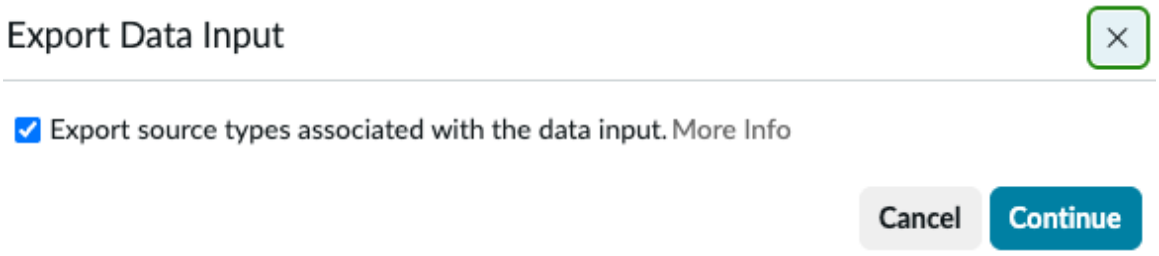

- Leave the check box selected to include the source types that are associated with the data input in the update set.

  When the box is selected, the update set will include the following source type-related tables:

  - sn_occ_source_type
  - sn_occ_source_type_structure
  - sn_occ_sources
  - sn_occ_timestamp_formats
  - sn_occ_lexical_keywords
  - sn_occ_log_correlator
  - sn_occ_log_source_retention
  - sn_occ_quick_replacements
  - sn_occ_reconstructors
  - sn_occ_replacements
- Clear the box to export the data input without the associated source types. For example, if you want to reuse the data input configuration, but configure the source types from scratch.

  When the box is clear, the update set will not include the source type-related tables.

  Note: You can export source types separately. For more information, see Export source types to an update set in Health Log Analytics.

c. Select **Continue**.

The data input, with or without the source types, is exported to an update set.

Note: An update set can act as the parent for multiple child update sets.

4. Select **Download update set** to download the update set to the instance.
The update set is downloaded as an XML file.

What to do next

Import the update set to the required ServiceNow instance. For more information, see Import a Health Log Analytics data input to a target instance.

Import a Health Log Analytics data input to a target instance

Import a data input configuration that you've exported from another instance as an update set.

Before you begin

For an overview of this feature, see Migrating a Health Log Analytics data input configuration between instances.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before performing this procedure, export the data input to an update set. For more information, see Export a Health Log Analytics data input configuration.

Role required: admin

Procedure

1. On the target instance, import the update set that contains the data input configuration.
   a. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**.
   b. In the Related Links, select **Import Update Set from XML**.
   c. Select **Choose file** and open the update set XML file from the Downloads folder.
   d. Select **Upload**.
   The update set is uploaded to the instance and appears in the **Retrieved Update Sets** list in the Loaded state.
2. Identify and address errors in the update set.
   a. In the **Retrieved Update Sets** list, open the update set record.
   b. Select **Preview Update Set Batch** to detect errors.
   The **Batch Update Set Preview** dialog indicates failure, which is expected behavior.
   c. Close the dialog and resolve the errors in the **Update Set Preview Problems** related list.

      For more information, see Preview a remote update set.

      You can ignore errors that start with 'Could not find a record in [table]' for the following tables:

      - aws_credentials
      - azure_sas_credentials
      - azure_service_principal
      - basic_auth_credentials
      - cmdb_ci_service_auto - Service instance
      - ecc_agent - MID
      - ecc_agent_cluster - MID Server Cluster
      - kafka_ssl_credentials
      - sn_occ_data_input_preprocess_example
      - sn_occ_data_input_preprocess_example_row
      - sn_occ_gcp_credentials
      - sn_occ_mapper_example
      - sn_occ_source_type_examples
      - sn_occ_source_type_structure_data
   d. Select the errors that you ignored and then select the **Select All** check box.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. From the **Actions on selected rows...** drop-down menu, choose **Accept remote update**.

f. In the **Preview Problems for Batch** related list, select all remaining errors and accept the remote update as explained in the earlier steps.

3. When all errors in the update set have been resolved, select **Commit Update Set Batch**.
   The imported data input appears on the Data Inputs list. The import includes mapping settings and all other relevant configuration.
4. In the imported data input, configure the following.
   - The MID Server to which the logs stream.
   - Credentials.

     Note: Whether you need to configure credentials depends on the data input that you've imported. For example, you must configure Amazon Web Services (AWS) credentials for Amazon CloudWatch and Amazon S3 data inputs.

   - (Optional) The service instance to which to bind the log data.

   For more information, see the configuration documentation for the relevant data input: Set up data inputs in Health Log Analytics manually.

Result
You can begin using the data input for streaming and processing log data.

Export source types to an update set in Health Log Analytics

Export source types to an update set separate from the Health Log Analytics data input configuration. You can then import the update set to the target environment.

Before you begin
For an overview of this feature, see Migrating a Health Log Analytics data input configuration between instances.

Role required: evt_mgmt_admin

Procedure
1. Navigate to **All** > **Health Log Analytics** > **Source Types**.
2. In the Source Types list, select one or more source types that you want to export.
3. From the **Actions on selected rows** drop-down menu, choose **Export Source Types**.
   The source types are exported to an update set.

   Note: An update set can act as the parent for multiple child update sets.

4. Select **Download update set** to download the update set to the instance.
   The update set is downloaded as an XML file.

What to do next
Import the update set to the required ServiceNow instance. For more information, see Import Health Log Analytics source types to a target instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Export source types by log source

Export all source types related to one or more selected log sources to an update set together in Health Log Analytics. You can then import the update set to the target environment.

Before you begin

For an overview of this feature, see Migrating a Health Log Analytics data input configuration between instances.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Log Sources**.
2. In the Log Sources list, select one or more log sources for which you want to export the related source types.
3. From the **Actions on selected rows** drop-down menu, choose **Export Related Source Types**.
   The source types are exported to an update set.

   Note: An update set can act as the parent for multiple child update sets.

4. In the **Related Source Type Exporter** screen, select **Download update set** to download the update set to the instance.
   The update set is downloaded as an XML file.

What to do next

Import the update set to the required ServiceNow instance. For more information, see Import Health Log Analytics source types to a target instance.

Note: Only the related source types and tables for the selected log sources will be imported, not the data inputs.

Import Health Log Analytics source types to a target instance

Import source types that you've exported separate from the Health Log Analytics data input configuration.

Before you begin

For an overview of this feature, see Migrating a Health Log Analytics data input configuration between instances.

Before performing this procedure, export the source types to an update set. For more information, see Export source types to an update set in Health Log Analytics.

Role required: admin

Procedure

1. On the target instance, import the update set that contains the source types.

   a. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**.

   b. In the Related Links, select **Import Update Set from XML**.

   c. Select **Choose file** and open the update set XML file from the Downloads folder.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Select **Upload**.
The update set is uploaded to the instance and appears in the **Retrieved Update Sets** list in the Loaded state.

2. Identify and address any errors in the update set.

a. In the **Retrieved Update Sets** list, open the update set record.

b. Select **Preview Update Set Batch** to detect errors in the update set.
The **Batch Update Set Preview** dialog indicates failure, which is expected behavior.

c. Close the dialog and resolve the errors in the **Update Set Preview Problems** related list.
You can ignore errors that start with 'Could not find a record in [table]' for the following tables:

d. Select the errors that you ignored and then select the **Select All** check box.

e. From the **Actions on selected rows...** drop-down menu, choose **Accept remote update**.

f. In the **Preview Problems for Batch** related list, select all remaining errors and accept the remote update as explained in the earlier steps.

3. When all errors in the update set have been resolved, select **Commit Update Set Batch**.

Result
The imported source types appear in the Source Types list.

Check gRPC and network connectivity for Health Log Analytics on the MID Server

Verify that Health Log Analytics can function properly by performing diagnostic tests on the MID Server to check gRPC and network connectivity.

Before you begin
Role required: evt_mgmt_admin

About this task
Your organization's network and security configuration may be keeping Health Log Analytics from functioning properly. For example, complex configurations with proxy and firewall settings might block outbound connections. To avoid operational issues, you can test gRPC and network connectivity on the MID Server.

Procedure

1. Navigate to **All** > **MID Server** > **Servers**.
2. Open the relevant MID Server record.
3. Under Related Links, select **Test gRPC connection**.
The system performs gRPC connectivity tests, such as domain name resolution, TCP/IP, and HTTP connection. When the testing is complete, the test results are displayed. Successful tests are preceded by a green check-mark, while failed tests are preceded by a red x mark.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

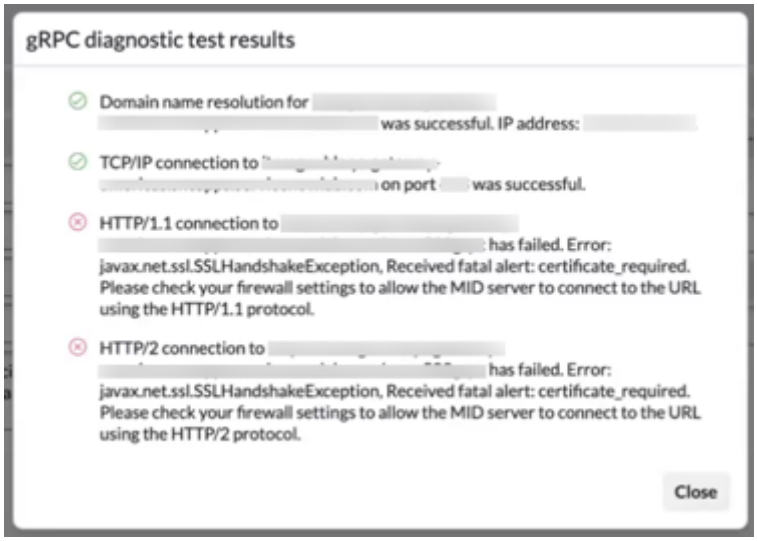

4. Fix the issues that caused any tests to fail by following the instructions on the **gRPC diagnostic test results** screen.
5. Select **Test gRPC connection** again and confirm that all the tests are completed successfully.

Administering Health Log Analytics

This section covers tasks involved in Health Log Analytics administration. It provides you with information needed to keep Health Log Analytics running efficiently.

Configure Health Log Analytics system properties for a single log source

Configure system properties for a specific log source if you need to alter the default values, which should not usually be necessary.

Before you begin

Role required: evt_mgmt_admin

About this task

Changing Health Log Analytics system properties affects back-end processes and can produce unexpected results. Therefore, if you must make changes to these properties, keep a record of your adjustments. For example, note the value before and after a modification and the time of the change.

Note: To modify a Health Log Analytics system property whose **Requires restart** value is **true**, contact ServiceNow support to modify it for you.

For a complete description of all the Health Log Analytics system properties, see the HLA System Properties [KB0869842] article in the Now Support Knowledge Base.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Data Input** > **Log Sources**.
2. On the **Source properties** tab, right-click a property.
3. Select **Override property value** from the pop-up menu.
4. In the **Value** field, enter the appropriate value.
5. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The new value overrides the previous setting for the selected log source.

Related topics
Configure global Health Log Analytics system properties

Configure global Health Log Analytics system properties

Configure global Health Log Analytics system properties if you need to alter the default values, which typically should not be necessary.

Before you begin
Role required: evt_mgmt_admin

About this task
Changing Health Log Analytics system properties affects back-end processes and can produce unexpected results. Therefore, if you must make changes to these properties, keep a record of your adjustments. For example, note the value before and after a modification and the time of the change.

Note: To modify a Health Log Analytics system property whose **Requires restart** value is **true**, contact ServiceNow support to modify it for you.

For a complete description of all the Health Log Analytics system properties, see the HLA System Properties [KB0869842] article in the Now Support Knowledge Base.

Procedure
1. Navigate to **All** > **Health Log Analytics** > **Health Log Analytics Administration** > **System Properties**.
2. Select a property with the scope Source.
3. In the **Value** field, enter the appropriate value.
4. Select **Update**.
   The setting will apply to log sources created from this point forward.
5. Determine the scope of application of the new setting.
   - To apply the setting only to log sources created from this point forward, select **Update**.
   - To override values for this property in all existing log sources and apply the modified setting to all new log sources, select **Override source property**.

Related topics
Configure Health Log Analytics system properties for a single log source

Enable or disable Health Log Analytics system features

Customize the basic configuration of Health Log Analytics by enabling or disabling system features.

Before you begin
Role required: evt_mgmt_admin

Procedure
1. Navigate to **All** > **Health Log Analytics** > **Health Log Analytics Administration** > **Features**.
   By default, the table shows the system features that administrators modify most commonly. To view the complete list of features, use the filtering options.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Select a feature record from the table.
   The record contains the following information:
   - The name of the feature
   - A description of the feature's purpose
   - The default state of the feature
   - The current state of the feature
3. Set the state of a system feature to ON to enable the feature or to OFF to disable it.
4. Select **Update**.

View log anomaly metrics in Health Log Analytics

View and query all metrics that Health Log Analytics tracks from a single location.

Before you begin

Role required: evt_mgmt_admin

About this task

The Metric Data table provides a centralized view of every metric that HLA collects and monitors. It includes pattern metrics, custom metrics, and raw metrics. Use the table for the following tasks:

- Query metrics by log source, integration, source type, and metric type to narrow down which metrics HLA is tracking for a data input or environment.
- Review anomaly detection statistics to see how often HLA detects anomalies for a given metric.
- View metric baselines to understand what HLA records as normal behavior.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Log Anomaly Detection** > **All Metrics**.
   The **Metric Data** table displays a record for each metric that HLA is collecting. For a description of the columns in the table, see Metric Data table columns in Health Log Analytics.
2. **Optional:** Filter the table by log source, integration, source type, or metric type.
   Filtering enables you to narrow the results to the metrics relevant to a specific data input or environment.
3. Select a metric record to view its details.

Health Log Analytics alert notification in Slack or Microsoft Teams channels

Health Log Analytics sends notifications for new anomaly alerts in real time to Slack or Microsoft Teams channels. This built-in functionality frees you from having to continuously monitor logs to prevent incidents.

The notifications include the number of the alert, its description, severity, configuration item, state, and priority group. You can filter the alerts for which notifications are sent to your channel. For example, you may only want to receive alerts of a Major severity or a High priority group. By default, no notifications are sent on secondary alerts.

To support channel notification, the channel must be configured to receive Health Log Analytics notifications.

Starting with Version 33.0.27 - August 2024, Health Log Analytics uses the Microsoft Teams Graph spoke application to send notifications on new anomaly alerts. The application replaces the Microsoft Teams spoke plugin, which will be deprecated. If you're using an HLA version earlier than 33.0.27, follow the setup procedure in Set up Health Log Analytics alert notifications

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

in channels. Otherwise, see Set up Health Log Analytics alert notifications in a Microsoft Teams channel.

Sample HLA alert notification in Microsoft Teams

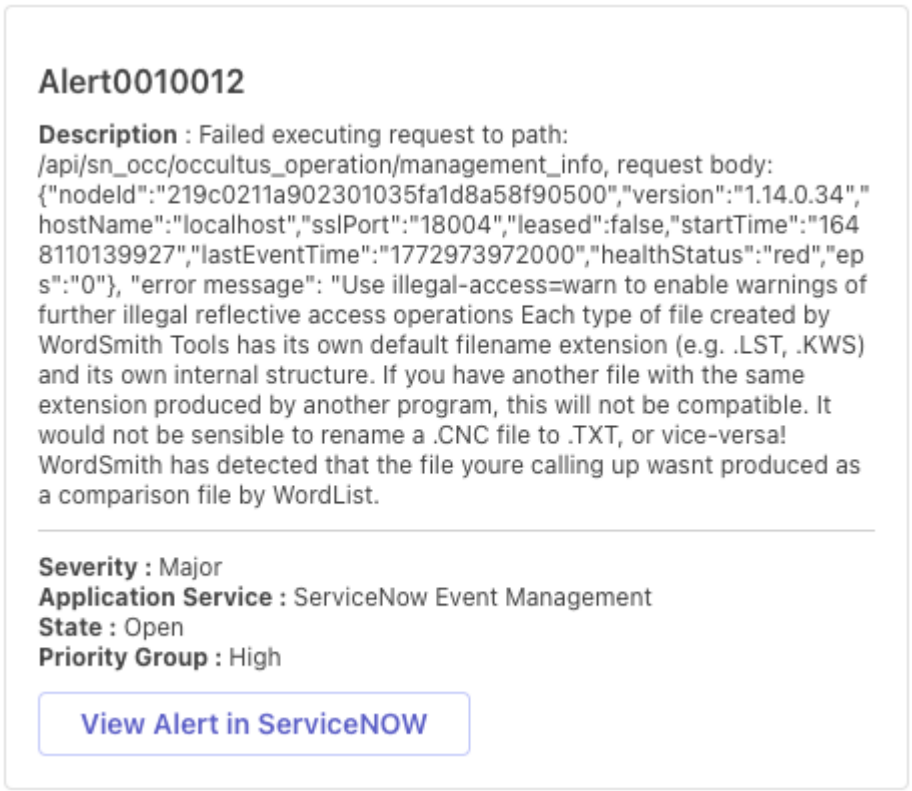

Set up Health Log Analytics alert notifications in channels

Set up notifications for new Health Log Analytics anomaly alerts in either a Slack or Microsoft Teams channel.

Before you begin

Starting with Version 33.0.27 - August 2024, Health Log Analytics uses the Microsoft Teams Graph spoke application to send notifications for new anomaly alerts. The new application replaces the Microsoft Teams spoke plugin, which will be deprecated. If you're using an HLA version earlier than 33.0.27, follow the setup procedure below. Otherwise, refer to Set up Health Log Analytics alert notifications in a Microsoft Teams channel.

Note: If you upgrade from an earlier version to Version 33.0.27 - August 2024 or later, redo the setup using the procedure in Set up Health Log Analytics alert notifications in a Microsoft Teams channel.

For alert notifications in Slack channels, verify that the plugin Slack Webhooks Spoke for ServiceNow Integration Hub (com.sn.slack.ah) is installed:

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel. For more information, see Request a plugin.

Role required: evt_mgmt_operator or evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create a Webhook URL for a channel in Slack or Microsoft Teams.

   Health Log Analytics uses the Webhook to notify about the alerts in the channel.

2. Map the Webhook URL for the channel to the service instance associated with the alert.

   a. Navigate to **Health Log Analytics Administration** > **Webhook Notification**.

   b. Select **New**.

   c. On the form, fill in the fields.

   Alert Notification Webhook form

Field	Description
Webhook	The Webhook URL that you created for the channel.
Active	Select to activate the Webhook.
Configuration item	The service instance associated with the alert. If this field is empty, the system uses the default property for the channel: default.slack.webhook.url for Slack or default.teams.webhook.url for Microsoft Teams.

   d. Select **Submit**.

3. Activate the Health Log Analytics Alert Management rule for submitting alerts to the appropriate channel.

   a. Navigate to **Event Management** > **Rules** > **Alert Management Rules**.

   b. Locate the appropriate rule:
      - For Slack: Submit HLA alerts to slack.
      - For Microsoft Teams: Submit HLA alerts to teams.

   c. In the **Active** column for the rule, set the value to `true`.

4. **Optional:** Filter the alerts on which notifications will be sent to the channel.
   a. In the Alert Management Rules table, select the rule for your channel.
   b. Select **Alert Filter**.
   c. On the form, set the conditions that must be met for sending a notification on the alert to the channel.
   d. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up Health Log Analytics alert notifications in a Microsoft Teams channel

Set up notifications for new Health Log Analytics anomaly alerts in a Microsoft Teams channel.

Before you begin

Starting with Version 33.0.27 - August 2024, Health Log Analytics uses the Microsoft Teams Graph spoke application to send notifications on new anomaly alerts. For more information, see Microsoft Teams Graph Spoke. If you're using an HLA version earlier than 33.0.27, refer to the setup procedure in Set up Health Log Analytics alert notifications in channels.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Install the latest version of the Microsoft Teams Graph spoke from the ServiceNow Store.
2. Set up the Microsoft Teams Graph spoke.
   For instructions, see: Set up the Microsoft Teams Graph spoke.
3. Activate the Health Log Analytics Alert Management rule for submitting alerts to the Microsoft Teams channel.
   a. Navigate to **Event Management** > **Rules** > **Alert Management Rules**.
   b. Locate the appropriate rule: Submit HLA alerts to teams.
   c. In the **Active** column for the rule, set the value to `true`.
4. Map the service instance to the Microsoft Teams channel.
   a. In the **All** navigation menu, type `sn_occ_teams_alert_notification_config` and then press Enter.
   b. Select **New**.
   c. On the form, fill in the fields.

      Note:

      - You can obtain the Teams ID and Channel ID from the Microsoft Teams channel.
      - The **Configuration Item** field must correspond to the Service instance. If no CI is needed, leave this field empty. In this case, the record is used for alerts when there's no record that matches their CI.

   d. Select **Submit**.
   e. **Optional:** Repeat steps 1 to 4 to add additional configurations.

Storage space for log retention in Health Log Analytics

The Health Log Analytics base system typically provides 1,000 GB storage space for log source retention, although the storage capacity may vary depending on your license. The default retention time for logs is three days, but you can modify that period.

This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. To activate it, contact your ServiceNow account manager or ServiceNow Support.

When the storage space used by all log sources exceeds 80 percent of the total, you receive an email notification. You can change the default threshold percentage by modifying the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

sn_occ.elastic_storage.utilization_limit.percent system property in the System Properties table.

Automatic reduction of log retention days

By default, logs are retained for three days. You can adjust this period to 3, 7, 14, or 30 days. For more information, see Modify the log source retention period in Health Log Analytics.

If the storage space used by all log sources reaches 100 percent, the number of retention days allocated per source is automatically lowered by one level and you receive an email notification with the specifics. The automatic reduction process starts with the sources that were allocated the highest number of retention days. For example, for a source with 30 retention days, the number is reduced to 14. If the source has 14 retention days, the number is reduced to 7, and so on.

You can change the threshold percentage for the automatic reduction of log retention days by modifying the *sn_occ.log_retention.automatic_reduction_utilization_percent* system property in the System Properties table.

Related topics

Modify the log source retention period in Health Log Analytics

Modify the log source retention period in Health Log Analytics

Modify the period that Health Log Analytics retains logs from a specific source. You can calculate the impact of your intended change on storage to help you make an informed decision.

Before you begin

This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. To activate it, contact your ServiceNow account manager or ServiceNow Support.

To display **Log source retention** in the filter navigator:

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Enter the property name *sn_occ.activate_log_source_retention* in the Search box.
3. Ensure that the property value is set to **true**.

Role required: evt_mgmt_admin

About this task

By default, the log retention duration per source is set to three days. You can change this number to 7, 14, or 30 days. For more information, see Storage space for log retention in Health Log Analytics. Note that longer log retention results in higher storage usage.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Log source retention**.
2. **Optional:** Calculate how changing the retention period for logs from a particular source affects storage.
   a. Open a log source record.
   b. In the Related Links section, select **Calculate how retention impacts storage**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. From the New retention list, select the proposed number of retention days for logs from this source.

d. Select **Calculate**.
The **Projected storage usage (GB)** field displays the amount of storage space in GB that the source would use if the new policy were implemented.

> **Note:** Health Log Analytics bases its calculation of projected storage usage on the volume of log data that was streamed in the past. The estimate could change as the volume of streamed log data varies.

e. Compare the projected usage to the amount of storage needed for the current policy that is shown in the **Current storage usage (GB)** field.

f. Select **Cancel** to close the dialog box.

3. Update the log retention policy per source or for several sources together.
   - To update log retention for a single source:
     a. Open a log source record.
     b. From the Retention days list, select the required number of retention days for logs from this source.
     c. Select **Update**.
   - To update log retention for multiple sources:
     a. Select several log source records.
     b. From the Actions on selected rows choice list, select **Update multiple source retentions**.
     c. From the New retention policy list, select the required number of retention days for logs from the selected sources.
     d. Select **Update**.

   The updated policy takes effect in up to 15 minutes.
4. **Optional:** Restore the default log retention policy per source or for all log sources.
   - To restore the default policy for a single source:
     a. Open a log source record from the table.
     b. In the Related Links section of the record, select **Restore default**.

     The log retention policy for this log source reverts to the default value of three days.
   - To restore the default policy for all sources:
     a. In the Related Links section, select **Restore defaults**.
     b. In the confirmation dialog, select **Continue**.

     The log retention policy for all log sources reverts to the default value of three days.

Scaling Health Log Analytics to stream logs at a higher rate

Stream log data to Health Log Analytics in a scalable, more stable way using the advanced ServiceNow infrastructure.

The new ServiceNow architecture leverages queuing technology in IT Operations Management Cloud Services for authentication and log ingestion. The Health Log Analytics AI engine has

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

been enhanced to scale dynamically in response to increased log ingestion by your organization. Health Log Analytics can be scaled to support more than 50,000 events per second (EPS) or more than 10,000 subscription units (SUs).

You can switch to the advanced ServiceNow infrastructure by submitting a scaling request through your ServiceNow account manager. When your scaling request is approved, ServiceNow provisions the number of Elasticsearch and AI Engine nodes required to support your subscription units.

Enhancing the CMDB with host data Health Log Analytics found in logs

When Health Log Analytics streams logs, it extracts host data from the log events. If host data is discovered that doesn't match the information in the Configuration Management Database (CMDB), the system creates a configuration item (CI) candidate. You can review the candidate based on the data found in the logs.

After reviewing a CI candidate, you can choose to either add it to the CMDB or dismiss it as irrelevant. The **Log-Based CI Candidates** table lists all CI candidates that Health Log Analytics has created, sorted by system update time. The **Operation** column shows the suggested action to take on each CI candidate. You are automatically notified when there are new CI candidates for your consideration. For more information, see Add log-based CIs to the CMDB in Health Log Analytics.

For a description of the related system properties, see the "logsourceinfo (CMDB)" section in the HLA System Properties [KB0869842] article in the Now Support Knowledge Base.

This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. You can use it on the ServiceNow AI Platform using the Rome Patch 6 release (scheduled to be released on January 27, 2022) or later.

Add log-based CIs to the CMDB in Health Log Analytics

Keep the Configuration Management Database (CMDB) updated with host data Health Log Analytics discovered in your logs by adding log-based configuration items (CIs) to the database.

Before you begin

This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. You can use it on the ServiceNow AI Platform using the Rome Patch 6 release (scheduled to be released on January 27, 2022) or later.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Health Log Analytics Administration** > **Log-Based CI Candidates**.
2. Handle a single CI candidate or multiple CI candidates together.

Option	Procedure
Take action on a single CI candidate	**a.** Select a CI candidate record and review the information on the form. For field descriptions, see Log-based CI candidates fields in Health Log Analytics. **b.** Add the CI candidate to the CMDB or delete it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Procedure
	▪ If the CI candidate is appropriate, either add it to the CMDB as is or include more information in it before adding it to the CMDB. ▪ To add the CI candidate to the CMDB as is, select **Add to CMDB** and then se lect **Continue**. ▪ To add more information, select **Edit ad ditional CI fields**. Enter data that was not extracted from the logs, such as the name of the company associated with the CI, and then select **Submit**. The CI is added to the CMDB. In the CMDB table for the relevant class type, the discovery source for the new CI is Health Log Analytics. The CI candidate record is automatically deleted from the Log-based CI candidates table. ▪ If the CI candidate is irrelevant, delete it by selecting **Delete**.
Take action on multiple CI candidates at the same time	**Note:** This option does not enable you to add extra information to the selected CIs. ○ Add multiple CIs to the CMDB. **a.** In the Log-based CI candidates table, se lect the CI candidates you want to add to the CMDB. **b.** From the Actions on selected rows list at the bottom of the page, select **Add to CMDB**. The selected CIs are added to the CMDB. The CI candidate records are automatical ly deleted from the Log-based CI candi dates table. ○ Delete multiple CI candidates. **a.** In the Log-based CI candidates table, se lect the CI candidates you want to delete. **b.** From the Actions on selected rows list at the bottom of the page, select **Delete**.

Analyzing and resolving Log Analytics alerts

Analyze and resolve Log Analytics alerts by investigating log data and taking action to resolve the underlying issue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview of analyzing and resolving a Log Analytics alert

As an Operator, you're responsible for analyzing and resolving the alerts that Health Log Analytics generates. When HLA creates an alert, you review the alert's severity, the affected configuration item (CI), the log data associated with the anomaly, and the impacted services. You try to identify the root cause by investigating the logs that surround the anomaly.

In the Express List, review alert details and use ServiceNow Otto to get an in-depth analysis of the alert and potential resolutions in straightforward, human-readable language. By drilling down into the alert, you can quickly identify the issue and proceed to resolve it before it affects your users.

Using the Log Viewer, you can browse the alert logs by timestamp or range for further investigation. You can visualize the frequency of anomalous log lines in a chart.

More detailed information on tasks and procedures for analyzing and resolving Log Analytics is available via the following links.

- Start remediation of a Log Analytics alert from the Overview tab

  Begin the remediation process of a Log Analytics alert from the alert **Overview** tab. This tab provides information on the alert, log data associated with the anomalous behavior, CIs associated with the alert, and services impacted by it.

- Analyze log lines around an anomalyAnalyze the logs that surround the anomaly

  Review the log lines surrounding the anomaly for clues about the state of faulting systems. This information can help you narrow down the root cause of the alert.

- Use log correlators to identify relationships in log data

  Identify relationships between alerts to help you determine whether an alert is part of a larger issue.

- Navigate to the Express List and select an alert from the Alerts list. Use ServiceNow Otto to get an in-depth analysis of the alert and potential resolutions. By drilling down into the alert, you can quickly identify the issue and proceed to resolve it.
- Review the logs for an alert on the Log viewer

  For further investigation, navigate to the **Log Viewer** to browse the alert logs by timestamp or time range. You can visualize anomaly frequency within a time period for a comprehensive view of log data over a specified time range.

- Add a KB article to a Log Analytics alert

  When you have resolved an alert that Health Log Analytics generated, you can add a knowledge base (KB) article to it. For example, you can provide information that might help others resolve similar issues.

For a brief explanation of key terms and concepts used in HLA, see the Health Log Analytics terminology.

Use cases

Use Case: Proactive monitoring of your ServiceNow instance in Health Log Analytics - Use Health Log Analytics to detect and resolve emerging issues in your organization's ServiceNow instance before they affect platform users.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Start remediation of a Log Analytics alert from the Overview tab

Begin the remediation process of a Log Analytics alert from the alert **Overview** tab. This tab provides information on the alert, the log data associated with the anomalous behavior, CIs associated with the alert, and services impacted by it.

Before you begin

Role required: evt_mgmt_operator, or evt_mgmt_user, or evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. In the Service Operations Workspace, select the lists icon ( ).
3. Select the appropriate list in the Alerts sub-list and navigate to the relevant alert or group.

   All Health Log Analytics alerts have the value **Log Analytics** in the **Source** column. The value in the **Group** column identifies the type of Log Analytics alert. See Types of Health Log Analytics alerts for a more detailed description of each type of alert or group.

   Tip: To preview an alert in the list, select its info icon ( ).

4. Select the alert number.

   The **Overview** tab displays.

5. View the section of the **Overview** tab that provides the information that you need.

   For a description of the sections and cards, see Sections and cards on the alert Overview tab in Health Log Analytics.

   Note: Because some sections on the **Overview** tab show only a portion of the information, many sections include a link that displays different or more complete information.

6. **Optional:** Run predefined remediation tasks to resolve alert issues.
   a. Select **Launch playbook**.
   b. Select the **Playbook** tab.

      The **Run remediation** card displays the available remediation actions.

   c. Select the tile corresponding to the action you want to run.
   When the action is complete, it appears in the **Completed** card.

Review logs surrounding an anomaly

When Health Log Analytics identifies an anomaly, viewing the logs that surround the anomaly provides clues about the state of faulting systems. This information can help you narrow down the root cause of an alert.

While viewing a Log Analytics alert, the **Surrounding logs** tab lists the log lines that were generated one minute before and one second after the anomaly occurred. The log lines are related to the metric or pattern that triggered the alert. The list is filtered to the relevant component.

Logs generated immediately before the anomaly might indicate warning signs or conditions that led to the issue. By reviewing these logs, you can trace the sequence of events that led to the anomaly, helping you understand what went wrong. Logs created right after the anomaly can

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

help you evaluate the impact of the event by analyzing how the system responded and whether other components were affected. The timeline of the logs around the anomaly can help clarify what might have caused the issue, for example a configuration change or a software update.

You can modify the time range of the displayed log lines. For example, investigate logs from more than one minute before the anomaly or more than a second afterward.

Note: Logs that surround the anomaly are retained and available for 30 days after the creation of the alert. The system does not delete these logs when the global retention period of logs expires. When the retention period expires, the surrounding logs are available only on the **Surrounding logs** tab and not in the Log viewer.

Health Log Analytics enables you to view the anomalous log data graphically on the Log viewer, accessed from the Service Operations Workspace. The Log viewer displays a chart of the frequency of anomalous log lines during one minute before and one minute after the Log Analytics alert. It lists the associated log lines and automatically shows the query that relates to the anomaly, the selected component, and the appropriate time filter. For more information, see Review alert-related logs on the Log Viewer.

Related topics

Analyze log lines around an anomaly

Analyze log lines around an anomaly

Review the log lines around an anomaly to help you identify the root cause of a Log Analytics alert.

Before you begin

For background information, see Review logs surrounding an anomaly.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. In the Service Operations Workspace, open a Log Analytics alert.
2. Select the **Surrounding logs** tab and review the information.
   The tab displays the list of log lines that were generated one minute before and one second after the Log Analytics alert. For an explanation of the information on the tab, see Surrounding logs tab fields.
3. **Optional:** View a different time span of the log lines using one of the following methods:
   - Update the time range relative to the alert to a predefined time range in the **Select range** list.
   - Specify a **Start time** or **End time** using the time picker.
4. **Optional:** View the anomalous log data graphically as a function of time by selecting **Log viewer**.

Log correlators for identifying related alerts

In Health Log Analytics, log correlators are keys or values in log data that detect correlations between alerts. These correlations help you determine whether an alert is part of a larger issue.

For example, a log correlator could detect when the interface ID of a particular network device occurs simultaneously in multiple warnings across service instances.

You can identify related alerts in your log data by using log correlators. The base system includes several log correlators. You can define custom correlators for a specific log source, all log sources, or only log sources created after the correlator is activated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Most log lines include a metadata portion plus a message portion. Some log lines, however, include only message text with metadata included in the text. The two types of log correlators, free text correlators and log property correlators, analyze different portions of each log to identify relationships between data from multiple sources.

Free text correlators

Free text correlators analyze the text within the log message portion of log lines that are associated with an anomaly. The system uses free text correlators to identify correlations between alerts. You use free text correlators to add a term that you expect to appear within log messages. A good choice is a term that is not structured and would not otherwise be extracted as a log property. For example, “policy-id” or “ thread-id”.

You also typically add free text correlators for the names of systems, applications, and services that are unique to your environment. Because such a value can be referred to by multiple sources, layers, middleware, or databases, the free text correlator can be an effective detector of correlated alerts. For example, if your organization's service is called TeaTime, then you might add "teatime" as a free text correlator. The correlator would identify alerts that are related because they were generated for resources that support the TeaTime service. For example, a database lock or a connection failure between TeaTime components.

Log property correlators

Log property correlators analyze the metadata portion of log lines. For example, the correlator can analyze the name of a service instance, the interface ID of a network device, or the request ID of a web-facing component. A log property correlator could flag a correlation when the interface ID of a network device simultaneously occurs in multiple warnings in different log sources. Log property correlators are specific to the business context of your environment.

You can specify the set of log sources whose log data are analyzed by a log correlator. Choices are as follows:

- **Only new sources**: The system applies the log correlator only to log lines from log sources that were created after this log correlator is activated.
- **All sources**: The system applies the log correlator to log lines from all log sources.
- **Specified source**: For a log correlator, the system analyzes only log lines from the log source that you specify.

Add a log correlator to identify related alerts

In Health Log Analytics, detect related alerts in log data by adding log correlators. The base system includes several log correlators and you can define custom log correlators.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task

For information about the types and functions of log correlators, see Log correlators for identifying related alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Use one of the following methods to add a log correlator.

Option	Procedure
Add a log correlator for a specific log source	**a.** Navigate to **Health Log Analytics** > **Log Anomaly Detection** > **Log Correlators**. The list of existing log correlators opens. **b.** Select the name of a log correlator. The names appear in the **Correlation indicator** column. **c.** Select **New**.
Add a log correlator that applies either to all log sources or to only those log sources that become active after you define this log correlator	**a.** Navigate to **Health Log Analytics** > **Data Input** > **Log Sources**. **b.** Select the name of the log source. The Log correlators related list displays the list of existing log correlators that analyze log data from the selected log source. **c.** On the **Log correlators** tab, select **New**.

2. Fill in the Log correlator form.
   For a description of the fields, see Log correlators form fields.
3. Select **Active** and then click **Submit**.

Exclude a source from a log correlator in Health Log Analytics

Prevent Health Log Analytics from analyzing log lines from a specific source by excluding that source from the log correlator.

Before you begin

Role required: evt_mgmt_operator, evt_mgmt_admin

Procedure

1. Navigate to **Health Log Analytics** > **Data Input** > **Log Sources**.
2. Open the relevant source on the Log Sources form.
3. On the **Log Correlators** tab, select the option for the log correlator.
4. In the Actions on selected items list, select **Exclude from this source**.

Review alert-related logs on the Log Viewer

The **Log Viewer** tab in Health Log Analytics lets you browse the logs for an alert by timestamp or time range, and visualize anomaly frequency within a specific time period. Customizing the displayed data and adjusting time filters enables you to better understand the framework in which the anomaly occurred, helping you find the root cause faster.

The Log Viewer presents all data connected with the Log Analytics alert. It shows the query that relates to the anomaly, the selected component, and the appropriate time filter. You can personalize the displayed data, and manually adjust the time range without affecting the other settings. The applied filters appear in the **Filters** pane. You can add or remove filters as needed to show only the data you want to view.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Log Viewer** displays a chart of the frequency of anomalous log lines during one minute before and one minute after the Log Analytics alert and lists the associated log data. This information helps you identify trends leading up to and following the event, providing context for root cause analysis.

As you analyze the logs for an alert on the **Log Viewer**, you can modify the query to fine-tune the search, save useful searches, and share them with others. For a description of the information displayed in the **Log Viewer** table, see Log Viewer table fields.

You can perform the following tasks on the **Log Viewer**:

- View log data for an alert in Health Log Analytics

  View a chart of the frequency of anomalous log lines and the associated log data.

- Define, save, and share a log data search

  Fine-tune the search query to help determine the causes of the issue, and save and share useful searches.

- Use or modify a saved log data search in Health Log Analytics

  Use a saved search. As the owner of a saved search, you can modify the search values and save your changes.

- Filter search results on the Log Viewer in Health Log Analytics

  Apply filters to show only the data you want to view.

- Customize the Log Viewer table in Health Log Analytics

  Add or remove columns in the table to show only the data you want to view.

If you discover an important metric in the log data, you can use it to define a new Log Analytics alert rule. For more information, see Define a custom Log Analytics alert rule in Health Log Analytics.

Related topics

Review logs surrounding an anomaly

View log data for an alert in Health Log Analytics

View a chart of the frequency of anomalous log lines and the associated log data on the Log Viewer.

Before you begin

Role required: evt_mgmt_operator, or evt_mgmt_admin

Procedure

1. Open the **Log Viewer** tab using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. **Optional:** Personalize the displayed data and how it is presented on the Log Viewer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Filter search results on the Log Viewer in Health Log Analytics to show only the data you want to view.
- Customize the Log Viewer table in Health Log Analytics by adding or removing columns.

These features are supported in the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store.

3. Analyze log-based anomalies and added context to help you achieve faster mean time to repair (MTTR).
For example, say you find a log line with errors found in a configuration file. On the **Anomaly** card, a sudden spike on the chart indicates where the configuration file is generating the errors. You can view additional key-value pairs extracted from the logs on the **Meaningful log properties** card and find the affected application services and their related service offerings on the **Impacted services** card. For greater context, you may want to investigate the logs that surround the anomaly, locate the configuration file, and find the time frame where the issues occurred.

What to do next

Use additional features on the Log Viewer to refine your search or define alerts.

- Modify the search query to fine-tune the search and save useful searches. See Define, save, and share a log data search.
- If you discover important relationships in the log data, select **Define alert** to define the kind of alert that should be triggered by the data. See Define a custom Log Analytics alert rule in Health Log Analytics for instructions.

Define, save, and share a log data search

Define, save, and share searches of log data to help determine the causes of Log Analytics alerts.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Define a search.
   a. Select the selection icon ( ) and then select **New search**.
   b. Set the values of the search parameters in the search fields.

Search fields

Search field	Description
Query	Search query.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Search field	Description
	Tip: The Log viewer uses the Elasticsearch search engine, so you can use any supported search term structure in the **Query** field.
Component	Logical component of the service instance that generated the event. Multiple CIs can sometimes perform the same function.
Time range	Time range to apply to the X-axis when displaying the returned data. The setting that you specify appears in the **Start time** and **End time** fields. Use one of the following methods: ▪ Select a time period from the list. ▪ Click **Custom range** to use the date and time picker to specify a range. **Note:** You can modify the settings in the **Start time** and **End time** fields manually. The selected time range shown in **Select range** then changes to Custom range. This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store. **Note:** Saved searches do not include time range settings.

c. Select **Search**.

The system returns the full list of log lines that match the search values. The information is displayed in the Results over time chart.

3. **Optional:** Save the search.
The saved search includes any selected filters. For information about filters, see Filter search results on the Log Viewer in Health Log Analytics.

Note: Saved searches do not include time range settings.

a. Select **Save As**.

b. In the **Search name** field, specify a unique and descriptive name for the search and then click **Save**.

Note: If you are using Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store, you can define an alert rule without saving the search. For more information, see Define a custom Log Analytics alert rule in Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. **Optional:** Share the saved search with an assignment group.
   a. Select **Share**.
   b. Select an assignment group from the list.
   c. Select **Save**.

Related topics

Use or modify a saved log data search in Health Log Analytics

Define a custom Log Analytics alert rule in Health Log Analytics

Use or modify a saved log data search in Health Log Analytics

Use a saved search of log data to better understand the causes of an alert. As the owner of a saved search, you can modify the search values and save your changes.

Before you begin

Note: If you're not the owner of the saved search, save the search with a different name. You can then update search values, save your changes, and share the search with others.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Use a saved search.
   a. Select the selection icon ( ) and then select **Load search**.
   b. In the Load search dialog box, select the name of the search to load.
   The system returns the full list of log lines that match the search values and displays the information in the Results over time chart.
3. **Optional:** Update the saved search.
   a. Select the selection icon ( ) and then choose **Manage my searches** from the drop-down list.
   b. Modify the settings.

Manage my searches form

Field	Description
Name	Name of the saved search.
Query	Search query. (Optional) The Log viewer uses the Elasticsearch search engine, so you can use any supported search term structure in the **Query** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Assignment group	Assignment groups that can access the search. The members of the groups can use the search.
Filter	Column filter in standard format (`field1=value1, field2=value2, field3=value3, ...`).
Updated	Date and time the search was updated. (Optional) This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store.

To revert changes you have made to the search values, select the selection icon ( ) and then select **Discard Changes**. The changes that you made to the search values are discarded. You can continue to update the search settings.

c. Save the updated search.

i. Select **Save as**.

ii. In the **Search name** field, specify a unique and descriptive name for the search and then select **Save**.

Filter search results on the Log Viewer in Health Log Analytics

Apply filters on the **Log Viewer** to show only your desired data.

Before you begin

Role required: evt_mgmt_operator, or evt_mgmt_user, or evt_mgmt_admin

About this task

This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store.

Procedure

1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Define and run a search.
   For more information, see Define, save, and share a log data search.
3. Filter the search results in one of the following ways.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
Add filters using the fields list	**a.** Select the filter icon ( ). The Selected fields list includes the fields that currently display as columns in the **Log Viewer** table. By default, the table includes the following columns: Application service, Component, Host, Level, Message, Raw message, and Time. The Available fields list includes all remaining fields that the sys tem has extracted from the log. **b.** Select a field in the list. **Note:** You can search for a specific field using the search option. The top five values the system has found in the records for this field are displayed, along with the percentage of their occur rence. 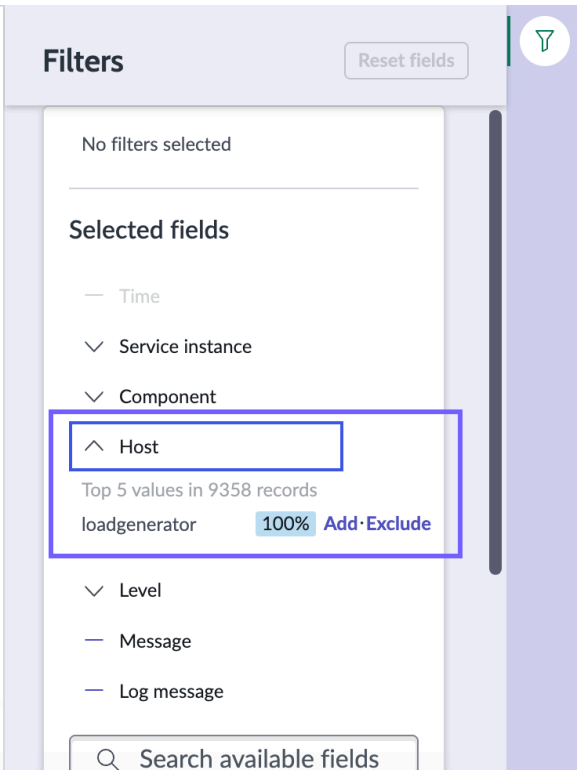 **c.** Define a filter for a value in the field. ▪ To display only data that contains a val ue, select **Add** for it. For example, to set the filter **[Level][is] [critical]**, select **Add** for the value "criti cal" in the Level field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	▪ To create negative filters that exclude da ta that contains a value, select **Exclude** for it. For example, to set the filter **[Level][is not][critical]**, select **Exclude** for the val ue "critical" in the Level field.
Add filters from the Log Viewer table	**a.** In a column header, select the more actions icon (⋮). **b.** In the dialog box, select the operator and specify the filter terms for the field. 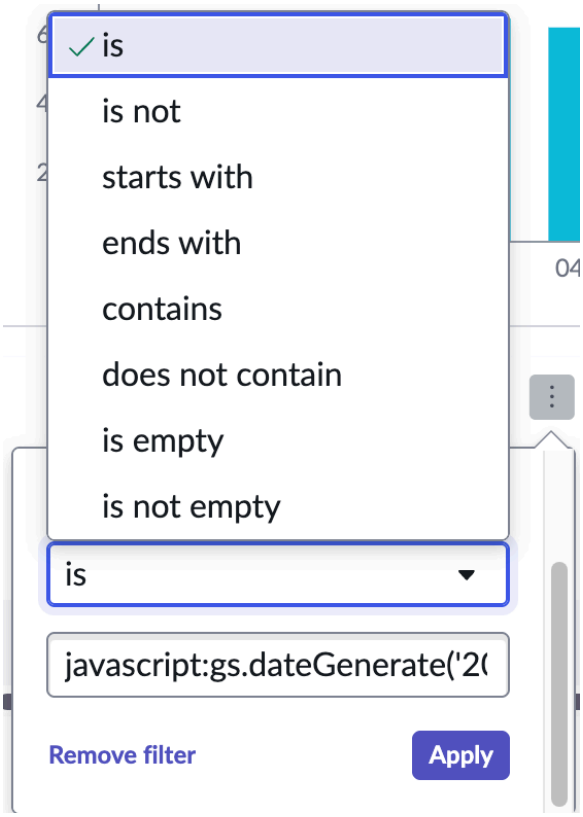 **c.** Select **Apply**. A filter icon in the column header indicates that a filter applies for this field.

The applied filter appears at the top of the **Filters** pane. The total number of applied field value filters in the **Filter** icon adjusts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

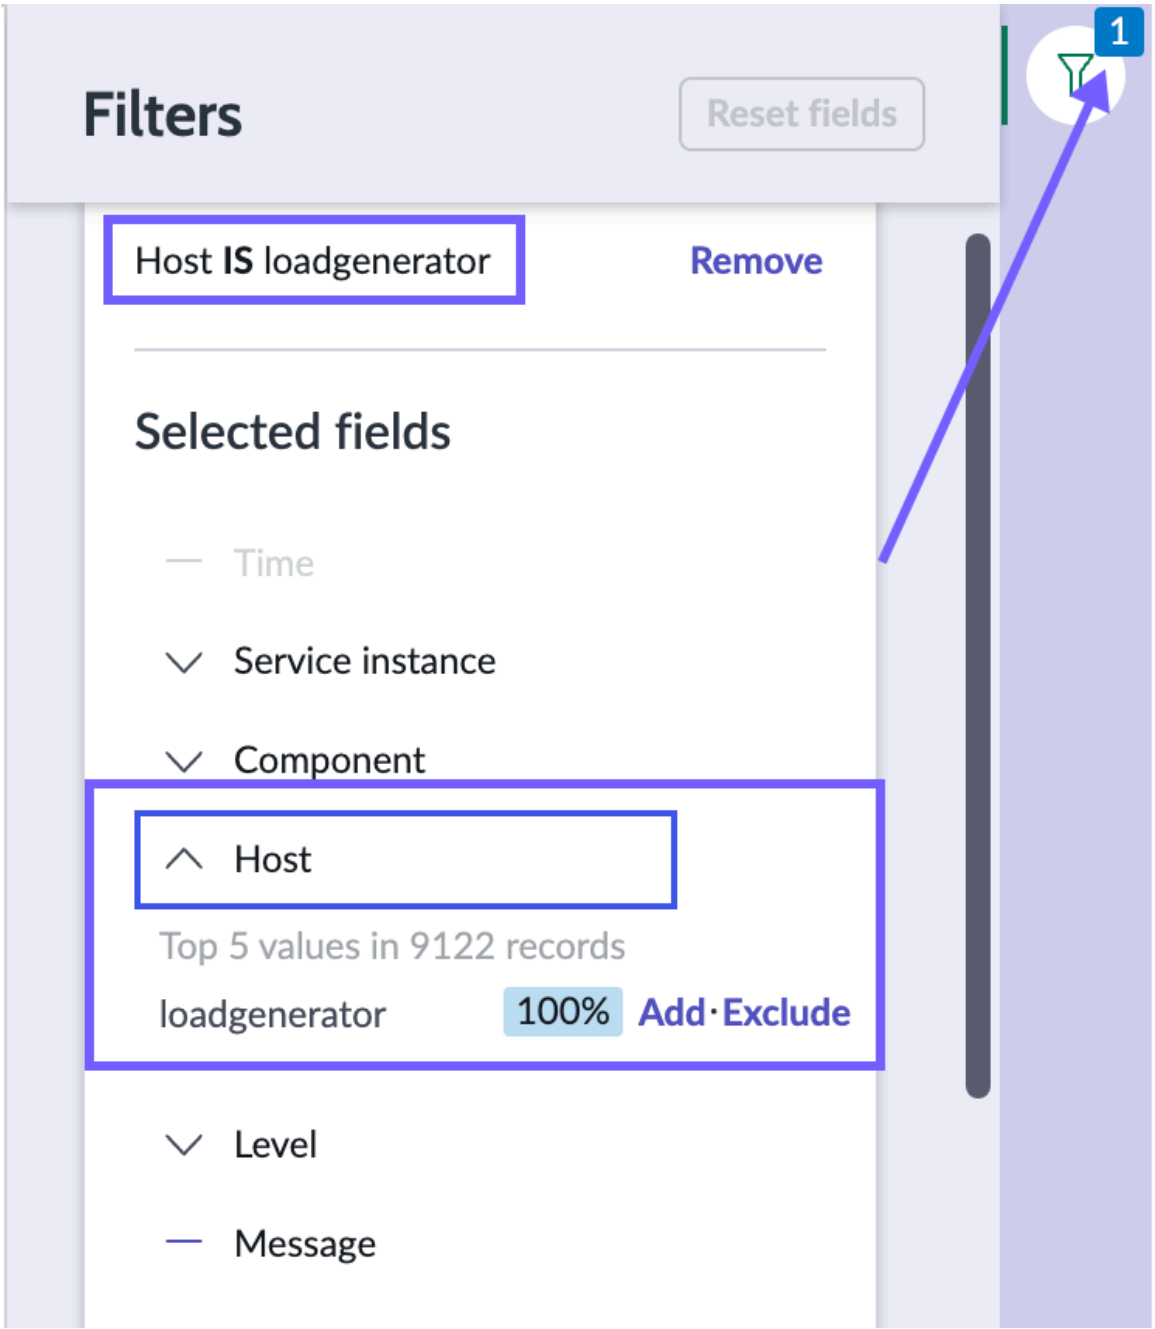

4. **Optional:** Remove a filter in one of the following ways.
   - In the **Filters** pane:
     a. Locate the filter you want to remove.
     b. Select **Remove**.
   - In a column header:
     a. Select the more actions icon (⋮).
     b. Select the filter you want to remove.
     c. Select **Remove filter**.

Customize the Log Viewer table in Health Log Analytics

Add or remove columns in the **Log viewer** table to show only the data you want to view.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: evt_mgmt_operator, or evt_mgmt_admin

About this task
This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store.

Procedure
1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Select the filter icon ( ) to display the **Filters** pane.
   The Selected fields list includes the fields that currently display as columns in the **Log viewer** table. By default, the table includes the following columns: Application service, Component, Host, Level, Message, Raw message, and Time. The Available fields include all remaining fields that the system has extracted from the log.

   In the Service Operations Workspace Log Analytics application, Version 21.2.7 - November 2022 and later, available from the ServiceNow Store, the Log message column is added. The Raw message column no longer displays by default.

   You can add the Raw message column to the table by selecting it from the **Filters** pane.
3. Add or remove columns in the **Log viewer** table.
   - To add a column to the table, select the field of the same name in the Available fields list and then select the add icon ( ).

     The added field displays as a column in the table. It is now listed in the Selected fields list.
   - To remove a column, select the field of the same name in the Selected fields list and then select the remove icon ( ).

     The corresponding column is removed from the table. The removed field is added to the Available fields list.

     If you have filtered the data for specific values in a field, that data still displays even when the field is no longer selected. For more information about filtering data, see Filter search results on the Log Viewer in Health Log Analytics.

     Note: The Time column can't be removed from the table.
4. **Optional:** Adjust the maximum number of characters per field in the table.
   a. Navigate to **Health Log Analytics** > **Health Log Analytics Administration** > **System Properties**.
   b. Open the *sn_occ.log_viewer.max.characters* property record.
   c. In the **Value** field, edit the default value.
   d. Select **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

The **Log viewer** table shows the customized view from the current session onward until you change your preferences.

Add a KB article to a Log Analytics alert

Add your own knowledge base (KB) article to an alert that was generated by Health Log Analytics. For example, you can provide additional information that might help to resolve the underlying issue.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

Use Case: Proactive monitoring of your ServiceNow instance in Health Log Analytics

Use Health Log Analytics to detect and resolve emerging issues in your organization's ServiceNow instance before they negatively impact users.

Using the ServiceNow System Logs Retriever integration to detect platform issues automatically

Starting with version 35.0.26 of Health Log Analytics, you can use the ServiceNow System Logs Retriever integration to automatically detect errors and anomalies in your ServiceNow instance's system logs. As an operator you can view, manage, and respond to the generated alerts in the Express List. Filter the list to show only alerts for the impacted service. Then, use ServiceNow Otto to get an in-depth analysis of the alert and potential resolutions in straightforward, human-readable language. This AI-driven tool demonstrates a deep understanding of the Glide system and offers actionable insights.

Combining HLA anomaly detection with ServiceNow Otto AI capabilities provides a comprehensive toolkit for identifying and resolving platform system issues more effectively. This collaboration enables Health Log Analytics to help prevent system issues from impacting your platform users.

Using the MID Server and System Logs Retriever data inputs together to monitor your instance

You can also stay ahead of disruptions to your ServiceNow environment by using the MID Server data input and the ServiceNow System Logs Retriever data input together.

As an admin, you select a MID Server from the MID Servers list. Under Related Links, choose **Create MID Server Data Input using HLA** to create a data input specific to the selected MID Server. Then activate the System Logs Retriever data input and configure it to collect your ServiceNow instance's system logs.

As an operator, you navigate to the Express List and select an alert from the Alerts list. Review the alert details, including duration and impacted services, and analyze the logs that surround the anomaly. Use ServiceNow Otto to get an in-depth analysis of the alert and potential resolutions in straightforward, human-readable language. If needed, for further investigation, you can navigate to the Review alert-related logs on the Log Viewer for a visual presentation of the system logs. Select System Logs Retriever from the components drop-down list to view all the errors that have occurred in the ServiceNow instance. You can customize the visualization by using filters.

By drilling down into the alert and MID Server logs, you can quickly identify the issue and resolve it before it affects your platform users.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Managing alert generation, prioritization, anomaly detection

Influence how Health Log Analytics generates and prioritizes alerts and identifies anomalies. Use alert filters to reduce noise from unimportant issues.

Overview of impacting alert generation, prioritization, and anomaly detection

Explore how to influence alert generation and anomaly detection in Health Log Analytics. You can create and modify rules to generate alerts for specific metrics, ensuring that you're informed of important issues. You can also assign different levels of significance to alerts, helping Health Log Analytics prioritize them based on their relevance. In addition, you can influence anomaly detection by managing the keywords that Health Log Analytics looks for in your log data, and use advanced log alert filters to reduce noise.

- Using alert rules to control when Health Log Analytics generates alerts

  Manage alert rules to generate Log Analytics alerts for specified metrics.

- Assign higher or lower significance to an alert in Health Log Analytics

  Assign higher or lower significance to an alert to help prioritize the alerts that Health Log Analytics generates.

- Influence anomaly detection in Health Log Analytics with lexical keywords

  Influence how Health Log Analytics finds anomalies by managing keywords it looks for in your log data.

- Reduce noise by adding advanced log alert filters

  Use advanced log alert filters to reduce noise by making Health Log Analytics drop alerts that don't indicate a significant issue.

Using alert rules to control when Health Log Analytics generates alerts

Alert rules enable you to set conditions that determine when Health Log Analytics (HLA) triggers an alert.

HLA detects anomalies automatically by learning from your log data. However, some log types need custom alert rules to generate alerts reliably. For more information, see Alert rules in Health Log Analytics.

You can define, change, or delete custom alert rules:

- Define a custom Log Analytics alert rule in Health Log Analytics

  Define a Log Analytics alert rule for log data that is not generating alerts automatically.

- Change a custom Log Analytics alert rule in Health Log Analytics

  Edit a Log Analytics alert rule that you have defined.

- Delete a custom Log Analytics alert rule in Health Log Analytics

  Delete a Log Analytics alert rule that you have defined if it’s no longer needed.

Alert rules in Health Log Analytics

Health Log Analytics (HLA) detects anomalies automatically by learning from your log data. However, some log types require a custom alert rule to generate alerts reliably.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can use custom alert rules to specify the metric, threshold, and alert properties for generating alerts that HLA might not detect automatically.

Anomaly detection logic by log pattern

HLA classifies incoming logs into three patterns before applying anomaly detection. This classification determines which detection logic is used.

Log patterns

Pattern	Description
Lively	Logs arrive frequently and consistently: at least once in 20 seconds. For example, an application writes hundreds of log entries per hour. The ML Engine has enough volume to build a reliable baseline, so it applies standard anomaly scoring to detect deviations.
Sparse	Logs arrive infrequently or irregularly: less than once in 1 minute (60 seconds). For example, a batch job runs every night and writes a small number of log entries. Standard anomaly scoring would produce unreliable results here, so the ML Engine applies probability distribution analysis instead. Sparse logs might not generate alerts if the volume is too low to establish a baseline.
Stopped	Logs have not arrived for more than the configured period. Default is 5 minutes (300 seconds). To modify the default value, update the HLA system property `detective.resolution.signal_dead`. For a log stream to be alerted as stopped or dead it must first be considered alive by running continuously for a minimum period of time. You can set this time in the HLA system property `detective.alive_period_seconds_for_signal_d` For information about setting and changing HLA system properties, see Configure global Health Log Analytics system properties.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When to create custom alert rules

- For high-frequency logs with a lively log pattern, there is no need for a custom rule. However, you can add a rule to generate alerts under specific conditions.
- For low-frequency logs with a sparse log pattern, the system might not generate alerts automatically. If these logs should still generate alerts, define a custom alert rule.
- For known critical conditions that HLA might not flag automatically, define a custom rule. For example, if a specific log message indicates that a critical service has failed, define a rule that generates an alert every time that message appears.

Using custom alert rules

Scenario	ML detection	Custom rule
High-frequency logs with a lively log pattern	Likely sufficient	Optional
Low-frequency or periodic logs	Unreliable	Suggested
Known critical conditions	Insufficient	Required

Related topics

Define a custom Log Analytics alert rule in Health Log Analytics

Change a custom Log Analytics alert rule in Health Log Analytics

Delete a custom Log Analytics alert rule in Health Log Analytics

Define a custom Log Analytics alert rule in Health Log Analytics

Define a custom Log Analytics alert rule for log data that might not generate alerts automatically. A custom rule enables you to specify the metric, threshold, and alert properties directly.

Before you begin

For background information, see Alert rules in Health Log Analytics.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Define and run a search.
   For more information, see Define, save, and share a log data search.
3. When a search returns log data that should generate an alert, select **Save** or **Save as** to save the search.

   Note: If you're using Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store , you don't have to perform this step.

4. Access the form for creating an alert rule by selecting **Define alert** at the top right of the **Log viewer** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. On the form, provide the name and description for the alert rule.
   The name appears on the Anomaly card for generated alerts.
6. Determine whether the rule generates alerts when the threshold conditions are met in the log data by setting the **State** field.
   To activate the rule, select **Enabled**.
7. If you have installed the Health Log Analytics application, Version 20.0.11 - July 2021 , set the Severity value for the alert generated by this alert rule.
   - **Low**: Attention is required even though the resource is still functional.
   - **Medium**: Either performance has degraded or a partial, non-critical loss of functionality has occurred.
   - **High**: Major functionality is severely impaired or performance has degraded.
   - **Critical**: Immediate action is required. Either the resource is not functional or critical problems are imminent.
8. If you have installed the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, fill in the details for the query.

Query details

Field	Description
Your query	The query to search for in the log data. By default, this is the query that was defined on the Log viewer. You can modify the query by adding search requirements to it. To search for all queries, specify an asterisk (*).
Active filters	The filters defined for the query on the Log viewer. You can delete the filters.
Component	The component to which this alert rule applies. By default, this is the component that was defined on the Log viewer. You can choose a different component by selecting the search icon ( ) in the Component row and then selecting the required component from the list.

9. Fill in the fields related to the threshold that will trigger the alert rule.

Threshold

Field	Description
Operator	Mathematical comparison operator that qualifies how the hit count triggers an alert of this kind. This field is automatically set to Bigger Than.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Hit count	Number of matching log entries that serve as the trigger for an alert of this kind.
Time period	Period of time over which the hit count is measured. The period is measured in the units specified by the **Time unit** value.
Time unit	Units of the time period. This field is automatically set to minutes.
Persistence	Time period over which the specified hit count per unit time must persist to trigger an alert of this kind. The field is automatically set to 1 minute.

10. Select **Save**.

Related topics

Change a custom Log Analytics alert rule in Health Log Analytics

Delete a custom Log Analytics alert rule in Health Log Analytics

Change a custom Log Analytics alert rule in Health Log Analytics

Modify a Log Analytics alert rule that you have defined.

Before you begin

For background information, see Alert rules in Health Log Analytics.

This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Open a defined alert.
   a. Select the selection icon ( ) at the top right of the **Log viewer** tab, and then choose **Manage my defined alerts** from the drop-down list.
   b. From the **Defined alerts** table, select the alert rule that you want to modify.
3. On the form, update the settings.
   For a description of the settings, see Define a custom Log Analytics alert rule in Health Log Analytics.
4. Select **Save** to save your changes.

Delete a custom Log Analytics alert rule in Health Log Analytics

Delete a Log Analytics alert rule that you have defined if it is no longer necessary.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

For background information, see Alert rules in Health Log Analytics.

This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, and the Health Log Analytics Viewer application, Version 20.0.4 - July 2021, available from the ServiceNow Store.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Open a defined alert.
   a. Select the selection icon ( ) at the top right of the **Log viewer** tab, and then choose **Manage my defined alerts** from the drop-down list.
   b. From the **Defined alerts** table, select the alert that you want to delete.
3. Select **Delete**.

Assign higher or lower significance to an alert in Health Log Analytics

Label an alert in Health Log Analytics as meaningful or insignificant, or restore normal importance to the metric involved in generating it.

Assigning higher or lower significance to an alert helps you prioritize the alerts Health Log Analytics generates. Here are some common use cases for adjusting the importance of an alert:

- Labeling alerts as meaningful: You can make a significant alert more likely to be included in a Log Analytics group when its associated metric behaves anomalously. This ensures that important alerts receive appropriate attention.
- Muting distracting alerts: You can mute alerts for insignificant issues to reduce noise and avoid being overwhelmed by alerts for minor problems. This enables you to focus on more critical alerts. You can view a list of muted metrics to confirm you're aware of metrics that are no longer generating alerts.
- Reducing sensitivity to certain anomalies: You can lower the sensitivity of anomaly detection to a specific anomaly. This can help you focus on more significant anomalies.
- Restoring normal significance: If you previously marked an alert as significant, muted it, or if you lowered sensitivity to certain anomalies, you can return the associated metric to its normal significance. This is useful when you no longer want to treat the alert specially.

You can do the following to adjust the importance of an alert or reduce sensitivity to specific anomalies:

- Mark an alert as significant in Health Log Analytics

  Make an alert more likely to be included in a Log Analytics group when the associated metric behaves anomalously by labeling the alert as meaningful.

- Mute an unimportant alert in Health Log Analytics

  Eliminate distracting new alerts for insignificant issues by muting them. View the list of metrics that were muted and no longer generate alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Decrease sensitivity to similar anomalies in Health Log Analytics

  Make Health Log Analytics anomaly detection less sensitive to anomalies like the one that triggered the current Log Analytics alert.

- Restore normal importance to an alert metric in Health Log Analytics

  Return normal significance to the metric involved in generating a Log Analytics alert.

For detailed information about alert metrics, see How Health Log Analytics generates alerts.

Mark an alert as significant in Health Log Analytics

Make an alert more likely to be included in a Log Analytics group when the associated metric behaves anomalously by labeling the alert as meaningful.

Before you begin
Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. In the Service Operations Workspace, open a Log Analytics alert and then select **Apply ML feedback** > **Mark as significant**.
2. Confirm the action in the dialog box.

   Note: Processing this feedback might take a few seconds.

   The alert is marked as significant. The **Restore ML feedback** button appears instead of the **Apply ML feedback** drop-down list.
3. **Optional:** Restore the significant alert.
   If you no longer want the significant alert to be treated specially, you can restore normal significance to the alert metric. For more information, see Restore normal importance to an alert metric in Health Log Analytics.

Mute an unimportant alert in Health Log Analytics

Eliminate distracting new alerts for insignificant issues by muting them.

Before you begin
Role required: evt_mgmt_operator or evt_mgmt_admin

About this task
When you mute an alert that is based on the keyword metric or the pattern severity metric, the system mutes the metric in the dominant pattern. The dominant pattern is the main pattern that contributes to the anomaly. For the specified source only, anomalous behavior for the metric that is associated with the muted alert no longer generates alerts. You can also mute the metric that is associated with an unimportant alert for all components of the Application Service. A component can be multiple CIs that perform the same function.

Procedure

1. In the Service Operations Workspace, open the Log Analytics alert that you don't want to see again.
2. Select **Apply ML feedback**
3. Mute the metric associated with the alert by doing one of the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Select **Mute alert metric** to mute the metric for the specified source only.
- Select **Mute alert metric for all components** to mute the metric for all components of the Application Service.

4. Confirm the action in the dialog box.

   Note: Processing this feedback might take a few seconds.

   The alert is muted. The **Restore ML feedback** button appears instead of the **Apply ML feedback** drop-down list.

5. **Optional:** Restore the muted alert.
   If you no longer want the muted alert to be treated specially, you can restore normal significance to the alert metric. For more information, see Restore normal importance to an alert metric in Health Log Analytics.

Related topics

View the list of muted metrics in Health Log Analytics

View the list of muted metrics in Health Log Analytics

View the list of metrics that were muted so that they no longer generate alerts.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

Navigate to **All** > **Health Log Analytics** > **Log Anomaly Detection** > **Muted Metrics**.

Muted Metrics

Field	Description
Service instance	Service instance where the muted alert with which the metric is associated appears.
Component	Logical component of the service instance that generated the event. Multiple CIs can sometimes perform the same function.
Pattern text	Text for the metric that is associated with the muted alert.
Metric	Log text or key that represents a meter or gauge value in the data. For example, patterns, raw metrics, keyword, severity, all events metrics.
Created	Date and time when the metric was added to the list.

Related topics

Mute an unimportant alert in Health Log Analytics

Restore normal importance to an alert metric in Health Log Analytics

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Decrease sensitivity to similar anomalies in Health Log Analytics

Make anomaly detection less sensitive to anomalies like the one that triggered the current Log Analytics alert.

Before you begin

Note: This functionality is not available for customer alerts and new signal alerts.

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task

Selecting the Raise feedback option raises the threshold for generating alerts when Health Log Analytics identifies future anomalies similar to the one that led to this alert.

Procedure

1. In the Service Operations Workspace, open a Log Analytics alert and then select **Apply ML feedback** > **Raise**.

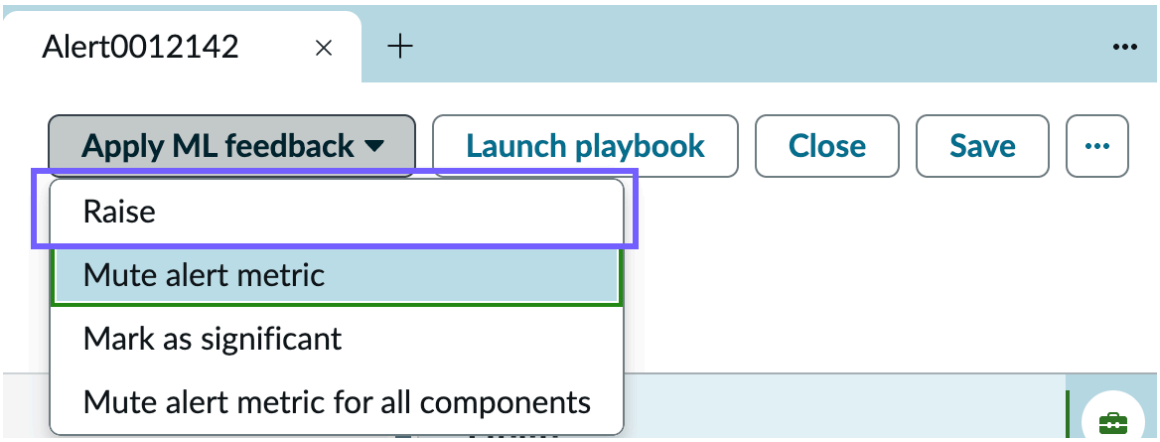

2. Confirm the action in the dialog box.

   Note: Processing this feedback might take a few seconds.

   Anomaly detection is now less sensitive to anomalies like the one that triggered the current alert. The **Restore ML feedback** button appears instead of the **Apply ML feedback** drop-down list.

3. **Optional:** Restore normal sensitivity to anomalies like the one that triggered this alert.
   a. While viewing the alert, select **Restore ML feedback**.
   b. Confirm the action in the dialog box.

Suppress alerts from short-lived anomalies

Apply sustained alert feedback to generate alerts only when anomalies persist for a set duration in Health Log Analytics.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task

Short-lived anomaly spikes that resolve quickly often don't require alerts. If a metric jumps briefly and resolves on its own, the spike is noise and not actionable. The sustained alert feature helps reduce this noise. When you apply sustained alert feedback, the anomaly must persist continuously for 15 minutes before Health Log Analytics generates an alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, the system calculates this duration automatically based on the metric's signal amplitude, historical anomaly patterns, and previous sustained alert applications (adaptive mode). You can override it and set a fixed duration yourself, instead (manual mode).

Procedure

1. Open a Log Analytics alert.

   a. In the Service Operations Workspace, select the list icon ( ).

   b. Select the appropriate list in the Alerts sub-list and open the alert you want to configure for sustained alert feedback.
   In the **All Alerts** list, alerts that were generated by Health Log Analytics have the value **Log Analytics** in the **Source** column.

   c. Select the alert number.

2. Select **Apply ML feedback** and then choose **Sustain Alert**.
   A message confirms that the anomaly must persist for 15 minutes before an alert is generated.

   Health Log Analytics applies adaptive mode by default and calculates a recommended duration automatically. A confirmation message states the calculated duration and that it's based on signal amplitude and patterns, for example, "Sustained Alert applied: 12 min (Smart)".

3. **Optional:** To set a fixed duration instead of using the calculated adaptive duration, switch to manual mode.

   a. Select the **[OXY-PLACEHOLDER]** advanced option next to **Sustain Alert**.

   b. Select a preset duration (5, 10, 15, 20, or 30 minutes), or specify a custom duration between 1 and 120 minutes.

   c. Confirm the duration.

   (Optional) Health Log Analytics applies the duration you set instead of the calculated adaptive duration.

4. Confirm sustained alert feedback application.Review the applied sustained alert feedback.

   Sustained alert feedback is applied to the metric. Anomalies that persist for less than the continuous duration threshold are suppressed. The sustained alert configuration applies across user sessions and system restarts.

   The value **Sustained** is displayed underneath the **Applied feedback** label in the header strip at the top of the alert record.A value showing the applied duration and mode—for example, **Sustained: 12 min (Smart)** for adaptive mode or **Sustained: 20 min (Custom)** for manual mode—is displayed underneath the **Applied feedback** label in the header strip at the top of the alert record.

5. **Optional:** Remove the sustained duration threshold for anomalies like the one that triggered this alert.

   a. While viewing the alert, select **Restore ML feedback**.

   b. Confirm the action in the dialog box.

   (Optional) The system removes the sustained alert configuration for the metric associated with this alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Restore normal importance to an alert metric in Health Log Analytics

Return normal significance to the metric involved in generating a Log Analytics alert. Use this option if you no longer want an alert to be treated specially that was muted, marked as significant, or made less sensitive to similar anomalies.

Before you begin
Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. While viewing a Log Analytics alert in the Service Operations Workspace, select **Restore ML feedback**.
2. Confirm the action in the dialog box.

Result
The metric that is associated with the alert returns to normal importance.

Influence anomaly detection in Health Log Analytics with lexical keywords

Influence how Health Log Analytics finds anomalies by managing keywords it looks for in your log data. When text in log data for a source matches a lexical keyword that exceeds a specified count threshold, the system identifies an anomaly and generates an alert.

Health Log Analytics scans your logs for words that can indicate important issues. Lexical keywords such as 'crashed' or 'failed' signal conditions that may require attention. The application sets a threshold for each lexical keyword based on its normal occurrence pattern and frequency in your logs.

When Health Log Analytics scans your logs, it identifies all occurrences of the keyword. If the frequency of a lexical keyword in the log data for a source exceeds the specified threshold, the system identifies an anomaly and generates an alert.

Health Log Analytics comes with many default global keywords. You can add, edit, and delete these global keywords or phrases, which apply to all source types.

Note: To add a specified keyword associated with a specific source type, see Configure source type capabilities in Health Log Analytics.

You can manage lexical keywords as follows:

- Add, edit, or delete lexical keywords in Health Log Analytics

  Add, edit, and delete global keywords or phrases.

- View the lexical keywords that generate alerts in Health Log Analytics

  View the list of lexical keywords that can indicate important issues in log entries. By default, the table lists only global keywords.

Add, edit, or delete lexical keywords in Health Log Analytics

Manage the keywords that Health Log Analytics looks for in your log data.

Before you begin
Role required: evt_mgmt_operator or evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

In log data, terms like "crash" or "failed" are called lexical keywords because they indicate issues that can merit attention. When text in log data for a source matches a lexical keyword that exceeds a specified count threshold, the system identifies an anomaly and generates an alert.

Important: A lexical keyword differs from a key in a `key : value` pair in a log line. For example, *`Hostname`* is a key that takes on a value: the name or IP address of the host. In contrast, a keyword like *`Failed`* is important by itself and does not take on a value.

The application comes with many default global keywords. You can add, edit, and delete global keywords or phrases. These keywords apply to all source types.

Note: To add a specified keyword that is associated with a specific source type, see Configure source type capabilities in Health Log Analytics.

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Log Anomaly Detection** > **Lexical Keywords**.
   By default, the Lexical Keywords table lists only global keywords.
2. **Optional:** Add a global keyword.
   a. Select **New**.
   b. On the form, fill in the fields.

Lexical keyword form

Field	Description
Name	Unique and descriptive name for the keyword.
Regular expression	Regular expression (regex) that defines matches.
Exact match	Option to make Health Log Analytics match the exact regex. For example, 'NullPointerException' in a message would not match the regex 'exception'. This field is automatically set to **True**.
Case-sensitive	Option to make Health Log Analytics look for a case-sensitive match of the regex. This field is automatically set to **False**.
Range of analysis	Range of sources types where Health Log Analytics looks for the keyword in the log data. Choices are as follows: ▪ **All source types** ▪ **Specified source type**
Excluded source types	Source types that are not associated with the keyword. Health Log Analytics does not

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	look for the keyword in the log data of these source types.

c. Select **Submit**.

3. **Optional:** Edit a global keyword.
   a. Select the keyword that you want to edit in the list.
   b. On the form, edit the relevant fields.
   c. Select **Update**.
4. **Optional:** Delete one or more global keywords.
   a. Select the rows of the keywords that you want to delete.
   b. From the Actions on selected rows list at the bottom of the page, select **Delete**.
   c. Select **OK**.

Related topics

View the lexical keywords that generate alerts in Health Log Analytics

View the lexical keywords that generate alerts in Health Log Analytics

View the list of lexical keywords that can indicate important issues in log entries.

Before you begin
Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

Navigate to **All** > **Health Log Analytics** > **Log Anomaly Detection** > **Lexical Keywords**.
The Lexical Keywords table appears. By default, the table lists only global keywords.

For information about adding, editing, or deleting global keywords, see Add, edit, or delete lexical keywords in Health Log Analytics.

To add a specified keyword that is associated with a specific source type, see Configure source type capabilities in Health Log Analytics.

Reducing noise by adding advanced log alert filters in Health Log Analytics

Use advanced log alert filters to determine whether to allow an alert or to drop it. These filters reduce noise by dropping alerts that don't indicate a significant issue.

You can add advanced log alert filters to scan alerts for your defined conditions. For example, you can define a filter that drops alerts coming from specific log sources, or alerts for anomalies that do not cross the specified threshold.

Some examples of the actions that advanced filters can enable:

- Alert only on anomalies that are shared across multiple hosts.
- Do not alert on anomalies that happen outside of working hours.
- Do not alert if the anomaly does not cross the specified threshold.
- Alert only on anomalies that are part of a correlation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can manage advanced log alert filters as follows:

- Create advanced log alert filters

  Add advanced log alert filters to scan alerts for conditions that you specify. The filters reduce noise by dropping alerts that do not indicate a significant issue. While developing a filter, you can test, update, publish, or activate the filter at any time.

- Continue modifying the filter by reopening the filter record from the filters list. You can then edit, test, publish, and activate the filter.

For more technical information on log alert filters, see the Advanced Log Alert Filtering [KB0863538] article in the Now Support Knowledge Base.

Log Analytics in Service Operations Workspace for ITOM

View and address the alerts that Health Log Analytics generates in the Service Operations Workspace.

Create advanced log alert filters

Add advanced log alert filters to scan alerts for conditions that you specify. The filters reduce noise by dropping alerts that do not indicate a significant issue. While developing a filter, you can test, update, publish, or activate the filter at any time.

Before you begin

The Disable Alert Rule Engine feature and the Disable Detection Rule Engine feature must be in the OFF state. You set the values by navigating to **Health Log Analytics** > **Health Predictive AI Ops Administration** > **Features**.

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Navigate to **Health Log Analytics** > **Log Anomaly Detection** > **Advanced Log Alert Filter**.
2. Select **New**.
3. Enter a unique and descriptive name for the filter.
4. **Optional:** Enter a description of the operation of the filter.
5. In the **Script template** field, select the script that most closely matches your intended logic. The template can act as a starting point for your custom script code.
   After you select a template, the Custom JS function text box is populated with the appropriate JavaScript function. The JavaScript function applies the filter to the alert payload data and either allows or drops the alert. The alert payload is the text and metadata for the kind of alert that the filter will analyze.
6. Save the record by selecting **Submit**.
   To continue modifying the filter, you must reopen the record from the filters list. You can then edit, test, publish, and activate the filter.
7. Edit the default **Alert payload** text in preparation for testing your intended logic.

   To remove your changes and revert to the default text for the Alert payload text box, click **Reset**.

   For your convenience, Health Log Analytics provides sample alerts with preconfigured payloads. Select a sample alert in the **Example alert** field to display its payload in the Alert payload text box.
8. **Optional:** Save the current values of the filter without testing by selecting **Update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. When your content for the alert payload and your JavaScript function are complete, select **Test**.
   To simulate the alert operation, the system saves the filter values, applies the filter to the **Alert payload** text, and then displays one of the following results:
   - `Alert will be dropped.`
   - `Alert will be allowed.`

Testing the JavaScript to determine whether to allow or drop the alert

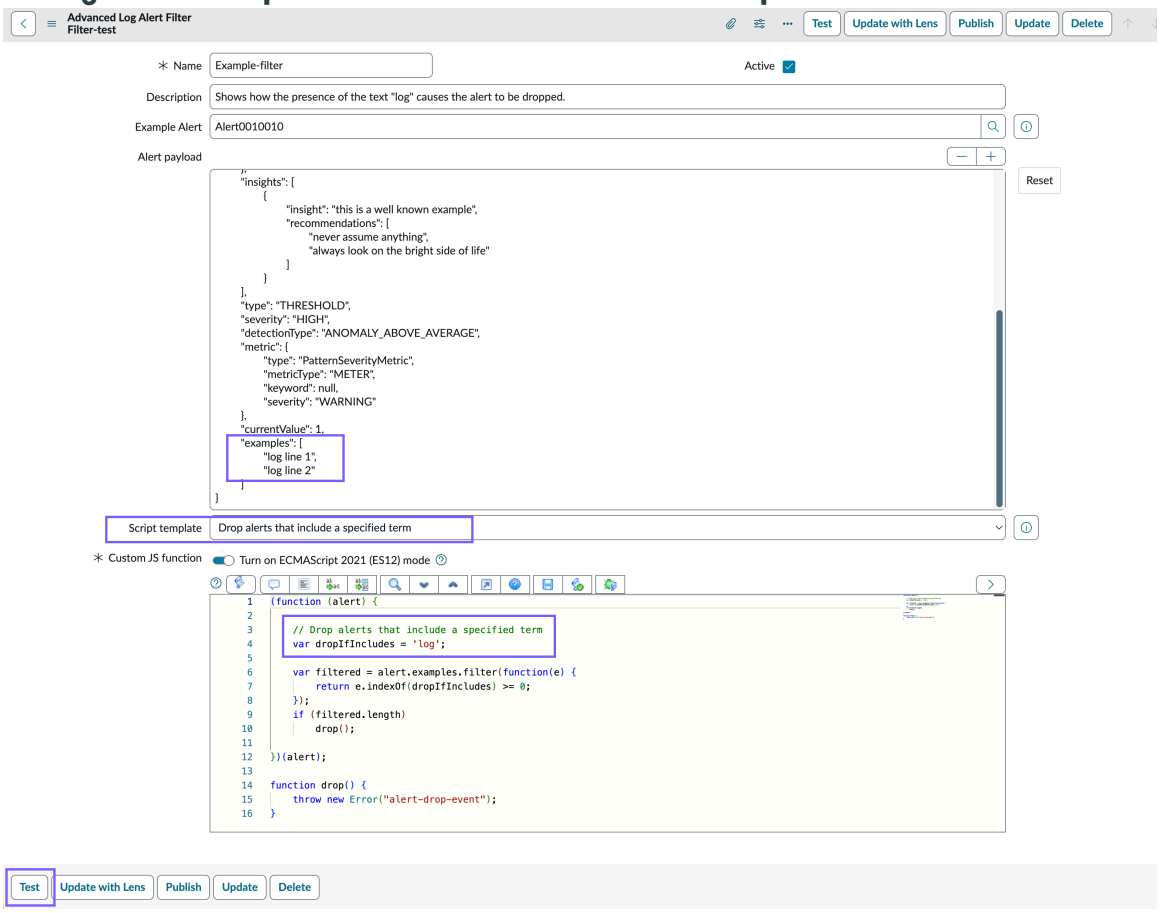

Note: If your new JavaScript function is not behaving as expected, you can revert to the last published one by selecting the **Revert JS Function** related link.

10. Repeat the process of updating the alert payload and testing the JavaScript function as often as needed.
11. When you're satisfied with the filter, save its values and determine whether to apply it to the log stream.
    - To save the values and apply the filter to the log stream, make sure the **Active** check box is selected and select **Publish**.
    - To save the filter without applying it to the log stream, clear the **Active** check box and select **Publish.**

Note: If you modify a published filter, you must publish the modified filter to apply it to the log stream.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The values of the active filter are saved. If you selected the Active option, the filter is applied to the log stream.

Dashboards for real-time visualization of log data in Health Log Analytics

Health Log Analytics enables you to create log data dashboards and visualizations in real time for enhanced data monitoring and faster identification of deviations.

Log data dashboards and visualizations can give you a high-level visual understanding of important issues as they occur in the system. The Health Log Analytics data source of your choice is used to ingest the data that is provided to the visualization. The log data is taken directly from your stored logs, enabling you to create dashboards for any aspect of the data that interests you. For example, by visualizing the components that are generating logs, you can get an indication of the number of components that are fully integrated.

The data can be presented in various types of graphs and charts. You can adjust the time range and the other settings of the visualization as needed.

This functionality is supported in the Service Operations Workspace Log Analytics application, Version 21.2.7 - November 2022 and later, and has a dependency on the Platform Analytics Workspace, available from the ServiceNow Store.

As an alternative to defining your own dashboard, you can use the AIOps dashboards, which offer comprehensive visualizations of critical operational data. For more information, see AIOps Dashboards in Service Operations Workspace for ITOM.

Example of a log data dashboard

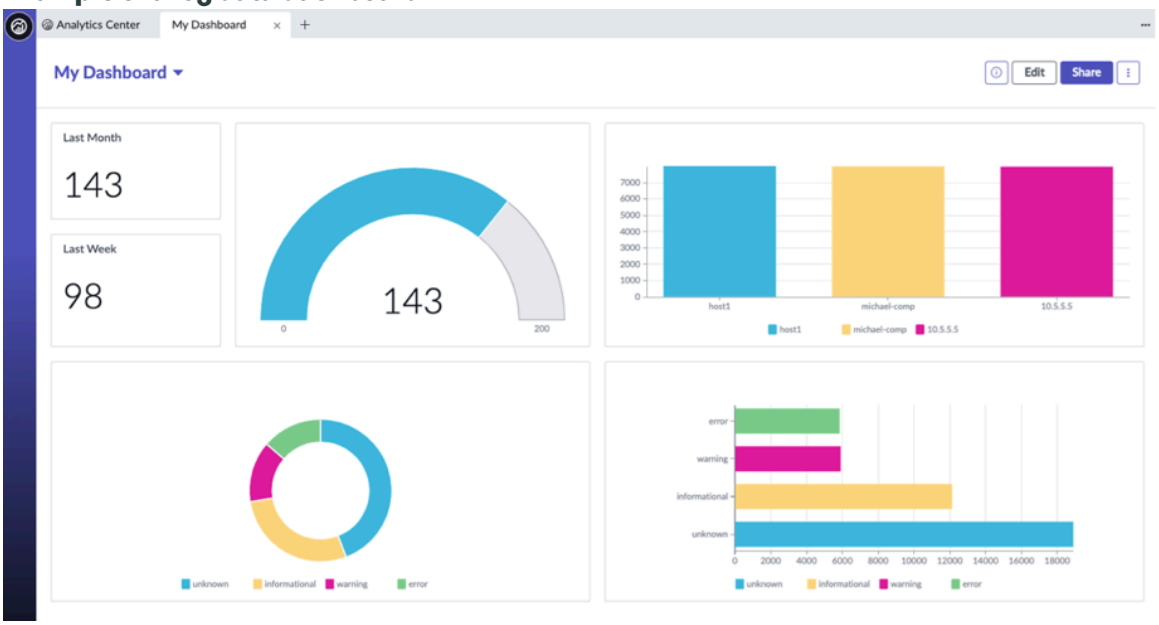

Create log data dashboards and visualizations in Health Log Analytics

Build Health Log Analytics log data dashboards and visualizations in real time for easier identification and understanding of issues.

Before you begin
This functionality is supported in the Service Operations Workspace Log Analytics application, Version 21.2.7 - November 2022 and later, and has a dependency on the Platform Analytics Workspace, available from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For detailed information about creating and working with dashboards, see Dashboards in the Analytics Center.

Note: As an alternative to defining your own dashboard, you can use the AIOps dashboards, which offer comprehensive visualizations of critical operational data. For more information, see AIOps Dashboards in Service Operations Workspace for ITOM.

Role required: evt_mgmt_operator or viz_admin

Procedure

1. Navigate to **Workspaces** > **Platform Analytics Workspace**.
2. On the Analytics Overview page, select the **Dashboards** tab.
   This tab enables you to view all the dashboards you have rights to access on your instance. You can filter the dashboard tiles to limit the number that are visible.

   Note: If you want to use an existing dashboard, select the dashboard here to open it. Otherwise, create a dashboard as explained in the remainder of this procedure.

3. Create a dashboard.
   a. Select **Create new dashboard**.
   b. In the New dashboard dialog, choose **In-line editor** and specify a name and optional description for the dashboard.

      Note: The in-line editor does not automatically save your work. Ensure that you save regularly.

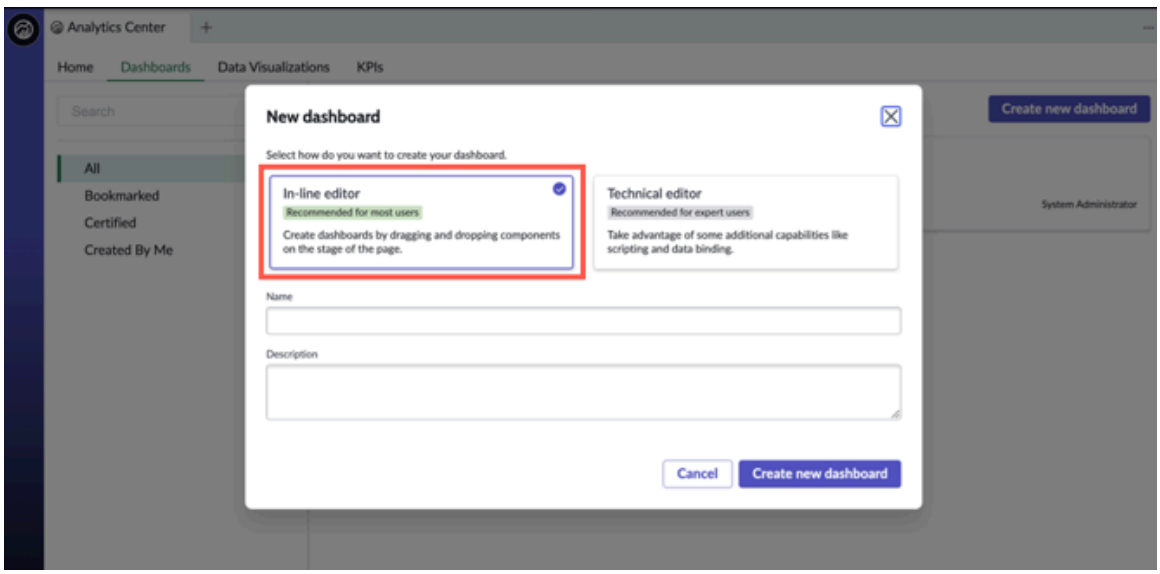

   c. Select **Create new dashboard**.
4. Add an element to the dashboard.
   Various elements are available to help you organize and populate the information on the dashboard. For more information, see Analytics Center dashboard elements.

   a. Select **Add new element**.

   b. Select **Data visualization** from the available elements.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Data visualization** element appears on the dashboard, where you can edit it.

Note: You can click and drag the corners of the element to resize it.

5. Edit the element.

   a. On the element, select the Settings icon:

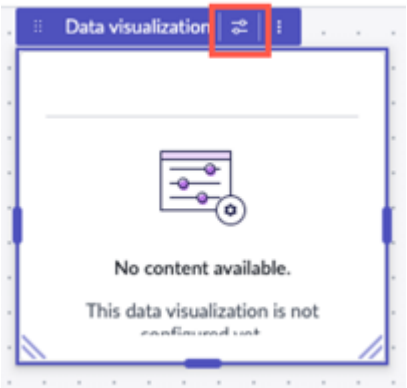

   The Configuration panel opens.

   b. From the **Visualization type** list, choose the type of visual representation you want for the log data.
   You can choose from various types of graphs and charts.

   For detailed information, see Data visualizations in the Analytics Center.

   c. In the Data section, select **+ Add data source**.

   The Health Log Analytics data source that you select is used to ingest the log data that is provided to the visualization.

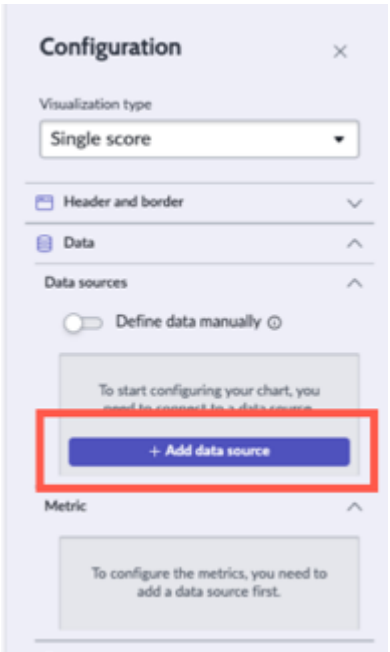

   d. Select the Health Log Analytics Log components data source.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. Choose the required data source from these options:

- Saved searches – a list of all the searches you saved in the Log viewer.
- Select component – a list of all the components that stream data to Health Log Analytics.

f. Use Conditions, a Lucene query, or both to specify the data segment you're interested in and then select **Run**.

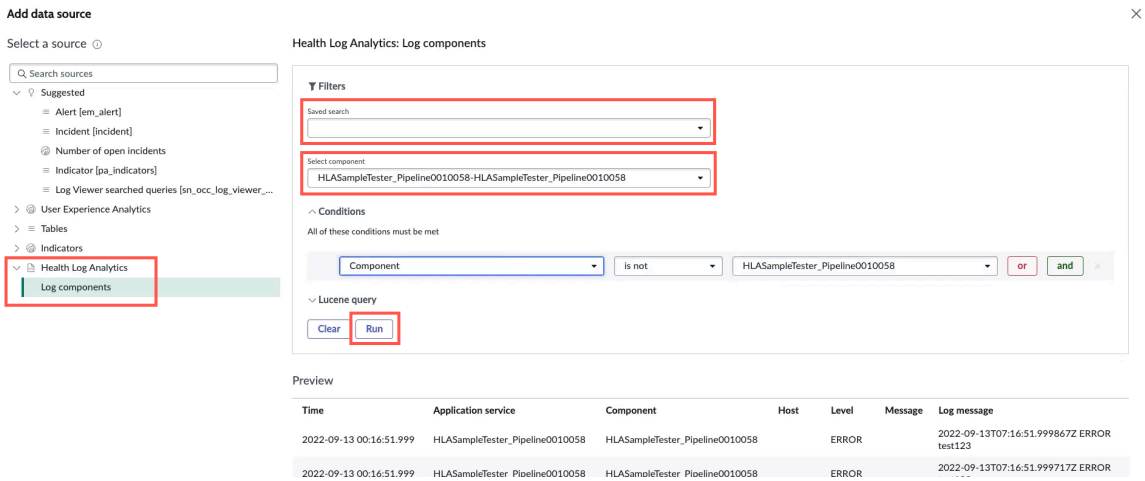

g. Select **Add this source** to add the selected Health Log Analytics data source. The specified visualization displays in the **Data visualization** element.

h. **Optional:** Set fixed start and end dates to customize the time range of the presented data.

(Optional)

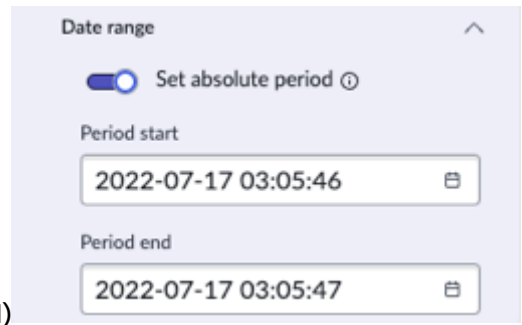

i. In the Date range section, select **Set absolute period**.

ii. Specify the start and end date for the absolute period in the appropriate calendars.

i. **Optional:** Adjust the other settings of the visualization as needed.

(Optional) For detailed information on settings for all the visualization types, select the link for the relevant type in Data visualizations in the Analytics Center.

j. **Optional:** Add more elements as needed.

k. Select **Save** to save your changes and continue editing.

l. Select **Exit editing mode** to stop editing and display the dashboard.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Share a Platform Analytics dashboard

Monitor visualized log data on a predefined dashboard in Health Log Analytics

Use visualizations to monitor log data on the built-in Operational Dashboard in Service Operations Workspace Log Analytics and address issues as they occur in the system.

Before you begin

Role required: evt_mgmt_operator or viz_admin

Important: From the Vancouver release onward, this feature is deprecated and replaced with the AIOps HLA Operational dashboard. For more information about AIOps dashboards, see AIOps Dashboards in Service Operations Workspace for ITOM. You can also build a custom dashboard per your needs. For more information, see Create log data dashboards and visualizations in Health Log Analytics.

Procedure

1. Open the **Log Viewer** using one of the following methods:
   - Navigate to **Workspaces** > **Service Operations Workspace** and select the Log Viewer icon ( ).
   - While viewing log entries for an alert on the **Surrounding logs** tab, select **Log Viewer**.
2. Select the selection icon at the top right of the **Log viewer** tab, and then choose **Manage my dashboards** from the drop-down list.
   The **Operational Dashboard** tab displays.

   By default, the dashboard visualizes log data that was ingested in the last 24 hours.

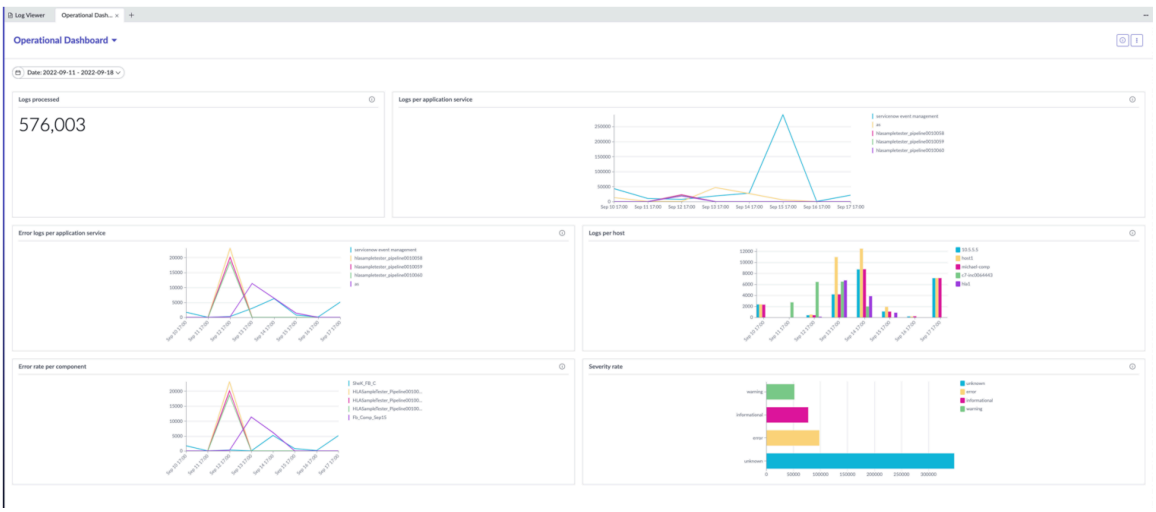

3. **Optional:** Change the time range for which data is shown by selecting the required range from the **Date** drop-down.
   You can specify a custom range by selecting **Custom** and then selecting the required start and end dates.
4. Monitor the visualized log data and use it to identify and address issues.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Operational Dashboard

Name	Description	How to use the data
Logs processed	The number of logs that were processed and analyzed in the selected period.	Ensure that the log data is streaming as expected.
Logs per Application Service	The number of logs streamed for each application service in the selected period.	Ensure that log data is streaming as expected for each application service.
Error rate per Application Service	The number of logs with errors for each application service in the selected period. **Note:** Health Log Analytics automatically creates alerts for Error anomalies.	Identify errors in the logs.
Logs per Host	The number of processed logs for each host in the selected period.	Identify hosts that are logging more data, or less data, than expected.
Error rate per Component	The number of logs with errors for each component in the selected period.	Ensure that the components have a normal Error rate.
Severity rate	The severity level rate of the logs for the streaming sources.	Ensure that Critical and Error level rates are as expected.

Viewing system features and system health alerts in Health Log Analytics

View all Health Log Analytics system features that your admin can activate or deactivate for you. In addition, you can view a model of the Health Log Analytics core components and the system health alerts that affect them.

Overview of Health Log Analytics system features and system health alerts

Health Log Analytics enables you to view all of its system features and their description. The admin can enable or disable a feature for you. You can also view a model of the core components of Health Log Analytics and the system health alerts that could impact them. Health Log Analytics and the Event Management application together form a single service. Examining the configuration items (CIs) that comprise Health Log Analytics and their relationships helps you understand the impact of HLA system health issues on the service.

- View the settings of Health Log Analytics system features

  View the settings for all Health Log Analytics system features. These features configure many operations that affect your work. The admin can enable or disable features for you.

- Identify Health Log Analytics system health issues on the service map

  The ServiceNow Event Management service instance map includes a model of the Health Log Analytics core components and the system health alerts that affect them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View the settings of Health Log Analytics system features

View the settings for all Health Log Analytics system features. These features configure many operations that affect your work. The admin can enable or disable features for you.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

Procedure

1. Navigate to **All** > **Health Log Analytics** > **Health Log Analytics Administration** > **Features**.

   A full definition of each feature appears in its **Description** field.

   By default, the table shows the system features that administrators modify most commonly. To view the complete list of features, use the filtering options.

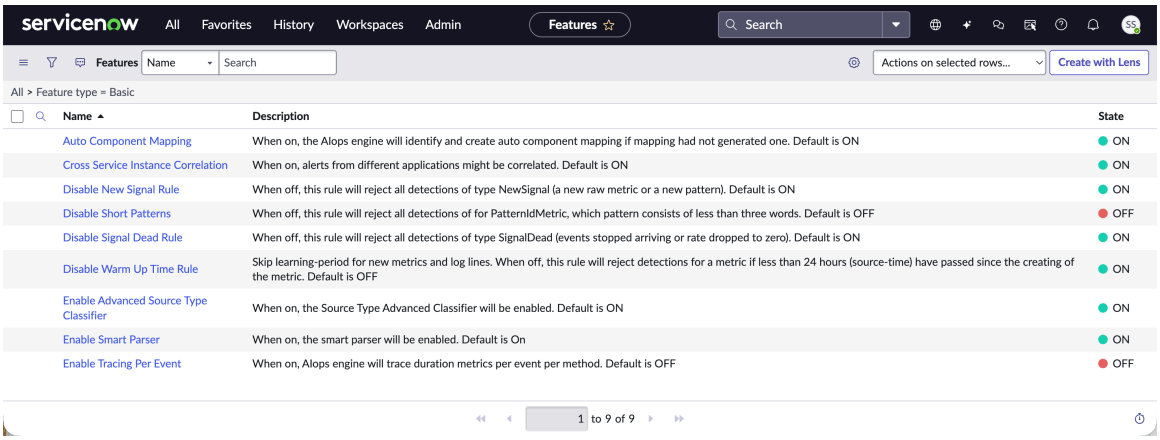

2. **Optional:** Ask your admin to activate or deactivate a feature by updating the **State** setting.

Viewing Health Log Analytics system health alerts on the service map

The ServiceNow Event Management service map includes a model of the Health Log Analytics core components and the system health alerts that affect them.

Health Log Analytics and the Event Management application form a single service. Viewing the configuration items (CIs) that comprise Health Log Analytics and their relationships helps you visualize the impact of HLA system health issues on the service.

HLA components on the Event Management service instance map

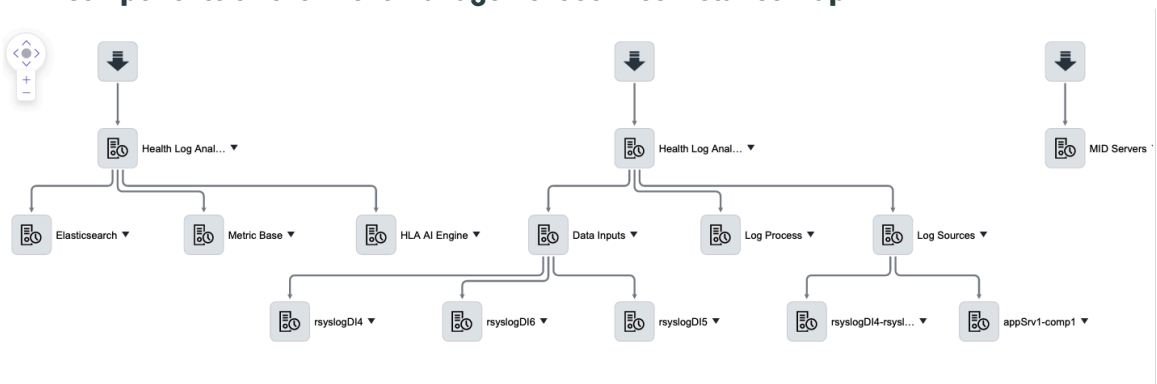

Health Log Analytics uses self-health checks to monitor the health of its components. When it discovers a system health issue, a notification triggers an alert. System health alerts display on the affected CIs on the service map. A colored bar on an affected CI indicates the criticality

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

level of the alert: A red bar signifies a critical alert, an orange bar a major alert, and a yellow bar a minor alert.

For example, the red bar indicates a critical alert on this CI:

If the alert impacts the parent object, the parent object has a bar of the same color. For more information, see View an alert impact on CIs in a service map.

You can drill down into an alert to view details about the issue and then address it by performing the action proposed in the notification. For more information, see Health Log Analytics self-health checks, notifications, and proposed actions.

Monitored components

Component	What is monitored
Health Log Analytics core	• Elasticsearch • Metric Base • HLA AI Engine
Health Log Analytics	• Data inputs • Log sources **Note:** New data inputs and log sources are automatically added to the service map.
Log ingestion process	All elements that affect the log ingestion process.

Related topics

Identify Health Log Analytics system health issues on the service map

Identify Health Log Analytics system health issues on the service map

View Health Log Analytics (HLA) system health issues on the ServiceNow Event Management service map.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task

Viewing the configuration items (CIs) that comprise Health Log Analytics and their relationships helps you visualize the impact of HLA system health issues on the service. For more information, see Viewing Health Log Analytics system health alerts on the service map.

Procedure

1. Navigate to **All** > **Event Management** > **Services** > **Application Services**.
2. On the **Mapped Application Services** page, locate ServiceNow Event Management and select **View Service**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

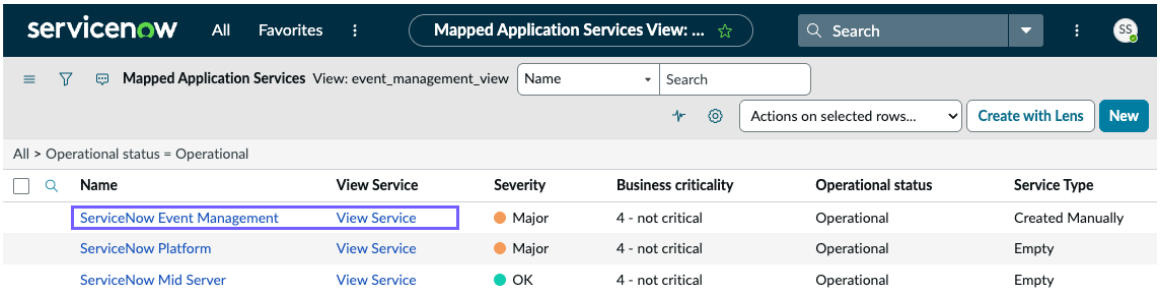

The Event Management service map shows the Health Log Analytics core components and the system health alerts that affect them.

3. On the service map, identify Health Log Analytics system health issues.
4. Drill down into an alert to view details about the issue.

What to do next

Address the alert by performing the action proposed in the notification. For more information, see Health Log Analytics self-health checks, notifications, and proposed actions.

Related topics

Health Log Analytics system health notifications

Analytics and Reporting in Health Log Analytics

View the number of IT issues predicted before users were affected and the money saved by using Analytics and Reporting in Health Log Analytics.

Platform Analytics Solutions for Health Log Analytics

Platform Analytics Solutions contain prepackaged Performance Analytics and Reporting content for use with other ServiceNow AI Platform products. This Platform Analytics Solution provides valuable insights into the number of IT issues Health Log Analytics predicted before your users were affected. It also shows the money saved by preventing critical outages.

Required roles

The following roles and tasks are associated with this solution:

- ServiceNow AI Platform administrator (admin): Install and activate this solution and make any necessary changes to system properties.
- Performance Analytics administrator (pa_admin): Review the indicators, breakdowns, widgets, and dashboards. Set up and start data collection. Share the dashboards with appropriate stakeholders.

Activate the Platform Analytics Solution for Health Log Analytics

Activate the content pack ServiceNow Store application for the Platform Analytics Solution.

Before you begin

Activate the full, subscription version of Performance Analytics on the instance where you are activating this Platform Analytics Solution. You can activate Performance Analytics on a non-production instance without purchasing a subscription. For more information, see Activate your Performance Analytics subscription.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Applications** > **All Applications** > **ServiceNow Store**.
2. Browse for Health Log Analytics and verify that you have the necessary valid entitlements.

   You need a HI account and permission to request applications for the instances under consideration. The HI account and the instance need to be part of the same company. For more information, see Install a ServiceNow Product on the ServiceNow Store.
3. Search for `Health Log Analytics (sn.occ)`.
4. Install the application.

   When you install this application, you also activate the Performance Analytics - Premium (com.snc.pa.premium) plugin.

   Note: This plugin requires an additional fee.

Configure the Platform Analytics Solution for Health Log Analytics

Run diagnostics, review and customize components, and start collecting data.

Before you begin

A ServiceNow AI Platform administrator must have installed the content pack plugins or ServiceNow Store application for this Platform Analytics Solution.

Role required: pa_admin

Procedure

1. Run diagnostics on all records.

   These diagnostics can catch many mismatches between the configuration of your Platform Analytics Solutions and your tables.
2. Navigate to **All** > **Performance Analytics** > **Administration Console**.
3. Verify the name of the dashboard that has been installed with this Platform Analytics Solution: Health Log Analytics Overview dashboard.
4. Open the **Indicator Sources** tab.
5. For each dashboard included in this Platform Analytics Solution:

   a. Filter the indicator sources on the dashboard name.

   b. Review the facts table, conditions, and frequency of the indicator sources compared to the data structure on your own instance.

   Important: Pay particular attention to time stamp fields such as Resolved or Created that are referenced in the **Conditions** field. You may need to use different time stamp fields.

   c. If necessary, open an indicator source and make corrections.

   d. If you are editing an indicator source record, go to the Indicators related list and review the Conditions field for each indicator.
   Changing the indicator source can also affect the additional conditions on the individual indicators.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. If you changed the time field stamps in any indicator sources, then change any related Performance Analytics scripts.
7. Open the **Breakdown Sources** tab.
8. For each dashboard included in this Platform Analytics Solution:
   a. Filter the breakdown sources on the dashboard name.
   b. Review the conditions on the breakdown source as described in Review the breakdown sources.
9. Navigate to **All** > **Performance Analytics** > **Jobs**.
10. Set up and run the historical job for this Platform Analytics Solution, *[PA HLA] Historic Data Collection*.
11. Edit and activate the scheduled data collection job for this Platform Analytics Solution, *[PA HLA] Historic Data Collection*.

Health Log Analytics Operational dashboard

The Health Log Analytics Operational dashboard uses Performance Analytics to monitor log data, alerts, and error rates in Service Operations Workspace.

For more information, see HLA Operational dashboard.

Note: The Health Log Analytics Operational dashboard replaces the Health Log Analytics Overview dashboard, which has been deprecated.

Health Log Analytics reference

Reference topics provide additional information about the Health Log Analytics application.

Components installed with Health Log Analytics

Activating the Health Log Analytics, Health Log Analytics Viewer and Health Log Analytics Core plugins adds several components.

Note: The Application Files table lists the components that are installed with this application. For instructions on how to access this table, see Find components installed with an application.

Plugins

Plugin	Description
Health Log Analytics	Health Log Analytics is an AI-powered log analysis solution capable of preventing IT incidents before users are affected. It automatically identifies your environment and logs, predicts IT incidents, generates alerts, and provides expert suggestions for fixing issues. Health Log Analytics is a ServiceNow Store app. Once installed, the module is available, but you can choose not to interact with it. Dependency:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugin	Description
	• Health Log Analytics Core • The ServiceNow® Event Management application
Health Log Analytics Viewer	Event Management operators use the Health Log Analytics Viewer plugin to browse logs by timestamp or time range. Health Log Analytics Viewer enables Event Management operators to search for a specific log text and visualize the frequency of hits in a particular time period. Dependency: The ServiceNow® Event Management application.
Health Log Analytics Core	This plugin contains essential components for Health Log Analytics, such as the AI Engine and Elasticsearch. Health Log Analytics Core is a ServiceNow Store app. Once installed, the module is available, but you can choose not to interact with it. Dependency: The ServiceNow® Event Management application.

Roles

Health Log Analytics uses the following roles.

Role title [name]	Description
Event Management Administrator [evt_mgmt_admin]	Has read and write access to all Event Management features to configure Health Log Analytics.
Event Management Operator [evt_mgmt_operator]	Has read access to all Event Management features. Has write access to alerts to manage the alert life cycle. Has the itil role to manage incidents created from alerts. Can also activate operations on alerts, such as to acknowledge, close, and open an incident and run remediation.

Lists

Health Log Analytics adds the following lists.

Table [name]	Description
Advanced Log Alert Filtering [sn_occ_advanced_alert_filtering]	Controls processing of alert metrics with feedback options. Enables providing more complex feedback to the AI engine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table [name]	Description
	Beginning with the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store, this representation no longer displays.
Advanced Log Alert Filter [sn_occ_advanced_alert_filter]	Controls processing of alert metrics with feedback options. Enables providing more complex feedback to the AI engine. **Note:** This feature is supported in the Health Log Analytics application, Version 20.0.11 - July 2021, available from the ServiceNow Store.
Data Inputs [sn_occ_base_data_input_config]	Enables setting up data input connectors to stream log data into the instance.
Data Input Mapping [sn_occ_event_mapper]	Enables mapping raw log data that streams into the instance to determine how Health Log Analytics processes it.
Data Input Preprocessor [sn_occ_data_input_preprocess_screen]	Enables editing raw log data before the system maps and structures it.
Define Alert [sn_occ_custom_alert]	Enables creating custom alerts based on a predefined set of rules.
Excluded Values [sn_occ_auto_service_normalizer]	Enables configuring Health Log Analytics to stop extracting specific log data that is not descriptive enough or contains redundant text or information.
Extracted Values [sn_occ_auto_service_mapper]	Enables configuring Health Log Analytics to extract a specific term found in a field and map it to a specified component.
Features [sn_occ_features]	Lists all Health Log Analytics features. The features can be activated and deactivated.
Lexical Keywords [sn_occ_lexical_keywords]	Shows all default lexical keywords that Health Log Analytics looks for in log data.
Log Correlator [sn_occ_log_correlator]	Shows all correlators that determine whether to correlate different alerts to the same group. Enables adding correlators to Health Log Analytics.
Log Sources [sn_occ_sources]	Provides a view of all data input sources when the log data has been mapped. Enables controlling properties and correlators for these sources.
Muted Metrics [sn_occ_metric_state]	Shows all muted alert metrics together with the muted components. Enables reactivating the muted alert metrics.
Source Type [sn_occ_source_type]	Contains the source types extracted in the mapping process.
Source Type Structure [sn_occ_source_type_structure]	Enables fine-tuning the source type structure. Provides a view of all source types when the log data has been mapped. Source types define how Health Log Analytics treats specific sources and how it parses the logs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table [name]	Description
Streaming Sources [sn_occ_log_streaming_source_stats]	Enables monitoring all entities that stream log data to the instance.
System Health – Notifications [sn_occ_health_check]	Shows all system health notifications together and enables updating their status.
System Properties [sn_occ_system_settings]	Displays system properties and enables customizing their values.
Timestamp Formats [sn_occ_timestamp_formats]	Displays predefined timestamp formats. Enables defining and adding timestamp formats.
Time Zones [sn_occ_time_zones]	Enables deleting records in sn_occ_time_zones for users with the evt_mgmt_admin role.

Forms

Health Log Analytics adds the following forms.

Table [name]	Description
Beats Data Input [sn_occ_data_input_config_filebeat]	Enables configuring advanced settings for Beats data input connectors.
Beats Linux (Filebeat) [sn_occ_data_input_connector_filebeat]	Enables configuring Filebeat data input connectors.
Elasticsearch Data Input [sn_occ_data_input_config_elastic]	Enables configuring an Elasticsearch data input.
MID Server Data Input [sn_occ_data_input_config_midserver]	Enables configuring a MID Server data input.
Rsyslog [sn_occ_data_input_connector_rsyslog]	Enables configuring Rsyslog data input connectors.
Rsyslog Data Input [sn_occ_data_input_config_rsyslog]	Enables configuring advanced settings for Rsyslog data input connectors.
TCP Data Input Configuration [sn_occ_data_input_config_tcp]	Enables configuring advanced settings for TCP data input connectors.
Windows Beats [sn_occ_data_input_connector_winbeat]	Enables configuring Windows Beats data input connectors.
WinLog Beats [sn_occ_data_input_connector_winlogbeat]	Enables configuring WinLog Beats data input connectors.

Supported data inputs for Health Log Analytics

Health Log Analytics (HLA) enables you to connect your ServiceNow instance to several types of data input.

HLA supports the following data input types:

- Passive data inputs (listeners), which wait for log data to be pushed to them. These data inputs require a network port to be open on the MID Server:
  - Rsyslog
  - Beats

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Splunk
- TCP
- UDP
- MID Server
- GCP PubSub
- REST API

The Agent Client Collector data input is supported for use with the Agent Client Collector Log Analytics application, available from the ServiceNow Store.

- Active data inputs (pullers), which pull data from repositories:
  - Elasticsearch
  - Splunk Polling
  - Amazon CloudWatch
  - Amazon S3
  - Microsoft Azure Log Analytics
  - Microsoft Azure Event Hubs
  - Apache Kafka

  For all active data inputs, Health Log Analytics supports MID Server clusters for failover protection. The active data input runs on a single MID Server in the cluster. If that MID Server fails, the system moves its tasks to the next available MID Server in the cluster in a configured order.

  The Elasticsearch data input fetches data from a data repository or database using credentials. If your data is in Elasticsearch, Health Log Analytics must have the following:

  - Permissions to query Elasticsearch

    One of the following types of credentials:
    - Basic authentication (user and password)
    - AWS, for Elasticsearch on Amazon AWS Cloud
  - Network connectivity to the relevant Elasticsearch cluster

  Note: Health Log Analytics must be pointed to the correct index to start pulling the data.

In addition, Health Log Analytics supports Crible, Edge Delta, and Vector Agent data inputs. These data inputs enable HLA to process log messages that are streaming from these tools into your instance.

Native ServiceNow data inputs

Streaming logs from Cloud Observability to Health Log Analytics

Health Log Analytics can process log data it ingests from the ServiceNow® Cloud Observability application, formerly#Lightstep. HLA automatically sets up the configuration needed to enable log streaming from Cloud Observability as part of its native integration. Setting up the connection from Cloud Observability to HLA must be done in the Cloud Observability application. In HLA, you handle log records from Cloud Observability in the same way as any other Data Input Mapping and Source

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type Structure records, as explained in Log data auto-mapping and mapping in Health Log Analytics and Source type structure adjustment in Health Log Analytics. For more information about Cloud Observability, see Explore Cloud Observability documentation.

Related topics

Set up Health Log Analytics on your ServiceNow instance

General guidelines for streaming data to Health Log Analytics

This section helps you identify the appropriate data input for streaming your specific data type to the Health Log Analytics application.

Data inputs for streaming log data to Health Log Analytics

Data type	Data input	Details
Flat files	ACC agent (Agent Client Collector Framework)	Use the ServiceNow ACC agent to stream flat files directly to Health Log Analytics. This method is suitable for log files, CSVs, and other structured file formats.
Cloud applications	Depends on the cloud: Microsoft Azure Event Hubs, Amazon Data Firehose, other methods	• Microsoft Azure Event Hubs: Configure your Microsoft Azure Event Hubs to send data to Health Log Analytics using the Microsoft Azure Event Hubs data input. • Amazon Data Firehose: Set up an Amazon Data Firehose delivery stream to send data to Health Log Analytics using the Amazon Data Firehose data input. • More examples.
Data lakes	Apache Kafka, Cribl	• Apache Kafka: Export data from your data lake to a Kafka topic and use the Apache Kafka data input to stream it to Health Log Analytics. • Cribl: Use Cribl to export data from your data lake and use the Cribl data input to stream it to Health Log Analytics.
Databases	Treat as flat files	Treat database logs, such as Oracle error logs, as flat files and use the ACC agent

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data inputs for streaming log data to Health Log Analytics (continued)

Data type	Data input	Details
		to stream them to Health Log Analytics.
Third-party tools	Varies by tool	• VMware: vCenter can use UDP to stream data to the MID Server and Health Log Analytics using the UDP data input. • More examples.
APM tools	Varies by tool	• New Relic: Use the New Relic plugin to stream events directly to Health Log Analytics. Export raw logs to a flat file and use the ACC agent to stream them to Health Log Analytics. • Dynatrace: Set up the Dynatrace integration to stream data directly from Dynatrace to Health Log Analytics. • Nagios: Use the Nagios plugin to stream events directly to Health Log Analytics. Export logs to a flat file and use the ACC agent to stream them to Health Log Analytics. • Datadog: Use the Vector Agent data input to stream data from Datadog to Health Log Analytics.

MID Server proxy preconditions for streaming logs to Health Log Analytics

Requirements for using a MID Server proxy to stream log data to Health Log Analytics.

- Use a Squid MID Server proxy to enable streaming log data to Health Log Analytics. For more information, see Proxy server configuration for MID Servers used for Cloud Discovery and Cloud Provisioning and Governance.
- Install and configure the MID Server proxy with the log ingestion capability enabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

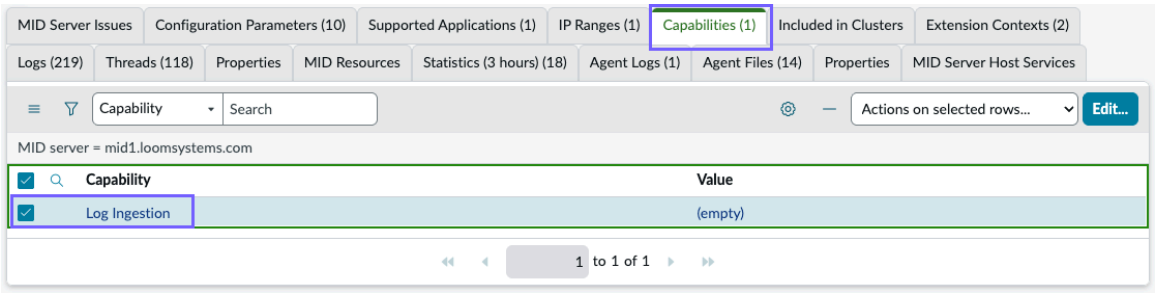

In addition, configure the following properties:

- mid.proxy.use_proxy = true
- mid.proxy.host = <proxy_ip>
- mid.proxy.port = <proxy_port>

Health Log Analytics configuration preferences

Commonly used settings for Health Log Analytics properties and general configuration.

MID Server settings

- The MID Server log ingestion capability must be enabled.

  Note: Enabling **All capabilities** on the MID Server includes enabling the log ingestion capability.

- Use dedicated MID Servers for log ingestion whenever possible.
- To enable MID Servers to run multiple products, Health Log Analytics must have at least the Java Virtual Machine (JVM) memory setting for the standard product for each MID Server thread configuration.

The preferred MID Server settings for Health Log Analytics are:

- - CPUs: 8
  - RAM: 32 GB
  - Network Bandwidth: Up to 10 Gbps
  - EBS Bandwidth: Up to 4,750 Mbps
  - Maximum Java heap size for MID Server: 8,192 MB

  With the above specifications, the expected log ingestion throughput on a Washington DC instance is as follows:

  - For a log message of 300 bytes: 20,000
  - For a log message of 1.1 KB: 12,300
  - For a log message of 2 KB: 7,970

  The minimum requirements for streaming logs to Health Log Analytics are:

  - CPUs: 4
  - RAM: 16 GB
  - Java heap size for MID Server: 8 GB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For general information, see: MID Server system requirements.

- To increase log ingestion throughput, you can either increase the ulimit or the network bandwidth, or decrease the size of the logs being streamed. The ulimit setting can be configured on an individual MID Server. However, the correlation between the ulimit and the throughput can't be modified.

  The following table lists the ulimit settings for open files relating to network throughput on the MID Server. It shows the size of the logs being streamed from the MID Server to the agent, and the gRPC streaming rate equivalent to the throughput.

  Ulimit settings in relation to throughput

Queue Type	Log line size	gRPC rate
In Memory Queue	300 bytes	18,000
In Memory Queue	1.1 KB	13,000
In Memory Queue	2 KB	10,000
Disk-based Queue	300 bytes	11,000
Disk-based Queue	1.1 KB	5,000
Disk-based Queue	2 KB	3,000

  Starting from the August 2024 release, you can enhance MID Server communication with the ServiceNow instance by using the Lightning gRPC client. This can increase log streaming speeds to Health Log Analytics by up to six times. The Lightning gRPC client requires manual configuration to activate. For more information, see the Lightning gRPC client - Enabling the new MID gRPC streaming architecture [KB1648419] article in the Now Support Knowledge Base.

- By default, the number of data inputs per MID Server is limited to 10. You can configure this limitation for an individual MID Server or for all MID Servers.
- Both in FIPS and non-FIPS mode, MID Servers with Health Log Analytics capability must run on the Java Runtime Environment (JRE) 11 or above.

  Note: To support BC-FIPS version 2.0, Health Log Analytics requires an upgrade to version 34.0.37, December 2024.

Log source retention settings

By default, log retention per source is set to three days. This setting can't be modified.

When using Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store, you can modify the log retention policy per source or for multiple sources together. For more information, see Modify the log source retention period in Health Log Analytics.

Health Log Analytics system health notifications

Health Log Analytics system health notifications help to ensure that you're aware of potentially dangerous conditions. You can manually adjust the status of a notification as the issue is being addressed.

Health Log Analytics uses self-health checks to monitor the health of its components. When it discovers a system health issue, a notification is created and an email sent to all subscribers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Health Log Analytics Administration **System Health - Notifications** list displays all HLA system health notifications, including those about internal errors. You can ask your admin to subscribe you to the system health notifications.

Health Log Analytics System Health - Notifications

Field	Description
Name	Notification header.
Description	Description of the notification.
Created	Date when the notification was created.
State	State of the notification. • **New**: The notification was newly added. • **Acknowledged**: The notification record was read. • **Work in progress**: The issue described in the notification is being addressed. Adjust the value manually. • **Closed**: The issue is no longer active. Adjust the value manually. **Note:** Changing the state of a notification does not affect the state of the issue described in it. A notification in the Closed state will still display in the table. The notification is automatically filtered out of the table only when the issue is detected as fixed (that is, when the Healthy status is **True**).
Healthy	Field that indicates the status of the issue mentioned in the notification. The table shows only notifications with a Healthy value of **False**. When the issue is detected as fixed, the value is automatically changed to **True** and the notification is filtered out of the table.

Health Log Analytics self-health checks, notifications, and proposed actions

Health Log Analytics provides self-health checks, notifications, and proposed problem-solving steps.

Health Log Analytics system health notifications

Self-health check	Notification	Proposed action
Maximum AI engine queue size exceeded	The AI engine is running at overcapacity (size=%s, threshold=%s, duration=%s). As a result, the accuracy of	Contact ServiceNow support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics system health notifications (continued)

Self-health check	Notification	Proposed action
	anomaly detection might deteriorate.	
High rate of unclassified errors in %s %s	"%s %s has a high rate of unclassified errors and is failing frequently.	Contact ServiceNow support.
Errors invoking glide API	Errors invoking '%s' for %s.	Contact ServiceNow support.
Log size limitation was exceeded	Log messages from %s: %s are being dropped because the log size limitation was exceeded.	If the difference is minor, increase the log size limitation by adjusting the property *mid.data_inputs.max_length_bytes* on the MID Server Properties form. If the logs are too long or multi-line, you can also modify the log size at the source of the logs.
Excessive data volume for data input %s.	The data volume for data input %s is too high. Log messages are being discarded.	Contact ServiceNow support.
Log messages are being dropped due to failed auto-structuring	Log messages are being dropped because auto-structuring of Source Type %s failed.	In the Source Type, find out what could be causing auto-structuring to fail. Inspect the error logs to determine which component or parsing action is failing.
Excessive number of script errors per second in %s:%s	The rate of script errors in %s: %s is too high.	Review the JavaScript Parsing Error in the source given in the message: Data Input Mapping Screen or Source Type Structure Screen.
Log messages that exceed the permitted number of properties are being discarded	The source type %s contains multiple log messages with too many properties. Such log messages are discarded. Review the Source Type Structure record.	On the relevant Source Type Structure record, review the properties that are being extracted. Check for script loops or conditions that might cause too many properties to be created. Avoid using a dynamic property name. For more information, see Refine the source type structure in Health Log Analytics.
Too many clusters are being generated. The accuracy of anomaly detection might deteriorate.	Too many clusters are being generated for Source Type %s. As a result, the accuracy of anomaly detection might deteriorate.	Inspect the created patterns and verify that the Message label is being assigned correctly. On the Source Type Structure form, review the Message property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics system health notifications (continued)

Self-health check	Notification	Proposed action
Sources are being created at a very high rate	Sources are being created at a very high rate. Data input %s was stopped automatically.	On the Data Input Mapping form, inspect the Mapping script and fix any JavaScript errors that might cause too many sources to be created.
Source Types are being created at a very high rate	Source Types are being created at a very high rate. Data input %s was stopped automatically.	On the Data Input Mapping form, inspect the Mapping script and fix any JavaScript errors that might cause too many Source Types to be created.
Metrics are being created at a very high rate	Source Type %s is generating too many metrics. The properties %s were automatically classified as invalid.	On the Source Type Structure form, observe the metrics being generated. Make sure that the Metric label is assigned only to the required properties. A dynamic property name or inaccurate JavaScript function can cause too many properties to be generated.
Processing duration of REST calls in the AI Engine is too high	Processing of REST calls in the AI Engine is taking too long.	Contact ServiceNow support.
The number of log messages being discarded is too high	Too many log messages are being discarded from source %s.	Review the log retention settings. Verify that the log lines timestamp is correct and does not represent a future or past time outside the retention period.
Auto-structuring is slowing down the AI engine	Auto-structuring is slowing down the AI engine.	Contact ServiceNow support.
Log messages are being dropped due to a timestamp parsing failure	Log messages from %s %s are being dropped due to a timestamp parsing failure. Search for 'Failed parsing timestamp from log message' in the logs for the relevant Service instance and Component and fix the issue.	On the Data Input Mapping form, locate the extracted timestamp. If it is missing, navigate from the form to the relevant Source Type Structure and open the Timestamp Parsing Error tab. The tab displays the issue that led to the parsing error. Fix the issue.
No memory consumption drop	Memory consumption hasn't dropped in the last few minutes.	Contact ServiceNow support.
Pings the Elasticsearch clusters' health endpoints to	Elasticsearch connectivity issues.	Contact ServiceNow support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Health Log Analytics system health notifications (continued)

Self-health check	Notification	Proposed action
determine the health of the connection		
Uses the Metric base client connection check to determine the health of the connection	Metric base connectivity issues.	Contact ServiceNow support.
Uses the Glide client connection check to determine the health of the connection	Glide connectivity problems.	Contact ServiceNow support.
The data input is in Failed state	The data input %s failed to stream logs.	On the Data Input form, check the connection to the MID Server and review the error message.

Health Log Analytics language and character support

Health Log Analytics provides international language support. The default language is US English.

Language support coverage

The current available languages for Health Log Analytics are as follows: US English, UK English, French, German, Italian, Japanese, and Spanish.

Note: To receive keyword-based alerts for languages other than English, keywords in the relevant language must be added to the system before you start streaming. Contact ServiceNow Support for assistance.

Character support

Health Log Analytics currently supports ASCII and Windows-1252 (commonly referred to as ANSI) character encodings, regardless of the text source. It doesn't support encodings for non-Latin scripts, such as Japanese.

Domain separation and Health Log Analytics

Domain separation is supported for Health Log Analytics. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application’s service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

Domain separation and Health Log Analytics overview

Domain separation is present in all aspects of the Health Log Analytics application. Users belonging to a specific domain see only the data existing in their own domain.

How domain separation works in Health Log Analytics

When data is domain separated using a single Health Log Analytics server, each Managed Service Provider (MSP) can see log data only in its own domain. MSPs can also see data in child domains below their domain.

Each MSP can view alerts that Health Log Analytics generates only in their domain. Actions to remediate the alerts apply only for the scope of that domain.

By default, all users and records are set to the parent domain. The admin can assign them to a specific domain.

The Health Log Analytics Domain Separation plugin must be installed before you configure your data inputs in the Health Log Analytics application. There is no setup procedure for the plugin. Install the plugin with the Health Log Analytics application Version 21.0.1 - September 2021, and then activate it. Make sure that you map your data into logical silos and configure rules and entities.

You define the domain-separated environment when you configure your data inputs. Users can use data inputs that are only available in their own domain. Health Log Analytics creates alerts only for logs that arrive in those data inputs. All relevant records and all data processing in the Health Log Analytics program flow reside in the same domain as the data input. A data input's domain name is shown in the **Domain** column displayed in the tables in your instance.

Using domain separation in your instance is transparent to Health Log Analytics. The application manages all aspects of the data, such as system settings and custom operations, separately. When a property is changed, the new value affects new sources only in the specific domain. System properties affecting the server are common to all domains because all domains use the Health Log Analytics server.

Note: Health Log Analytics supports up to 60 kilobytes events per second (EPS) across all domains, without the ability to provide a service level agreement (SLA) to a specific domain and without fairness. If a domain streams a large amount of data, the Health Log Analytics server processes it. Other domains might suffer latency, drops, or other issues as a result, even if they stream a low number of logs. On the MID Server side, 10 kilobytes EPS is expected.

Use cases

- An MSP wants to provide the Health Log Analytics application to multiple customers in a similar environment with a single instance.
- An organization with many tenants wants to isolate its sensitive data, such as security logs.
- An administrator of a tenant organization wants to define a data input only for their own domain.
- An operator in a tenant organization wants to view logs only in their own domain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- An operator in a tenant organization wants to provide feedback for alerts only in their own domain.
- An MSP Admin wants to view log data from all of their organization's tenant domains.

Related topics

Domain separation for service providers

JavaScript function test outcome fields

Description of the JavaScript function test outcome fields.

The following fields are read-only:

Field	Description
Duration (ms)	The processing time of all samples, in milliseconds.
Source types	The number of source types predicted or created for this data input.
Sources	The number of sources predicted or created for this data input.
Errors	The number of JavaScript errors that occurred in all the samples.
Unmapped	The number of log lines that could not be mapped. These log lines are mapped to default tags.
Samples	The number of log samples on which the test is run. The default value is 1,000. **Note:** You can increase the number of log samples up to a maximum of 5,000 samples in the system property data_input_mapping.max_examples. For more information, see Configure global Health Log Analytics system properties.

In the related lists, view the predicted or created source types and sources.

List/Tab	Description
Outcome Key Value Mapping	The results of running the JavaScript function on the log samples.
Source Types (from samples)	The predicted or created source types. **Note:** You can open source types that have already been created by selecting the appropriate link.
Sources (from samples)	The predicted or created sources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

List/Tab	Description
	Note: You can open sources that have already been created by selecting the appropriate link.
Test Outcome	The log samples used for testing your JavaScript function, ordered by sources and source types. The **Test Outcome** list displays the number of log samples for each source-source type combination, and is updated when the log samples are refreshed and you start testing your script. The list links to samples in the **Mapping Samples** table, which shows the outcome of the JavaScript function for each sample. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.
Source Type Structures	The source type structures for this data input. Reviewing the source type structures can help you gain a better understanding of the data flow. The **Last event time** column on this tab shows the last date and time the AI engine processed a log for the source type structure. The system refreshes this information every 5 minutes. Knowing when the last log was processed for this source type structure makes it easier to verify that log data is streaming.

Health Log Analytics severity mapping

Health Log Analytics (HLA) uses common severity values to identify severity labels in the Source Type Structure.

The **Log severity mapping** table shows how HLA identifies severity words from raw logs (left column) and maps them to its own severity values (middle column). HLA then presents them in Event Management alerts (right column).

Log severity mapping

Raw log severity	HLA log severity	Event Management alert severity
TRACE	TRACE	Keywords-based
VERBOSE	TRACE	Keywords-based
FINE	TRACE	Keywords-based
FINER	TRACE	Keywords-based

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Log severity mapping (continued)

Raw log severity	HLA log severity	Event Management alert severity
FINEST	TRACE	Keywords-based
DEBUG	DEBUG	Keywords-based
DBG	DEBUG	Keywords-based
CONFIG	DEBUG	Keywords-based
INFORMATIONAL	INFORMATIONAL	Keywords-based
INFORMATION	INFORMATIONAL	Keywords-based
INF	INFORMATIONAL	Keywords-based
INFO	INFORMATIONAL	Keywords-based
UNKNOWN	UNKNOWN	Keywords-based
NOTICE	NOTICE	LOW
NOTIFICATION	NOTICE	LOW
WARNING	WARNING	MEDIUM
WARN	WARNING	MEDIUM
ERROR	ERROR	HIGH
ERR	ERROR	HIGH
CRITICAL	CRITICAL	CRITICAL
CRIT	CRITICAL	CRITICAL
FAT	CRITICAL	CRITICAL
FATAL	CRITICAL	CRITICAL
SEVERE	CRITICAL	CRITICAL
ALERT	ALERT	CRITICAL
EMERGENCY	EMERGENCY	CRITICAL
EMER	EMERGENCY	CRITICAL

Note: If your raw log contains numerical severities or unrecognized words, map these values to the HLA Log Severity column by creating a conditional statement using the JS function in the Source Type Structure. For more information, see Source type structure adjustment in Health Log Analytics.

Sections and cards on the alert Overview tab in Health Log Analytics

The **Overview** tab helps you understand Log Analytics alerts, Log Analytics alert groups, and component-based alerts.

Overview tab sections for Log Analytics alerts

The **Overview** tab in Health Log Analytics helps you understand Log Analytics alerts.

For a detailed description of Log Analytics alerts, see Types of Health Log Analytics alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Summary

Identified issue

This card describes the issue that led to the alert. The identified issue appears on the card and in the title for the alert. Information about the alert appears in the banner.

Identified issue

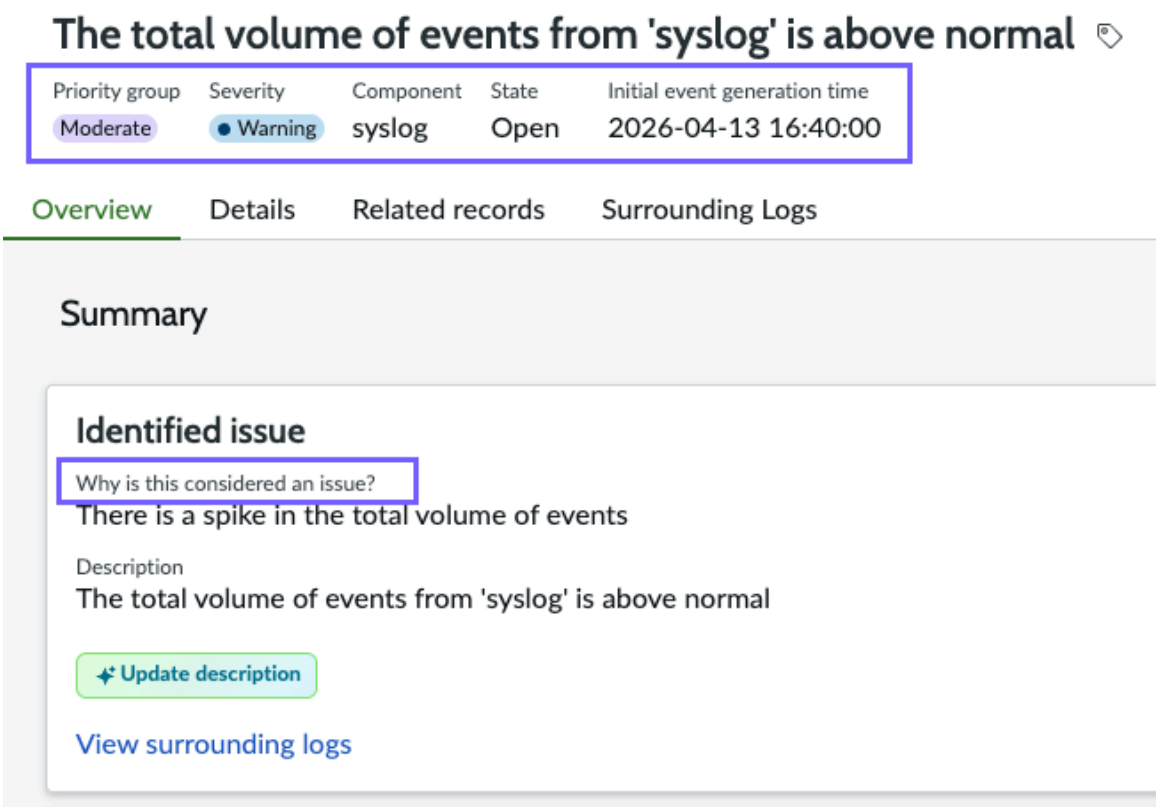

Select **Details** for more information about the alert.

Select **View surrounding logs** to view the log lines that were generated one minute before and one minute after the alert. See Analyze log lines around an anomaly.

Anomaly

This card illustrates the anomalous activity that led to the alert.

The chart shows:

- Recent anomalous activity
- Expected behavior (the learned baseline)
- Baseline values from one day earlier
- Baseline values from the previous week

In this example, the system tracks the baseline rate (the average number of events per minute) for a specific log pattern. When this typically inactive log generates a spike in events, the system detects the deviation from the baseline and generates an alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Anomaly card

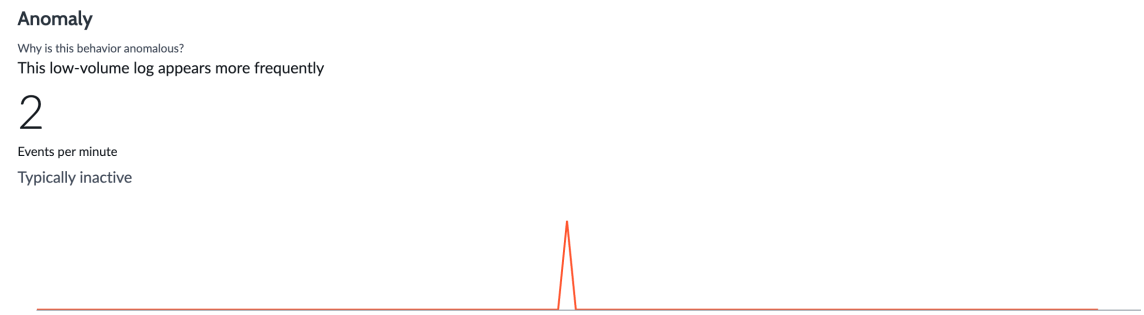

In this example, the blue line represents the current average number of events per minute. The orange-shaded area represents the baseline values for the same hour in the previous week.

Anomaly card with baseline values one week earlier

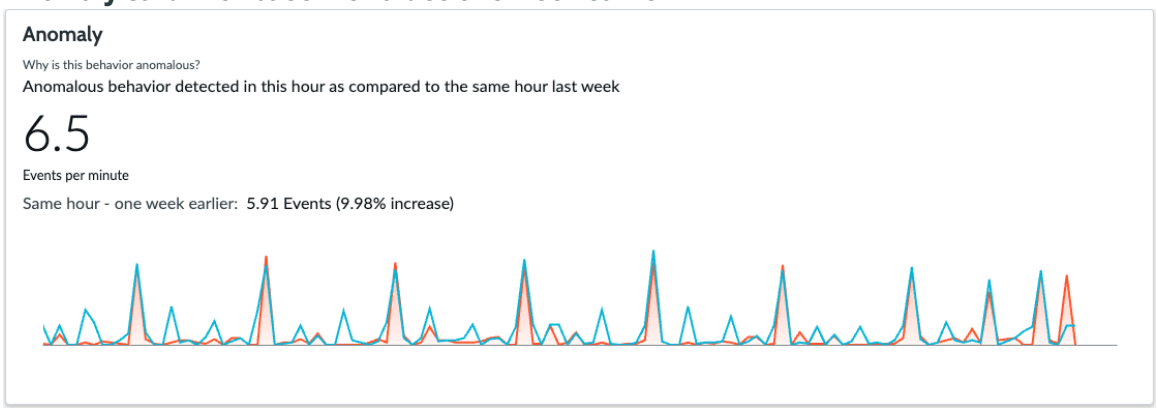

For more information on the kinds of anomalies that you might encounter, see Types of anomalous behavior in Health Log Analytics.

Impact

Configuration Items

This card provides information about the CIs that are impacted by the alert.

Impacted services

This card provides information about the services that are impacted by the alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Impact section

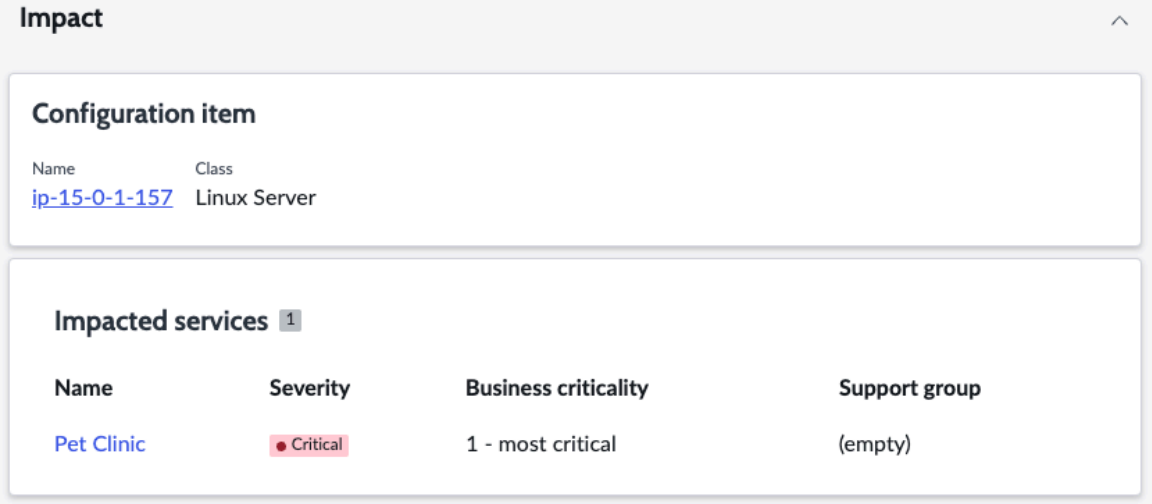

Cause

Meaningful log properties

On this card, each bar chart shows the distribution of values for a single log property that contributed to the anomaly. Each property value is associated with a color. The length of a color bar correlates to the percentage that the property value holds in comparison with all other values for the property.

Meaningful log properties

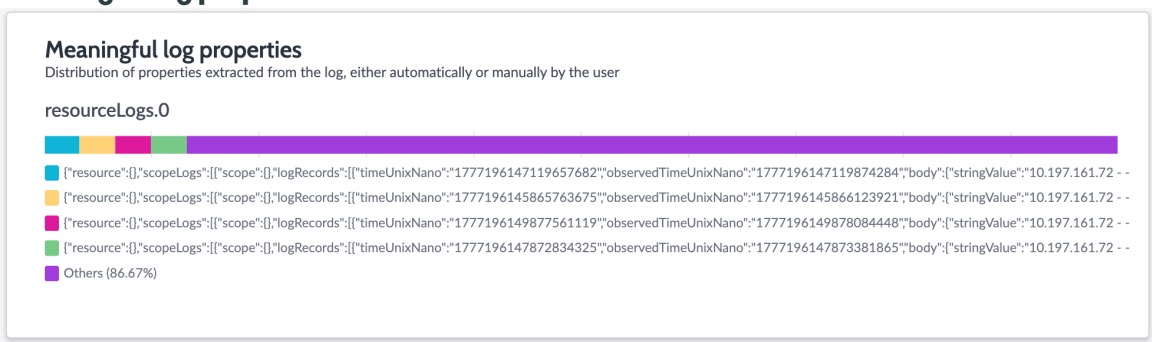

Overview tab sections for Log Analytics alert groups

The alert **Overview** tab in Health Log Analytics helps you understand Log Analytics groups.

For a detailed description of Log Analytics groups, see Types of Health Log Analytics alerts.

Summary

Identified issue

This card describes the issue that led to the alert. The identified issue appears on the card and in the title for the alert. Information about the alert appears in the banner.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

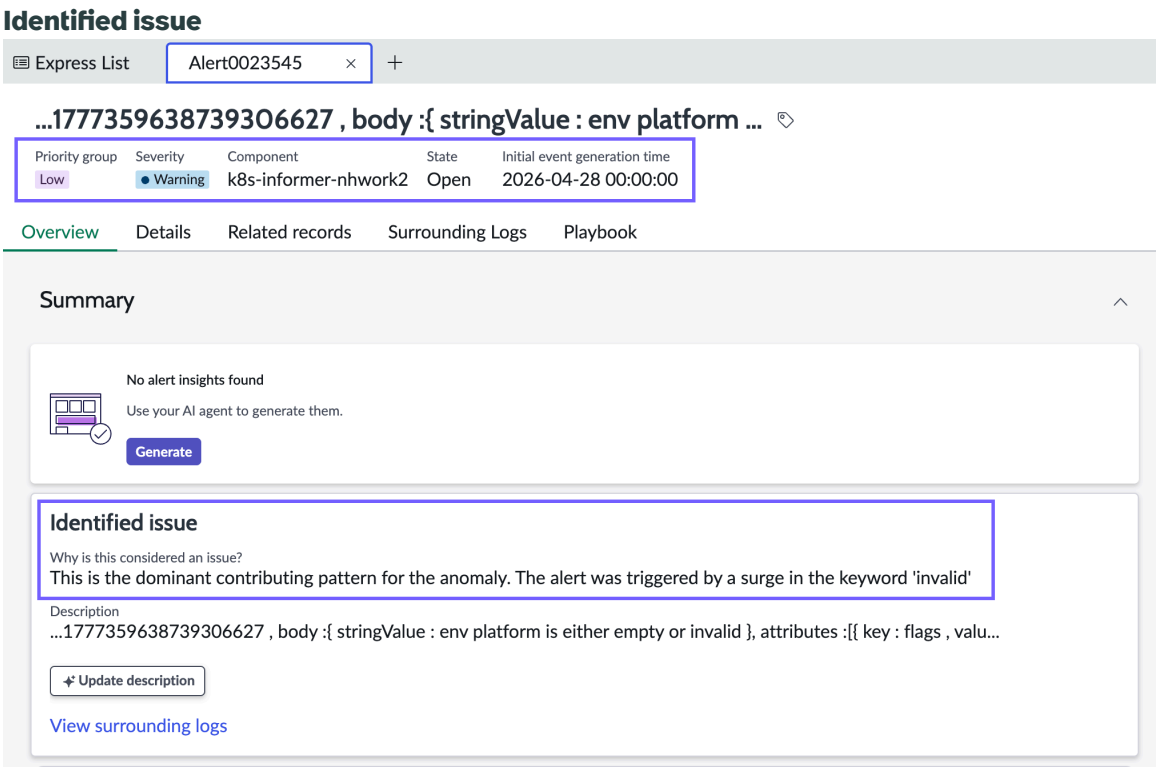

Select **View correlations** to view the list of correlations that relate the Log Analytics alerts.

Correlations list

During initial analysis, alerts are scored. Each correlation in the alert's log data with another alert contributes to the score. The higher the score, the more likely the alert is to be included as a Log Analytics alert in a Log Analytics group.

The following kinds of data are considered when determining whether alerts are correlated:

- Time: The events all occurred within a configured time interval.
- Metadata: The alerts have matching values in log-line metadata. For example, all alerts involve the same host.
- Message text: The message text in the log data is similar or identical between alerts.
- Trend: The alerts show a similar tendency in values or rates. For example, a particular metric value is increasing in all alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Correlations

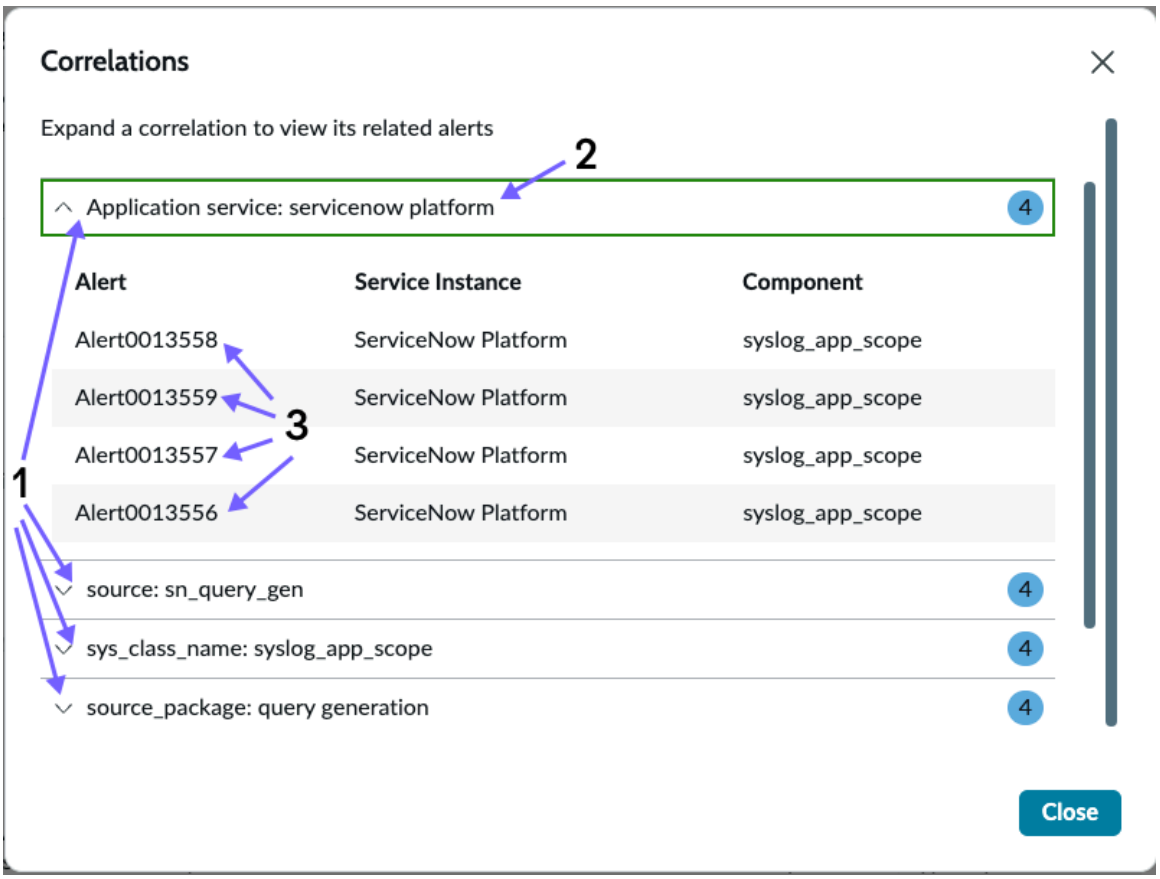

1. List of correlations: The first correlation in the list is expanded to show the individual Log Analytics alerts that are correlated and the log correlator that the alerts share.
2. An individual log correlator: The identifier for a group of correlated Log Analytics alerts. The alerts are grouped by the log-line data or metadata that is common to the alerts (for example, IP address, host name, or user name). The number in the blue square indicates the number of correlated alerts.
3. Log Analytics alerts that are correlated.

Alerts in group

For a Log Analytics alert, the Alerts in group card shows the Log Analytics alerts that are grouped under the Log Analytics alert. Select a Log Analytics alert to view its details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

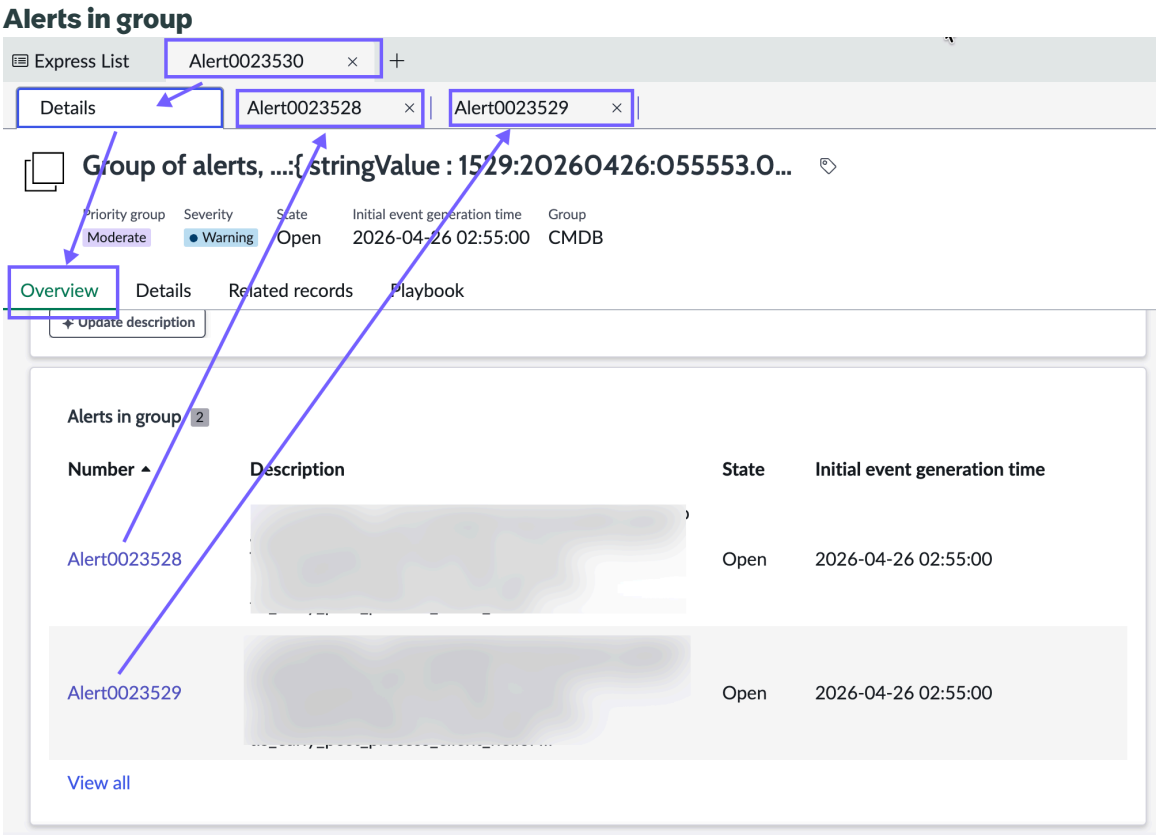

Select **View all** to the view the list of all Log Analytics alerts in the group and relevant information about them. You can also view the Alerts in group list by selecting the **Related records** tab and then selecting **Alerts in group**.

For each Log Analytics alert in the group, the following information is available.

Alerts in group

Column	Description
Number	The number of the alert. Select the number to view detailed information for an alert. This field is automatically set.
Initial event generation time	The time when the event that generated the alert first occurred. **Note:** Time here is the ServiceNow processing time, not the source system time.
Group	Type of group that the alert belongs to: a standalone Log Analytics alert or a Component-based alert.
Description	Anomalous pattern or metric that caused the alert to be generated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alerts in group (continued)

Column	Description
Severity	Severity value for the alert. The available values are: • **Critical**: Immediate action is required. Either the resource is not functional or critical problems are imminent. • **Major**: Major functionality is severely impaired or performance has degraded. • **Minor**: Either performance has degraded or there is a partial, non-critical loss of functionality. • **Warning**: Attention is required even though the resource is still functional. • **Info**: An informational message. An alert is created, but the resource is still functional. • **Clear or Resolved**: No action is required. An alert is not created from this event. Existing alerts are closed.
Priority group	Priority group that indicates the order in which to resolve alerts. Choices are as follows: • **Urgent** • **High** • **Moderate** • **Low** The priority group value is more important than severity alone. For example, a high priority and low severity alert should be addressed before a low priority and high severity alert. For information on how priority is calculated, see Alert priority.
State	Processing state of the alert. A newly generated alert is in the **Open** state. Other states are as follows: • **Reopen**: A previously closed alert is open again, and it requires your attention. • **Flapping**: The alert is receiving identical events from the same source at high frequency. This state can cause an alert to re-open from the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alerts in group (continued)

Column	Description
	Closed state, resulting in a high frequency of changes between Open and Closed states. • **Closed**: The alert is closed and does not require any further action. You close an alert when it is remediated.
Configuration item	CI in the CMDB. The CI is applied to by the alert.
Node	Node field that is received in the log message. The event described in the log message occurred on this node. Often, the node is the name of the CI that is associated with the alert. For example, a computer name, IP address, FQDN, or MAC address.
Source	All Health Log Analytics alerts have the value **Log Analytics** in the **Source** column to indicate that the Health Log Analytics app generated the alert.
Metric name	Name of the metric whose anomalous behavior led to the alert. For example, the *`I/O request`* in the case that the I/O request took longer than 15000 ms to complete.
Updated	Most recent time when the alert information or state was updated.

Impact

Configuration Items

This card provides information about the CIs that are impacted by the alert.

Impacted services

This card provides information about the services that are impacted by the alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Impact section

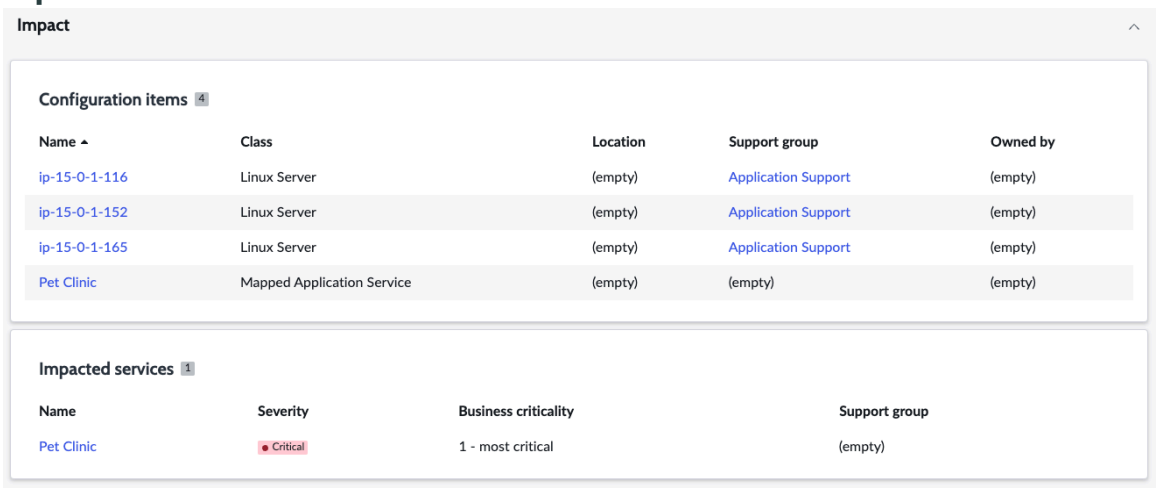

Overview tab sections for Component-based alerts

The alert **Overview** tab in Health Log Analytics helps you understand Component-based alerts.

For a detailed description of Component-based alerts, see Types of Health Log Analytics alerts.

Summary

Identified issue

This card describes the issue that led to the alert. The identified issue appears on the card and in the title for the alert. Information about the alert appears in the banner.

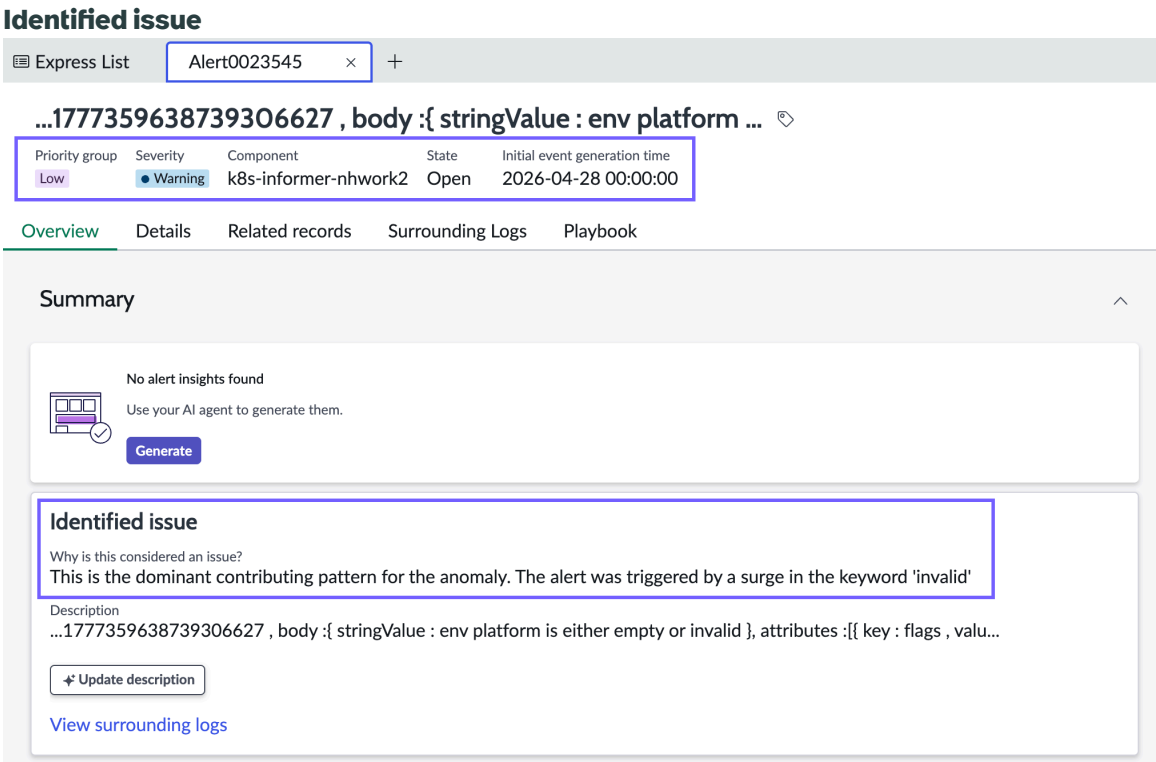

Select the **Details** tab for more information about the alert.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select the **View surrounding logs** link to view the log lines that were generated one minute before and one minute after the alert. For more information, see Analyze log lines around an anomaly.

Anomaly

The **Anomaly** card illustrates the anomalous activity that led to the alert. The chart shows:

- Recent anomalous activity
- Expected behavior (the learned baseline)
- Baseline values from one day earlier
- Baseline values from the previous week

In this example, the system tracks the baseline rate (the average number of events per minute) for a specific log pattern. When this typically inactive log generates a spike in events, the system detects the deviation from the baseline and generates an alert.

Anomaly card

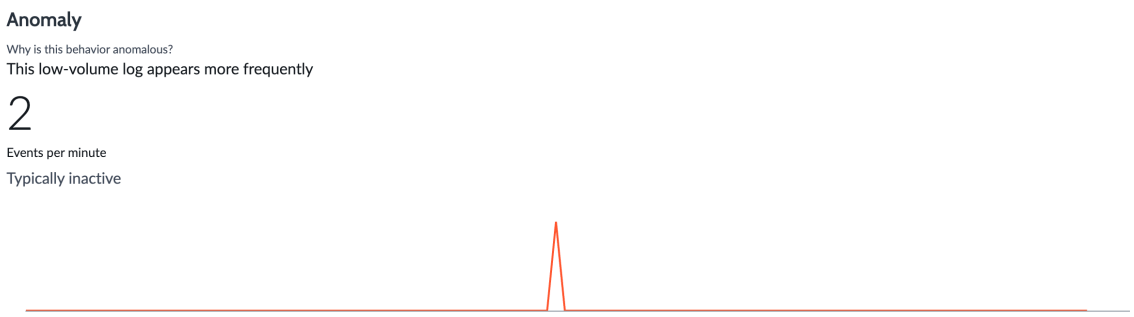

In this example, the blue line represents the current average number of events per minute. The orange-shaded area represents the baseline values for the same hour in the previous week.

Anomaly card with baseline values one week

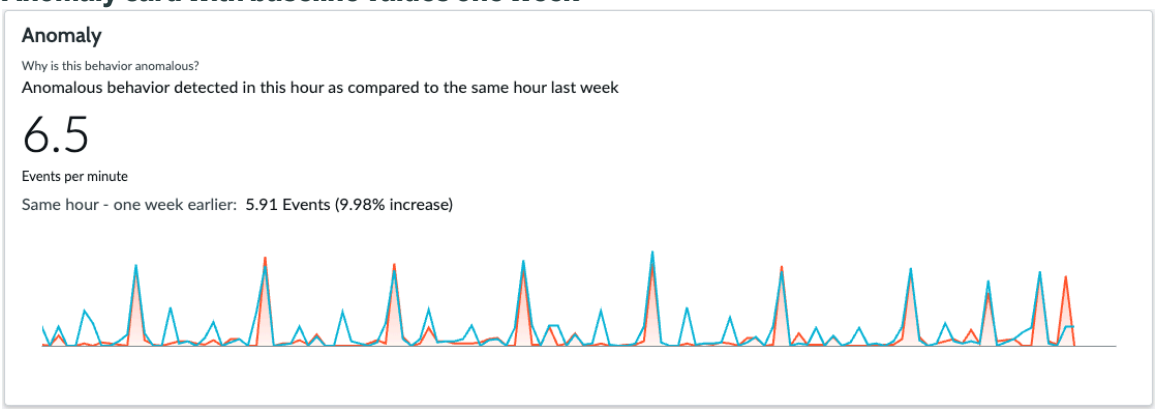

For more information on the kinds of anomalies that you might encounter, see Types of anomalous behavior in Health Log Analytics.

Impact

Configuration Items

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This card provides information about the CIs that are impacted by the alert.

Impacted services

This card provides information about the services that are impacted by the alert.

Impact section

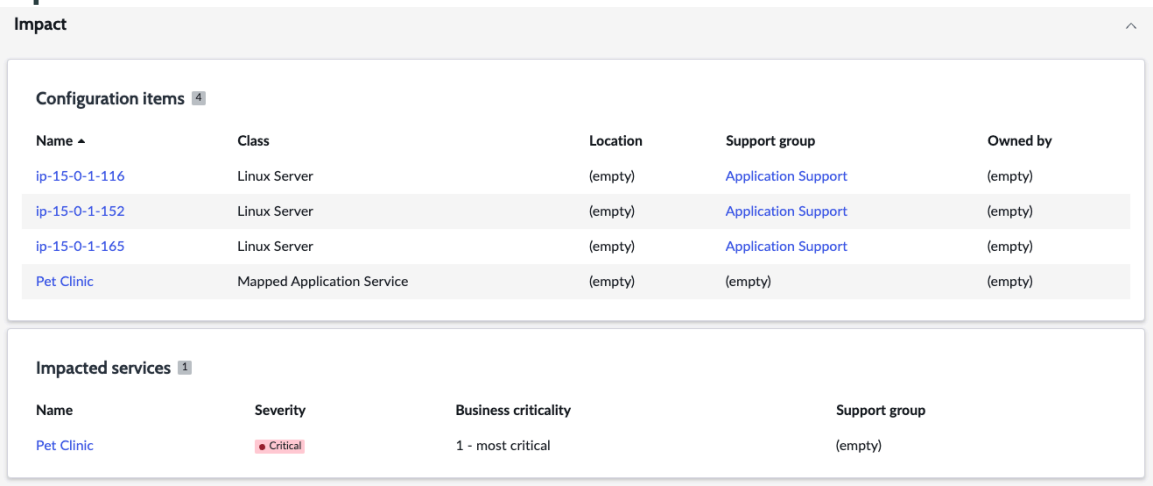

Cause

Meaningful log properties

On this card, each bar chart shows the distribution of values for a single log property that contributed to the anomaly. Each property value is associated with a color. The length of a color bar correlates to the percentage that the property value holds in comparison with all other values for the property.

Meaningful log properties

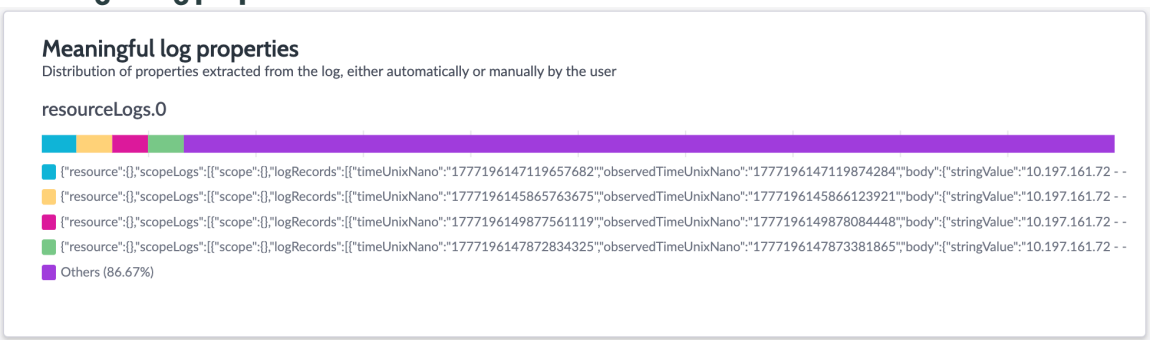

Integration configuration fields

This section provides descriptions of the fields on the integration configuration forms for Health Log Analytics.

ACC Log Analytics (ACC-L) integration configuration fields

Description of the fields on the ACC Log Analytics (ACC-L) integration configuration forms for Health Log Analytics.

For the ACC Log Analytics (ACC-L) integration setup procedure, see Set up an ACC Log Analytics integration for Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details

Field	Description
Transport	The protocol used to send the log data. This field is read-only. The ACC-L integration uses a ServiceNow Agent to send data.
Port	The port on the MID Server to which the ACC-L agent connects. This field is required. **Important:** This port can't be the same as the MID Web Server port.
Description	Option to add a brief description of the integration to help identify it.
ACC Listener	
MID Server	The MID Server to which the logs stream. This field is required. **Note:** • Only one ACC-L integration can be defined per MID Server. You can only select MID Servers with AgentClientCollector capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. Exception: if 10 non-ACC-L integrations are already defined for a MID Server, you can still add one ACC-L integration, for a total of 11. • When you submit the form, this field is set to read-only. • **Important:** The configured MID Server can't be changed after the integration is activated.
MID Web Server port	Port used by the MID Server to receive data from ACC agents. This field is auto-populated and read-only when a MID Server has been configured.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: The MID Web Server port must be reachable through your firewall.

Advanced settings

Field	Description
Lookup hostnames	Option to perform DNS lookup to resolve IPs to hostnames.
Use SSL	Option to use SSL.
Client inactivity timeout (seconds)	The timeout, in seconds, to close an inactive channel.
Sub sample drop ratio	The ratio of logs to drop. The default value is -1: no logs are dropped. For example: If you want one out of every five logs to be dropped, change the value to 5.
Max length in bytes	The maximum length of log messages, in bytes.
Character encoding	The character encoding for this integration: UTF-8. This field is read-only.
Worker thread count	The number of threads that handle incoming data.
Sub sample receive ratio	The ratio of logs to receive. The default value is -1: no logs are received. For example: If you want one out of every five logs to be received, change the value to 5.
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.
Drop if queue is full	Option to discard logs if many processes are waiting in the queue to access the MID Server.

Amazon CloudWatch integration configuration fields

Description of the fields on the Amazon CloudWatch integration configuration forms for Health Log Analytics.

For the Amazon CloudWatch integration setup procedure, see Set up an Amazon CloudWatch integration for Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My CloudWatch integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID server name	(Only when the **Execute on** field is set to Specific MID Server) The MID Server to which log data from Amazon CloudWatch is pulled. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of integrations streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Data source	The source of the log data that the integration pulls to your ServiceNow instance: CloudWatch. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Set data retrieval method

Field	Description
How to extract data from AWS CloudWatch?	
Authentication method	The AWS credentials needed to access resources in AWS service accounts. This field is required. Choose AWS credentials from the list, or add new AWS credentials to the list by selecting **Manage Credentials** and then defining new credentials on the **AWS Credentials** page.
Group Name(s)	The log group(s) to fetch log data from in Amazon CloudWatch. This field is required. If you want the integration to search multiple log groups, specify the relevant groups in a

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set data retrieval method (continued)

Field	Description
	comma-separated list. To fetch log data from all groups, use an asterisk (*) as a wildcard character.
AWS Region	The AWS region where the Amazon CloudWatch cluster is hosted. For example: us-west-1. This field is required. For a list of AWS regions, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions
Which data to send to ServiceNow?	
Filter pattern	The pattern by which to filter incoming events. Various types of filter patterns are supported. For example: • A pattern for fetching log events that contain a single term. • A pattern for fetching log events that contain multiple terms. • A pattern for fetching log events that include a term and exclude another. **Note:** Filter patterns are case sensitive. For more information, see https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/FilterAndPatternSyntax.html
Prefix	Option to define a name prefix for Amazon CloudWatch log streams to read from. The integration reads only from log streams with this name prefix. **Note:** Only a single log stream prefix per data item is supported. For multiple prefixes, create multiple integrations.

Advanced settings

Field	Description
Connection timeout (ms)	The number of milliseconds to wait before timing out the AWS connection attempt.
Socket timeout (ms)	The number of milliseconds to wait before timing out a data transfer over an established connection.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Sub sample drop ratio	The number of logs to batch together, out of which one will be dropped. The default value is -1: No logs are dropped. This setting is used to reduce the number of fetched logs. For example, if you want one log out of every five to be dropped, change the value to 5.
Polling interval	The interval, in seconds, to wait before polling for new logs.
Character encoding	The character encoding used for this integration: UTF-8.
Batch size	The maximum number of logs retrieved per query.
Max length in bytes	The maximum length of log messages, in bytes.
Sub sample receive ratio	The number of logs to batch together, out of which one log will be received and the remaining ones dropped. The default value is -1: No logs are received. This setting is used to reduce the number of received logs. For example, if you want to receive one log out of every five, change the value to 5.
Sleep interval (seconds)	The interval, in seconds, to wait before querying again after a query has returned no logs.
Default timezone	The default timezone if the log date and time doesn't include timezone information.
Drop if queue is full	Option to select discarding logs if there is a load on the MID Server.

Amazon S3 integration configuration fields

Description of the fields on the Amazon S3 integration configuration forms for Health Log Analytics.

For the Amazon S3 integration setup procedure, see Set up an Amazon S3 integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My S3 integration. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID server name	(Only when the **Execute on** field is set to Specific MID Server) The MID Server to which log data from Amazon S3 is pulled. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Data source	The source of the log data that the integration pulls to your ServiceNow instance: S3. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Set data retrieval method

Field	Description
Bucket name	The Amazon S3 bucket to stream. This field is required.
Region	The AWS region where the Amazon CloudWatch cluster is hosted. For example: us-east-2. This field is required. For a list of AWS regions, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions
Authentication method	The AWS credentials needed to access resources in AWS service accounts. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set data retrieval method (continued)

Field	Description
	Choose AWS credentials from the drop-down list, or add new AWS credentials to the list by selecting **Manage Credentials** and then defining new credentials on the **AWS Credentials** page.

Advanced settings

Field	Description
Include	The file or object to include when the system performs pattern matching.
Exclude	The file or object to exclude when the system performs pattern matching.
Bucket prefix	The prefix for the objects to be fetched. The integration fetches only objects with this prefix. **Note:** This field is typically used to fetch only specific folders within a bucket.
Buffer size	The size of the download buffer, in bytes.
Threadpool size	The number of concurrent threads that are downloading files from the bucket.
Anonymous	Option to force accessing the bucket without credentials, while ignoring default credentials if they exist.
Traverse	Option to access each directory in the bucket.
Sorting object by last modified time	Option to fetch objects if the objects in the bucket are sorted by time.
Sub sample drop ratio	The number of logs to batch together, out of which one will be dropped. The default value is -1: No logs are dropped. This setting is used to reduce the number of fetched logs. For example, if you want one log out of every five to be dropped, change the value to 5.
Sleep interval (seconds)	The interval, in seconds, to wait before querying again after a query has returned no events.
Max length in bytes	The maximum length, in bytes, of events.
Connection timeout (ms)	The number of milliseconds to wait before timing out the attempt to establish a connection to AWS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Max retry per file	The maximum number of times to retry file processing in case of failure.
Socket timeout (ms)	The number of milliseconds to wait before timing out data transfer over an established connection.
Default timezone	The default timezone if the log doesn't include timezone information.
Sub sample receive ratio	The number of logs to batch together, out of which one log will be received and the remaining ones dropped. The default value is -1: No logs are received. This setting is used to reduce the number of received logs. For example, if you want to receive one log out of every five, change the value to 5.
Polling interval	The interval, in seconds, to wait before polling for new events.
Character encoding	The character encoding used for this integration: UTF-8
Drop if queue is full	Option to select discarding logs if there is a load on the MID Server.

Apache Kafka integration configuration fields

Description of the fields on the Apache Kafka integration configuration forms for Health Log Analytics.

For the Apache Kafka integration setup procedure, see Set up an Apache Kafka integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Kafka integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID server name	(Only when the **Execute on** field is set to Specific MID Server)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	MID Server to which log data from Apache Kafka is pulled. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify the maximum number by adding the property `sn.occ.log_ingestion.max_datainputs_per` to the MID Server and then changing the default value. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Data source	The source of the log data that the integration pulls to your ServiceNow instance: Kafka. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Set data retrieval method

Field	Description
Kafka node names	A comma-separated list in the format HOST:PORT,HOST:PORT. This field is required. The list does not have to include all the Apache Kafka Cluster servers.
Topics	A comma-separated list of topics to which the data input must subscribe. This field is required.
Kafka credentials	The Apache Kafka credentials.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set data retrieval method (continued)

Field	Description
	You can select existing Kafka SSL credentials, or create new ones by selecting **Create Kafka credentials** from the drop-down list. For a description of the fields on the Kafka SSL credentials form, see Kafka SSL credentials fields.
Group ID	The name of the Apache Kafka Consumer Group.

Advanced settings

Field	Description	Default value
Timeout	The time, in milliseconds, spent waiting in the poll if data is not available in the topics.	500
Default timezone	The time zone of events that the system will use if a log does not specify the time zone. By default, the system uses GMT in such cases, but you can specify a different time zone.	GMT
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.	-1
Character encoding	The character encoding for this data input. This field is read-only.	UTF-8
Node discovery timeout	The time, in milliseconds, before node discovery times out.	30
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched events.	-1
Max length in bytes	The maximum length, in bytes, of events.	32766
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	False

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Amazon Data Firehose integration configuration fields

Description of the fields on the Amazon Data Firehose integration configuration forms for Health Log Analytics.

For the Amazon Data Firehose integration setup procedure, see Set up an Amazon Data Firehose integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Firehose integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Transport	The protocol used for streaming log messages to your instance: Amazon CloudWatch Logs. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Cribl integration configuration fields

Description of the fields on the Cribl integration configuration forms for Health Log Analytics.

For the Cribl integration setup procedure, see Set up a Cribl integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Cribl integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Service Instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
MID Server name	The MID Server to which the Cribl logs are streamed. This field is required.
Port	The port on the MID Server. This field is required. Choose a port within the range 6000-6200 from the array. The port must not be occupied by another process. Make sure that your

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	organization's security team opens the selected port on the MID Server.
Description	Option to add a brief description of the integration to help identify it.
Transport (for push integrations)	The protocol used for streaming log messages to your ServiceNow instance: TCP. This field is read-only.
Source (for pull integrations)	The source of the log data that the integration pulls to your ServiceNow instance: Cribl. This field is read-only.

Set up connection

Field	Description
Cribl URL	The URL of the Cribl instance. This field is required. HLA uses this URL to establish the connection with the Cribl instance and make the necessary API calls. For example: • On-Prem: https://my-cribl.company.com:9000 • Cloud: https://myorg-12345abcde.cribl.cloud • Sandbox: https://sandbox-xyz789.sandbox.cribl.io
Authentication method	The authentication method used by the Cribl integration. This field is required. The Cribl integration supports: • Token-based authentication: Required for Cribl Cloud and also supported for self-hosted (On-Prem) Cribl instances. • Basic authentication: Preferred for self-hosted (On-Prem) Cribl instances. Both of these authentication methods use the default credential alias: Cribl_Credential_Alias. For information about setting up credentials for the authentication method used by the Cribl integration, see Set up Cribl integration authentication credentials. **Important:** Verify that only one active credential record of each credentials type exists under the credential alias.
Worker Group	The Cribl Worker Group from which to stream log data to the instance. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up connection (continued)

Field	Description
	The drop-down list displays Worker Groups based on the provided Cribl cloud instance URL and authentication method. For example: • default • dev-workers • prod-workers • staging-workers
Cribl route	The Cribl route that directs log data to a specific destination. This field is auto-populated with a drop-down list of all available routes, each configured with a default destination. Select the desired route from the list. This field is required. For example: • default • demo-route
Cribl destination	The Cribl destination to which the log data is directed. This field is auto-populated and read-only. When HLA populates this field, it provides a link to the destination on the Cribl server. For example: sn_hla_cribl_tcp_json_abcdef

Edge Delta REST integration configuration fields

Description of the fields on the Edge Delta REST integration configuration form for Health Log Analytics.

For the Edge Delta REST integration setup procedure, see Set up an Edge Delta REST integration for Health Log Analytics.

Provide details

Field	Description
Integration name	Unique name of this integration. For example: My Edge Delta REST integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID web server	The MID Web Server to use. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
Service instance	The service instance to which to bind the log data. This field is required.
Port	The port for the MID Server. This field is read-only. It is auto-populated when the MID Web Server is chosen.
Transport	The transport method used for streaming log messages to your ServiceNow instance: REST. This field is read-only.
MID server name	The MID Server to which the log data is pushed. This field is auto-populated when the MID Web Server is chosen. It is read-only.
Description	Option to add a brief description of the integration to help identify it.

Edge Delta TCP integration configuration fields

Description of the fields on the Edge Delta TCP integration configuration forms for Health Log Analytics.

For the Edge Delta TCP integration setup procedure, see Set up an Edge Delta TCP integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Edge Delta TCP integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID server name	MID Server to which the Edge Delta logs are streamed. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.
Port	The port on the MID Server used for listening on logs. This field is required. Choose a port within the suggested range from the array. The port must not be occupied by another process. Ensure that your organization's security team opens the selected port.
Service instance	The service instance to which to bind the log data. This field is required.
Transport	The protocol used for streaming log messages to your ServiceNow instance: TCP. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Advanced settings

Field	Description
Lookup hostnames	Option to perform DNS lookup to resolve IPs to hostnames. The default value is false.
Sub sample drop ratio	The ratio of logs to drop. The default value is -1: no logs are dropped. For example: If you want one out of every five logs to be dropped, change the value to 5.
Sub sample receive ratio	The ratio of logs to receive. The default value is -1: no logs are received. For example: If you want one out of every five logs to be received, change the value to 5.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Max length in bytes	The maximum length of log messages in bytes. The default value is 32766.
Character encoding	The character encoding for this data input. Default is UTF-8.
Line breaker delimiters	The line break character separating the raw log lines. Splitting values must be separated by a comma followed by a space: ", ". For example: "\r, \n, , splitHere, #".
Boss thread count	The number of threads that manage connections.
Worker thread count	The number of threads that handle incoming data.
Read timeout seconds	The timeout in seconds since the last read. When the timeout expires, the system closes the channel.
Default timezone	The time zone of events that the system uses if a log does not specify the time zone. By default, the system uses GMT, but you can specify a different time zone.
Drop if queue is full	Option to discard logs if there is a load on the MID Server.

Elasticsearch integration configuration fields

Description of the fields on the Elasticsearch integration configuration forms for Health Log Analytics.

For the Elasticsearch integration setup procedure, see Set up an Elasticsearch integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Elasticsearch integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Execute on	Option to determine whether to use a specific MID Server or a specific MID Server cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
MID server name	(Only when **Execute on** is set to Specific MID Server) The MID Server to which log data from Elasticsearch indices is pulled. This field is required.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server Cluster) The MID Server cluster to which the log data is pulled. This field is required. When you select a cluster, the MID Servers in the selected cluster and their status are displayed. The integration runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the integration tasks to the next available MID Server in the cluster according to the configured order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the integration form, the MID Server Clusters list displays only failover clusters. • The MID Server cluster must include only MID Servers that support basic authentication. mTLS is not supported for log ingestion. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • If Elasticsearch uses client certificate or CA certificate authentication, all MID Servers in the cluster must have the appropriate certificates. • The default maximum number of integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 integrations running on it, even when that MID Server is down.
Service instance	The service instance to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.
Data source	The source of the log data that the integration pulls to your ServiceNow instance: Elastic. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Data retrieval method

Field	Description
Server URL	The URL used to access the cluster. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data retrieval method (continued)

Field	Description
Authentication method	The authentication method used to authenticate the integration to Elasticsearch. Default is none. When you select the authentication method, the corresponding credentials fields display on the form. **Note:** As an admin, you can create an authentication method by navigating to **All** > **Health Log Analytics** > **Authentication Methods** and selecting **New**.
Index prefix	Prefix prepended to the names of the Elasticsearch indices from which you want to read data. The integration only reads data from indices that match the configured prefix. For example: network-logs-* matches indices such as network-logs-2024.01.01. This field is required. This setting ensures that HLA only ingests data from the relevant indices. For example: only-read-these-indices-*
Document timestamp field	Timestamp field in documents stored in the read indices. This field is required.
Timestamp field format	Format of the timestamp field in the documents. If no format is specified, the default Unix epoch time format is used, in milliseconds. For example: 1684168407 (May 15, 2023 4:33:27 PM)
Term filters	JSON map of the terms to filter. **Note:** Avoid using the term query for text fields. If the target field is mapped as both text and keyword, reference the keyword by using fieldname.keyword. For example: {"severity": ["error", "warning"]}

Advanced settings

Field	Description
Max connections per route	The maximum number of connections to be opened per node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Max scroll slices	The number of shards configured for the relevant index in Elasticsearch. This number tells Elastic how many parallel queries to execute in each polling request.
Proxy host	Host name of the HTTP proxy through which requests are sent.
Proxy port	Port of the HTTP proxy through which requests are sent.
Use MID certificate policy check	Option to enable the MID certificate policy check. Select this option if you want to ship your logs encrypted using SSL TLS. Then navigate to **All** > **MID Server** > **MID Security Policy** and add the MID certificate policy check to the table. For more information, see MID Server certificate check policies.
Use cross-cluster search	Option to search for data across Elasticsearch clusters. When this check box is selected, the **Clusters to search** field displays. **Note:** Your settings in the **Use minimal privileges** check box and the **Delay in reading current timestamp (seconds)** field on the **Advanced configuration** form affect how data is collected across multiple clusters.
Clusters to search	The Elasticsearch clusters to search.This field displays only when the **Use cross-cluster search** check box is selected. Do one of the following: • Leave this field empty or enter "*" to search all the remote clusters defined in Elasticsearch. • Specify the clusters to search in a comma-separated list. **Note:** To also search the local cluster, add a comma at the beginning or end, or add two commas in succession to the list. For example: "east,,west" or ",east,west" or "*,"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Use minimal privileges	Option to read log data directly from the Elasticsearch indices with the configured prefix. • When selected, the integration reads the log data directly from the Elasticsearch indices with the configured prefix. To perform this task, it needs only read privileges. **Note:** When this check box is selected and you're using cross-cluster search, data is collected from all the clusters simultaneously. • When clear, the integration fetches all indices with the prefix, filters them, and reads the log data from the filtered indices. Performing this task requires additional privileges. **Note:** Leaving this check box clear when using cross-cluster search affects how data is collected from the clusters. For more information, see the Enabling and Using Cross-Cluster Search for Elasticsearch Data Inputs in Health Log Analytics [KB1556079] article in the Now Support Knowledge Base. For additional information about streaming logs using the Elasticsearch integration, see the Stream logs using Elasticsearch data input - Advanced guide [KB1080162] article in the Now Support Knowledge Base.
Max documents per query	Maximum number of documents fetched in a single query.
Sliced-scrolling tie breaker	Value used to slice the data. Each slice is scrolled in parallel. Default: _id
Search-after tie breaker	Unique value per document to use as tiebreaker when sorting log entries by timestamp.
Use search-after API	Option to toggle between using sliced-scrolling and search-after APIs. **Note:** Sliced-scrolling APIs are preferable when reading historical data, while search-after APIs are better for reading real-time data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Index time-suffix format	Format of the time suffix when using time-based index names, such as [logstash-]YYYY.MM.DD. When using aliases, leave this field empty. For example: uuuu.MM.dd
Data reading timeout (milliseconds)	The duration of time, in milliseconds, before a request to the Elasticsearch cluster times out.
Index discovery interval (seconds)	The number of seconds between intermittent MID Server requests to the Elasticsearch cluster for new indices from which to read data.
Scroll context time (milliseconds)	The lifetime of the created scroll when using the scroll API to read data from Elasticsearch. For more information, see the Elasticsearch scroll API documentation.
Event processor workers	The maximum number of CPU cores used in parallel to process events fetched from Elasticsearch. A higher setting increases the data input throughput at the cost of higher CPU usage.
Worker queue size	The maximum number of batches to queue for processing. A higher setting increases throughput, at the cost of higher RAM usage.
Default timezone	The time zone of events that the system will use if a log does not specify the time zone. By default, the system uses GMT in such cases, but you can specify a different time zone.
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched events.
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.
Character encoding	The character encoding for this data input.
Sleep interval (seconds)	The interval, in seconds, to wait before querying again after a query has returned no data.
Max length in bytes	The maximum length, in bytes, of log messages.
Delay in reading current timestamp (seconds)	The number of seconds before current time to query to include delayed data.The configured

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
	number of seconds is subtracted from the current time for reading the last timestamp. **Note:** If this value is 0 and data is collected from multiple clusters simultaneously, the query may not include data that was sent with a delay on one of the clusters,
Polling interval	How frequently the system polls for new log data.

GCP PubSub integration configuration fields

Description of the fields on the GCP PubSub integration configuration forms for Health Log Analytics.

For the GCP PubSub integration setup procedure, see Set up a GCP PubSub integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My GCP PubSub integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Service Instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID Server name	(Only when **Execute on** is set to Specific MID Server.) The MID Server to which the logs are streamed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: • You can only select a MID Server with log ingestion capability that supports basic authentication. MID Servers that support mTLS are not listed. • If log ingestion is not enabled on the selected MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. This field is required.
MID Server cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	For more information about MID Server clusters, see Configure a MID Server cluster.
Source	The source of the log data that the integration pulls to your ServiceNow instance: PubSub. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Set up connection

Field	Description
Project ID	The project ID of the Google Cloud project. For example, my-project-id. This field is required. You can obtain the project ID from your Google Cloud homepage by selecting My Project. The project ID will be displayed.
Authentication method	The credential alias ServiceNow® uses to access Google Cloud. The default value is sn_occ.HLA_Credential_Alias. This field is required. **Note:** You can add more credential aliases to the list of available aliases to choose from by selecting **Manage credential aliases** from the list and adding a credential alias.
Topic name	The Google Cloud Pub/Sub topic to which to subscribe Health Log Analytics. For example, my_topic. This field is required. You can obtain the topic name from your Google Cloud homepage by searching for Pub/Sub. All topic names will be displayed.
Subscription name	The subscription Health Log Analytics uses to receive log data that is published on the Google Cloud Pub/Sub topic. If you leave this field blank, Health Log Analytics uses ServiceNow-Subscription.

Advanced Settings

Field	Description	Default value
Subscriber thread pool size	The number of concurrent threads that are downloading	1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced Settings (continued)

Field	Description	Default value
	files from the Google Cloud Pub/Sub topic.	
Sub sample receive ratio	The ratio of logs to receive. This setting determines the number of logs to group together, out of which one is received. It reduces the number of logs you receive. The default value is -1, meaning no logs are received. For example, if you want to receive one log out of every five, change the value to 5.	-1
Max length in bytes	The maximum length of log messages, in bytes.	32,766
Character encoding	The character encoding for this data input. This field is read-only.	UTF-8
Sub sample drop ratio	The ratio of logs to discard. This setting determines the number of logs to group together, out of which one is dropped. It reduces the number of fetched logs. The default value is -1: no logs are discarded. For example, if you want one log out of every five to be dropped, change the value to 5.	-1
Rate limit	The maximum number of events per second that this data input processes.	-1
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Drop if queue is full	Option to discard logs if many processes are waiting in the queue to access the MID Server.	Clear

Microsoft Azure Event Hubs integration configuration fields

Description of the fields on the Microsoft Azure Event Hubs integration configuration forms for Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For the Microsoft Azure Event Hubs integration setup procedure, see Set up a Microsoft Azure Event Hubs integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Event Hubs integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Service Instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Source	The source of the log data that the integration pulls to your ServiceNow instance: PubSub. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID Server name	(Only when **Execute on** is set to Specific MID Server.) The MID Server to which the logs are streamed. **Note:** • You can only select a MID Server with log ingestion capability that supports basic authentication. MID Servers that support mTLS are not listed. • If log ingestion is not enabled on the selected MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. This field is required.
MID Server cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.

Set data retrieval method

Event Hubs namespace	The scoping container for the event hub or hubs. This field is required.
Event Hub name	The event hub from which to fetch log data. This field is required.
Consumer Group name	The Consumer Group to use. This field is required. Enter an existing Consumer group or create one in the event hub. The default value is $Default.
Event Hub Credentials	The Event Hub credentials to use. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set data retrieval method (continued)

	Select existing Azure Shared Access Signature (SAS) credentials or create the Azure SAS credentials to use. If you create new Azure SAS credentials, you must provide the SAS policy name and key.

Advanced settings

Field	Description
Restore point rate	The number of logs the integration can read before a restore point is saved.
Processor threads count	The total number of processor threads.
Sub sample drop ratio	The ratio of logs to discard. This setting determines the number of logs to group together, out of which one is dropped. It reduces the number of fetched logs. The default value is -1: no logs are discarded. For example, if you want one log out of every five to be dropped, change the value to 5.
Sub sample receive ratio	The ratio of logs to receive. This setting determines the number of logs to group together, out of which one is received. It reduces the number of logs you receive. The default value is -1, meaning no logs are received. For example, if you want to receive one log out of every five, change the value to 5.
Operation timeout (seconds)	The number of seconds to wait before timing out event hubs operations.
Receive idle timeout (seconds)	The number of seconds to wait before timing out receive operations.
Max length in bytes	The maximum length, in bytes, of log messages. The default value is 32766.
Prefetch count	The number of logs received in advance of event hubs operations.
Max batch size	The maximum number of events passed to a single process call. The default value is 100.
Default timezone	The default time zone of logs. The system uses this default when the log does not specify a time zone. The default value is GMT.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Character encoding	The character encoding for this integration. The default value is UTF-8.
Drop if queue is full	Option to select discarding logs if there is a load on the MID Server.

Microsoft Azure Log Analytics integration configuration fields

Description of the fields on the Microsoft Azure Log Analytics integration configuration forms for Health Log Analytics.

For the Microsoft Azure Log Analytics integration setup procedure, see Set up a Microsoft Azure Log Analytics integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Azure Log Analytics integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID server name	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Microsoft Azure Log Analytics is pulled. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of integrations streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
MID MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Data source	The source of the log data that the integration pulls to your ServiceNow instance: Azure Log Analytics. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set data retrieval method

Field	Description
Redirect url	The redirect URL of the access log application. This field is required. The URL refers to the Microsoft Azure redirect_uri authorization property. For more information, see the Authorization code URL (GET request) section in the Microsoft Azure documentation.
Authentication method	The credentials used to access Microsoft Azure resources. This field is required. If no credentials exist, select **Create new Azure Service Principal credential**. Create a credential by filling in the fields and then selecting **Submit**. For information on the **Tenant ID**, **Client ID**, and **Secret key** fields, refer to the Microsoft Azure documentation. When you have created a credential, you can select it from the drop-down list.
Workspace ID	The Customer ID used to call the Microsoft Azure Log Analytics REST API. This field is required.
Table name	The name of the table in Microsoft Azure Log Analytics where the data input fetches the log data. For more information, see the View table information section in the Microsoft Azure documentation. This field is required.
Event time property name	The Microsoft Azure Log Analytics field in which to detect the event time. This field is required.

Advanced settings

Field	Description
Event Processor workers	The number of concurrent event processing workers, where each worker processes a batch of events independently.
Workers queue size	The queue size of the Event Processor workers.
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched events.
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Max documents per query	The highest number of rows retrieved in each query.
Columns to select	Comma-separated list of column names to return. **Note:** This field is ignored when you provide a custom query.
Character encoding	The character encoding for this data input.
Sleep interval (seconds)	The interval, in seconds, to wait before querying again after a query has returned no events.
Polling interval	The interval, in seconds, to wait before polling for new events.
Drop if queue is full	Option to discard logs if there is a load on the MID Server.

Log Query	Option to define your own Log Analytics query. This field overrides the values configured in the other query settings fields.
	
	Note: If this field is empty, Health Log Analytics generates the query using the values set in the other fields.
	
	For your custom query, use the following JSON format:
	
	`{"query": "query | where TimeGenerated > %s | take 500"}`
	
	For example:
	
	`{ "query":"ContainerLog | where LogEntry contains 'cartservice' | where TimeGenerated > %s | take 500", "workspaces": ["defaultworkspace-3ab145ff-f9cd-433f-8533-d1b1ee24aee6-eus"], "project": ["TimeGenerated", "LogEntry", "LogEntrySource"] }`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MID Server integration configuration fields

Description of the fields on the MID Server integration configuration forms for Health Log Analytics.

For the MID Server integration setup procedure, see Set up a MID Server integration for Health Log Analytics.

Provide details

Field	Description
MID server name	The name of the MID Server from which the integration collects log messages and streams them to the instance. **Note:** You can only create one MID Server integration per available MID Server. If an existing MID Server integration or data input uses the same MID Server, you can't create an additional one for this MID Server. This field is required.
Integration Name	The name of the integration. When you select the MID Server name, this field is auto-populated with the selected MID Server. This field is read-only.
Service instance	The service instance to which to bind the log data. This field is required. By default, it is set to ServiceNow MID Server. **Note:** This field is only available for MID Server integrations created in version 36.0.17 or later versions.
Description	Option to add a brief description of the integration to help identify it.

Advanced settings

Field	Description
Default timezone	The time zone of events that the system will use if a log does not specify the time zone. By default, the system uses GMT in such cases, but you can specify a different time zone.
Sub sample drop ratio	The ratio of events to drop. The default value is -1.
Sub sample receive ratio	The ratio of events to receive. The default value is -1.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Max length in bytes	The maximum length of log messages, in bytes. The default value is 32,766.
Character encoding	The character encoding for this data input. The default value is UTF-8.
Drop if queue is full	Option to discard logs if many processes are waiting in the queue to access the MID Server.

REST API integration configuration fields

Description of the fields on the REST API integration configuration forms for Health Log Analytics.

For the REST API integration setup procedure, see Set up a REST API integration for Health Log Analytics.

Note: REST API integrations support only UTF-8 encoding for incoming data.

Provide details tab

Field	Description
Integration Name	Unique name of this integration. For example: My REST API integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID web server	The MID Web Server to use. This field is required.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, Create a service instance and add CIs (Configuration Items) to it. Set the status of the new service instance to Operational.
Transport	The protocol used for streaming log messages to your ServiceNow instance: REST. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.
MID server name	The MID Server to which the log data is pushed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details tab (continued)

Field	Description
	This field is auto-populated when the MID Web Server is chosen. It is read-only.
Port	The port for the MID Server. This field is read-only. It is auto-populated when the MID Web Server is chosen.

ServiceNow System Logs Retriever integration configuration fields

Description of the fields on the ServiceNow System Logs Retriever integration configuration form for Health Log Analytics.

For the ServiceNow System Logs Retriever integration setup procedure, see Set up a ServiceNow System Logs Retriever integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My System Logs integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required. The default service instance for this integration is ServiceNow System Logs.
Description	Option to add a brief description of the integration to help identify it.

Splunk Poller integration configuration fields

Description of the fields on the Splunk Poller integration configuration forms for Health Log Analytics.

For the Splunk Poller integration setup procedure, see Set up a Splunk Poller integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID server name	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Splunk is pulled. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Data source	The source of the log data that the integration pulls to your ServiceNow instance: Splunk. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Set data retrieval method

Field	Description
REST API Server URL	The URL used to access the Splunk REST API. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set data retrieval method (continued)

Field	Description
	Note: • By default, the Splunk REST API endpoint is available on port 8089. Therefore, the default endpoint for the Splunk REST API is: `http://<splunk-server>:8089`. • The REST API endpoint is different from the front-end endpoint.
Authentication method	The credential alias to be used. This field is required. The Splunk integration uses credential aliases instead of direct references to credentials for authentication. The credential aliases are listed by type: aliases that contain basic auth credentials, and those that contain token auth credentials. If an alias holds both credential types, it appears in both categories. You can select **Manage credential aliases** to manage your credential aliases and create new ones in the **Connection & Credential Aliases** list.
Query	The query Splunk uses to search your data. This field is required. For example, the query `sourcetype="adc_access_log"` instructs the Splunk Poller integration to retrieve all logs with the source type adc_access_log.

Advanced settings

Field	Description
Max documents per query	The maximum number of documents retrieved each time log data is fetched from Splunk. Default: 10,000.
Splunk Request Timeout (seconds)	The maximum time, in seconds, allowed for data retrieval before the request times out.

Splunk TCP integration configuration fields

Description of the fields on the Splunk TCP integration configuration forms for Health Log Analytics.

For the Splunk TCP integration setup procedure, see Set up a Splunk TCP integration for Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Splunk TCP integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID server name	MID Server to which log data from Splunk is pulled. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.
Port	The port for the MID Server. This field is required. Make sure that your organization's security team opens the selected port on the MID Server.
Description	Option to add a brief description of the integration to help identify it.
Transport	The protocol used for streaming log messages to your ServiceNow instance: TCP. This field is read-only.
Use Cooked Data	Option to ingest log data from Splunk in the preprocessed ("cooked") format that Splunk uses on the forwarder. Ingesting data into HLA in this format ensures that each log line retains the relevant contextual information that Splunk embeds into it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings

Field	Description
Use SSL/TLS	Option to use SSL/TLS, for enhanced data security and protection. **Note:** SSL/TLS must be enabled if you want to send logs in a compressed format.
Lookup hostnames	Option to perform DNS lookup to resolve IPs to hostnames. The default value is false.
Use Forwarder TimeZone	Option to pass information about the time zone in which the forwarder is located. The MID Server uses this information to adjust for the time zone from which the logs arrive. This option is displayed when **Use Cooked Data** is selected. It is relevant when using Splunk Universal Forwarders.
Enable Compression	Option to send logs in compressed format. Sending logs in a compressed format minimizes the size of the data being transferred, which is important when dealing with large volumes of log data. This option is displayed when **Use Cooked Data** and **Use SSL/TLS** are selected. It is relevant when using Splunk Universal Forwarders.
Sub sample drop ratio	The ratio of logs to drop. The default value is -1: no logs are dropped. For example: If you want one out of every five logs to be dropped, change the value to 5.
Sub sample receive ratio	The ratio of logs to receive. The default value is -1: no logs are received. For example: If you want one out of every five logs to be received, change the value to 5.
Max length in bytes	The maximum length of log messages in bytes. The default value is 32766.
Character encoding	The character encoding for this data input. Default is UTF-8.
Boss thread count	The number of threads that manage connections.
Worker thread count	The number of threads that handle incoming data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Read timeout seconds	The timeout in seconds since the last read. When the timeout expires, the system closes the channel.
Default timezone	The time zone of events that the system will use if a log does not specify the time zone. By default, the system uses GMT in such cases, but you can specify a different time zone.
Drop if queue is full	Option to discard logs if there is a load on the MID Server.

Splunk UDP integration configuration fields

Description of the fields on the Splunk UDP integration configuration forms for Health Log Analytics.

For the Splunk UDP integration setup procedure, see Set up a Splunk UDP integration for Health Log Analytics.

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Splunk UDP integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID server name	MID Server to which log data from Splunk is pulled. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
Port	The port for the MID Server. This field is required. **Note:** Make sure that your organization's security team opens the selected port on the MID Server.
Description	Option to add a brief description of the integration to help identify it.
Transport	The protocol used for streaming log messages to your ServiceNow instance: UDP. This field is read-only.

Advanced settings

Field	Description
Lookup hostnames	Option to perform DNS lookup to resolve IPs to hostnames. The default value is false.
Sub sample receive ratio	The ratio of logs to receive. The default value is -1: no logs are received. For example: If you want one out of every five logs to be received, change the value to 5.
Character encoding	The character encoding for this data input. The default value is UTF-8. This field is read-only.
Drop if queue is full	Option to discard logs if there is a load on the MID Server.
Sub sample drop ratio	The ratio of logs to drop. The default value is -1: no logs are dropped. For example: If you want one out of every five logs to be dropped, change the value to 5.
Max length in bytes	The maximum length of log messages in bytes. The default value is 32766.
Default timezone	The time zone of events that the system will use if a log does not specify the time zone. By default, the system uses GMT in such cases, but you can specify a different time zone.

TCP integration configuration fields

Description of the fields on the TCP integration configuration forms for Health Log Analytics.

For the TCP integration setup procedure, see Set up a TCP integration for Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details tab

Field	Description
Integration Name	Unique name of this integration. For example: My TCP integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID server name	The MID Server to which the logs are streamed. This field is required. • You can select only MID Servers with log ingestion capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties.
Port	The port for the MID Server. This field is required. Select a unique port from the array. The placeholder shows the range of ports from which to choose. Make sure that your organization's security team opens the selected port.
Service instance	The service instance to which to bind the log data. This field is required.
Transport	The protocol used for streaming log messages to your ServiceNow instance: TCP. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Advanced settings

Field	Description
Use SSL/TLS	Option to use SSL/TLS.
Lookup hostnames	Option to perform DNS lookup to resolve IPs to hostnames.
Boss thread count	The number of threads that manage connections.
Worker thread count	The number of threads that handle incoming data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings (continued)

Field	Description
Read timeout seconds	The time zone of events that the system will use if a log does not specify the time zone. By default, the system uses GMT in such cases, but you can specify a different time zone.
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.
Sub sample drop ratio	The ratio of events to drop.
Sub sample receive ratio	The ratio of events to receive.
Max length in bytes	The maximum length of log messages in bytes.
Character encoding	The character encoding for this data input.
Drop if queue is full	Option to discard logs if there is a load on the MID Server.
Line breaker delimiters	The line break character separating the raw log lines. Splitting values must be separated by a comma followed by a space: ", ". For example: "\r, \n, , splitHere, #".

UDP integration configuration fields

Description of the fields on the UDP integration configuration forms for Health Log Analytics.

For the UDP integration setup procedure, see Set up a UDP integration for Health Log Analytics.

Provide details tab

Field	Description
Integration Name	Unique name of this integration. For example: My UDP integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID server name	The MID Server to which the logs are streamed. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details tab (continued)

Field	Description
	Note: • You can select only MID Servers with log ingestion capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties.
Service instance	The service instance (formerly the application service) to which to bind the log data. This field is required.
Port	The port for the MID Server. This field is required. The port must be within the range 6000-6200. Make sure that your organization's security team opens the selected port on the MID Server.
Transport	The protocol used for streaming log messages to your ServiceNow instance: UDP. This field is read-only.
Description	Option to add a brief description of the integration to help identify it.

Advanced

Field	Description
Default timezone	The time zone of events that the system will use if a log does not specify the time zone. By default, the system uses GMT in such cases, but you can specify a different time zone.
Sub sample drop ratio	The ratio of events to drop. The default value is -1.
Sub sample receive ratio	The ratio of events to receive. The default value is -1.
Rate limit	Limitation on the number of events per second that the data input processes. The default value is -1.
Max length in bytes	The maximum length of log messages, in bytes. The default value is 32,766.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced (continued)

Field	Description
Character encoding	The character encoding for this data input. The default value is UTF-8. This field is read-only.
Drop if queue is full	Option to discard logs if many processes are waiting in the queue to access the MID Server.

Vector Agent integration configuration fields

Description of the fields on the Vector Agent integration configuration forms for Health Log Analytics.

For the Vector Agent integration setup procedure, see Set up a Vector Agent integration for Health Log Analytics .

Provide details

Field	Description
Integration Name	Unique name of this integration. For example: My Vector Agent integration. This field is required. **Note:** When you fill in this field, the generic name displayed on the form adjusts automatically to match the name you entered.
MID server name	The MID Server to which the Vector Agent logs are streamed. This field is required. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of integrations streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. • If the MID Server is down or the connection is lost, the Vector pipeline may get blocked.
Port	The port on the MID Server. This field is required. Choose a port within the range 6000-6200 from the array. The port must not be occupied

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provide details (continued)

Field	Description
	by another process. Make sure that your organization's security team opens the selected port on the MID Server.
Transport	The protocol used for streaming log messages to your ServiceNow instance: TCP.
Description	Option to add a brief description of the integration to help identify it.

Data input configuration fields

This section provides descriptions of the fields on the Health Log Analytics data input configuration forms.

Amazon CloudWatch data input configuration fields

Description of the fields on the Amazon CloudWatch data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Amazon CloudWatch is pulled. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Service instance	The service instance to which to bind the log data. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational. This field is required.

The following fields show read-only information:

Field	Description
Status	Status of the data input.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Transport	Protocol used to stream the log data. This data input uses Amazon CloudWatch to stream log data to your instance.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Query settings tab

Field	Description	Example
From	The date and time to start reading the data. Data older than this date and time is not read. **Note:** Setting this value to a past date might require the system to read large amounts of data, causing congestion. This field is required.	Now -1 week
Group Name(s)	The log groups to search. If multiple log groups must be searched, specify the groups in a comma-separated list. To fetch log data from all groups, use an asterisk (*) as a wildcard character. This field is required.	hla-cw-loggroup1,hla-cw-loggroup2
Prefix	Name prefix for the Amazon CloudWatch log streams to read from. The data input reads only from log streams with this name prefix.	hla-cw

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query settings tab (continued)

Field	Description	Example
	Note: Only a single log stream prefix per data item is supported. For multiple prefixes, create multiple data inputs.	
Filter pattern	Pattern by which to filter incoming events. Various types of filter patterns are supported. For example: • A pattern for fetching log events that contain a single term. • A pattern for fetching log events that contain multiple terms. • A pattern for fetching log events that include a term and exclude another. **Note:** Filter patterns are case sensitive. For more information, see https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/FilterAndPatternSyntax.html	• Pattern for fetching log events that contain a single term, "STATUS_CODE": "STATUS_CODE" • Pattern for fetching log events that contain multiple terms, “STATUS_CODE” and “200”: “STATUS_CODE 200” • Pattern for fetching log events that include a term, “STATUS_CODE,” and exclude another term, “200”: “STATUS_CODE” – “200”

Transport tab

Field	Description
AWS credentials	Field that refers to the AWS Credentials list (aws_credentials.list). The list contains the AWS access and secret access keys.
AWS region	The AWS region where the Amazon CloudWatch cluster runs, for example, us-west-1. For a list of AWS regions, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration

Advanced configuration form

Field	Description	Default value
Connection timeout	The number of milliseconds to wait before timing out the AWS connection attempt.	1000
Batch size	The maximum number of logs retrieved per query.	2500
Socket timeout	The number of milliseconds to wait before timing out a data transfer over an established connection.	10000
Default timezone	The default timezone if the log date and time doesn't include timezone information.	GMT
Sub sample drop ratio	The number of logs to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched logs.	-1
Sub sample receive ratio	The number of logs to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received logs.	-1
Max length in bytes	The maximum length of log messages, in bytes.	32766
Sleep interval	The interval, in seconds, to wait before querying again after a query has returned no logs.	60
Polling interval	The interval, in seconds, to wait before polling for new logs.	0
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	False

Amazon S3 data input configuration fields

Description of the fields on the Amazon S3 data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
MID	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Amazon S3 is pulled. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. This field is required.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Service instance	The service instance to which to bind the log data. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational. This field is required.

The following fields show read-only information:

Field	Description
Status	Status of the data input.
Transport	Protocol used to stream the log data. This data input uses Amazon S3 to stream log data to your instance.
Sources count	The number of log sources this data input has created.
Disabled since	Time when the data input stopped or failed.
Last log time	Time when the last log streamed in the data input.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Query settings tab

Field	Description	Example
From/To	Starting and ending dates and time for reading the data. • From: Data older than this date and time is not read. **Note:** Setting this value to a past date might require the system to read large amounts of data, causing congestion. • To: Data newer than this date and time is not read. For live data, set this date far into the future. This field is required.	From: Now -1 week To: 2300-01-01 15:59:59
Bucket name	The Amazon S3 bucket to stream. This field is required.	my-s3-bucket
Prefix	Prefix for the objects to be fetched. The data input fetches only objects that begin with this prefix. This field is typically used to fetch only specific folders in a bucket.	/only/this/folder/
Include	File or object to include when the system performs pattern matching.	.*include-me.*
Exclude	File or object to exclude when the system performs pattern matching.	.*skip-me.*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Transport tab

Field	Description
AWS credentials	Field that refers to the AWS Credentials list (aws_credentials.list). The list contains the AWS access and secret access keys.
AWS region	The AWS region where the Amazon S3 bucket is located, for example, us-west-1. For a list of AWS regions, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions This field is required.

Advanced configuration

Advanced configuration form

Field	Description	Default value
Anonymous	Option to force accessing the Amazon S3 bucket without credentials, while ignoring default credentials if they exist.	False
Traverse	Option to access each directory in the bucket.	False
Object sorted by Last Modified Time	Option to fetch objects if the objects in the bucket are sorted by time.	False
Connection timeout	The number of milliseconds to wait before timing out the attempt to establish a connection to AWS.	100
Threadpool size	The number of concurrent threads that are downloading files from the bucket.	1
Buffer size	Size of the download buffer, in bytes.	100MB
Socket timeout	The number of milliseconds to wait before timing out data transfer over an established connection.	10000
Max retry for file	The maximum number of times to retry file processing in case of failure.	10

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration form (continued)

Field	Description	Default value
Default timezone	The default timezone if the log doesn't include timezone information.	UTC
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched events.	-1
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.	-1
Rate limit	The maximum number of events per second that this data input processes.	-1
Max length in bytes	The maximum length, in bytes, of events.	32766
Character encoding	The character encoding for this data input.	UTF-8
Sleep interval	The interval, in seconds, to wait before querying again after a query has returned no events.	60
Polling interval	The interval, in seconds, to wait before polling for new events.	0
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	False

Apache Kafka data input configuration fields

Description of the fields on the Apache Kafka data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
MID	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Apache Kafka is pulled. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. This field is required.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Service instance	The service instance to which to bind the log data. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational. This field is required.

The following fields show read-only information:

Field	Description
Status	Status of the data input.
Transport	Protocol used to stream the log data. This data input uses Apache Kafka to stream log data to your instance.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Query Settings tab

Field	Description	Example
From	Starting date and time for reading the data. Data older than this date and time is not read. **Note:** Setting this value to a past date might require the system to read large amounts of data, causing congestion. This field is required.	Now -1 week

Transport tab

Field	Description	Example
Kafka node names	A comma-separated list in the format HOST:PORT,HOST:PORT. The list does not have to include all the Apache Kafka Cluster servers.	123.4.5.6:9092,123.3.4.5:9093
Topics	A comma-separated list of topics to which the data input must subscribe. This field is required.	FirstTopic,SecondTopic,ThirdTopic
Kafka credentials	Reference to the Apache Kafka credentials. You can display the Kafka SSL credentials form by selecting the Search icon ( ) and opening a record. The form enables you to choose the security protocol used for authenticating with Apache Kafka from the following options:	None

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Transport tab (continued)

Field	Description	Example
	• **SSL** - SSL channel. • **SASL_SSL** - SASL authenticated, SSL channel. • **SASL_PLAINTEXT** - SASL authenticated, non-encrypted channel. For a description of the fields on the Kafka SSL credentials form, see Kafka SSL credentials fields.	
Group Id	The name of the Apache Kafka Consumer Group.	logs

Advanced configuration

Advanced configuration form

Field	Description	Default value
Timeout	The time, in milliseconds, spent waiting in the poll if data is not available in the topics.	500
Node discovery timeout	The time, in milliseconds, before node discovery times out.	30
Default timezone	The default timezone if the log doesn't include timezone information.	GMT
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched events.	-1
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.	-1
Max length in bytes	The maximum length, in bytes, of events.	32766
Character encoding	The character encoding for this data input.	UTF-8
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	False

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kafka SSL credentials fields

Description of the fields on the Kafka SSL credentials form.

SSL security protocol

Field	Description
SSL keystore	The content of an actual server.keystore.jks file, encoded as a Base64 string.
SSL keystore type	Format of the SSL keystore file. Default: Java KeyStore (JKS).
SSL keystore password	Password to the keystore, identical to the value of the ssl.keystore.password property in Apache Kafka.
SSL key password	Password, identical to the ssl.key.password property value in Apache Kafka.
Use SSL truststore	Option to use SSL truststore. When checked, an SSL truststore is mandatory.
SSL truststore	The content of an actual client.truststore.jks file, encoded as a Base64 string. This field appears only when **Use SSL truststore** is checked.
SSL truststore type	Format of the truststore file. Default: Java KeyStore (JKS). This field appears only when **Use SSL truststore** is checked.
SSL truststore password	Password to the truststore, identical to the value of the ssl.truststore.password property in Apache Kafka. This field appears only when **Use SSL truststore** is checked.
SSL provider	Name of the security provider used for SSL connections.
Disable hostname verification	Option to disable verification of the hostname.
Additional Kafka consumer properties	Semicolon-separated string of additional Kafka consumer properties. For example: ssl.protocol=TLSv1.2;ssl.enabled.protocols=TLSv1.2,TLSv1.1

SASL_SSL security protocol

Field	Description
SASL Mechanism	The SASL authentication mechanism to be used by the Kafka client. The options are: • PLAIN • SCRAM-SHA-256 • SCRAM-SHA-512 • OAUTHBEARER Default: PLAIN.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SASL_SSL security protocol (continued)

Field	Description
Token endpoint URL	The token provider URL where credentials are sent to fetch access tokens. This value is defined in the client properties file. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
Client ID	Unique, public identifier for the application. The authorization server uses it to recognize the calling application. This value is defined in the client properties file. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
Client Secret	Private password known only to the application and the authorization server. Used to prove the application's identity. This value is defined in the client properties file. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
Scope	Defines the level of access or permissions requested by the application. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
OAUTH extensions	Custom metadata or additional parameters for the token request or validation process, in key-value pairs. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
Use SSL truststore	Option to use SSL truststore. When checked, an SSL truststore is mandatory.
SSL truststore	The content of an actual client.truststore.jks file, encoded as a Base64 string. This field appears only when **Use SSL truststore** is checked.
SSL truststore type	Format of the SSL truststore file. Default: Java KeyStore (JKS). This field appears only when **Use SSL truststore** is checked.
SSL truststore password	Password to the truststore, identical to the value of the ssl.truststore.password property

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SASL_SSL security protocol (continued)

Field	Description
	in Apache Kafka. This field appears only when **Use SSL truststore** is checked.
SSL provider	Name of the security provider used for SSL connections.
Disable hostname verification	Option to disable verification of the hostname.
User name	The user name, if any, associated with this credential. **Note:** The specified user name is used for SASL authentication using JAAS. The following JAAS configuration is passed to the Kafka Server: `org.apache.kafka.common.security.plain.P required username=<user_name> password=<password>` This field does not appear when the **SASL Mechanism** is **OAUTHBEARER**.
Password	The password associated with this credential. **Note:** • The specified password is used for SASL authentication using JAAS. The following JAAS configuration is passed to the Kafka Server: `org.apache.kafka.common.security.plain. required username=<user_name> password=<password>` • The password you specify is invisible. It is stored in the database in encrypted form. This field does not appear when the **SASL Mechanism** is **OAUTHBEARER**.
SASL JAAS config template	The JAAS configuration to be used by the Kafka client, with the placeholders $[username] for the username, and $[password] for the password.
Additional Kafka consumer properties	Semicolon-separated string of additional Kafka consumer properties. For example: ssl.protocol=TLSv1.2;ssl.enabled.protocols=TLSv1.2,TLSv1.1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SASL_PLAINTEXT security protocol

SASL mechanism	The SASL authentication mechanism to be used by the Kafka client.The options are: • PLAIN • SCRAM-SHA-256 • SCRAM-SHA-512 • OAUTHBEARER Default: PLAIN.
Token endpoint URL	The token provider URL where credentials are sent to fetch access tokens. This value is defined in the client properties file. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
Client ID	Unique, public identifier for the application. The authorization server uses it to recognize the calling application. This value is defined in the client properties file. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
Client Secret	Private password known only to the application and the authorization server. Used to prove the application's identity. This value is defined in the client properties file. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
Scope	Defines the level of access or permissions requested by the application. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
OAUTH extensions	Custom metadata or additional parameters for the token request or validation process, in key-value pairs. This field appears only when the **SASL Mechanism** is **OAUTHBEARER**.
User name	The user name, if any, associated with this credential.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SASL_PLAINTEXT security protocol (continued)

	Note: The specified user name is used for SASL authentication using JAAS. The following JAAS configuration is passed to the Kafka Server: `org.apache.kafka.common.security.plain.P required username=<user_name> password=<password>` This field does not appear when the **SASL Mechanism** is **OAUTHBEARER**.
Password	The password associated with this credential. **Note:** • The specified password is used for SASL authentication using JAAS. The following JAAS configuration is passed to the Kafka Server: `org.apache.kafka.common.security.plain. required username=<user_name> password=<password>` • The password you specify is invisible. It is stored in the database in encrypted form. This field does not appear when the **SASL Mechanism** is **OAUTHBEARER**.
Additional Kafka consumer properties	Semicolon-separated string of additional Kafka consumer properties. For example: ssl.protocol=TLSv1.2;ssl.enabled.protocols=TLSv1.2,TLSv1.1

Cribl data input configuration fields

Description of the fields on the Cribl data input configuration form.

Getting Started

Field	Description
Data input name	Name of the new data input. This field is required.
Description	Description of the data input.
MID server	The MID Server to which the Cribl logs are streamed. This field is required.
Port	The port on the MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Getting Started (continued)

Field	Description
	Choose a port within the suggested range from the array. The port must not be occupied by another process. Make sure that your organization's security team opens the selected port. This field is required.

Config

Field	Description
Cribl instance URL	The URL of the Cribl instance. HLA uses this URL to establish the connection with the Cribl instance.
Cribl API token	The Cribl API token HLA uses to authenticate with Cribl and make API calls.

Process

Field	Description
Cribl route	The Cribl route that directs log data to a specific destination. This field is auto-populated with a drop-down list of all available routes, each configured with a default destination. You must select the desired route from this list. This field is required.
Cribl destination	The created Cribl destination. The Cribl destination points to the MID Server host and port configured in the data input. This field is read-only.

Edge Delta data input configuration fields

Description of the fields on the Edge Delta data input configuration form.

Edge Delta data input configuration form

Field	Description
Data input name	Name of the new data input. This field is required.
Description	Description of the data input.
Transport	The Edge Delta destination used for streaming log messages to the ServiceNow instance. The options are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Edge Delta data input configuration form (continued)

Field	Description
	• TCP - Frequently used for handling large volumes of log data. If you select TCP, you must provide the MID Server and port. When choosing the TCP transport option, create a TCP destination in Edge Delta to send log data to the instance from any source. Use the same MID Server host and port you have configured for the Edge Delta data input. **Note:** For TCP, TLS/SSL is not supported. • HTTP: Frequently used for sending smaller logs via RESTful APIs, HTTP, or HTTPS for secure transmission. **Note:** For the HTTP protocol, authentication is used. If you select HTTP, you must provide the MID Web Server that is running on the MID Server. The system then populates the MID Server and port fields automatically, based on the provided MID Web Server. Depending on the MID Web Server configuration, you can send logs using either HTTP or HTTPS. If the **Secure Connection** check box on the MID Web Server extension form is clear, use HTTP. If it is selected, use HTTPS. For more information, see Event Management MID Web Server extension form. In addition, if the **Authentication Type** on the MID Web Server extension form is Basic, provide Basic authentication details in the Authorization header. If the **Authentication Type** is API Key, provide the key as “Key Api_Key". When choosing the HTTP transport option, create an HTTP destination in Edge Delta and enter the Endpoint URL in the following format: http://{MID_SERVER_IP}:{MID_WEB_SERVER_PORT}/api/mid/hla/raw. In Advanced Settings, create two headers to be attached to the HTTP request: ○ An Authorization header with the MID Web Server API key value or basic credentials as MID Web Server configured. ○ A Content-Type header named with the application/json value.
MID web server	The MID Web Server to which the log data is streamed.
MID server	The MID Server to which the Edge Delta logs are streamed. This field is required. If you selected HTTP in the **Transport** field, this field is populated automatically based on the selected MID Web Server.
Port	The port on the MID Server used for listening to logs. Choose a port within the suggested range from the array. In the case of TCP transport, the port must not be occupied by another process. Make sure that your organization’s security team opens the selected port.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Edge Delta data input configuration form (continued)

Field	Description
	This field is required. If you selected HTTP in the **Transport** field, this field is populated automatically based on the selected MID Web Server.
Service instance	The service instance to which to bind the log data. This field is required. Choosing the service instance enables HLA to correlate CIs and service instances in the CMDB for root cause analysis. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.

Elasticsearch data input configuration fields

Description of the fields on the Elasticsearch data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Execute on	Option to determine whether to use a specific MID Server or a MID Server cluster. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.
MID	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Elasticsearch indices is pulled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. This field is required.
MID Server Cluster	(Only when the **Execute on** field is set to Specific MID Server Cluster) The MID Server cluster to which the log data is pulled. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. This feature is supported in the Health Log Analytics application, Version 26.0.17 - February 2023 and later, available from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input form, the MID Server Clusters list displays only failover clusters. • The MID Server cluster must include only MID Servers that support basic authentication. mTLS is not supported for log ingestion. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • If Elasticsearch uses client certificate or CA certificate authentication, all MID Servers in the cluster must have the appropriate certificates. • The default maximum number of data inputs streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs running on it, even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster. This field is required.
Service instance	The service instance to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.

The following fields show read-only information:

Field	Description
Status	Status of the data input.
Transport	Protocol used to stream the log data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	This data input uses Elastic to stream log data to your instance.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.

Transport tab

Field	Description
Server URL	URL used to access the cluster. This field is required.
Max connections per route	Maximum number of connections to be opened per node. Default: 2.
Max scroll slices	The number of shards configured for the relevant index in Elasticsearch. This number tells Elastic how many parallel queries to execute in each polling request.
Proxy host	Host name of the HTTP proxy through which requests are sent.
Proxy port	Port of the HTTP proxy through which requests are sent.
Authentication method	The authentication method used to authenticate the data input to Elasticsearch. The options are: Basic auth, apiKey, or client certificate. **Note:** When you select the required authentication method, the corresponding credentials fields display on the form.
Basic auth credentials	User name and password used to connect to the Elasticsearch search engine. **Note:** Fill in either this field or the **AWS credentials** field.
AWS credentials	AWS credentials to use to connect to the AWS-hosted Elasticsearch search engine. **Note:** Fill in either this field or the **Basic auth credentials** field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Transport tab (continued)

Field	Description
AWS region	AWS region where the Elasticsearch cluster runs.
API key credentials	The API key used to connect to the Elasticsearch search engine.
Client certificate	The client certificate used to connect to the Elasticsearch search engine.
Use MID certificate policy check	Option to enable the MID certificate policy check. Select this option if you want to ship your logs encrypted using SSL TLS. Then navigate to **All** > **MID Server** > **MID Security Policy** and add the MID certificate policy check to the table. For more information, see MID Server certificate check policies.

Query Settings tab

Field	Description	Example
From/To	From and to dates and time for reading the data. • From: Don't read data older than this date. **Note:** Setting this value to a past date might require the system to read large amounts of data, causing congestion. • To: Don't read data newer than this date. For live data, set this date far into the future.	From: 1970-01-01 15:59:59 To: 2300-01-01 15:59:59
Use cross-cluster search	Option to search for data across Elasticsearch clusters. When this check box is selected, the **Clusters to search** field displays. **Note:** Your settings in the **Use minimal privileges** check box and the **Delay in reading current timestamp (seconds)** field on the **Advanced configuration** form affect how data is collected across multiple clusters.	
Clusters to search	The Elasticsearch clusters to search.	east,west,south

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query Settings tab (continued)

Field	Description	Example
	This field displays only when the **Use cross-cluster search** check box is selected. Do one of the following: • Leave this field empty or enter "*" to search all the remote clusters defined in Elasticsearch. • Specify the clusters to search in a comma-separated list. **Note:** To also search the local cluster, add a comma at the beginning or end, or add two commas in succession to the list. For example: "east,,west" or ",east,west" or "*,"	
Index prefix	Prefix for the Elasticsearch indices to read from. The data input reads only from indices with this prefix. This field is required.	only-read-these-indices-*
Use minimal privileges	Option to read log data directly from the Elasticsearch indices with the configured prefix. • When selected, the data input reads the log data directly from the Elasticsearch indices with the configured prefix. To perform this task, it needs only read privileges. **Note:** When this check box is selected and you're using cross-cluster search, data is collected from all the clusters simultaneously. • When clear, the data input fetches all indices with the prefix, filters them, and reads the log data from the filtered	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query Settings tab (continued)

Field	Description	Example
	indices. Performing this task requires additional privileges. **Note:** Leaving this check box clear when using cross-cluster search affects how data is collected from the clusters. For more information, see the Enabling and Using Cross-Cluster Search for Elasticsearch Data Inputs in Health Log Analytics [KB1556079] article in the Now Support Knowledge Base. For additional information about streaming logs using the Elasticsearch data input, see the Stream logs using Elasticsearch data input - Advanced guide [KB1080162] article in the Now Support Knowledge Base.	
Document timestamp field	Timestamp field in documents stored in the read indices. This field is required.	
Timestamp field format	Format of the timestamp field in the documents. If no format is specified, the default Unix epoch time format is used, in milliseconds. For example: 1684168407 (May 15, 2023 4:33:27 PM)	yyyy-MM-dd'T'HH:mm:ss.SSSSSSS'Z'
Term filters	JSON map of the terms to filter. **Note:** Avoid using the term query for text fields. If the target field is mapped as both text and keyword, reference the keyword by using fieldname.keyword.	{"severity": ["error", "warning"]}
Max documents per query	Maximum number of documents fetched in a single query.	
Sliced-scrolling tiebreaker	Value used to slice the data. Each slice is scrolled in parallel. Default: _id	
Search-after tiebreaker	Unique value per document to use as tiebreaker when sorting log entries by timestamp.	
Use search-after API	Option to toggle between using sliced-scrolling and search-after APIs.	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query Settings tab (continued)

Field	Description	Example
	Note: Sliced-scrolling APIs are preferable when reading historical data, while search-after APIs are better for reading real-time data.	
Index time-suffix format	Format of the time suffix when using time-based index names, such as [logstash-]YYYY.MM.DD. When using aliases, leave this field empty.	uuuu.MM.dd

Advanced configuration

Advanced configuration form

Field	Description
Data reading timeout (milliseconds)	The duration of time, in milliseconds, before a request to the Elasticsearch cluster times out.
Index discovery interval (seconds)	The number of seconds between intermittent MID Server requests to the Elasticsearch cluster for new indices from which to read data.
Scroll context time (milliseconds)	The lifetime of the created scroll when using the scroll API to read data from Elasticsearch. For more information, see the Elasticsearch scroll API documentation.
Event processor workers	The maximum number of CPU cores used in parallel to process events fetched from Elasticsearch. A higher setting increases the data input throughput at the cost of higher CPU usage.
Worker queue size	The maximum number of batches to queue for processing. A higher setting increases throughput, at the cost of higher RAM usage.
Default timezone	The default timezone if the event date and time doesn't include timezone information.
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched events.
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.
Character encoding	The character encoding for this data input.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration form (continued)

Field	Description
Sleep interval (seconds)	The interval, in seconds, to wait before querying again after a query has returned no data.
Max length in bytes	The maximum length, in bytes, of log messages.
Delay in reading current timestamp (seconds)	The number of seconds before current time to query to include delayed data. The configured number of seconds is subtracted from the current time for reading the last timestamp. **Note:** If this value is 0 and data is collected from multiple clusters simultaneously, the query may not include data that was sent with a delay on one of the clusters,

GCP PubSub data input configuration fields

Description of the fields on the GCP PubSub data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Port	The port for the MID Server. Select a unique port from the array. The placeholder shows the range of ports from which to choose. Make sure that your organization's security team opens the selected port. This field is required.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID	(Only when the **Execute on** field is set to Specific MID Server) The MID Server to which the logs are streamed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • You can select only MID Servers with log ingestion capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. This field is required.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Service instance	The service instance to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.

The fields in the following table show read-only information.

Field	Description
Transport	Protocol used to send the log data: PubSub.
Sources count	The number of log sources this data input has created.
Status	Status of the data input.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.

Transport tab

Field	Description
Project ID	The project ID of the Google Cloud project. For example, my-project-id.
Topic Name	The Google Cloud Pub/Sub topic to which to subscribe Health Log Analytics. For example, my_topic.
GCP Credential Alias	The credential alias to be used. Specify one GCP credential alias by selecting the magnifying glass icon. Select an existing credential alias from the **Connection & Credential Aliases** list, or select **New** to create a new record. The default is sn_occ.HLA_Credential_Alias. For information about creating a credential alias, see Credential aliases for Discovery.
Subscription Name	The subscription Health Log Analytics uses to receive log data that is published on the Google Cloud Pub/Sub topic. If you leave this field blank, Health Log Analytics uses ServiceNow-Subscription.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration

Advanced tab

Field	Description	Default value
Subscriber Thread Pool Size	The number of concurrent threads that are downloading files from the Google Cloud Pub/Sub topic.	1
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Max length in bytes	The maximum length of log messages, in bytes.	32,766
Character encoding	(Read-only) The character encoding for this data input.	UTF-8
Sub sample drop ratio	The ratio of events to drop.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Rate limit	The maximum number of events per second that this data input processes.	-1
Drop if queue is full	Option to discard logs if many processes are waiting in the queue to access the MID Server.	Clear

Microsoft Azure Event Hubs data input configuration fields

Description of the fields on the Microsoft Azure Event Hubs data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Microsoft Azure Event Hubs is pulled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. This field is required.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Service instance	The service instance to which to bind the log data. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational. This field is required.

The following fields show read-only information:

Field	Description
Status	Status of the data input.
Transport	Protocol used to stream the log data. This data input uses Microsoft Azure Event Hubs to stream log data to your instance.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Query settings tab

Field	Description	Example
From	Starting date and time for reading the data. Data older than this date and time is not read. **Note:** Setting this value to a past date might require the system to read large amounts of data, causing congestion. This field is required.	Now -1 week

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query settings tab (continued)

Field	Description	Example
Max batch size	The maximum number of events passed to a single process call.	100

Transport tab

Field	Description
Event Hubs namespace	The scoping container for the event hub or hubs. This field is required.
Event Hub name	The event hub from which to fetch log data. This field is required.
Event Hub credentials	Shared access authorization policy with which to authenticate to the event hub. This field is required.
Consumer Group name	The Consumer Group to use. This field is required.

Advanced configuration

Advanced configuration form

Field	Description	Default value
Restore point rate	The number of events the data input can read before a restore point is saved.	100
Processor threads count	The total number of processor threads.	2
Operation timeout	The number of seconds to wait before timing out event hubs operations.	120
Receive idle timeout	The number of seconds to wait before timing out receive operations.	60
Prefetch count	The number of events received in advance of event hubs operations.	500
Default timezone	The default timezone if the log doesn't include timezone information.	GMT
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This	-1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration form (continued)

Field	Description	Default value
	setting is used to reduce the number of fetched events.	
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.	-1
Max length in bytes	The maximum length, in bytes, of events.	32766
Character encoding	The character encoding for this data input.	UTF-8
Sleep interval	The interval, in seconds, to wait before querying again after a query has returned no events.	60
Polling interval	The interval, in seconds, to wait before polling for new events.	0
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	False

Microsoft Azure Log Analytics data input configuration fields

Description of the fields on the Microsoft Azure Log Analytics data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID	(Only when the **Execute on** field is set to Specific MID Server) MID Server to which log data from Microsoft Azure Log Analytics is pulled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. • If log ingestion is not enabled for the selected MID Server, Health Log Analytics enables it automatically. This field is required.
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Service instance	The service instance to which to bind the log data. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational. This field is required.

The following fields show read-only information:

Field	Description
Status	Status of the data input.
Transport	Protocol used to stream the log data. This data input uses Microsoft Azure Log Analytics to stream log data to your instance.
Disabled since	The time when the data input stopped or failed.
Sources count	The number of log sources this data input has created.
Last log time	The time when the last log streamed in the data input.
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Query settings tab

Health Log Analytics uses the values set in these fields to generate the query for receiving log data from Microsoft Azure Log Analytics. The Log query field enables you to configure a custom query

Field	Description	Example
From	Starting date and time for reading the data. Data older than this date and time is not read.	Now -1 week

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query settings tab

Health Log Analytics uses the values set in these fields to generate the query for receiving log data from Microsoft Azure Log Analytics. The Log query field enables you to configure a custom query

(continued)

Field	Description	Example
	Note: Setting this value to a past date might require the system to read large amounts of data, causing congestion. This field is required.	
Data source name	The name of the table in Microsoft Azure Log Analytics where the data input fetches the log data. For more information, see the View table information section in the Microsoft Azure documentation. This field is required.	ContainerLog
Event time property name	The Microsoft Azure Log Analytics field in which to detect the event time. This field is required.	TimeGenerated
Max documents per query	The highest number of rows retrieved in each query.	500
Columns to select	Comma-separated list of column names to return. **Note:** This field is ignored when you provide a custom query.	LogEntry,LogEntrySource

Log query	A custom query for receiving log data from Microsoft Azure Log Analytics.
	
	The settings in this field override those in all other fields on the **Query settings** tab except the **From** field. If the **Log query** field is empty, Health Log Analytics
		`{ "query":"ContainerLog`
		
		`| where LogEntry contains 'cartservice'`
		
		`| where TimeGenerated > %s`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query settings tab

Health Log Analytics uses the values set in these fields to generate the query for receiving log data from Microsoft Azure Log Analytics. The Log query field enables you to configure a custom query

(continued)

Field	Description	Example

	generates the query using the values set in the other fields.
	
	For the custom query, use the following JSON format:
	
	`{"query": "query | where TimeGenerated > %s | take 500"}`
		<pre>| take 500",
		"workspaces":
		["defaultworkspace
		-3ab145ff-f9cd-433f
		-8533-d1b1ee24ae
		e6-eus"],
		"project":
		["TimeGenerated",
		"LogEntry",
		"LogEntrySource"]
		}</pre>

Transport tab

Field	Description
Redirect Url	The redirect URL of the access log application. The URL refers to the Microsoft Azure redirect_uri authorization property. For more information, see the Authorization code URL (GET request) section in the Microsoft Azure documentation. This field is required.
Azure service principal credentials	The credentials used to access Microsoft Azure resources. Choose **Client Secret** from the drop-down list. **Note:** You do not need Azure Enterprise Agreement (EA) credentials. This field is required.
Workspace Id	The Customer ID used to call the Microsoft Azure Log Analytics REST API.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration

Advanced configuration form

Field	Description	Default value
Event Processor workers	The number of concurrent event processing workers, where each worker processes a batch of events independently.	4
Workers Queue Size	The queue size of the Event Processor workers.	5
Sub sample drop ratio	The number of events to batch together, out of which one will be discarded. This setting is used to reduce the number of fetched events.	-1
Sub sample receive ratio	The number of events to batch together, out of which all but one will be discarded. This setting is used to decrease the number of received events.	-1
Character encoding	The character encoding for this data input.	UTF-8
Sleep interval	The interval, in seconds, to wait before querying again after a query has returned no events.	60
Polling interval	The interval, in seconds, to wait before polling for new events.	0
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	False

MID Server data input configuration fields

Description of the fields on the MID Server data input configuration form.

Basic configuration

Field	Description
Description	Description of the data input.
MID	The name of the MID Server from which the logs are streamed.

The following fields show read-only information:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Status	Status of the data input.
Transport	Protocol used to stream the log data. This data input uses the MID Server to stream log data to your instance.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.

Advanced configuration

Field	Description	Default value
Default timezone	The default time zone of events. The system uses this default when the log does not specify a timezone.	GMT
Sub sample drop ratio	The ratio of events to discard.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Max length in bytes	The maximum length of log messages, in bytes.	32766
Character encoding	The character encoding for this data input.	UTF-8
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	

REST API data input configuration fields

Description of the fields on the REST API data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
MID Web Server	MID Server to which the log data is pulled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. When you select the MID Server, the **MID** and **Port** fields are completed automatically. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. This field is required.
Service instance	The service instance to which to bind the log data. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational. This field is required.

The following fields show read-only information:

Field	Description
Status	Status of the data input.
Transport	Protocol used to stream the log data. This data input uses REST to stream log data to your instance.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.

Rsyslog, Filebeat, or Winlogbeat data input configuration fields

Description of the fields on the Rsyslog, Filebeat, and Winlogbeat data input configuration forms.

Basic configuration

Getting Started tab

Field	Description
Data input name	Name of the new data input. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Getting Started tab (continued)

Field	Description
Description	Description of the data input.
MID Server	The MID Server to which the logs stream. **Note:** • You can select only MID Servers with log ingestion capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. This field is required.
Port	The port on the MID Server. Choose a port within the suggested range from the array. The port must not be occupied by another process. Make sure that your organization's security team opens the selected port. This field is required.
Content pack	(Linux using Filebeat only) The content pack to use. Content packs contain default source types and mapping script templates. Health Log Analytics activates the selected pack automatically and uses its mapping script for mapping the data input sources. For more information, see Health Log Analytics content packs for quicker time to value.

Tagging and Binding tab

Field	Description
Path	The full path from which to stream logs. You can use a wildcard. This field is required.
Service instance	The service instance to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tagging and Binding tab (continued)

Field	Description
Component	The device type or stack layer as context for the logs that is used for anomaly detection and correlation. For example: Tomcat. Components typically represent CIs in the CMDB. Several components are often clustered together in a single service instance.
Source Type	The source type, which defines how Health Log Analytics handles a specific application and parses the log data. For example: Tomcat Catalina. Each data input can have multiple source types, based on the diversity of its log formats. Service instances and components can have any number of source types.

Advanced configuration

For Rsyslog data inputs:

Rsyslog advanced configuration form

Field	Description	Default values
Use SSL/TLS	Option to use SSL/TLS.	
Look up hostnames	Option to perform DNS lookup to resolve IPs to hostnames.	false
Boss thread count	The number of threads that manage connections.	1
Worker thread count	The number of threads that handle incoming data.	4
Read timeout seconds	The timeout in seconds since the last read. When the timeout expires, the system closes the channel.	30
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Sub sample drop ratio	The ratio of events to drop.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Max length in bytes	The maximum length of log messages in bytes.	32766
Character encoding	The character encoding for this data input.	UTF-8

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Rsyslog advanced configuration form (continued)

Field	Description	Default values
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	

For data inputs that use Beats agents:

Beats advanced configuration form

Field	Description	Default value
Client inactivity timeout (sec)	The timeout, in seconds, to close an inactive channel.	15
Worker thread count	The number of threads that handle incoming data.	4
Default time zone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Sub sample drop ratio	The ratio of events to drop.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Max length in bytes	The maximum length of log messages, in bytes.	32766
Character encoding	The character encoding for this data input.	UTF-8
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	false

ServiceNow Log Export data input configuration fields

Description of the fields on the ServiceNow Log Export data input configuration form.

ServiceNow Log Export data input configuration fields

Field	Description
Name	Name of the new data input. This field is required.
Service instance	The service instance to which to bind the log data. **Important:** The service instance field is pre-populated with a value. Changing this value is necessary only if there is a compelling reason to do so. Consider the implications carefully before making any changes.
Certificate password alias	The credential alias you created in **Basic Auth Credentials**. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow Log Export data input configuration fields (continued)

Field	Description
	For more information, see the Before you begin section in Configure a ServiceNow Log Export data input in Health Log Analytics manually.
Description	Option to add a brief description of the data input to help identify it.
Log level	The log level indicating the severity of the log messages. The available log levels are: • INFO: Information messages representing service actions. This is the default log level. • ERROR: Non-fatal error messages. • WARN: Warning messages indicating potential issues.
Status	Status of the data input: Active or Not Active. This field is read-only.
Disabled since	The time when the data input stopped or failed. This field is read-only.

ServiceNow System Logs Retriever data input configuration fields

Description of the fields on the ServiceNow System Logs Retriever data input configuration form.

ServiceNow System Logs Retriever data input form fields

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Service instance	The service instance to which to bind the log data. The default value is ServiceNow Glide Sys Log. This field is required. **Note:** Health Log Analytics comes with the source type glide_sys_log for log traffic from the ServiceNow Glide Sys Log service instance.
From	From date and time for reading the data. The data input doesn't read data older than this date and time. This field is required.
Status	Status of the data input.
Disabled since	The time when the data input stopped or failed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow System Logs Retriever data input form fields (continued)

Field	Description
Syslog table filter	The criteria used to filter log data from the System Log table before streaming it to the AI engine.

Splunk data input configuration fields

Description of the fields on the Splunk data input configuration form.

Basic configuration

Getting Started tab

Field	Description
Data input name	Name of the new data input. This field is required.
Description	Description of the data input.
MID Server	The MID Server to which the logs stream. **Note:** • You can select only MID Servers that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. This field is required.
Port	The port for the MID Server. Make sure that your organization's security team opens the selected port in the MID Server. This field is required.
Transport Protocol	The protocol used for streaming log messages to your ServiceNow instance. • TCP - When using the Transmission Control Protocol (TCP) protocol, all logs will reach the instance. However, the Splunk pipeline might be blocked if the MID Server is down or the connection to it is lost. TCP is the default transport protocol. • UDP - When using the User Datagram Protocol (UDP) protocol, the Splunk pipeline will never be blocked. However, some logs might be dropped before they reach the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Getting Started tab (continued)

Field	Description
	For more information about streaming log data using the TCP or UCP transport protocol, see the Streaming Splunk data using Heavy Forwarder: Selecting TCP or UDP [KB0998928] article in the Now Support Knowledge Base.
Use Cooked Data	Option to ingest log data from Splunk in the preprocessed ("cooked") format that Splunk uses on the forwarder. Ingesting data into HLA in this format ensures that each log line retains the relevant contextual information that Splunk embeds into it. **Note:** If you select this option, there is no need to edit the `props.conf` and `transforms.conf` files during Splunk data input configuration.
Use Forwarder TimeZone	Option to pass information about the time zone in which the forwarder is located. The MID Server uses this information to adjust for the time zone from which the logs arrive. This option is relevant when using Splunk Universal Forwarders.
Enable Compression	Option to send logs in compressed format. Sending logs in a compressed format minimizes the size of the data being transferred, which is important when dealing with large volumes of log data. This option is relevant when using Splunk Universal Forwarders and can only be used when SSL/TLS is enabled.

Advanced configuration

Advanced configuration form

Field	Description	Default values
Use SSL/TLS	Option to use SSL/TLS. **Note:** To send logs in a compressed format, SSL/TLS must be enabled.	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration form (continued)

Field	Description	Default values
Look up hostnames	Option to perform DNS lookup to resolve IPs to hostnames.	false
Boss thread count	The number of threads that manage connections.	1
Worker thread count	The number of threads that handle incoming data.	4
Read timeout seconds	The timeout in seconds since the last read. When the timeout expires, the system closes the channel.	30
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Sub sample drop ratio	The ratio of events to drop.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Max length in bytes	The maximum length of log messages in bytes.	32766
Character encoding	The character encoding for this data input.	UTF-8
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	

Splunk Polling data input configuration fields

Description of the fields on the Splunk Polling data input configuration form.

Basic configuration

Getting Started tab

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Execute on	Option to select whether to use a specific MID Server or a MID Server cluster. This field is required.
MID	(Only when the **Execute on** field is set to Specific MID Server.) The MID Server to which the logs are streamed. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Getting Started tab (continued)

Field	Description
MID Server Cluster	(Only when **Execute on** is set to Specific MID Server cluster.) The MID Server cluster to which the log data is pulled. This field is required. The data input runs on a single MID Server in the cluster until that MID Server fails. The system then moves all the data input tasks to the next available MID Server in the cluster according to the configured order. **Note:** • Health Log Analytics supports only failover MID Server clusters. In these clusters, multiple MID Servers are grouped together for failover protection. When selecting a cluster from the data input or integration form, the MID Server clusters list displays only failover clusters. • Log ingestion must be enabled for each MID Server in the cluster. If log ingestion is not enabled for the active MID Server, Health Log Analytics enables it automatically. • The default maximum number of data inputs or integrations streaming logs to a single MID Server is 10. A cluster passes capacity validation if it contains at least one MID Server with fewer than 10 data inputs or integrations running on it. This is true even when that MID Server is down. For more information about MID Server clusters, see Configure a MID Server cluster.
Service instance	The service instance to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, create a service instance and add CIs to it. Set the status of the new service instance to Operational.
Transport	The protocol used for streaming log messages to your ServiceNow instance.
Sources count	The number of log sources this data input has created.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Getting Started tab (continued)

Field	Description
Status	Status of the data input.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.

Advanced configuration

Transport tab

Field	Description
Server URL	The URL used to access the Splunk REST API.
Query	The query Splunk uses to search your data.
Authentication Type	The authentication type. • Basic authentication: Sends a username and password with each HTTP request. Basic authentication is simpler than token-based authentication, but less secure. • Token authentication: The client obtains a token from an authentication server and uses that token to authenticate against Splunk.
Splunk Poll Credential Alias	The credential alias to be used. Select the magnifying glass icon to specify a Splunk Poll credential alias. Select an existing credential alias from the **Connection & Credential Aliases** list, or select **New** to create a new record. The selected credential alias can hold one Basic Auth credential and one Token Auth credential. For information about creating a credential alias, see Credential aliases for Discovery.
From	The date and time from which Splunk searches the data.
To	The date and time until which Splunk searches the data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced tab

Field	Description
Max documents per query	The maximum number of documents retrieved each time log data is fetched from Splunk. Default: 10,000.
Splunk request timeout (seconds)	The maximum time, in seconds, allowed for data retrieval before the request times out.

TCP data input configuration fields

Description of the fields on the TCP data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Description	Description of the data input.
Port	The port for the MID Server. Select a unique port from the array. The placeholder shows the range of ports from which to choose. Make sure that your organization's security team opens the selected port. This field is required.
MID	The MID Server to which the logs are streamed. **Note:** • You can select only MID Servers with log ingestion capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. This field is required.
Service instance	The service instance to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.

The following fields show read-only information:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Status	Status of the data input.
Transport	Protocol used to send the log data. Rsyslog and Splunk send data using the TCP protocol.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Advanced configuration

Advanced configuration form

Field	Description	Default value
Use SSL/TLS	Option to use SSL/TLS.	
Look up hostnames	Option to perform DNS lookup to resolve IPs to hostnames.	false
Boss thread count	The number of threads that manage connections.	1
Worker thread count	The number of threads that handle incoming data.	4
Read timeout seconds	The timeout in seconds since the last read. When the timeout expires, the system closes the channel.	30
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Sub sample drop ratio	The ratio of events to drop.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Max length in bytes	The maximum length of log messages in bytes.	32766
Character encoding	The character encoding for this data input.	UTF-8
Drop if queue is full	Option to discard logs if there is a load on the MID Server.	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration form (continued)

Field	Description	Default value
Line breaker delimiters	The line break character separating the raw log lines. Splitting values must be separated by a comma followed by a space: ", ". For example: "\r, \n, , splitHere, #".	

UDP data input configuration fields

Description of the fields on the UDP data input configuration form.

Basic configuration

Field	Description
Name	Name of the new data input. This field is required.
Port	The port for the MID Server. Select a unique port from the array. The placeholder shows the range of ports from which to choose. Make sure that your organization's security team opens the selected port. This field is required.
MID	The MID Server to which the logs are streamed. **Note:** • You can select only MID Servers with log ingestion capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. This field is required.
Service instance	The service instance to which to bind the log data. This field is required. **Note:** If no relevant service instance exists, Create an service instance and add CIs to it. Set the status of the new service instance to Operational.

The fields in the following table show read-only information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Status	Status of the data input.
Transport	Protocol used to send the log data: UDP.
Sources count	The number of log sources this data input has created.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Advanced configuration

Field	Description	Default value
Default timezone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Sub sample drop ratio	The ratio of events to drop.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Rate limit	Limitation on the number of events per second that the data input processes.	-1
Max length in bytes	The maximum length of log messages, in bytes.	32,766
Character encoding	The character encoding for this data input.	UTF-8
Drop if queue is full	Option to discard logs if many processes are waiting in the queue to access the MID Server.	false

Vector Agent data input configuration fields

Description of the fields on the Vector Agent data input configuration form.

Getting Started tab

Field	Description
Data input name	Name of the new data input. This field is required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Getting Started tab (continued)

Field	Description
Description	Description of the data input.
MID Server	The MID Server to which the Vector Agent logs are streamed. This field is required. **Note:** If the MID Server is down or the connection is lost, the Vector pipeline may get blocked.
Port	The port on the MID Server used for listening to logs. This field is required. Choose a port within the range of 6000-6200.
Transport	The protocol used for streaming log messages to your ServiceNow instance. Currently, only TCP is available.
Status	Status of the data input.

Surrounding logs tab fields

This section describes the information displayed on the **Surrounding logs** tab.

Surrounding logs tab

Column	Description
Time	Timestamp of the log line in the format that the source uses. If no value appears, then check the source type structure of the raw data.
Application service	Service instance in which the metric was found.
Component	Logical component of the service instance that generated the event. Multiple CIs can sometimes perform the same function.
Message	Inner message of the raw log line that contains the text of the system-generated log message regarding the nature of the occurrence.
Level	Type of event. The available values, in order of importance, are: • **Emergency** • **Alert** • **Critical** • **Error** • **Warning** • **Notice**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Surrounding logs tab (continued)

Column	Description
	• **Informational** • **Debug**
Host	Host identifier from the log line that consists of the hostname or IP address of the endpoint.
Log message	The raw log message without the header.

Log Viewer table fields

This section describes the information displayed in the **Log Viewer** table.

Log Viewer table fields

Column	Description
Time	Timestamp of the log line in the format that the source uses. If no value appears, then check the source type structure of the raw data.
Service instance	Service instance in which the metric was found.
Component	Logical component of the service instance that generated the event. Multiple CIs can sometimes perform the same function.
Message	Inner message of the raw log line that contains the text of the system-generated log message regarding the nature of the occurrence.
Level	Type of event. The available values, in order of importance, are: • **Emergency** • **Alert** • **Critical** • **Error** • **Warning** • **Notice** • **Informational** • **Debug**
Host	Host identifier from the log line that consists of the hostname or IP address of the endpoint.
Log message	The raw log message without the header.

Note: The Raw message column does not appear in the Log Viewer table by default. You can display this column by selecting it from the **Filters** pane. For more information, see Customize the Log Viewer table in Health Log Analytics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Log correlators form fields

This section describes the fields on the Log correlators form.

Log correlators form

Field	Description
Type	Type of log correlator. Choices are as follows. • **Free text correlator**: The log correlator analyzes text in the log message. • **Log-key correlator**: The log correlator analyzes log metadata. For example, the name of a service instance in your infrastructure. Log property correlators are specific to the business context of your environment. For more information, see Log correlators for identifying related alerts.
Free text term or Identifier	Text that the log correlator isolates when parsing log lines. The label for this field is **Free text term** for free text log correlators and **Identifier** for log property correlators. **Note:** The value of this field can be the text for either type of log correlator. As a result, the label for this column in the resulting list of log correlators is **Correlation indicator**.
Active	Option to apply the log correlator. When you select this field, the system applies the log correlator to the log stream.
Range of analysis	Set of sources whose log data are analyzed by this log correlator. Choices are as follows: • **All sources**: The log correlator is applied to log lines from all sources in the data stream. • **Only new sources**: The log correlator is applied to log lines from all sources created after this log correlator is activated. • **Specified source**: The log correlator is applied only to the log lines from the source that you specify on this form.
Excluded sources	Sources that are excluded from the log correlator. The system does not analyze log lines from sources that are listed in this field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Log correlators form (continued)

Field	Description
	Note: This field applies only when the **Range of analysis** field is set **All sources** or **Only new sources**. For information about excluding a source from a log correlator, see Exclude a source from a log correlator in Health Log Analytics.
Source	Source of the log correlator. The log correlator analyzes the data of this log source. This field appears only when you're adding the log correlator from the Log Sources form. The field is automatically set to the log source that you're working on.

Log-based CI candidates fields in Health Log Analytics

Field descriptions for the Log-based CI candidates table, the Log-based CI candidate form, and the Add CI candidate to CMDB form.

Log-based CI candidates table

Field	Description
Host name	The host name that Health Log Analytics discovered in the log data.
CI class type	The CI class type discovered in the log data. The CI class type is determined by the data input type: • Rsyslogs – CI class type: Linux Server • Linux Filebeat - CI class type: Linux Server • Windows Filebeat - CI class type: Windows Server • Winlogbeat - CI class type: Windows Server
Last log time	(Read only) The last time this host was found in the log data.
Operation	(Read only) The suggested action to take: Add - Add the CI to the CMDB.
Status	(Read only) The status of the suggested action. • Pending - No action has been taken. • If the taken action failed, this field contains an error message.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Log-based CI candidate form

Field	Description
CI class type	The CI class type discovered in the log data. The CI class type is determined by the data input type: • Rsyslogs – CI class type: Linux Server • Linux Filebeat - CI class type: Linux Server • Windows Filebeat - CI class type: Windows Server • Winlogbeat - CI class type: Windows Server
CI name	The host name that Health Log Analytics discovered in the log data.
Operation	(Read only) The suggested action to take: Add - Add the CI to the CMDB.
Last log time	(Read only) The last time this host was found in the log data.
Status	(Read only) The status of the suggested action: • Pending - No action has been taken. • If the taken action failed, this field contains an error message.

Add CI candidate to CMDB form (Linux Server)

Field	Description
Name	The name of the CI.
Asset tag	Alphanumeric information to help track the asset associated with the CI.
Manufacturer	The device manufacturer.
Asset	(Read-only) The asset associated with the CI.
Class	The asset group. For example, Linux Server.
Comments	Information about the asset.
Company	The company associated with the CI. This field is required.
Serial number	The serial number associated with the CI.
Model ID	The device model ID.
Assigned to	The person using or primarily responsible for this item.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add CI candidate to CMDB form (Linux Server) (continued)

Field	Description
Host name	The host name of the CI.
OS Domain	The OS domain of the CI.
Operating System	The OS running on the CI.
OS Version	The version of the OS running on the CI.
OS Service Pack	The service pack installed on the OS.
DNS Domain	The DNS domain of the CI.
IP Address	The IP address of the CI.
Disk space (GB)	The amount of disk space, in GB.
Description	Textual description of the CI.
RAM (MB)	The amount of RAM on the computer, in MB.
CPU manufacturer	The CPU manufacturer.
CPU type	The CPU type.
CPU speed (MHz)	The speed of the CPU, in MHz.
CPU count	The number of CPUs.
CPU core count	The number of cores per CPU.
IsVirtual	Option to indicate if the device is running on a virtual machine instance.

Metric Data table columns in Health Log Analytics

The Metric Data table lists all metrics that Health Log Analytics collects. Use the table to query metrics by log source, integration, source type, and metric type, review anomaly detection statistics, and view metric baselines.

For more information about using the Metric Data table, see View log anomaly metrics in Health Log Analytics.

Column descriptions

Each row in the Metric Data table represents one metric that HLA collects.

Metric Data columns

Column	Description
Metric Name	Unique identifier of the metric being tracked.
Service	Service instance associated with the metric.
Source	Log source from which the metric originates.
Component	Component within the log source that generates the metric.
Metric Type	Type of measurement, such as meter or gauge.
Dimension	AI Engine metric identifier associated with this record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric Data columns (continued)

Column	Description
Subject	Subject identifier that groups related metrics.
Class	Metric classification, such as Keyword Metric, Pattern ID Metric, Raw Metric, or All Events Metric.

Sustained alert table schema

Reference documentation for the sn_occ_sustained_alerts table that stores sustained alert feedback configurations per metric.

Table overview

The `sn_occ_sustained_alerts` table stores per-metric sustained alert feedback configurations. Each record represents a sustained alert setting applied to a specific metric. The table supports multi-tenant deployments through domain separation and provides audit trail support through standard system fields.

Table access and permissions

Access to the sustained alert table is controlled by role-based access control (RBAC) lists. Standard read, write, create, and delete operations are supported. Users with alert feedback permissions can create and modify sustained alert entries for metrics they have access to.

Field definitions

Field name	Type	Description
`metric_id`	String (indexed)	Full metric identifier and foreign key reference to the metric being configured. Required field. Indexed for fast lookup during detection evaluation.
`required_duration_minutes`	Integer	Minimum sustained duration threshold in minutes. Valid range: 1–120 minutes. The system requires anomalies to persist continuously for this duration before emitting a detection. Updates without creating a new record when feedback is reapplied.
`duration_mode`	Choice (enum)	Feedback mode indicator. Values: `adaptive` (system-calculated duration based on signal characteristics) or `manual` (operator-selected duration). Default: `adaptive`.
`sustain_count`	Integer	Counter tracking how many times sustained alert feedback has been applied to this metric. Auto-increments on each reapplication. Allows the system to learn from repeated operator feedback patterns.
`applied_by`	User reference	User who initially applied the sustained alert feedback. Captured for audit trail purposes.
`applied_at`	Timestamp	Date and time when the sustained alert feedback was first applied. Captured for audit trail purposes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field definitions (continued)

Field name	Type	Description
`expires_at`	Datetime (nullable)	Optional expiration timestamp. When set, the feedback automatically becomes inactive after this date. Default: null (feedback does not expire).
`active`	Boolean	Active status flag. Values: `true` (feedback is evaluated) or `false` (feedback is deactivated). Default: `true`. The system skips inactive entries during detection evaluation.
`sys_domain`	Domain reference	Multi-tenancy support field. Indicates which domain (partition) the record belongs to. Enforces domain isolation for all operations and queries.
`sys_created_by`	User reference	System-generated field capturing the user who created the record. Set automatically on record creation.
`sys_updated_by`	User reference	System-generated field capturing the user who last modified the record. Updated automatically on any change.
`sys_created_on`	Timestamp	System-generated field capturing the creation timestamp. Set automatically on record creation.
`sys_updated_on`	Timestamp	System-generated field capturing the last modification timestamp. Updated automatically on any change.

Constraints and validation

The sustained alert table enforces these constraints:

- Metric uniqueness per domain: Each metric can have only one active sustained alert entry per domain. Reapplying feedback updates the existing record rather than creating a duplicate.
- Duration range validation: `required_duration_minutes` must be between 1 and 120 minutes. Invalid values are rejected.
- Domain isolation: All queries and operations are scoped to the active domain. Cross-domain access is blocked by ACL enforcement.
- Expiration enforcement: Records with `expires_at` dates in the past are treated as inactive during detection evaluation, even if `active` is `true`.

Performance characteristics

The sustained alert table is optimized for detection-time lookups:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Lookup performance: P95 query latency ≤ 50 milliseconds per metric lookup, enabled by indexing on `metric_id`.
- Caching: The system caches sustained alert entries in memory with automatic invalidation when table records change, reducing database query load.
- Feedback application: Single-alert feedback application responds in P95 ≤ 500 milliseconds; bulk feedback (up to 50 alerts) in P95 ≤ 2 seconds.

Audit trail and history

All changes to sustained alert records are logged in the audit table. The following changes are tracked:

- Record creation (user, timestamp, initial configuration)
- Duration threshold modifications (`required_duration_minutes` changes)
- Mode changes (switches between adaptive and manual)
- Deactivation or removal (record deletion or `active` flag changes)
- Expiration date updates

The `sys_created_by`, `sys_updated_by`, `sys_created_on`, and `sys_updated_on` fields provide basic audit information. For detailed change history, query the audit table using the sustained alert record's `sys_id`.

Integrations Launchpad in Service Operations Workspace for ITOM

The Integrations Launchpad showcases all Event Management connector integrations, allowing users to pull or push events from external or internal devices. This centralized hub simplifies Event Management by consolidating integration options in one place.

The Integrations Launchpad streamlines the process of managing and utilizing event connectors, enhancing operational efficiency and ensuring seamless integration of event data from various external sources. This centralization helps users maintain better control and visibility over their Event Management ecosystem.

Browse Integrations tab

Displays all pull, push, and custom connector integrations, including connectors available from the ServiceNow store and third parties. This tab features:

- Search Box: Helps you quickly find a specific data source by entering keywords, improving efficiency and saving time.
- All Integrations list: Organizes integrations into categories such as All Integrations, Events (Push, Pull, Custom), Metrics (ACC-M), and Logs (HLA). This categorization makes it easier for users to browse the list and select the relevant type of integration they need.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Browse integrations tab

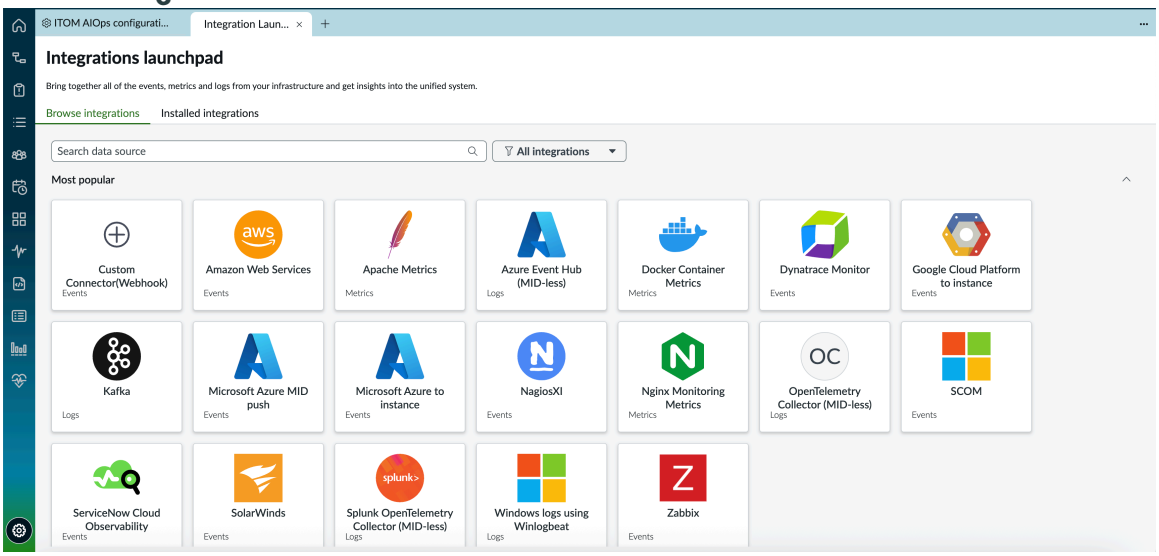

Installed Integrations tab

Active and inactive installed integrations

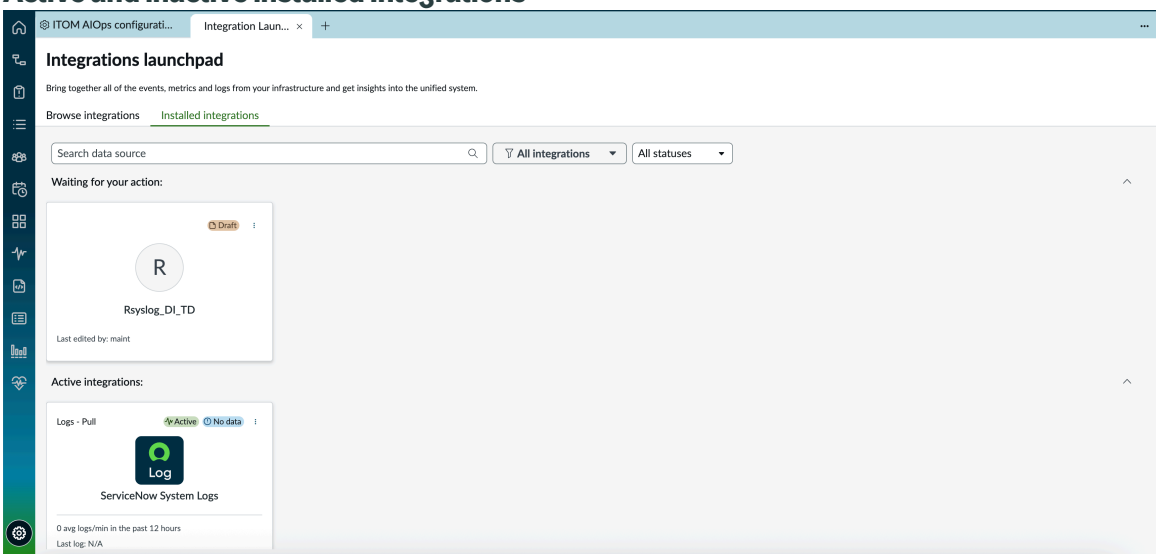

Lists all user-configured integrations, both active and inactive. This tab features:

- Search Box: Helps users quickly find a specific data source by entering keywords, improving efficiency and saving time.
- All Integrations list: Organizes integrations into categories such as All Integrations, Events (Push, Pull, Custom), Metrics (ACC-M), and Logs (HLA). This categorization makes it easier for users to browse the list and select the relevant type of integration they need.
- All statuses list: Allows you to filter integrations by their status—All Statuses, Active, Inactive, Draft, or No Data. This helps users quickly identify and manage integrations based on their current state.

Integration status and description

Status	Description
All statuses	All integration statuses
Active	Active integrations that were successfully installed

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Status	Description
Inactive	Inactive integrations that were successfully installed
Draft	In-progress integrations with a configuration that was started but not completed
No data	Integrations that were successfully installed with no data

Configure an event pull connector

Configure event pull connectors that require a script, connector definition, and connector instance to pull events from external devices. These connectors automate the data retrieval process, ensuring the seamless integration of external events into your system for efficient monitoring and management.

Before you begin

Role required: evt_mgmt_admin

About this task

Starting with the Brazil release, you can use ServiceNow Otto® to configure the following connectors: vCenter, NagiosXI, SolarWinds, SCOM, and Zabbix. Each connector tile displays the **Install with Otto** option. When you select it for any connector, the Otto panel opens on the right side of the screen. In the Otto panel, you can chat with Otto to configure the connector, and later deactivate or delete it if required. You need the ServiceNow Otto for ITOM plugin for the **Install with Otto** option to appear.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. In the **Browse Integrations** tab, select the **All integrations** drop-down list and select **Pull**.
   Only pull connector tiles are displayed.

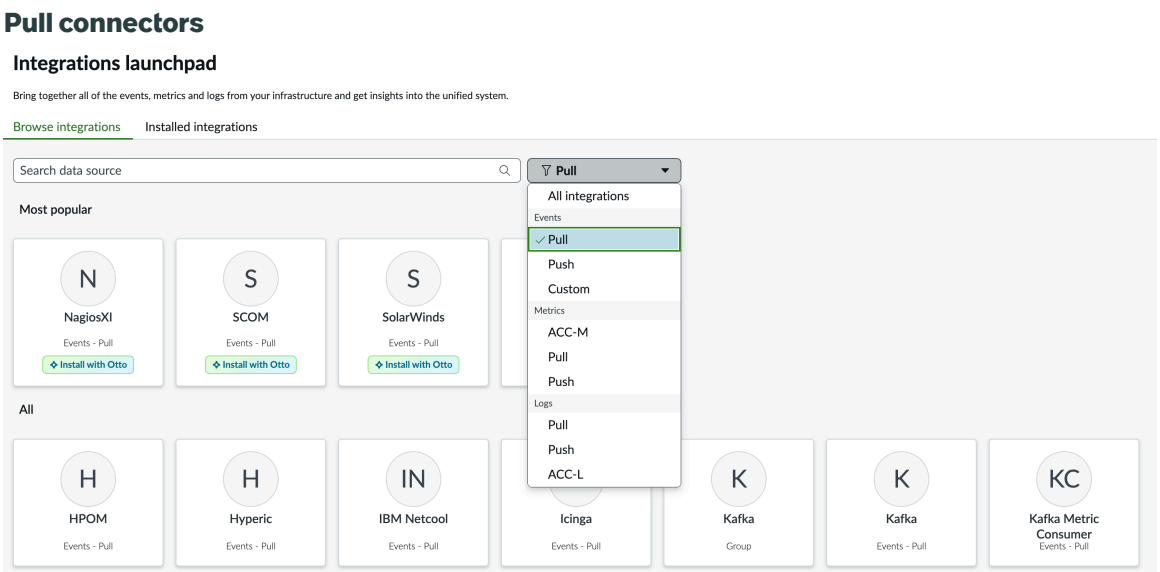

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Select a pull connector tile.
   If a pop-up menu opens, select the data to track from that connector and select **Continue**.

Pull connectors details page

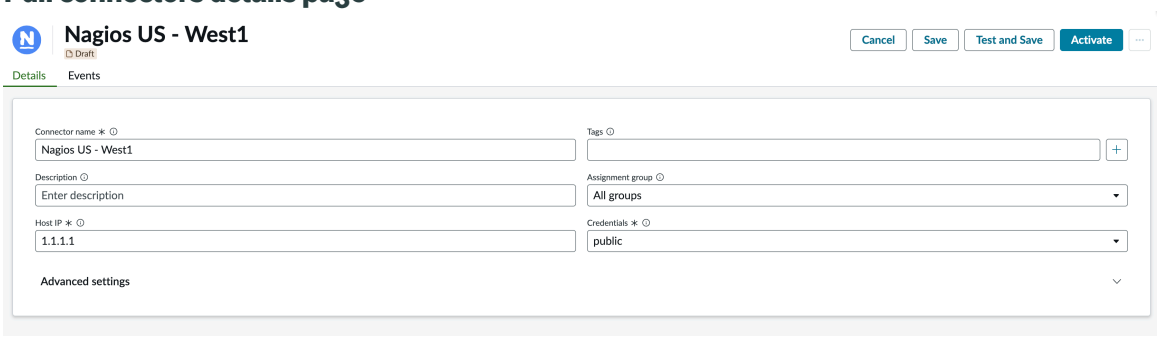

6. In the **Details** tab, in the **Connector name** field, enter a unique name for the connector type.
7. **Optional:** In the **Tags** field, enter tags to help locate and identify connectors of this type on the Express List.

   To add additional tags, select the plus icon ( ) next to the **Tags** field. The Add tags window opens.

   a. In the **Key** field, enter a tag key.

   b. In the **Value** field, enter a tag value.

   c. Select **Add**.
8. **Optional:** In the **Description** field, enter brief information about this connector.
9. **Optional:** From the **Assignment group** drop-down list, select the group or team that is responsible for managing and maintaining the connector.
10. In the **Host IP** field, enter the IP address that is used to select the appropriate MID Server for communicating with the event source host.
    You can use the host name instead of the host IP.
11. From the **Credentials** drop-down list, select the valid credentials to access the event source host.
    To create a new credential, select **Create new credentials** and continue with the process.
12. **Optional:** If you want to customize configurations to meet specific requirements, select **Advanced settings** and fill in the form.

Field	Description
Connected MID Server	Determine whether you want to have the system choose the optimal MID Server or choose a specific MID Server yourself. Perform only if you must make modifications to the MID Server otherwise leave it empty and let the system choose the optimal MID Server.
Event collection schedule (seconds)	Enter the frequency in seconds that indicates how often the system should check for new events from the external event source. The value can't be lower than the minimum schedule property, which by default is 120 seconds.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Protocol	Defines the communication protocol used for retrieving events from external devices.
severity_by_state_type	Determines the severity level assigned to events based on their state type.
time_zone	Specifies the time zone used for recording and displaying event timestamps.
offset_min	Sets the time offset for event retrieval to ensure accurate synchronization.
Port	Verify the default values for your connector type such as port details or time formats and modify them as necessary. Default values from your connector are filled-in. Specifies the port number used for communication with external devices.
max_fetch_interval_min	Limits the maximum interval between event fetch operations.
initial_sync_in_days	Sets the duration for the initial synchronization process when retrieving historical data.

13. **Optional:** Test the connector before activating it by selecting **Test and Save**.
    When the test succeeds, a pop-up appears where you can activate the pull connector by selecting **Activate integration**.
14. To save the connector, select **Save**.
15. To activate the pull connector, select **Activate**.
    The pull connector is activated and its tile is displayed in the **Installed Integrations** tab.

What to do next

To confirm the successful configuration of the pull connector and the flow of events into the system, select the **Related events** tab. If the setup is successful, proceed to **Add automation** to establish an alert automation. For details on **Related events** tab, see Review integration and configuration health and for information on alert automation, see Alert automation in Service Operations Workspace for ITOM.

Configure an event push connector

Integrate with an event push connector to connect to an external event source and push event information to your ServiceNow instance.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. In the **Browse Integrations** tab, select the **All integrations** drop-down list and select **Push**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Only push connector tiles are displayed.

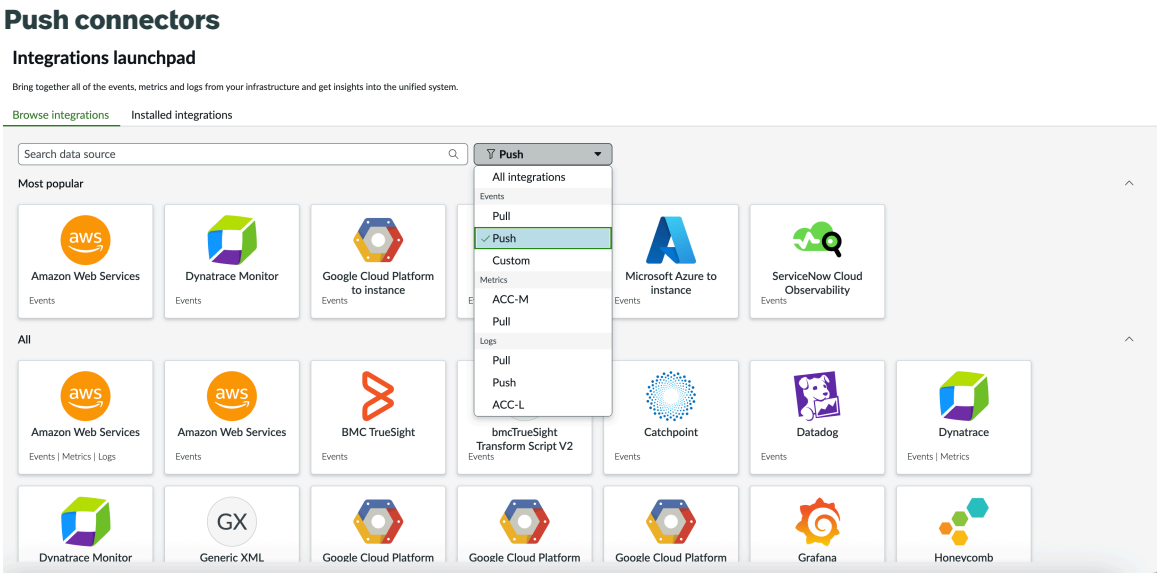

5. Select a push connector type tile.
6. Provide details about the connector type on the **Provide details** screen.

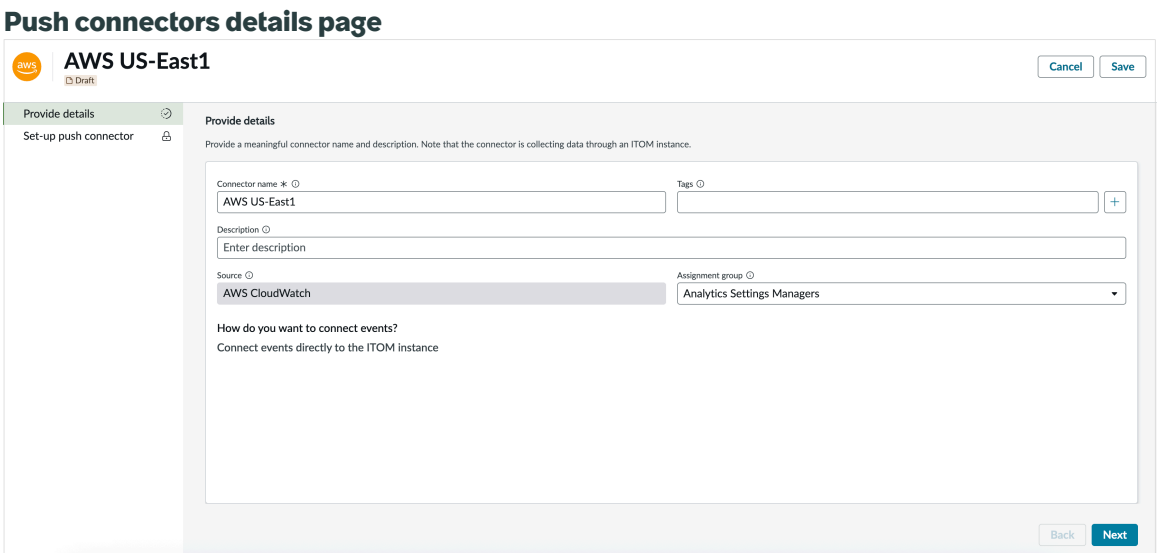

a. In the **Connector name** field, enter a unique name for the connector type.

b. **Optional:** In the **Tags** field, enter tags to help locate and identify connectors of this type on the Express List.
To add additional tags, select the plus sign next to the **Tags** field. The Add tags window opens.

   i. In the **Key** field, enter a tag key.

   ii. In the **Value** field, enter a tag value.

   iii. Select **Add**.

c. **Optional:** In the **Description** field, enter information to help identify this connector type.

d. Verify that the vendor name in the **Source** field is correct.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

e. **Optional:** From the **Assignment group** drop-down list, select the group or team that is responsible for managing and maintaining the Push connector.

f. Select **Next**.

The connector is saved in draft mode, and the process advances to the Set-up push connector section.

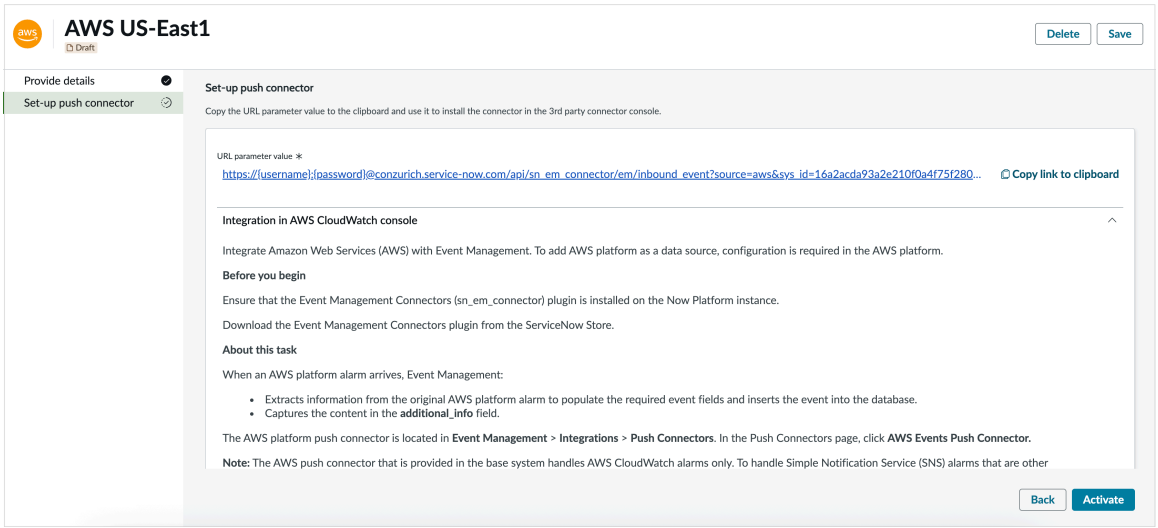

7. In the **URL parameter value** field, select **Copy to clipboard** to copy the auto-generated URL parameter value to the clipboard.

   This URL is generated by the code that depends on the ServiceNow instance and connector type to receive events from the third-party connector. Pasting the URL in your third-party connector configuration creates a webhook that enables it to send events to the ServiceNow instance.

8. In the Select authentication method, select the authentication type.

   Integration Launchpad shows authentication options based on the supported authentication types. You can either create new credentials, if you have the required roles, or select from existing credentials.

   - **Basic**: Sends a username and password (Base64-encoded) with each request. The user_admin role is required to create new Basic Auth credentials. The drop-down list displays users who have the evt_mgmt_integration role. Users with the evt_mgmt_admin role are excluded from the drop-down.
   - **OAuth**: Grants apps limited, token-based access to resources without sharing credentials, using scoped, expiring access tokens. The oauth_admin role is required to create new OAuth client credentials.
   - **API Keys**: Uses a unique secret key to identify and authorize an application or user when calling an API. Creates a new API key credential. Associates the API key with the sn_em_connector.integration user, which is exported as part of the plugin.

9. Select **Activate** and then select **OK**.
10. In the **Details** tab, view the connector details and the URL link.

    Integration Launchpad uses these fields to render the latest error details for the push connector instance:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Last error time**: Stores the timestamp of the latest error. This uses the glide_date_time field type, allowing the Details page to show when the last failure occurred.
- **Last error message**: Stores the latest error message returned for the push connector instance. This read-only string field has a maximum length of 4,000 characters.
- **Last error payload**: Stores the detailed error payload associated with the latest failure. This read-only string field has a maximum length of 10,000 characters.

11. Follow the step-by-step instructions to integrate your third-party connector with Event Management.
    The connector is saved and a notification is sent after the connector is active.

What to do next

To confirm the successful configuration of the pull connector and the flow of events into the system, go to the **Related events** tab. If the setup is successful, proceed to **Add automation** to establish an alert automation. For details on **Related events** tab, see Review integration and configuration health and for information on alert automation, see Alert automation in Service Operations Workspace for ITOM.

Configure an event custom connector

Authenticate your event custom connector as a data source to enable the pushing of event information to your ServiceNow instance. Define the custom connector by providing its details and mapping fields. This allows the extracted information from your custom connector's event messages to populate the required event fields. Finally, create and supply a URL for the custom connector webhook.

Before you begin

- Ensure the Event Management connectors (sn_em_connector) plugin is installed on the ServiceNow AI Platform instance. For more information, see Configure Event Management connectors.
- Ensure that the configuration items for the hosts managed by your custom connector exist in the ServiceNow AI Platform instance. These CIs can be physical or virtual and can be either manually created or discovered via IP discovery or Cloud Discovery. For more information, see Access to cloud environments for ITOM products.

Role required: evt_mgmt_admin

About this task

To start pushing event information to your ServiceNow instance, you must set your ServiceNow AI Platform instance as the REST endpoint using a standard webhook on your custom connector.

By providing connector details and mapping fields, you define the custom connector. The extracted information from the custom connector's event messages then populates the required event fields and inserts the event into the database.

The URL is generated by the code based on the ServiceNow instance and the custom connector. Use this URL to create a webhook that enables your custom connector to send events to the ServiceNow instance.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. In the **Browse Integrations** tab, select the **All integrations** drop-down list and select **Custom**. Only custom connector tiles are displayed.

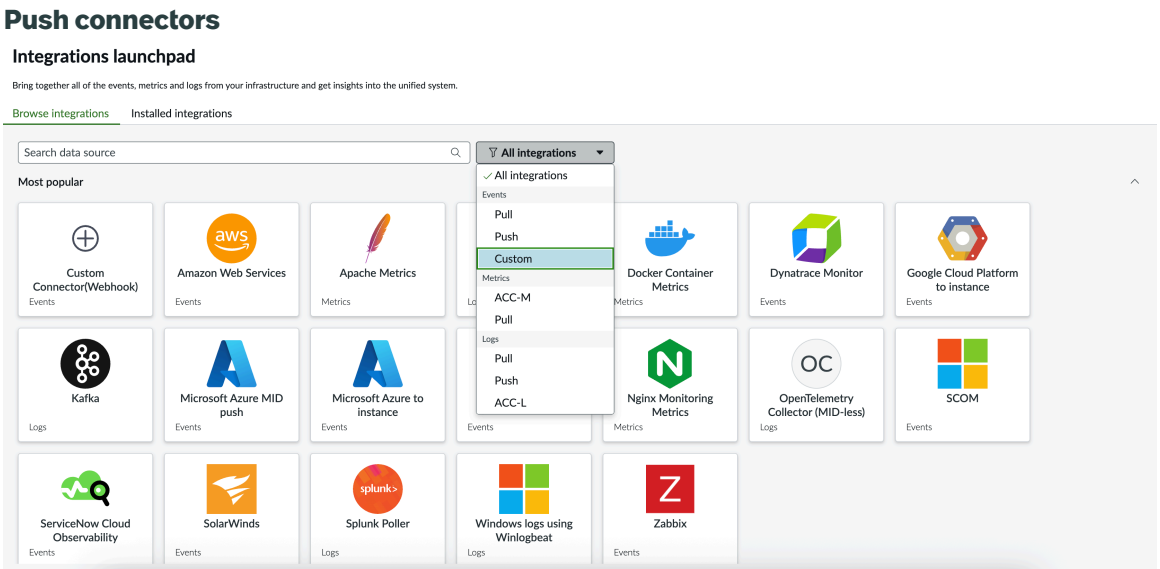

5. Select a custom connector type tile.
6. On the **Provide details** page, provide the details for the connector.

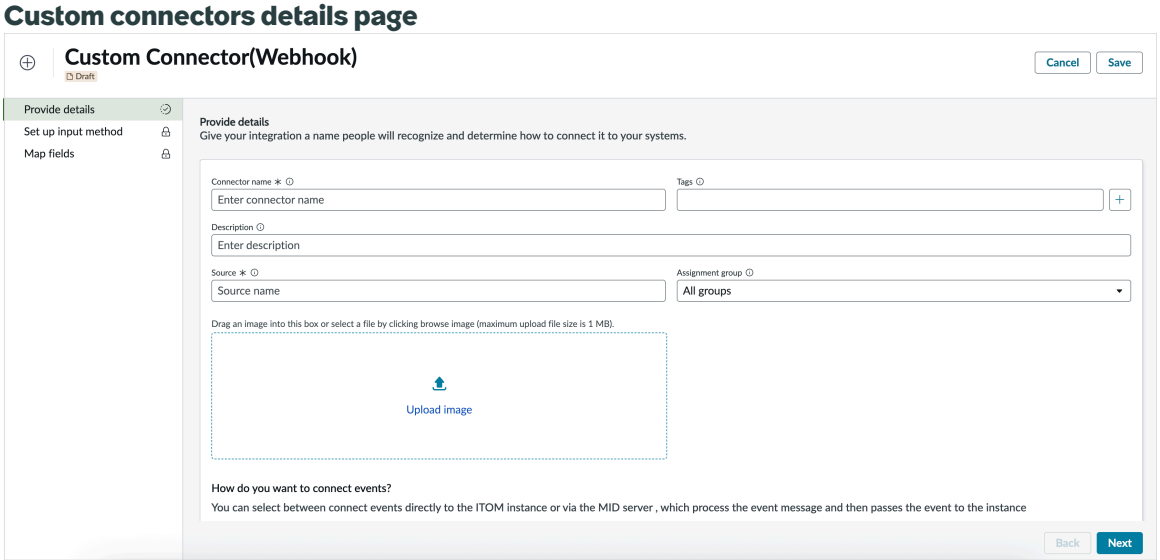

   a. In the **Connector name** field, enter a unique integration name for the connector.

   b. **Optional:** In the **Tags** field, enter tags to help locate and identify connectors of this type on the Express List.

      To add additional tags, select the plus icon ( ) next to the **Tags** field. The Add tags window opens.

      i. In the **Key** field, enter a tag key.

      ii. In the **Value** field, enter a tag value.

      iii. Select **Add**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. **Optional:** In the **Description** field, enter brief information about this connector.

d. In the **Source** field, enter the vendor used to gather events from the external event source.

The **Source** field identifies the connector and serves as an identifier to recognize the source of events and alerts once the connector is active.

e. **Optional:** From the **Assignment group** drop-down list, select the group or team that is responsible for managing and maintaining the connector.

7. **Optional:** Drag or upload an image to display in your connector tile.
   - Drag an image into the box.
   - Browse to the image file by selecting **Upload image** and then choose the file.

   The maximum upload file size is 1 MB.
8. Determine whether you want to connect events directly to the ITOM instance or through a MID Server.
   - To send events directly to the ITOM instance, select **Directly to the ITOM instance**.
   - To send events through a MID Server, which processes and then passes the event to the instance, select **Via the MID server**.

     If a MID Server exists, from the **Select MID server** field, select an available MID Server connected to an events listener.

     If you must create a MID Server, then select one of the following.
     - **Configure MID web server extension**
     - **Configure MID web service event collector context**

     For more information, see Configure the MID Web Server extension and Configure the MID WebService Event Collector Context.
9. Select **Next**.

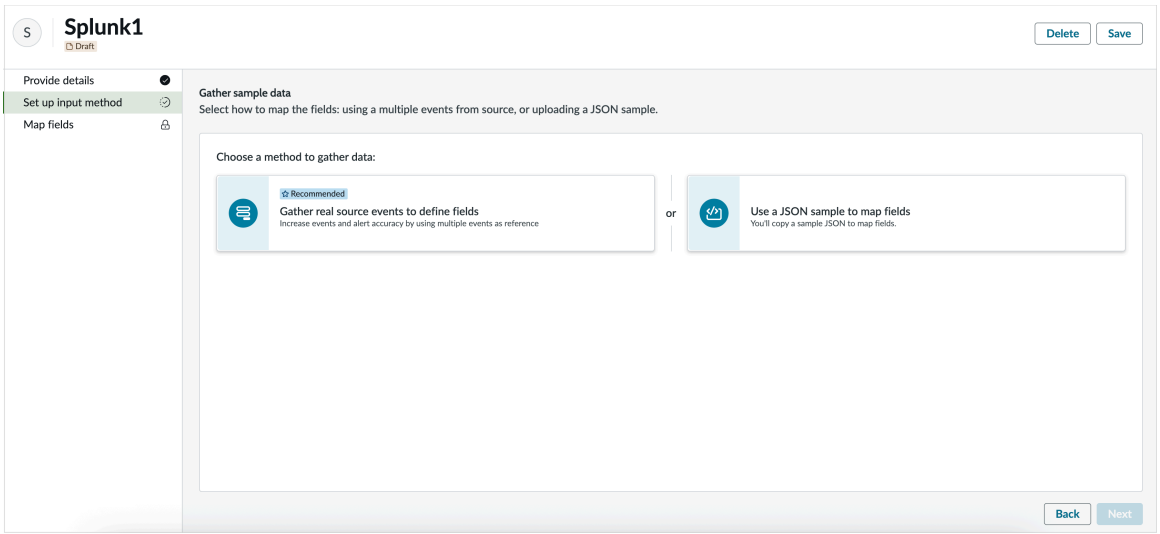

10. On the **Set up input method** page, select a method to gather event data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Description
Gather real source events to define fields	**a.** Select **Gather real source events to define fields**. **Note:** By gathering real-source events from your custom connector, you increase events and alert accuracy by using multiple events as a reference to define the fields. **b.** Select **Copy to clipboard** and copy the URL. **c.** Paste the URL into your custom connector's web hook setup to create a webhook that enables your custom connector to send events to the Service Now instance. Enter your basic authentication cre dentials into the monitoring system's webhook con figuration (such as the username and password). The webhook enables your custom connector to send events to the instance. **d.** Select **Check for events**. The number of events are displayed. The event creates JSON files that you can browse after selecting **Next**.
Use a JSON sample to map fields	**a.** Select **Use a JSON sample to map fields**. **Note:** Only if you have a sample JSON se lect this option. **b.** In the **Enter payload** dialog box, copy and paste your JSON sample payload to create and view the mapping to map the fields. 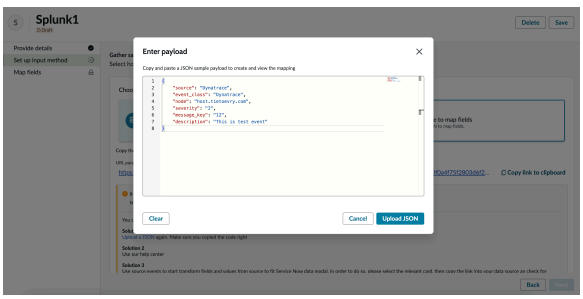 **c.** Select **Upload JSON**. **d.** Select **Edit or paste new JSON** to edit the JSON that you uploaded or paste a new one. **e.** Select **Copy to clipboard** and copy the URL. **f.** Paste the URL into your custom connector's web hook setup to create a webhook that enables your custom connector to send events to the Service Now instance. Enter your basic authentication cre dentials into the monitoring system's webhook con figuration (such as the username and password). The webhook enables your custom connector to send events to the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Select **Next**.
    The page opens where you can map fields.

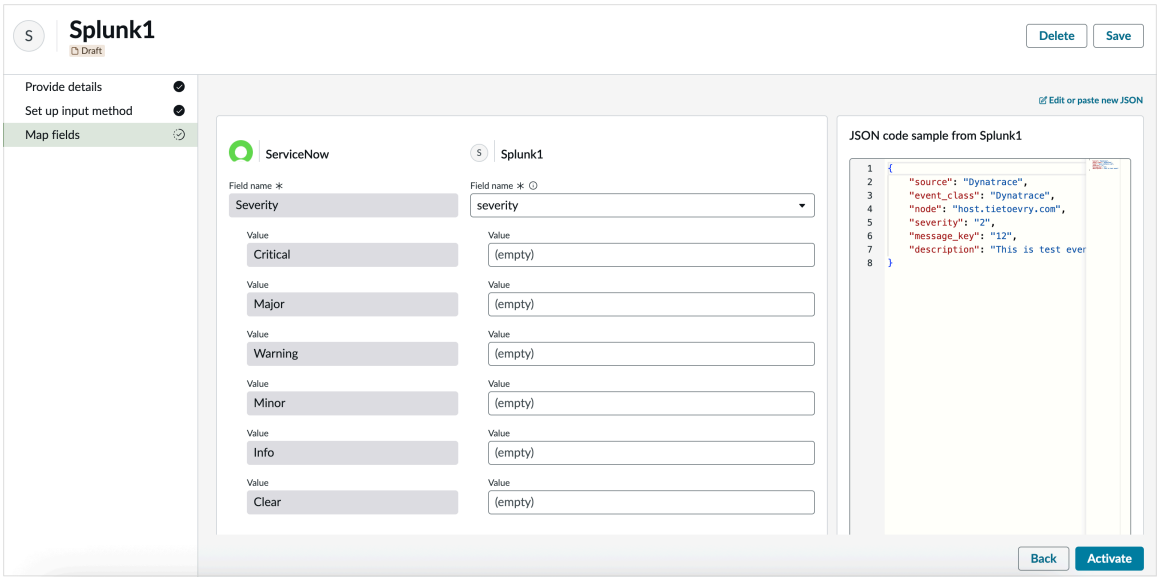

12. On the **Map fields** page, in the connector field names column, enter at least one value that maps to the severity values listed in the ServiceNow **Severity** field column.
13. On the **Map fields** page, in the connector field names column, enter at least one value that maps to the severity values listed in the ServiceNow **Severity** field column.

    You can enter the severity value from the JSON code sample. To set a severity value, you must be familiar with your custom connector's payload and the types of severity it supports. The severity values of the received events must match the generated events in your system.

    Enable the system to create events correctly from incoming data by mapping the fields in your JSON code sample to the ServiceNow fields. The **Severity** and **Message key** information are automatically filled in to save you time and effort. The **Message key** provides the correct identifier for the events, ensuring de-duplication works correctly and redundant alerts are not generated.
14. **Optional:** To view multiple JSON files, select the arrow above the JSON sample.
15. **Optional:** In the Other fields section, in your connector column, specify custom event parameters that are optional but provide additional information about the event.
    If you choose not to provide the optional information, leave the field empty.

    a. In the **Description** field, enter a meaningful alert description.

    b. In the **Metric name** field, enter a metric name that allows you to associate events with the performance metrics or KPIs that triggered the alert.

    c. In the **Node** field, enter a node value to be used for CI binding.

    d. In the **Event time** field, enter the creation time of the custom connector event.
    If the event time isn't mapped, the system will use the default event time.

    e. In the **Event time settings** field, select the time format and the event time zone.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Time format**: A GlideDateTime field in UTC or GMT format, with a maximum length of 40 characters.
- **Time zone**: Times are stored in Coordinated Universal Time (UTC) and are displayed globally based on the system time zone.

16. Activate and save the connector by selecting **Activate** and then select **OK** on the confirmation pop-up window.
    You can view the connector details on the **Details** tab and your field mappings on the **Fields mapping** tab.

Result

After the connector is activated, events start flowing from your custom connector console into the ServiceNow instance. The plug-in extracts information from the original custom connector event message to populate the required event fields and inserts the event into the database. To see the events in your ServiceNow AI Platform instance, navigate to All Events. For more information, see View events.

What to do next

To verify that the process was successfully completed and events are flowing into the system, view the related events on the **Related events** tab. If the setup is successful, proceed to **Add automation** to establish an alert automation. For information on alert automation, see Alert automation in Service Operations Workspace for ITOM.

Configure Solarwinds metric pull connector

Configure a metric pull connectors that require a script, connector definition, and connector instance to pull metrics from external sources. These connectors automate the data retrieval process, ensuring the seamless integration of external metrics into your system for efficient monitoring and performance analysis.

Before you begin

Role required: evt_mgmt_admin

About this task

Note: The new UI for configuring metric pull connector is currently available only for the SolarWinds Metric Connector. For all other metric connectors, the configuration continues to use the Platform Table UI.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. In the **Browse Integrations** tab, select the **All integrations** drop-down list and select **Metrics** > **Pull**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

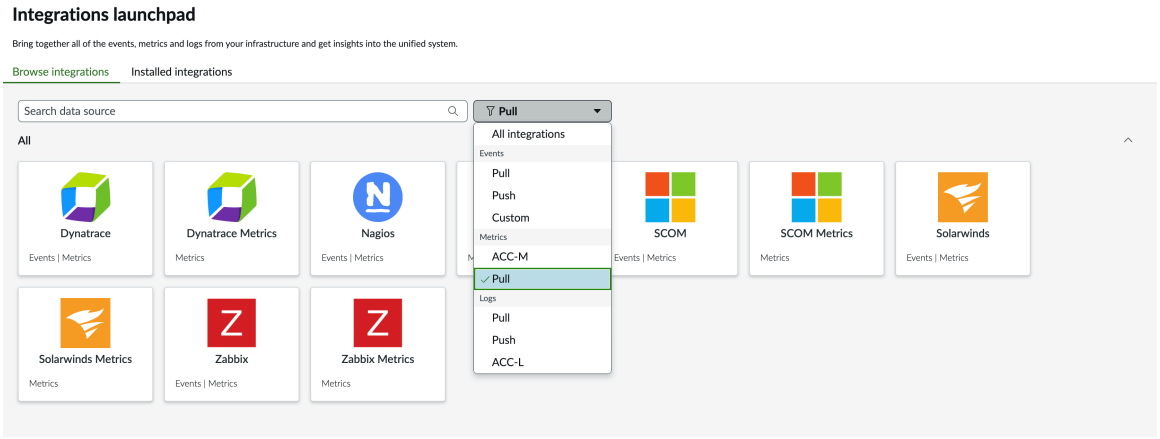

Only metric pull connector tiles are displayed.

5. Select a pull connector tile.
   If a pop-up menu opens, select the data to track from that connector and select **Continue**.

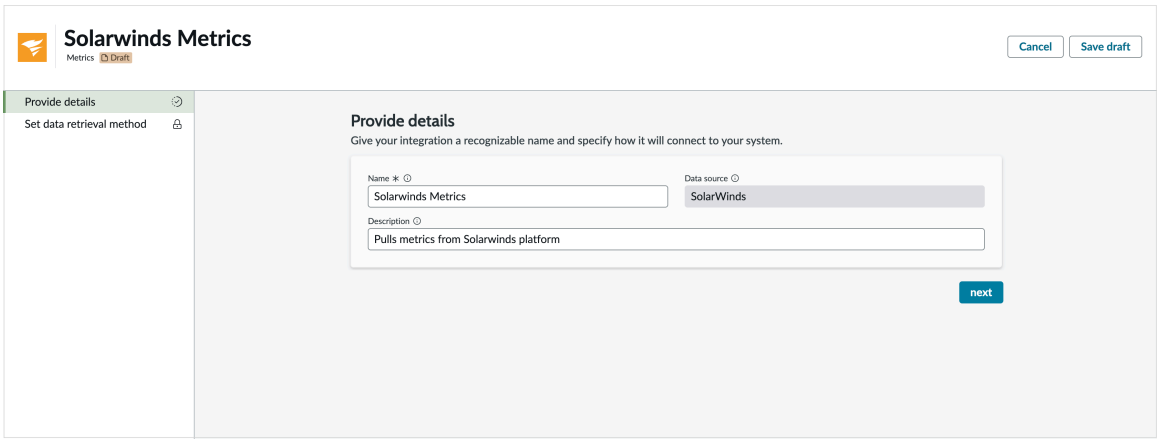

6. In the **Provide details** section, enter the following information and select **Next**.
   - In the **Name** field, enter a unique name for the connector type.
   - In the **Data source** field, the external source for retrieving metrics is auto-populated based on the tile type you select.
   - In the **Description** field, enter brief information about this connector.
7. In the **Set data retrieval method** section, enter the following information:
   - In the **Hostname/Host IP** field, enter the IP address that is related to the Solarwinds server.
   - From the **Credentials** drop-down list, select the valid credentials to access the event source host.
   - In the **Metrics collection schedule (seconds)** field, enter the time interval (in seconds) at which the connector retrieves metrics from the source host.
   - In the **MID Server** field, select a MID Server that collects metric data from external tools and securely sends it to ServiceNow.

     If none are available, ensure one is installed, running, and has the required capabilities.

     Note: You can select multiple MID Servers. The system randomly assigns one as the primary and uses the others as backups.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. **Optional:** If you want to customize configurations to meet specific requirements, select **Advanced settings** and fill in the form.

Field	Description
Max rows per query	Frequency in seconds that indicates how often the system should check for new metrics from the external event source. The value can't be lower than the minimum schedule property, which by default is 120 seconds.
Max fetch interval (in minutes)	Limits the maximum interval between event fetch operations.
Offset (in minutes)	Sets the time offset for event retrieval to ensure accurate synchronization.
Port	Specifies the port number used for communication with external devices. (Optional) Verify the default values for your connector type such as port details or time formats and modify them as necessary. Default values from your connector are filled-in.
Protocol	Defines the communication protocol used for retrieving events from external devices.
Debugging	○ Debug: Provides detailed logs for troubleshooting. ○ Log payload: Displays raw log data. Use only for debugging, as it can quickly fill the MID Server logs.

9. **Optional:** Test the connector before activating it by selecting **Test and Save**.
10. To save the connector, select **Save**.
11. To activate the pull connector, select **Activate**.
    The metric pull connector is activated and its tile is displayed in the **Installed Integrations** tab. When you open the tile, you can find the **Overview**, **Details**, and **Data retrieval method** tabs.

    The **Overview** tab allows you to verify that the connector is running and collecting data. It also offers options to improve system settings, such as refining alert definitions or enhancing binding configurations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

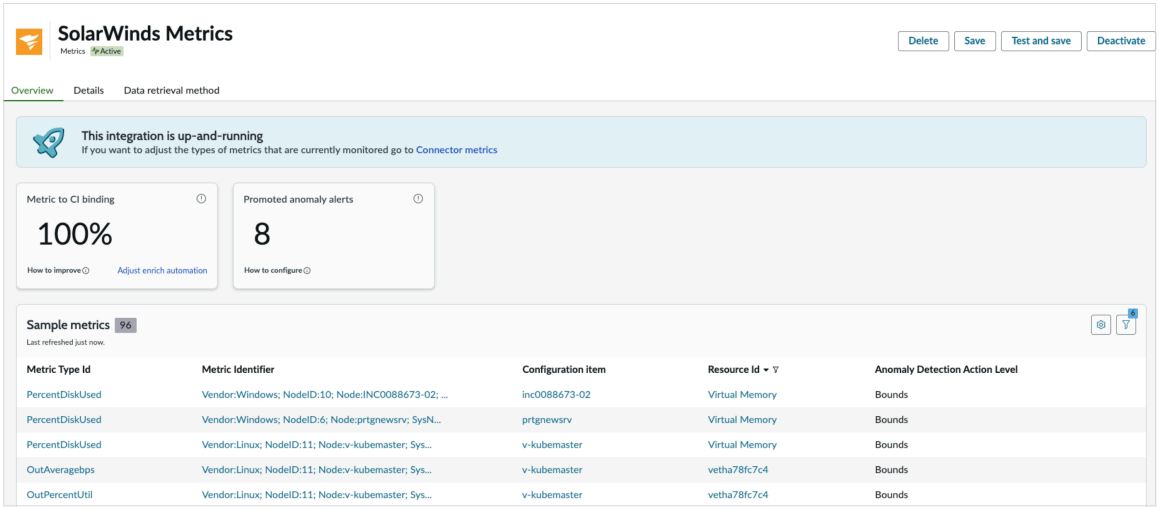

The **Details** tab allows you to update the connector's general information (such as name or description) and displays any errors encountered during its operation.

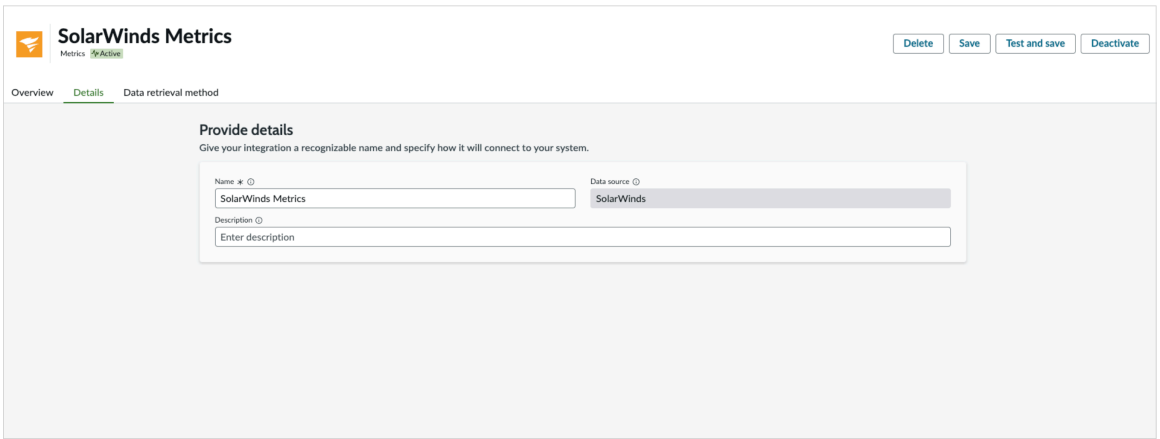

The **Data Retrieval Method** tab contains the connector settings (such as host, port, and credentials) and may provide additional details for issues like MID Server errors.

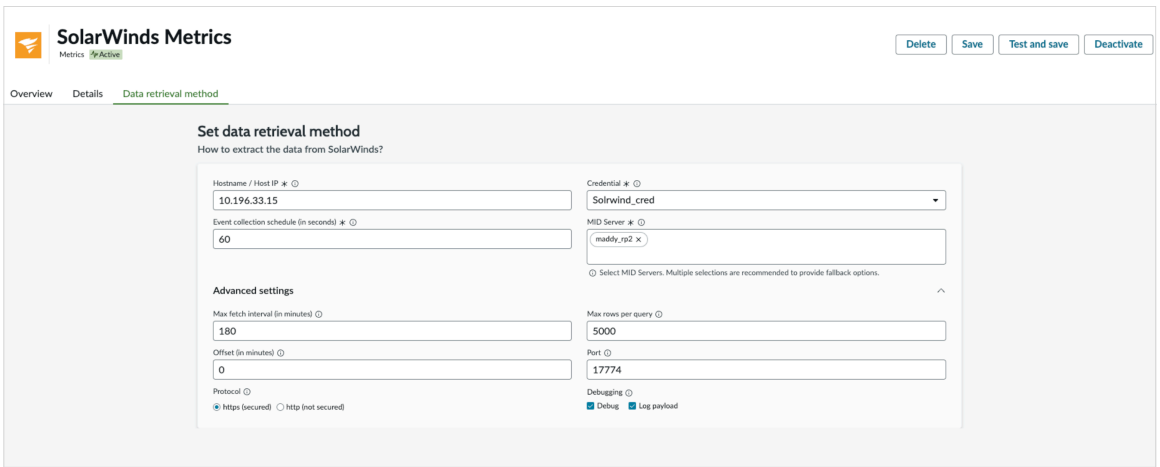

Configure Dynatrace metric pull connector

Configure metric pull connectors that require a script, connector definition, and connector instance to pull metrics from external sources. These connectors automate the data retrieval

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

process, ensuring the seamless integration of external metrics into your system for efficient monitoring and performance analysis.

Before you begin

Role required: evt_mgmt_admin

Configure the Dynatrace API Key credential from **All** > **Connections and Credentials** > **Credentials**. For more information on how to configure Dynatrace API Key credential, see API key credentials.

About this task

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon.
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. In the **Browse Integrations** tab, select the **All integrations** drop-down list and select **Metrics** > **Pull**.

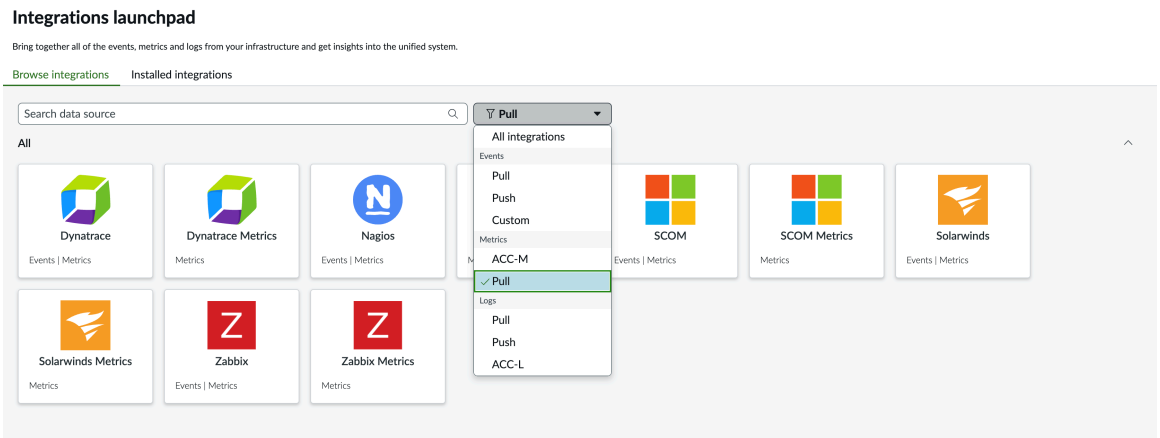

   Only metric pull connector tiles are displayed.
5. Select the **Dynatrace Metrics** connector tile.
6. In the **Provide details** section, enter the following information and select **Next**.
   - In the **Name** field, enter a unique name for the connector type.
   - In the **Data source** field, the external source for retrieving metrics is auto-populated based on the tile type you select.

     In this case, select **Dynatrace**.
   - In the **Description** field, enter brief information about this connector.
7. In the **Set data retrieval method** section, enter the following information:
   - In the **Hostname/Host IP** field, enter the IP address that is related to the Dynatrace server.
   - From the **Credentials** drop-down list, select the valid credentials to access the metrics source host.

     Select the credential.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- In the **Metric collection schedule (in seconds)** field, enter the time interval (in seconds) at which the connector retrieves metrics from the source host.

  The default value is 60 seconds.

- In the **MID Server** field, select a MID Server that collects metric data from external tools and securely sends it to ServiceNow.

  If none are available, ensure one is installed, running, and has the required capabilities. If none are available, ensure one is installed, running, and has the required capabilities. For example, the **Metric Intelligence Context** must be running on the same MID Server that you are configuring.

  Note: You can select multiple MID Server. The system randomly assigns one as the primary and uses the others as backups.

8. **Optional:** If you want to customize configurations to meet specific requirements, select **Advanced settings** and fill in the form.
   For more information on fields, see Dynatrace advanced settings fields.
9. **Optional:** Test the connector before activating it by selecting **Test and Save**.
10. To save the connector, select **Save**.
11. To activate the pull connector, select **Activate**.
    The metric pull connector is activated and its tile is displayed in the **Installed Integrations** tab. When you open the tile, you can find the **Overview**, **Details**, and **Data retrieval method** tabs.

    The **Overview** tab allows you to verify that the connector is running and collecting data. It also offers options to improve system settings, such as refining alert definitions or enhancing binding configurations.

    The **Details** tab allows you to update the connector's general information (such as name or description) and displays any errors encountered during its operation.

    The **Data Retrieval Method** tab contains the connector settings (such as host, port, and credentials) and may provide additional details for issues like MID Server errors.

Configure Datadog metric pull connector

Configure OAuth-based integration between your Datadog instance and ServiceNow Event Management to securely ingest metric from the Datadog Instance. This setup enables authenticated metric delivery using client credentials, maps discovered server CIs, and allows anomaly alerts to create actionable events in ServiceNow through standard webhooks.

Before you begin

Configure the Datadog API Key credential from **All** > **Connections and Credentials** > **Credentials**. For more information on how to configure Datadog API Key credential, see API key credentials.

Ensure you have the application key configured in the Datadog. For more information, see https://docs.datadoghq.com/account_management/api-app-keys/.

To use container metrics: First, discover your Docker Container CIs (through Discovery, Cloud Discovery, or Agent Client Collector) so that the container_id field on each cmdb_ci_docker_container record is populated. The shipped Datadog Metrics Container Binding event rule matches container metric events to these CIs based on container_id. Datadog reports the container_id as the container's short ID (the first 12 characters of the Docker container

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ID), which is what Discovery stores on the CI. Once the CIs are discovered, enable the Enable container metrics advanced setting (set it to true).

Role required: evt_mgmt_admin

About this task

The Datadog connector can collect Docker container metrics and bind them to Docker Container CIs (cmdb_ci_docker_container). This capability is off by default and is enabled through the **Enable container metrics** advanced setting.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. In the **Browse Integrations** tab, select the **All integrations** drop-down list and select **Metrics** > **Pull**.

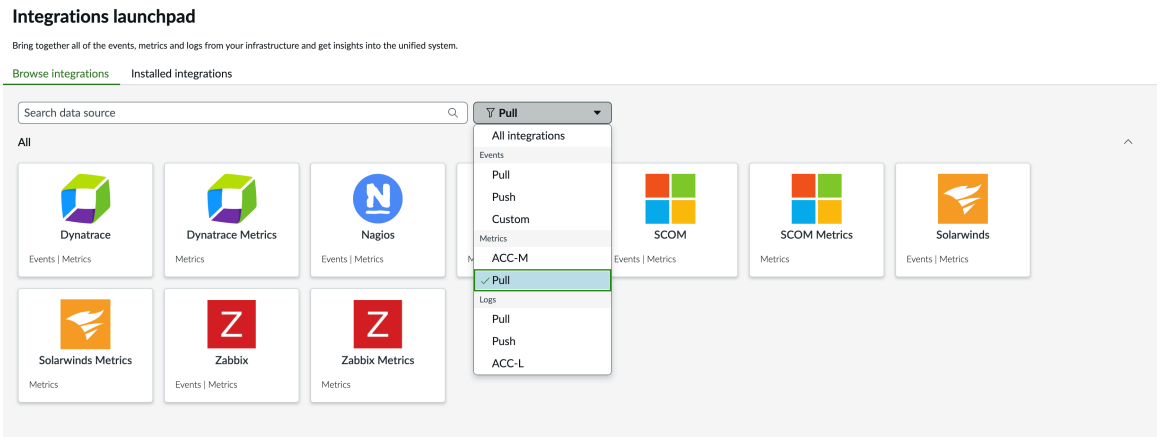

   Only metric pull connector tiles are displayed.
5. Select the **Datadog Metrics** connector tile.
6. In the **Provide details** section, enter the following information and select **Next**.
   - In the **Name** field, enter a unique name for the connector type.
   - In the **Data source** field, the external source for retrieving metrics is auto-populated based on the tile type you select.

     In this case, select **Datadog**.
   - In the **Description** field, enter brief information about this connector.
7. In the **Set data retrieval method** section, enter the following information:
   - In the **Hostname/Host IP** field, enter the IP address that is related to the Datadog server.
   - From the **Credentials** drop-down list, select the valid credentials to access the metrics source host.

     Select the credential.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- In the **Metric collection schedule (in seconds)** field, enter the time interval (in seconds) at which the connector retrieves metrics from the source host.

  The default value is 60 seconds.

- In the **MID Server** field, select a MID Server that collects metric data from external tools and securely sends it to ServiceNow.

  If none are available, ensure one is installed, running, and has the required capabilities. If none are available, ensure one is installed, running, and has the required capabilities. For example, the **Metric Intelligence Context** must be running on the same MID Server that you are configuring.

  Note: You can select multiple MID Server. The system randomly assigns one as the primary and uses the others as backups.

- in the **Application key** field, enter the application key that is associated with the user account.

8. **Optional:** If you want to customize configurations to meet specific requirements, select **Advanced settings** and fill in the form.
   For more information on fields, see Datadog advanced settings fields.
9. **Optional:** Test the connector before activating it by selecting **Test and Save**.
10. To save the connector, select **Save**.
11. To activate the pull connector, select **Activate**.
    The metric pull connector is activated and its tile is displayed in the **Installed Integrations** tab. When you open the tile, you can find the **Overview**, **Details**, and **Data retrieval method** tabs.

    The **Overview** tab allows you to verify that the connector is running and collecting data. It also offers options to improve system settings, such as refining alert definitions or enhancing binding configurations.

    The **Details** tab allows you to update the connector's general information (such as name or description) and displays any errors encountered during its operation.

    The **Data Retrieval Method** tab contains the connector settings (such as host, port, and credentials) and may provide additional details for issues like MID Server errors.

Configure Azure Monitor Issue in Integrations Launchpad

An Azure Issue is a unified case that aggregates related alerts and signals from Azure Monitor into a single, trackable operational problem. It preserves investigation context and serves as a durable record for incident management, so you can resolve problems faster instead of tracking scattered alerts.

Before you begin

Verify you already have an Azure Monitor Workspace.

Role required: evt_mgmt_admin

About this task

ServiceNow represents an Azure Issue as a single alert. View the alerts that Azure grouped into the Issue in the **Related Records** tab from the Express List side panel or the alert record in Service Operations Workspace. ServiceNow retains each related alert as an individual alert. This approach lets ServiceNow group and correlate alerts across multiple resources and monitoring systems. ServiceNow is not limited to the alerts that Azure grouped into the Issue. As

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a result, ServiceNow identifies broader incidents, reduces duplicate alerts, and provides a more comprehensive operational view. For this reason, ServiceNow represents an Azure Issue as a single alert instead of an alert group.

Note the following before you start:

- If you already use the Azure push connector to send Microsoft Azure alerts to ServiceNow, use the same push connector URL to send Azure Issues.
- If you have different tenants with different Azure Service Principal credentials (per tenant or per subscription), create separate Azure push connector instances. Use them to push Azure alerts or issues into your ServiceNow instances.

Procedure

1. If you don't already have a Azure monitor alert integration then create a push connector instance by perform the following steps:
   a. Navigate to **Workspaces** > **Service Operations Workspace**.
   b. From the bottom of the navigation pane, select the AIOps Configuration Center icon .

      The ITOM AIOps Configuration Center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
   c. On the ITOM AIOps Configuration Center page, under the Setup Integrations section, select **Add integrations**.

      The Integrations Launchpad page opens.
   d. Search for **Microsoft Azure to instance** and create a push connector instance by providing the following details:
      i. In the Provide details page, provide the connector name and description for the connector in the **Connector name** and **Description** field.
      ii. Select **Next**.

         The Set-up push connector page opens.
      iii. From the **URL parameter value** field, copy the URL by selecting the **Copy link to clipboard**.

         The URL can be used in the Microsoft Azure Action group Securewebhook URL to send Azure alerts and Azure Issues into ServiceNow instance. The URL must be in the following format:

         ```
         https://<instance_name>.service-now.com/api/sn_em_connector/em/inbound_event?source=azuremonitor&sys_id=<sys_id_push_connector_instance>
         ```
      iv. Select **Activate** to activate the push connector instance.
2. In the Azure Monitor portal, perform the following steps:
   a. Navigate to your Azure Monitor Workspace.
   b. Navigate to **Settings** > **Action groups** and select the existing action group that you use to send Azure alerts to ServiceNow.

      This creates an alert, and in turn, the Alert Management rule creates an incident. If you're configuring Azure issues or alerts for the first time, first create an action group by using the above URL you have copied. Then follow the steps in Integrate Azure Monitor with OAuth authentication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you're configuring Azure issues or alerts for the first time, first create an action group by using the above URL you have copied. Then follow the steps in Integrate Azure Monitor with OAuth authentication.

c. The pull connector shows related alerts and resource information for an issue in the Azure portal under the **Alerts** and **Resources** tab, and enables bi-directional synchronization. To configure it, perform the steps below.

3. Verify that the pull connector instance is active by performing the following step:

> **Note:** If you used the push connector to bring Azure alerts or issues into ServiceNow, the pull connector instance may already be active.

In Integration Launchpad, on the **Installed Integrations** tab, search for `Azure Monitor` and confirm that it is already active.

> **Note:** If you have already configured a pull connector instance for Azure Alerts bi-directional sync, it also works for Azure Issues bi-directional sync and for fetching Azure Issue-related alerts.

4. If you used the push connector to bring Azure alerts or issues into ServiceNow but the Azure Monitor pull connector is not yet active, perform the following steps:
   a. On the **Installed Integrations** tab, search for `Microsoft Azure` and select the **Microsoft Azure (Events – pull)** tile to open the pull connector instance.
   b. In the **Credentials** field, enter your Azure Service Principal credential.
   c. Leave the **Host** field as `1.1.1.1`, because it is not used.
   d. To use a specific MID Server, go to **Advanced Settings** and select the mid server in the **Connected mid server** field.
   e. Select **Test and Save**.

      If the test connection succeeds, activate the connector.

      If Azure bi-directional connector is enabled, when an Azure issue is closed in ServiceNow, it automatically closes in the Azure portal as well.

5. If you used the push connector to bring Azure alerts or issues into ServiceNow, configure a pull connector instance by performing the following steps:
   a. On the **Browse Integrations** tab, search for `Microsoft Azure` and select the **Microsoft Azure (Events – pull)** tile to create a new pull connector instance.
   b. On the **Details** tab, provide the following information:

      i. In the **Connector name** field, enter a name for the pull connector instance.

      > **Note:** The pull connector instance name must match with the push connector instance name.

      ii. In the **Credentials** field, enter your Azure Service Principal credential.
      iii. Leave the **Host** field as `1.1.1.1`, because it is not used.
      iv. To use a specific MID Server, go to **Advanced Settings** and enter the mid server name in the **Connected mid server** field.
      v. Configure the following parameters to control how the system fetches related alerts for an issue:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Set *fetchIssueRelatedAlerts* to control whether the system fetches related alerts for the issue.

  The default is true. Set it to false if you want to disable this behavior.

- Set *fetchIssueRelatedAlertsLimit* to the maximum number of related alerts the system fetches for the issue. The default is 10.

vi. Select **Test and Save**.

If the test connection succeeds, activate the connector.

Review integration and configuration health

Review your events to ensure alerts are created as intended. Start this process by checking the health status of your integration. This provides useful statistics to highlight where you can make improvements. You can enrich alerts with relevant business context, group them for better noise reduction, escalate critical alerts promptly, and notify your team for swift response and resolution.

Before you begin

Role required: evt_mgmt_operator or evt_mgmt_admin

About this task

Health state metrics indicate how effectively alert automations add context and reduce noise, helping operators lower the Mean Time to Repair (MTTR). When adding a new integration, review these metrics, assess the rates, and develop new rules to further enhance noise reduction and improve overall system performance.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. In the **Browse Integrations** tab, select the **All integrations** drop-down list and select **Pull** or **Push** as required.
5. Select a pull, push, or a custom connector type tile.
   If a pop-up menu opens, select the data to track from that connector and select **Continue**.
6. Select the **Events** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

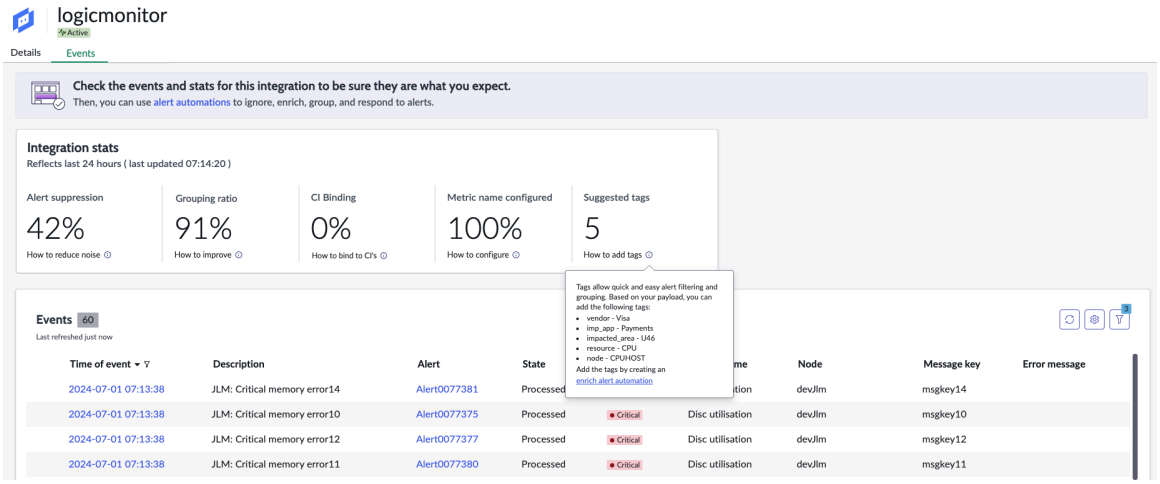

7. In the Health state section, view the data for alerts and their associated attributes.

Alert attributes

Alert attribute	Description
Alert Suppression	The suppression ratio indicates the proportion of alerts created to events received. Reducing unwanted alerts can be achieved by de-duplicating or ignoring them. ○ De-duplication: The **message_key** field is a unique identifier for an alert in the source monitoring system. It is used for de-duplication and alert flapping detection purposes. If the source event does not provide reliable unique values, you can use the default value or create your own. If the message key is not defined, the default message key will be <Source + Node + Type + Resource + Metric Name>. The guideline is for the event source to automatically populate the <Source + Node + Type + Resource + Metric Name> fields and generate the message key. This practice facilitates better distribution of event processing across instance workers and nodes. If the source event does not provide values for these fields, ensure they are populated using transform rules. This action, which does not impact event processing, is essential for de-duplication. Populate as many of these fields as possible before sending the event to the instance. Doing so enhances the distribution of events across processor workers, resulting in improved throughput and scalability. To learn more about de-duplication, see Event Management configuration preferences. To set the message key in your custom connector integration, see Configure an event custom connector. ○ Ignore alert event rules: Filter out low-quality alerts by creating an ignore automation.
Alert compression rate	When you group alerts, a high compression rate is desirable because it means fewer alerts remain for your operations team to review and

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert attribute	Description
	act on. This can lead to improved team efficiency. The compression rate is calculated as 1 – (the number of alert groups + the number of ungrouped alerts) / total alerts. To learn about alert aggregation, see Alert grouping. One way to improve your alert compression rate is by defining a new grouping automation.
CI Binding	The percentage of alerts successfully bound to a Configuration Item (CI) in your CMDB is crucial. Aim for this number to be at or near 100% to enhance alert grouping, root cause analysis, assignment, and more. To bind a host, machine, or any device with an IP, populate the **Node** field with a unique hostname, FQDN, IP, or MAC address. This can be done within the integration itself or by creating an enrich automation. For more complex scenarios, you can override node-based binding using event rules. For more information on CI binding, see Binding alerts to CIs.
Metric name configured	Shows the percentage of alerts that have a metric name defined. It's crucial to aim for near 100% because it's used by default in our automated correlation algorithm. This can be achieved within the integration itself or by creating an enrich automation.
Suggested tags	Provides suggested alert tags based on the integration payload. Alert tags are structured fields that you can define to facilitate easier filtering and grouping of alerts. The tags are stored in *sn_itom_integ_app_recommendation_tags_suggestion*, and only the top 5 most popular tags are displayed. To refine the results of the suggestion mechanism, you can use the *sn_itom_integ_app.evt_mgmt_tag_suggestion_ignore_keys* property. This property displays a list of lowercase field names or keywords, separated by commas, to be ignored when scanning additional alert info for tag suggestions. The default value is time.

8. In the Events section, ensure that your events are generating alerts as expected.

   Verify that alert fields and tags are populated correctly to make the alerts easier to read and filter in Express List. Confirm that severity values align with your operations team's expectations and that fields and tags follow a standardized format across sources for easier grouping.

   By verifying alert generation, you can be confident that your system is properly configured to detect and notify you of noteworthy occurrences. For information on the fields in the Events section, see Integration and configuration health event fields.

What to do next

Now that you've confirmed alerts from the events, you can initiate the alert automation process by selecting **Add automation**. This takes you to the ignore automation step, where you can disregard irrelevant or false-positive alerts. Following that, proceed to enrich, group, and automate responses. This process involves transforming raw alerts into a format understandable

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

by ServiceNow, grouping the alerts, and escalating an alert. This ensures an immediate response from the appropriate team or individual. For more information on alert automation, see Alert automation in Service Operations Workspace for ITOM.

Activate integration

Activate integration allows you to enable a previously deactivated connector integration, restoring its functionality. This feature helps you seamlessly resume data flow and integration activities, ensuring continuity and efficient event management without the need to reconfigure the integration.

Before you begin

Ensure the integration is deactivated.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. Select the **Installed integrations** tab.
5. For the integration that you want to activate, select the Ellipses icon (•••) and then select **Activate**.

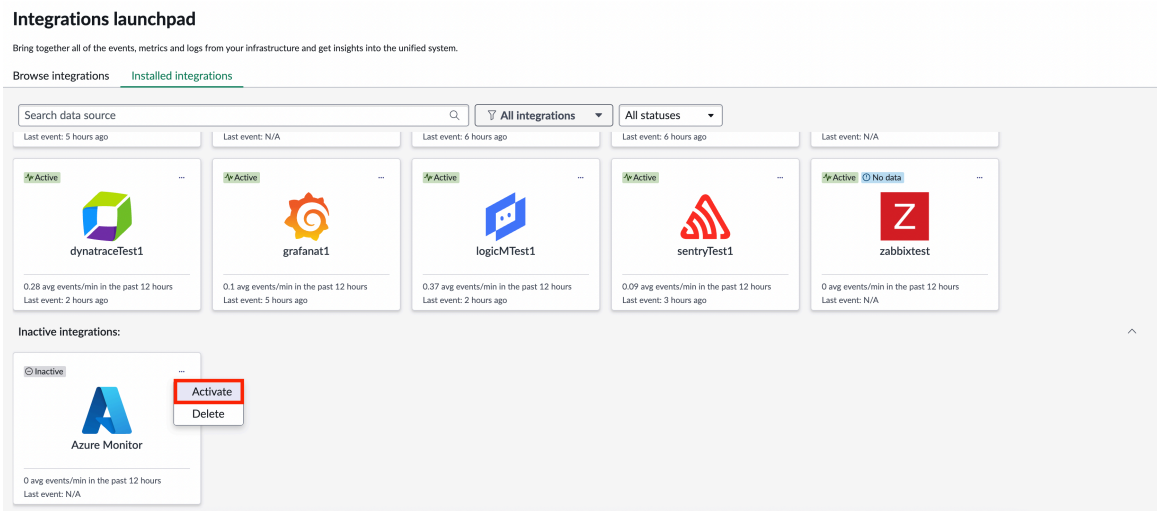

The selected integration is activated, and its status is shown as **Active** on the installation tile.

Deactivate integration

Deactivate integration allows you to temporarily disable an active connector integration without deleting it. This feature helps you manage integrations more effectively by pausing data flow when needed, reducing unnecessary event processing and improving system performance.

Before you begin

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. Select the **Installed integrations** tab.
5. For the integration that you want to deactivate, select the Ellipses icon (•••) and then select **Deactivate**.

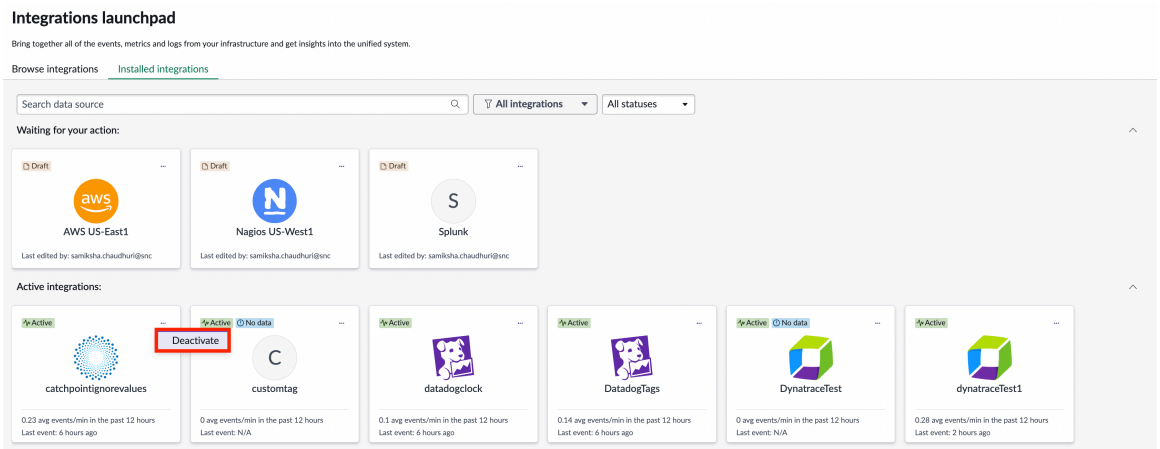

The selected integration is deactivated, and its status is shown as **Inactive** on the installation tile.

Delete integration

Delete integration allows you to permanently remove a connector integration from your system. This feature helps you maintain a clean and organized system by eliminating unused or obsolete integrations, freeing up resources and reducing potential clutter.

Before you begin

Ensure the integration is deactivated.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. From the navigation pane, select the AIOps configuration center icon .
   The ITOM AIOps configuration center page appears. The configuration center is a centralized workspace. Use it to configure and manage AIOps features from a single place.
3. On the ITOM AIOps configuration center page, under the **Setup** > **Integrations** section, select **Manage Installed integrations**.
4. Select the **Installed integrations** tab.
5. For the integration that you want to delete, select the Ellipses icon (•••) and then select **Delete**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

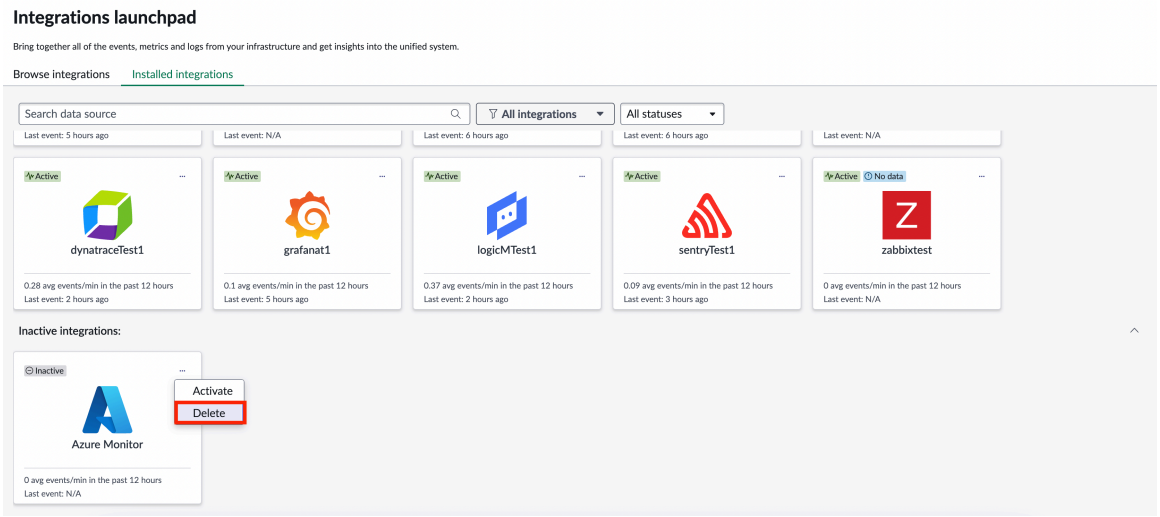

6. On the confirmation window, select **Delete**.
   The integration is deleted.

Service Level Objective Management

Service Level Objective Management (SLO Management) helps your organization define, track, and meet agreed-upon service quality standards. It works alongside Service Level Agreements (SLAs) to help make sure that services meet customer expectations.

Overview of how SLO Management works

- SLAs set the expectations

  SLAs are formal agreements that outline the level of service a customer can expect. SLAs can include response times, resolution times, and uptime targets.

- SLOs defines how to meet expectations

  SLOs are specific, measurable objectives that an organization sets to achieve the SLAs. They translate the broad goals of the SLA into actionable metrics.

- ServiceNow SLO Management helps manage SLOs. It offers functionalities to:
  - Define SLOs based on various metrics.
  - Monitor SLO performance in real-time.
  - Generate alerts when SLOs are at risk of being breached.
  - Track historical data to identify trends and improve processes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore Learn about Service Level Objective Management concepts and features.	Use Using Service Level Objective Management
Reference Get details about Service Level Objective Management components, such as fields, tables, and properties.	

Troubleshoot and get help

- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring SLO Management

Service Level Objective Management (SLO Management) helps IT services meet customer expectations.

SLO Management overview

SLO Management is a framework for setting clear expectations and measuring the performance of IT services. It helps organizations deliver consistent services and identify areas for improvement. SLOs define the target service level for a specific service, such as incident resolution or service request fulfillment. Effective SLO management involves setting realistic objectives, monitoring performance, and continuously improving services to meet customer needs.

SLO Management users

Users

Users	Description	Contains Roles
admin	A ServiceNow administrator is responsible for the administration,	All

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Users (continued)

Users	Description	Contains Roles
	development, operation, education, and maintenance of the ServiceNow platform. Responsible for installation and can perform Service Operations Workspace Admin Center configuration of SRM.	
Administrator [srm_admin] **Note:** Not the ServiceNow *admin* role	SRM administrators can manage account settings, configurations, and users. Administrators can perform the following actions: • Access, create, edit, or delete all SRM configurations. • Add or manage integrations. • Create Integrations with Application Performance Monitoring (APM) tools • Set up and maintain Reliability Indicators. • Set up and maintain Error Budget Policies.	• Manager • Responder
Manager [srm_manager]	Managers oversee a team of service reliability engineers (SREs). Managers assign SREs to the team on-call schedule, monitor their performance, create procedures to deal with incidents, and develop solutions. Managers help ensure resilience across all the systems and the DevOps workflows. Managers can perform the following actions within the context of their teams: • Define and set up and teams, on-call schedules, and services. • Add and delete users such as responders, and managers for the teams the are a part of. • Add or manage integrations. • Create Integrations with Application Performance Monitoring (APM) tools	Responder

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Users (continued)

Users	Description	Contains Roles
	• Set up and maintain Reliability Indicators. • Set up and maintain Error Budget Policies.	
Responder [srm_responder]	An SRE that uses SRM to perform everyday tasks. Responders are the individuals who are on call and diagnose and remediate incidents. Responders can only access configurations that they're a part of. They can only access the alerts or incidents for which they have permissions. SREs can perform the following actions, within the context of their teams: • Set up services, teams, and integrations • Confirm their on-call schedules • Manage incident and alert records • Update teams that they've created • Add other responders • Create Integrations with Application Performance Monitoring (APM) tools • Set up and maintain reliability metrics • Set up and maintain error budget actions	Inherits 17 roles including the following: • cmdb_read • sn_sow.sow_user • sn_sow_srm.srm_responder • workspace_user • slo_operator

Note: If you're using the SLO plugin without SRM, make sure that the required alert and cmdb_ci read access is granted to team members.

SLO Management workflow

1. Define SLOs - Identify critical services and define SLOs based on customer expectations and business requirements.

   Note: SLOs can also be generated automatically using the SLO creator agent. For details, see Generating service level objectives.

2. Establish SLIs - Develop Service Level Indicators (SLIs) to measure SLO performance.
3. Monitor and analyze - Track SLI data and analyze performance against SLO targets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Identify gaps and improve - Determine areas where SLOs aren't being met and implement changes to improve service performance.
5. Review and refine - Regularly review SLO performance and refine SLOs as needed.

SLO Management benefits

- Improved service quality - SLOs help IT services meet customer expectations, leading to increased satisfaction and loyalty.
- Increased transparency - Clear SLOs provide a shared understanding of service expectations between IT and customers.
- Better resource allocation - SLOs help prioritize resources and focus on areas that need improvement.
- Enhanced collaboration - SLO management encourages collaboration between IT teams and customers to achieve common goals.
- Data-driven decision making - SLO performance data informs decisions and drives continuous service improvement.
- Automated SLO generation - The SLO creator agent analyzes operational data to generate SLOs, helping teams get started with SLO-based monitoring. For details, see Generating service level objectives.

Reliability metrics in SLO Management

Learn about the reliability metrics and features that can help you track service health, respond to issues, and support business goals.

Service reliability dashboard

The Service reliability dashboard displays a customizable, high-level view of service performance. It helps you monitor and manage reliability using visualizations that track service states, error budgets, and service level objectives (SLOs) over time.

The dashboard displays information about all services in Service Reliability Management (SRM). You can access the dashboard in Service Operations Workspace in the following ways:

- Navigate to **Services (** **)** > **Service reliability**.
- Navigate to **Home (** **)** > **Service reliability**

For more details, see Visualizations in the Service reliability dashboard.

Note: You can also view SLO information for all services on the Services Overview tab. See Working with SRM services for more information.

Notification destinations

Notification destinations help keep teams informed about service reliability. Attach them to error budget policies to send notifications when a policy is breached.

To view and manage notification destinations in Service Operations Workspace, navigate to **Teams** > **[Your team]** > **SLO Notification destinations**.

Visit the following links to learn more about creating and working with notification destinations:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Create a notification destination in SRM
- Create SLOs, SLIs, and error budget policies
- Notifications for breached error budgets

Reliability metrics tab

The Reliability metrics tab shows how well a specific service is meeting its reliability goals. Use it to track SLOs, service level indicators (SLIs), and error budgets for a service.

To view the Reliability metrics tab in Service Operations Workspace, navigate to **Services ( )** > **[Your service]** > **Reliability Metrics**.

See these links to learn more about what you can do in the Reliability metrics tab:

- Create SLOs, SLIs, and error budget policies
- Edit a reliability metric
- View SRM reliability metrics

Service level objectives table

On the Reliability metrics tab, the Service level objectives table includes the following details about the service SLOs:

- **Name** - Name of the SLO. You can select the name to view the SLO record.
- **Reliability** - Current state of the SLO. For example, stable, at risk, or critical.
- **% Error budget remaining** - Percentage of the error budget still available in the current compliance period.
- **Compliance period** - Time window used to calculate performance:
  - Month - Current month, for example, if the current date is January 26, the month is January 1 through January 31.
  - Rolling 7, 30, or 90 days - Number of days from the current date. For example, for rolling 7 days, the duration is 7 days back from the current date.
- **SLI type** - Performance category being measured:
  - Availability - Percentage of time your service or configuration item is available, also known as uptime.
  - Errors - Frequency of your service errors.
  - Latency - Time that it takes to service a request.
  - Saturation - Fullness of your system, focusing on resource usage.
- **Source type** - Origin of the data used to calculate the SLIs for this SLO:
  - Alert - Uses alerts from integrated monitoring tools.
  - Outage - Uses outages detected by monitoring tools and reported by users. An outage indicates when the service was unavailable. This source type excludes planned outages, such as scheduled maintenance.
- **Updated** - The date and time the SLO was last edited.
- **Updated by** - The user name of the person who last edited the SLO.
- **State** - Status of the SLO. For example, running or retired.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: For performance purposes, SLO and SLI records ([sn_sow_srm_slo_history] and [sn_sow_srm_sli_metric]) are archived after one year and deleted five years later. Archived data is omitted from tables and visualizations.

AI in SLO Management

AI capabilities in SLO Management help teams adopt service level objectives (SLOs) faster and monitor service performance.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

AI features in SLO Management

AI feature	Description	Use case	Resources
SLO generation with the SLO creator agent	Analyzes CI data and automatically generates baseline SLOs for eligible configuration items (CIs). You can review, update, and manage the generated SLOs and adjust the agent settings.	An operator needs baseline SLOs for new or existing CIs without manually defining thresholds.	• Generating service level objectives • Manage SLO creator agent settings • View generated SLOs • SLO creator agent system properties
Conversational analysis and SLO creation with the ITOM MCP Server	Connects an MCP client application, such as Claude or Moveworks, to your ServiceNow ITOM environment. You can have natural language conversations about CI reliability, topology, and incident impact, and create SLOs through the MCP client.	An operator wants to investigate CI reliability and topology, assess incident impact, or create SLOs through an MCP client.	• Using the ITOM MCP Server Console to perform ITOM tasks • Activate the ITOM MCP Server Console • Review CI reliability with an MCP Client

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generating service level objectives

Use the service level objective (SLO) creator agent to auto-generate SLOs for services and configuration items (CIs). The agent can help teams adopt SLOs faster and monitor service performance.

How the SLO creator agent works

The SLO creator agent requires ServiceNow Otto for ITOM and Service Reliability Management (SRM). It automatically generates SLOs by analyzing operational data for services and CIs.

1. Every 14 days, the agent processes up to 25 service or CIs that don't already have an auto-generated SLO.
2. For each service or CI, the agent analyzes critical alerts, incidents, and outages, including data from related CIs.
3. The agent uses existing availability commitments, when available, or business criticality to determine the SLO target.
4. For every service or CI that it analyzes, the agent generates one SLO with up to 10 SLIs.
5. The relevant team manager receives an email when the SLO is generated.
6. The generated SLO is active by default and the assigned team receives an email when the error budget falls to 25% or below.

   SRM managers and responders can review, edit, or turn off generated SLOs.

You can manage the SLO creator agent, including turning the agent on or off and customizing system properties. For more information, see Manage SLO creator agent settings and SLO creator agent system properties. To view SLOs generated by the agent and review details about their creation, see View generated SLOs.

Important: This AI agent is turned on by default. For more information, see AI agents, skills, and agentic workflows on by default.

Access the SLO creator agent in AI Agent Studio

Access the SLO creator agent in AI Agent Studio to manage security controls or duplicate it for further customization.

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select the **AI agents** tab.
3. Select **SLO Creator Agent**.

The SLO Creator Agent page lets you manage the agent, including defining security controls. To modify the agent, duplicate it and adjust the settings.

Important: When you modify an agentic workflow, AI agent, or tool, make sure that you update all instructions accordingly.

Configuring SLO Management

Configure SLO Management to help teams monitor and improve service reliability.

Manage SLO creator agent settings

Activate or deactivate the service level objective (SLO) creator agent and configure the notification settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

You must have ServiceNow Otto for ITOM and Service Reliability Management (SRM) installed on your instance. For more information, see Install plugins for ServiceNow Otto and Install Service Reliability Management.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with your applications have specific roles that come included with the application. If you select *Users with specific roles* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

For an overview of how and when the SLO creator agent generates SLOs, see Generating service level objectives.

Role required: srm_admin

Procedure

1. Navigate to **All** > *Service Operations Workspace* > **Configurations**.
2. Select **Service Reliability Management** > **Activate** > **AI-generated SLOs**.
3. Configure the SLO creator agent settings.

Setting	Description
Allow the agent to automatically generate SLOs from service alerts and outages.	Turns on or off the SLO creator agent. When the agent is active, it generates SLOs based on alerts, incidents, and outages for eligible services and configuration items (CIs).
Include a policy to email the assigned team if the error budget is less than or equal to 25%.	Sends an email notification to the team assigned to the service or CI when the SLO error budget reaches 25% or lower.
Email team managers when an SLO is generated, or the team if no manager is assigned.	Sends an email notification when the agent generates a new SLO. If a team manager is assigned to the service or CI, the manager receives the notification. Otherwise, the notification is sent to the group email. If no group email exists, all group members with an email address receive the notification.

4. Select **Save** to implement your changes.
   The UI displays `Settings saved`.

What to do next

To view generated SLOs and learn how to modify them, see View generated SLOs and Edit a reliability metric. To configure advanced agent behavior, see SLO creator agent system properties.

Using SLO Management

Using SLO Management ensures IT services meet customer expectations, improves service quality and transparency, and enables data-driven decision making.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View generated SLOs

View service level objectives (SLOs) generated by the SLO creator agent and review information about how they were created.

Before you begin

You must have ServiceNow Otto for ITOM and Service Reliability Management (SRM) installed on your instance. For more information, see Install plugins for ServiceNow Otto and Install Service Reliability Management.

For an overview of how and when the SLO creator agent generates SLOs, see Generating service level objectives.

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`* > **Services**.
2. Select a service and then select the **Reliability Metrics** tab.
3. Select an SLO generated by the SLO creator agent.
   SLOs generated by the agent display `AI` in the Created by column.
4. Select the **Details** tab.
5. In the Activity panel, review the work notes.
   Generated SLOs have work notes that summarize the data and configuration used to create them.

What to do next

To learn more about SLOs, including SLO types and compliance periods, see Reliability metrics in SLO Management. To edit or deactivate an SLO in SRM, see Edit a reliability metric.

Activate SLO-based change approval decisions

Activate Service Level Objective (SLO) change approval decisions to evaluate change requests based on SLO performance.

Before you begin

Review change approval policies and their terminology, including policy inputs, decisions, and approval definitions. For more information, see Change approval policies.

Role required: admin

About this task

Activate the Service Level Objective Change policy's decisions to include SLO performance in your change management process. The policy is active by default, however, its decisions are inactive.

When you activate the policy's decisions, the policy evaluates SLOs as inputs and implements the following:

- Require group manager approval if these conditions are true:
  - The lowest SLO error budget is less than 25%.
  - The highest SLO burn rate is greater than 1%.
- Auto-approve change requests if these conditions are true:
  - The lowest SLO error budget is greater than 25%.
  - The highest SLO burn rate is less than 1%.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Change** > **Change Policy** > **Change Approval Policies**.
2. Select the **Service Level Objective Change policy**.
3. In the **Decisions** tab, set **Active** to true for these decisions:
   - **Budget < 25% and Burn Rate > 1**
   - **Budget > 25% and Burn rate < 1**
4. Select **Update** to save your changes.
   The decisions are active and part of your change management process.

What to do next

Use a test change request to verify that change requests are either auto-approved or routed to a manager as expected.

You can also customize the Service Level Objective Change policy, including its name and conditions, as needed.

For more advanced setups, use Flow Designer to customize the SLOs and services included as policy inputs:

Important: Customizing the flow prevents it from receiving future release updates. That behavior is intentional to avoid overwriting your customizations.

1. Navigate to **Process Automation** > **Flow Designer**.
2. In the list of flows, select **Change - SLO – Authorize** and then select **SLO - Get worst performing SLOs**.
3. Customize the **Action Inputs** section.
4. Select **Done** to implement your changes.

Create SLOs, SLIs, and error budget policies

Define service level objectives (SLOs), service level indicators (SLIs), and error budget policies to monitor service health. These tools help you and your teams track performance and respond when needed.

Before you begin

To send notifications when an error budget policy is breached, your team must have a notification destination configured. You can check by navigating to **Teams** > **[Your team]** > **SLO Notification destinations** in Service Operations Workspace. For more information, see Create a notification destination in SRM.

Note: By default, the outage data source for SLIs includes unplanned and degradation outages. Administrators can include planned outages by adding `planned` to the `sn_sow_slo.outage_supported_types` system property value.

Role required: srm_admin, srm_manager, or srm_responder

About this task

Create SLOs, SLIs, and error budget policies to monitor service health and help make sure that your services support business goals.

In Service Reliability Management (SRM), an SLO must have at least one SLI. SLIs can be filtered to a service or configuration item (CI) within the parent service hierarchy. Filtering an SLI to a CI can help you more accurately track service health and identify root causes faster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Each SLI can be associated with only one SLO, helping prevent conflicting performance expectations.

Note: SLOs can also be generated automatically using the SLO creator agent. For details, see Generating service level objectives.

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace* > **Services ( )**.
2. On the **Services** page, select the service you want to monitor, then select the **Reliability metrics** tab and select **Add SLO & SLI**.
3. In the Service Level Objective (SLO) form, fill in the fields for your SLO and select **Next**.
   For information about the SLO fields, see Create SLO form.
4. In the Service Level Indicators (SLI) form, select the data source for your SLI.

Data source	Description
Alerts	Base the SLI on alerts from integrated monitoring tools. Use alerts when you want your SLI to include early warning signs and a broader view of reliability.
Outages	Base the SLI on outages detected by monitoring tools and reported by users. Use outages when you want to measure actual downtime and customer impact.

Note: If you create multiple SLIs for an SLO, they must use the same data source.

5. In the same form, select **Add indicators for <data source>**, fill in the fields for your SLI, and then select **Next**.
   For more information about the SLI fields, see Create SLI form.
6. **Optional:** In the Error budget policies form, select **Add policy**, and fill in the fields for your error budget policy.
   For more information, see Add an error budget policy form.

   Note: If you don't add an error budget policy, the SLO objective percentage is only informational. Adding a policy helps you take corrective actions based on the performance data. If you're not ready to add a policy, you can add one later.

7. Review your SLO, SLI, and error budget policy, and then select **Activate**.

   Your SLO, SLI, and error budget policy are now active. The SLO record opens, displaying its reliability state and error budget information.

What to do next

See the following to learn how to monitor and manage your SLOs:

- Reliability metrics in SLO Management - Learn where and how to view SLO performance.
- Edit a reliability metric - Update your SLOs to keep them aligned with your reliability goals.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a notification destination in SRM

Set up notification destinations to help keep teams informed about service reliability.

Before you begin

Before creating notification destinations, set up the Microsoft Teams Graph spoke. When registering an application in the Microsoft Azure portal, grant the following Microsoft Teams Graph API permissions:

- `ChannelMessage.Send`
- `Group.ReadWrite.All**`

For setup instructions, see Set up the Microsoft Teams Graph spoke.

Role required: srm_manager, srm_admin, or admin

About this task

Create a notification destination to deliver messages to a specific Microsoft Teams channel. After you set up a notification destination, team members can attach it to error budget policies to receive messages about policy breaches.

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
2. From the primary navigation, select **Teams** ( ).
3. On the Teams page, select the team you want to use the notification destination.
4. Select the **SLO Notification destinations** tab, select **Create a destination**, and then fill out the form.

Field	Description
Channel URL	Deep link to the relevant Microsoft Teams channel. You can find deep links to channels in Microsoft Teams. They follow this general format: `https://teams.microsoft.com/l/channel/19%abcdef`.
Destination name	Display name of the notification destination. Teams see this name when adding notification destinations to error budget policies.

5. **Optional:** Select **Send test notification** to verify your setup.
   When successful, a message appears in the relevant Microsoft Teams channel. The message is titled `Test Notification` and includes the following text:

```
This is a test notification from the Service Reliability
 Management (SRM) Team '<team-name>'.
```

6. Select **Save**.
   The notification destination appears in the SLO Notification destinations tab. The table includes details about the Microsoft Teams channel and when the notification destination was last updated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

To use and learn more about notification destinations, visit the following pages:

- Create SLOs, SLIs, and error budget policies
- Notifications for breached error budgets

You can also edit and delete notification destinations from the SLO Notification destinations tab. Deleting a notification destination removes it from all error budget policies where it was previously used.

Edit a reliability metric

Update a reliability metric to keep it relevant and aligned with your team's goals.

Before you begin

Updating reliability metrics, such as Service Level Objectives (SLOs) or Service Level Indicators (SLIs), during active measurement periods can result in graph inconsistencies.

Role required: srm_admin, srm_manager, or srm_responder

About this task

To help you manage reliability-metric history, Service Reliability Management (SRM) follows specific naming conventions. When you edit a reliability metric, SRM does the following:

- Retires the existing SLO.
- Creates a copy of the SLO with your edits in a draft state, which you can either save as a draft or activate.
- Appends a number to the copy name, for example, Uptime (2) or Uptime (3).

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
2. Follow these steps to access the **Reliability metrics** tab:
   a. Select the **Services** page icon ( ).
   b. Select the service that uses the reliability metric you want to edit.
   c. Select the **Reliability metrics** tab.
3. To edit a reliability metric, select the relevant reliability metric, and then select **Edit** and **Continue to edit**.
4. Make your changes in the Details, Service level indicators, or Error budget policies tabs.

   Note: If you change the SLI data source, all existing SLIs for the SLO are removed and you must create new ones. An SLO can use only one data source type, and each data source has unique settings that can't be preserved or automatically converted.

   While you can't recover the removed SLIs while editing the SLO, you can view the earlier versions in the Reliability metrics tab.

5. To save and implement your changes, select **Activate**.

   Your edited reliability metric appears in the Reliability metrics tab. SRM sets the state to running and includes a version number in its name, for example, Uptime (3).

   The Activity panel on the Details tab records changes, providing a history of updates to the reliability metric.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Managing SLOs requires ongoing updates to make sure that they reflect your reliability goals. You can edit or retire existing SLOs and, if needed, reactivate retired SLOs.

Note: If you reactivate a retired SLO, SRM creates and activates a new copy of it. For example: If there are two versions of an SLO, Uptime (1) and Uptime (2), and you reactivate Uptime (1), SRM creates an active version called Uptime (3).

View SRM reliability metrics

Learn about an existing service level objective (SLO) and service level indicator (SLI).

Before you begin

Role required: srm_manager, srm_responder, or srm_admin

Note: Administrators can view any Service Reliability Management (SRM) SLO.

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
2. From the primary navigation, select the **Services** icon ( ).
3. Select a service, select the **Reliability metrics** tab, and then open an SLO.
4. View the SLO header to get key information about the SLO, SLI, and state.

Header element	Description
SLI type	○ **Availability**: Percentage of time the service is available. ○ **Errors**: Frequency of service errors. ○ **Latency**: Time taken to service a request. ○ **Saturation**: Fullness of the system, focusing on resource use.
SLO type	○ **Duration**: Measures the time that the service spends without breaching. The associated metrics are: ▪ **Objective (percentage)**: The target percentage of the SLI that the service is expected to meet. ▪ **Error budget**: The remaining error budget, shown in days and time, before the objective is breached. ○ **Count**: Measures the number of occurrences during a compliance period. The associated metrics are: ▪ **Limit (occurrences)**: The number of occurrences that triggers a breach. ▪ **Remaining breach occurrences**: The number of occurrences left before a breach occurs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Header element	Description
State	○ **Draft**: The SLO isn't active yet. You can add new SLIs, update existing ones, or delete the SLO. ○ **Running**: The SLO is active. You can edit, retire, or delete the SLO. Editing a running SLO automatically retires it and creates a new copy. See Edit a reliability metric for more information. ○ **Retired**: The SLO is no longer active. You can reactivate it if needed.
Service	The service associated with the SLO.
Reliability	○ **Stable**: All SLOs have more than 25% of their error budget remaining. ○ **At risk**: All SLOs have some error budget remaining, but at least one SLO has less than 25% left. ○ **Critical**: At least one SLO has fully used its error budget.

Note: From the header, you can delete, retire, or edit the SLO. Deleting an SLO also deletes the associated SLIs. Retiring an SLO changes its state, but you can reactivate it later from the SLO page or the Reliability metrics tab.

5. Select the **Overview** tab to view summary metrics and performance charts for the SLO.
   - Historic periods: Select a time range. The available periods span from the SLO creation date to the present in defined increments.
   - Summary metrics: View details about SLO performance depending on the configured SLIs and error budget policy thresholds.
   - Service level indicators: View the SLIs associated with the SLO.
   - Service level objective (SLO) history:
     - For duration SLOs, view charts for error budget used, error budget remaining, and burn rate.
     - For count SLOs, view charts for cumulative breach occurrences, burn rate, and alerts, incidents & changes impacting this service.

   You can point to the charts to view daily values.

6. Select the **Details**, **Service level indicators**, and **Error budget policy** tabs to view key configuration details.

   For information about the details shown, see Create SLO form, Create SLI form, and Add an error budget policy form. On the Error budget policies tab, you can select **Add policy** to define actions based on performance data.

   The Details tab also includes a Work notes field and Activity panel to record and track changes to the reliability metric.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

To learn how to create or edit SLOs, see Create SLOs, SLIs, and error budget policies and Edit a reliability metric.

SLO Management reference

Reference topics provide additional information about Service Level Objective Management.

Add an error budget policy form

Learn about the available fields for setting up error budget policies.

Fields in the Add an error budget policy form

The following table describes the available options in the Add an error budget policy form. For step-by-step instructions, see Create SLOs, SLIs, and error budget policies.

Field	Description
Policy name	Display name of the error budget policy.
Threshold type	Select how you want to track your error budget. • **Burn rate (count)**: Tracks how quickly the error budget is being consumed. A higher burn rate means faster consumption. • **Error budget remaining (%)**: Tracks the percentage of the error budget that's still available during the specified period. **Note:** You can add multiple thresholds to a policy, but each must be unique. If you try to add a duplicate threshold, an error message appears.
Threshold value	The point at which you get notified based on the selected SLO indicator. • Burn rate: Set a value that indicates the burn rate is too high. For example, to be notified when your service or configuration item is burning through its error budget twice as fast as expected, set the threshold to 2. • Error budget remaining: Set a value that indicates the remaining error budget is too low. For example, to be notified when only 25% or less of the error budget remains, set the threshold to 25%. **Note:** The threshold values support up to 4 decimal places and then round up.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Policy action	[Optional] If a breach occurs, select the action that you want to take: • Create incident. • Send an email. • Notify Microsoft Teams. If you select this option, you must also select a relevant notification destination. To set up a notification destination for your team, see Create a notification destination in SRM. **Note:** To learn about the notifications that teams receive for breached error budget policies, see Notifications for breached error budgets.

Create SLO form

Learn about the available fields for adding a service level objective (SLO) to Service Reliability Management.

Service level objective form

The following table describes the available options in the Service level objective form. For step-by-step instructions, see Create SLOs, SLIs, and error budget policies.

Fields in the Service Level Objective form

Field	Description
Name	Name of the SLO.
SLI type	• **Availability**: Percentage of time the service is available. • **Errors**: Frequency of service errors. • **Latency**: Time taken to service a request. • **Saturation**: Fullness of the system, focusing on resource use.
How do you want to measure this objective?	You can measure this objective by: • **Duration**: The amount of time the service spends without breaching. It's the only value available. • **Count**: The number of periods or occurrences in a given compliance period.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Measuring by Duration

Field	Description
Objective (%)	The percentage of how available (uptime) you want the service to be over a length of time. If measuring by duration, an objective percentage of Five Nines (99.999%) means that the service wasn't available for 26 seconds per month, or about 5 minutes and 35 seconds per year. This is used to calculate your error budget.
Compliance period	Period for which the metrics are calculated. The available options are: • **Month**: The duration is considered to be the current month. For example, if the current date is 26th January, the duration will be considered from 1st January until 31st January. • **Rolling 7 days**: The duration is considered to be 7 days from the current date. • **Rolling 30 days**: The duration is considered to be 30 days from the current date. For example, if the current date is 26th January, the duration will be considered from 25th December. • **Rolling 90 days**: The duration is considered to be 90 days from the current date. For example, if the current date is 26th January, the duration will be considered from 25th October.
Error budget	Auto-populated. The maximum level of errors (downtime) allowed over a certain period. A percentage of total service availability. Once you set your objective (%), your error budget is automatically calculated using this formula: 1 – (Desired) Availability. The error budget of a service will be calculated in days, hours, minutes, and seconds for a desired availability of 99.99% (Four Nines) of the year by: **1.** Calculating the total time in seconds in a year: 31536000 seconds. **2.** Determining the error budget as 1 – (Desired) Availability: 1 - 0.9999 = 0.0001. **3.** Converting the error budget to time units. For example, if the error budget is 0.0001:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Measuring by Duration (continued)

Field	Description
	• Error Budget in Seconds = 0.0001 * 31,536,000 seconds = 3,153.6 seconds. • Error Budget in Minutes = 3,153.6 seconds / 60 = 52.56 minutes. • Error Budget in Hours = 52.56/60 = 0.875 hours. • Error Budget in Days = 52.56 minutes / (24 * 60) ≈ 0.0364 days. **Note:** Your error budget appears in units of days, hours, minutes, and seconds only when you measure your SLO by duration.

Measuring by Count

SLO type	Count by periods or Count by occurrences
Limit occurrences	The number of occurrences after which a breach occurs. Limit occurrences act as an error budget.
Compliance period	Period for which the metrics are calculated. The available options are: • **Month**: The duration is considered to be the current month. For example, if the current date is 26th January, the duration will be considered from 1st January until 31st January. • **Rolling 7 days**: The duration is considered to be 7 days from the current date. • **Rolling 30 days**: The duration is considered to be 30 days from the current date. For example, if the current date is 26th January, the duration will be considered from 25th December. • **Rolling 90 days**: The duration is considered to be 90 days from the current date. For example, if the current date is 26th January, the duration will be considered from 25th October.

Note: The **Assignment group** is auto-populated.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create SLI form

Learn about the available fields for adding a service level indicator (SLI) to a service level objective (SLO).

Add SLI form

The following table describes the available options in the Add SLI form. For step-by-step instructions, see Create SLOs, SLIs, and error budget policies.

Fields in the Add SLI form

Field	Description
SLI name	Enter a name for your SLI. For example, Critical router issue.
Data source	This field is auto-populated. The alerts or outages are used to calculate the error budget.
Service or configuration item	• <Service-name> - Select this option to filter your SLI to the entire service. This option helps you measure overall service health. • Related configuration item - If your service includes one or more configuration items (CIs) in its hierarchy, select this option to filter your SLI to specific CIs. This option provides a more granular view of service health. **Note:** By default, only level-one CIs appear in the CI list. Administrators can change the number of levels displayed by updating the `sn_sow_slo.slo_rollup_depth` system property value.
Set conditions (optional) Customize your SLI by filtering it to specific alerts or outages. The fields in this section are optional.	
Add condition set	Select this option to specify the alerts or outages included in the error budget. For example, for alerts, if you set the condition to **Status is Critical**, only alerts with a critical status count toward the error budget.

Notifications for breached error budgets

Details about the notifications that teams receive for breached error budget policies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When notifications are sent

When a service or configuration item breaches its error budget policy and a notification destination is set, Service Reliability Management (SRM) sends a notification to the configured Microsoft Teams channel.

Notification content

The message contains details about the breach to help team members learn about the issue and act. In Microsoft Teams, the message is titled `Error budget policy violation` and lists the following information:

- Service level objective - Name of the service level objective (SLO), for example, `Navigation service availability 99.95`.
- Service/CI - Name of the associated service or configuration time. Select the name to open the relevant service in SRM.
- Policy name - Name of the error budget policy that was breached.
- Error budget remaining - Percentage of error budget left at the time of the breach.
- View SLO in ServiceNow - Link to view the SLO details in SRM.

Sample notification

The following image is a sample notification in Microsoft Teams. It shows an error budget policy violation for the Cart service and details about the SLO and error budget remaining.

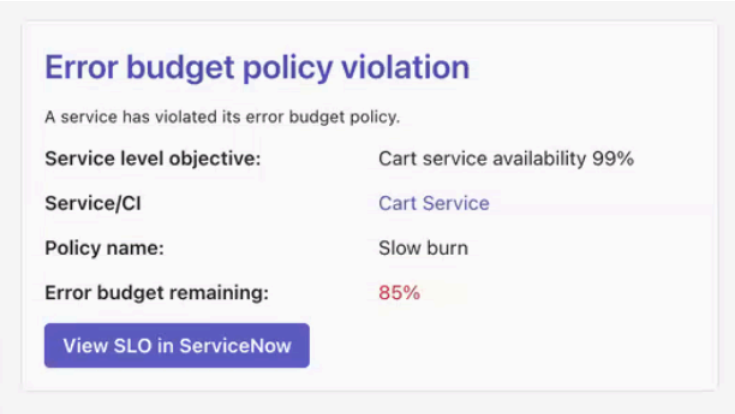

Related documentation

See Create a notification destination in SRM and Create SLOs, SLIs, and error budget policies.

SLO creator agent system properties

Configure how the SLO creator agent generates and activates SLOs.

To turn automatic SLO generation on or off and manage email notifications for generated SLOs, see Manage SLO creator agent settings.

Note: Property changes apply only to newly generated SLOs. Existing SLOs aren't recalculated.

sn_ai_agents_slo.commitment_buffer

Description

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Adds a buffer to an availability commitment when the agent calculates the SLO target. Applies only when the agent finds an availability commitment.

Values

Default: `0.09`

When to adjust

Increase the buffer to set stricter SLO targets and provide earlier warning before the service falls below its availability commitment. Decrease it if the generated targets are too strict.

sn_ai_agents_slo.create_slo_mode

Description

Controls if the agent activates generated SLOs automatically or saves them as drafts.

Values

Default: `auto_activate`

Other value: `draft`

When to adjust

Use `draft` to review and edit generated SLOs before activation. Use `auto_activate` when generated SLOs can take effect without review. Error budget notifications begin after activation.

sn_ai_agents_slo.linked_alerts_cap

Description

Limits the number of critical alerts associated with an incident that the agent considers when generating SLOs.

Values

Default: `5`

Maximum: `10`

When to adjust

Increase the value when incidents commonly produce more relevant critical alerts than the default includes. Decrease it when numerous alerts represent the same underlying issue.

sn_ai_agents_slo.linked_outages_cap

Description

Limits the number of outages associated with an incident that the agent considers when generating SLOs.

Values

Default: `5`

Maximum: `10`

When to adjust

Increase the value when a single incident commonly spans more outage records than the default includes. Decrease it when multiple outage records represent the same underlying issue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

sn_ai_agents_slo.orphan_alert_window_days

Description

Specifies how far back the agent looks for critical alerts that aren't associated with an incident when generating SLOs.

Values

Default: `30`

Range: `7 - 90`

When to adjust

Use a shorter window when older alerts no longer reflect current service reliability. Use a longer window to account for less frequent or recurring alert patterns.

sn_ai_agents_slo.other_incident_window_days

Description

Specifies how far back the agent looks for other P1 and P2 incidents associated with the service or CI when generating SLOs.

Values

Default: `90`

Range: `30 - 180`

When to adjust

Use a shorter window to focus on recent conditions. Use a longer window to account for severe incidents that occur less frequently.

Visualizations in the Service reliability dashboard

List of visualizations and options on the Service reliability dashboard in Service Reliability Management (SRM).

Service state charts

Top-level charts show the number of services in critical, at-risk, and stable states. Their states are based on the error budget remaining on their service level objectives (SLOs). You can select the charts to view service names, adjust the time range, and access additional chart options.

Note: An error budget is the amount of failure a service can experience before breaching its SLO.

Each visualization also includes a trend line showing changes in service count over the past 12 months. Smaller figures indicate how the count has changed compared to a week ago, for example, ↓25 (22%) since Jun 11.

Chart	What it is	How to use it
Critical	Displays the number of services in a critical state. Critical services have 0% error budget remaining on their SLOs.	View how many services have consumed their error budgets and identify the services needing immediate attention.
At risk	Displays the number of services at risk. At-risk services have <= 25% error	Monitor how many services are approaching critical

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart	What it is	How to use it
	budget remaining on their SLOs.	thresholds and find issues early.
Stable	Displays the number of stable services. Stable services have more than 25% error budget remaining on their SLOs.	Get insights into overall service health and identify if services are staying reliable over time.

Risk trends over time

Line charts track the number of SLOs with high burn rates and low error budget remaining over the past 12 months. You can use them to find recurring patterns and potential reliability risks.

Chart	What it is	How to use it
High burn rate (>=1)	Shows the number of SLOs with a burn rate >= 1 over time. A high burn rate indicates that the service linked to the SLO is likely to breach its error budget before the compliance period ends. For example, if a service has 30 days to meet its SLO but is using up its error budget in 15 days, the burn rate is 2.	• Find risks early by seeing when services begin consuming error budgets too quickly. • Identify emerging or recurring reliability issues by tracking burn rates over time. • Point to the chart to see the number and percentage of SLOs with a high burn rate at that time. • Select the chart to view SLO details, adjust the time range, and access additional chart options.
Low budget remaining (<=25%)	Shows the number of SLOs with low or no error budget remaining over time.	• Monitor how many services are nearing or have breached their SLOs. • Track rising trends, which might indicate declining reliability or recurring issues that need investigation. • Point to the chart to see the number and percentage of SLOs with little or no error budget remaining at that time. • Select the chart to view SLO details, adjust the time range, and access additional chart options.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service level objectives (SLOs) table

The SLOs table lists the SLOs defined in Service Reliability Management (SRM), and it's sorted by SLO name by default. Use the table to monitor overall reliability, identify services at risk, and find the assigned teams.

The SLO table includes the following columns by default. To customize the columns shown, select the gear icon.

- Name - Name of the SLO. You can select the arrow to sort the table by SLO name, and you can select the name to view the SLO record.
- Reliability - Current state of the SLO. For example, stable, at risk, or critical.
- Measured reliability - Percentage showing the actual performance of the service. For example, if your SLO is 99.9% success, and the actual performance for the month is 99.7%, the measured reliability is 99.7%.
- Objective (percentage) - Target SLO value.
- Burn rate - Numeric value showing how quickly the service is consuming its error budget.
- % Error budget remaining - Percentage of the error budget still available in the current compliance period.
- Service - Name of the service associated with the SLO. You can select the service name to view the service record.
- Assigned - Team responsible for the service.

Dashboard filters and actions

The Service reliability dashboard is built with Platform Analytics and includes standard dashboard features. For details on customizing, duplicating, or sharing dashboards, see the Dashboards in Platform Analytics documentation.

Note: Changes to the Service reliability dashboard affect all SRM users on your instance. To create a personalized version, either build a new dashboard or duplicate the existing dashboard and edit it. Learn more in Create a dashboard with the in-line editor and Duplicate a Platform Analytics dashboard.

Service Observability

The ServiceNow® Service Observability application enables operators to triage and manage incidents in a complex and distributed system. Application and platform telemetry insights from multiple observability tools are combined into a centralized workflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview of Service Observability

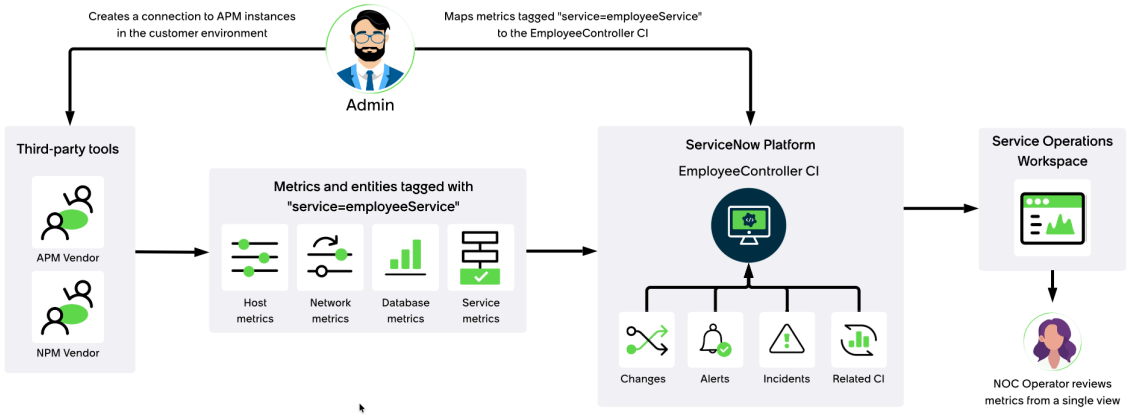

Get started

Explore

Learn how Service Observability brings disparate data together to help operators find issues in your system

Configure

Learn how to install and configure Service Observability

Monitor and investigate

Learn how to find the source of issues and view the health of services and their related infrastructure entities

Reference

Find reference information about Service Observability

Exploring Service Observability

Service Observability helps operations teams triage and manage incidents in a complex and distributed production system. It combines external observability monitoring systems' telemetry with related data from the Configuration Management Database (CMDB) and displays both in a single workflow in the Service Operations Workspace (SOW).

Service Observability overview

Service Observability displays application, infrastructure, and network health metrics in the SOW related to a given service. Metrics can be ingested from an external observability vendor

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

(application, network, and cloud monitors) and displayed alongside information for related configuration items in the CMDB.

Service Observability supports the following observability vendors:

- Amazon CloudWatch
- AppDynamics
- Cisco ThousandEyes synthetic tests
- Datadog
- DynatraceSaaS and on-premise (both Classic and Grail environments)
- LogicMonitor
- Microsoft Azure Monitor
- New Relic
- Prometheus on-premise
- SolarWinds on-premise
- Splunk Observability and logs from Splunk Enterprise
- Zabbix on-premise

Service Observability also supports MetricBase as a data source. For setup instructions, see Create and manage MetricBase data mappings.

Service Observability supports the following databases:

- MySQL
- PostgreSQL (not supported with Splunk)
- RDS (Relational Database Service) for Amazon CloudWatch

After connecting an observability vendor to Service Observability, you map services in the CMDB to observability metrics using existing tags.

With this data mapping, Service Observability displays metrics for entities such as host, database, or network components, along with details about related CI information. Operators use these metrics and contextual information, including current incidents and alerts, to assess service health.

Note: Operators can also use the Analyze a dashboard in Service Observability and the Analyze service health in Service Observability AI skills to generate an analysis for them.

For example, you use Dynatrace to monitor your `checkout` service, databases, and hosts. Metrics from all these entities use the tag `checkout-service` to denote requests coming from that service. By mapping the `checkout` service CI to the Dynatrace data tagged with `checkout-service`, Service Observability retrieves metrics for those databases, hosts, and related CIs. It then displays them together. Operators can pinpoint issues on entities related to the service and narrow down the mitigation process without having to leave the SOW.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Observability users

Users

User	Description
System admin	System admins configure users and teams, register services to be monitored, connect Service Observability to observability vendors, and then map those services to that data. They can also view the data in the SOW
Service Observability admin	Service Observability admins can configure users and teams, connect Service Observability to observability vendors, and then map services to that data. They can also view the data in the SOW. Admins can also customize dashboard templates used to display metrics and related information.
Operator/ operations manager **Note:** These users must belong to an `srm` group type to see all data.	Operators use Service Observability when triaging incidents in the SOW. They can view basic health metrics for a service, along with related incidents, alerts, and changes. They can get more detailed information by navigating to the **Observability** tab to view additional service metrics, along with metrics from related entities, such hosts, networks, or databases.

Service Observability workflow

Admins configure Service Observability by creating a connection to an observability vendor and then mapping CI services to that data. Operators use Service Observability to determine if another related entity is causing issues surfaced by the service's performance.

As an admin, you:

1. Determine the services to be monitored by Service Observability based on business criticality.
2. Connect existing observability vendor instances to Service Observability.
3. Map services to observability metric data using vendor-based tags attached to that data.
4. Customize the templates used to display metric charts.

As an operator or manager, you:

1. Spot an issue with a service while working in the SOW, for example, from an alert, the Service dashboard, or Express List, then navigate to the Service Details page.
2. View overall health metrics for the service, along with related incidents, alerts, and changes. If one of the metrics seems unhealthy, navigate to the Observability tab.
3. View more detailed service metrics, as well as information from related entities, to start root cause investigation. When finding that the issue is further down the system's stack, identify the ownership for that entity to start remediation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Observability benefits

Benefit	Feature	Users
Consolidate data from existing monitoring tools, network health tools, cloud providers, ServiceNow agents, and third-party tools for a full-stack view of service health: • Connect data from external Observability vendors. • Map that data to CMDB services • View combined data in the SOW	• Connect a Service Observability data source • Create and manage data mappings .	Admins
Increase efficiency and reduce mean time to resolution (MTTR). View combined metrics from entities associated with a service to begin to determine blast radius and ownership of an incident.	View service health metrics	Operators
See related changes to the system and alerts associated with a service in one place.	View overall service health.	Operators
Use generative AI to analyze metric data and find insights to help determine service health.	• Analyze a dashboard in Service Observability • Analyze service health in Service Observability	Operators
See Service Observability data as part of Incident Management workflows	Digital End-User Experience and Service Observability UI experience on investigate tab	Operators
Customize dashboard templates.	Customize Service Observability dashboard templates	Admins

What to explore next

To learn more about configuring and using Service Observability, see:

- Configuring Service Observability
- View overall service health
- View service health metrics
- Service Observability reference

AI in Service Observability

Generative AI skills in Service Observability help operators analyze metric data and assess service health faster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

AI features in Service Observability

AI feature	Description	Use case	Resources
Analyze dashboard skill	Uses generative AI to summarize a single Service Observability dashboard and call out insights found in the charts within the selected time period.	An operator wants a quick orientation to an unfamiliar dashboard or needs to identify anomalies in chart data without manually reviewing each chart.	• Activate the analyze Service Observability dashboard skill • Analyze a dashboard in Service Observability
Analyze service health skill	Uses generative AI to analyze all available Service Observability dashboards for a selected service, detect insights, and report on overall service health. When activated, the skill also runs automatically during incident investigation from the **Investigate** tab.	An operator is triaging an incident and needs to quickly assess service impact, identify root causes, and determine next steps across all dashboards for the affected service.	• Activate the analyze service health skill • Analyze service health in Service Observability

Configuring Service Observability

After you install Service Observability, you must connect external observability systems and map that data to services. You can also customize the out-of-the-box dashboards to show different information, including data from other ServiceNow products.

Configuration overview

Install Service Observability

If you have the system admin role, you can install the Service Observability application (`sn_sow_svcobs`). The application installs related ServiceNow® Store applications and plug-ins if they aren't already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Service Observability application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: system admin [admin]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The following items are installed with Service Observability:

- Plugins
- Store applications
- Roles
- Tables

For more information, see Components installed with Service Observability.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Service Observability application (`sn_sow_svcobs`) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you can't find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

Activate the analyze service health skill

The analyze service health skill uses generative AI to provide an analysis of all Service Observability dashboards for the selected service. You need to activate this skill before using it.

Before you begin

Before activating the AI skills, you must install the following:

- ServiceNow Otto for ITOM plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.
- Service Observability. For more information, see Install Service Observability.

Role required: admin

Procedure

1. Navigate to **All** > *AI Admin Hub* > **Skills**
2. On the navigation panel, select **ITOM**.
3. On the **Analyze service health** tile, select **Activate**.

Activate the analyze Service Observability dashboard skill

The analyze Service Observability dashboard skill uses generative AI to provide an analysis of a Service Observability dashboard. You need to activate this skill before using it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Before activating the AI, you must install the following:

- ServiceNow Otto for ITOM plugin. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.
- Service Observability. For more information, see Install Service Observability.

Role required: admin

Procedure

1. Navigate to **All** > **AI Admin** > **Skills**
2. On the navigation panel, select **ITOM**.
3. On the **Analyze Service Observability dashboard** tile, select **Activate**.

Connect a Service Observability data source

Connect Service Observability to an external observability system. Service Observability displays metrics in the Service Operations Workspace (SOW) from that observability instance.

Before you begin

You must enter a connection URL and credentials for your observability vendor. You can either use an existing connection or create one. Note the following for specific vendors:

Vendor-specific connection information

Observability Vendor	Connection information
AWS	The connection URL must be specific to the region to be monitored, for example, `https://monitoring.us-east-2.amazonaws.com`.
Azure	The connection URL should be `https://management.azure.com`.
Splunk Observability	The connection URL must include the Realm, for example, https://stream.[REALM].signalfx.com
Zabbix	The connection URL must be the IP on which the Zabbix server is hosted

If you're using existing credentials, you need the credential name. If you're creating credentials, collect the following information.

Vendor-specific credential requirements

Observability vendor	Credential type	Information needed
AppDynamics	OAuth 2.0	• OAuth Client ID • OAuth Client secret

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Vendor-specific credential requirements (continued)

Observability vendor	Credential type	Information needed
		• OAuth token URL • AppDynamics account name
Azure	OAuth 2.0	• OAuth Client ID • OAuth Client secret • OAuth token URL • Subscription ID
AWS	AWS credentials	• Access key • Secret key If the user associated with this key is restricted to specific permissions, they must include the following, scoped to all resources: • `cloudwatch:Describe` scoped to all resources • `cloudwatch:Get` scoped to all resources • `cloudwatch:List` scoped to all resources • `cloudtrail:LookupEvents` • `tag:GetResources` • `tag:GetTagKeys` • `tag:GetTagValues` • `apigateway:Get`
Datadog	API	• API key • Application key

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Vendor-specific credential requirements (continued)

Observability vendor	Credential type	Information needed
		If an application key is restricted to a specific scope, those scopes are required for the following endpoints: • https://api.datadoghq.com/api/v2/catalog/entity • https://api.datadoghq.com/api/v1/hosts • https://api.datadoghq.com/api/v1/query • https://api.datadoghq.com/api/v2/query/timeseries
Dynatrace SaaS Classic	API	API Tenant Access Token The token must have `entities.read` and `metrics.read` scope permissions.
Dynatrace SaaS Grail	API	API Platform Token The token must have `storage:metrics.read` and `storage:bucket.read` scope permissions.
Dynatrace on-premise	API	• API key. The key must be an environment scoped token. • The name of the MID Server your Dynatrace instance is hosted on
LogicMonitor	API	Access token
New Relic	API	• Account ID • API key
Prometheus on-premise	Basic Auth	• Basic Auth credential user name • Basic Auth credential password • The name of the MID Server your Prometheus instance is hosted on

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Vendor-specific credential requirements (continued)

Observability vendor	Credential type	Information needed
SolarWinds	Basic Auth	• Basic Auth credential user name • Basic Auth credential password • The name of the MID Server your SolarWinds instance is hosted on
Splunk Enterprise	API	Bearer Token
Splunk Observability	API	Access token
ThousandEyes	API	API key
Zabbix	Basic Auth	• Basic Auth credential user name • Basic Auth credential password • The name of the MID Server your Zabbix instance is hosted on. The MID Server must be in an active state.

Role required: sn_sow_svcobs.admin

About this task

The first step to ingesting observability data is to create a connection to the observability instance.

Service Observability supports the following observability vendors:

- Amazon CloudWatch
- AppDynamics
- Cisco ThousandEyes synthetic tests
- Datadog
- DynatraceSaaS and on-premise (both Classic and Grail environments)
- LogicMonitor
- Microsoft Azure Monitor
- New Relic
- Prometheus on-premise
- SolarWinds on-premise
- Splunk Observability and logs from Splunk Enterprise
- Zabbix on-premise

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Service Operations Workspace** > **Configurations**, then navigate to **Service Observability** > **Data sources**.
2. Select your observability vendor.
3. On the **Add connection** form, fill in the fields.
   You can use an existing HTTP or HTTPs connection and an existing credential or create one.
4. Select **Save**.

Result

A new connection appears in the list of observability data sources. You use this connection when you create a data mapping.

What to do next

Map your connection to the services that you want to monitor in Service Observability. See Create and manage data mappings.

Create and manage data mappings

Map your services to the data from a connected external observability vendor instance, and view it in charts for the service.

Before you begin

- You can map any of the following service types:
  - Service instance
  - Mapped application service
  - Calculated application service
  - Dynamic CI group
  - Technology management service
  - Tag-based service
  - Offerings
  - Business service
- Connect a Service Observability data source

Role required: sn_sow_svcobs.admin

About this task

Tip:

If you want to use MetricBase as your data source, follow the procedure described in Create and manage MetricBase data mappings instead.

When a service CI is mapped to specific observability vendor metric tag key/value pairs, Service Observability displays metrics containing those tags for that service. These metrics are grouped by the following entity types:

- Application: Metrics associated with the selected service
- Compute: Metrics associated with the hosts the service is run on
- Databases: Metrics associated with databases used by the service
- Networking: Metrics associated with network devices used by the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, say you have metrics from New Relic that are used to monitor services, databases, and hosts. On those metrics, you have used the tag `service_name`. Say you have also mapped the `checkout` service CI to the observability metadata key/value pair `service_name = checkout`. Service Observability shows you all the database and host metrics tagged with `service_name = checkout` in the context of the `checkout` service CI.

If some of your metrics on different entities use a different key name or a different key value, you can create exceptions. For example, if your host metrics use `service` instead of `service_name`, you can create an exception to include those metrics.

Each entity type has its own dashboard and each metric or related CI has its own default chart. You can customize the templates used for the dashboards and charts. See Customize Service Observability dashboard templates for more information.

Note: To understand how Service Observability creates these mappings, see Observability vendor entity mappings for Service Observability.

For a description of the field values, see Service Observability data mapping form.

Procedure

1. Navigate to **All** > **Service Observability** > **Configure Data mappings**.
2. Select **Create your first mapping** or **Create mapping**.
3. On the Observability data mapping page, enter a name for the mapping.
4. Choose the services that should use this mapping.
   a. Choose **Select services**.

      The top 30 recommended services display in dialog. These are services that would most benefit from using Service Observability based on business criticality and frequency of alerts, incidents, and SLOs over the past 90 days.

      You can instead manually choose services by choosing **All Services** from the left navigation.

      Note: Use the filter to narrow down the service list.

   b. Select the services to add to the mapping. You can add services from any combination of service type.

   c. Select **Add services** to add them to the mapping.
   The Mapping preview pane updates to show the selected service, or the first service in a list of services.
5. Create a rule to determine how Service Observability should map internal services to the observability entities using tag keys and values.

   a. Select the data source for this map.

   b. Enter the tag key that you want to use to map the observability data to the service.
   The key can be any tag you use on your observability metadata, for example `service`. The entered key must match the tag you use in the observability metadata and is case-sensitive.

      Note: If you use more than one tag to represent a service, you can create an exception. For example, if most tags use `service` except for database metrics, which use `service_name`, you can create an exception for the database metrics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Manually enter a tag value or choose one from the drop-down menu.

Note: Values are case-sensitive.

Values can be strings that are from your observability metadata, like a service name, or they can be a variable used to represent multiple instances, for example `$(service)`. When your key represents a service, a drop-down menu shows fields from the corresponding CI that can be used as a variable. If you choose to use a variable, be sure that the values resolve to values sent in the observability data.

Note: If you selected more than one service to map, the drop-down displays only the fields that are common to all selected services.

You can now test your mapping or add exceptions.

6. **Optional:** If some of your entities use different key names or values, create exceptions to the default policy by configuring them in the Exceptions card.
7. Test the mapping.
   The **Mapping preview** panel shows the test results for the first service mapped. Service Observability sends a request through the connection and returns the entities found using the provided mapping. If you created an exception using a different source, the panel shows the results from each source separately. If the results don't seem correct, try reconfiguring the mapping by changing the tag key or value.

   If you have more than one service mapped, use the **Service** drop-down menu to select and test a different service.
8. Select **Save mapping**.

Result

On the **Observability** tab of the Service details page, dashboards and charts show metrics for the entities found from the mapping.

Create and manage MetricBase data mappings

Map your services to metrics from MetricBase, and view them in charts for the service.

Before you begin

- You can map any of the following service types:
  - Service instance
  - Mapped application service
  - Calculated application service
  - Dynamic CI group
  - Technology management service
  - Tag-based service
  - Offerings
  - Business service
- Your instance has metrics in MetricBase.

Role required: sn_sow_svcobs.admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

When a service CI is mapped to MetricBase, Service Observability can display metrics from MetricBase for that service and its related entities (such as hosts, databases, and network devices). The system automatically identifies the entities through CMDB relationships.

When you select MetricBase as your data source, generic dashboards appear that display ServiceNow-native charts. You can customize these dashboards by adding MetricBase charts for the specific metrics you want to view. See Add MetricBase charts to Service Observability dashboard templates for more information.

Procedure

1. Navigate to **All** > **Service Observability** > **Configure Data mappings**.
2. Select **Create your first mapping** or **Create mapping**.
3. On the Observability data mapping page, enter a name for the mapping.
4. Choose the services that should use this mapping.
   a. Choose **Select services**.

      The top 30 recommended services display in dialog. These are services that would most benefit from using Service Observability based on business criticality and frequency of alerts, incidents, and SLOs over the past 90 days.

      You can instead manually choose services by choosing **All Services** from the left navigation.

      Note: Use the filter to narrow down the service list.

   b. Select the services to add to the mapping. You can add services from any combination of service type.
   c. Select **Add services** to add them to the mapping.
   The Mapping preview pane updates to show the selected service, or the first service in a list of services.
5. Select **MetricBase** as the data source for this mapping.
   Service Observability automatically detects and displays the entities related to your services. No additional configuration is required.
6. **Optional:** If some of your entities require different configurations, create exceptions to the default policy by configuring them in the Exceptions card.
7. Test the mapping.
   The **Mapping preview** panel shows the test results for the first service mapped. Service Observability sends a request and returns the entities found using the provided mapping. If the results don't seem correct, try reconfiguring the mapping or adjusting your exception rules.

   If you have more than one service mapped, use the **Service** drop-down menu to select and test a different service.
8. Select **Save mapping**.

Result

On the **Observability** tab of the Service details page, dashboards for the selected service(s) show ServiceNow-native charts. You can customize these dashboards by adding MetricBase charts for the specific metrics you want to view. See Add MetricBase charts to Service Observability dashboard templates for more information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Customize Service Observability dashboard templates

You can customize the Service Observability dashboards on both the Overview and Observability tabs of the Service Details page. You can change or add metrics and related data to fit your business needs.

Service Observability dashboards are built using templates specific for each observability vendor and entity tab. They're built using Platform Analytics, letting you customize templates to meet your business needs. For example, you might want to display metrics from your observability vendor that aren't available by default.

The Service Observability dashboards use vendor-specific templates. You can see what each template displays by default from its respective reference topic.

- Amazon CloudWatch templates for Service Observability
- AppDynamics templates for Service Observability
- Azure Monitor templates for Service Observability
- Cisco ThousandEyes templates for Service Observability
- Datadog templates for Service Observability
- Dynatrace templates for Service Observability
- LogicMonitor templates for Service Observability
- New Relic templates for Service Observability
- Prometheus templates for Service Observability
- SolarWinds templates for Service Observability
- Splunk Observability templates for Service Observability
- Zabbix templates

Note: When you make a customization, you're changing the template and not the individual service's dashboard. This means that dashboards for any services that also use the same data source are also changed.

Along with observability data, you can also add charts for data from ServiceNow tables and also from data stored in the MetricBase and Health Log Analytics applications.

When you customize a template, a copy of the original is saved so that you can reimplement it if needed. Default dashboards display a `Certified` tag.

Edit data charts on Service Observability dashboard templates

Edit Service Observability dashboard templates to view different observability metrics on the Overview or Observability tabs' charts. Metrics are scoped to the selected service.

About this task

When you edit or add an observability chart on a dashboard, you use the Edit data source page of Platform Analytics to create a query. Use the built-in query editor for simple metric queries. Use Advanced mode for more complex queries or to paste a query copied from a chart in your observability vendor instance.

Use simple queries to return metrics associated with services, hosts, and databases. The returned metrics are aggregated across the time window set by the user.

Use advanced query mode to cut and paste queries from existing observability charts into Service Observability. For most observability vendors, you can simply copy a query from an existing chart in your observability instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Both query types support template variables that let you use the same query to return metrics for multiple entity instances.

Before you begin

Role required: sn_sow_svcobs.admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.
2. Open the template in editing mode.
   - If you're editing a certified template, select **Duplicate**.
   - If you're editing a custom template, select **Edit**.

   Note: Duplicating a certified template keeps you from overwriting it and also lets you reinstall it.

   The new dashboard is titled with the words `- Copy` appended. Use the pencil icon to change the dashboard name.
3. To do basic editing, such as rearranging, resizing, or deleting charts, follow the instructions for Edit in-line Platform Analytics dashboard elements.
4. To add a vendor query for a new chart, follow these steps:

   a. In Edit mode, select **Add new element** and choose **Data visualization** > **New visualization**.

   b. In the **Configuration panel**, select `Line` for **Visualization type** and choose **Add data source**.

   c. In the **Add data source** page, open **Service Observability** in the navigation panel and select your observability vendor.
   Be sure to select the same vendor used on the template you're editing.

   d. Enter a query in the **Metric** card.

   Use the query builder to edit a simple vendor-specific query.

   To copy a full query from your vendor chart, activate the **Advanced** switch and paste in a query from your vendor chart.

   For AWS and Azure, instead of entering a query in advanced mode, you can import selected charts from an existing dashboard. Select **Import from AWS/Azure Dashboard**. Select a data source name, then select a dashboard and from the list of charts, select the one to import, then select **Import query**. The query for the chart is pasted into the Query box.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use template variables to replace hard-coded entities or time periods. This makes the query return results that match the service and time period selected in the UI.

Note: Use the **Docs** link to refer to the vendor-specific documentation for help with creating the query.

e. To preview your query, select **Run** and then choose a service and entity.

f. When satisfied with your query, select **Add this source** to add the chart to the template.

5. To edit an existing query, follow these steps.

a. In Edit mode, select the chart to edit.

b. In the **Data sources** section of the Configuration panel, use the **More actions** menu to select **Edit**.

c. In the **Edit data source** page, edit the query in the **Metric** card.

Use the query builder to create a simple vendor-specific query.

To copy a full query from your vendor chart, activate the **Advanced** switch and paste in a query from your vendor chart.

For AWS and Azure, instead of entering a query in advanced mode, you can import selected charts from an existing dashboard. Select **Import from AWS/Azure Dashboard**. Select a data source name, then select a dashboard and from the list of charts, select the one to import, then select **Import query**. The query for the chart is pasted into the Query box.

Use template variables to replace any hard-coded entities or time periods.

Note: Use the **Docs** link to refer to the vendor-specific documentation for help with creating the query.

d. To preview your query, select a service and an entity and select **Run**.

e. When satisfied with your query, select **Apply** to add the chart to the template.

6. When done customizing, select **Save** and then **Exit editing mode.**
7. To return the template to the default (Certified) version, use the **More actions** menu to select **Return to certified**.

Edit ServiceNow data on Service Observability dashboard templates

Edit Service Observability dashboard templates to view data from problem and business app records on the Overview or Observability dashboards.

Before you begin

Role required: sn_sow_svcobs.admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
You can access a service record from these pages in the SOW:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Services** list: Choose a service from the list.
- **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
- **List**: Navigate to **Application Services** > **Services** and select a service.
- **Express list** alert: Select a service from the **Impacted services** column.

The Service Details page opens and the Overview tab is displayed.

Trouble?
If charts are displaying error messages, see Chart error states.

2. Open the template in editing mode.
   - If you're editing a certified template, select **Duplicate**.
   - If you're editing a custom template, select **Edit**.

   Note: Duplicating a certified template keeps you from overwriting it and also lets you reinstall it.

   The new dashboard is titled with the words `- Copy` appended. Use the pencil icon to change the dashboard name.
3. To do basic editing, such as rearranging, resizing, or deleting charts, follow the instructions for Edit in-line Platform Analytics dashboard elements.
4. To add a new ServiceNow table, follow these steps:
   a. In Edit mode, select **Add new element** and choose **List**.
   b. In the **Configuration panel** in the **Data sources** section, choose **Add data source**.
   c. In the **Add data source** page, find the **Table** whose data you want to display and choose **Add this source**.
   d. In the **Columns and rows** section, add and arrange the columns to display.
5. To edit an existing chart, follow these steps.
   a. In Edit mode, select the chart to edit.
   b. In the **Data** section of the Configuration panel, verify that the correct table is selected in **Table** field.
   c. In the **Columns and rows** section, add and arrange the columns to display.
6. When done customizing, select **Save** and then **Exit editing mode.**
7. To return the template to the default (Certified) version, use the **More actions** menu to select **Return to certified**.

Add MetricBase charts to Service Observability dashboard templates

Add MetricBase data to charts on Service Observability dashboard templates when you want to view those metrics in context of Service Observability.

About this task

You can add metrics that are stored in the MetricBase application to your Service Observability dashboards. These charts display with data scoped to the selected service.

Before you begin

Role required: sn_sow_svcobs.admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.
2. Open the template in editing mode.
   - If you're editing a certified template, select **Duplicate**.
   - If you're editing a custom template, select **Edit**.

   Note: Duplicating a certified template keeps you from overwriting it and also lets you reinstall it.

   The new dashboard is titled with the words `- Copy` appended. Use the pencil icon to change the dashboard name.
3. To do basic editing, such as rearranging, resizing, or deleting charts, follow the instructions for Edit in-line Platform Analytics dashboard elements.
4. To add a new chart with MetricBase data, follow these steps:

   a. In Edit mode, select **Add new element** and choose **Data visualization**, select **New Visualization**, and then select **Line**.

   b. In the **Data sources** section of the Configuration panel, select **Add data source**.

   c. In the **Add data source** page, navigate to **MetricBase** > **Monitorable Item** and select the data to display, and then choose **Add this source**.
   The data is added to the dashboard.

   d. Use the controls in the **Data** section of the Configuration panel to select the metric and configure the query.
5. When done customizing, select **Save** and then **Exit editing mode.**
6. To return the template to the default (Certified) version, use the **More actions** menu to select **Return to certified**.

Add Health Log Analytics data

Add a graph showing logs from Health Log Analytics (HLA) in a Service Observability dashboard.

About this task

You can add logs from HLA to your Service Observability dashboards. These charts display with data scoped to the selected service.

Before you begin

Role required: sn_sow_svcobs.admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.
2. Open the template in editing mode.
   - If you're editing a certified template, select **Duplicate**.
   - If you're editing a custom template, select **Edit**.

   Note: Duplicating a certified template keeps you from overwriting it and also lets you reinstall it.

   The new dashboard is titled with the words `- Copy` appended. Use the pencil icon to change the dashboard name.
3. To do basic editing, such as rearranging, resizing, or deleting charts, follow the instructions for Edit in-line Platform Analytics dashboard elements.
4. To add a new chart with HLA data, follow these steps:

   a. In Edit mode, select **Add new element** and choose **Data visualization**, select **New Visualization**, and then select **Line**.

   b. In the **Data sources** section of the Configuration panel, select **Add data source**.

   c. In the **Add data source** page, navigate to **Health Log Analytics** > **Log components**.

   d. Enter a saved search or select a component and then optionally add conditions or a Lucene query, as described in Step 5 of Create log data dashboards and visualizations in Health Log Analytics.

   e. Once the filter is correct, choose **Add this source**.
   The data is added to the dashboard.

   f. Use the controls in the **Data** section of the Configuration panel to further configure the chart.
5. When done customizing, select **Save** and then **Exit editing mode.**
6. To return the template to the default (Certified) version, use the **More actions** menu to select **Return to certified**.

Add Splunk Enterprise data to Service Observability dashboard templates

Add Splunk Enterprise data to charts on Service Observability dashboard templates when you want to view those metrics in context of Service Observability.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

You can add metrics that are stored in Splunk Enterprise to your Service Observability dashboards. These charts display with data from the query you add to the chart.

Note: The results from these queries are not automatically scoped to the selected service.

Before you begin

You need a connection to your Splunk Enterpise instance. See Connect a Service Observability data source for more information.

Role required: sn_sow_svcobs.admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.
2. Open the template in editing mode.
   - If you're editing a certified template, select **Duplicate**.
   - If you're editing a custom template, select **Edit**.

   Note: Duplicating a certified template keeps you from overwriting it and also lets you reinstall it.

   The new dashboard is titled with the words `- Copy` appended. Use the pencil icon to change the dashboard name.
3. To do basic editing, such as rearranging, resizing, or deleting charts, follow the instructions for Edit in-line Platform Analytics dashboard elements.
4. To add a new chart with MetricBase data, follow these steps:

   a. In Edit mode, select **Add new element** and choose **Data visualization**, select **New Visualization**, and then select **Line**.

   b. In the **Data sources** section of the Configuration panel, select **Add data source**.

   c. In the **Add data source** page, navigate to **Service Observability** > **Splunk Enterprise Metrics**, paste in a query from an existing Splunk chart, and then choose **Add this source**. The data is added to the dashboard.

   d. Use the controls in the **Data** section of the Configuration panel to select the metric and configure the query.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. When done customizing, select **Save** and then **Exit editing mode.**
6. To return the template to the default (Certified) version, use the **More actions** menu to select **Return to certified**.

Use synthetic monitoring with Service Observability

Create a synthetic monitor to test your service's endpoints. Then display the results by customizing a Service Observability dashboard.

Overview of using synthetic monitoring with Service Observability

Synthetic monitoring simulates user transactions on a service's HTTP endpoints, enabling you to identify performance bottlenecks, help confirm up-time, and optimize user experiences. It continually calls a service's endpoint to test availability, response times, or the presence of a defined string in the response body. When a test fails the configured criteria, it can send an alert. For more information about synthetic monitoring, see Exploring synthetic monitoring.

The results of the monitor's test can be added to the Service Observability Overview dashboard and to the Observability Service dashboard. To view synthetic data related to a service on a dashboard, you first create a monitor that tests the endpoints of your service. Then you customize a dashboard template to add the monitor's test results.

Note: When you customize the template to add synthetic monitoring, you're adding it to all dashboards that use that template. Services that don't have a configured monitor display an empty chart.

1. Create a synthetic monitor to test a service.

   Install and configure synthetic monitoring and then create a monitor to continually test your service's endpoints using status code, response time, or response body as criteria.

2. Add synthetic monitor results to a Service Observability dashboard

   Customize a Service Observability dashboard template to add synthetic monitoring test results.

Create a synthetic monitor to test a service

Install and configure synthetic monitoring and then create a monitor for your service.

Before you begin

Your service must have an HTTP-reachable endpoint, configured as a CI in the HTTP(S) Endpoints [cmdb_ci_endpoint_http] table of the Configuration Management Database (CMDB).

Role required: n_sow_svcobs.admin, sn_sow_synthetics.synthetics_admin

Procedure

1. Install synthetic monitoring.
   See Install synthetic monitoring for more information.
2. If needed, create a location to run the monitor from.
   Monitors can be run locally from a ServiceNow Glide instance or from your environment using either a cluster of Agent Client Collector (ACC) agents or from a MID Server. If you choose to run them from your environment, you must create a location. See Create synthetic monitoring locations for more information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Create the monitor.
   Follow the steps in Create and edit a synthetic monitor. Be sure to select the service you want to monitor as the **Related service CI**.

What to do next

Once you create the monitor, it begins running tests. View the results of those tests to verify the results before embedding synthetic monitoring into a Service Observability dashboard. For more information about viewing monitor test results, see Identifying system issues with synthetic monitoring.

Once you have expected results, move on to Add synthetic monitor results to a Service Observability dashboard.

Add synthetic monitor results to a Service Observability dashboard

Add a synthetic monitor to a dashboard in Service Observability by customizing an existing dashboard template.

About this task

Once you have created monitors for your services, you can display those results on your Service Observability dashboards by customizing the template to include a chart that queries synthetic monitoring. You can add synthetic monitoring to either the Overview dashboard or the Observability > Service dashboard.

Note: When you customize the template to add synthetic monitoring, you're adding it to all dashboards that use that template. Services that don't have a configured monitor display an empty chart.

Before you begin

Role required: sn_sow_svcobs.admin

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.
2. Open the template in editing mode.
   - If you're editing a certified template, select **Duplicate**.
   - If you're editing a custom template, select **Edit**.

   Note: Duplicating a certified template keeps you from overwriting it and also lets you reinstall it.

   The new dashboard is titled with the words `- Copy` appended. Use the pencil icon to change the dashboard name.
3. To add a new ServiceNow table, follow these steps:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In Edit mode, select **Add new element** and choose **List**.

b. In the **Configuration panel** in the **Data sources** section, choose **Add data source**.

c. In the **Add data source** page, search for and select the synthetic monitoring data source, depending on the data you want to display.

- If you want to display the status of all tests associated with a service, select the `Single API` [sn_synthetics_check_single_api] table.
- If you want to display the results of a single synthetic test for a service, select the `Single API Result` [sn_synthetics_check_single_api_result] table.

Note: Leave the **Follow filters** check box selected. This control enables the chart to display information for the selected service.

4. In the **Default display** section, choose **Add** to begin adding columns from the table to the new list.
5. Choose **Save** and then **Exit editing mode** to return to the dashboard.
The synthetic monitoring data now appears on the dashboard.

What to do next

For detailed instructions about editing dashboards, see Edit in-line Platform Analytics dashboard elements.

Monitoring services and investigating issues

Operators use Service Observability to monitor the health of their services and to triage and manage incidents. They can view performance metrics from external observability monitoring systems in the context of an application service and its related CIs (Configuration Item).

Overview of using Service Observability

Service Observability can display metrics from an observability monitoring instance in context of a service. Along with metrics for the service and information about related CIs, Service Observability can infer entities related to the service, like a host or database, and display their metrics. This functionality gives an operator a more holistic view of what might be affecting performance.

The **Overview** tab on the Service Details page shows high-level performance metrics for the service and other contextual information such as open alerts or incidents affecting the service. The **Observability** tab displays more detailed metrics from the observability system for the service, along with metrics for inferred entities that are found through data mappings configured during setup. It also shows information about CIs related to those inferred entities. Instead of having to switch to an external tool, an operator can view those external metrics in the context of the Service Operations Workspace (SOW).

Note: Operators can also use the Analyze a dashboard in Service Observability and the Analyze service health in Service Observability AI skills to generate an analysis for them.

- View overall service health

  In the **Overview** tab of the Service Details page in the SOW, view overall service health and related information, like alerts and changes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- View service health metrics

  In the **Observability** tab, view more detailed service metrics, as well as metrics from related entities such as databases, hosts, and network components.

Analyze a dashboard in Service Observability

Use the analyze Service Observability dashboard skill to view an analysis of a single Service Observability dashboard.

About this task

The analyze Service Observability dashboard skill uses generative AI to summarize a single Service Observability dashboard, including calling out insights found in the charts. This summary helps operators quickly orient to and understand a particular dashboard.

When you ask the skill to analyze a dashboard, it generates an optimized version of the Service Observability charts on the dashboard and then analyzes the data, taking the time period shown into account. It then reports any insights it finds.

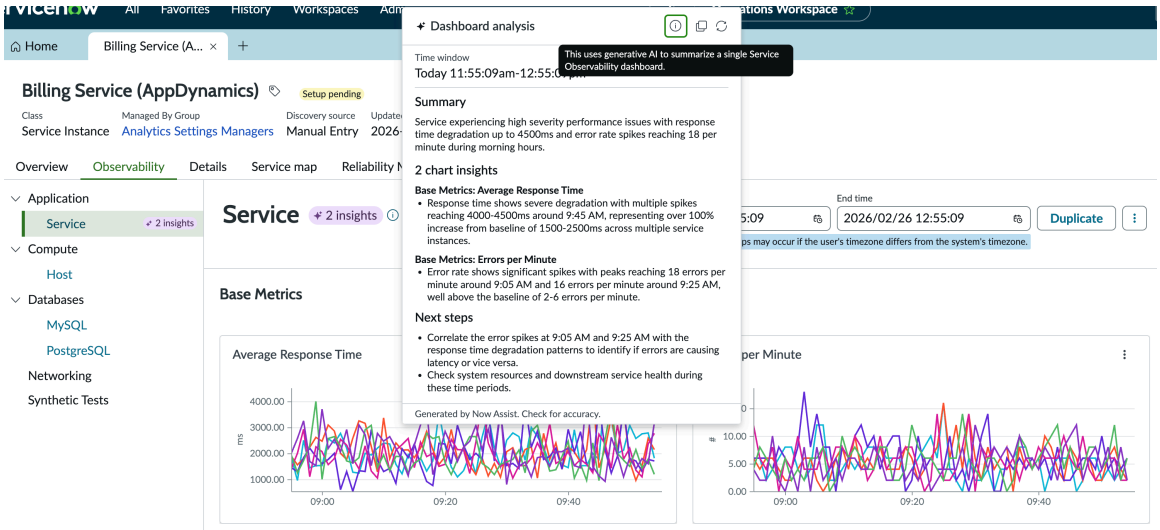

The skill's results remain available for an hour and don't automatically refresh. You need to rerun the skill to view a new analysis.

Note the following:

- The analysis captures the time range as reflected by the user's set time zone. When sharing analysis data across time zones, be aware that the analysis time period always reflects the time zone originally captured and may not be in sync with your time zone settings.
- The analysis doesn't weigh insights with regards to possible impact to an alert, incident, or overall dashboard health.
- The skill has a maximum number of 50 images that can be analyzed.

Before you begin

You need to manually activate this skill before it can be used. See Activate the analyze Service Observability dashboard skill for more information.

Role required: admin, operator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.
2. To the right of the dashboard name, select **Analyze dashboard**.
   The skill analyzes the data in Service Observability charts, looking for insights in the data that happened within the currently selected time period. Once complete, a badge displays one of the following:
   - **Analysis**: The skill didn't find any insights, but will display a summary of the dashboard.
   - **Insights**: A count of insights found is shown.
   - **Analysis unavailable**: The skill wasn't able to run an analysis.

   Note: Because the skill only analyzes the data in the time period captured by the screen shot, it may produce more accurate results if you lengthen the time period before analyzing. Note that longer time periods result in longer response time. For best results, use time periods less than 48 hours.
3. Select the info icon next to the badge to review the analysis and possible insights.
   The analysis tells you where it found the insights and the next steps you can take to determine the source of the insights.

   The analysis badge remains on the page for 60 minutes, after which the button resets so you can run a new analysis. After five minutes of running or if you change the time period, the analysis shows you a warning that it's out of date. Select the refresh icon to rerun it.

   Note: Select the **Copy** icon to copy the analysis, for example to paste it into work notes for an incident or alert.

Analyze service health in Service Observability

Use the Analyze Service Health skill to view an analysis of all Service Observability dashboards for the selected service.

About this task

The analyze service health skill uses generative AI to analyze all available Service Observability dashboards for the selected service.

When you ask the skill to analyze a service's health, it generates an optimized version of the Service Observability charts for a service, whether or not they are currently displayed. It then analyzes the data, and taking the time period shown into account, detects reports any insights it finds, and analyzes general service health. With these insights on service health, operators can assess service impact, identify root causes, and streamline incident response.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

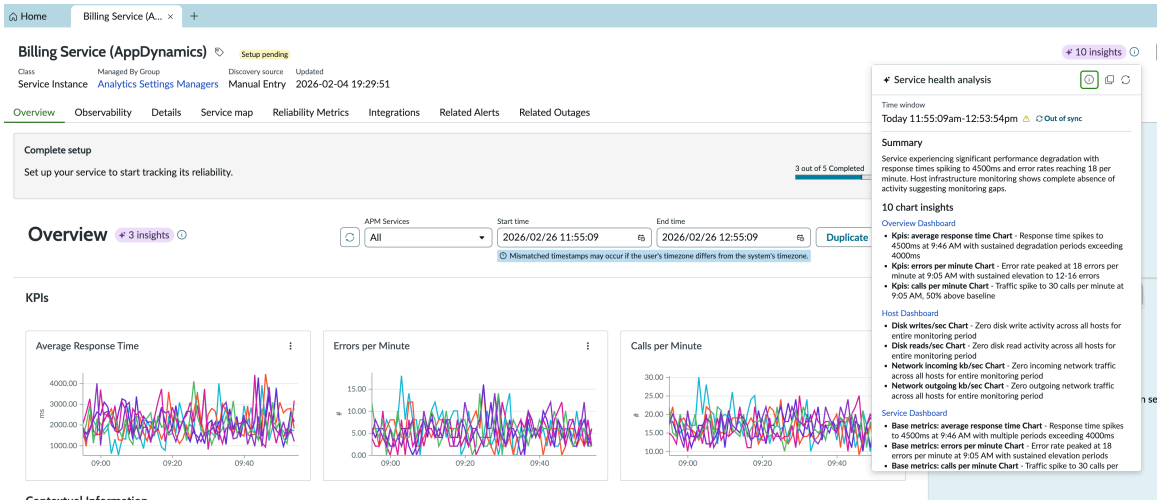

The skill's results remain available for an hour and don't automatically refresh. You need to rerun the skill to view a new analysis.

Once the skill is activated, it automatically runs for you in incident investigation from the **Investigate** tab in Incident Management. See Digital End-User Experience and Service Observability UI experience on investigate tab for more information.

Note the following:

- The analysis captures the time range as reflected by the user's set time zone. When sharing analysis data across time zones, be aware that the analysis time period always reflects the time zone originally captured and may not be in sync with your time zone settings.
- The analysis doesn't weigh insights with regards to possible impact to an alert, incident, or overall dashboard health.
- The skill has a maximum number of 50 images that can be analyzed.

Before you begin

You need to manually activate this skill before it can be used. See Activate the analyze service health skill for more information.

Role required: admin, operator

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.
2. At the top-right of the record page, select **Analyze service**.
   The skill optimizes the data on all existing Service Observability charts in dashboards and analyzes the data, looking for insights in the data that happened within the currently selected time period. Once complete, a badge displays one of the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Analysis**: The skill didn't find any insights, but will display a summary of the dashboard.
- **Insights**: A count of insights found is shown.
- **Analysis unavailable**: The skill wasn't able to run an analysis.

Note: Because the skill only analyzes the data in the time period captured by the screen shot, it may produce more accurate results if you lengthen the time period before analyzing. Note that longer time periods result in longer response time. Additionally, dashboards with more than 10,000 points can't be included as part of the analysis. Reducing very large time windows can help keep this from occurring. For best results, use time periods less than 48 hours.

3. Select the info icon next to the badge to review the analysis.
   The analysis tells you where it found the insights and the next steps you can take to determine the source of the insights.

   The analysis badge remains on the page for 60 minutes, after which the button resets so you can run a new analysis. After five minutes of running or if you change the time period, the analysis shows you a warning that it's out of date. Select the refresh icon to rerun it.

   Note: Select the **Copy** icon to copy the analysis, for example to paste it into work notes for an incident or alert.

View overall service health

View overall service health and related events, like alerts and changes, with the **Overview** tab on the Service Details page in the SOW.

Before you begin

A data mapping must be configured. For more information, see Create and manage data mappings.

Role required: operator [snc_sow_svcobs.manager]

About this task

When you have been notified that a service is experiencing performance issues, you can begin your investigation with the **Overview** tab on the Service Details page.

Note: Operators can also use the Analyze a dashboard in Service Observability and the Analyze service health in Service Observability AI skills to generate an analysis for them.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:
   - **Services** list: Choose a service from the list.
   - **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
   - **List**: Navigate to **Application Services** > **Services** and select a service.
   - **Express list** alert: Select a service from the **Impacted services** column.

   The Service Details page opens and the Overview tab is displayed.

   Trouble?
   If charts are displaying error messages, see Chart error states.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Use the charts in the Overview section to help determine where there might be performance issues.
   The Overview section shows the basic rate, error, and duration (RED) metrics along with metrics for key transactions found on the service. Hover over a metric to view more details.

   Note: If more than one service is found with the same name, select an individual service using the **All services** drop-down menu.

3. Change the time period using the time picker.
4. Use the other charts on the page to view related information about the service, such as recent changes, open incidents, and trends for incidents and alerts.
   These charts help explain whether recurring incidents or alerts might be caused by performance issues on the service. For detailed information about using this page, refer to the corresponding Overview templates for your observability vendor described in Service Observability templates.

What to do next

To view more detailed metrics and metrics for related entities, select the **Observability** tab. For more information, see View service health metrics.

View service health metrics

View detailed metrics for a service, as well as metrics from related entities, such as databases and hosts, on the **Observability** tab in SOW.

About this task

The **Observability** tab on the Service details page displays metrics for related entities inferred from external observability instances. You can view related metrics without leaving the ServiceNow platform.

The metrics are grouped by the following entity types:

- Application metrics: Metrics associated with the application services
- Computer metrics: Metrics associated with the hosts of the service
- Databases: Metrics associated with databases used by the service

For example, say you have metrics from New Relic used to monitor services, databases, and hosts. Say that you also have an issue with the `checkout` service and want to view performance information and metrics related to that service. Service Observability can show you the service, database, and host metrics related to that service that are coming from New Relic, along with related CI information from the ServiceNow platform. Each entity type has its own dashboard and each metric and related CI information has its own chart.

Note: Operators can also use the Analyze a dashboard in Service Observability and the Analyze service health in Service Observability AI skills to generate an analysis for them.

Before you begin

A data mapping must be configured. For more information, see Create and manage data mappings.

Role required: operator [snc_sow_svcobs.manager]

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** and then navigate to a service record.
   You can access a service record from these pages in the SOW:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Services** list: Choose a service from the list.
- **Service dashboard**: Choose a service in the dashboard and select **Service Details**.
- **List**: Navigate to **Application Services** > **Services** and select a service.
- **Express list** alert: Select a service from the **Impacted services** column.

The Service Details page opens and the Overview tab is displayed.

Trouble?
If charts are displaying error messages, see Chart error states.

2. Select the **Observability** tab.
   The **Observability** tab shows the related metrics and CI information available for the service. Use the navigation pane to view the entity dashboards.

   Note: If you don't see navigation options for an entity, then the metrics for that entity haven't been discovered or mapped.

3. Select an entity to view its dashboard.
   You can do the following with the charts:
   - Display metrics for a single entity by electing the **All [entity type]** drop-down list and choosing an entity.

     Note: If more than one entity returned metrics, charts display multiple lines with one line for each entity.

   - Display metrics from a different time period using the time picker. The default time period is the last hour.
   - Refresh the charts to the current time period of the last hour by selecting the **Refresh** button.
   - Hover over a point to view detailed information.

   For detailed information about using this page, refer to the corresponding Observability templates for your observability vendor described in Service Observability templates.

Service Observability reference

Reference topics provide additional information about administering and using Service Observability.

Components installed with Service Observability

Several types of components are installed with activation of the Service Observability plugin, including tables and user roles.

Roles installed

Role title [name]	Description	Contains roles
Operator/ Operations manager [snc_sow_svcobs.manager]	Triages incidents in the SOW. Views basic health metrics for a service and more detailed service metrics from related entities.	• sn_sow_srm.srm_responder • connection_admin • evt_mgmt_connector_instance_write

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description	Contains roles
Service Observability admin [sn_sow_svcobs.admin]	Configures data sources and adds data mappings. Customizes dashboard templates.	• sn_sow_srm.srm_admin • connection_admin • evt_mgmt_connector_instance_write • sn_sow_itsm_admin.sow_admin_user

Store applications installed

The following are installed when you install Service Observability.

Store application [plugin]	Description
Service Observability UI [sn_service_obs_ui]	The Overview and Observability tabs in the Service Details page of the SOW.
Service Reliability Management	Service Reliability Management and Service Observability display service performance metrics together to give a holistic view into service health.
Platform Analytics	Enables customizing the dashboards.

Tables installed

Table	Description
Dashboard fragments [sn_sow_svcobs_dashboard_fragments]	Contains the definitions of dashboards in JSON format.
Service to data mapping configuration map [sn_sow_svcobs_svc_to_config_map]	Associates data mapping configurations with a registered service using the CMDB Group Contains Configuration Item [cmdb_group_contains_ci] table.
Data mapping configurations [sn_sow_svcobs_mapping_config]	Defines a data mapping configuration for a service.
Data mapping configuration rules [sn_sow_svcobs_mapping_config_rules]	Defines a single tag-based mapping rule.
Observability data source [sn_sow_svcobs_data_source]	Defines a third-party observability vendor data source (an `http_connections`) that can be used to get telemetry for services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Service Observability entity records [sn_sow_svcobs_entity_records]	Records configuration items and unknown monitored objects for entities that appear in dashboards.
Service Observability PA Dashboards sn_sow_svcobs_pa_dashboards	Contains reference information for customized dashboards.

Domain separation and Service Observability

Domain separation is supported for Service Observability. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

Domain separation and Service Observability overview

Domain separation is present in all aspects of the synthetic monitoring application. Users belonging to a specific domain see only the data existing in their own domain.

How domain separation works in Service Observability

When data is separated by domains in Service Observability, users can only configure and see the observability data in their own domain. By default, all users and records are set to the parent domain unless the admin assigns them to a specific domain.

You define the domain-separated environment when you configure your data inputs. All relevant records and data processing in the Service Observability program flow reside in the same domain as the data input. The domain name of a data input appears in the Domain column that is displayed in the tables in your instance.

Service Observability manages all aspects of the data, such as the data mappings and the data displayed on dashboards. When a property is changed, the new value affects new sources only in the specific domain.

Related topics

Domain separation for service providers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Observability vendor entity mappings for Service Observability

Understand how Service Observability maps service, host, and database entities to your observability vendor resources.

Service Observability displays metrics from your observability vendor for services, hosts, databases, and network components on the Observability dashboards based on the key:value pairs in the mapping rules you create during configuration. Service Observability sends a request to the observability vendor using that mapping as a filter to find related entities. Any additional filtering needed to find the entities is noted in the following sections.

Amazon CloudWatch entity mapping

Resources are returned using the AWS `GetResources` API.

Service Observability entity category	Service Observability entity dashboard	AWS resource
Application	API Gateway - HTTP	API Gateway HTTP APIs
	API Gateway - REST	API Gateway REST APIs
	ELB	ELB application load balancers
	Lambda	Lambda functions
Compute	Host	EC2 instances
Databases	RDS	RDS database instances

AppDynamics entity mapping

Resources are returned using the value of the `entityName` property.

Service Observability entity category	Service Observability entity dashboard	AppDynamics resource
Application	Service	AppDynamics applications returned by the `/controller/rest/applications/` API
Compute	Host	Server nodes for applications returned by the `/controller/sim/v2/user/machines/keys/` API
Databases	MySQL	MySQL database instances returned by the `/controller/rest/databases/collectors/`
	PostgreSQL	MySQL database instances returned by the `/controller/`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Observability entity category	Service Observability entity dashboard	AppDynamics resource
		`rest/databases/collectors/`

Azure entity mapping

Resources are returned using the Azure `ResourceGraph` API.

Service Observability entity category	Service Observability entity dashboard	Azure resource
Application	Service	• `Microsoft.Compute/cloudServices` • `Microsoft.Web/sites`
Compute	Host	`Microsoft.Compute/virtualMachines`
Databases	MySQL	• `Microsoft.DBforMySQL/servers` • `Microsoft.DBforMySQL/flexibleServers`
	PostgreSQL	• Microsoft.DBforPostgreSQL/servers • `Microsoft.DBforPostgreSQL/flexibleServers`

Datadog entity mapping

Service Observability entity category	Service Observability entity dashboard	Datadog resource
Application	Service	Entities returned from the Software Catalog: List Entities API
Compute	Host	Hosts returned from the Hosts: List Hosts API
Databases	MySQL	Databases returned by filtering the `mysql.net.max_connections`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Observability entity category	Service Observability entity dashboard	Datadog resource
		metric, filtered by the provided key:value pair in the data mapping. **Note:** If your databases don't emit this metric, they aren't mapped.
	PostgreSQL	Databases returned by filtering the `postgresql.connections` metric, filtered by the provided key:value pair in the data mapping. **Note:** If your databases don't emit this metric, they aren't mapped.

Items of note:

- Service entities: The `Software Catalog list entities` API only returns data for services that include metadata. If you want to map services that don't include metadata, you must create a mapping using `service` as the tag and the name of the service as the value.

  For example, say you have a service named `internet-banking-4` that you want to use in a mapping and it doesn't contain metadata. You set up the mapping as shown in this screenshot.

  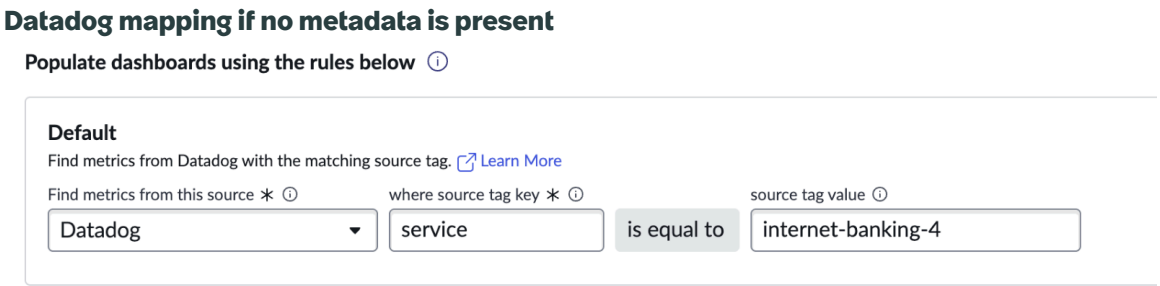

- Default dashboard templates: The Requests, Errors, and Latency charts on the Overview and Observability dashboard templates are created using the Datadog `trace.http.request` trace metric. If a service isn't emitting that metric, no data is found. You can customize the template to use a different trace metric query. See Customize Service Observability dashboard templates for more information.

Dynatrace entity mapping

You can use either the Dynatrace Classic query or Grail (DQL) query, however both require separate data connections. For Grail queries, Service Observability first queries the Grail data source and if not found, will fall back to a Classic query.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Observability entity category	Service Observability entity dashboard	Dynatrace resource
Application	Service	Services
Compute	Host	Hosts
Databases	MySQL	MySQL database instances
	PostgreSQL	PostgreSQL database instances

LogicMonitor entity mapping

Resources are returned using the Get Devices API, using specific property values on the entities as a filter.

Service Observability entity category	Service Observability entity dashboard	LogicMonitor resource
Application	Service	Entities with the `predef.externalResourceType` auto property set to `Service`
Compute	Host	Entities with the `predef.externalResourceType` auto property set to `PhysicalServer`, `Container`, or `VirtualMachine`
Databases	MySQL	Entities with the `system.categories` custom property set to `MySQL`
	PostgreSQL	Entities with the `system.categories` custom property set to `PostgreSQL`

New Relic entity mapping

Service Observability entity category	Service Observability entity dashboard	New Relic resource
Application	Service	Application services
Compute	Host	Hosts
Databases	MySQL	MySQL database instances
	PostgreSQL	PostgreSQL database instances

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prometheus entity mapping

Service Observability entity category	Service Observability entity dashboard	Prometheus resource
Application	Service	Applications
Compute	Host	Server nodes for applications
Databases	MySQL	MySQL database instances
	PostgreSQL	PostgreSQL database instances

SolarWinds entity mapping

Service Observability entity category	Service Observability entity dashboard	SolarWinds resource
Application	Service	Application services
Compute	Host	Hosts
Networking	Firewall	By default, data mappings find and display metrics using Palo Alto Firewall keys. If you want metrics from a different firewall, follow the instructions for customizing dashboard templates.
	Load Balancer	By default, data mappings find and display metrics using F5 Big-IP Load Balancer keys. If you want metrics from a different load balancer, follow the instructions for customizing dashboard templates.
	Interface	Network interfaces
	Other Network Devices	Other network devices, such as switches and routers

Items to note:

- You can use custom properties for key/value pairs in a data mapping.
- If you are using SolarWinds as an exception in your mapping to ingest networking metrics, using custom properties is optional. If you do use them, only entities with those properties in their metadata are returned. If you leave the tag and key values blank, all supported entities are returned.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Splunk entity mapping

Resources are returned using the Splunk Metric time series Metadata API. Service Observability searches for matching key:value pairs in custom properties and then falls back to searching dimensions.

The returned payload is then filtered by the presence of a specific metric in the metadata that corresponds with an entity type.

Service Observability entity category	Service Observability entity dashboard	Splunk Property or dimension	Splunk metric used for filtering
Application	Service	• Property: `sf_service` • Dimension fallback: `service.name`	`sf_metric:service.request`
Compute	Host	• Property: `host.name` • Dimension fallback: `host.id`	`sf_metric:disk.summary_util`
Databases	MySQL	• Property: `mysql.instance.name` • Dimension fallback: `mysql.instance.endpoint`	`sf_metric:mysql.threads`
	PostgreSQL	Not supported	

Zabbix entity mapping

In Zabbix, every monitored entity is categorized as a host. To enable differentiation between entities, Service Observability searches for keywords in attributes on metadata tags, including host tags, inherited tags, templates, and host groups. Any Zabbix host that does not match a keyword are classified and displayed as Host entities.

The following table shows the keywords used to find distinct entities. Keywords are case insensitive.

Service Observability entity category	Service Observability entity dashboard	Metadata keyword
Application	Service	• Apache • nginx

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Observability entity category	Service Observability entity dashboard	Metadata keyword
		• Kubernetes • cluster • application
Compute	Host	Any non-matching keyword
Databases	MySQL	• Mysql • Msql • mariadb
	PostgreSQL	• Postgres • Psql • Pgsql • Postgresql
Networking	Firewall	• Firewall • palo-alto • palo alto • pan • fortigate • fortinet • fgt • sophos • sonicwall • watchguard • zscaler By default, data mappings find and display metrics using Palo Alto Firewall keys. If you want metrics from a different firewall, follow the instructions for customizing dashboard templates.
	Load Balancer	• Loadbalancer • load_balancer • alb • nlb • f5

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Observability entity category	Service Observability entity dashboard	Metadata keyword
		• netscaler • big-ip • bigip By default, data mappings find and display metrics using F5 Big-IP Load Balancer keys. If you want metrics from a different load balancer, follow the instructions for customizing dashboard templates.
	Other Network devices	• Network • router • rt • gateway • gw • juniper • junos • switch • catalyst • nexus • arista

Item to note:

The default dashboards for each entity type display metrics using Zabbix standard keys. If your implementation uses custom item keys, you need to customize your dashboard templates accordingly.

Advanced query support for AWS and Azure

Understand how to create advanced queries for Service Observability dashboard charts.

Advanced queries in Service Observability support most vendor-specific query languages. Simply copy a query from the observability chart and paste it into the Advanced query field in Service Observability. However, queries to AWS and Azure must be JSON-formatted and follow the guidelines outlined in this topic.

Note: As of version 1.10, you can import charts directly from AWS and Azure. See Edit data charts on Service Observability dashboard templates for more information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS advanced queries

Data from AWS is accessed using the `GetMetricData` API. Queries can be either a direct request in JSON format to that API or they can be in the form of a `DashboardBody` metric widget. For example, you can use the widget representation for a chart that exists on an AWS CloudWatch dashboard.

Example: Direct GetMetricData request

The following example shows a `GetMetricData` JSON query for a chart that displays the average CPU Utilization for an ELB instance.

```
{

    "StartTime": 1637061900, // can be replaced with “${START}”
    "EndTime": 1637074500, // can be replaced with “${END}”
    "MetricDataQueries": [
        {
            "Expression": "SELECT AVG(CPUUtilization) FROM
 SCHEMA(\"AWS/EC2\", InstanceId)",
            "Id": "q1",
            "Period": 300,
            "Label": "Cluster CpuUtilization"
        },
        {
            "Id": "m1",
            "Label": "Unhealthy Behind Load Balancer",
            "MetricStat": {
                "Metric": {
                    "Namespace": "AWS/ApplicationELB",
                    "MetricName": "UnHealthyHostCount",
                    "Dimensions": [
                        {
                            "Name": "LoadBalancer",
                            "Value":
 "app/EC2Co-EcsEl-EXAMPLE69Q/fdd2210e799e4376" // can be
 replaced with “${ENTITIES}” or “${ENTITIES_ELB}”
                        }
                    ]
                },
                "Period": 300,
                "Stat": "Average"
            }
        }
    ]
}
```

Example: Copied source for a single object in an AWS CloudWatch dashboard

This example shows the JSON copied directly from a single object in the widgets list of the Source view for a CloudWatch dashboard.

```
{
  "type": "metric",
  "x": 9,
  "y": 0,
  "width": 15,
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
  "height": 5,
  "properties": {
    "view": "timeseries",
    "metrics": [
      [ "AWS/ApplicationELB", "RequestCountPerTarget",
 "TargetGroup", "targetgroup/api-tg/1e3dc9z72fe21ca2",
 "AvailabilityZone", "us-west-1a" ],
      ["...", "targetgroup/api-tg-cs/82a1db5f950073e1", ".",
 "." ]
    ],
    "region": "us-west-1"
  }
}
```

Both types of queries can use template variables to represent a set of entities or the start and end time selected in the UI. For example, instead of querying for a single ELB instance, you can replace the `Value` field for the `LoadBalancer` dimension with `${ENTITY_ELB}`. The query is run against each ELB instance associated with the selected service in the Service Observability UI. The chart displays a time series for each one.

Note: JSON objects copied from the source of a CloudWatch dashboard don't include explicit time fields. The values from the time picker in the Service Observability UI are used automatically.

Template variables

AWS field	Template variable
`Dimensions[n].Value`	• `ENTITIES` • `ENTITIES_HOST` • `ENTITIES_LAMBDA` • `ENTITIES_API_GATEWAY_HTTP` • `ENTITIES_API_GATEWAY_REST` • `ENTITIES_ELB` • `ENTITIES_RDS` The `Dimensions[n].Value` field must be a resource identifier key.
`StartTime`	`START`
`EndTime`	`END`

Azure advanced queries

You run an advanced query for an Azure-based chart by copying the source from the Azure UI and pasting it in the **Advanced query** field.

To copy the source:

1. In the Azure UI, export and download the dashboard.
2. Copy the JSON for a single chart. Charts are located at `properties.lenses[i].parts[j].metadata.settings.content.options.chart`.
3. Paste the JSON into the query field. The JSON must include at least the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
```
{
“metrics”: [
  “name”: string,
  “resourceMetadata”: {
    “id”: string
    }
  ]
}
```
```

The `id` field must be either an Azure resource ID or a template variable.

Service Observability template variables

Understand the template variables that you can use in your queries when editing Service Observability dashboards and charts.

You can use the following template variables in your query. Variables enable the query to be used for multiple services, hosts, and database instances, as well as for the time period currently selected for the dashboard.

Note: Amazon CloudWatch uses slightly different template variables. See Advanced query support for AWS and Azure for more information.

Aside from the variables listed below, you can also use any tag key used in the chart's data mapping as a template variable.

Template variables

Variable	Description
`{$ENTITIES_HOST}`	Returns metrics for the selected host
`{$ENTITIES_MYSQL}`	Returns metrics for the selected MySQL instance
`{$ENTITIES_POSTGRESQL}`	Returns metrics for the selected PostgreSQL instance
`{$ENTITIES_SERVICE}`	Returns metrics for the selected service.
`{$ENTITIES}`	Returns metrics for the selected entity; either service, host or database.
`{$START_TIME}`	Returns metric time series using the selected start time.
`{$END_TIME}`	Returns metric time series using the selected end time.

Example query

Let's say you're using NewRelic and you create a data mapping that maps the key `service` to the value `checkout-service`.

A simple query built using the query builder might look like this:

```
rate(count(apm.service.transaction.duration), 1 minute
```

It would return the rate of transactions for the `checkout-service` service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Instead you might use a full vendor query like this:

```
SELECT rate(count(apm.service.transaction.duration), 1
 minute) as 'Web throughput' FROM Metric WHERE (entity.guid
 = 'NDc2NDMyNXxBUE18QVBQTE1DQVRJT058MTA3NjIyODQwMw') AND
 (transactionType = 'Web') LIMIT MAX SINCE 30 minutes ago
 TIMESERIES UNTIL now
```

It would return the duration for the hard-coded `NDc2NDMyNXxBUE18QVBQTE1DQVRJT058MTA3NjIyODQwMw` entity from the last 30 minutes.

To create a query that would return the time series for any selected service at any time period, you can replace the entity and times with template variables:

```
SELECT average(convert(apm.service.transaction.duration, unit,
 'ms')) as metricValue,
 average(convert(apm.service.transaction.duration, unit, 'ms'))
 - 100 as loop FROM Metric WHERE entity.guid IN (${ENTITIES})
 FACET entity.guid, entity.name SINCE ${START} UNTIL ${END}
 TIMESERIES LIMIT 25
```

Service Observability templates

View the templates for your observability vendor to understand the default dashboards provided by Service Observability.

See Customize Service Observability dashboard templates for information about customizing these templates.

Amazon CloudWatch templates for Service Observability

Templates used to create Amazon CloudWatch dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

Amazon CloudWatch Overview tab for Service Observability

Information that is displayed on the Amazon CloudWatch **Overview** tab of the Service Details page in the SOW.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Invocations	Aggregate rate of transactions through the service.	AWS
Errors	Number of transactions that have an error.	AWS
Duration	Aggregate duration of transactions through the service, in milliseconds.	AWS

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Status indicators section

The Status indicators section displays charts for related Configuration Items (CIs) that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
All active alerts	Listing of active alerts that affect this service.	CMDB

Amazon CloudWatch Observability tab for Service Observability

Dashboard and charts on the Amazon CloudWatch tab of the Service Details page in the SOW.

Application dashboards

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

API Gateway - HTTP

Chart	Description	Data source
Request Count	Aggregate rate of transactions through the service, per minute.	AWS
Latency	Aggregate duration of transactions through the service, in milliseconds.	AWS
4xx Errors	Number of errors returned with a status of 4xx.	AWS
5xx Errors	Number of errors returned with a status of 5xx.	AWS

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

API Gateway - REST

Chart	Description	Data source
Request Count	Aggregate rate of transactions through the service, per minute.	AWS
Latency	Aggregate duration of transactions through the service, in milliseconds.	AWS
4xx Errors	Number of errors returned with a status of 4xx.	AWS
5xx Errors	Number of errors returned with a status of 5xx.	AWS

ELB

Chart	Description	Data source
Request Count	Aggregate rate of transactions through the service, per minute.	AWS
Latency	Aggregate duration of transactions through the service, in milliseconds.	AWS
4xx Errors	Number of errors returned with a status of 4xx.	AWS
5xx Errors	Number of errors returned with a status of 5xx.	AWS

Lambda

Chart	Description	Data source
Errors	Number of errors reported by the Lambda function.	AWS
Duration	Aggregate duration of transactions through the service, in milliseconds.	AWS
Invocations	Number of Lambda invocations executed by this service.	AWS

Compute dashboards

This dashboard displays metrics for hosts related to the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Host

Chart	Description	Data source
CPU utilization	Percent of the host processing power being consumed by this service.	AWS
Disk Write Bytes	Number of bytes written to the database by this service.	AWS
Disk Read Bytes	Number of bytes read from the database by this service.	AWS
Network In Bytes	Amount of incoming network traffic in bytes for this service.	AWS
Network Out Bytes	Amount of outgoing network traffic in bytes for this service.	AWS
All EC2 VM Instances	Information for all EC2 VM instances the service is actively using. Select a link to view more detailed information.	CMDB
All Active Servers	Information for all servers the service is actively using. Select a Server link to view more detailed information.	CMDB

Database dashboards

This dashboard displays metrics for RDS databases related to the service.

RDS

Chart	Description	Data source
CPU utilization	Percent of the database's processing power being consumed by this service.	AWS
Free Memory	Amount of memory in bytes available on this database.	AWS
Free Storage	Amount of storage in bytes available on this database.	AWS
Current Connections	Number of connections from this service to the database.	AWS
Read Latency	Time taken in milliseconds for read operation from the database to complete.	AWS
Write Latency	Time taken in milliseconds for write operation from the database to complete.	AWS
All RDS Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

AppDynamics templates for Service Observability

Templates used to create AppDynamics dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AppDynamics Overview tab for Service Observability

Information that is displayed on the AppDynamics **Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Average Response Time	Average rate of transactions through the service.	AppDynamics
Errors per Minute	Rate of errors per minute on transactions for the service.	AppDynamics

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

AppDynamics Observability tab for Service Observability

Dashboard and charts on the AppDynamics Observability tab of the Service Details page in the SOW.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service

Chart name	Description	Data source
Average Response Time	Average rate of transactions through the service.	AppDynamics
Errors per Minute	Rate of errors per minute on transactions for the service.	AppDynamics
Calls per Minute	Number of calls to the server, per minute.	AppDynamics
Number of slow calls	Number of calls to the server that have a response time slower than the set slow threshold.	AppDynamics
Number of Very Slow Calls	Number of calls to the server that have a response time slower than the set very slow threshold.	AppDynamics

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU %Busy	Percent of the host processing power being consumed.	AppDynamics
Memory Used %	Percent of memory the host is using.	AppDynamics
Disk Reads/sec	Percent of disk space being used, per second.	AppDynamics
Disk Writes/sec	Percent of disk space being written to, per second.	AppDynamics
Network Incoming KB/sec	Number of kilobytes coming into the service, per second.	AppDynamics
Network Outgoing KB/sec	Number of kilobytes being sent by the service, per second.	AppDynamics
All Active Servers	Information for all servers the service is actively using. Select a host link to view more detailed information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL

Chart	Description	Data source
Calls Per Minute	Number of database calls made by the service, per minute	AppDynamics
Time Spent in Executions	Number of seconds spent executing database calls by the service.	AppDynamics
DB Inserts	Number of database inserts by the service.	AppDynamics
DB Reads	Number of database reads by the service.	AppDynamics
All MySQL Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

PostgreSQL

Chart	Description	Data source
Calls Per Minute	Number of database calls made by the service, per minute	AppDynamics
Time Spent in Execution(s)	Time in seconds the database in execution mode for this service.	AppDynamics
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

Azure Monitor templates for Service Observability

Templates used to create Azure Monitor dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

Azure Monitor Overview tab for Service Observability

Information that is displayed on the Azure Monitor **Overview** tab of the Service Details page in the SOW.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Requests	Aggregate rate of transactions through the service.	Azure
Response Time	Aggregate duration of transactions through the service, in milliseconds.	Azure
Http 5xx	Number of errors returned with a status of 5xx.	Azure

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Status indicators section

The Status indicators section displays charts for related Configuration Items (CIs) that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
All Active alerts	Listing of active alerts that affect this service.	CMDB

Azure Monitor Observability tab for Service Observability

Dashboard and charts on the Azure Monitor Observability tab of the Service Details page in the SOW.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service

Chart	Description	Data source
Requests	Aggregate rate of requests through the service.	Azure
Http 4xx	Number of transactions that have an HTTP 4xx response.	Azure
Response time	Aggregate duration of transactions through the service, in milliseconds.	Azure
Data Out	Volume in bytes of outgoing network traffic from this service.	Azure

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service (continued)

Chart	Description	Data source
Http 2xx	Number of transactions that have an HTTP 2xx response.	Azure
Data In	Volume in bytes of incoming network traffic to this service.	Azure

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
Available Memory Percentage	Percent of the memory available to the service.	Azure
CPU %	Percent of the host processing power being consumed.	Azure
VM Cached Bandwidth Consumed Percentage	Percent of the virtual machine's cached bandwidth that the service is consuming.	Azure
All Active VM Instances	Information for all virtual machines (VM) the service is actively using. Select a VM link to view more information.	CMDB
All Active Servers	Information for all hosts the service is actively using. Select a host link to view more detailed information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
Storage %	Percentage of storage currently being consumed by the service.	Azure
CPU %	Percentage of CPU currently being consumed by the service.	Azure
Memory %	Percentage of memory currently being consumed by the service.	Azure
Active Connections	Number of connections currently used by the service.	Azure
Aborted Connections	Number of connections from the service terminated by the database	Azure
MySQL Uptime	Total time in seconds that this MySQL instance has been running since its last restart.	Azure
Slow Queries	Number of queries that take longer than the specified threshold to execute. This threshold is typically defined by the `long_query_time` parameter in MySQL.	Azure

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PostgreSQL

Chart	Description	Data source
Storage %	Percentage of storage currently being consumed by the service.	Azure
CPU %	Percentage of CPU currently being consumed by the service.	Azure
Memory %	Percentage of memory currently being consumed by the service.	Azure
Active Connections	Number of connections currently used by the service.	Azure
Aborted Connections	Number of connections from the service terminated by the database	Azure
Succeeded Connections	Number of successful connections from the service to the database.	Azure
Deadlocks	Number of deadlocks encountered by the service.	Azure
Is Database Alive	1 if database is up, 0 if it is down.	Azure
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

Cisco ThousandEyes templates for Service Observability

Templates used to create Cisco ThousandEyes synthetics dashboard in Service Observability. You can edit this template as needed.

See Customize Service Observability dashboard templates for more information.

Cisco ThousandEyes Observability tab for Service Observability

Dashboard and charts on the Cisco ThousandEyes Observability tab of the Service Details page in the SOW.

Synthetic Tests dashboard

This dashboard displays synthetic monitoring metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Network Test

Chart	Description	Data source
Latency	Network latency measurements over time.	Cisco ThousandEyes
Jitter	Variation in latency measurements over time.	Cisco ThousandEyes
Packet loss %	Percentage of network packets lost during transmission.	Cisco ThousandEyes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HTTP Test

Chart	Description	Data source
Throughput	HTTP throughput measurements showing data transfer rates.	Cisco ThousandEyes
Response Time	HTTP response time measurements for the monitored service.	Cisco ThousandEyes

Page Load Test

Chart	Description	Data source
Time to complete	Time required to fully load the monitored page.	Cisco ThousandEyes
Error Count	Number of errors encountered during page load tests.	Cisco ThousandEyes

Datadog templates for Service Observability

Templates used to create Datadog dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

Datadog Overview tab for Service Observability

Information that is displayed on the Datadog**Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Request	Aggregate rate of transactions through the service.	Datadog
Errors	Percent of transactions that have an error.	Datadog
Latency	Aggregate duration of transactions through the service, in milliseconds.	Datadog

Transactions charts

The Transactions charts show metrics for specific key transactions on the service. The key transactions are determined by your observability system configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart name	Description	Data source
Requests	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	Datadog
Errors	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	Datadog
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	Datadog

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

Datadog Observability tab for Service Observability

Dashboard and charts on the Datadog Observability tab of the Service Details page in the SOW.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service KPIs

Chart	Description	Data source
Requests	Aggregate rate of requests through the service.	Datadog
Errors	Number of transactions that have an error.	Datadog
Latency	Aggregate duration of transactions through the service, in milliseconds.	Datadog

Top Transactions

Chart	Description	Data source
Requests	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	Datadog
Errors	Each line represents the number of errors on a key transaction in the service. Hover over a point to view the different transaction names.	Datadog
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	Datadog

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU Utilization	Percent of the host processing power being consumed.	Datadog
Memory Utilization	Percent of memory the host is using.	Datadog
Disk Utilization	Percent of disk space being used.	Datadog
All Active Servers	Information for all hosts the service is actively using. Select a host link to view more detailed information.	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Host (continued)

Chart	Description	Data source
All Active VM Instances	Information for all virtual machines (VM) the service is actively using. Select a VM link to view more information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
Max Connections	Maximum number of concurrent connections that can be established with the MySQL database server.	Datadog
Failed Connections	Number of failed attempts to connect to the MySQL server.	Datadog
Slow Queries Rate	Number of database queries that have been categorized as slow by Datadog.	Datadog
Data Reads	Number of read operations performed on the MySQL database.	Datadog
Data Writes	Number of write operations performed on the MySQL database.	Datadog
Cache Utilization	Percentage of the cache memory being used by the MySQL database.	Datadog
All MySQL Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

PostgreSQL

Chart	Description	Data source
Connections	Total number of connections to the database server, including all connection states such as active, idle, and others.	Datadog
Connections Used	Number of connections to the database used by this service.	Datadog
Active Waiting Queries	Number of queries in the queue waiting to be executed.	Datadog
Deadlocks	Number of deadlocks encountered by the service.	Datadog
Rollbacks	Number of transactions rolled back by the service.	Datadog
Buffer Hits	Number of times disk blocks were found in the buffer cache, so that a read wasn't necessary.	Datadog
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dynatrace templates for Service Observability

Templates used to create Dynatrace dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

Dynatrace Overview tab for Service Observability

Information that is displayed on the Dynatrace **Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Rate	Aggregate rate of transactions through the service.	Dynatrace
Error %	Percent of transactions that have an error.	Dynatrace
Latency	Aggregate duration of transactions through the service, in milliseconds.	Dynatrace

Key Transactions charts

The Transactions charts show metrics for specific key transactions on the service. The key transactions are determined by your observability system configuration.

Chart name	Description	Data source
Rate	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	Dynatrace
Error %	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	Dynatrace
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	Dynatrace

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

Dynatrace Observability tab for Service Observability

Dashboard and charts on the Dynatrace Observability tab of the Service Details page in the SOW.

Dynatrace dashboards are the same for both managed and SaaS environments.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service Basic metrics

Chart	Description	Data source
Rate	Aggregate rate of transactions through the service.	Dynatrace
Error %	Percent of transactions that have an error.	Dynatrace
Latency	Aggregate duration of transactions through the service, in milliseconds.	Dynatrace
Client error rate	Rate of errors the service clients experienced.	Dynatrace
Client throughput	Rate of client requests for this service.	Dynatrace

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service Basic metrics (continued)

Chart	Description	Data source
Client response time	Response time for client requests on this service, in milliseconds.	Dynatrace
Total DB connections	Number of databases that this client has connections to. **Note:** Not available with Grail (DQL) queries.	Dynatrace
DB connection errors	Number of errors received by the client when trying to access databases. **Note:** Not available with Grail (DQL) queries.	Dynatrace
Time spent in database calls	Duration of requests to databases, in milliseconds. **Note:** Not available with Grail (DQL) queries.	Dynatrace

Service Transactions

Chart	Description	Data source
Key Request Count	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	Dynatrace
Key Request Error %	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	Dynatrace
Key Request Response Time	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	Dynatrace

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU utilization	Percent of the host processing power being consumed.	Dynatrace
Memory utilization	Percent of memory the host is using.	Dynatrace
Disk utilization	Percent of disk space being used.	Dynatrace
Network RX - bytes received	Count of bytes received.	Dynatrace
Network TX - bytes transferred	Count of bytes sent.	Dynatrace
All Active Servers	Information for all hosts the service is actively using. Select a host link to view more detailed information.	CMDB
All Active VM Instances	Information for all virtual machines (VM) the service is actively using. Select a VM link to view more information.	CMDB

Process

Chart	Description	Data source
CPU Usage	Percent of the host processing power being consumed.	Dynatrace
Memory Usage	Percent of memory the host is using.	Dynatrace
Requests	Rate of requests per second	Dynatrace
Throughput	Size of requests, in megabytes per second.	Dynatrace
Latency	Aggregate duration of transactions, in milliseconds.	Dynatrace

Process Group

Chart	Description	Data source
CPU Usage	Percent of the host processing power being consumed.	Dynatrace
Memory Usage	Percent of memory the host is using.	Dynatrace
Requests	Rate of requests per second	Dynatrace
Throughput	Size of requests, in megabytes per second	Dynatrace
Latency	Aggregate duration of transactions, in milliseconds.	Dynatrace
All Active Servers	Information for all hosts the service is actively using. Select a host link to view more detailed information.	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
CPU utilization	Percent of the database's processing power being consumed.	Dynatrace
Memory utilization	Percent of memory the database is using.	Dynatrace
Availability	Percent of time the database is available.	Dynatrace
Connections	Number of connections to the database.	Dynatrace
Slow queries	Number of database queries that have been categorized as slow by Dynatrace.	Dynatrace
All MySQL Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

PostgreSQL

Chart	Description	Data source
Deadlocks	Number of deadlocks detected in this database.	Dynatrace
Committed Transactions	Number of transactions in this database that have been committed.	Dynatrace
Active Backend Connections	Number of backends currently connected to this database.	Dynatrace
Query Conflicts	Number of queries cancelled due to conflicts with recovery in this database.	Dynatrace
Active Time	Time spent executing SQL statements in this database.	Dynatrace
Fatal Sessions	Number of database sessions to this database that were terminated by fatal errors	Dynatrace
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

LogicMonitor templates for Service Observability

Templates used to create LogicMonitor dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

LogicMonitor Overview tab for Service Observability

Information that is displayed on the LogicMonitor **Overview** tab of the Service Details page in the SOW.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Response Time	Response time metrics for the service, displayed in milliseconds over time.	LogicMonitor
Status	Status metrics for the service. A value of `1` means the service is up. A value of 2 means it's down.	LogicMonitor

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

LogicMonitor Observability tab for Service Observability

Dashboard and charts on the LogicMonitor Observability tab of the Service Details page in the SOW.

Application dashboard

This dashboard displays performance metrics for the selected service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service basic metrics

Chart	Description	Data source
Response Time	Response time metrics for the service, displayed in milliseconds over time.	LogicMonitor
Status	Status metrics for the service. A value of `1` means the service is up. A value of `2` means it's down.	LogicMonitor

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU Busy %	Percentage of time the CPU is actively busy processing tasks, showing periods of high activity versus idle time.	LogicMonitor
Memory Usage	Percent of memory the host is using.	LogicMonitor
CPU Utilization	Overall CPU load and usage patterns across the host, showing continuous utilization levels over time.	LogicMonitor
Network RX - bytes transferred	Amount of data received by the host over the network, measured in bytes.	LogicMonitor
Network TX - bytes transferred	Amount of data transmitted by the host over the network, measured in bytes.	LogicMonitor
All Active Servers	Information for all servers the service is actively using. Select a host link to view more detailed information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
Response Time	Response time metrics for the database, displayed in milliseconds over time.	LogicMonitor
CPU Busy %	Percent of the database's processing power being consumed.	LogicMonitor
Memory Used %	Percent of memory the database is using.	LogicMonitor

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL (continued)

Chart	Description	Data source
Status	Status metrics for the service. A value of `1` means the database is up. A value of 2 means it's down.	LogicMonitor
All MySQL Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

PostgreSQL

Chart	Description	Data source
Response time	Response time metrics for the database, displayed in milliseconds over time.	LogicMonitor
CPU Busy %	Percent of the database's processing power being consumed.	LogicMonitor
Memory Used %	Percent of memory the database is using.	LogicMonitor
Status	Status metrics for the service. A value of `1` means the database is up. A value of 2 means it's down.	LogicMonitor
Busy servers	Number of busy servers for the database over time.	LogicMonitor
Maximum servers	Maximum number of servers available for the database.	LogicMonitor
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

MetricBase templates for Service Observability

Templates used to create MetricBase dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

MetricBase Overview tab

Information that is displayed on the MetricBase **Overview** tab of the Service Details page in the SOW.

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM).	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart name	Description	Data source
	For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

MetricBase Observability tab for Service Observability

Dashboard and charts on the MetricBase Observability tab of the Service Details page in the SOW.

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
All Active VM Instances	Information for all virtual servers the service is actively using. Select a host link to view more detailed information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
All MySQL Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

PostgreSQL

Chart	Description	Data source
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Networking dashboards

These dashboards display metrics for network devices related to the service.

Other Networking Devices

Chart	Description	Data source
All Router Instances	Information of all routers the service is actively using. Select a router link to view more detailed information.	CMDB
All Switch Instances	Information of all switches the service is actively using. Select a switch link to view more detailed information.	CMDB

New Relic templates for Service Observability

Templates used to create New Relic dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

New Relic Overview tab for Service Observability

Information that is displayed on the New Relic **Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Rate	Aggregate rate of transactions through the service.	New Relic
Error %	Percent of transactions that have an error.	New Relic
Latency	Aggregate duration of transactions through the service, in milliseconds.	New Relic

Transactions charts

The Transactions charts show metrics for specific key transactions on the service. The key transactions are determined by your observability system configuration.

Chart name	Description	Data source
Rate	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	New Relic
Error %	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	New Relic

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart name	Description	Data source
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	New Relic

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

New Relic Observability tab for Service Observability

Dashboard and charts on the New Relic Observability tab of the Service Details page in the SOW.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service basic metrics

Chart	Description	Data source
Rate	Aggregate rate of transactions through the service, per minute.	New Relic

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service basic metrics (continued)

Chart	Description	Data source
Error %	Percent of transactions that have an error.	New Relic
Latency	Aggregate duration of transactions through the service, in milliseconds.	New Relic
Time spent in database calls	Duration of requests to databases, in milliseconds.	New Relic
Database throughput	Rate of database requests for this service, per minute.	New Relic
Time spent in database queries	Time spent querying the database, in milliseconds.	New Relic

Service transactions

Chart	Description	Data source
Rate	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	New Relic
Error %	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	New Relic
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	New Relic

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU utilization	Percent of the host processing power being consumed.	New Relic
Memory utilization	Percent of memory the host is using.	New Relic

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Host (continued)

Chart	Description	Data source
Disk utilization	Percent of disk space being used.	New Relic
All Active Servers	Information for all servers the service is actively using. Select a host link to view more detailed information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
CPU utilization	Percent of the database's processing power being consumed.	New Relic
Memory utilization	Percent of memory the database is using.	New Relic
Availability	Percent of time the database is available.	New Relic
Connections	Number of connections to the database.	New Relic
Slow queries	Number of database queries that have been categorized as slow by New Relic.	New Relic
All MySQL Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

PostgreSQL

Chart	Description	Data source
Scheduled Checkpoints	Average number of scheduled checkpoints that were performed.	New Relic
Buffers Written for Checkpoint	Average number of buffers written during checkpoints.	New Relic
Buffers Written by Background Writer	Average number of buffers written by the background writer.	New Relic
Buffers Allocated	Average number of buffers allocated.	New Relic
Checkpoint Write Time	Time in milliseconds spent writing files to disk.	New Relic
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prometheus templates for Service Observability

Templates used to create Prometheus dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

Prometheus Overview tab for Service Observability

Information that is displayed on the Prometheus **Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Request Rate	Aggregate rate of transactions through the service, per second.	Prometheus
Error %	Percent of transactions that have an error.	Prometheus
Average Response Time	Average duration of transactions through the service, in seconds.	Prometheus

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
All Active alerts	Listing of active alerts that affect this service.	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Prometheus Observability tab for Service Observability

Dashboard and charts on the Prometheus Observability tab of the Service Details page in the SOW.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service KPIs

Chart	Description	Data source
Request Rate	Aggregate rate of transactions through the service, per second.	Prometheus
Error %	Percent of transactions that have an error.	Prometheus
Average Response Time	Aggregate duration of transactions through the service, in seconds.	Prometheus
Throughput	Count of successful transactions, per second.	Prometheus
Active Requests	Count of current in-flight requests to the service.	Prometheus
Response Time (95th Percentile)	95th percentile response time for HTTP requests, in seconds.	Prometheus
Average Request Size	Average size of requests to the service, in bytes.	Prometheus
Average Response Size	Average size of responses from the service, in bytes.	Prometheus

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU utilization	Percent of the host processing power being consumed.	Prometheus
Memory utilization	Percent of memory the host is using.	Prometheus
Disk utilization	Percent of the host disk being used.	Prometheus

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Host (continued)

Chart	Description	Data source
All Active Servers	Information for all servers the service is actively using. Select a host link to view more detailed information.	CMDB
All Active VM Instances	Information for all virtual machines the service is actively using. Select a host link to view more detailed information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
Thread Utilization %	Percent of the database's threads being consumed.	Prometheus
Queries Per Second	Number of queries made to the database by the service, per second.	Prometheus
InnoDB Memory %	Percent of InnoDB buffer pool being consumed.	Prometheus
Active Connections	Number of connections to the database.	Prometheus
Aborted Connections	Rate of aborted connections.	Prometheus
Uptime	Average number of hours the database has been available.	Prometheus
Inbound Network Traffic	Number of bytes per second of inbound traffic to the database.	Prometheus
Outbound Network Traffic	Number of bytes per second of outbound traffic to the database.	Prometheus
Slow Queries	Rate of slow queries detected by MySQL, per second.	Prometheus
All MySQL Instances	Information for all databases the service is actively using. Select a database link to view more detailed information.	CMDB

PostgreSQL

Chart	Description	Data source
Active Sessions	Count of sessions currently active.	Prometheus
Idle Sessions	Count of idle sessions.	Prometheus
Committed Transactions	Total committed transactions since server start.	Prometheus
Failed Transactions	Total rolled back transactions since server start.	Prometheus

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PostgreSQL (continued)

Chart	Description	Data source
Deadlocks	Total deadlocks detected since server start.	Prometheus
Locks Count	Current count of locks given out by the database.	Prometheus
Cache Hit Ratio	Percentage of database requests that can be served by the cache.	Prometheus
All PostgreSQL Instances	Information of all databases the service is actively using. Select a database link to view more detailed information.	CMDB

SolarWinds templates for Service Observability

Templates used to create SolarWinds dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

SolarWinds Overview tab for Service Observability

Information that is displayed on the SolarWinds **Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Rate	Aggregate rate of transactions through the service.	SolarWinds
Error %	Percent of transactions that have an error.	SolarWinds
Latency	Aggregate duration of transactions through the service, in milliseconds.	SolarWinds

Transactions charts

The Transactions charts show metrics for specific key transactions on the service. The key transactions are determined by your observability system configuration.

Chart name	Description	Data source
Rate	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	SolarWinds
Error %	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	SolarWinds
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	SolarWinds

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

SolarWinds Observability tab for Service Observability

Dashboard and charts on the SolarWinds Observability tab of the Service Details page in the SOW.

Note: The X axis on SolarWinds charts use Universal Time Code (UTC) instead of the user's timezone.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service basic metrics

Chart	Description	Data source
Rate	Aggregate rate of transactions through the service, per minute.	SolarWinds
Latency	Aggregate duration of transactions through the service, in milliseconds.	SolarWinds

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service basic metrics (continued)

Chart	Description	Data source
Time spent in database calls	Duration of requests to databases, in milliseconds.	SolarWinds
Time spent in database queries	Time spent querying the database, in milliseconds.	SolarWinds
Database throughput	Rate of database requests for this service per minute.	SolarWinds

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU utilization	Percent of the host processing power being consumed.	SolarWinds
Memory utilization	Percent of memory the host is using.	SolarWinds
Disk utilization	Percent of disk space being used.	SolarWinds
All Active Servers	Information for all hosts the service is actively using. Select a host link to view more detailed information.	CMDB
All Active VM Instances	Information for all virtual machine instances the service is actively using. Select a VM link to view more detailed information.	CMDB

Database dashboards

Database metrics are currently not supported for SolarWinds.

Networking dashboards

These dashboards display metrics for network devices related to the service.

By default, the Firewall dashboard displays Palo Alto metrics and the Load Balancer dashboard displays F5 Big-IP metrics.

Firewall

Chart	Description	Data source
CPU Utilization	Percentage of CPU utilization on the firewall device.	SolarWinds
Session Utilization	Percentage of session capacity currently in use.	SolarWinds
Drop Rate	Percentage of network packets/connections that the device has intentionally blocked or failed to process.	SolarWinds

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Firewall (continued)

Chart	Description	Data source
Availability	Percentage of time the firewall is available.	SolarWinds
Recent Changes	Table displaying recent configuration changes made to the firewall device.	CMDB
Current Open Incidents	Table listing active incidents affecting the firewall device.	CMDB
All Active Alerts	Table listing all active alerts related to the firewall device.	CMDB

Load Balancer

Chart	Description	Data source
Connection Utilization	Percentage of available connections currently in use.	SolarWinds
Inbound Rate	Rate of bytes per second received from the configured nodes.	SolarWinds
Outbound Rate	Rate of bytes per second sent to the configured nodes.	SolarWinds
Pool Member Status	Pool Availability Status: • 0: None (no status available) • 1: Available, all members healthy • 2: Degraded, some members down • 3: Unavailable, all members down • 4: Unknown, status cannot be determined	SolarWinds
Error	Percentage of inbound errors on the load balancer.	SolarWinds
Latency	Aggregate duration of transactions through the load balancer, in milliseconds.	SolarWinds

Interface

Chart	Description	Data source
Error	Percentage of inbound errors on the interface.	SolarWinds
Traffic	Rate in bytes per second of traffic through the interface.	SolarWinds
Bandwidth Utilization	Percentage of available bandwidth.	SolarWinds

Other Networking Devices

Chart	Description	Data source
CPU Utilization	Percentage of CPU utilization on the networking device.	SolarWinds
Packet Loss	Percentage of packets dropped.	SolarWinds
Error	Percentage of errors on the networking device.	SolarWinds

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Splunk Observability templates for Service Observability

Templates used to create Splunk Observability dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

Splunk Observability Overview tab for Service Observability

Information that is displayed on the Splunk **Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
Rate	Aggregate rate of transactions through the service.	Splunk
Error %	Percent of transactions that have an error.	Splunk
Latency	Aggregate duration of transactions through the service, in milliseconds.	Splunk

Key Transactions charts

The Key Transactions charts show metrics for specific key transactions on the service. The key transactions are determined by your observability system configuration.

Chart name	Description	Data source
Rate	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	Splunk
Error %	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	Splunk
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	Splunk

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
All Active alerts	Listing of active alerts that affect this service.	CMDB

Splunk Observability tab for Service Observability

Dashboard and charts on the Splunk Observability tab of the Service Details page in the SOW.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service KPIs

Chart	Description	Data source
Rate	Aggregate rate of transactions through the service, per minute.	Splunk
Error %	Percent of transactions that have an error.	Splunk
Latency	Aggregate duration of transactions through the service, in milliseconds.	Splunk

Service top transactions

Chart	Description	Data source
Rate	Each line represents the throughput of a key transaction through the service. Hover over a point to view the different transaction names.	Splunk

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service top transactions (continued)

Chart	Description	Data source
Error %	Each line represents the percentage of errors on a key transaction in the service. Hover over a point to view the different transaction names.	Splunk
Latency	Each line represents the response time for a key transaction on the service. Hover over a point to view the different transaction names.	Splunk

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU utilization	Percentage of CPU currently being consumed by the service.	Splunk
Memory utilization	Percentage of memory currently being consumed by the service.	Splunk
Disk utilization	Percentage of disk space currently being consumed by the service.	Splunk
All Active Servers	Information for all hosts the service is actively using. Select a host link to view more detailed information.	CMDB
All Active VM Instances	Information for all VM instances the service is actively using. Select a host link to view more detailed information.	CMDB

Database Dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
Queries Per Second	Number of queries per second the service is sending to the database.	Splunk
Row Locks Rate	Frequency at which row locks are encountered by transactions from this service.	Splunk
Buffer Pool Usage	Usage of the buffer pool to access cached data by this service.	Splunk
Table I/O Wait Time	Time in milliseconds that the database is waiting on Input/Output operations on tables sent by this service to complete.	Splunk

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL (continued)

Chart	Description	Data source
Threads	Number of threads used by this service	Splunk
Uptime	Amount of time in (?) that the database has been running since its last restart.	CMDB

Zabbix templates

Templates used to create Zabbix dashboards in Service Observability. You can edit these templates as needed.

See Customize Service Observability dashboard templates for more information.

Zabbix Overview tab for Service Observability

Information that is displayed on the Zabbix **Overview** tab of the Service Details page in the SOW.

KPIs section

The KPIs section displays basic performance metrics for the service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Chart name	Description	Data source
CPU Usage (System)	Displays the system CPU usage percentage over time.	Zabbix
CPU Load	Shows the CPU load average for the system.	Zabbix
Running Processes	Number of processes currently running on the system.	Zabbix
Available Memory (%)	Percentage of memory currently available on the system.	Zabbix

Contextual information section

The Recent changes table shows changes made to the system that might have an impact on the service. This data comes from the related configuration items in the CMDB.

Status indicators section

The Status indicators section displays charts for related CIs that might be affecting the service.

Chart name	Description	Data source
SLO Error Budget	Information about your current error budget for this service, as configured in Service Reliability Management (SRM). For more information on error budgets, see Create SLOs, SLIs, and error budget policies.	CMDB
Systems impacted	Services or service offerings that have a relationship to the selected service.	CMDB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart name	Description	Data source
Open incidents last 30 days	Day-by-day occurrence of open incidents affecting this service over the last month.	CMDB
Open alerts last 30 days	Day-by-day occurrence of alerts affecting this service over the last month.	CMDB
Current open incidents	Listing of active incidents that affect this service.	CMDB
Current open alerts	Listing of active alerts that affect this service.	CMDB

Zabbix Observability template for Service Observability

Dashboard and charts on the Zabbix Observability tab of the Service Details page in the SOW.

Application dashboard

This dashboard displays performance metrics for the selected service.

Note: If the dashboards have been customized, they might not reflect the information shown on this page.

Service

Chart	Description	Data source
CPU Usage (System)	CPU usage percentage over time.	Zabbix
CPU Load	CPU load for the system: rolling one minute average.	Zabbix
Running Processes	Number of processes currently running on the system.	Zabbix
Available Memory (%)	Percentage of memory currently available on the system.	Zabbix

Compute dashboard

This dashboard displays metrics for hosts related to the service.

Host

Chart	Description	Data source
CPU Load	CPU load average for the host on a rolling one minute average.	Zabbix
CPU Idle Time	Percentage of time the CPU is idle.	Zabbix
Memory Available	Percentage of available memory on the host.	Zabbix

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Host (continued)

Chart	Description	Data source
Running Processes	Number of processes currently running on the host.	Zabbix
Network RX	Network data received in bytes per second.	Zabbix
Network TX	Network data transmitted in bytes per second.	Zabbix
All Active Servers	Information for all active server instances with their operating system, CPU count, RAM, disk space, and support group information.	CMDB
All Active VM Instances	Information for all virtual machine instances the service is actively using. Select a VM link to view more detailed information.	CMDB

Database dashboards

These dashboards display metrics for databases related to the service.

MySQL

Chart	Description	Data source
Active Connections	Number of active database connections.	Zabbix
Queries Per Second	Rate of database queries executed per second.	Zabbix
Rows Deleted By Queries Per Second	Number of rows deleted by queries per second.	Zabbix
Rows Returned By Queries Per Second	Number of rows returned by queries per second.	Zabbix
Rows Inserted By Queries Per Second	Number of rows inserted by queries per second.	Zabbix
Rows Updated By Queries Per Second	Number of rows updated by queries per second.	Zabbix
All MySQL Instances	Table listing all MySQL database instances with their operational status, class, and support group information.	CMDB

PostgreSQL

Chart	Description	Data source
Number of Connections	Total number of active database connections.	Zabbix
Number of Waiting Connections	Number of database connections currently waiting.	Zabbix
Connected Backends	Number of connected backend processes.	Zabbix

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PostgreSQL (continued)

Chart	Description	Data source
Rows Returned By Queries Per Second	Number of rows returned by queries per second.	Zabbix
Rows Inserted By Queries Per Second	Number of rows inserted by queries per second.	Zabbix
Rows Deleted By Queries Per Second	Number of rows deleted by queries per second.	Zabbix
All Postgresql Instances	Table listing all PostgreSQL database instances with their operational status, discovery source, category, subcategory, and support group information.	CMDB

Networking dashboards

These dashboards display metrics for network devices related to the service.

By default, the Firewall dashboard displays Palo Alto metrics and the Load Balancer dashboard displays F5 Big-IP metrics.

Firewall

Chart	Description	Data source
CPU Utilization	Percentage of CPU utilization on the firewall device.	Zabbix
Session Utilization	Percentage of session capacity currently in use.	Zabbix
Memory Utilization	Percentage of memory utilization on the firewall device.	Zabbix
Availability	Uptime and availability status of the firewall device: • 0: Unavailable • 1: Available	Zabbix
Recent Changes	Table displaying recent configuration changes made to the firewall device.	CMDB
Current Open Incidents	Table listing active incidents affecting the firewall device.	CMDB
All Active Alerts	Table listing all active alerts related to the firewall device.	CMDB

Load Balancer

Chart	Description	Data source
Current Connections (To Node)	Number of current connections from server-side to the configured nodes.	Zabbix

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Load Balancer (continued)

Chart	Description	Data source
Requests TX	Rate of bytes per second sent to server-side from the configured nodes.	Zabbix
Requests RX	Rate of bytes per second received server-side from the configured nodes.	Zabbix
Pool Availability State	Pool Availability Status: • 0: None (no status available) • 1: Available, all members healthy • 2: Degraded, some members down • 3: Unavailable, all members down • 4: Unknown, status cannot be determined	Zabbix
Packets TX	Number of packets transmitted per second from the server-side to the configured nodes.	Zabbix
Packets RX	Number of packets received per second by the server-side from the configured nodes.	Zabbix
Inbound Error %	Percentage of inbound errors on the load balancer.	Zabbix
Outbound Error %	Percentage of outbound errors on the load balancer.	Zabbix

Other Networking Devices

Chart	Description	Data source
CPU Utilization	Percentage of CPU utilization on the networking device.	Zabbix
Availability	Uptime and availability status of the networking device: • 0: Unavailable • 1: Available	Zabbix
Inbound Error %	Percentage of inbound errors on the networking device.	Zabbix
Outbound Error %	Percentage of outbound errors on the networking device.	Zabbix
Inbound Traffic	Amount of inbound network traffic in bytes per second.	Zabbix
Outbound Traffic	Amount of outbound network traffic in bytes per second.	Zabbix

Service Observability data mapping form

Field descriptions for the Observability data mapping form. Use this form to map activated services in Service Observability to metrics from your observability instance. Each service can be mapped only once, but can contain exceptions to include all related entities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Mapping name	Name for the mapping. The name should be descriptive so that users understand where the data is coming from.
Services	Service to use for this mapping. Each service can be mapped only once.
Default	
Find metrics from this source	Existing data connection to use for metrics. To add a connection, see Connect a Service Observability data source.
where source tag key	Tag key that you want to map to the activated service. If you use more than one tag to represent a service, you can create an exception. For example, most tags use `service_name` except for database metrics, which use `service`.
source tag value	Value that represents the service name. Values can be strings or a variable that represents fields on the CI for the service. You can use variables on your observability data as the value to map multiple entities at once. For example, if your observability data uses the `$service_name` variable and you enter that as the **source tag value**, all services using that value are mapped.
Exceptions	
For entity type	Entity that uses a different key or value than the default.
Find metrics from this source	Existing data connection to use for metrics. To add a connection, see Connect a Service Observability data source.
where source tag key	Tag key that you want to map to the activated service for this entity.
source tag value	Value that represents the service name. Values can be strings or a variable that represents fields on the CI for the service.

Chart error states

Understand the different error states that the charts in Service Observability might display and how to fix them.

Error message	Fix
Set up this service's observability data mapping	Observability mapping hasn't been configured for the service by an admin. See Create and manage data mappings.
Data mapping is not configured	Observability mapping hasn't been configured for the service by an admin. See Create and manage data mappings.
Authentication failed	Service Observability can't connect to the data source. See Connect a Service Observability data source.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Error message	Fix
No results found	No metrics were found for the chart.
Something went wrong	Connection to the data source was successful, but the observability instance didn't respond. This error might be due to a connection timeout on the observability vendor side.
No impacted systems found (Impacted systems chart on the Overview page)	Only services found on your service map can be displayed. Update your service map to include all services that should be included in the chart. See Service Mapping.
Set up objectives (SLO Error Budget chart on the Overview page)	SLOs haven't been configured for Service Reliability Management. See Create SLOs, SLIs, and error budget policies.

Service Reliability Management

Service Reliability Management (SRM) helps your organization respond, collaborate, track, and self-remediate when working on alerts and incidents. SRM is available when you have both ITSM Standard and ITOM Operator Professional subscriptions.

Overview of Service Reliability Management

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore Learn about SRM concepts and features.	Configure Install and configure SRM.
Use Work with services, teams, and reliability tasks in SRM.	Reference Get details about SRM components, such as fields, tables, and properties.

Install from the Store

Install SRM from the ServiceNow Store. Visit the ServiceNow Store to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.

Note: When installing from the store, first confirm that you have all the dependencies needed by selecting the Dependencies tab. For more information, see Install Service Reliability Management.

Troubleshoot and get help

- Visit our community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Service Reliability Management

Service Reliability Management (SRM) provides a self-serve, guided experience for teams to manage service health. The experience is built using the Service Operations Workspace application and combines ITOM and ITSM capabilities into a single workflow.

SRM overview

Optimize service health with site reliability engineering (SRE) practices. SRM is a single operations workspace that empowers teams to improve the reliability of digital services with SRE.

- Use on-call escalations to respond to issues in a timely manner.
- Reduce setup friction with guided self-service to onboard distributed teams with separated data, empowered access, and minimal governance from central IT.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When SRM is installed, several plugins and applications are also activated. For more information, see Plugins or applications installed with ITOM AIOps.

SRM users

Users

Users	Description	Contains Roles
admin	ServiceNow administrators manage, configure, and maintain the ServiceNow platform. In SRM, they can access and work in the Service Operations Workspace Admin Center. Only administrators can do the following: • Install SRM. • Add and manage SRM administrators. • Create and manage integration users.	All
SRM administrator [srm_admin] **Note:** This role differs from the ServiceNow admin role.	SRM administrators can manage account settings, configurations, and users. Administrators can perform the following actions: • Access, create, edit, or delete all SRM configurations. • Add or manage integrations. • Create integrations with Application Performance Monitoring (APM) tools. • Set up and maintain reliability metrics. • Set up and maintain error budget policies.	• Manager • Responder
SRM manager [srm_manager]	Managers oversee a team of SREs. Managers assign SREs to the team on-call schedule, monitor their performance, and create procedures to handle incidents and develop solutions. Managers promote resilience across all the systems and the DevOps workflows. Managers can perform the following actions within the context of their teams:	Responder

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Users (continued)

Users	Description	Contains Roles
	• Define and set up teams, on-call schedules, and services. • Add and delete users such as responders and managers for the teams they're a part of. • Add or manage integrations. • Create Integrations with Application Performance Monitoring (APM) tools. • Set up and maintain reliability metrics. • Set up and maintain error budget policies.	
SRM responder [srm_responder]	A Service Reliability Engineer (SRE) that uses SRM to perform everyday tasks. Responders are the individuals who are on call and diagnose and remediate incidents. Responders can only access configurations that they're a part of. They can only access the alerts or incidents for which they have permissions. SREs can perform the following actions, within the context of their teams: • Set up services, teams, and integrations. • Confirm their on-call schedules. • Manage incident and alert records. • Update teams that they've created. • Add other responders. • Create integrations with Application Performance Monitoring (APM) tools. • Set up and maintain reliability metrics. • Set up and maintain error budget actions.	Inherits 17 roles including the following: • cmdb_read • sn_sow.sow_user • sn_sow_srm.srm_responder • workspace_user • slo_operator

For more information, see SRM roles and responsibilities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SRM workflow

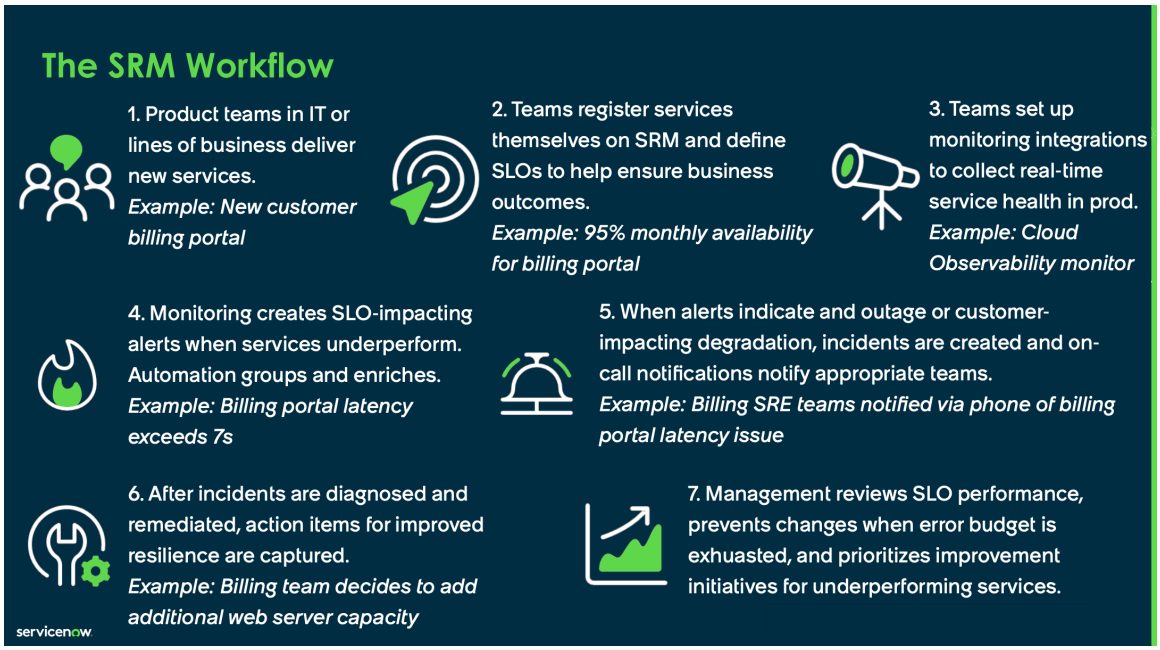

1. Product teams in IT or Lines of Business continuously deliver new service instances and technology management services. Example: New customer billing portal.
2. Along with SLO Management, teams can register services and define service level objectives (SLOs), helping them reach business outcomes. Example: 95% monthly availability for the billing portal.
3. Monitoring integrations are set up by the teams to collect the real-time health of these services. Example: Cloud Observability.
4. Monitoring creates service level indicators (SLIs) impacting alerts when services are underperforming. Automation groups and enriches. Example: Billing portal latency is exceeding 7 s.
5. When the alerts indicate an outage or customer-impacting degradation, incidents are created and on-call notifications notify appropriate team resources. Example: A Billing SRE team is notified via phone of a latency issue on the billing portal.
6. After teams collaboratively diagnose and remediate incidents, they identify action items for improving the system's resilience. Example: The Billing team decides to add additional web server capacity.
7. Management continually reviews SLO performance, helps to prevent changes when the error budget is exhausted, and prioritizes improvement initiatives for underperforming services.

SRM benefits

Benefit	Feature	Users
Team-based experience	Working with SRM teams	SRM administrators, managers, and responders
Service registration	Working with SRM services	SRM administrators, managers, and responders

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefit	Feature	Users
Prebuilt integrations	Working with integrations in SRM	SRM administrators, managers, and responders
Measure service health	Reliability metrics in SLO Management	SRM administrators, managers, and responders
On-call coverage	Create an SRM on-call schedule	SRM administrators, managers, and responders
Remediate high severity alerts and incidents	Working with SRM reliability tasks	SRM administrators, managers, and responders

What to explore next

To learn more about configuring and using SRM, see:

- Configuring Service Reliability Management
- Using Service Reliability Management
- Service Reliability Management reference

Get started with Service Reliability Management

Service Reliability Management (SRM) accelerates your path to viewing service health in the context of service level objectives and incident resolution. Helps IT Operations and DevOps teams deliver on the promise of agility, performance, and uptime.

Get started with SRM to understand the different sections in the SRM interface.

For more information on roles, see SRM roles and responsibilities.

Basic SRM tasks

Step	Description	Reference
Set up teams and services	Setup guides for the Home, Services, and Teams pages show how to add a team or service to SRM.	Add an SRM team
Visit and learn about your Home page	The Home page is where you find the things most important to you. For example, the services with critical alerts and incidents and work assigned to you and your team.	SRM Home page
Learn how to navigate SRM	Get familiar with the different sections and elements of the SRM interface. The sections and elements are used throughout the documentation.	SRM interface

SRM helps you when you must create and administer teams, services, and integrations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SRM tasks

Step	Description	See this
Manage a service instance or technology management service in SRM	Define the basic tasks and parameters that make up your service and how it should behave.	Add a service to SRM
Set up an SRM team	Set up a team. Teams are responsible for the issues that occur in their associated services.	Add an SRM team
Set up on-call schedules and escalation policies	Create an on-call schedule for your team to make sure they're available to resolve issues. You can set up an escalation policy to make sure at least one team member is engaged in incident response.	Create an SRM on-call schedule
Configure service relationships	Use a map canvas to add, configure, and arrange services. You can add child services that depend on parent services.	View impact of child service on parent service
Integrate services with third-party monitoring tools	Set up a third-party integration, such as Datadog or ServiceNow Cloud Observability, with SRM so that alerts are available to your teams within SRM.	Working with integrations in SRM
Establish SLOs, SLIs & error budgets for services	Establish goals for how well your service should operate. Also, specify the maximum amount of time that a technical system can fail without contractual consequences.	Service Level Objective Management
Set up alert automations **Note:** This functionality is only available if you have installed the Alert automations application.	Alert automations enable you to define alert conditions. Set up alert rules for each APM tool to define the conditions when the APM tool should send notifications to SRM.	The Alert automation application is available from the ServiceNow Store.

SRM Home page

The Service Reliability Management (SRM) Home page provides relevant information about your work, services, schedules, maintenance reminders, and teams. View the Home page by navigating to **Service Reliability Management**.

Service Reliability Management Home page sections

In the workspace, select **Home** ( ) to view the Home page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

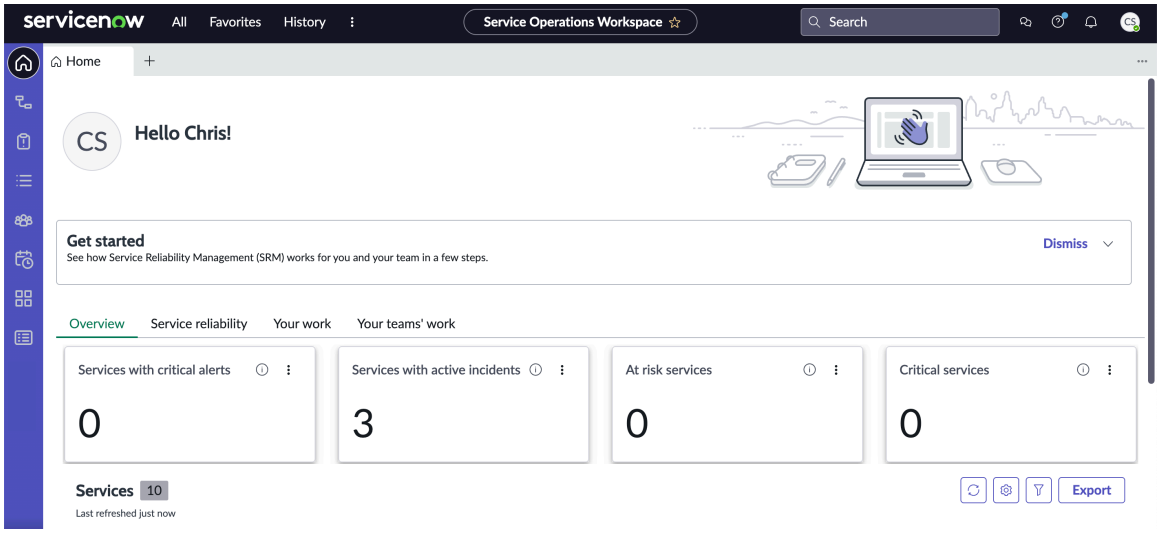

When you're on-call, your schedule is listed.

Home page sections

Get started	Contains setup tasks and their progress for teams and services. For more information on completing setup, see Add a service to SRM and Edit service details in SRM. **Note:** You can skip any section and come back to it later. Progress is marked in percentages by the actions that you take. As long as a step remains incomplete, the tiles are available. You can minimize them, at any point, by selecting the Up icon . Use **Dismiss** to remove it from your page. The checklist states are as follows: • **None**: No action has been taken. • **Done**: All required fields have been added. • **Pending**: Some fields or approvals still must be added or granted.
Overview	Contains an overview of all service alerts, incidents, and at-risk and critical services. **Services with critical alerts** are the default list view. Select a card to see the list view for that area.
Service reliability	Displays a dashboard for monitoring and managing service performance. The dashboard offers multiple visualizations to help you track error budgets, monitor service level objectives, and identify issues across services.
Your work	Contains incidents and alerts assigned to you. Select a card to see the list view for that area.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Home page sections (continued)

Your teams' work	Contains incidents and alerts assigned to your team. Select a card to see the list view for that area.

SRM interface

The different components on the Service Reliability Management (SRM) interface provide easy access to different functions. The features and functions that appear depend on your role and permissions.

The SRM interface has the following different sections:

Interface sections

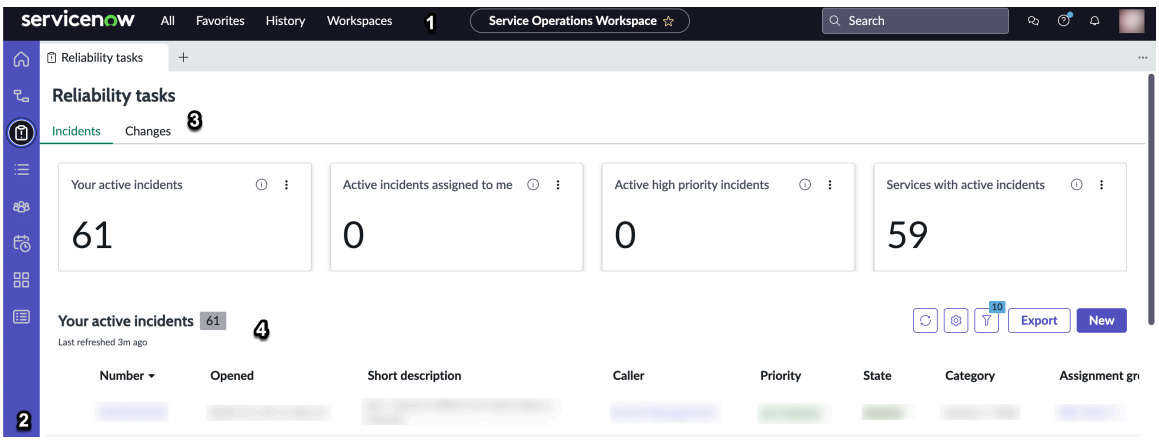

Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics. Depending on your role, you may see other icons in the primary navigation.

If you see the SOW Home page, select an SRM icon to view SRM-specific information.

Sections	Description
1 - Header	A home for your organization logo and system-level utilities present everywhere in the interface. • All • Favorites • History • Workspaces • Current workspace • Search • Sidebar discussions • Help

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sections	Description
	• Notifications • Profile icon: Represents your profile
2 - Primary navigation	Easy access to specific pages in SRM. The primary navigation uses icons, including icons for Services, Reliability tasks, Teams, and Schedules. Use tooltips to identify them.
3 - Content headers	Tabs, cards, or buttons that provide the means to accomplish product-level tasks. The elements in the content header differ from one page to another.
4 - Content area	Canvas or the main area where you can gather information, read metrics, and perform tasks. **Note:** If you want to minimize the side panel on forms, by default, navigate to **Preferences** > **Workspace** and disable the **Show sidebar** option. That action disables it across the Service Operations Workspace, not just SRM.

Search, filters, lists, forms, and guides

An overview of the additional elements in the Service Reliability Management (SRM) interface that help you find items, view available forms and lists, and follow setup processes.

Table List view

View services and incidents as a list.

- Search using the header text box.
- Order items using the list columns.
- Filter list results using the list columns.

Note: View the active filters using the filter icon in the list view header.

Each view contains its own unique items. See those pages for details.

Form view

When you open any of the table records from the list view, you see the details of the record as a form.

Each view contains its own unique items. See those pages for details.

Note: Grayed out fields can't be changed.

The Attachments panel is often visible by default, to hide it select the Attachment icon .

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Setup checklists

Use the setup checklists to visualize the process that you go through to reach your goals. Each step is highlighted as you complete it.

Export list information to a file

You can export live information to a file in Service Reliability Management (SRM).

Before you begin

Note: Bulk selection doesn't impact exports. The export includes all records that match the current filter criteria. It doesn't consider the selections you make in the list view.

Role required: srm_admin, srm_manager, or srm_responder

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. Select the list view in any landing page section in SRM with an **Export button**.

3. Determine which filters you want to use and select them from the Filters or More actions menu in the list columns.

   Note:

   The Export button only becomes available if there are records in the list view.

   The export contains all list data that matches your filtered/grouped lists.

4. Select **Export**.
   In the pop-up window select:
   - File Type:
     - Excel
     - CSV
     - JSON
     - PDF
   - Delivery type:
     - Download
     - Email: Provide an email address

   For **Email**: Once the export is complete, a success message appears with your email address at the top of the screen. It's useful for larger files because, depending on the size of the file, exporting might take some time. Email runs in the background, so you're free to do something else while waiting.

   For **Download**: The pop-up begins exporting and the file appears in your Download directory. Depending on the size of the file, it might take some time. You can open it in any application that supports the file format.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Manually create SRM tags

Add tags to classify, categorize, and add context to data. Service Reliability Management (SRM) imports tags and attaches them to services, alerts, and incidents. You can also manually create tags.

Before you begin

Role required: srm_admin, srm_manager, or srm_responder

About this task

There are two ways that tags can enter the system and attach to an alert or incident. The first is automated but the last one is manual.

- Imported from a third-party integration such as Datadog.
- Using the tag icon in the header of a service, alert, or incident.
- Added using the **Modify field values** action in an automation rule.

This task describes how to add tags to services and incidents in SRM. For information about managing alerts, see Alerts in Service Operations Workspace.

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Services** or **Reliability tasks**.
3. Open a service or incident.
4. In the header, select the **Tag** icon to add tags.
5. Enter your tag in the pop-up window.
   Your tag is saved and appears below the Add Tag box alongside any existing tags.
6. Select anywhere in the form to close the pop-up.

Note:

Tags are also shown under the description in the incident list view.

SRM incidents

Track and collaborate on incidents in the Incidents tab, helping you and your teams resolve issues efficiently.

The Incidents tab appears in the Service Reliability Management (SRM) Reliability tasks page. Its top header shows description information, tags, priority information, state, service, and assignment group.

Note: Tags are imported from third-party integrations and attached to incidents. However, you can also create tags to categorize data and drive system logic using the Tag icon . For more information on creating and viewing tags, see Manually create SRM tags.

When selecting an incident, you're taken to an overview. This section contains three content panels: Summary, Activity, and Recommended actions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: When an editable field is updated, notifications are sent based on user notification preferences.

In the Impact Summary, you find the following information:

Affected CIs

Field	Description
Configuration item	Name of and link to the configuration item associated with this incident.

Impacted Services

Field	Description
Configuration item	Name of and link to the configuration item associated with this incident.
Managed by	Team responsible for the incident.
Owned by	Incident owner.
Approval group	Team responsible for approvals.
Location	Geographical location of the incident if defined.
Operational status	Status of the incident.
Manually added.	Whether this incident was manually added or the result of an error budget policy.

Details

Contains incident details and related alerts

Field	Description
Short description	Brief description of the incident.
Description	Detailed description of the incident.
Number	Incident identifier.
Impact	Impact is a measure of the effect of an incident on business processes. Choices are: • 1 - High • 2 - Medium • 3 - Low
Assignment group	Team assigned to the incident.
Priority	Priority is based on **Impact** and **Urgency**, and it identifies how quickly the incident should be addressed. Choices are: • High Impact = Critical Priority • Medium Impact = High Priority • Low Impact = Moderate

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Details

Contains incident details and related alerts

(continued)

Field	Description
	Note: See Priority Data Lookup rules below.
Category	Category of the incident. Choices are: • Inquiry/Help • Software • Hardware • Network • Database
Urgency	Urgency is a measure of how long the resolution can be delayed until an incident has a significant business impact.
Service	Affected service associated with the incident. You can edit this field from a list menu. **Note:** When you set a service, and that service has an assigned team, the assigned team of that service is automatically assigned to the incident.
Severity	Imported severity on the incident. Values P1–P5.
Configuration item	Asset associated with the incident.
Assigned to	Individual responsible for resolving the incident.
State	State of the incident. The state moves and tracks incidents through several stages of resolution. Choices are: • New • In Progress • Ob Hold • Resolved • Canceled
Watch list	Individuals interested in the incident. Subscribes them to notifications according to their preferences.
Resolution code	Code associated with the resolution. Choices are: • None (default) • Duplicate • Known error • No resolution provided

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Details

Contains incident details and related alerts

(continued)

Field	Description
	• Resolved by caller • Resolved by change • Resolved by problem • Resolved by request • Solution provided
Resolution notes	Detailed information on the resolution. Resolution notes are posted to the **Activity** timeline.

Note: Priority is calculated according to the following sample data lookup rules:

Priority Data lookup rules

Impact	Urgency	Priority
1 - High	1 - High	1 - Critical
1 - High	2 - Medium	2 - High
1 - High	3 - Low	3 - Moderate
2 - Medium	1 - High	2 - High
2 - Medium	2 - Medium	3 - Moderate
2 - Medium	3 - Low	4 - Low
3 - Low	1 - High	3 - Moderate
3 - Low	2 - Medium	4 - Low
3 - Low	3 - Low	5 - Planning

When a service is deleted, its integrations, alerts, incidents, and automations are removed. Deleting a service isn't a recoverable action so consider deactivating the service instead.

The Related Records tab contains any alerts associated with this incident.

Note: All field values are imported from the related records and aren't editable from SRM.

Child Incidents

Field	Description
Number	Identifier for the related incident.
Short description	Brief description of the related incident.
Priority	Priority of the related incident.
State	State of the related incident.
Assigned to	Team member assigned to the related incident.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Child Incidents (continued)

Field	Description
Assignment group	Team assigned to the related incident.
Service	Service associated with the related incident.

Alerts

Field	Description
Number	Identifier for the related alert.
Severity	Severity of the related alert.
Short description	Brief description of the related alert.
Acknowledged	Whether the related alert has been acknowledged or not.
Assigned to	Team member assigned to the related alert.
Configuration item	Asset associated with the related alert.
Parent	Primary alert associated with the related alert.

Recent Impactful Changes

Field	Description
Number	Identifier for the related changes associated with this incident.
Short description	Brief description of the related changes associated with this incident.
Type	Type of the recent changes associated with this incident.
State	State of the recent changes associated with this incident.
Planned start date	Date and time to begin changes associated with this incident.
Planned end date	Date and time to begin changes associated with this incident.
Assigned to	Team member responsible for the change request associated with this incident.

Change Requests

Field	Description
Short description	Brief description of the change request associated with this incident.
Type	Type of the change request associated with this incident.
State	State of the change request associated with this incident.
Planned start date	Date and time to begin change request work associated with this incident.
Planned end date	Date and time change request work associated with this incident ended.
Assigned to	Team member responsible for the change request associated with this incident.

Postmortem tab

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About Postmortem owner:

- If a Scribe has already been chosen in the Collaboration panel Responders list, they're automatically made the owner.
- If there's more than one scribe, the first one assigned is chosen. Otherwise, you can select an owner from the list menu.
- If a new scribe is added or an existing user persona is changed to Scribe, that user is made the owner.
- If you manually remove an owner, no new owner is automatically selected.

Use the drag-and-drop feature to edit postmortem sections. You can exclude sections from PDFs, delete sections, or reorder sections.

To add new sections, select the **Add** icon .
Compose panel

Work notes are added to the Activity timeline automatically. Visible to responders and above. Add comments using the **Compose** text box.

Comments are added to the Activity timeline automatically. Visible to everyone in SRM.

Activity section

Stream that contains all the system activity, and resolution notes. You can filter, sort, expand or collapse all posts using the header icons.

Note: The filter option includes **Post types**, **Field changes**, and **Filter sets**. All choices are selected by default. Flagged entries are either all or none.

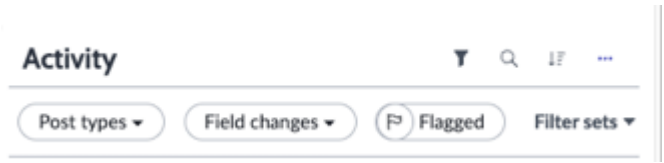

Actions panel

Note:

The **Notifications**, **Agent assist**, **Expert on-call**, and **Email templates** panels aren't currently supported in SRM.

Attachments section

Open the **Attachments** section using the Attachment icon .

On the **Attachments** section, you can initially **Browse** your local hard drive for files.

Once you have added attachments, you can:

- Search for an attachment based on its name or extension

  Note: You must have two or more attachments to search.

- Upload a new attachment using the ( ) icon.
- Download, remove, or rename existing attachments using the More actions icon ( ).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Template section

Open the **Template** panel using the Template icon .

Note: Existing templates are displayed, but creating templates isn't available in SRM.

Collaboration panel

Displays the available Zoom meetings, Microsoft Teams, or Slack channels to start or join.

To start or join a Zoom meeting:

Note: The Zoom integration and recognized IDs must be active in your instance. If Start Channel is unavailable, contact your administrator.

1. Select **Start Zoom**.
   - A meeting pop-up appears.
   - Add participants.
   - Select **Create Meeting**.
   - The meeting is created, invitations are sent, and you join the meeting.

     Note: Any issues found when starting a meeting are shown in a banner message.
2. Select **Join Meeting**
   - **Sign in** to your Zoom account.
   - You join the Zoom meeting.

1. Select **Start Channel**.
   - Sign in to your Microsoft Teams account.
   - Select **Create**. You are redirected to the Microsoft Teams application.
   - Back in the Collaboration column, the **Join Channel** button is activated.

     Note: Any issues found when joining a channel are shown in a banner message.
2. Select **Join Channel** in the Collaboration panel in SRM.
3. Sign into your Microsoft Teams account.

   You are taken to the channel in Microsoft Teams.

Start or join a Slack channel:

Note: The Slack integration and recognized IDs must be active in your instance. If Start Channel is unavailable, contact your administrator.

1. Select **Start Slack**.
   - Add participants in the Create a Slack channel pop-up.
   - Select **Create**. You are redirected to the Slack application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Back in the Collaboration column, the **Join Slack** button is activated.

  Note: Any issues found when starting a channel are shown in a banner message.

2. Select **Join Channel** in the Slack application window.
3. Sign into you Slack account.

   You are taken to the Slack channel.

Service Operations Workspace setup for SRM

Install and configure Service Reliability Management (SRM) within the Service Operations Workspace Admin Center so that SRM admins, managers and responders can provide service reliability efficiently.

You can Install Service Reliability Management from the Service Operations Workspace Admin Center.

Note: The Data visualizations app v24.1.8 is also required. See Data visualizations in the ServiceNow Store to install the app.

Once installed, configuring SRM within the Service Operations Workspace Admin Center sets up the SRM workspace with options and actions.

Using the Service Reliability Management menus in the Service Operations Workspace Admin Center, you can:

- Assign an administrator to Service Reliability Management.
- Activate teams and services in SRM
- Customize service approval settings.
- Customize team approval settings
- Access and publish the request catalog item for users.
- Configure error budget actions for breached SRM SLOs.

Alert automations are conditional triggers that execute response actions automatically, based on the contents of the alert. This functionality is available to SRM upon request, from the Alert automations application in the ServiceNow Store.

Configuring Service Reliability Management

Configuring Service Reliability Management (SRM) involves installing it from the ServiceNow Store or Service Operations Workspace Admin Center. You can also assign administrators and import services and teams.

Using guided setup for SRM

The guided setup provides a sequence of tasks to help you configure SRM on your instance. To open the guided setup, navigate to **All** > *Service Operations Workspace* **Admin Center** > **Overview**. Next, find the Service Reliability Management tile and select **Configure**.

SRM configuration documentation

Visit the following pages to set up and configure SRM. The content is primarily for srm_admin or admin users.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Install Service Reliability Management

Install the Service Reliability Management (SRM) application from the Admin Center or ServiceNow Store.

Before you begin

Subscriptions required: ITSM Standard and ITOM Operator Professional subscriptions.

Note: The data visualizations store app v24.1.8 is also required. See Data visualizations in the ServiceNow Store to install the app.

To create and manage alert automations, install the Alert automation application available on request in the ServiceNow Store.

Role required: admin

About this task

Roles and demo data are installed with SRM. The demo data includes templates that help accelerate implementation and SRM use.

Procedure

1. Navigate to **All** > *Service Operations Workspace* **Admin Center** > **Overview**.
2. Locate the *Service Reliability Management* card and select **Install**.
3. In the pop-up window, select **Start Install**.
   You're taken to the ServiceNow Store portal.
4. Complete the necessary steps.

   Note: If any dependencies are red on the app store listing, select the row to open the dependent app and install it. If the dependent app is already installed but still appears in red, you can get or reopt it as needed.

   Demo data comprises the sample records that describe application features for the common use cases. Demo data is loaded when you first install the application on a development or test instance.

   Following installation, the *Service Reliability Management* card can be used to configure SRM.

Assign an administrator to Service Reliability Management

Add people from your ServiceNow instance to administer Service Reliability Management (SRM) and your service operations.

Before you begin

Role required: admin

About this task

The SRM admin can add and create teams. They can also manage data across all SRM teams.

Procedure

1. Navigate to **All** > *Service Operations Workspace* **Admin Center** > **Overview**.
2. In the *Service Reliability Management* tile, select **Configure**.
3. Navigate to **Assign admins** > **Service reliability management admins**.
4. Enter user names in the field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Users must have an email address associated with their user record to appear in the search results.

5. Select **Save**.

Activate teams and services in SRM

Activate teams and services to be managed in the Service Reliability Management (SRM) application.

Before you begin

Role required: srm_admin or admin

Procedure

1. Navigate to **All** > *Service Operations Workspace* **Admin Center** > **Configurations**.
2. Select *Service Reliability Management* > **Activate** > **Teams and services**.
3. Select **Activate teams and services** to begin importing or adding teams.
4. Use the condition builder to filter the teams shown in the Assignment groups list.

   Note:

   Filters on the same field use AND logic, and filters on different fields use OR logic. For example, filtering for multiple `Business Criticality` levels and `State = Open`, evaluates as `(Criticality = 1 OR 2 OR 3 OR 4) AND (State = Open)`.

   List view filters are saved automatically and persist when you navigate away and return.

5. In the Assignment groups list, select the teams you want to add and then select **Activate**.
   You reach the **Activate services** option to add services.
6. From the **Select a service class** list, select a service.
   The following service types are available:
   - Application Services
   - Mapped Application Services
   - Calculated Application Services
   - Dynamic CI Group
   - Technical Services
   - Tag-based Services
7. Use the condition builder to add new services to the list view.
8. Select the services that you want to add and then select **Activate**.
   The **Activate Services** pop-up opens.
9. From the **Support Group** drop-down, select the relevant support group and select **Confirm**.

Result

On the **Activate teams and services** page, the following options appear:

- Existing teams
- Existing services

You can select the options to view the teams or services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Customize service approval settings

Add an approval process to your services for governance.

Before you begin
Role required: srm_admin or admin

About this task
Require approval when Service Reliability Management (SRM) admins, managers, or responders request to create or add a service in SRM.

Procedure

1. Navigate to **All** > *Service Operations Workspace* **Admin Center** > **Configurations**.
2. Select *Service Reliability Management* > **Governance and autonomy** > **Service governance**.
3. Choose one or both processes by selecting the relevant check box.

   Approval requirement types:

   - Approval required when creating a new service:

     Requires approval when creating a service in SRM, and assigns it to the team responsible for new services.

   - Approval required when associating an existing Configuration Management Database (CMDB) service:

     Requires approval to add an existing service to SRM, and assigns it to the team responsible for existing services.

4. Enter or select the team responsible for approvals.
5. Select **Save**.
   This action initiates the approval process and assigns approval to the team responsible for approving the new service.

   Note: Approvals take two steps. First, the request is approved, which appears in the **Stage** field of the request record. Next, the request remains **Pending** until the request item is approved in the **Request** record.

Customize team approval settings

Customize how team approvals work when non-admin users create or add a team in Service Reliability Management (SRM). This feature, also known as team governance, helps you control access and align with internal policies.

Before you begin
Role required: srm_admin or admin

About this task
By default, team governance is active in SRM. When non-admin users request a team, SRM creates two records: a parent request and a requested item. The approval team is notified, and any member can approve the request from My approvals. Once approved, the team is created or added, the parent request is set to `Approved`, and the requested item is set to `Closed Complete`.

You can customize team governance by turning off approvals or changing the approval team.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Non-admin users are SRM managers and SRM responders. When team governance is active, admins and SRM admins can still create and add teams without approval.

Procedure

1. Navigate to **All** > *Service Operations Workspace* **Admin Center** > **Configurations**.
2. Under **Service Reliability Management**, select **Governance and autonomy** and then select **Team governance**.
3. Review or update the approval options.

Option	Description
Approval required when creating a new team	Clear this option to let non-admin users cre ate teams in SRM without approval.
Approval required when associating an ex isting team	Clear this option to let non-admin users add existing teams to SRM without approval.

4. Review or update the approval team for one or both options.
   By default, the approval team is set to `SRM Team approvers`. You can view its members in `sys_user_group.list`.

   Note: Any member of the approval team can approve requests, but they need the `sn_request_read` role to view request details. Without that role, they only see the approval task.

5. Select **Save** to update your team governance settings.

Configure error budget actions for breached SRM SLOs

Configure the actions that the team can select when the error budget of a service level objective (SLO) is breached.

Before you begin

Role required: srm_admin or admin

Procedure

1. Navigate to **All** > *Service Operations Workspace* **Admin Center** > **Configurations**.
2. Select *Service Reliability Management* > **Governance and autonomy** > **Error budgets**.
3. Select **Configure error budget actions** to edit the Process Error Budget Policy Threshold subflow in Workflow Studio.

   Note: Working in Workflow Studio requires advanced ServiceNow expertise. For more detailed information on editing workflows or creating custom workflows, see Flow Designer.

Approve a change request for SRM

Approve a service change request in Service Reliability Management (SRM) to allow the requested action to proceed.

Before you begin

Only open change requests are available on the **Reliability tasks** page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: SRM approval team or admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Select the **Changes** tab.
4. Select a change request and then select the **Details** tab.
5. On the **State** field, select the list and select **Approved**.
6. Select **Save**.
   The change request state updates from `Assess` to `Approved`.

   On approval, the new service is created and added to the SRM team. The owner is set to the person shown in the Assigned to field. And the Managed by group is set to the SRM team associated with the assignee.

Configure fields for escalation triggers

Configure additional fields for on-call escalation triggers in Service Reliability Management (SRM).

Before you begin
Role required: admin

About this task
Configure additional fields that teams can use when setting up escalation triggers. By default, escalation triggers run only when the Assignment group field changes. Completing the following procedure makes other fields available, giving teams more flexibility and helping them respond to issues quickly.

Procedure

1. Navigate to **All** > **On-Call Scheduling** > **Administration** > **Trigger Rule Table Config**.
2. Select **em_alert** or **incident** to configure additional trigger fields for alerts or incidents.
3. In **Trigger condition fields**, add fields from the following table.

   Note: Selecting fields not listed in the table might cause performance issues.

Table	Supported fields
em_alert	○ Assigned to ○ Flap count ○ Flap last state ○ Initial event generation time ○ Last event generation time

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Supported fields
	○ Number of Work Notes ○ Overall Event Count
incident	○ Assigned to ○ Category ○ Escalation ○ Impact ○ Reopen count ○ State ○ Urgency

4. Select **Update** to save your changes.
   The selected fields now appear as options for teams configuring escalation triggers.

What to do next
To set up an escalation trigger, see Create an escalation trigger.

Configure Twilio to send notifications in SRM

Connect your Service Reliability Management (SRM) instance and Twilio account to send notifications to teams.

Before you begin
Before you begin, you must have a String Identifier (SID) and an authentication token for an active Twilio account.

Note: Performing this action requires ServiceNow expertise.

Role required: srm_admin or admin

About this task
You can adjust what gets sent from Twilio by editing the on-call workflow in the system.

Procedure
1. Navigate to **All** > *Service Operations Workspace* **Admin Center** > **Configurations**.
2. Select *Service Reliability Management* > **Integrations** > **Twilio API Key**.
3. Select **Twilio Config**.
4. In the new window, select **here** to edit the record.
5. Enter the **Account SID** value.
6. **Optional:** Enter a **Validity Period** value.
7. Select **Connect** to activate the key or **Save** to save your changes without activating.

Install the Alert Automation application

Install the **Alert Automation** application and any supporting applications.

Before you begin
Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Subscriptions required: ITSM Professional or ITOM Operator Professional subscriptions.

About this task

Before creating management rules, you must install and configure the Service Reliability Management (SRM) application, available in the ServiceNow Store.

Procedure

1. To access the ServiceNow Store and download the application onto your SRM instance, select the link in the email you received.
2. In your ServiceNow AI Platform instance, navigate to **All** > **System Applications** > **All Available Applications** > **All**.
3. Locate the *Service Operations Workspace* **Alert Automation** application and select **Install**.
   Alert automation rules are ready to use in SRM.

Using Service Reliability Management

Service Reliability Management (SRM) enables you to register services, monitor service health, and respond to service degradations with on-call shifts and escalation policies.

Basic Tasks in SRM

Working with SRM services

A service represents a functional outcome like networking, payments, or HR services, that is owned by a team. To deliver that outcome, a service can contain one or more technical components like a user authentication service, or a piece of shared infrastructure like a database.

Service Reliability Management (SRM) works with integrations to prioritize and route alerts to the relevant responders. It follows up with escalations until the alert is acknowledged and someone responds. When you create or add a service in SRM, it must reflect a service in your SRM infrastructure.

Note: You might want multiple tool integrations to monitor each technology management service and receive events from those tools. See Working with integrations in SRM for more information.

In addition, you can create reliability metrics for the service. See Reliability metrics in SLO Management.

Tying a team and policies to that service makes it easier to divide responsibilities and track technical outcomes. It also makes it easier to automate response routines and focus on who you notify and when.

The state of an exiting service is inherited. The state of a created service in SRM is **None**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Services Overview

Information on the Overview tab

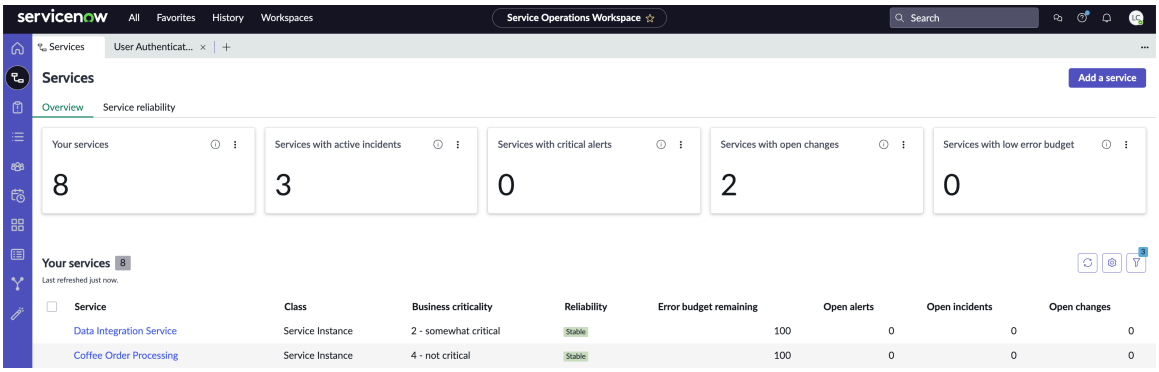

The cards on the Overview tab display the following metrics. By default, the list view shows information related to the Your services card. Select a different card to view different information in the list view.

- **Your services**: Count of all the services you or your team manages and monitors for reliability.
- **Services with active incidents**: Services with open incidents sorted in the following order:
  - Business criticality - most critical first.
  - Number of active incidents - highest first.
  - Percentage of error budget remaining - lowest first.
- **Services with critical alerts**: Services with open alerts sorted in the following order:
  - Business criticality - most critical first.
  - Number of alerts - highest first.
  - Percentage of error budget remaining - lowest first.
- **Services with open changes**: All the services your team manages and monitors.
- **Services with low error budget**: Services with less than 25% error budget remaining.

  The error budget metric is represented as the amount of service level objective (SLO) that you can spend over a specified time. It can be used to manage release velocity.

Note: To refresh the card and list values, select Refresh .

You can interact with the list in the following ways:

- Group or filter columns to customize the view.
- Edit, sort, or export the list as needed. See Export list information to a file.

For more information about individual service details, see Edit service details form.

Services list view definitions

The columns include the following details:

- **Service**: Name of the service.
- **Class**: Service instance or technology management service.
- **Business criticality**: Importance of the service to the business.
- **Open alerts**: Number of open alerts assigned to the service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: For technology management services, the counts for **Open alerts** and **Open critical alerts** include alerts for related configuration items (CIs). To view the individual alerts and their associated CIs, open the service and select the **Related alerts** tab. The alert counts include related CIs up to 3 levels deep, with a maximum of 2,000 CIs.

- **Open incidents**: Number of open incidents assigned to the service.
- **Error budget remaining**: Percentage of error budget remaining for the service.

Service reliability

The Service reliability tab is a customizable dashboard showing high-level service performance. For more information about the dashboard, see Visualizations in the Service reliability dashboard.

Add a service to SRM

Add services to Service Reliability Management (SRM) to help your teams manage service health.

Before you begin

If you're adding a technology management service to SRM, make sure it doesn't have a support group. You assign the support group when you add the service to SRM.

Role required: srm_manager, srm_responder, or srm_admin

About this task

Add a service to SRM to assign it to your teams and let them access its incidents and alerts. You can either create a service or import an existing one. SRM supports both service instances and technology management services:

- Service instances - Business-facing services, such as an HR portal. Service instances often have multiple infrastructure components.
- Technology management services - IT-focused services that support a business service or service instance. For example, an authentication service or database service.

When you add a technology management service to SRM, you also add its offerings. Offerings are specific components within a parent technology management service. For example, Spam filtering is an offering within the Email technology management service. In SRM, you can assign a parent technology management service and its offerings to different teams.

Note: When you import an existing technology management service, its offerings don't inherit its support group. Offerings either keep their existing support group or, if they don't have one, you must assign it before setting up Service Level Objectives (SLOs).

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. Select the Services icon and then select **Add**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the Add services for your teams modal, you can either create a service or add existing services.

Option	Steps
Create a new service	**a.** Select **Create a new service**. **b.** Fill in the fields. For field descriptions, see Add service form. **c.** Select **Next**.
Select existing services	**a.** Select **Select existing services**. **b.** Search for and select your service. You can add more than one. **c.** Select **Next**. **Tip:** Select the search box to view up to 25 services not yet in SRM, ordered by creation time. After you enter three characters, it shows up to 25 matching results. Both lists can include services with or without a support group. If a ser vice has one, the support group can be an SRM team or a group outside SRM.

4. In the Assign a team step, assign your service to an existing or new support group. In SRM, support groups are also known as teams.

Option	Steps
Select an existing team	Enter your team's name in the search box.
Create a team	Select **Create a team**. For more information about creating teams, see Add an SRM team.

5. Select **Add service**.
If approvals aren't enabled, your service appears in the service list. If approvals are enabled, your request is sent for approval. After it's approved, the service appears in the services list.

Note: If you added a technology management service, its offerings appear in the Offerings tab on the parent technology management service's page. Offerings also appear in the services list. To view an offering without an assigned support group, remove all filters from the services list.

What to do next

If an offering doesn't have a support group, assign it to a support group before setting up SLOs. For how to assign support groups, see Edit service details in SRM.

You can also add integrations to your service to start monitoring its health.

Edit service details in SRM

Edit an existing service owned by your team. For each service that you support in SRM provide general information about the service and the SRM team that supports it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: srm_admin, srm_manager, or srm_responder

Procedure
1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select the **Services** icon ( ).
3. From your services list, select the service that you want to edit or update.
   To view and select a technology management service offering without an assigned support group, remove all filters from the services list.
4. On the Overview page, select the **Details** tab.
5. Fill in or edit the fields.
   For more information on the field descriptions, see Edit service details form.
6. To save the updated service details after editing, select **Save**.

Remove a service from SRM

Remove a service from Service Reliability Management (SRM) when you no longer need to track or monitor its reliability.

Before you begin
SRM administrators can remove any service from SRM. SRM managers can remove only the services assigned to their teams.

Note: Administrators and SRM administrators have a delete option on service pages. That option deletes the service from both SRM and the Configuration Management Database (CMDB). The remove option, described below, only removes the service from SRM.

Role required: srm_manager or srm_admin

About this task
Remove one or more services from SRM when you no longer need to track their reliability. Removing a service from SRM does the following:

- Removes the service from SRM views, including the Service page and Service reliability dashboard.
- Deletes all draft service level objectives (SLOs) associated with the service. Those SLOs aren't recoverable.
- Retires all active SLOs associated with the service. Those SLOs are no longer active but remain visible in SRM.

Removing a service from SRM doesn't affect the service reliability tasks, such as alerts, change tasks, or incidents. The service also remains in the CMDB, and you can add it back to SRM at any time.

Note: The following procedure shows how to remove services from SRM using the Services page. You can also remove services using the Teams page by selecting a team and then the **Services managed** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
2. In the primary navigation, select **Services**.
3. Point to a service you want to remove, and then select its check box.
   You can select one or more services.
4. Select **Remove**.
5. In the dialog, review the details about the affected SLOs and then select **Remove service**.
   The page displays `<service-name> removed from SRM`. The service no longer appears in the SRM application, but it remains in the CMDB. You can add it back to SRM at any time.

Note: If you receive an error when removing a service, see Troubleshoot SRM.

What to do next

To learn how to monitor and manage services that remain in SRM, see Working with SRM services. To add a service that you want to monitor or that was removed by mistake, see Add a service to SRM.

Working with integrations in SRM

Connect your services to monitoring tools using the Integrations Launchpad . Integrations send information to Service Reliability Management (SRM), helping you track alerts, manage incidents, and maintain service health.

Access the Integrations Launchpad

As an SRM admin, manager, or responder, you can access the Integrations Launchpad in the following ways:

- From the primary navigation, select the ITOM Admin Experience icon, and then select either **Add integration** or **Manage installed integrations**.
- From the Services page ( ), select a service, select the **Integrations** tab, and then select **Add an integration**.
- From the Services page ( ), select a service, expand **Complete setup**, and then select **Add integrations**.

Integration details specific to SRM

Be aware of the following when using the Integrations Launchpad for SRM:

Connector type support

SRM supports custom, pull, and push connectors. Pull connectors retrieve data from external sources, and push connectors send data from external sources to your instance.

Note: For information about setting up connectors, see Integrations Launchpad in Service Operations Workspace for ITOM.

Configuration item field

When you set up an integration for SRM, the Provide details step has an additional field called Default configuration item.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set this field to the service your integration monitors. If an incoming alert doesn't include service details, SRM uses this field to connect the alert to the relevant service. This field helps make sure that your teams receive alerts for the services they support.

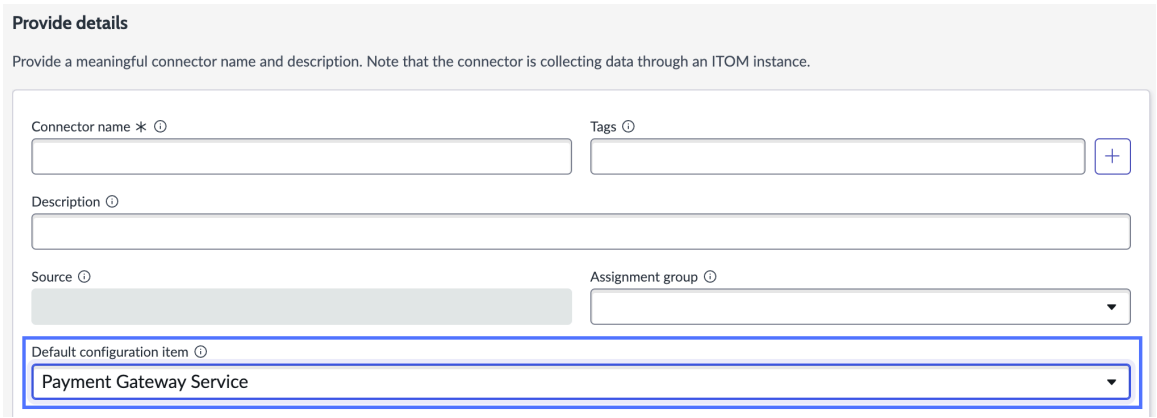

Learn more

Visit the following links to set up and manage integrations with the Integrations Launchpad. The pages are in a different section because the Integrations Launchpad is part of the broader Service Operations Workspace.

- Activate integration
- Configure an event custom connector
- Configure an event pull connector
- Configure an event push connector
- Deactivate integration
- Delete integration

Working with SRM teams

Manage schedules and define escalation policies for your team. That way, your team sees who is on call and accountable and can have the confidence that critical alerts or incidents are acknowledged in a timely manner.

On the Teams landing page, view and manage your teams and team requests.

Teams landing page

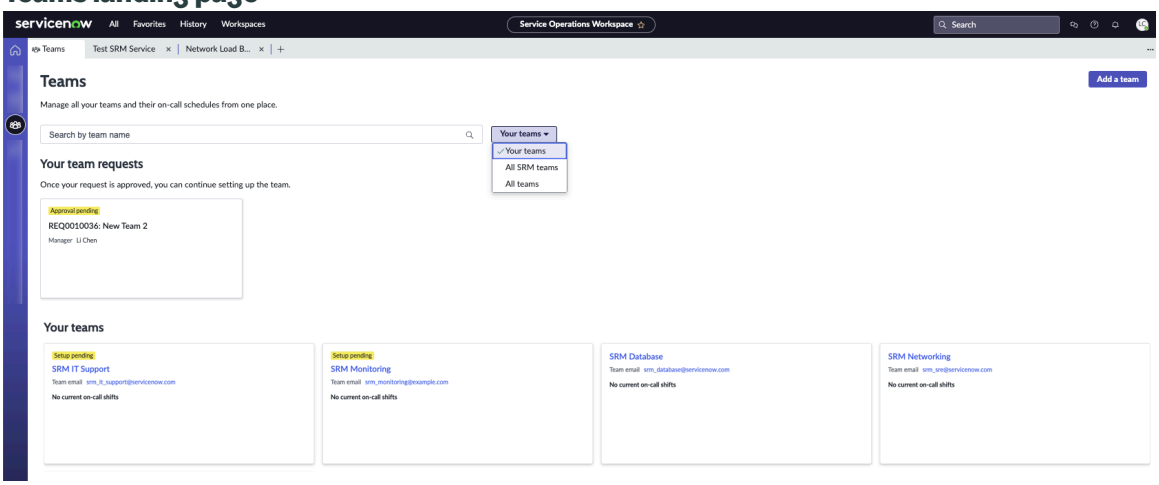

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Search by team name among teams you belong to, all SRM teams, or all teams in your organization.

Your selection determines the cards that you're shown on the Teams landing page.

Choices are:

- **Your teams** (default)
- **All SRM teams**: All the SRM teams.

  Note: Your ability to take actions in each team depends on permissions for that team.

- **All teams**: All teams you have access to in your organization.

  Note: Your ability to take actions in each team depends on permissions for that team.

The Your team requests section contains team cards for any pending requests, if there are any.

The Your teams section contains cards that provide information about a specific team, such as team status, name, manager, email, current shift name and duration, and primary responder. See Manage your SRM team details for more information.

Add an SRM team

Request to create a team to start monitoring your services. Teams are responsible for the issues that occur in the associated services.

Before you begin

Role required: srm_manager, srm_responder, or srm_admin

About this task

If you have the srm_admin or admin role, you can set up teams using a different process. See Activate teams and services in SRM for more information.

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. Select **Teams** ( ), select **Add a team**.
3. Create a new team or add an existing one.

Option	Steps
Add a new team	**a.** In the **New or existing** field, select **I am creating a new team**. **b.** On the form, fill in the fields. For more information on the field descriptions, see Create team form. **c.** Select **Add team**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Steps
Register an existing team	**a.** In the **New or existing** field, select **I want to give a group access to SRM app**. **b.** Select an existing team from the **Select team** list. **c.** Select **Add team**.

Result
When the request is approved, the team is created. Your team appears in the Teams landing page under **Your teams**.

What to do next
- Manage your SRM team details
- Set up escalation policies for your team in SRM
- Create an SRM on-call schedule
- Add team members to SRM
- Manage or add a service to your SRM team

Manage your SRM team details

Add or edit team members, description, contact information of you team.

Before you begin
Role required: srm_manager, srm_responder, or srm_admin

Procedure
1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Teams** ( ).
3. Select the team card whose details you want to update.
4. On the form in the **Details** tab, update the details as required.

Field	Description
Name	Unique name of the team.
Manager	Person who manages the team.
Group email	Email contact for the team that includes all team members.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Description	Description of the team that helps others to understand its responsibilities, purpose, or location.

5. Select **Save**.

Add team members to SRM

Add new team members to your team as required to plan, manage, resolve issues, and increase the efficiency of your business operations.

Before you begin

Role required: srm_manager or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Teams** ( ).
3. Select the team card to open.
4. Select **Teams members** tab.
5. Select **Add**.
6. In the pop-up window select members to add from the list menu.

   Note: Member details are imported from user records and not editable.

7. Select **Add**.

Manage or add a service to your SRM team

Add a service to an incident to increase the amount of data for the incident.

Before you begin

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Teams** ( ).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select a team card to open any team you want to view or add a service to.
4. From the team page, select the **Services managed** tab.
   You can export the list of services. See Export list information to a file for more information.
5. **Optional:** Add a service to your team by selecting **Add**.
6. Add a new service or register an existing one.

Option	Steps
Add a new service	**a.** In the **New or existing** field, select **Create a new service in Service Reliability Management**. **b.** On the form, fill in the fields. For more information on the field descriptions, see Add service form. **c.** Select **Next**.
Register an existing service	**a.** In the **New or existing** field, select **Associate existing CMDB services to Service Reliability Management**. **b.** Select one or more services. **c.** Select **Next**.

7. In the Assign a team step, fill in the fields.
   - In the **Service owner** field, select the user who will own the services.
   - In the **Support Group** field, select a team that will monitor services.

   The service record has been added.
8. Finish setting up your service or do it later.

Option	Steps
Finish setting up your service	**a.** Select **Continue to setup**. **b.** Select **Add integrations** and connect your service to a data source to monitor events. **c.** Select **Add SLO & SLI** and create a service level objective and measure your service health with a service level indicator.
Do it later	Select **Do it later**.

See Edit service details in SRM and Edit service details formfor information on the service and completion steps for setup.

Create an SRM on-call schedule

Set up an on-call schedule to make sure that someone is available to respond to incidents and critical alerts.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Before creating an on-call schedule, a team must exist in Service Reliability Management (SRM). The Alert table [em_alert] must also be added to the Trigger Rule Table Config list. See Add an SRM team and Add alert table to enable on-call scheduling for Service Operations Workspace for ITOM for more information.

Role required: srm_manager, srm_responder, or srm_admin

About this task

Shifts in SRM let you define variable time periods for your shift-based workforce. Create shifts for particular dates and times when you need coverage, and assign them to team members.

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace** > **Teams (** **)**.
2. On the Teams page, select your team.
3. Select the **Schedule** tab and then select **Create shift**.
4. In the Create new dialog, select either **Create new** or **Use template**.
   Your system admin can create predefined templates for you to use to create shifts.
5. Fill in the fields.
   For more information on the field descriptions, see Create a shift form in Service Reliability Management.
6. Select **Create shift**.

Set on-call team preferences in SRM

Set on-call preferences for team members to increase the efficiency of your business operations.

Before you begin

Role required: srm_manager or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Teams** ( ).
3. Select the team card to open.
4. Select the **On-call team preferences** button.
5. On the form, fill in the fields.

Field	Description
On-call shift managers	Team managers for the on-call shift.
PTO requests	Requests for Paid Time Off (PTO) made by team members. Choices are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ System default (Not allowed) ○ With approval ○ Without approval ○ Not allowed
First day of the week	First day of your work week. **Note:** This option is available for Washington DC and later releases only.
Allow shift schedule overlap	Option to allow schedules to overlap
Escalation rule on shift overlap	Choices are: ○ System default (Escalate to all shifts) ○ Escalate to outgoing shift ○ Escalate to incoming shift ○ Escalate to all shifts

6. Select **Save**.

Set up escalation policies for your team in SRM

Set up an escalation policy for your team to ensure that alerts or incidents are resolved in a timely manner by the appropriate team member.

Before you begin

Before creating an escalation policy, you must create an on-call shift for your team.

Note: Only the responder who created the team or the manager of the team can set up an escalation policy for that team.

Role required: srm_manager, srm_responder, or srm_admin

About this task

An escalation policy is a set of guidelines that outlines how and when incidents should be elevated to higher levels of authority. It delineates clear procedures and responsibilities for escalation to help responders promptly address problems. It defines the order in which notifications should be escalated if an alert or incident isn't acknowledged or resolved within a defined amount of time. If no one responds, the escalation is based on your defined policies until someone responds. After a team member responds, the escalation policy stops escalating and no further notifications are sent.

Note: Escalation workflow is triggered only if there's an active shift.

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. From the primary navigation, select **Teams** ( ).
3. Select the **Escalation triggers and policies** tab.
4. Under **Escalation policies** in the left panel, select **Create policy**.
5. On the Escalation policy form, fill in the fields.
   For more information, see Set up escalation policies form in SRM
6. **Optional:** Select **Add escalation step** to define the additional steps for escalation and escalation path.
   For more information on the Add escalation step form, see Set up escalation policies form in SRM.
7. Select **Save changes**.
8. **Optional:** Add another escalation policy to the team for another category of alert or incident by selecting **Add a policy**.

Create an escalation trigger

To address issues promptly, define the conditions that trigger a team's escalation policy in Service Reliability Management (SRM).

Before you begin

If your team needs trigger fields other than Assignment group, confirm that an admin has already configured them. For more information, see Configure fields for escalation triggers.

Role required: srm_manager or srm_admin

About this task

This procedure shows how to create escalation triggers in SRM. If your organization manages on-call scheduling using IT Service Management (ITSM), you can also use ITSM On-Call Scheduling to define escalation triggers. See Create an escalation trigger rule for more information. While that procedure lists ITSM roles, SRM users only need the srm_manager or srm_admin role.

For a broader overview of escalation processes in ITSM, see Designing an escalation process.

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
2. From the primary navigation, select **Teams** ( ).
3. Select your team, and then select the **Escalation triggers and policies** tab.
4. In the **Escalation triggers** panel, select **Create trigger**, and fill in the Escalation trigger form.

Field	Description
Escalation trigger name	Name of the trigger. Example: `P1 incidents`.
Order	Order of execution for the trigger. When a team has multiple escalation triggers, the triggers with lower values are evaluated first.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Active	Option to activate or deactivate the trigger.
Conditions	Choose the table (**Alert** or **Incident**) and set the conditions for the trigger. When conditions are met, the system triggers an escalation for the relevant team. Example: `Priority is 1 - Critical`.
Advanced options	Define what happens when the trigger conditions are met and when the trigger runs. For **Trigger action**, select **Workflow**, **Subflow** or **Script**, then choose the specific item. Example: Select **Workflow**, and **On-Call: Assign by Acknowledgment** to auto-assign the incident to the person who acknowledges it. Select **Run trigger rule every time the selected field changes** to run the trigger when one of the chosen fields updates. If not selected, the trigger runs only when the Assignment group changes.

5. Select **Save**.
   The UI displays `The Escalation trigger has been added successfully`. The trigger is now active and runs when the defined conditions are met.

Working with SRM reliability tasks

Alerts, incidents, and change requests are reliability tasks. From creation to resolution, SRM helps you manage your alerts throughout the response life cycle.

The Reliability Tasks pane ( ) has several tabs:

- **Incidents** are created manually or based on Service Level Objectives (SLOs) and Service Level Indicators (SLIs). For more information on incidents, see Working with incidents in SRM. For more information on SLO, SLI, and error budget actions, see Create SLOs, SLIs, and error budget policies.
- **Changes** can be created within SRM. See Manually create an SRM change task.

The Express List pane ( ) provides a focused view of alerts to help you monitor, evaluate, and resolve health issues efficiently. You can create alerts from SRM service pages or by using integrations. For more information on alerts, see Express List in the Service Operations Workspace for ITOM and Manually create an SRM alert.

Responders and above are notified for updates to reliability tasks based on their notification preferences.

Working with alerts in SRM

Monitor, evaluate, and resolve health issues in the Express List.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Event Management Express List feature provides a customizable, sortable view for accessing and managing alerts. Use the Express List feature to prioritize alerts, understand their impact, and resolve issues faster.

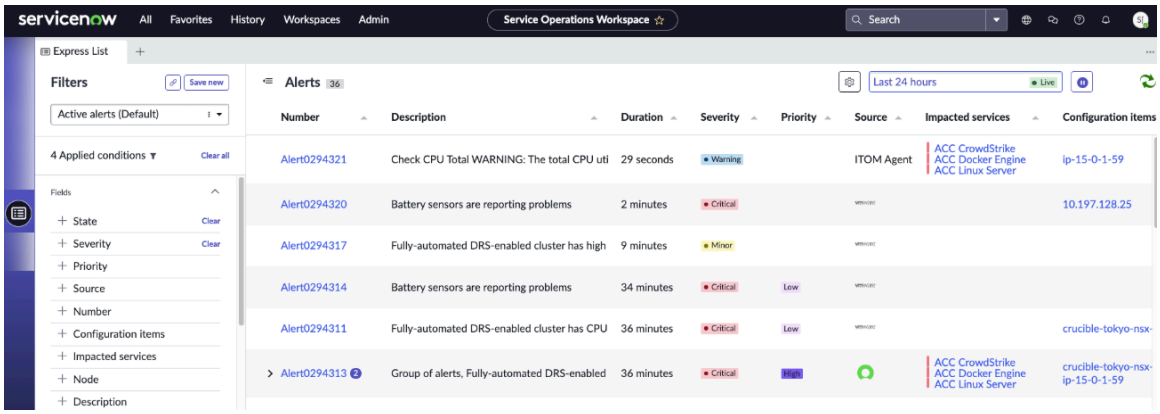

As a Service Reliability Management (SRM) user, you automatically have the evt_team_operator role, which lets you access the Express List feature ( ). To learn more about that role, see Roles used by Express List.

Visit the following to work with alerts in the Express List feature:

- Acknowledge viewed alerts in Express List
- Assign alerts in Express List
- Close an alert in Express List
- Create an incident from an alert in Express List
- Monitor incoming alerts
- Place an alert in the Acknowledged state
- Place an alert in the Maintenance state

In SRM, integrations create alerts or you create alerts manually. Visit the following link to learn how to create an alert from the **Services** page:

Manually create an SRM alert

Create an alert for testing purposes or if you think an issue poses a risk.

Before you begin

Role required: srm_admin, srm_manager, or srm_responder

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. In the navigation pane, select **Services**.
3. Select the service that you want to trigger an alert for.
4. Select the **Related alerts** tab.
5. Select **Trigger alert** and fill out the **Trigger an alert** form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To trigger an alert, you must first create an event to get processed. Only **Service** is required.

Trigger an alert

Field	Description
Service	Select a service for the event.
Source	Manual.
Node	Enter the name of the node (for example, FQDN, IP address, or MAC address) associated with the event.
Type	Enter a type of event.
Resource	Enter a node resource (process or service) associated with the event.
Metric Name	Enter the name of the metric associated with the event.
Source instance	Enter a data source for the event.
Message key	Enter an identifier for multiple events related to the same event.
Severity	Select the expected impact of this event. Choices are: ○ Critical: Immediate action is required. The resource is either not functional or critical problems are imminent. ○ Major: Major functionality is severely impaired or performance has degraded. ○ Minor (default): Partial, non-critical loss of functionality or performance degradation occurred. ○ Warning: Attention is required, even though the resource is still functional. ○ OK: The resource is still functional.
Time of the event	Select the date and time that the event occurred in the source system.
Description	Enter information describing the event.
Additional information	Add more information relevant to the event.
Processing notes	Enter relevant processing information.

6. Select **Trigger**.
   Service Reliability Management (SRM) creates an event to generate the alert. Refresh the screen to see the event processing results.

   Express List automatically displays your alert. If enabled, SRM triggers the applicable automations and on-call escalation workflows.

What to do next

Consider making the following optional adjustments to your alert for testing purposes or team collaboration:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Acknowledge the alert or assign it to someone on your team.
- Add comments or work notes in the **Compose** panel.
- Add an attachment related to the alert using the attachment icon .

Working with incidents in SRM

Plan ahead of service disruptions and have SRM send notifications and create status when incidents occur. Distractions are minimized and teams stay focused on remediation.

You can manually create incidents within SRM or create an error budget action to do so. See Manually create an SRM incident or Create SLOs, SLIs, and error budget policies for more information.

- Manual creation from an alert.
- Manual creation from the incident view.
- Automatically from Alert automations, if Alert automations are installed.

The **Assigned to** field on an incident specify who should be notified. When a team is selected as a responder, team automations are checked to determine which schedule to use for the notifications.

Note:

The **Assigned to** field is cleared when the **Assigned-team** or **Service** field is updated on an incident. Escalation policies for the newly assigned teams run. The field remains cleared until a user on the new team acknowledges an escalation notification.

If the Service is changed and the new Service does not have an assigned team, no changes occur.

When a service is deleted, its integrations, alerts, incidents, and automations are removed. Deleting a service isn't a recoverable action so consider deactivating the service instead.

Responders and above are notified for updates to incidents based on their notification preferences. If you made the update, you won't be notified.

You can resolve incidents by selecting specific incidents or in bulk. See Resolve an SRM incident, or Close an SRM incident or Cancel an SRM incident.

Note:

Incidents in the **Resolved** state are automatically closed after 3 days.

For more information on the areas and fields available in an incident, see SRM incidents.

Export list information to a file from the list view.

Respond to an incident in the following ways:

- Tag an incident.
- Update the priority of an SRM incident.
- Set the state of the incident.
- Copy an SRM incident.
- Add a service to an SRM incident.
- SRM incidents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Resolve an SRM incident.
- Cancel an SRM incident.
- Close an SRM incident.
- Reopen an SRM incident.

Manually create an SRM incident

Create an incident if you think an issue poses a serious risk and should be taken care of as soon as possible.

Before you begin

Role required: srm_manager, srm_responder, or srm_admin

About this task

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Select the Incidents tab.
4. Select **New**.
5. Fill in the fields on the Create New Incident form, as appropriate.

   Only the **Short Description** field is required.

Incident Field mapping

Field	Description
Short Description	Add brief description of the incident.
Description	Add detailed description of the incident.
Number	Identifier assigned to the incident.
Impact	Select is a measure of the effect of the incident. Impact affects **Priority**. Choices are: ○ 1 - High ○ 2 - Medium ○ 3 - Low
Assignment group	Select a team to work on the incident by typing or searching.
Priority	Automatically set based on **Impact** and **Urgency**. It identifies how quickly the incident should be addressed. See **Priority Data lookup rules** table below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Values are: ○ 3 - Moderate ○ 4 - Low ○ 5 - Planning
Category	Select a category for the incident. Choices are: ○ --None-- ○ Inquiry/Help (default) ○ Software ○ Hardware ○ Network ○ Database
Urgency	Select a measure of how long the resolution can be delayed until an incident has a significant business impact. Urgency affects **Priority**.
Service	Select a service associated with the incident.
Severity	Select a severity to associate with the incident by typing or searching. Choices are: ○ 1 - High ○ 2 - Medium ○ 3 - Low
Configuration item	Select a asset to associate with the incident by typing or searching.
Assigned to	Assign team member to work on the incident by typing or searching.
State	Choose a state for the incident. Choices are: ○ New (default) ○ In Progress ○ On Hold ○ Resolved ○ Canceled See SRM incident states for more information.
Watchlist	Select individuals to watch the incident by typing or searching.
Resolution code	Select a resolution code when you have one. Choices are: ○ None (default) ○ Duplicate ○ Known error ○ No resolution provided ○ Resolved by caller ○ Resolved by change

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ Resolved by problem ○ Resolved by request ○ Solution provided
Resolution notes	Enter any relevant resolution information.

Note: Priority is calculated according to the following sample data lookup rules:

Priority Data lookup rules

Impact	Urgency	Priority
1 - High	1 - High	1 - Critical
1 - High	2 - Medium	2 - High
1 - High	3 - Low	3 - Moderate
2 - Medium	1 - High	2 - High
2 - Medium	2 - Medium	3 - Moderate
2 - Medium	3 - Low	4 - Low
3 - Low	1 - High	3 - Moderate
3 - Low	2 - Medium	4 - Low
3 - Low	3 - Low	5 - Planning

6. Select **Save**.
   The incident is created and the new incident form displayed.
7. Add an attachment related to the incident in the **Actions** panel using the **Attachments** icon .
8. Create an incident template using the Template icon .
9. Start or join a Zoom meeting, Microsoft Teams, or Slack channel in the **Collaboration** panel..
10. Select **Save**.
    The incident appears with comments, work notes, and incident activity captured in the **Activity** stream. For more information on incident fields, see SRM incidents.

Reassign an SRM incident

Reassign an incident to a responder when the incident tasks should be addressed by a particular team member.

Before you begin

This action is available when the **Assigned to** field in the incident is populated.

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Select the **Incidents** tab.
4. You have two options.

Option	Description
Option	Description
In the incident list view	Double-click the **Assigned to** field for the inci dent and replace the team member from the list menu.
In the Details tab of incident form	Delete the name in the **Assigned to** field and reassign. Select **Save**.

The incident is now assigned to the new responder and an email notification sent. The assignee is displayed in the list view and on the form.

Update the priority of an SRM incident

If the priority of an incident should be changed, you can manually update it to reflect its new criticality.

Before you begin

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
   You're taken to your SRM Home page.

Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Select the **Incidents** tab.
4. In the Incident list view, select a **Priority** value from the list menu.
   - 1 - Critical
   - 2 - High
   - 3 - Moderate
   - 4 - Low
   - 5 - Planning (default)
5. Select **OK**.
   The new priority is reflected in the incident list view and incident header and form field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add a service to an SRM incident

Add a service to an incident to increase the amount of data for the incident.

Before you begin
Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. You have two options.

Option	Description
In the incident list view	Double-click the **Service** field for the incident and replace the team member using search. Select **OK**.
In the incident form	Select a new **Service** value from the list menu. Select **Save**.

   For value choices see SRM incidents.
   The new service is displayed in the incident list view and incident form.

Add resolution code and notes to an SRM incident

Add a resolution code and notes to an incident when you have resolved the incident.

Before you begin
Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Open the incident.
4. In the Incident **Details** panel, select a **Resolution code** from the list menu.
5. Enter text in the **Resolution notes** text box.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The resolution notes are displayed in the **Activity** timeline panel.

6. Select **Save**.

Copy an SRM incident

Copy an incident, if you have a similar issue.

Before you begin

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Open the incident you want to copy.
4. Using the list menu under the More option button, select **Copy incident**.
5. Select a team member from the list menu for the **Caller** field.
   After the incident is copied, the **Activity** timeline of the new incident is updated with the original incident as a work note.

Resolve an SRM incident

Resolve an incident, if you need it reviewed, or you think it has already been handled, but aren't ready to close it.

Before you begin

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Select the **Incidents** tab.
4. You have two options.

Option	Steps
In the incident list view	**a.** Double-click the **State** field (or use the key board shortcut) for the incident and select from the list menu. **b.** Select **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Steps
	Tip: If you have the srm_admin role, you can update multiple incidents at once. First, select the relevant incidents and select **Edit**. Then use the **State** drop-down to choose the appropriate value and select **Update**.
In the incident form	**a.** Select a new **State** field value from the list menu. **b.** Select **Save**.

For value choices, see SRM incidents.
The new state is displayed in the incident list view, incident form, and form header.

Result
Resolved incidents remain in the Incidents list view. Check your filters if you don't see them.

Note:

When you resolve an incident, the related alerts remain open until you close the incident.

Cancel an SRM incident

Close an incident, if you think it is not an issue, or if it has already been handled.

Before you begin
Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Open the incident you want to cancel.
4. You have two options.

Option	Steps
In the incident list view	**a.** Double-click the **State** field (or use the key board shortcut) for the incident and select from the list menu. **b.** Select **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Steps
	Tip: If you have the srm_admin role, you can update multiple incidents at once. First, select the relevant incidents and select **Edit**. Then use the **State** drop-down to choose the appropriate value and select **Update**.
In the incident form	**a.** Select a new **State** field value from the list menu. **b.** Select **Save**.

For value choices, see SRM incidents.
The new state is displayed in the incident list view, incident form, and form header.

Result
Canceled incidents remain in the Incidents list view. Check your filters if you don't see them.

Note:

When you cancel an incident, the related alerts are closed.

Close an SRM incident

Close an incident when it's not an issue or has already been handled.

Before you begin
Role required: srm_responder, srm_manager, or srm_admin

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. You have two options.

Option	Steps
In the incident list view	**a.** Double-click the **State** field (or use the key board shortcut) for the incident and select from the list menu. **b.** Select **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Steps
	Tip: If you have the srm_admin role, you can update multiple incidents at once. First, select the relevant incidents and select **Edit**. Then use the **State** drop-down to choose the appropriate value and select **Update**.
In the incident form	**a.** Select a new **State** field value from the list menu. **b.** Select **Save**.

For value choices, see SRM incidents.
The new state is displayed in the incident list view, incident form, and form header.

Result
Closed incidents remain in the incidents list view. If they don't appear, try adjusting your filters.

Note:

When you close an incident, any related alerts are closed.

Reopen an SRM incident

Reopen a closed incident if you think the issue is not yet resolved.

Before you begin
Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *Service Operations Workspace*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. From the primary navigation, select **Reliability tasks** ( ).
3. Find and open the incident you want to reopen.
   Check your filters if you don't see them.
4. You have two options.

Option	Steps
In the incident list view	**a.** Double-click the **State** field (or use the key board shortcut) for the incident and select from the list menu. **b.** Select **OK**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Steps
	Tip: If you have the srm_admin role, you can update multiple incidents at once. First, select the relevant incidents and select **Edit**. Then use the **State** drop-down to choose the appropriate value and select **Update**.
In the incident form	**a.** Select a new **State** field value from the list menu. **b.** Select **Save**.

For value choices, see SRM incidents.
The new state is displayed in the incident list view, incident form, and form header.

Manually create an SRM change task

Create a change task to implement a controlled process for modifying approved and supported actions for SRM.

Before you begin

For general information on Change requests, see Using Change Management.

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Navigate to **Workspaces** > *`Service Operations Workspace`*.
   You're taken to your SRM Home page.

   Note: If you use other Service Operations Workspace (SOW) applications, you may see the SOW Home page instead of the SRM Home page. The SOW Home page includes SRM alerts and incidents in its metrics.

2. Select **Reliability tasks** ( ).
3. Select the **Changes** tab.
4. Select **New** from the **Changes** list header.
5. In the **Select a change** window, search for a change record or select an SOW change creation page card to select the type of change request.
6. Select **Next**.
7. In the form Details tab, fill in the fields on the **Change request**, **Impact**, **Assignment**, **Planning**, **Schedule**, **Notes**, and **Closure information** forms.
   For more information, see Create change request forms.
8. **Optional:** Select **Assign to me**.
9. Select **Save**.

   The change request is created and a **Related Records** tab appears for the request.

10. **Optional:** You can take the following steps in **State** workflow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

(Optional) If the change request requires approval, you'll see the **Request approval** button. Select **Request approval** and **State** moves to **Assess**. See Approve a change request for SRM for instructions on approving a change request.

If the change request is pre-approved or doesn't require approval, you can move the **State** to **Scheduled**.

Once approved and scheduled, you can select **Implement** to put the change request into action.

Once implemented, you only after reviewing the details on the change request, select Review. The change request is moved to the **Review** state.

Note: All open change tasks are set to **Canceled**.

For more information on **State** in Change requests, see, SRM change request states

Service Reliability Management reference

Reference topics provide additional information about administering Service Reliability Management.

SRM roles and responsibilities

Roles grant users access to different parts of the SRM console. Roles determine the actions that users can or can't perform in Service Reliability Management.

SRM has the following predefined roles:

Note: The SRM experience was optimized for you to having just the SRM roles. Additional roles could change the UI and have unintended consequences.

Roles List

Role	Tasks
SRM Administrator [srm_admin]	• Update a service managed by any SRM team • Assign new incidents • Create and edit incidents • Create and edit changes • Create teams • Update any SRM team • Create and edit on-call shifts • Delete on-call shifts • Create and edit escalation policies and triggers • Delete escalation policies and triggers • Schedule coverage request • Submit time off • Create event

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles List (continued)

Role	Tasks
SRM Manager [srm_manager]	• Update a service managed by your SRM team • Assign new incidents • Create and edit incidents • Create and edit changes • Create teams • Update SRM teams that you manage • Create and edit on-call shifts • Delete on-call shifts • Create and edit escalation policy and trigger • Delete escalation policies and triggers • Schedule coverage request • Submit time off • Create event
SRM Responder [srm_responder]	• Update a service if you are in the support group tied to the service • Assign new incidents • Create and edit incidents • Create and edit changes • Create teams • View on-call shifts • View escalation policies and triggers • Edit contact preferences • Schedule coverage request • Submit time off • Create event
Administrator [admin]	• Update a service managed by your SRM team • Delete services • Assign new incidents • Create and edit incidents • Create and edit changes • Delete incidents • Create teams

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles List (continued)

Role	Tasks
	• Update SRM teams that you manage • Delete on-call shifts • Override contact preferences • Manually create an alert • Delete alerts • Submit time off

Create team form

Request to create a team to start monitoring your services.

Create a team form

For more information, see Add an SRM team.

Field	Description
Team name	Name the team.
Team manager	Manager for the team.
Team email	Email contact for the team.
Team members	Members for your team.
Tags	Words that describe or relate to your team.
Briefly describe the team responsibilities	Description of your team. For example, team responsibilities or geographic location.

Create a shift form in Service Reliability Management

Set up an on-call schedule for your team to ensure that the right person is always available to quickly respond to incidents and critical alerts.

Create a shift form

For more information, see Create an SRM on-call schedule.

Create shift

Field	Description
Shift name	Name for the shift.
All day shift	If selected, the shift is an all-day commitment.
Shift starts	Date and time the shift begins.
Shift ends	Date and time the shift ends.
Repeats	Shift repetition. Choices are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create shift (continued)

Field	Description
	• Never • Daily • Weekly
Every	How often to repeat. Display depends on the time frame chosen.
More options	More options to customize your schedule.
Timezone	User time zone.
Holiday schedule	Schedule to use for the holiday.
Repeats until	Date and time to end repetition.

Create shift using shift template

Field	Description
Shift name	Name for the shift.
Template/Schedule	Existing templates for shifts. Choices are: • 24 x 7 • Off hours (workday 8:00 - 5:00) • Off hours (workday 9:00 - 5:00) • Workday 8:00 - 5:00 • Workday 9:00 - 5:00

Add service form

Create or add an existing service to your instance so that alerts and incidents on that service are available to your teams within SRM.

Add a service form

For more information, see Add a service to SRM.

Add services for your teams

Field	Description
Service name	Enter a name for your service. For example, Cloud storage service. Your teams see this name in SRM.
Service class	Select either service instance or technology management service .
Business criticality	Specify how important the service is to the business:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add services for your teams (continued)

Field	Description
	• 1 - most critical • 2 - somewhat critical • 3 - less critical • 4 - not critical
Operational status	Select the status of your service: • Operational • Non-operational • Repair in Progress • DR Standby • Ready • Retired • Pipeline • Catalog
Short description	Describe your service to help your teams understand its purpose and other relevant information.

Edit service details form

Edit an existing service owned by your team.

Service Details form

For more information, see Edit service details in SRM.

Field	Description
Name	Name of the service.
Support group	SRM team managing the service.
Business criticality	How important this service is the business.
Owned by	Service owner. May not be part of an SRM team.
Operational Status	Status for the service's ability to perform operational functions.
Stakeholders	Product owners or service consumers who determine the target availability of a service.
Environment	Choices are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• None • Development • Production • Test
Managed by group	Team manager responsible for the service. Add by search or edit using the text editor.
Change group	Team responsible for changes. Add by search or edit using the text editor.
Status	Status of the service.
Life Cycle Stage	Choices are: • Deploy • Design • End of Life • Ideation • Operational
Phase	Phase of the service. Choices are: • None • Pipeline • Catalog • Retired
Location	Address where the service resides. Add by search or edit using the text editor.
Used for	Purpose for the service. Choices are: • Production (default) • Staging • QA • Test • Development • Demonstration • Training • Disaster recovery
Short description	Additional information about the service.
Support by	SRM team managing the service. Add by search or edit using the text editor.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up escalation policies form in SRM

Escalation policies define the order and conditions in which members or devices receive notifications.

Escalation policy form

Field	Description
Name	Name of the escalation policy.
Active on shift	Team that receives the notifications on their shift..
Create from scratch	Option to create your own policy.
Create from template	Option to create a policy by using a template created by your admin.
Description	Text describing the policy.
Active	Option to activate or deactivate the policy.
Use as default	Option to use this team as the default
Order	Order of execution. This is the hierarchy or chain of command. When there are multiple escalation policies with the same conditions for a team, the escalation policy with the smaller order number is executed first.
Conditions	Conditions for the policy. The table and conditions that define the escalation policy These conditions are checked after the escalation is started by an escalation trigger. If the conditions are met, the policy is run.
Escalation steps	Escalation steps to add. See the Escalation steps form for more information about the fields.
Escalation notifications	Notification conditions for the policy. **Note:** When user preference override is on, your escalation policy overrides any policy that the on-call agents set for their own notification preferences. When user preference override is off, the on-call agents' notification preferences take priority over this escalation policy.
Add notification step	Additional notification steps to add. You can select email, call, or SMS. You can add as many attempts as you like.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Escalation steps form

Add escalation step

Field	Description
Escalation step name	Name of the escalation step.
Order	Order of execution. This is the hierarchy or chain of command.
Responder level	Level of authority or expertise needed to respond to the incident.
Rotate through members	Option to notify all members, one by one, until one of them acknowledges the incident or until all of them are notified.
Additional audience	Specific users, teams, groups, or devices that you want to notify.
Time to next step after last notification	The maximum amount of time that will pass before the issue is escalated to the next step. For example, if a responder can't resolve the issue within 30 minutes of the notification, the issue is escalated to the manager.

Domain separation and SRM

Domain separation is supported for Service Reliability Management (SRM).

Support level: Basic*

The support level is Basic but has some exceptions or special conditions.

- Business logic: Ensure that data goes into the proper domain for the application’s service provider (SP) use cases.
- The user interface, cache keys, reporting, rollups, and aggregations all use the domain at production run time.
- The owner of the instance must be able to set up the application to function across multiple tenants.

Sample use case: When an SP uses chat to respond to a tenant-customer’s message, the client must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview of domain separation

- Domain separation is present in all metadata tables in the application. Instance level tables also have domain separation explicitly. Other internal tables, such as the internal Team Request table, are domain-separated implicitly by referencing domain-separated records.
- System properties are not domain-separated, which has implications for Team Management features. The properties are shared by multiple domains and are set at the instance level. Domain-specific setup for these is not supported.
- Where another application is being leveraged (for example the Event Management Connectors) domain separation is determined by the domain-separation capabilities of that application.

How domain separation works in the Service Reliability Management application

The specific conditions indicated by the Basic* support level rating above relate to team management:

- New teams created through a catalog item backed by an Integration Hub flow are setup by the Service Operations Workspace administrator. They can initiate it from the Service Portal. An SRM admin, manager or responder can do it through the SRM workspace. Instructions appear in the SRM documentation.
- Integration Hub subflows and actions control management of SRM teams.
- Only one catalog item and flow can be set up for each instance. The customer is responsible for setting up team properties to support domain separation as a customization of the existing flows.
- The requestor of a team catalog item is associated with the domain and is available as part of the request item. As a result, if needed, the requester can create the team in a specific domain or the catalog Item and extend it to capture the domain another way. In either case, the Service Operations Workspace must make changes to the Integration Hub flows and actions to support this.

For more information, see Domain separation and On-Call Scheduling.

Related topics

Domain separation for service providers

Create change request forms

Fill out the fields in the forms to create a change task to implement a controlled process for modifying approved and supported actions for SRM.

Change request form

Create a new change request. For more information, see Manually create an SRM change task.

Field	Description
Short description	Brief description of the change request.
Description	Description based on Short description.
Number	Identifier for the change request.
Category	Category for the change request.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Choices are: • Hardware • Software • Service • System Software • Application Software • Network • Telecom • Documentation • Other (default)
Model	Model to use for the change request. Standard (default)
Risk	Risk the issue poses. Choice are: • High • Moderate (default) • Low
Type	Type of change request. Standard (default). Read-only.
Impact	Impact the issue poses. Choices are: • 1 - High • 2 - Medium • 3 - Low (default)
State	State of the change request. Choices are: • New (default) • Assess If approval is required further states are unavailable. • Scheduled • Implement • Canceled
Work notes list	Team members to receive notifications when work notes are added.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: An administrator must have already created an email notification for the work notes list for this to be available.
Work notes	Any additional information related to the request.

Impact, Assignment, Planning, Schedule, Notes, and Closure information

Field	Description
Configuration item	Asset affected by the issue.
Service	Service affected by the issue.
Service offering	Service offering affected by the issue.
Assignment group	Team assigned to the change request
Assigned to	User assigned to the change request.
Justification	Reasoning for the change request.
Implementation plan	Imported text.
Risk and impact analysis	Analysis of risk and impact of the issue.
Backout plan	Imported text.
Test plan	Imported text.
Planned start date	Expected date to begin change.
Planned end date	Expected date to finish change.
Actual start date	Date change implementation began.
Actual end date	Date change implementation ended.
Conflict status	Not run. Read-only.
Conflict last run	Based on Conflict status. Read-only
Watch list	Interested users.
Work notes list	Team members to receive notifications when work notes are added. **Note:** An administrator must have already created an email notification for the work notes list for this to be available.
Additional comments	Additional notes for the request.
Work notes	Additional notes on the work done for the request
Close code	Code for the change request. Choices are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	• None (default) • Successful • Successful with issues • Unsuccessful
Close notes	Notes on the change request closure.

SRM incident states

Incidents can be in a few different states depending on how they are acted on.

An incident is created in SRM when a third-party monitoring tool has an event. Incidents can be manually created or assigned. Edit actions from the incident form using the state button in the form. Use the **Save** button in the form to keep any changes you've made. Changes to incidents in the list view are automatically saved.

Incident states

State	Description
New	State on incident creation. The incident is logged but not yet worked on.
In Progress	The incident is assigned being worked on.
On Hold	The responsibility for the incident shifts temporarily to another entity to provide further information, evidence, or a resolution.
Resolved	A satisfactory fix is provided for the incident to ensure that it does not occur again.
Closed	The incident is closed when it is confirmed that the incident is satisfactorily resolved.
Canceled	Incident was triaged but found to be a duplicate incident, an unnecessary incident, or not an incident at all. To reopen an incident. Select **New**.

SRM alert states

Alert can be in a few different states depending on how they are acted on.

An alert is created in SRM when a third-party monitoring tool has an event. Alerts are manually assigned. An alert is one of four states. You can edit actions from either the Alerts list or the Alert form. Use the **Save** button in the form to keep any changes you've made. Changes to Alerts in the list view are automatically saved.

Note: You can acknowledge and assign an alert in any state.

Alert actions

State	Description
Open	State on alert creation. From this state you can acknowledge and set to: **Flapping**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Alert actions (continued)

State	Description
	Flapping occurs when the event source continues to generate events even after its associated alert has been closed. See Configure alert flapping for more information on event flapping. **Closed** Closes the alert.
Reopen	From this state you can acknowledge and set to: **Open** Sets the alert to the **Open** state. **Flapping** Flapping occurs when the event source continues to generate events even after its associated alert has been closed. **Closed** Closes the alert.
Flapping	Flapping occurs when the event source continues to generate events even after its associated alert has been closed. From this state you can set to: **Open** Sets the alert to the **Open** state. **Reopen** Reopens the alert. **Closed** Closes the alert.
Closed	From this state you can set to: **Open** Sets the alert to the **Open** state. **Reopen** Reopens the alert. **Flapping** Flapping occurs when the event source continues to generate events even after its associated alert has been closed.

SRM change request states

Incidents can be in a few different states depending on how they are acted on.

Change requests can be manually created in SRM. Edit actions from the change request form using the state field in the form or using buttons in the change request record. Use the **Save** button in the form to keep any changes you've made. Changes to change requests in the list view are automatically saved.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Change request states

State	Description
New	State on change request creation. The request is logged but not yet worked on.
Assess	The change request approval is being worked on.
Scheduled	The change request is scheduled for work.
Implement	The change request is taking action.
Review	The change request is moved to the **Review** state. **Note:** All open change tasks are set to **Canceled**.
Closed	The change request is closed.
Canceled	Change request was triaged but found to be a duplicate, or unnecessary. **Note:** You cannot reopen a canceled changed request.

CMDB service classes

Service Reliability Management (SRM) typically supports several Configuration Management Database (CMDB) service classes.

Types of CMDB service classes that SRM supports:

- Service instances (cmdb_ci_service_auto)
- Technology management services (cmdb_ci_service_technical)
- Query-based services (cmdb_ci_query_based_service)
- Mapped application services (cmdb_ci_service_discovered)
- Calculated application services (cmdb_ci_calculated)
- Tag-based application services (cmdb_ci_service_by_tags)

Troubleshoot SRM

Find answers to issues that you may encounter when using Service Reliability Management (SRM).

Condition

When trying to remove a service from SRM, you see one of the following messages:

- `<service-name> couldn't be removed from SRM. Check your permissions. Only SRM team managers can remove their team's services.`
- `<x> services couldn't be removed from SRM. Check your permissions. Only SRM team managers can remove their team's services.`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cause

You're not assigned the required role to remove the selected services. As an SRM manager, you can only remove services that belong to your teams. You can still view and filter services across teams, but you can't remove services owned by another team.

Remedy

To remove the service from SRM, follow these steps to identify the manager to contact:

1. Navigate to **Workspaces** > *`Service Operations Workspace`* > **Services**.
2. Select the filter icon ( ) and remove the existing filters, letting you view services not assigned to your teams.
3. Find and select the service you want to remove.
4. Select the **Details** tab and then select the information icon ( ) in the Support group field.

The manager's name appears below the team name. Contact the manager to request the removal of the service from SRM.

ITOM Mobile Agent

Stay connected and keep services reliable with ITOM Mobile Agent, which lets you access Service Reliability Management (SRM) features on iOS and Android devices. Use it to track alerts, manage incidents, work on change and catalog tasks, and manage on-call shifts on the go.

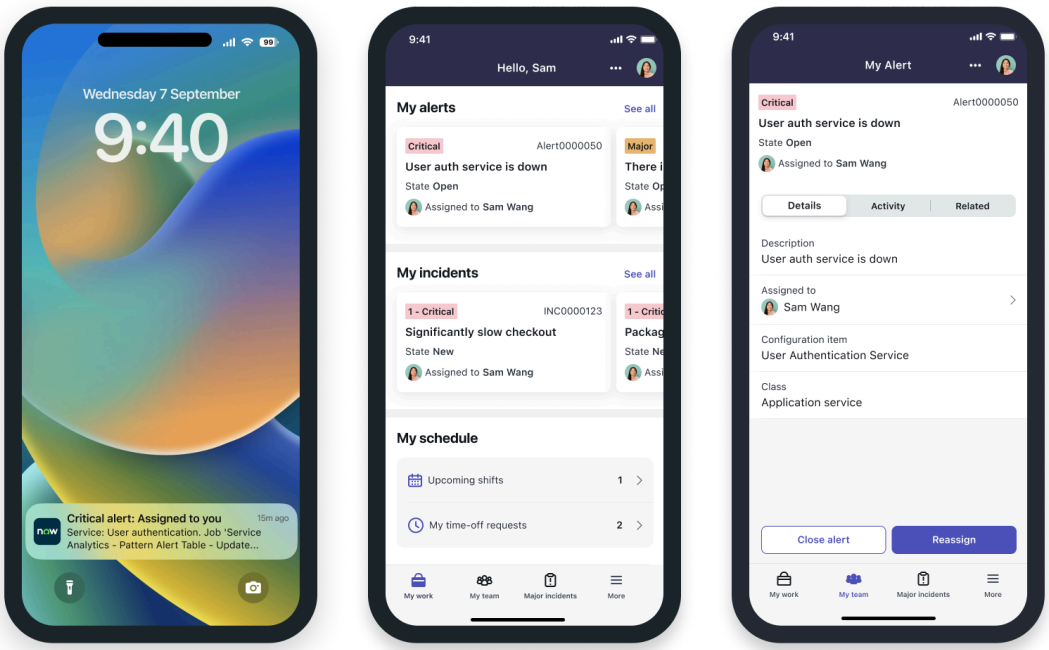

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore	Configure
Learn about ITOM Mobile Agent, including its benefits and common workflows.	Install ITOM Mobile Agent and configure notifications.
Use	**Reference**
Log in and work with alerts, incidents, tasks, and on-call schedules.	Get details about ITOM Mobile Agent settings.

Explore ITOM Mobile Agent

Learn how you can use ITOM Mobile Agent to manage service reliability on the go, including its capabilities, benefits, and a sample workflow.

ITOM Mobile Agent overview

Use ITOM Mobile Agent to access Service Reliability Management (SRM) features on iOS and Android devices. With ITOM Mobile Agent, you can track alerts, manage incidents, and work on change and catalog tasks on the go. You can also manage on-call schedules by checking shifts, requesting time off, and filling scheduling gaps.

To help you manage service reliability in one place, ITOM Mobile Agent auto-integrates with existing ServiceNow mobile apps. It runs in the ServiceNow Agent app and automatically includes ITSM Mobile Agent features. With that setup, you can manage alerts, incidents, on-call schedules, and tasks in one mobile app.

Note: If you already use ITSM Mobile Agent for incidents, on-call schedules, and tasks, install ITOM Mobile Agent to add alerts to your mobile experience. With ITOM Mobile Agent, you get and configure alert notifications and view alerts in the My work, My team, and Notifications tabs. See Install ITOM Mobile Agent for installation instructions.

ITOM Mobile Agent users

The following table lists user roles and what each user can do in ITOM Mobile Agent:

User	Description
Admin	Admins install ITOM Mobile Agent so teams can access service reliability features on the go.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

User	Description
SRM admin	SRM admins customize alert and incident notifications for ITOM Mobile Agent. They make sure that the relevant people are aware of critical issues.
SRM manager	SRM managers track team progress, monitor unassigned issues, and review performance metrics. They also handle on-call schedules, including shift rescheduling, time-off requests, and scheduling gaps.
SRM responder	SRM responders manage alerts, incidents, and tasks on the go. They accept, assign, and close alerts; create and collaborate on incidents; and process critical issues. They can also view upcoming shifts and request time off.

ITOM Mobile Agent workflow

This sample workflow shows how ITOM Mobile Agent helps you manage services and the teams assigned to them:

1. Monitoring tools generate critical alerts when they detect issues, and the ITOM Mobile Agent system sends mobile alert notifications to on-call responders.
2. As an on-call SRM responder, you acknowledge, reassign, or close the alerts. If needed, you create incidents from the alerts, and then the ITOM Mobile Agent system assigns them to the appropriate team.
3. You troubleshoot and collaborate on incidents, escalating them to major incidents when necessary, and then resolve and close the incidents, providing relevant details as needed.
4. As an SRM manager, you adjust shifts, approve time-off requests, and resolve scheduling gaps for the on-call responders. Use dashboards to view alert trends, incident resolution times, and team performance.

ITOM Mobile Agent benefits

Working with ITOM Mobile Agent has the following benefits:

Improve service reliability

Minimize downtime by managing alerts and incidents on the go.

Respond to issues faster

Keep on-call engineers connected, letting them accept, assign, and close alerts from anywhere.

Notify the right people

Customize push notifications to make sure that the relevant people are aware of critical issues.

Boost collaboration

Help teams work efficiently by managing on-call schedules and filling coverage gaps on mobile.

What to explore next

To learn more about configuring and using ITOM Mobile Agent, see:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Configure ITOM Mobile Agent
- Use ITOM Mobile
- ITOM Mobile Agent reference

Configure ITOM Mobile Agent

Set up and configure ITOM Mobile Agent.

Configuration overview

Use the pages in this section to set up and configure ITOM Mobile Agent. The information covers required and optional setup steps, including installing the app, managing alerts, and customizing alert recipients.

ITOM Mobile Agent auto-integrates with ITSM Mobile Agent, so you can manage alerts, incidents, and on-call schedules from the same mobile app. Where relevant, this section links to ITSM Mobile Agent content, providing additional configuration options for your mobile experience.

Required configuration steps

If you're an admin, start by installing ITOM Mobile Agent. If you're a Service Reliability Management admin, manager, or responder, you can start by downloading the ServiceNow Agent app and then logging in.

- Install ITOM Mobile Agent
- Download the ServiceNow Agent app
- Log in to ITOM Mobile Agent

Optional configuration steps

After you log in to ITOM Mobile Agent, you can work with the core features. Visit the following links to configure and customize ITOM Mobile Agent further to meet your team's needs:

- Configure when to send alert notifications for ITOM Mobile
- Customize ITOM Mobile Agent email recipients
- Customizing other mobile settings for ITOM Mobile Agent

Install ITOM Mobile Agent

Install and give your teams access to ITOM Mobile Agent.

Before you begin

Service Reliability Management (SRM) services, service integrations, teams, and on-call schedules must be set up. ITOM Mobile Agent connects to your existing SRM setup, helping your teams manage alerts, incidents, tasks, and on-call schedules.

Role required: admin

About this task

ITOM Mobile Agent runs in the ServiceNow Agent app. When you install ITOM Mobile Agent, your teams can access it in the ServiceNow Agent app on their devices.

This procedure is for admins managing mobile access for their teams. If you aren't an admin and want to use ITOM Mobile Agent on your mobile device, see Download the ServiceNow Agent app.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Search for and select ITOM Mobile Agent.
   If you can't find ITOM Mobile Agent, request it from the ServiceNow Store.
3. Review the listed dependencies and, if needed, install the required plugins and applications.
4. Select **Install**.

   After the application is installed, your teams can access ITOM Mobile Agent by downloading the ServiceNow Agent app and logging in to the relevant instance.

What to do next

To start using ITOM Mobile Agent, see Download the ServiceNow Agent app and Log in to ITOM Mobile Agent.

You can also customize ITOM Mobile Agent to meet your team's needs:

- Configure when to send alert notifications for ITOM Mobile
- Customize ITOM Mobile Agent email recipients
- Customizing other mobile settings for ITOM Mobile Agent

Download the ServiceNow Agent app

Download the ServiceNow Agent app to use ITOM Mobile Agent on your device.

Before you begin

ITOM Mobile Agent must be installed. See Install ITOM Mobile Agent for more information.

Role required: srm_manager, srm_responder, or srm_admin

About this task

The ITOM Mobile Agent application runs in the ServiceNow Agent app. To access ITOM Mobile Agent, download the ServiceNow Agent app on your device.

Procedure

1. Download the ServiceNow Agent app for your device:
   - For iOS, use the App Store.
   - For Android, use Google Play.
2. Confirm that the app is on your device.

   Look for the ServiceNow Agent app icon ( ). The app appears on your device with the name Agent.

What to do next

To start using ITOM Mobile Agent, see Log in to ITOM Mobile Agent. You must log in to receive push notifications about alerts and incidents.

Log in to ITOM Mobile Agent

Log in to your ITOM Mobile Agent instance on the ServiceNow Agent app to manage alerts, incidents, and on-call schedules on the go.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

ITOM Mobile Agent must be installed, and the ServiceNow Agent app must be downloaded on your device. See Install ITOM Mobile Agent and Download the ServiceNow Agent app for more information.

You must also have your instance address, for example `nebula-dev.service-now.com`, or a QR code for the instance. If you don't have one, contact your admin.

Role required: srm_manager, srm_responder, or srm_admin

Procedure

1. Tap the ServiceNow Agent icon ( ) on your device.
2. Enter the address of your instance manually or with a QR code.

Option	Description
Manual entry	Replace `example.service-now.com` with the address of your instance. You don't need to include `service-now.com`. For example, enter `nebula-dev` for `nebula-dev.service-now.com`.
QR code	Tap the QR code icon and scan the QR code.

3. **Optional:** In the **Nickname** field, enter an alternative name for your instance.
   Alternative names can help you quickly navigate between different instances on your device.
4. Tap **Save and log in**, enter your credentials, and then tap **Log in**.
   The app displays `Logged into <instance-name> as <username>` and the ITOM Mobile Agent **My work** tab.

> **Note:** You must be logged in to get push notifications about alerts and incidents. If you aren't, you only get email notifications.

What to do next

Visit these links to start using ITOM Mobile Agent:

- Work with alerts in ITOM mobile
- Managing incidents and major incidents on mobile
- Working with On-Call Scheduling on mobile
- Working with tasks on mobile

Configure when to send alert notifications for ITOM Mobile

Keep your teams informed by configuring when ITOM Mobile Agent sends a notification about an alert.

Before you begin

Role required: admin

About this task

By default, ITOM Mobile Agent sends a push notification and email when an alert meets these conditions:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Assigned to is empty** - The alert isn't assigned to anyone.
- **State is not Closed** - The alert state isn't closed. It's either empty, flapping, open, or reopen.
- **Severity is Critical** - The alert severity is critical. It isn't clear, empty, OK, major, minor, or warning.

You can update those conditions as needed. For example, the ITOM Mobile Agent app can send notifications for both critical and major alerts. However, if you customize the conditions, the conditions don't receive future release updates. That behavior is intentional to avoid overwriting your work.

Note: To get push notifications, you must log in to ITOM Mobile Agent. If you log out, you only get email notifications.

This procedure is for admins configuring alert push notifications and emails. To manage other notification features, see Customizing other mobile settings for ITOM Mobile Agent.

Procedure

1. Navigate to **All** > *Mobile App Builder*.
2. Search for and select **ITOM Mobile Agent**.
3. Select **Mobile notifications**.
4. In the **When to send** section, update or add to the assigned to, state, and severity conditions. For example, adding the condition **OR [Severity][is][Major]** under the condition **[Severity][is][Critical]** sends a push notification and email when the **Severity** field of an alert is Critical or Major.
5. Implement your changes by selecting **Save**.
   The page displays `Your changes have been saved.`

What to do next

To customize who receives emails about alerts, see Customize ITOM Mobile Agent email recipients. You can also customize other notification features, such as setting custom notification tones and overriding the Do Not Disturb setting for critical issues. See Customizing other mobile settings for ITOM Mobile Agent for more information.

Customize ITOM Mobile Agent email recipients

Customize who receives emails about alerts to help reduce notification overload and efficiently delegate responsibilities.

Before you begin

Role required: srm_admin or admin

About this task

By default, ITOM Mobile Agent sends emails about alerts to everyone in a team's group email. You can customize the email recipients as follows:

- Exclude managers from alert emails, reducing notification overload.
- Include individual team members on alert emails, making sure that the right people stay informed.

Procedure

1. Navigate to **All** > **Service Reliability Management** > **Teams**.
2. In the list, find the team that you want to customize.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Choose one or more options to include or exclude certain recipients from email alerts.

Option	Instruction
Exclude managers from alert emails	In the **Exclude manager** column, set **false** to **true**.
Include individual team members on alert emails	In the **Include members** column, set **false** to **true**.

4. Save your changes by selecting the check mark.
   Your updates appear in the **Teams** list, and ITOM Mobile Agent sends emails accordingly.

Customizing other mobile settings for ITOM Mobile Agent

Learn how to customize your mobile experience further. Configure actionable notifications, set custom notification tones, and override the Do Not Disturb setting for critical issues to help your teams stay informed.

Explore other settings

ITOM Mobile Agent auto-integrates with ITSM Mobile Agent, letting you manage alerts, incidents, and on-call schedules from the same mobile app. Visit the following ITSM Mobile Agent documentation to learn about customizing mobile notifications.

- Configurable actionable notifications for ITSM Mobile Agent
- Enable Push notification categories for ITSM Mobile Agent
- ITSM Mobile email notification redirection via web
- Set custom alert tones in ITSM Mobile Agent for different incident priority
- Enable Override do not disturb to receive critical alerts

Use ITOM Mobile

Work with alerts, incidents, on-call schedules, and tasks on the go.

ITOM Mobile Agent tasks

Work with alerts in ITOM mobile

Manage alerts on the go. Learn how to assign, reassign, and close alerts in ITOM Mobile Agent. You can also add work notes to alerts and create incidents from alerts.

Add work notes to an alert

Add work notes to an alert to update your team, share important information, and document your actions.

Role required: srm_manager, srm_responder, or srm_admin

1. On your device, tap the ServiceNow Agent icon ( ).
2. Tap the **Notifications** tab and select the relevant alert.
3. Tap the More options icon ( ).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Select **Add work notes** and enter your notes.
5. Save and add your notes to the alert by tapping **Submit**.

ITOM Mobile Agent displays `Success! Comments added`. The assigned team member also receives a push notification that the alert has had comments added.

To view an alert's work notes on mobile, navigate to the alert and select the **Activity** tab.

Assign an alert

Assign an alert to make sure the correct team member acts on it. You can assign an alert to yourself.

Role required: srm_manager, srm_responder, or srm_admin

1. On your device, tap the ServiceNow Agent icon ( ).
2. Tap the **Notifications** tab and select the relevant alert.
3. Tap the More options icon ( ) and select **Assign alert**.
4. Tap **Assigned to** and select a team member to assign the alert to.
5. Implement your changes by tapping **Submit**.

The assigned team member receives a push notification that the alert has been assigned to them. The alert also appears in their **My work** tab.

Close an alert

Close an alert to show that the issue is resolved or that no further action is needed.

Role required: srm_manager, srm_responder, or srm_admin

1. On your device, tap the ServiceNow Agent icon ( ).
2. Tap the **Notifications** tab and select the relevant alert.
3. Tap the **More options** icon ( ) and select **Close alert**.

ITOM Mobile Agent displays Alert closed successfully. The alert no longer appears in the assigned team member's **My work** tab.

To view closed alerts that were assigned to you, tap **My work** and then tap **See all** in the My Alerts section. The **Closed** tab shows your closed alerts.

Create an incident from an alert

Create an incident when an alert signifies a critical issue that disrupts or could disrupt your system.

When you create an incident from an alert, ITOM Mobile Agent assigns the incident to the same team but not the same responder. For example, if the alert is assigned to the Network Ops team and responder Alex, the incident is assigned to the Network Ops team.

Role required: srm_manager, srm_responder, or srm_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. On your device, tap the ServiceNow Agent icon ( ).
2. Tap the **Notifications** tab and select the relevant alert.
3. Tap the **More options** icon ( ) and select **Create incident**.

ITOM Mobile Agent displays a notification similar to `INC0013579 was created from Alert0002468`. You can view the incident in the **My team** tab. For more information about managing incidents in ITOM Mobile Agent, see Managing incidents and major incidents on mobile.

Note: If you don't see the option to create an incident from an alert, make sure that the alert is assigned to a team. To assign the alert to a team, see Assign alerts in Express List. In the Express List, teams are also known as assignment groups.

Reassign an alert

Reassign an alert to make sure the correct team member or an available team member acts on it. You can reassign an alert to yourself.

Role required: srm_manager, srm_responder, or srm_admin

1. On your device, tap the ServiceNow Agent icon ( ).
2. Tap the **Notifications** tab and select the relevant alert.
3. Tap the **More options** icon ( ) and select **Reassign alert**.
4. Tap **Assigned to** and select a team member to reassign the alert to.
5. Implement your changes by tapping **Submit**.

ITOM Mobile Agent displays Alert reassigned successfully. The assigned team member receives a push notification and the alert appears in their **My work** tab.

Managing incidents and major incidents on mobile

Manage and resolve incidents on the go using ITOM Mobile Agent. Learn how to create, process, and collaborate on incidents on your mobile device. You can also track and collaborate on major incidents.

ITOM Mobile Agent auto-integrates with ITSM Mobile Agent, letting you manage alerts, incidents, and on-call schedules from the same mobile app. Refer to the following ITSM Mobile Agent documentation to learn about managing incidents on your mobile device.

Manage an incident

- Collaboration with incidents
- Collaboration with incident tasks
- Create an incident from app landing page
- Generating task summarization for incident records
- My team dashboard
- Performance metrics
- Process a breached incident

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Process an incident at risk
- Unassigned incidents

Manage a major incident

- Active major incidents
- Collaboration with major incidents
- Major incident dashboard
- Review and work on a major incident candidate

Working with On-Call Scheduling on mobile

Manage on-call schedules on the go using ITOM Mobile Agent. As a manager, learn how to manage shifts, time-off requests, and scheduling gaps. As a responder, learn how to request time off, view who is on call, and see upcoming shifts.

ITOM Mobile Agent auto-integrates with ITSM Mobile Agent, letting you manage alerts, incidents, and on-call schedules from the same mobile app. Refer to the following ITSM Mobile Agent documentation to learn about managing on-call schedules on your mobile device.

On-Call Scheduling for managers

- Manage gaps and conflicts with On-Call Scheduling
- Manage time-off requests with On-Call Scheduling
- Manage upcoming shifts with On-Call Scheduling

On-Call Scheduling for responders

- Request time off with On-Call Scheduling
- View on-call group members
- View upcoming shifts with On-Call Scheduling

Working with tasks on mobile

Manage change and catalog tasks on your mobile device using ITOM Mobile Agent. Learn how to respond to issues quickly by viewing, closing, and adding comments to tasks on the go.

ITOM Mobile Agent auto-integrates with ITSM Mobile Agent, letting you manage alerts, incidents, and on-call schedules from the same mobile app. Refer to the following ITSM Mobile Agent documentation to learn about managing tasks on your mobile device.

Manage change tasks

- Change tasks
- Overdue changes

Manage catalog tasks

View catalog tasks

ITOM Mobile Agent reference

Learn about ITOM Mobile Agent settings. You can view the app version, give feedback, and adjust your preferences.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Explore reference documentation

ITOM Mobile Agent auto-integrates with ITSM Mobile Agent, letting you manage alerts, incidents, and on-call schedules from the same mobile app. Visit the following ITSM Mobile Agent documentation to learn more about your mobile app settings.

Settings

Synthetic monitoring

The ServiceNow® synthetic monitoring application empowers organizations to proactively manage and enhance the performance and availability of critical services. By simulating user transactions on API endpoints, this solution identifies performance bottlenecks, helps ensure up-time, and optimizes user experiences.

Get started

Explore Learn about synthetic monitoring	Configure Add users and create and manage a synthetic monitor
Identify Use synthetic monitoring to identify performance issues	Reference Find reference information about synthetic monitoring

Exploring synthetic monitoring

Learn how synthetic monitoring provides proactive, automated testing of service endpoints. By simulating user interactions, it can identify bugs, performance issues, and outages before they impact real users.

Synthetic monitoring overview

Synthetic monitors continually call HTTP endpoints in the Configuration Management Database (CMDB) to test availability, response times, or the presence of a defined string in the response body. When a test fails the configured criteria, it can alert you to endpoint issues before your customers do.

For example, you might create a monitor that tests the *GET* HTTP endpoint of a service. You could configure that monitor to check for a status code of 200 and a response time of under 500

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ms. If any code other than 200 is returned or the request takes longer than 500 ms, the test fails. An alert can then be generated and sent through Event Management.

The synthetic monitoring tests can be run by an Agent Client Collector (ACC) deployed in your network, on a MID Server, or on the ServiceNow platform.

Note: The HTTP endpoints must exist in the HTTP Endpoints [cmdb_ci_endpoint_http] table. Endpoints tested from the ServiceNow platform must be publicly accessible.

Monitor creation options

You can create synthetic monitors in the following ways:

Manual monitor creation

Create monitors for any HTTP endpoint by manually configuring the endpoint details. See Create and edit a synthetic monitor.

Monitors for discovered APIs

Create monitors for APIs discovered through API Insights. Each discovered API is represented in the Configuration Management Database (CMDB) as an API component CI, and you create the monitor directly against that existing CI. See Create a synthetic monitor for a discovered API.

Monitors from incidents

Create monitors directly from incident records to prevent future occurrences. See Create synthetic monitors from incidents.

Synthetic monitoring users

Synthetic monitoring users

User	Description
Synthetic monitoring administrator (Admin role)	Assigns roles, creates and edits monitors, and views monitor tests and their results.
Service owner (Editor role)	Creates and edits monitors and can also view monitor tests and their results.
Network operation center (NOC) operator (Viewer role)	Monitors tests and their results.

Synthetic monitoring customer-hosted workflow

Synthetic monitoring workflow

Admins, service owners, and NOC operators use synthetic monitoring in the following way:

1. As a service owner, you determine which endpoints to monitor based on business criticality.
2. As an admin or service owner, you configure locations to run the monitor from.

Note: The ServiceNow hosted location is available by default and doesn't need to be configured. However, endpoints tested by that location must be publicly available. You can also run the monitor from an ACC Proxy cluster or a MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. As an admin or service owner, you configure the monitor and then after it runs once, review the results of the tests. If the results aren't successful, you troubleshoot the configuration. You can also create an alert that fires when a test fails.
4. Synthetic monitoring runs tests based on the monitor's configuration and reports the results of each test.
5. As a service owner or operator, you monitor the test results and also respond to any generated alerts.

Synthetic monitoring benefits

Synthetic monitoring benefits

Benefit	Feature	Users
View aggregate monitor information. View the synthetic monitoring home page where you can: • See an aggregate view of all monitors • View the status of all monitors. • Drill down into a monitor's details	View aggregate information about the monitors.	Operators, service owners
Visualize synthetic test results. View the details page for a monitor where you can see: • The status and configuration of the monitor. • Charts that display test success and response time. • A historical log of synthetic test runs, including details and the response body. • A listing of alerts associated with the monitor.	View a monitor and its tests.	Operators, service owners
Get real-time notifications for outages before they impact users.	Optionally configure alerts when tests don't succeed.	Operators, service owners
Share insights with stakeholders.	View aggregate monitor information.	Operators, service owners
Update monitors to match your business needs.	Edit existing monitors, including deactivation.	Service owner
Embed monitor results in a Service Observability dashboard.	Use synthetic monitoring with Service Observability	Service owner

What to explore next

To learn more about configuring and using synthetic monitoring, see:

- Configuring synthetic monitoring
- Identifying system issues with synthetic monitoring

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Synthetic monitoring reference
- Use synthetic monitoring with Service Observability

API discovery and synthetic monitoring integration

Synthetic Monitoring integrates with API Insights to enable proactive monitoring of discovered APIs.

How the integration works

API Insights discovers APIs across your estate and represents each one in the Configuration Management Database (CMDB) as an API component CI. You create a synthetic monitor directly against that existing API component CI — Synthetic Monitoring does not create a new CI.

The integration eliminates manually configuring endpoint details, as the API component CI already captures the endpoint URL and HTTP method. If the API is related to an application service in the CMDB, that relationship is used to populate the monitor's service automatically. Service relationships are optional and may not be present for every discovered API.

Benefits

- The API component CI already carries the endpoint URL and HTTP method, so you don't configure them manually
- Proactive monitoring identifies issues before users are impacted
- Continuous validation of API availability and performance
- Integration with existing CMDB relationships and service mapping
- Consistent monitoring across all discovered APIs

Workflow

The following workflow describes how API discovery integrates with synthetic monitoring:

1. API Insights discovers an API endpoint in your environment and represents it in the CMDB as an API component CI.
2. You create a synthetic monitor directly against that existing API component CI. See Create a synthetic monitor for a discovered API.
3. The monitor continuously tests the API at the specified frequency.
4. Test results are displayed on the monitor details page.
5. Alerts are generated when tests fail based on your configured criteria.

Discovery sources

Synthetic monitors can be created for APIs discovered through the following sources:

API Insights

APIs discovered and cataloged through the API Insights application. The application provides visibility into API usage, performance, and dependencies.

Multi-service monitoring for shared endpoints

A single HTTP endpoint or discovered API is tagged to more than one application service so every team that depends on it is notified when it fails.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Why tag more than one service

In earlier releases, an HTTP endpoint check or a discovered API check could relate to only one application service. If a shared endpoint served several teams, each team had to create its own check against the same endpoint. This was necessary to see that endpoint's health reflected against their own service.

You can now tag a single HTTP endpoint or discovered API to multiple application services. One check now represents every team that depends on that shared endpoint. The underlying CMDB relationships between the endpoint and each application service are preserved rather than replaced when a new service is added.

How it works

When you select more than one application service for an HTTP endpoint or discovered API, synthetic monitoring creates a Connects to::Connected by relationship between the endpoint and each service you select. Adding a new service to an endpoint does not remove the relationships to services that were already tagged.

The endpoint's related services list shows every application service tagged to it, with duplicate entries removed, so each team can see which other services share the same endpoint.

When a shared endpoint fails a test, the resulting event is evaluated against every application service tagged to that endpoint, and each of those services can show the impact on the Service Health dashboard.

Note: Alert routing to each team individually is not part of this release. All teams tagged to a shared endpoint currently see the same alert configuration for the check.

otherprops="store-2026-09">How the Service Health dashboard resolves impacted services depends on how the endpoint was added to the CMDB:

Endpoints created directly in synthetic monitoring

Impacted services display on the Service Health dashboard through an automatic fallback mechanism.

Important: When you upgrade to a higher platform version, manually created entries must be repopulated using upgrade scripts. Contact ServiceNow support to request these scripts.

Endpoints discovered and mapped through Discovery or Service Mapping

Impacted services display on the Service Health dashboard because they're resolved through the platform's Discovery/Service Mapping model.

Where multi-service tagging applies

- HTTP endpoint checks. See Create and edit a synthetic monitor.
- API Insights discovered API checks. See Create a synthetic monitor for a discovered API.

Configuring synthetic monitoring

Plan and configure your implementation of synthetic monitoring.

Configuration overview

- Install synthetic monitoring
- Create synthetic monitoring locations

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Create and edit a synthetic monitor
- Create synthetic monitors from incidents
- Create a synthetic monitor for a discovered API
- Manage synthetic monitors

Install synthetic monitoring

You can install the synthetic monitoring application (`com.snc.uib.sow_synthetics`) with the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the synthetic monitoring application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Confirm that the Configuration items (CI) that you want to monitor exist in the CMDB HTTP Endpoints [cmdb_ci_endpoint_http] table.
- Confirm that configured credentials exist in the Credentials [discovery_credentials] table if your endpoint requires authentication.

> **Note:** Synthetic monitoring supports API key credentials, basic authentication credentials, and OAuth.

Role required: sn_sow_synthetics.synthetics_admin

About this task

The following items are installed with synthetic monitoring:

- Plugins
- Roles
- Tables

For more information, see Components installed with synthetic monitoring.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the synthetic monitoring application (com.snc.uib.sow_synthetics) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you can't find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you are displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

What to do next

Create synthetic monitoring locations

Create synthetic monitoring locations

Create a synthetic monitoring location to run a synthetic monitor. If you plan to run monitors from a ServiceNow hosted location, you can skip this procedure.

Before you begin

You need at least one of the following:

- A MID Server
- One or more Agent Client Collector (ACC) agent that monitors the health and performance of the HTTP endpoints to be monitored using synthetic monitoring. If you plan to run tests only from the ServiceNow hosted location, you don't need to set up a location.

Note:

- If your synthetic monitor is testing a private HTTP endpoint, your location must only include ACC agents that are within the same network boundaries as the HTTP endpoint.
- If your monitor is testing an HTTP endpoint from different datacenters, cloud providers, or geographic regions, a different location must be used for each.

Role required: sn_sow_synthetics.synthetics_admin

About this task

Synthetic monitors can be hosted from:

- A MID Server
- A cluster of Agent Client Collector agents
- A ServiceNow hosted location, as long as all monitored endpoints are publicly available. You don't need to configure a local instance as a location. It's available by default.

You can create as many locations as needed, but you must have at least one location to create a synthetic monitor.

Procedure

1. Navigate to **All** > *Service Operations Workspace* > **Configurations**.
2. Navigate to **Synthetic monitoring setup** > **Manage Locations**.
3. Select **Create a new location**.
4. Expand the Location details card and enter a unique name and an optional description for this location.
   The description should clearly state where the servers or agents in this location are running. For example, `US-1-east-cluster`.
5. In the **Location type** field, select either **MID Server** or **Agent Client Collector**, depending on where you want the monitor to be hosted.
6. Expand the **Select** card and select one or more MID Servers or ACC agents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Select **Save**.

> **Note:** You can delete locations from the Synthetic Locations menu, which you can navigate to by selecting **All** > **Synthetic Monitoring** > **Synthetic Locations**. However, you must first remove any monitors running on that location, either by editing the monitors to use a different location or deleting the monitors. See Create and edit a synthetic monitor for more information.

What to do next

Create and edit a synthetic monitor

Create and edit a synthetic monitor

Create or edit a synthetic monitor to test the availability and performance of your HTTP endpoints before your users discover issues.

Before you begin

- A configuration item (CI) for the endpoint you want to test.
- If you're testing private endpoints or want to run the monitors from your environment, one or more locations must be created to host the monitor. To create a location, see Create synthetic monitoring locations.

> **Note:** If you're monitoring an API discovered through API Insights or discovery patterns, see Create a synthetic monitor for a discovered API.

> **Note:** For information about how to bulk create synthetic monitors, see the Synthetic Monitoring Developer Guide.

Role required: sn_sow_synthetics.synthetics_editor or sn_sow_synthetics.synthetics_admin

1.

Procedure

1. Navigate to **All** > *Service Operations Workspace* and select the synthetic monitoring icon ( ).
2. On the Overview page, either create a synthetic monitor or edit an existing one.
   - To create a synthetic monitor, select **New**.
   - To edit an existing one, select the link for the monitor you want to edit and then select the **Details** tab.
3. On the form, fill in the fields.

   a. In the Monitor details section, enter a unique name for the monitor and a description.

   b. In the HTTP settings section, fill in the fields.

HTTP settings section

Field	Description
HTTP method	HTTP method that you want to monitor for the endpoint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: HEAD and OPTIONS methods aren't supported from a MID Server.
HTTP endpoint CI	CI for the endpoint that you want the monitor to test. If the endpoint doesn't exist, choose **Manage HTTP endpoints in CMDB** to create a CI.
Related service CI	Application service CI or CIs related to the HTTP endpoint. You can select one or more application services. Saving the monitor replaces the endpoint's application-service relationships with exactly what's selected here — deselecting a service removes its CI relationship..
Support group	Group who is notified and assigned alerts when this monitor fails. **Note:** To view this field in the HTTP Endpoints table, use the Synthetic Monitoring view.
Query parameters	String of query parameters to include in the request URL. **Note:** The string shouldn't include the leading question mark. For example, use `id=xyz&customer=myco` instead of `?id=xyz&customer=myco`.
Credentials	Credential needed if the endpoint requires authentication. OAuth, Basic Auth, and API keys are supported.
Headers	Key-value pairs for the HTTP request headers. You can add multiple pairs.
Body	Text that is used if the API expects a message body, such as a JSON snippet.

c. In the Locations section, select one or more locations.

Choose a MID Server or Agent Client Collector location, or to run this monitor's test from your ServiceNow instance, choose ServiceNow hosted location.

Note: When running tests from your instance, only six tests can be run every minute for performance reasons.

Choose **Create new location** to create a location.

For more information on how to create a location, see Create synthetic monitoring locations.

d. In the Assertion section, define a successful test by selecting one or more criteria, an operator, and a value.

Note: Multiple criteria are joined with an `AND` phrase, so the test must pass all criteria to be successful.

Criteria type	Description of success
Status code	Endpoint returns the given status code.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Criteria type	Description of success
Response time (in ms)	Endpoint response time is less than the given value.
Response body	Endpoint sends a string value in the response body that matches the selected operator and given value.

e. In the Frequency section, enter the number of minutes between each test run.

f. **Optional:** To have an alert sent through Event Management when a test fails:

- In the Alert settings section, activate the toggle switch.
- Select an alert severity for a test failure.
- Add tags to the alert. For more information about using tags in alerts, see Tag cluster alert grouping.

4. Select **Save**.

Result

The Overview page for the synthetic monitor displays the results of the tests that the monitor runs. See Identifying system issues with synthetic monitoring for more information.

Create synthetic monitors from incidents

Create synthetic monitors from incidents to proactively track endpoint failures and reduce repeat occurrences.

Before you begin

Role required: sn_sow_synthetics.synthetics_admin

- Recommended Action for ITSM and Synthetic Monitoring application is installed on your ServiceNow instance
- Incident should have a valid configuration item (CI) and a service associated with it
- At least one execution location (self hosted MID Server) is available and configured

About this task

This feature helps operations teams quickly assess and respond to incidents by automatically correlating incidents with the impacted business service and related Configuration Items (CIs). It provides instant visibility into all affected endpoints and clearly highlights monitoring coverage, enabling teams to understand impact, identify gaps, and take action. Administrators currently need to use the Synthetic monitoring application to configure monitors. The application now recommends and enables synthetic monitor creation directly from the Incidents page, helping improve incident actionability and strengthening service resilience.

When incident analysis shows failing endpoints were not monitored, synthetic admin users access the incident and set up synthetic monitors to prevent the occurrence in the future.

While creating these synthetic monitors from a post-incident recommended action, the system automatically discovers HTTP endpoints associated with the incident's affected services. Users can select from the discovered endpoints to create monitors without manually entering URLs. The system automatically selects the mid server location based on the endpoints as part of the recommendations from incidents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > *Service Operations Workspace*, and to the incidents page.
2. Select the incident that has a valid configuration item (CI) and a service associated with the incident.
3. On the incident details page, locate the Recommendations panel.
   If Synthetic monitoring is available and there are no monitors for the affected endpoints, the **Add monitors** action card is displayed.

   Note:

   You must have either the administrator role or admin privileges to install the applications. If you don't, contact your system administrator to install the application.

4. Select **Add monitors** to open the monitor creation form.
5. Select the endpoints for which you want to create synthetic monitors.

   Synthetic Monitoring automatically selects an available MID server as the execution location when creating monitors. To configure a default MID server for auto-selection, set the following system property:

   `sn_sow_synthetics.default_mid_location`

   Note: Default values are configured for response code, frequency, and trigger status for an alert on the incident.

6. Select **Create monitor**.
   The following template is used to create synthetic monitors from Recommendations.
   ```
   name: autogenerated type: "single_api", interval: 5,
   should_create_alert: true, alert_severity: 1, method: "GET",
   valid_http_code: "200", valid_http_code_type: "equals",
   locations: Randomly selected from available Synthetics locations
   of type Self hosted MID
   ```

   Note: If you select the same endpoints and click **Create monitor** again, the application prevents duplicate monitors from being created.

7. Select the link on the success message to view the monitors created in Synthetic monitoring landing page.

Create a synthetic monitor for a discovered API

Create a synthetic monitor to test the availability of APIs discovered through API Insights.

Before you begin

- API Insights
- CMDB CI Class Models. For more information, see API extension classes.
- An existing API component (cmdb_ci_api_component) record in the CMDB for the API Insights to discover APIs.
- One or more locations must be created to host the monitor to test private endpoints or run the monitors from your environment. To create a location, see Create synthetic monitoring locations.

Role required: sn_sow_synthetics.synthetics_editor or sn_sow_synthetics.synthetics_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

When API Insights discovers APIs through Service Graph Connectors, you can proactively monitor their availability and performance by creating synthetic monitors. Synthetic monitors identify issues before they impact users and verify continuous monitoring of critical API endpoints.

For more information about how API discovery integrates with synthetic monitoring, see API discovery and synthetic monitoring integration.

Procedure

1. Navigate to **All** > *Service Operations Workspace* and select the synthetic monitoring icon ( ).
2. On the Overview page, create a synthetic monitor by selecting **New**.
3. In the Monitor details section, enter a name and description for the monitor.

   Use a descriptive name that identifies the discovered API being monitored, such as the API name or endpoint path.
4. In the **Endpoint settings** section, configure how the monitor interacts with the discovered API.

   a. For **Endpoint type**, select **Discovered API**.

   b. In the **API** field, select the discovered API that you want to monitor.

   Each option in the list identifies a discovered API by its HTTP method, path, and URL. The method is part of the selected API component and does not require separate configuration.

   Note: The **API** field lists only API components that meet the following criteria:
   - The endpoint's operational status is **Operational**.
   - The URL starts with `http`.
   - The URL does not contain placeholder or template characters (for example, `/users/{id}`).

   If a known API does not appear in the list, verify that it meets these criteria.

   To browse the available API component records, select the info icon beside API, and then select **View API components**.

   c. Verify the service association.

   If the discovered API is related to one or more application services in the CMDB, those services and their support groups are populated automatically and shown as read-only. If no service relationship exists, these fields are hidden and the monitor is created without a service association.

   d. **Optional:** To configure optional request details, expand **Advanced settings**.

   (Optional) Under **Advanced settings**, you can configure the following optional details:

   - **Query parameters**: Query parameters to append to the request URL.

     Note: Do not include the leading question mark. For example, use `id=xyz&customer=myco` instead of `?id=xyz&customer=myco`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Credentials**: The credential to use if the discovered API requires authentication. OAuth, Basic Auth, and API keys are supported.
- Select **Add headers and body** to add custom request headers and, if the API expects a message body, the request body.

  Note: For example, enter a JSON snippet for REST APIs that require a request body.

5. In the Locations section, select one or more locations where the monitor tests run.

   Select a MID Server or Agent Client Collector location, or to run this monitor's test from your ServiceNow instance, select ServiceNow hosted location.

   Note: When running tests from your instance, only six tests can be run every minute for performance reasons.

   To create a location, select **Create new location**. For more information, see Create synthetic monitoring locations.

6. In the Assertion section, define success criteria for the API tests.

   Select one or more criteria, a condition for the criteria, and a value to determine when a test passes.

   Note: Multiple criteria use an AND phrase, so the test must pass all criteria to be successful.

   Common assertions for discovered APIs include:

   - **Status code**: Verify the API returns the expected HTTP status code (for example, 200 for successful requests).
   - **Response time**: Verify the API responds within an acceptable time threshold.
   - **Response body**: Validate that the API returns expected data in the response body.

7. In the Frequency section, enter the number of minutes between each test run.

   Consider the criticality of the discovered API when setting the frequency. Critical APIs may require more frequent monitoring.

8. **Optional:** To receive alerts when the discovered API monitor fails, configure alert settings.
   - In the Alert settings section, activate the toggle switch.
   - Select an alert severity for test failures.
   - Add tags to the alert to help with alert management and routing. For more information about using tags in alerts, see Tag cluster alert grouping.
9. Select **Save**.

Result

The synthetic monitor begins testing the discovered API endpoint at the specified frequency. The Overview page displays test results, including availability, response time, and assertion status. See Identifying system issues with synthetic monitoring for more information about viewing and analyzing test results.

What to do next

After creating the monitor, verify that tests run successfully and adjust assertions or frequency as needed based on the API's performance characteristics and business requirements.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Manage synthetic monitors

Edit or delete synthetic monitors after creating them.

Before you begin
Role required: sn_sow_synthetics.synthetics_editor

Procedure
1. Navigate to **All** > *Service Operations Workspace* and select the synthetic monitoring icon ( ).
2. Follow these instructions to edit a monitor.

Option	Description
Edit multiple monitor configurations	From the synthetic monitoring landing page, select one or more monitors and select **Edit**. See Create and edit a synthetic monitor. **Note:** Current values aren't shown in the fields. If you don't edit a field, the original values remain. If you want to edit properties other than cre dentials, location, or alert settings, you must edit a single monitor at a time from the **De tails** tab.
Edit a single monitor configuration	From the synthetic monitoring landing page, select a monitor to open it and then select the **Details** tab to edit. See Create and edit a syn thetic monitor for more information.
Delete monitors	From the Synthetic monitoring landing page, choose one or more monitors to delete and select **Delete**.

Enable ECC queue cleanup for synthetic monitoring

Grant the ECC Queue auto-purge job permission to remove old synthetic monitoring messages from the ECC queue.

Before you begin
Role required: admin

About this task
The ECC Queue auto-purge job requires **Can Delete** application access on the ECC Queue [ecc_queue] table. This access is not enabled by default, so the job can evaluate messages against its purge categories but can't remove any of them until an admin grants this access.

Consider enabling this access as a performance-tuning step if you observe contention or delays in responses from synthetic monitoring checks, since an unpruned queue can contribute to that behavior.

Procedure
1. Navigate to **System Definition** > **Tables**.
2. Open the ECC Queue [ecc_queue] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the **Application Access** related list, select **Can Delete**.
4. Save the table record.

Result
On its next scheduled run, the ECC Queue auto-purge job removes and flags messages according to its purge categories.

What to do next
Note: Enabling this access is a manual step in this release. Automating it is planned for a future release, so this procedure may be removed once that work ships.

Identifying system issues with synthetic monitoring

Tests run by synthetic monitors enable service owners and operators to view service endpoint performance at scale.

Overview of using synthetic monitoring

As a service owner, you can use synthetic monitoring to monitor your service's endpoints, verifying that they're available and performing as expected. You can be notified when synthetic tests fail, enabling you to mitigate issues quickly. You can see trends in HTTP API success rates and response times.

As an operator, you can use synthetic monitoring as part of triaging issues. When you learn a service is reported to have issues, you can view test results for that service's endpoints. If tests are failing or response times are slow, that might be something to investigate.

View aggregate information about the monitors

The synthetic monitoring landing page shows an overview of all created monitors, including inactive monitors.

From here, you can see the status for all your monitors. Selecting a card at the top filters the list of monitors. For example, you can view only the monitors that have failed or that are in an unknown state.

By default, the list of monitors is sorted by the timestamp in the **Updated** column. You can select a different column header to sort by that category.

Selecting a monitor lets you view details about the monitor's tests, including details, configuration, and associated alerts.

View a monitor and its tests

To view a monitor's details and the results of its tests, select a monitor from the synthetic monitoring landing page. The Monitor details page provides key information needed to understand how the monitor is performing.

Use the information in the header to understand the basic health of the monitor, including its status and when it last ran a test.

View individual tests

The Metrics card displays two charts to help you understand the health of each test the monitor has run. The **Failed tests** chart displays each test with a value of 0 when it was successful and a value of `1` when it failed. Hover over a point on the chart to view further details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Response time** chart shows the amount of time in milliseconds that it took to receive a response from the endpoint. Hover over a point on the chart to view details.

Note: Test results are retained for 21 days.

The **Monitor result history** table lists each test. By default, the table is sorted from newest to oldest by the timestamp. You can sort by any column to help find issues. For example, you can sort by **Result** to see all failures together or you can sort by **Response time** to view the tests with the highest latency.

If a monitor was configured to generate an alert when it fails, a link to the alert displays in the **Alert** column. Select an alert link to view its details.

When you see a failure or high latency, select a test to view details, including the response body, which might help explain any issues.

Response body in the details modal

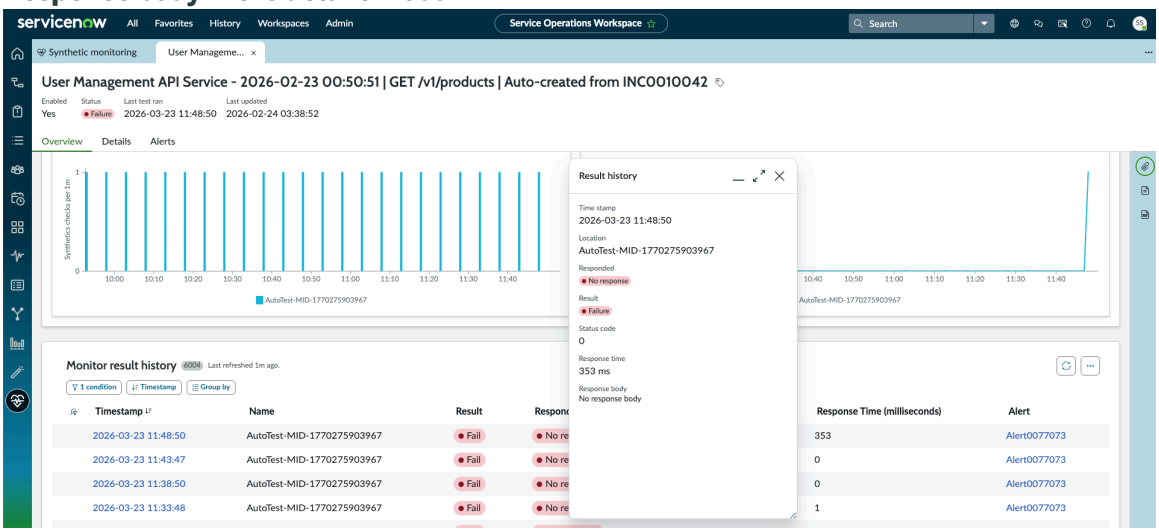

View aggregate monitor information

View information about all your synthetic monitors.

Before you begin

Role required: sn_sow_synthetics.synthetics_viewer, sn_sow_synthetics.synthetics_editor, sn_sow_synthetics.synthetics_admin

Procedure

1. Navigate to **All** > *Service Operations Workspace* and select the synthetic monitoring icon ( ).
2. To filter the list of monitors, select a card at the top of the page.
   For example, to view only monitors that have failed tests, select the **Failed** card.
3. To view details about each test run by the monitor, select a monitor name.
   The monitor's Overview tab opens. For more information about this page, see View a monitor's test results.
4. To export the list of monitors, select **Export**.
   You can export to the following file formats:
   - XLSX
   - CSV

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- JSON
- PDF

View a monitor's test results

View the results of tests run by a monitor, along with configuration details, and related configuration items (CIs).

Before you begin

Role required: sn_sow_synthetics.synthetics_viewer, sn_sow_synthetics.synthetics_editor, sn_sow_synthetics.synthetics_admin

Procedure

1. Navigate to **All** > *Service Operations Workspace* and select the synthetic monitoring icon ( ).
2. Select a monitor from the list of all monitors.
3. View the overall health of the monitor in the header, which shows its status, and when it was last run and updated.
4. View health metrics for the monitor from the **Metrics** card.
   The **Failed tests** chart displays each test with a value of `0` when successful and a value of `1` when it failed. The **Response time** chart shows the amount of time in milliseconds that it took to receive a response from the endpoint. Hover over a point on the charts to view details.
   - To change the time period, use the drop-down menu.
   - To refresh the chars, use the context menu.
5. View the status of each test in the **Monitor result history** card.
   To view the details for a test, select its timestamp. For more information about the test details, see View test details.

   Note: If your monitor uses OAuth and a test fails with a 4xx error, it may be because your refresh token expired. Re-authenticate by visiting the OAuth credential page and clicking the credential link.

   To view active alerts created by the monitor, select the Alert link to open its details.

   Note: Filtering the list affects the data displayed in the metric charts.

View test details

View the results of an individual synthetic test run by a monitor.

Before you begin

Role required: sn_sow_synthetics.synthetics_viewer, sn_sow_synthetics.synthetics_editor, sn_sow_synthetics.synthetics_admin

Procedure

1. Navigate to **All** > *Service Operations Workspace* and select the synthetic monitoring icon ( ).
2. Select a monitor from the list of all monitors.
3. Select a test from the list of tests.
   The modal displays information about the test including a response body. In this example, it shows that the request was sent to a bad gateway.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If your monitor uses OAuth and a test fails with a 4xx error, it may be because your refresh token expired. Re-authenticate by visiting the OAuth credential page and clicking the credential link.

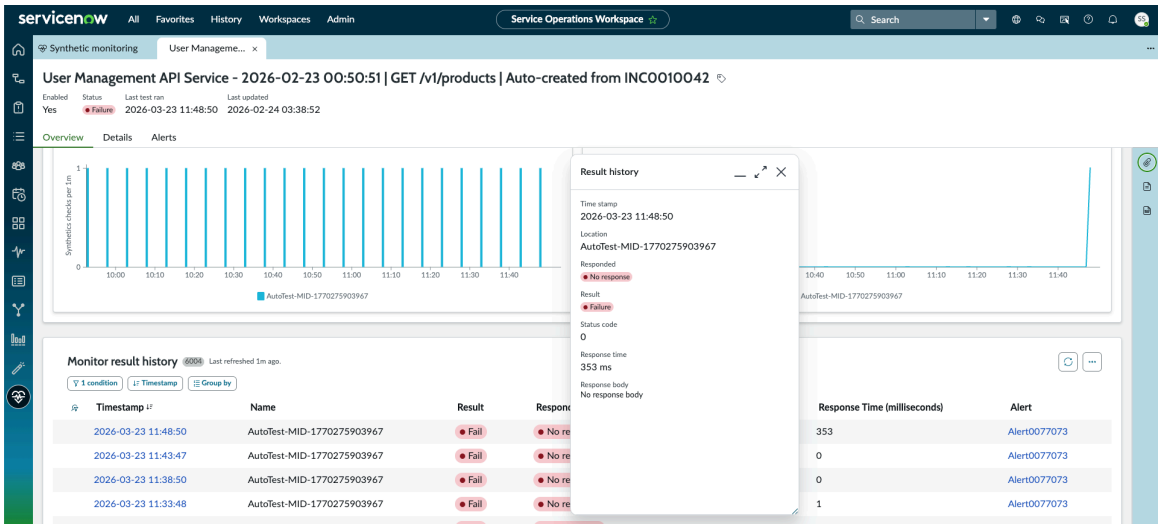

View alerts for a synthetic monitor

View all the current active alerts for your monitors.

Before you begin

Role required: sn_sow_synthetics.synthetics_viewer, sn_sow_synthetics.synthetics_editor, sn_sow_synthetics.synthetics_admin

About this task

You can view open alerts from all your monitors by selecting the **Open Alerts** card on the synthetic monitoring home page. To view a single alert, select the alert link from the resulting list.

To view alerts for a specific monitor, follow this procedure.

Procedure

1. Navigate to **All** > *`Service Operations Workspace`* and select the synthetic monitoring icon ( ).
2. Select a monitor from the list of all monitors and then select the **Alerts** tab.
   The **Alerts** tab shows the alerts for the monitor. Select an alert to view details.

Note: The alert shows the impacted service associated with the monitor. If for some reason this field isn't populated, see Impacted Service not appearing in alerts.

Synthetic monitoring reference

Reference topics provide additional information about the synthetic monitoring application, such as components installed and page details.

Components installed with synthetic monitoring

Several types of components are installed with activation of the synthetic monitoring plugin, including tables and user roles.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles installed

Role	Description	Contains roles
Synthetic monitor viewer [sn_sow_synthetics.synthetics_viewer]	Views test configurations and their results.	• cmdb_read • sn_sow.sow_user • agent_client_collector_user • sn_sow_synthetics.credential_reader
Synthetic monitor editor [sn_sow_synthetics.synthetics_editor]	Creates and edits tests. Views test configurations and their results.	• credential_admin • sn_cmdb_editor • sn_sow_synthetics.synthetics_viewer
Synthetic monitor admin [sn_sow_synthetics.synthetics_admin]	Assigns roles. Creates and edits tests. Views test configurations and their results.	• agent_client_collector_admin • sn_sow_synthetics.synthetics_editor • sn_sow_itsm_admin.sow_admin_user
Synthetic monitor credential reader [sn_sow_synthetics.credential_reader]	Used internally to read from the Credentials [discovery_credentials] table.	None
Synthetic monitor event writer [sn_sow_synthetics.synthetics_event_writer]	Used internally to communicate with the MID Server.	None

Tables installed

Table	Description
Check [sn_sow_synthetics_check]	Configuration information for synthetic checks.
Check Execution Location [sn__sow_synthetics_check_execution_location]	Execution location configuration for synthetic checks.
Single API [sn_sow_synthetics_check_single_api]	Additional configuration information and last run results for single API synthetic checks.
Single API Header [sn_sow_synthetics_check_single_api_header]	Key value pairs of requests headers for single API synthetic checks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Single API to Policy [sn_sow_synthetics_check_single_api_policy]	Single API synthetic test data and its corresponding Agent Client Collector (ACC) policies.
Single API Result [sn_sow_synthetics_check_single_api_result]	Results from each synthetic check run.
sn_sow_synthetics_recommendation_state	Per-endpoint creation state, idempotency_key, error tracking
sn_sow_synthetics_recommendation_telemetry	Aggregate workflow metrics: card_shown, endpoints_discovered, monitors_created, workflow_status
Synthetic Private Locations [sn_sow_synthetics_private_location]	Definitions of the ACC Agent Proxy Clusters available for synthetic monitoring. Extends the Proxy Agent Cluster [sn_agent_cluster] table.
sn_sow_synthetics_bulk_job	Provides async bulk job tracking.
sn_sow_synthetics_bulk_job_governance	Defines governance metadata for rate limiting and auditing for bulk jobs.

Synthetic monitoring properties

Synthetic monitoring includes the following properties.

These properties are available for synthetic monitoring.

Note: To open the System Properties [sys_properties] table, enter `sys_properties.list` in the navigation filter.

Properties for synthetic monitoring

Property	Description
sn_sow_synthetics.logger	Controls logging output. Valid values are: • debug • info • notice (default) • warning • err • crit
sn_sow_synthetics.auto_mid_selection.skip_phases	Determines the number of comma-separated phases to skip in selection cascade. Default value is empty.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties for synthetic monitoring (continued)

Property	Description
sow.recommendation.enabled	Feature flag to enable recommendation tile on incidents and by default it is set to 'false'.
sow.recommendation.max_endpoints	Determines the max endpoints to be shown in recommendation panel. The default value is 50.
sn_sow_synthetics.default_mid_location	Displays default location sys_id for auto MID selection phase 5. The default value is 'null'.
sn_sow_synthetics.auto_mid_selection.enabled	Flag to enable/disable auto MID selection and be default it is set to 'true'.
sn_sow_synthetics.auto_mid_selection.dns_cache_ttl	Determines the DNS cache time for IP affinity phase. The default value is 300 seconds.

Domain separation and synthetic monitoring

Domain separation is supported for synthetic monitoring. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

Domain separation and synthetic monitoring overview

Domain separation is present in all aspects of the synthetic monitoring application. Users belonging to a specific domain see only the data existing in their own domain.

How domain separation works in synthetic monitoring

When data is separated by domains in synthetic monitoring, users can only configure and see the monitoring data in their own domain. By default, all users and records are set to the parent domain unless the admin assigns them to a specific domain.

You define the domain-separated environment when you configure your data inputs. All relevant records and data processing in the synthetic monitoring program flow reside in the same domain

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

as the data input. The domain name of a data input appears in the Domain column that is displayed in the tables in your instance.

Using domain separation in your instance is transparent to synthetic monitoring. The application manages all aspects of the data, such as the monitor configurations and results. When a property is changed, the new value affects new sources only in the specific domain.

When setting up Agent Client Collector (ACC) for non-hosted synthetic checks, each domain that supports non-hosted synthetics requires its own ACC cluster and MID Server. The ACC cluster and MID Server must be created in a leaf domain.

Related topics

Domain separation for service providers

Synthetic monitor status definitions

Learn what each synthetic monitor test status means, so you can distinguish a genuine endpoint failure from a test that never ran.

Overview of Synthetic monitor status

Each test run by a synthetic monitor reports one of the following statuses. Use these definitions to distinguish a monitor that detected a genuine problem with the monitored endpoint from a monitor whose test could not run at all.

Status definitions

Synthetic monitor test statuses

Status	Definition
Passed	The test executed and the response met all configured assertion criteria.
Failed	The test executed but the response did not meet one or more configured assertion criteria. The monitor successfully reached the endpoint and received a response, but that response failed evaluation. Examples include an unexpected status code, a response time over the configured threshold, or missing expected content.
Unknown	The test could not run, so no result is available for evaluation. An Unknown status means the monitor could not attempt the check at all, rather than reaching the endpoint and receiving an unsuccessful response. Unknown is expected in this case, and does not by itself indicate a problem with the monitored endpoint.

Important: A Failed status indicates the monitored endpoint returned an unexpected result. An Unknown status indicates the monitor was unable to run the test at all. Treat these differently when triaging alerts: an Unknown status should prompt you to investigate the monitor's execution environment, such as its location or MID Server, rather than the health of the monitored endpoint itself.

Common causes of an Unknown status

An Unknown status can occur whenever the test cannot run to completion. Common causes include the following.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Causes of an Unknown status

Cause	Description
MID Server removed from the synthetic monitoring location	If the MID Server assigned to a monitor's location is removed from that location, the test cannot run and the status becomes Unknown. This is expected behavior.
MID Server is down or not validated	If the MID Server assigned to a monitor's location exists but is down or has not been validated, the test cannot run and the status should become Unknown, since the check never actually executed. **Note:** In earlier releases, this scenario could incorrectly report a Failed status and generate an alert, making it appear that the monitored endpoint was down. If you see a Failed status for a monitor and suspect its MID Server is down or unvalidated, check the MID Server's status before troubleshooting the monitored endpoint.

Related topics

Upgrade issues

Identifying system issues with synthetic monitoring

ECC queue auto-purge for synthetic monitoring

A scheduled job removes old Synthetic monitoring messages from the ECC queue. The job also flags stuck or orphaned messages with status as error, so that the queue doesn't grow indefinitely and stale entries aren't mistaken for pending work.

Enable the cleanup job

The ECC Queue auto-purge job requires **Can Delete** application access on the ECC Queue [ecc_queue] table. This access is not enabled by default, so the job can't remove any records until an admin grants it.

For information about enabling this access, see Enable ECC queue cleanup for synthetic monitoring.

How the job runs

The ECC Queue auto-purge job runs once daily, overnight. On each run, the job evaluates synthetic monitoring messages in the ECC Queue [ecc_queue] table against the categories in the following table. The job then either removes each matching message or flags it as error.

Important: Only messages that belong to synthetic monitoring are evaluated. Messages that other applications place in the same queue table are never removed or modified, regardless of their state or age.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Purge categories

ECC queue auto-purge categories

Category	Applies to	Condition	Result
Completed requests	Output-queue messages	State is **Processed** and the message was created more than 30 days ago.	Removed from the queue
Completed responses	Input-queue messages	State is **Processed** and the message was created more than 30 days ago.	Removed from the queue
Failed executions	Messages in either queue	State is **Error** and the message was created more than 30 days ago.	Removed from the queue
Orphaned ready	Messages in either queue	State is **Ready** and the message was created more than 1 day ago, meaning it was never picked up.	Flagged as error
Orphaned processing	Messages in either queue	State is **Processing** and the message was last updated more than 7 days ago, meaning it is stuck in progress.	Flagged as error
Disabled monitor	Messages in either queue	State is **Ready** or **Processing** and the monitor that the message belongs to has been disabled. Age isn't considered.	Flagged as error

Messages that are flagged as errored remain in the queue with a state of **Error**. They are then eligible for removal under the failed executions category once they are more than 30 days old.

When cleanup can't be completed

If the job can't remove records for reasons such as a table permission restriction, then the job writes a warning to its log instead of reporting success. Review the job's log output to confirm that each run completed the cleanup it reported.

For more information, see Troubleshoot ECC queue growth.

Troubleshoot synthetic monitors

Read these topics to understand how to fix common issues in synthetic monitoring.

Upgrade issues

What to do when monitors don't work after upgrading synthetic monitoring.

Condition

After upgrading synthetic monitoring, monitors remain in unknown state.

Cause

After upgrading synthetic monitoring, if the monitor is hosted on a MID Server, the server must be restarted to re-recognize the hosted monitors.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remedy

Restart the MID Server. See Manually start, stop, and restart a MID Server.

Related topics

Synthetic monitor status definitions

OAuth issues

Learn how to fix OAuth token refresh issues.

Condition

OAuth token doesn't refresh as expected.

Cause

Your OAuth provider and server have different refresh time periods. Refresh time periods are based on the OAuth provider. When creating an OAuth connection, you configure it to have an x-amount of time before needing reauthentication. If that configured amount of time isn't the same as your provider's, your token might require a refresh sooner or later than expected.

Remedy

Refer to your OAuth provider's documentation when configuring OAuth. See OAuth Inbound and Outbound authentication for more information.

Impacted Service not appearing in alerts

When a test fails for a monitor and you have configured an alert to trigger, the alert record should show the impacted service. There are a few reasons why this field might not be populated.

Condition

The Impacted Services field isn't populated on an alert record from a failed synthetic monitor. Following are the causes and remedies for the different reasons the field isn't populated.

Cause: Service isn't a mapped service

The service isn't a mapped service.

Remedy

This field only populates when the associated service is configured as a mapped application service in the Configuration Management Database (CMDB).

Cause: The field didn't have time to populate

Impacted Services are populated at alert creation.

Remedy

Close and reopen the alert so new mappings are applied.

Cause: Race condition encountered

If an alert triggers immediately, the platform might not have time to establish the necessary CMDB relationships before the alert attempts to populate impacted services.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Remedy

Try reloading the page or consider running tests from the MID Server.

Cause: Delayed population

The relationship might take up to a minute to appear, which means the field might not populate until after the initial alert creation event.

Remedy

Check back after a few minutes.

Bulk monitor import issues

Learn to fix issues that might occur when you create monitors in bulk.

Common bulk monitor import issues

Issue	Cause	Solution
cmdb_ci is required	Missing HTTP endpoint	Create endpoint in CMDB first
Location not found	Invalid location sys_id	Verify location exists
Authentication failed	Invalid credentials	Check username/password
Job stuck in "pending"	Event queue backed up	Wait or contact admin
Timeout errors	Large payload, slow processing	Increase timeout, reduce batch size

Troubleshoot ECC queue growth

If the ECC queue keeps growing or contains large numbers of stale synthetic monitoring messages, use this information to identify why the auto-purge job isn't clearing them.

The job reports success but the queue doesn't shrink

The most common cause of this condition is that the ECC Queue [ecc_queue] table doesn't have **Can Delete** application access enabled. This access isn't enabled by default, so until an admin grants it, the job can evaluate messages against the purge categories but can't actually remove any of them.

Review the job's log output for warnings from the most recent run. If a warning indicates that records couldn't be removed, confirm that the job has delete access to the ECC Queue [ecc_queue] table. To grant that access, see Enable ECC queue cleanup.

Messages remain for monitors that are disabled

When a monitor is disabled, messages that it already placed in the queue can remain there. On its next run, the auto-purge job flags queued messages in the **Ready** or **Processing** state as errored when the monitor they belong to is disabled. Because the job runs once daily, these messages can persist until the next run.

Stale entries appear to be pending work

Messages left in the **Ready** state that were never picked up, and messages stuck in the **Processing** state, can look like pending work when they are reviewed in the queue. The auto-purge job flags both as errored so that they are distinguishable from messages that are

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

genuinely waiting to be processed. For the age thresholds that apply to each state, see ECC queue auto-purge for synthetic monitoring.

Synthetic monitoring Landing Page

Read the following to understand what is displayed on the Synthetic monitoring landing page where you can view aggregate information about your synthetic monitors.

Aggregate information

The cards show aggregate information for all monitors, whether they're active or not. Select a card to filter the Monitors list.

Aggregate cards

Card title	Description
Total monitors	Total number of monitors created.
Passed	Number of total monitors that have returned the value that was set as the monitor's assertion, for example a status code of 200.
Failed	Number of total monitors that returned a value other than what was configured as the assertion.
Pending	Number of total monitors that are in the Pending state. Monitors are set to **Pending** before the first test runs.
Paused	Number of total monitors that are in the Paused state. Paused monitors are monitors that have been set to **Monitor disabled** on the Configuration page.
Unknown	Number of total monitors that aren't running due to an issue with the related Agent Client Collector (ACC) proxy.

Monitor list

By default, this list displays all monitors that have been created. You can select a card at the top of the page to filter to a list of monitors with that state. Select a monitor to view its details page.

Details page of a synthetic monitor

Read the following to understand how to use the Details page of a synthetic monitor to view the results of its tests.

Header

The header provides an overview of the monitor.

Item	Description
Monitor name	Name given to the monitor when it was created.
Enabled	Indication of whether the monitor is active.
Status	Status of the latest test run. Values can be one of the following: • Pending: The monitor is waiting for the first test to run, based on the configured frequency. • Success: The test ran and passed the assertion criteria that you set.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Item	Description
	• Unknown: The test failed to run, likely due to an issue with the related Agent Client Collector (ACC) proxy. • Failure: The test ran but didn't pass the assertion criteria. Select a timestamp link to view the response body for more information about the test.
Last test ran	Timestamp for when the latest test finished running.
Last updated	Timestamp for when the monitor was most recently edited and saved.

Overview tab

The **Overview** tab provides an aggregated view of all the tests that the monitor has run. This tab shows the following:

- Metrics card: Charts that provide performance of each test for the time period selected in the time picker.
- Monitor result history: All tests run by the monitor and its results.

Details tab

You can use the **Details** tab to create a monitor or edit an existing one. See Create and edit a synthetic monitor for more information.

Alerts tab

You can use the Alerts tab to view alerts generated by a monitor.

ITOM AIOps reference

Reference topics provide additional information about about using ITOM AIOps

Components installed with AIOps Experience

Several types of components are installed with activation of the AIOps Experience [sn_sow_aiops] application, including user role and plugins.

Note: The Application Files table lists the components that are installed with this application. For instructions on how to access this table, see Find components installed with an application.

Demo data is available for this feature.

Roles installed

Role title [name]	Description	Contains roles
evt_mgmt_user	Manages the alert lifecycle and performs basic operations, such as viewing alerts.	itil

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description	Contains roles
evt_mgmt_operator	Manages alerts, including tasks such as closing and acknowledging them.	evt_mgmt_user
evt_team_operator	Manages Event Management operations within a specific team as defined in the **Assignment Group** field. This role allows the operator to read and write alerts exclusively for their assigned team. Additionally, the operator can make configuration changes specific to their team, including updates to Alert Automation and the Integrations Launchpad. **Note:** The evt_team_operator role must be assigned to an assignment group to view and manage alerts for that group. If the role is created but not associated with any groups that have alerts assigned, the operator cannot see any alerts.	sn_sow_alert_rules.automation_create_nonglobal sn_sow_alert_rules.automation_delete_nonglobal
evt_mgmt_admin	Configures and sets up Event Management properties and rules.	evt_mgmt_user and evt_mgmt_operator

Plugins installed

Plugin	Description
Event Management [com.glideapp.itom.snac]	Provides service-oriented availability and monitor events through an intuitive console.
Performance Analytics - Premium [com.snc.pa.em]	Provides core out-of-the-box Key Performance Indicators (KPIs) to help monitor and enhance event management processes. Please note that activating this plugin on production instances may require a separate Performance Analytics license.
Performance Analytics - Content Pack - Event Management [com.snc.pa.premium]	Includes core out-of-the-box Key Performance Indicators (KPIs). Activating this plugin on production instances may require a separate Performance Analytics license.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To return to the procedure, see Install Service Operations Workspace for ITOM application.

Application service details page

Provides comprehensive information about an application service, including its name, business criticality, ownership, process and operational status, traffic-based discovery details, discovery status, and relevant comments. This enables thorough analysis of the application service's current state and usage, provides clarity on ownership and criticality, and aids in prioritizing resources and actions effectively.

Application service details page fields

Field	Description
Name	Name of the application service.
Business criticality	Importance level of the service to business operations.
Owned by	Person or team responsible for managing and overseeing the service.
Process Status	Current stage or status of the service's operational processes.
Operational status	Current stage or condition of the service's operational processes, showing whether it is functioning smoothly or encountering issues.
Traffic based discovery	Option to identify service instances based on network traffic patterns.
Discovery status	System's ability to locate and recognize the service within its network.
Comments	Additional notes or pertinent information regarding the service that may be helpful for understanding its management or usage.

To return to the procedure, see View application service details in Service Operations Workspace.

Create list fields

Create customized lists with fields tailored to your needs.

The form you complete varies depending on whether you choose to base your list on an existing one (**Start from existing**) or create a new list (**Create your own**).

Start From Existing form

Field	Description
List	Existing list that you want to modify.
List Name	Name for your list. The default name in this field is the name of the existing list in the **List** field with `- Copy` appended.
Select columns	Record fields to include in list view. The columns from the list you selected appear. You can add or remove columns to modify the list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Start From Existing form (continued)

Field	Description
Add Filters	Condition used to filter the list. By default the conditions applied to the list selected appear. You can modify this condition.

Create Your Own form

Field	Description
List Name	Name for your list.
Select Source	Source table for the records.
Select columns	The columns to display in the list. This field is populated with columns from a Workspace list view if one exists. If a Workspace list doesn't exist, the columns are populated with the Default list view of the table selected.
Add filters	Condition used to filter the list.

To return to the procedure, see Create a list in Service Operations Workspace for ITOM

Metrics View Configuration form fields

Define how metrics are displayed within Service Operations Workspace.

Metrics View Configuration form

Field	Description
Title	Name of the metric.
Active	Option to activate the metric.
Order	Numerical value indicating the order of the metric on the page.
Configuration Item Type	Configuration item type. Available options are: • cmdb_ci_app_server_tomcat • cmdb_ci_metric_extension • cmdb_ci_app_win_server
Source Metrics Types	Metric types you want the metric to monitor. Select the Lock icon ( ) to unlock the list of source metric types and select metric types to be monitored.
Chart Type	Type of chart in which the metric is presented in Service Operations Workspace, either **Line** or **Stacked Bar**.

To return to the procedure, see Configure alert metrics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integration and configuration health event fields

Populate the event fields to capture details such as event time, description, alert number, state, severity, metric name, node, message key, and error message.

Event fields

Field	Description
Time of Event	Indicates when the event occurred.
Description	Provides details about the event or incident.
Alert	Specifies the alert number that is triggered based on the event.
State	Describes the current state of the event (e.g., open, closed).
Severity	Indicates the severity level of the event (e.g., critical, warning).
Metric Name	Identifies the specific metric associated with the event.
Node	Specifies the node or device related to the event.
Message Key	Unique identifier for the event message.
Error Message	Provides additional information about any errors associated with the event.

To return to the procedure, see Review integration and configuration health.

Datadog advanced settings fields

Populate the Datadog advanced settings fields to define API paths, batching limits, host filters, sync duration, connection details, and debugging options. These values control how the pull connector retrieves and processes metrics and host data from Datadog, ensuring the connector runs with the desired scope and behavior.

Datadog advanced settings fields

Field	Description
Base API path	The specific API path for querying metrics. The default is `/api/v2/query/timeseries`.
Enable container	Enables collection of Docker container metrics (CPU, memory, network) and binds them to Docker Container CIs (cmdb_ci_docker_container). Set this to true only after your Docker Container CIs have been discovered — if enabled before the CIs exist, container metrics are still collected but remain unbound. Default: false.
Host API Path	The specific API path for retrieving the list of hosts. The default is `/api/v1/hosts`.
Host batch size	The number of hosts to retrieve in a single API call.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Datadog advanced settings fields (continued)

Field	Description
	The default is 500, and any value above 1000 is automatically capped at 1000.
Host filter	(Optional) A comma-separated list of regular expressions (regex) to filter for specific hosts to monitor. or example, ^prod-db-.* is a regex pattern that means "match any host whose name starts with prod-db-". Example Use Case: To monitor only hosts starting with app-server- and the specific host critical-db-01, you would set the hostFilter to: ^app-server-.*, ^critical-db-01$.
Initial sync in mins	Time, in minutes, that the system takes to complete the first synchronization of data with the Datadog Metrics API. For example, if the value is 15, the connector fetches metrics from the past 15 minutes. If the value is above 120, it is capped to 120. Default: 15
Metric batch size	The number of metrics to query in a single API call. The default is 50, and any value above 50 is automatically capped at 50.
Port	Port number for the Datadog server to connect to. This is typically not needed when using https.
Protocol	The protocol for the connection. Default: https
Debugging	• Debug: Provides detailed logs for troubleshooting. • Log payload: Displays raw log data. Use only for debugging, as it can quickly fill the MID Server logs.

To return to the procedure, see Configure Datadog metric pull connector.

Dynatrace advanced settings fields

Populate the pull connector advanced settings fields to optimize event retrieval and synchronization based on specific requirements such as connection preferences, event frequency, and time zone settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dynatrace advanced settings fields

Field	Description
Base API path	Root URL or endpoint used to connect to the Dynatrace Metrics API. All API requests are made relative to this base path.
connectionAliasSysID	Sys ID of the connection alias record associated with the Dynatrace connection.
debug	• Debug: Provides detailed logs for troubleshooting. • Log payload: Displays raw log data. Use only for debugging, as it can quickly fill the MID Server logs.
Initial sync in mins	Time, in minutes, that the system takes to complete the first synchronization of data with the Dynatrace Metrics API. For example, if the value is 15, the connector fetches metrics from the past 15 minutes. Default: 15
isGrailApiEnabled	Enables Grail-based metric collection for the Dynatrace connection. When enabled, this allows Grail API calls to fetch metrics data and populates the Connector Instance Metrics [sn_em_connector_instance_metrics] table with Grail-related metric IDs. When disabled, the table is populated with classic metric IDs. Default: false.
logPayload	Enables logging of the request and response payloads for the Dynatrace connection, useful for debugging.
Max metrics per query	Maximum number of metrics that can be requested in a single query to the Dynatrace Metrics API. The value cannot exceed 32, the maximum limit allowed by Dynatrace. Requests exceeding this limit must be split into multiple queries. Default: 32
Page Size	Number of records returned in a single response when querying the Dynatrace Metrics API. If the total results exceed this number, multiple pages are returned. Default: 500
Port	Port number used for connecting to the Dynatrace server.
Request per minute	Maximum number of API requests that can be sent to the Dynatrace Metrics API within one minute.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dynatrace advanced settings fields (continued)

Field	Description
	Exceeding this limit may result in throttling or temporary blocking of requests. **Note:** Relevant only when you enable **Throttle API calls**. Default: 1000
Throttle API calls	Limits the rate of API requests to the Dynatrace Metrics API, preventing overload, request failures, or temporary blocking. Enable this option to automatically limit the rate of API requests, preventing overloading the Dynatrace Metrics API and avoiding request failures. Default: false
Protocol	Communication protocol used for retrieving events from Dynatrace server. Default: https

To return to the procedure, see Configure Dynatrace metric pull connector.

Pull connector fields

Populate the pull connector advanced settings fields to optimize event retrieval and synchronization based on specific requirements such as connection preferences, event frequency, and time zone settings.

Advanced settings fields

Field	Description
Connected MID Server	Determine whether you want to have the system choose the optimal MID Server or choose a specific MID Server yourself. Perform only if you must make modifications to the MID Server otherwise leave it empty and let the system choose the optimal MID Server.
Event collection schedule (seconds)	Enter the frequency in seconds that indicates how often the system should check for new events from the external event source. The value can't be lower than the minimum schedule property, which by default is 120 seconds.
Protocol	Defines the communication protocol used for retrieving events from external devices.
severity_by_state_type	Determines the severity level assigned to events based on their state type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced settings fields (continued)

Field	Description
time_zone	Specifies the time zone used for recording and displaying event timestamps.
offset_min	Sets the time offset for event retrieval to ensure accurate synchronization.
Port	Verify the default values for your connector type such as port details or time formats and modify them as necessary. Default values from your connector are filled-in. Specifies the port number used for communication with external devices.
max_fetch_interval_min	Limits the maximum interval between event fetch operations.
initial_sync_in_days	Sets the duration for the initial synchronization process when retrieving historical data.

To return to the procedure, see Configure an event pull connector.

Outbound webhook parameters

The parameters to be configured when sending data to other systems using an outbound webhook.

Outbound webhook parameters

Parameter	Definition
Destination URL field	The webhook endpoint (URL) where the alert is to be sent.
JSON payload data tab	The alert data, composed in JSON format, being sent to the indicated URL. Enter the details you want to be notified about in the JSON object when the automation is triggered. You can insert alert fields by selecting **Insert Field Values** or add custom fields by typing {custom_field.JSON_path}. To add additional alert fields, select **Insert alert fields**.
Payload headers	Add predefined or custom metadata headers in the **Key** and **Value** fields. A key is the name of a header field, and a value is the corresponding data assigned to that header field.
Test payload	Select to verify that the webhook is functioning correctly and that the receiving system can process the data as expected.

To return to the procedure, see Create Respond automation.

Agent Client Collector

The Agent Client Collector is an agent that is installed on infrastructure components. The Agent Client Collector is built on a Sensu framework which enables you to adopt and extend monitoring checks from the community. It executes commands on the machines it is installed on and sends

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

output data to the ServiceNow® instance via the MID Server, storing events and metrics in the relevant database.

Watch this video for an overview of how the Agent Client Collector works.

https://player.vimeo.com/video/1096666842?h=77a5520958&badge=0&autopause=0&player_id=0&app_id=58479

Preparing for Agent Client Collector implementation Follow the steps described in the planning and deployment checklists to ensure successful Agent Client Collector deployment.	Exploring Agent Client Collector Analyze CI data sent to the instance, and provide data to other Agent Client Collector applications.	Agent Client Collector server deployment Connect agents directly to the ServiceNow® cloud using a MID Server or ITOM Cloud Services (ICS).	Agent Client Collector endpoint deployment Connect agents directly to the ServiceNow® cloud using ITOM Cloud Services (ICS).
Agent Client Collector shared deployment Deploy selected Agent Client Collector components on both servers and endpoints.	Agent Client Collector Log Analytics Stream log data from Linux and Windows hosts to a ServiceNow instance.		Live CI View View live data for incident related CIs to help resolve incidents.

Exploring Agent Client Collector

The Agent Client Collector Framework (ACC-F) is a powerful solution for monitoring the performance and health of infrastructure components by using agents installed on servers and devices. It collects and sends critical system data to ServiceNow for analysis, enabling proactive management and troubleshooting of Configuration Items (CIs).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector Framework overview

The Agent Client Collector Framework enables organizations to monitor and manage the health of their infrastructure through agents installed on key systems. These agents execute predefined commands on the machines they are installed on, sending the resulting data back to the ServiceNow® instance via a dedicated MID Server. The framework enables seamless management of both the Agent Client Collector and MID Server, storing event data and performance metrics in the appropriate database. By providing insights into CI performance, ACC-F helps organizations identify and resolve issues quickly, improving operational efficiency and system reliability.

Agent Client Collector Framework workflow

The following illustration describes the layout and data flow within the Agent Client Collector Framework application.

Collecting and distributing data with ACC-F

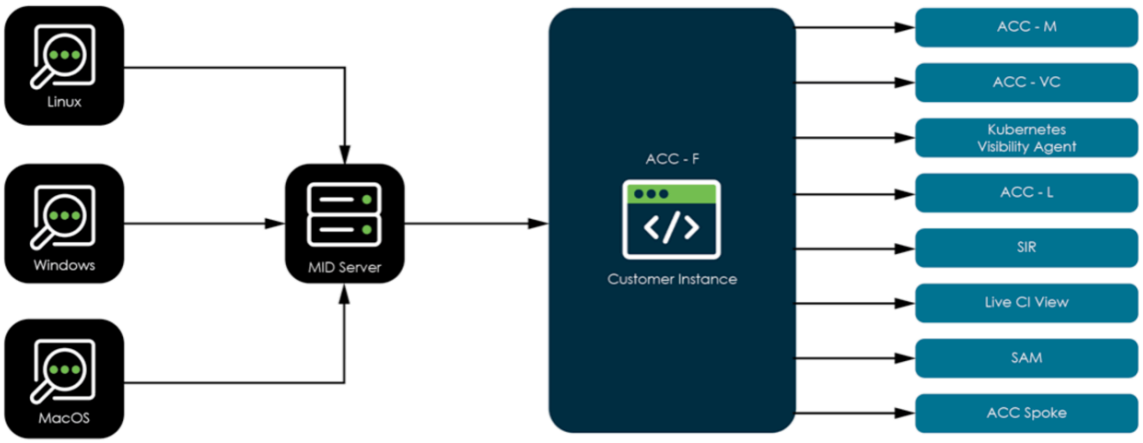

1. Agent Installation: The agent is installed on infrastructure components, such as servers, devices, and network equipment. These agents are responsible for executing system commands and gathering performance data from the host machine. ACC-F is deployed on the customer's ServiceNow® instance.
2. Data collection: The agent runs predefined scripts or queries (checks) on the infrastructure components to collect various system metrics and events. This includes performance data such as CPU usage, memory utilization, disk space, and network activity. The agent also collects error logs and system alerts.
3. MID Server communication: The agent sends the collected data to the ServiceNow instance through a dedicated MID Server. The MID Server acts as a secure communication bridge between the infrastructure and the ServiceNow platform, ensuring data is transmitted reliably and securely.
4. Data Storage: Upon receiving the data, the ServiceNow instance stores the events and performance metrics in the relevant database. This data is then associated with the respective Configuration Items (CIs) within the ServiceNow CMDB Configuration Management Database (CMDB), enabling efficient tracking and reporting.
5. Data Analysis and Reporting: The collected data is analyzed within the ServiceNow instance, where it is used for monitoring, troubleshooting, and reporting purposes. This analysis helps identify potential issues or areas of improvement, triggering alerts or actions to resolve problems proactively.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Feedback Loop: Based on the analysis, the system can trigger remediation actions, such as running corrective scripts, reconfiguring settings, or notifying responsible teams for further investigation. The feedback is looped back into the system, allowing for continuous monitoring and optimization.

Agent Client Collector Framework benefits

Agent Client Collector Framework provides data to other Agent Client Collector components.

Benefit	Feature	Users
Monitor your system's health, performance, and availability through automated collection of events and metrics, leveraging automated configurations.	Agent Client Collector Monitoring	NOC User, Event Management administrator
Track server inventory, software installations and usage continuously with non-admin access and minimal network communication.	How Agent Client Collector for Visibility Content works	CMDB/Discovery administrator
Gather detailed inventory data of devices not connected to your network or running in isolated environments (air-gapped).	Agent Client Collector Framework Air Gapped Configuration Item Management Solution	CMDB/Discovery administrator
Minimize triage time of incidents by direct access of live device details and interactions to remediate.	View live CI data with Agent Client Collector	ITSM user
Stream log data into your instance to predict problems and solve them before they happen, to minimize user impact.	Agent Client Collector Log Analytics	Agent Client Collector administrator

Agent Client Collector Framework use case

The Agent Client Collector Framework (ACC-F) use case demonstrates how a financial organization can use Agent Client Collector Framework to assist in IT asset discovery.

Use case overview

A leading financial organization needed a robust IT discovery solution to maintain compliance, secure critical systems, and manage a growing remote workforce. While the organization initially relied on agentless discovery methods — where MID Servers interrogate systems and populate the CMDB—it faced significant challenges. These limitations affected their ability to monitor secure environments and remote endpoints effectively, leading to potential security and compliance risks.

Challenges

- Restricted server access: The financial organization's security policies restricted MID Servers from initiating connections to core banking servers and sensitive financial databases. This made agentless discovery insufficient for these high-security assets.
- Remote workforce visibility: The company's growing remote workforce meant many employee laptops and mobile devices operated outside the corporate network. Agentless discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

couldn't capture data from these endpoints, creating blind spots in asset tracking and compliance monitoring.

- Regulatory compliance: Financial institutions face stringent regulations (such as PCI, DSS and SOX) requiring detailed and auditable asset data. Agentless discovery struggled to meet these requirements, especially for in-depth monitoring and policy enforcement.

Solutions

To address these challenges, the financial organization implemented an agent-based discovery approach.

- Agent deployment on critical servers: Agents were installed on core banking servers and sensitive financial databases. This allowed continuous, secure monitoring without violating access restrictions.
- Remote endpoint management: Agents deployed on remote employee laptops provided real-time data collection and monitoring, regardless of location. This ensured that all devices, even those outside the corporate network, were included in asset discovery.
- Enhanced data collection: The agents collected detailed system information and enforced security policies directly on each asset, ensuring compliance with regulatory standards and improving audit readiness.

Outcomes

- Comprehensive asset visibility: The organization achieved full visibility across its IT environment, including secure data centers and remote endpoints. This reduced blind spots and improved decision-making for IT operations and security teams.
- Improved compliance and audit readiness: Agent-based discovery provided granular, auditable data that met stringent regulatory requirements. This streamlined compliance reporting and reduced the risk of regulatory penalties.
- Enhanced security posture: By enforcing policies at the endpoint level, the organization strengthened its cybersecurity defenses, protecting sensitive financial data from potential threats.

This approach enabled the financial institution to maintain control over its IT environment, support its remote workforce effectively, and ensure compliance with industry regulations - critical factors in safeguarding sensitive financial operations and maintaining customer trust.

Agent Client Collector architecture

The Agent Client Collector is a ServiceNow agent installed on your Windows, Linux, and macOS devices to monitor your company's infrastructure and installed applications.

Agent Client Collector architecture - overview

The Agent Client Collector is built on a Sensu framework and comes installed with monitoring capabilities for servers, databases, application servers, and middleware. It also supports extending monitoring through additional checks from the Sensu community or Nagios-compatible plugins, allowing you to customize monitoring to your needs.

Monitoring with checks and policies

The agent runs checks on the host to gather relevant data, transforming it into events or metrics. These checks are defined within ServiceNow and are associated with monitoring policies. A policy is a combination of the Configuration Items (CIs) being monitored and the checks that run on those CIs. The checks are associated with policies to monitor different aspects of the system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can customize check instances - such as adjusting the frequency or specifying parameters like login credentials for databases - without affecting the original check definition. Customization of a check instance takes effect only on the check instance associated with the policy and does not change the global check definition. This customization ensures flexibility in monitoring various CIs under different scenarios.

Data collection and transmission

After installation, the agent collects information on its host and processes. The agent pushes the collected data to the ServiceNow instance through the MID Server. The MID Server serves as a bridge between the agent and the ServiceNow instance, ensuring the data is transformed and securely sent for processing.

On the instance, CIs are created for the host and applications classified from the running processes (such as Microsoft SQL Server). Once these CIs are created, the active monitoring policies associated with the CIs are downloaded to the agent's MID Server, which then pushes the policies to the agent for execution.

Data storage and use

The results from the checks are stored in the ServiceNow instance:

- CI data: CMDB for Configuration Items.
- Event data: Events table (for example, alerts triggered by threshold breaches).

This data can be leveraged for monitoring, alerting, and reporting. Integration ensures real-time visibility into the health of your infrastructure and applications, while also enabling proactive issue detection and remediation.

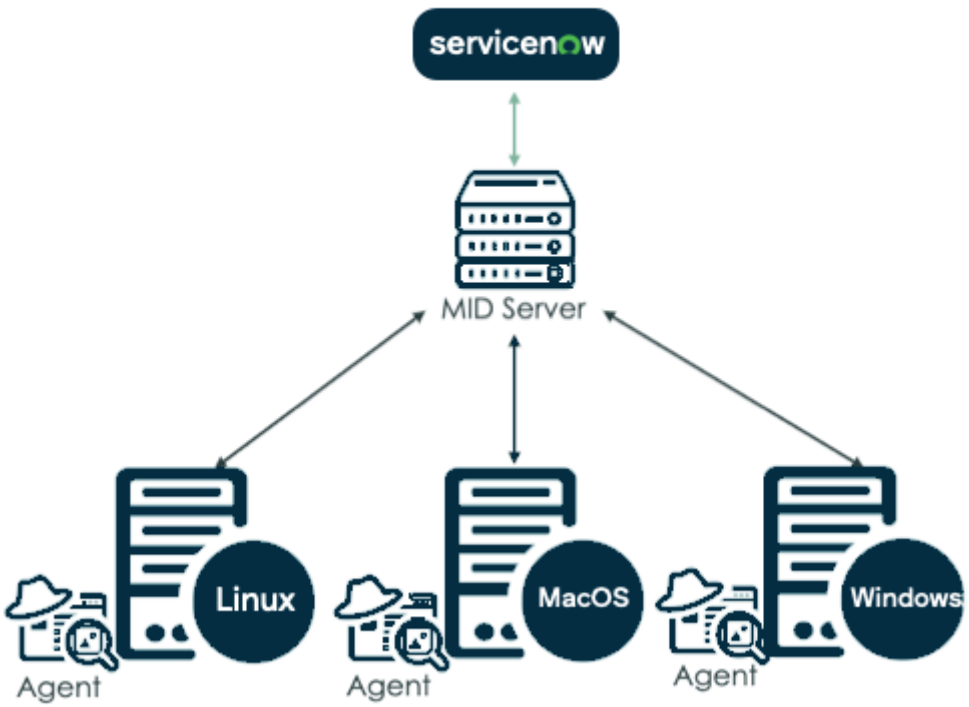

1. Define a monitoring policy in the ServiceNow instance.
2. The MID Server fetches the check instances from the instance and passes them to the agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. The agent runs the checks, collects data, and sends the results back to the MID Server.
4. The MID Server sends the collected data to the ServiceNow instance, where it is stored in the CMDB or events tables.

The commands and their configurations that run on the Agent are called **checks**. The Agent comes by default with **check definitions**, which determine a specific command and the default frequency with which it runs. Checks are defined on the instance and are passed to the Agent via the MID Server.

A **policy** is a combination of the CIs being monitored by the Agent Client Collector and the checks that run on those CIs. You associate check definitions with policies. Those check definitions are then referred to as **check instances**. You can customize check instances to meet your needs. For example, customize the running interval or the parameters specific to the policy, such as the login credentials to access a MySQL database. Customization of a check instance takes effect only on the check instance associated with the policy, which does not affect the original check definition or already created check instances in other policies.

Domain separation and Agent Client Collector

Domain separation is supported for Agent Client Collector (ACC). Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

How domain separation works in Agent Client Collector

Domain separation in ACC allows the isolation of data collected by the Agent Client Collector to a particular domain. As a result, only users within that domain can access and view this data. Similarly, agents and their associated configurations are tied to specific domains, ensuring that operations within one domain do not affect or leak into others. This creates a more organized and secure environment for managing multiple groups, business units, or regions.

- Data Security and Privacy: Restricting data access based on domains reduces the risk of unauthorized access and cross-domain data leakage, ensuring that sensitive information stays within its designated domain.
- Access Control: By isolating domains, organizations can enforce strict access controls, allowing only designated users or teams to view or interact with the data relevant to their domain.

Use cases for domain separation in Agent Client Collector

- Multi-department collaboration with data isolation: IT and Finance departments can independently monitor their respective systems without data overlap. For example, if there's an issue with a server in the payment processing system, Finance can run specific checks in their domain without interference from IT.
- Regulatory compliance: Domain separation helps comply with regulations like GDPR or SOX by ensuring that only authorized personnel can access sensitive data. For instance, the Compliance department can monitor regulatory checks while being restricted from accessing IT infrastructure data.
- Incident management and troubleshooting: IT can troubleshoot network issues, while Finance monitors transaction systems. Domain separation allows both departments to work independently, ensuring faster resolution and preventing data exposure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Domain separation for service providers

Use case: Data isolation and monitoring with domain separation

Optimizing data isolation and monitoring with domain separation ensures financial institutions protect sensitive information, improve operational efficiency, and maintain compliance by securely segregating departmental data.

Use case overview

In a large financial institution, departments like IT, Finance, Compliance, and HR need to manage their systems independently while keeping sensitive data protected. The Agent Client Collector (ACC) collects data from various agents on servers and devices, and domain separation ensures that each department's data is securely isolated. This allows departments to manage their own monitoring, troubleshooting, and reporting, greatly reducing the risk of cross-domain interference or unauthorized access.

Challenges without domain separation

- Data privacy and security risks: Sensitive data could be exposed to unauthorized users. For example, IT staff could inadvertently access financial or payroll data, violating privacy policies.
- Operational inefficiency: Departments would have to share configurations, making it harder to manage and troubleshoot issues, leading to confusion and inefficiencies.
- Compliance issues: Meeting regulatory requirements for data isolation would be difficult, risking compliance violations.
- Increased risk of data corruption: Shared environments can lead to accidental data corruption, as changes in one department could affect others.

Solution

- Data isolation: Each department has a dedicated domain, ensuring that data collected by agents is accessible only within that domain.
- Independent configuration and monitoring: Departments set their own checks, policies, and tasks without affecting others.
- Access control: Role-based access control (RBAC) restricts user access to data within their domain. For example, IT can access only infrastructure data, while Finance can access only financial transaction data.
- Enhanced security: Domain separation ensures that users in one department cannot access data from other departments, reducing the risk of unauthorized access.

Outcome

- Improved data security: Sensitive data is protected, reducing unauthorized access risks.
- Operational efficiency: Departments operate independently, managing configurations and monitoring without conflicts.
- Regulatory compliance: Easier adherence to data privacy regulations (GDPR, SOX) due to secure data segregation.
- Reduced data corruption: Isolated environments prevent unintended cross-departmental impacts, enhancing system stability.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use case: Regulatory compliance in financial services

A financial institution needs to ensure that its data is accessible only to the departments that have authorization for the indicated data.

Challenge

The financial institution must comply with regulations that require strict data separation between departments (such as HR, IT, or Finance).

Solution

1. Create separate domains for each department (IT, Finance, HR, Compliance, and so forth) in the Agent Client Collector (ACC) configuration.
2. Apply access controls so that sensitive data (such as employee records and financial transactions) is accessible only to authorized departments.
3. Use ACC health instance scans to ensure that data separation is maintained, especially when updating or troubleshooting systems.
4. Periodically audit domain separation settings to ensure compliance with data privacy regulations such as GDPR, SOX, or PCI-DSS.

Outcome

Compliance with legal and regulatory requirements is ensured, and sensitive data is protected from unauthorized access.

Agent Client Collector plugins

An Agent Client Collector (ACC) plugin is a script or group of scripts that extend the Agent Client Collector's capabilities. Plugins enhance monitoring by collecting metrics, performing specialized checks, and triggering events based on conditions, like monitoring an application's queue size when it reaches 60% or 80%. Plugins ensure scalable, customizable monitoring to adapt to evolving infrastructure or application needs.

Plugin-check dependency

Plugins are tightly integrated with checks in the Agent Client Collector Framework. A dependency is created between the plugin and the check to ensure that the plugin is executed alongside the associated check.

- Plugin Dependency: A plugin can be associated with one or more checks, and vice versa, where checks can depend on multiple plugins.
- Execution: Plugins and checks run on the same agent, working in unison to gather data, perform specialized tasks, or trigger necessary events

Plugin-agent compatibility

Agents must be the same version as their corresponding plugins or can be one version lower. For example, an ACC agent version 4.1.0 can work with an ACC-F store version 5.1.0.

Creating and managing plugins

You can create Agent Client Collector plugins as needed to customize the monitoring to your specific environment. These plugins are typically formatted as tar.gz files, which contain the scripts that define the plugin's functionality.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Plugin creation: Custom plugins are developed and packaged into tar.gz files.
- Integration with checks: Once created, the plugin is integrated with checks by establishing the appropriate dependencies, ensuring that the plugin and check run together.
- Customization: Plugins can be customized to meet the specific requirements of your monitoring needs, whether it's adding more detailed metrics, checking system thresholds, or generating custom events.

Pre-installed plugins

Agent Client Collector Framework includes a set of pre-installed plugins that provide basic functionalities for common monitoring tasks, such as gathering standard metrics or monitoring key system parameters. For more detailed information on the plugins that come with the Agent Client Collector as part of ITOM AIOps, see Plugins or applications installed with ITOM AIOps.

Secure parameters in Agent Client Collector

Secure parameters in Agent Client Collector (ACC) refers to securely passing sensitive data, such as user names, passwords, and API keys, during check execution, without exposing the sensitive data in the command line. Parameters are passed to the script through standard input (STDIN), hiding them from logs or any process that might capture command-line arguments.

Depending on the coding language in use, implementation of secure parameters may vary. For example, in Linux's Bash, the *read* command (*read <variable_name>*) is used, populating *<variable_name>* with the value entered in STDIN.

Secure parameters data flow

1. Create secure parameters: Define secure parameters in the instance, referencing credentials that need protection (such as usernames and passwords). For details, see Create secure parameters for a check.
2. Configure secure parameters: Set an **Order** value to determine the sequence in which credentials are passed to the standard input.
3. Pass credentials: Transfer credentials from the MID Server to the agent, based on the check to which they are assigned. When the check executes, the agent passes the secured parameters to the script using STDIN, in the order specified by the **Order** value configured in the instance.
4. Execute the check: Upon execution, the agents send the secure parameters to the script via STDIN, following the specified order.
5. Run the script: The script retrieves and uses the secure parameters from STDIN to execute with the secured credentials.

Secure parameters benefits

- Enhanced security: Passing sensitive data like usernames and passwords directly in the command line is a security threat, as they could be exposed to unauthorized access or malicious users. Secure parameters ensure that such data is securely handled and hidden from view.
- Prevent data exposure: When using secure parameters, the credentials are not written in plain text in the command line, avoiding the risk of them being captured in logs or process listings.
- Data flow control: Secure parameters enable controlled and customizable passing of credentials to the script, ensuring the correct handling and timing of sensitive information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector logs

Agent Client Collector (ACC) logs play a critical role in monitoring the activity and performance of the agent. Logs offer valuable feedback that helps identify potential issues, especially when the agent's performance is suboptimal. By providing insights into areas of concern, these logs are essential for troubleshooting and resolving issues, ultimately improving the overall effectiveness of the agent.

Agent Client Collector (ACC) logs offer key benefits for monitoring and maintaining agent performance.

- Configurable severity levels: Agent Client Collector logs enables you to configure the severity levels for monitoring, including *Error*, *Debug*, and *Info*. This flexibility helps focus on specific areas, such as tracking errors or diagnosing issues.
- Log retention management: You can set up log retention policies, determining how often to delete logs to preserve system resources. This ensures that the logs remain manageable and do not negatively impact system performance.
- Performance issue identification: Logs provide valuable insights to identify areas where the agent may be experiencing issues, helping to improve agent performance.

For details on locating and managing Agent Client Collector logs, see View the Agent Client Collector logs.

Operating system and application monitoring using Agent Client Collector

IT Operations Management (ITOM) monitoring for the Agent Client Collector is installed on the ServiceNow® instance. It monitors the specific operating system and applications that are installed on the host machine.

The Agent Client Collector monitors the following applications.

ACC Applications

Operating System/Application	Configuration description
Windows	No special configuration required.
Linux	No special configuration required.
Apache HTTP Server	See Configure Agent Client Collector Apache HTTP server monitoring.
Apache Tomcat	See Configure Agent Client Collector Apache Tomcat monitoring.
IBM WebSphere	No special configuration required. Accessed through JMX. Default management port is **9100**.
IIS	No special configuration required.
JBoss/WildFly	No special configuration required. Default port is **9990**. For troubleshooting information, see Monitoring performance in the Red Hat customer portal.
Microsoft SQL Server	No special configuration required. Ensure that the instance name is populated in the CI record, either by running Discovery or by manually populating the record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC Applications (continued)

Operating System/Application	Configuration description
My SQL	No special configuration required. Default port is **3306**.
Oracle Database	No special configuration required. Ensure that the SID and Oracle Home path are populated in the CI record, either by running Discovery or by manually populating the record. Default port is **1521**.
Oracle WebLogic	No special configuration required. Accessed via a JMX port. Default port is **7001**. Ensure that the WebLogic home (wls_home) parameter is populated in the CI record, either by running Discovery or by manually populating the record. Ensure that the user running the agent has sufficient permissions to access the JAR files in the **Weblogic** installation directories.

When using the **Microsoft SQL Server**, **Oracle Database**, and **Oracle WebLogic** applications, you must run Discovery to configure CI fields on the instance. Ensure that the Discovery plugin (com.snc.discovery) is installed to enable running Discovery.

When the application is running on a port other than the default, ensure that you update the port number in the check instance's parameters.

Configure Agent Client Collector Apache HTTP server monitoring

To configure the Agent Client Collector to perform Apache HTTP server monitoring, set the following configurations in the Apache HTTP server application.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. On the Apache HTTP Server host, navigate to one of the following locations:
   - In a Linux environment:
     - `/etc/httpd`
     - `/etc/httpd/conf`
     - `/etc/apache2`
     - `/etc/apache2/conf`
   - In a Windows environment: The `conf` directory in the Apache installation location.
2. Locate either the `httpd.conf` or `apache2.conf` file.
3. Ensure that the following line is uncommented:
   `LoadModule status_module modules/mod_status.so`
4. Configure the check instance port to be identical to the port on which Apache is listening (default for both is **80**).
5. To ensure a secure connection, configure SSL, as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Ensure that the following line is uncommented:
`LoadModule ssl_module modules/mod_ssl.so`

b. Activate the *ssl_secure_connection* parameter in the check instance.

c. Configure the check instance port to be identical to the port on which Apache is listening to secure connections (default is **443**).

Related topics

Apache Module

Apache SLL/TLS Strong Encryption: How-to

Configure Agent Client Collector Apache Tomcat monitoring

To configure the Agent Client Collector to perform Apache Tomcat monitoring, set the following configurations in the Apache Tomcat application.

Before you begin
Role required: agent_client_collector_admin

Procedure

1. Open a management JMX port when starting the server:

a. In a Linux environment:

i. In `bin` directory in the Tomcat home directory, create a `setenv.sh` file.

ii. Ensure that the file contains the following:

```
export CATALINA_OPTS=$CATALINA_OPTS"
-Dcom.sun.management.jmxremote=true
-Dcom.sun.management.jmxremote.port={PORT NUMBER}
-Dcom.sun.management.jmxremote.rmi.port={PORT NUMBER}
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.authenticate=false
-Djava.rmi.server.hostname=127.0.0.1"
```

iii. Run `chmod 755 setenv.sh` to give the `setenv.sh` file executable permissions.

b. In a Windows environment:

i. In `bin` directory in the Tomcat home directory, create a `setenv.bat` file.

ii. Ensure that the file contains the following:

```
@echo=off
if defined CATALINA_OPTS (
set CATALINA_OPTS=%CATALINA_OPTS% -
Dcom.sun.management.jmxremote=true
) else (
set CATALINA_OPTS=-
Dcom.sun.management.jmxremote=true
)
set CATALINA_OPTS=%CATALINA_OPTS% -
Dcom.sun.management.jmxremote.port={PORT NUMBER}
set CATALINA_OPTS=%CATALINA_OPTS% -
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Dcom.sun.management.jmxremote.rmi.port={PORT NUMBER}
set CATALINA_OPTS=%CATALINA_OPTS% -
Dcom.sun.management.jmxremote.authenticate=false
set CATALINA_OPTS=%CATALINA_OPTS% -
Dcom.sun.management.jmxremote.ssl=false
set CATALINA_OPTS=%CATALINA_OPTS% -
Djava.rmi.server.hostname=127.0.0.1
```

{PORT NUMBER} represents the number of the port you open for JMX RMI monitoring. The default value created by the check definition is **9000**.

2. Restart Tomcat.
   a. In a Linux environment: Run the `shutdown.sh` script, followed by the `startup.sh` script.
   b. In a Windows environment: Run the `shutdown.bat` script, followed by the `startup.bat` script.

Monitor HTTP points

Use the **Monitoring HTTP Entry Points** and **Monitoring HTTP Entry Points Metrics** policies that come with the ITOM Monitoring scoped app to monitor `http` and `https` entry points and entry points metrics. You can customize these policies as needed, or you can configure a new policy to monitor service entry points.

Before you begin

Role required: agent_client_collector_admin

About this task

An entry point is the access point for a service. In an Agent Client Collector policy, you specify the service you want to monitor, and the policy monitors the service's entry points.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Select either the **Monitoring HTTP Entry Points** or **Monitoring HTTP Entry Points Metrics** policy.
   The relevant **Policy - Monitoring HTTP Entry Points** page opens.
3. On the **Monitored CIs** tab:
   a. In the **Table** field, ensure that the **sn_itmon_http_entrypoint** database view table is selected for the Agent Client Collector to retrieve http and https entry points for monitoring.
   b. In the **Filter** section, configure a filter for the services and entry points to be taken from the **sn_itmon_http_entrypoint** database view table.
4. Click the **Proxy Settings** tab and select the type of proxy you want to enable for the entry point monitoring.
   - **Single proxy agent**: Select to enable monitoring a single agent, which you select in the **Proxy agent** field that appears after selecting this option.
   - **Multi-proxy agents by filter**: Select to enable monitoring multiple agents. Configure a filter in the **Proxy filter** section that appears after selecting this option to monitor only the agents that match the filter's conditions.
   - **Script based proxy**: Select to enable monitoring an agent based on a script.

   You must configure a proxy to run an HTTP policy.
5. Click the **Checks** tab and in the **Filter checks by groups** field, select **HTTP** to display all of the http checks in the **Available** cell.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Select the relevant checks and click the right arrow button to move them to the **Selected** cell.
7. Select the **Active** check box at the top of the page to activate the policy.
8. Click **Update**.
   The policy monitors the entry points specified by the policy's configurations.

Agent-based data flow

The Agent Client Collector gathers information on the health of the CI that the Agent is installed on, and the applications that run on the hosts. The Agent runs several checks and pushes the checks' results to the MID Server, where they are converted into events or metrics.

Events are stored in the instance's Events database, while metrics are stored in the Metrics database.

The following image displays the data flow from the Agent Client Collector to the instance.

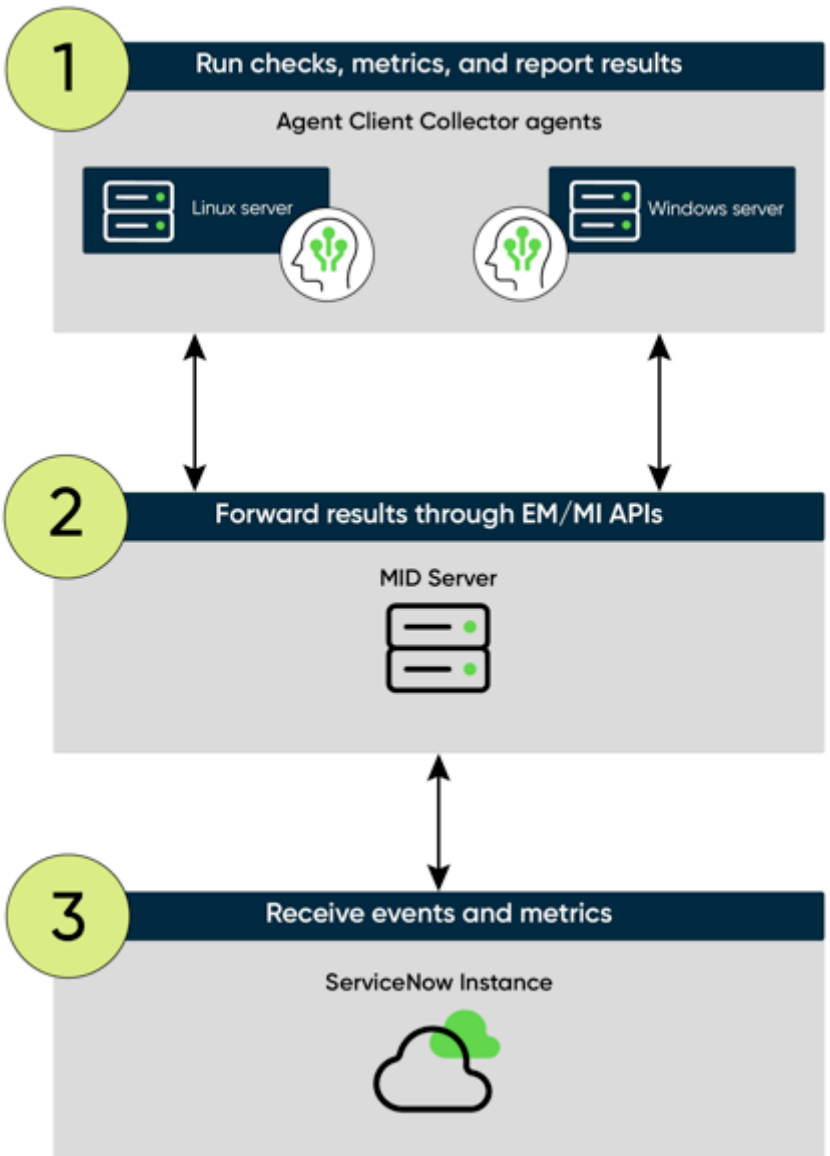

The numbered entries in the flow indicate the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Specified checks run on the agent. Data, including metrics, is pushed to the MID Server.
2. The MID Server analyzes the data and classifies the results as events or metrics. MID Server uses REST API to send the results to the instance.
3. The instance receives the data and stores it in the relevant database.

Agent Client Collector Monitoring use case

The Agent Client Collector Monitoring (ACC-M) use case demonstrates how organizations can achieve unified monitoring across hybrid IT environments.

Use case overview

Many enterprises today operate in a hybrid IT environment, where some resources are hosted on-premises, while others reside in the cloud. This complexity often leads to fragmented monitoring, making it difficult to obtain a unified view of IT service health. Agent Client Collector Monitoring (ACC-M) bridges this gap by providing a comprehensive and converged picture of the organization's services, regardless of where they're hosted. ACC-M is a lightweight monitoring solution that integrates seamlessly into its existing ServiceNow platform.

Challenge

An IT operations team struggled to cohesively monitor service health, due to the disparate nature of its infrastructure. On-premises resources, such as databases and ERP systems, were monitored using traditional methods, while cloud services had separate monitoring tools. This fragmentation resulted in:

- A lack of real-time visibility into overall service health.
- Delayed issue identification and troubleshooting.
- Multiple monitoring dashboards, increasing complexity for the IT team.
- Difficulty in correlating on-premises and cloud alerts, leading to inefficiencies.

Solution

ACC-M enables real-time, unified visibility across a hybrid IT environment. It improves operational efficiency, reduces downtime, and creates a more resilient IT infrastructure capable of supporting both on-premises and cloud-based services.

ACC-M streamlines the monitoring of on-premises and cloud resources into a single, unified display by offering the following solutions:

Solutions implemented by ACC-M

Solution	Description	Outcome
Unified monitoring across environments	ACC-M collects health metrics, performance data, and alert statuses from both on-premises and cloud infrastructure components. ACC-M uses lightweight agents installed on servers, VMs, and cloud resources to gather relevant data, providing:	Improved response: With a unified view of all resources, the IT operations team can quickly identify where issues are occurring and prioritize them based on overall service impact, reducing their mean-time-to-resolution (MTTR).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solutions implemented by ACC-M (continued)

Solution	Description	Outcome
	• CPU and memory utilization for on-premises servers. • Network performance metrics. • Availability and performance metrics for cloud services, such as databases and storage.	
Converged health view	ACC-M provides the company's IT team with a converged health dashboard within the ServiceNow platform, providing: • A unified view of both on-premises and cloud resources. • Real-time updates on service health, ensuring that hybrid services are monitored together and not in isolation. • Integrated insights into service dependencies, correlating incidents that impact both on-premises and cloud systems.	Proactive monitoring and maintenance: ACC-M's real-time health metrics enable proactive monitoring of services, avoiding outages by identifying and resolving issues before they escalate.
Hybrid Alerting and Automation	ACC-M enables setting up automated alerts for both on-premises and cloud environments, streamlining their incident response processes. When a service in the hybrid environment experiences an issue, ACC-M does the following: • Automatically generates alerts with enriched context, identifying whether the issue originated from the on-premises infrastructure or cloud services. • Triggers automated workflows for remediation	Operational efficiency: The convergence of on-premises and cloud monitoring into a single platform eliminates the need for multiple monitoring tools. This reduces complexity and frees IT resources to focus on more strategic tasks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solutions implemented by ACC-M (continued)

Solution	Description	Outcome
	based on the alert type, reducing manual intervention.	
Service Health Visualization	Using the ServiceNow Service Map enables visualizing how services are performing across the entire hybrid infrastructure. ACC-M feeds this with continuous, real-time data, helping the IT team do the following: • Track dependencies between on-premises and cloud services. • Identify potential points of failure before they impact the business. • See at a glance which services are healthy, degraded, or at risk, wherever they're hosted.	Enhanced Business Continuity: With a full view of the hybrid infrastructure health, IT can ensure that critical business services are up and running, improving overall business continuity.

Windows event log monitoring

Windows event log monitoring tracks, analyzes, and manages event logs generated by various Windows OS components.

The Windows OS Events policy includes the os.windows.check-event-log check, which measures the Windows event log against thresholds specified in the check's parameters. For details on the configurable check parameters, see Windows event log filter parameters.

How Agent Client Collector for Visibility Content works

Agent Client Collector for Visibility Content (ACC-VC) requires installation of ServiceNow Agent Client Collector (ACC) on the target host. ACC is a derivative of Sensu-Go, an open-source software.

ACC-VC use cases

ACC-VC applies Checks and Policies to schedule and collect host data. ACC-VC is triggered in the following cases:

- Periodic scheduling: A policy-based approach where Discovery is triggered on a periodic basis.
- On CI delete: When the computer or server CI record is deleted.
- MID Server cycle: When the MID Server goes down and comes back up.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Target host cycle: When the target host goes down and comes back up.
- Network break: When there is a break in the network link to the target.

Note:

- Discovery is triggered for those agents whose hosts are already present. Agents whose hosts are not present are discovered through ACC-F.
- Agent Client Connector for AI Control Tower is available on the ServiceNow Store. This app is an ACC connector and adds an AI proxy capability that collects AI-based traffic from managed devices, supporting AI Control Tower shadow AI use cases.

ACC-VC checks and policies

The ACC-VC assets are stored as Agent plugins with the main entry point [acc_visibility_main] and other modules for OS families. There is one main system Discovery Check definition, called *`Enhanced Discovery`*, which is used by the *`Enhanced Discovery Policy`*. This ACC-VC policy runs off a schedule, which is defaulted to 24 hours (86,400 seconds). This policy configuration is synced to all agents as defined in the ACC-VC policy.

When the payload is returned from the MID Server to the instance, the ACC-VC Check Type, *`EnhancedDiscovery`*, delegates tasks to the EnhancedDiscoveryHandler script include provided by ACC-VC. The script contains logic to process the data from the check and handles tasks such as:

- Data transformation into an identification and reconciliation engine (IRE) compatible payload
- Non-CI data reconciliation (cmdb_running_processes, cmdb_tcp_connections, and so on)

The ACC-VC Check Definition, *`Enhanced Discovery`*, is initiated by the ServiceNow Instance. Then, an ECC Queue record with topic, MonitoringProbe, is created on the output queue with relevant Check information. The MID Server then processes the check by sending a message to the ACC via WebSocket over TLS.

During this time, the MID Server also serves any relevant Assets or Plugins that the ACC requests. The MID Server ensures the Assets and Plugins are relevant to the particular Operating System, platform, OS version, and architecture on which the ACC is running.

You can edit and modify all parts of the ACC-VC application including check type, policy, and check definition. See Checks and policies for more information.

Virtual machines and cloud instances

ACC-VC associates a target, discovered via Discovery, with a pre-existing virtual machine (VM) Instance CIs. ACC-VC associates the discovered CI record with any pre-existing VM Instance record or Cloud Server Instance record with appropriate CMDB relationships.

The following variants of virtualization and cloud server vendors are supported for ACC-VC:

- vCenter
- Amazon AWS Cloud
- Google (GCP)
- Microsoft Azure

Agent Client Collector for Visibility Content use case

The Agent Client Collector for Visibility Content (ACC-VC) use case demonstrates how a financial organization can perform effective discovery across on-premise and remote environments.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use case overview

A leading financial organization needed to overcome the limitations of traditional agentless discovery to ensure comprehensive visibility, compliance, and security across both on-premise and remote environments.

Challenges

- Limited visibility into restricted environments: The organization's security policies prevented MID Servers from initiating connections to highly sensitive systems, such as core banking infrastructure and financial databases. This created visibility gaps in critical environments.
- Incomplete coverage of remote endpoints: With a rapidly expanding remote workforce, many devices operated outside the corporate network boundary, making them inaccessible to agentless discovery. This resulted in gaps in asset tracking and monitoring.
- Compliance and audit challenges: Strict regulatory requirements required granular, auditable, and continuously updated asset data. Agentless methods lacked the depth and consistency needed to meet these standards.

Solutions

To address these challenges, the financial organization implemented Agent Client Collector for Visibility Content (ACC-VC) as part of its ITOM strategy.

- Agent-based data collection for secure environments: Agents were deployed on restricted and high-security systems, enabling:
  - Secure, outbound communication without violating firewall policies.
  - Continuous monitoring without requiring inbound access.
- Remote endpoint visibility: Agents installed on employee devices provided continuous visibility, regardless of network location, and real-time data collection from off-network systems.
- Deep and continuous discovery: ACC-VC enabled:
  - Rich system telemetry and configuration data collection.
  - Continuous updates to the CMDB, improving data freshness and accuracy.
- Compliance and policy enforcement: Agents enforced and monitored endpoint level compliance by:
  - Capturing detailed, auditable system data.
  - Supporting compliance frameworks and audit requirements.

Outcomes

- End-to-end asset visibility: The organization achieved complete visibility across secure data center assets, cloud and hybrid environments, and remote workforce devices. This eliminated blind spots and improved operational decision making.
- Improved CMDB accuracy and reliability: Continuous data collection ensured up-to-date and trustworthy CI data, and improved service mapping and dependency insights.

> **Note:** Shadow AI detection is another key ACC-VC use case. The Agent Client Connector for AI Control Tower app, available on the ServiceNow Store, is an ACC connector that adds an AI proxy capability to collect AI-based traffic from managed devices, supporting AI Control Tower shadow AI use cases.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Enhanced compliance and audit readiness: The organization benefited from granular, auditable asset data, faster and more efficient compliance reporting, and reduced risk of regulatory violations.
- Increased operational efficiency: Reduced reliance of manual discovery and troubleshooting led to faster issue detection and lower operational overhead.

ACC-VC enabled the organization to transform its discovery strategy from periodic with limited visibility to continuous and intelligent, with comprehensive observability. This enabled the financial institution to:

- Secure critical systems
- Support a distributed workforce
- Maintain regulatory compliance
- Build a resilient and scalable IT operations framework

Configuring Agent Client Collector

Plan and configure your Agent Client Collector implementation.

Agent Client Collector configuration typically follows these steps:

1. Preparing for Agent Client Collector implementation
2. Deploying Agent Client Collector on servers
3. Deploying Agent Client Collector on endpoints
4. Deploying Agent Client Collector on both servers and endpoints

Preparing for Agent Client Collector implementation

Successful Agent Client Collector (ACC) implementations, especially involving large numbers of endpoints or servers, require careful deployment. When proceeding with a large-scale ACC deployment, follow the steps described in the planning and deployment checklists.

Agent Client Collector planning checklist

Successful Agent Client Collector (ACC) implementations, especially involving large numbers of endpoints or servers, require careful planning. Before proceeding with a large-scale ACC deployment, follow the steps described in the planning checklist.

Agent Client Collector planning

Step number	Action	Description
1	Define use cases, targeted devices, and success metrics	Typical ACC use cases include ITOM Visibility, ITAM (SAM/HAM), ITSM (DEX), ITOM Health (HLA/ Metric Intelligence). Many customers have success metrics targeting a percentage of devices with successful data collection for a use case over a 7 to 14-day window (for example, “95% of endpoints report software usage over a 14-day period”).
2	Identify Key Stakeholders	Involve network, security, and device/ software management teams (such

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector planning (continued)

Step number	Action	Description
		as jamf/Intune experts) as early as possible during planning.
3	Assess ServiceNow License, Store App and Instance Requirements	Ensure that you have all license entitlements for your use cases. Determine which software updates or configuration changes to your production and sub-production instances are needed to support your targeted number of devices.
4	Assess Software Requirements	1. Ensure that all CPU architectures and operating systems for your use cases are supported. If not, determine whether non-ACC solutions such as Service Graph Connector or Discovery Patterns can be used. 2. Determine whether your use cases require elevated system privileges on endpoints or servers (for example, root access on Linux for device serial number tracking).
5	Assess Network Requirements	Verify whether there are VPN, firewall, or proxy restrictions that prevent ACC from sending data to your ServiceNow instance.
6	Assess Device Management Requirements	Verify which internal solutions are available for deploying software within your organization on all targeted devices. For example, Intune and Jamf for endpoints; Chef, Puppet, or Ansible for servers.
7	Draft ACC Solution Architecture	Work with your ServiceNow account team to define how ACC agents send data to the ServiceNow instance, based on your use cases and differences in server and endpoint deployment techniques. Verify whether direct connections are made using ITOM Cloud Services or MID Servers.
8	Define and Segment your ACC Deployment Plan	Define evaluation, pilot, and broad rollout stages and the number of devices targeted in each phase. Align each segment with risk tolerance and business goals. ACC must be rolled out gradually, with no more than 5,000 agents per hour.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector planning (continued)

Step number	Action	Description
9	Internal Compliance and Security Review	Work with internal teams and your stakeholders to review your deployment plan, solution architecture, and to obtain approvals.

Agent Client Collector deployment checklist

Successful Agent Client Collector (ACC) implementations, especially involving large numbers of endpoints or servers, require careful deployment. When proceeding with a large-scale ACC deployment, follow the steps described in the deployment checklist.

Agent Client Collector deployment

Step number	Action	Description
1	Update production and sub-production instances	Ensure that all instances are updated to the most recent store applications, family releases, or both.
2	Assess instance capacity	Open support tickets if additional instance capacity is needed for the projected number of ACC agents.
3	Deploy and configure MID Servers when required	If connecting over 1,000 ACC agents to a single MID Server, assess performance and whether a load balancer in front of a MID Server is needed.
4	Verify whether ITOM Cloud Services (ICS) is required and, if so, whether it is configured	Some use cases (such as DEX) always require ICS.
5	Download appropriate installer files from the instance	Installer files may be in one of the following formats: • .rpm • .msi • .pkg
6	Define a deployment calendar	Share deployment groups and group size, assignments, start/stop, success gates, go/no-go signers, and rollback criteria with stakeholders.
7	Package ACC in appropriate software management tools	Software management tools include:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector deployment (continued)

Step number	Action	Description
		• Intune • jamf
8	Deploy pilot and test environments on sub-production instances using your preferred software management tool	In test and pilot environments, you can test multiple operating systems and hardware combinations (such as server vs. laptop).
9	Monitor the ServiceNow instance and the software management solution	• On the instance, monitor the performance and rate limit violations. • On the software management solution, monitor the deployment success rate.
10	Assess data collection, install success rate, and instance performance over the next 2-3 business days	Workloads and data collection can vary significantly, depending on the time of day; therefore, monitor deployment success over several days.
11	Update deployment calendar and continue with next install cohorts	After successful deployment, continue with data collection.

Deploying Agent Client Collector on servers

When deploying the Agent Client Collector, perform deployment and management tasks on your environment's servers.

Server deployment uses MID Servers to connect your agents to a ServiceNow instance. Agents deployed on data center servers open WebSocket connections to your MID Servers, which relay data to your instance. You manage the sizing, certificates, and load balancing of the MID Server.

Use server deployment when your agents are in controlled environments such as data centers, where you are already running MID Servers and want to keep traffic on your network. Agent Client Collector Monitoring (ACC-M) and Agent Client Collector Log Analytics (ACC-L) require server deployment, and you can deploy Agent Client Collector for Visibility Content (ACC-VC) with servers as well.

Configuring Agent Client Collector with a MID Server

Configure the Agent Client Collector with a MID Server to enable communication between network servers and the ServiceNow instance.

Configure an Agent Client Collector with a MID Server using a single-line command

Use a single-line command to set up an Agent Client Collector with a MID Server. You can also use the single-line command to migrate agents with MID-less configuration to be used with a MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required:

- In a Linux or macOS environment: root
- In a Windows environment: agent_client_collector_admin

About this task

Running the single-line command includes flags, which indicate the options to be added to the `acc.yml` configuration file. These options dictate the configuration of the Agent Client Collector with the MID Server. If you don't specify an option in the command, the existing option in the `acc.yml` file is used. If there's no existing option, an error message indicates that the parameter is required to enable connecting to an endpoint.

Running a command validates that a connection can be made with the specified parameter (for example, the specified api-key value). However, the command doesn't verify if the value of the parameter is valid.

Procedure

1. In a terminal window, navigate to the Agent Client Collector installation location:
   - Linux: `/usr/share/servicenow/agent-client-collector/bin/`
   - macOS: `/opt/servicenow/agent-client-collector/bin`
   - Windows: `C:\Program Files\ServiceNow\agent-client-collector\bin`
2. Run the following command: `acc gateway mid <flags>`.

   For example (replace with your actual values): `acc gateway mid -a fbhsbnxdbshs635oo -b wss://1.2.3.4:8800`.

   The available flags appear in the following table:

MID Server configuration flags

Flag	Type	Description
-y, <--acc-yml>	String	Location of the `acc.yml` being configured. Can also be used to modify and validate parameters if the file is in a different location that the default. Optional. If no value is entered, the default location of the `acc.yml` file is used.
-a, <--api-key>	String	The API key to be used by the MID Server.
-b, <--backend-url>	String	One or more backend connection URLs specifying a websocket endpoint to a MID Server. Multiple values are specified using commas between entries

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Flag	Type	Description
		with no spaces separating the values; for example, `<url1>,<url2>`.
-h, help	N/A	Explains what the specified command does.

3. Restart the agent for the configuration settings to be entered into the `acc.yml` file.

Configure the websocket server on the MID Server

Configure the websocket server on MID Servers to enable connections from agents to the MID Server. You can configure only one websocket server per MID Server.

Before you begin

- Ensure that you download and install curl (or similar software) on your MID Server Windows machine.
- Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **MID Server** > **Servers**.
2. Select the relevant MID Server name.
3. In the **Related Links** section, click **Setup ACC Listener**.
   The **Agent Client Collector Setup** dialog box appears.

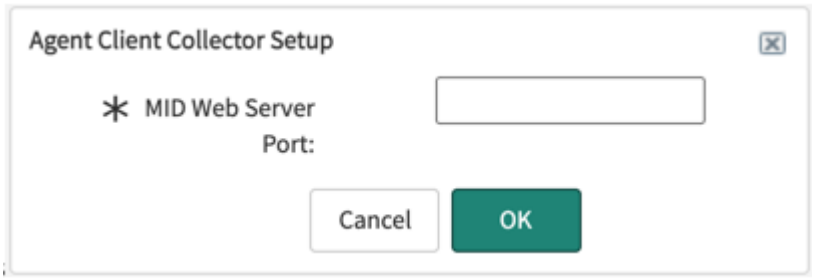

4. Enter the MID web server port (**8800**) in the displayed field.

   Note: You can change the default port, **8800**, as needed.

   - To invoke secure TLS configuration, configure the following in the `acc.yml` file: `wss://<mid_ip>:<websocket_port>/ws/events`
   - To invoke TLS 1.3 support on the MID Web Server, set the following parameter in the `wrapper-override.conf` file on the MID Server: `-Djdk.tls.server.protocols="TLSv1.3"`
   - In a Linux environment, port numbers under 1024 can be used only with superuser privileges.

5. Click **OK**.
   A confirmation message with the URL for the agent configuration file appears.

   The web server and websocket extension are configured for the MID Server. To verify the configuration, navigate to the following menu options:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Agent Client Collector** > **MID Web Server**. When editing the fields on the resulting form, you must restart the MID web server extension to save your changes.
- **Agent Client Collector** > **Websocket Endpoint**

6. Verify the validity of the websocket, using curl:

   a. Enter the following command into a cmd or terminal window: *curl -Lk -H "Authorization: Key <API_KEY>" https://<mid_server_address>:<port>/api/mid/mon*

   b. Verify that you receive a valid JSON response (and not an error message), such as the following:
   Successful installation JSON message:

```
{
                "connectedAgentsNum" : 1
}
```

Configure a websocket endpoint

Configure a websocket to enable the endpoint to provide web access from the MID Server to the Agent Client Collector (ACC).

Before you begin

Role required: agent_client_collector_admin

About this task

For added security, ensure that MID Servers deployed for the Agent Client Collector aren't deployed for other components.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **MID Servers**.
2. Select a MID Server.
3. Select the **ACC Websocket Endpoints** tab on the bottom of the page.
4. Select **New**.
5. Configure the fields on the page.

ACC Websocket Endpoint page

Field	Description
Name	Enter a descriptive name for the websocket endpoint.
MID Server	The name of the selected MID Server.
Short description	Optionally, enter a short description for the websocket endpoint.
Accessible IP address	Apply an external/public IP when the MID Server has an internal IP that is not accessible from an agent. Used for cloud solutions with an agent that has access only to the MID Server's public IP. If the MID Server and all agents are in the same Virtual

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	private cloud (VPC), this field can be left empty.

6. In the **Related Links** section, select **Start** to activate the websocket endpoint.

Configure the frequency of updating the agent MID Server list

By default, the list of MID servers connected to agents is updated once daily. If you have a dynamic environment that adds MID servers frequently, you may want to schedule updates more often, or execute the job on demand.

Before you begin
Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Locate the **Update Mid List for All Agents** job.
3. To run the job immediately, click **Execute Now**.
4. To edit the frequency by which the job runs:

   a. In the **Run** field, select **Periodically**.
   The **Repeat Interval** field appears.

   Repeat interval field

   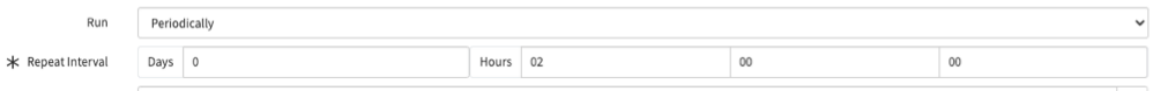

   b. Configure the frequency by which you want the job to run.
   The cells next to the **Hours** cell represent minutes and seconds, respectively.

   c. Click **Update** to update the job frequency.

Enable Agent Client Collector monitoring on the MID Server

Enable Agent Client Collector monitoring on the MID Server by configuring the Agent Client Collector framework (ACC-F) and Metric Intelligence extensions in a single action. Enabling Agent Client Collector monitoring also enables monitoring on agents connected to the MID Server.

Before you begin
Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **MID Servers**.
2. On the MID Server page, select a MID Server to view the MID Server's details.
3. In the **Related Links** section, select **Setup ACC Monitoring** and enter the relevant port number in the **MID Web Server Port** field of the dialog box.

Result
The MID Server is configured to support both the Agent Client Collector and Metric Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You are prompted to enter a port number only on initial ACC Monitoring setup.

Automatic MID Server selection

Automatic selection of MID Servers ensures that each agent uses the most efficient available MID Server.

The Agent Client Collector automatic MID Server selection feature set is disabled by default. There are several factors to consider before determining if automatic MID Server selection is right for your use cases.

Use case considerations

Automatic MID Server selection provides several benefits. MID Server capabilities can narrow the list of MID Servers an agent can potential connect to. Automatic MID Server selection keeps the back-end URL configuration list up to date on the agent host without additional work flows.

When automatic MID Server selection is enabled, the agent performs a connectivity test against the list of available MID Servers and determines the best connection based on latency and number of currently connected agents.

However, there are several additional considerations that may make automatic MID Server selection wrong for your use case.

The ACC automatic MID Server selection is not a true load balancer and should be considered a fail-over mechanism. There is higher complexity for choosing a back-end URL as it is derived from multiple input sources. If, instead, an application load balancer is used only one back-end URL needs to be specified. See Enable the Agent Client Collector load balancer for more information.

By default, each agent is sent the list of all MID Servers listening for ACC connections. Since that list includes MID Servers running in restricted areas, the agents can try to connect to them which can generate security alerts.

Your DNS configuration should be considered before using automatic MID Server selection because the resolution can lead to unexpected results. The DNS resolution is done periodically. End User Compute systems that alternate between internal corporate, Wi-Fi, VPN, or external networks may not be able to reconnect if the network context has changed.

Automatic selection process

When automatic selection is enabled, the ServiceNow instance sends periodic MID Server updates to existing agents. The agent selects a MID Server based on information it receives from the ServiceNow instance.

During automatic MID Server selection, the agent selects a primary MID and a secondary MID list, and writes them to the agent's config file. The agent uses this list for fail-over protection when the primary MID/secondary MID in the list loses connection.

When an agent detects that its IP address has changed, it automatically connects with the most efficient MID Server available from the agent's list of MID Servers.

To secure communication between the MID Server and the agent, use the MID Server's unified key store. For details on using the MID Server unified key store, see MID Server unified key store.

Related topics

Enable automatic MID Server selection

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Enable automatic MID Server selection

Enabling automatic MID Server selection lets the agent perform a connectivity test against the list of available MID Servers and determine the best connection based on latency and number of currently connected agents. Automatic MID Server selection is disabled by default.

Before you begin
Role required: agent_client_collector_admin

Procedure

- To enable sending periodic MID Server updates from the ServiceNow instance to existing agents, set the *sn_agent.enable_auto_mid_selection* system property to **true**.
- To enable automatic MID Server selection for individual agents:
  1. Set the *enable-auto-mid-selection* property in the agent's `acc.yml` file to **true**.
  2. Enable and execute the Update MID List for All Agents and Update MID List for All New Agents scheduled jobs.

What to do next
To disable automatic MID Server selection again, set each of the properties back to **false**.

Configure the Agent Client Collector capabilities for MID Servers

Configure the MID Servers with Agent Client Collector capabilities to enable the MID Servers to work with agents. Agents that have identical MID Server capabilities are eligible for automatic MID Server selection.

Before you begin
Role required: agent_client_collector_admin

About this task
A capability defines MID Server functions. Any MID Server that is assigned with a capability that has **Is ACC Capability=true** is enabled to work with an agent. For details on the MID Server capabilities, see MID Server capabilities.

Procedure

1. Navigate to **All** > **MID Server** > **Capabilities**.
2. Select a capability to enable the MID Server to work with the agent.
   a. Select the ACC Capability column next to the AgentClientCollector entry.
   b. Set the value to **true**.
3. Enable other capabilities for Agent Client Collector MID Servers by selecting the Is ACC Capability column next to the AgentClientCollector entry and setting the value to **true**.

Result
MID Servers with the same capabilities are eligible to be used during automatic MID Server selection.

Configure the automatic MID Server selection based on the Agent Client Collector capabilities

Configure the MID Servers that are ineligible for automatic MID Server selection with the agent. For example, you may want to block a MID Server from connecting to your agent due to firewall constraints.

Before you begin
Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Agents with different levels of Agent Client Collector (ACC) capability can't be considered for automatic MID Server selection. To block a MID Server from automatic selection, assign it different ACC capabilities from all other MID Servers in your environment. Do not assign different capabilities from the primary MID Server.

Procedure

1. Configure the Agent Client Collector capabilities for MID Servers, as described in Configure the Agent Client Collector capabilities for MID Servers.
2. Navigate to **All** > **Agent Client Collector** > **MID Servers**.
3. Select the MID Server that you want to be ineligible for automatic MID Server selection.
4. Select the additional actions icon and select **View** > **Default View**.
5. Select the **Capabilities** tab.
6. Ensure that the MID Server isn't assigned with the capabilities that have **Is ACC Capability=true** on the **Capabilities** tab.

Example: Configuring MID Servers to be ineligible for automatic MID Server selection

1. Assign the **AgentClientCollector** capability to MID Server A and MID Server C, but not to MID Server B.
2. On the MID Server Capabilities page:
   a. Assign a capability to MID Server A and MID Server C.
   b. Assign a different capability to MID Server B.
3. Install Agent Client Collector Framework and specify MID Server A as the MID Server during installation.

   After installation, the agent has **Status=Up** and its MID Server is MID Server A.
4. Execute the MID Server Selection scheduled job to determine which MID Servers are eligible for automatic MID Server selection.

   MID Server B is ineligible for automatic selection, because it doesn't have the same ACC capability as MID Server A. MID Server C is eligible, because it has the same ACC capability as MID Server A.

Agent Client Collector Installation

You can install the Agent Client Collector on any supported host machine. The Agent Client Collector connects to a MID Server using the HTTP/S protocol, and the connection remains active after being established. One MID Server may handle several agents simultaneously, while a single agent works with one MID Server at a time and switches to a different MID Server when necessary to provide failover protection.

When an agent's IP address changes, it selects a MID Server to connect to based on the agent's MID Server list.

Verify that your server's OS and version is supported. For a list of supported OS's and versions, see ServiceNow Store Page

The maximum number of agents that can be connected to a single MID Server is configurable in the *sn_agent.mid.max_allowed_agents* MID Server property. The default value is 4,000.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For the list of required OS versions, see For a list of supported OS's and versions ServiceNow Store Page

For ACC-VC, a default 1 GiB MID Server can support 700 agents concurrently. An 8 GiB configuration for a MID Server can support 8,000 agents concurrently. You can also scale out. For example, 5 MID Servers with 8 GiB of heap size can handle up to 40k agents.

Agent Client Collector is installed with ruby version 3.3.2.

The default user account is a local user called **servicenow**. This user has basic level permissions.

The following table describes the permissions required to work with various features associated with the agent, per OS. An asterisk (*) indicates that no special permissions are required.

Agent on Endpoints Deployment

Feature	Windows	Linux	macOS
Basic inventory	*	*	*
Serial number(s)	*	sudo dmidecode	*
Running processes	Debug programs	*	*
Mapping TCP connections to running processes	*	sudo ss	*
Storage devices	LOCAL SYSTEM	*	*
Logged-in users	LOCAL SYSTEM	*	*
Package self-upgrade	LOCAL SYSTEM	sudo rpm/dpkg	Not supported

If you completely reinstall the agent on a single host server, and you're using ITOM Cloud Services, you must delete the agent record to allow for re-registration. Delete the original agent on the **Agent Client Collectors** page (**All** > **Agent Client Collector** > **Agents**).

Agents whose **Status = Down** or **Disconnected** which haven't been deleted are deleted automatically after 30 days. You can modify this setting on the Autoflush form page (see Autoflush form).

Use the Manual Transport Layer Security protocol (mTLS) for secure authentication between your MID Web Server and the agent (the client). For details, see Connect the agent to the MID Server using mTLS.

For details on using Agent Client Collector in an air-gapped environment, see the Agent Client Collector Framework Air Gapped Configuration Item Management Solution [KB1585753] article in the Now Support Knowledge Base.

Golden image mode enables cloning of additional instances. Setting golden image mode is described in the installation procedure prerequisites for each OS. For information on the structure and modularity of the golden image plugin by operating system, see Golden image structure and modularity.

Agent Client Collector supports domain separation. The domain of the agent and the CIs it creates is determined by the domain of the MID Server that the agent is connected to. The

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

user's domain must be the lowest domain level (known as a leaf domain) to enable creating a Websocket endpoint extension for the MID Server.

Related topics

Agent Client Collector installation on a Linux OS system

Install the Agent Client Collector on a Windows machine manually

Agent Client Collector installation on a macOS system

Agent Client Collector system requirements

System requirements are the fundamental specifications and configuration needed to install and run Agent Client Collector (ACC) effectively.

Verify that your server's OS and version is supported. For a list of supported OS's and versions ServiceNow Store Page

Before installing the Agent Client Collector, you must do the following:

- When working with a MID Server:
  - Ensure that one or more MID Servers are properly registered and validated with your instance, to be available for an agent connection request. For details, see Configuring MID Servers. In a staging environment, you need only one MID Server. However, in a production environment, configure at least two MID Servers to support zero-touch configuration and ensure that a MID Server is always available if one fails.
  - Ensure that there's a validated connection from the designated server where you're installing the agent to the MID Server.
  - Ensure that the MID Server port on which the Agent Client Collector listener runs is accessible.
- When working with MID-less deployment, follow the instructions described in Configuring MID-less Agent Client Collector.
- Ensure that the following plugins are installed on your instance:

Application	Plugin
ACC-F	○ MID Server ○ CMDB for scoped applications ○ Agent Client Collector global utilities ○ MID Web Server ○ Service Error Management
ACC-VC	○ Agent Client Collector for Visibility Content ○ Agent Client Collector Framework ○ Discovery [com.snc.discovery] plugin
ACC-M	○ Agent Client Collector Framework ○ Agent Client Collector Monitoring

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	Plugin
	○ Event Management and Service Mapping Core ○ Metric Intelligence - [com.snc.sa.metric] plugin
ACC-L	○ Agent Client Collector Framework ○ Agent Client Collector Log Analytics ○ Health Log Analytics
DEX	○ Agent Client Collector Framework ○ ITOM Cloud Services Core [sn_itom_cloud_svc]

You can verify that these plugins are installed on the **System Definitions** > **Plugins** page.

All plugins that come with the base system are signed with the ServiceNow certificate. Optionally, you can use custom plugins and sign them using your own certificate, as described in Enable OpenSSL secure signing for plugins.

- Ensure that your system has the minimum requirements for agent installation.

Minimum installation requirements

Agent Client Collector component	Requirements
Agent Client Collector Framework	○ 1 GB RAM ○ 300 MB free disk space ○ 1 CPU core
Agent Client Collector Monitoring	○ 1 GB RAM ○ 1 GB free disk space ○ 1 CPU core
Agent Client Collector for Visibility Content	○ 1 GB RAM ○ 1 GB free disk space ○ 1 CPU core

- Ensure that you're installing Agent Client Collector on one of the supported operating system distributions. For details on the supported operating systems, see Agent Client Collector Framework in the ServiceNow® Store.

Important: All supported operating systems work only with the x86-64 CPU architecture type. ACC-F v3.1.1 also adds native support for ARM64 on Apple Silicon architectures. For the latest information, see the Agent Client Collection on non- x86-64 computers and embedded boards [KB1172387] article in the Now Support Knowledge Base.

When using ACC-L: The Agent Client Collector comes with the default *`servicenow`* user. Ensure that this user has read access to enable Agent Client Collector to view all the configured log paths. For example, the Agent Client Collector *`servicenow`* user that comes installed with the base system does not have permissions to view the paths to `/var/log/` in Linux and C :

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

`\Windows\System32` in Windows. For information about configuring permissions for the *`servicenow`* user, see the ACC-L Permission Denied issues [KB1117271] article in the Now Support Knowledge Base.

When determining the number of machines and servers needed to support your agents (such as the number of machines requiring proxy agents), consider the agents' performance KPIs. For details on agent performance KPIs, see Agent Client Collector performance and footprint for URL monitoring.

Deploy the Agent Client Collector using a custom base image

Deploy the Agent Client Collector on a virtual machine during mass deployment using the machine's base image. Mass deployment uses silent installation, which hides installation status.

Before you begin

Use Microsoft SCCM to manage the deployment and security of your devices, as described in the Create a Windows service account with Log on as a Service KB article.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agent Downloads**.
2. Download the installation file of the relevant OS.
3. Install the agent according to the relevant OS instructions.
4. Run the following command to stop the agent service:
   - In a Windows environment: Run the command `sc stop AgentClientCollector`
   - In a Linux environment: Run the command `sudo systemctl stop acc`
5. Ensure that the `acc.yml` file doesn't include an encrypted API key.
   For example: `api-key: "<insert plain text API key here>"`
6. Delete the contents of the agent's cache directory, keeping the directory structure intact:
   - In a Windows environment: `C:\ProgramData\ServiceNow\agent-client-collector\cache`
   - In a Linux environment: `/var/cache/servicenow/agent-client-collector`
7. When installing on a Windows machine, ensure that the installation was successful by doing one of the following:
   - Enable logging in silent installation and verify that the log file ends with the string, `Windows Installer installed the product.`
   - Ensure that the **AgentClientCollector** service has the status `Running`.
8. When installing on a Linux machine, access the machine where the agent is installed and modify the `/usr/share/servicenow/agent-client-collector/serial_number.txt` file in one of the following ways:
   - Change its content with the host's serial number, by running the following command: `sudo -n dmidecode -s system-serial-number`
   - Assign sudo permissions to the ServiceNow user to enable running the following `osqueryi` command: `sudo -E env "PATH=$PATH" osqueryi --logger_min_status=3 --line "select hardware_serial from system_info"`

> **Note:** You must change the virtual machine's host name after deployment to a unique value, to avoid damaging the host's CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. As a post-deployment step, restart the agent service.
   - In a Windows environment: Run the command `sc start AgentClientCollector`
   - In a Linux environment: Run the command `sudo systemctl start acc`

Enable the Agent Client Collector load balancer

Enable a load balancer to ensure that you have functional MID Servers. A load balancer distributes resources over multiple MID Servers to ensure that no single MID Server is overloaded.

Before you begin

1. Install the Agent Client Collector Framework.
2. Configure the Agent Client Collector web server.
3. Ensure that you have access to the internal application's load balancer configuration.
4. Verify that the file `<mid_server_install_path>/static/hello.txt` was created during MID Server setup.

   `hello.txt` is an empty file that is created manually, to retrieve an HTTP response 200.

Role required: agent_client_collector_admin

About this task

Load balancers are not configured by default. Follow this procedure to configure a load balancer in your environment.

Procedure

1. Configure the load balancer so that it can be used in the *`backend-url`* parameter in the `acc.yml` configuration file.
2. Perform the following setup actions on the load balancer:

   a. Increase the idle timeout to 300 seconds.

   b. Configure the load balancer target group with *`https`* and port 443.
   Configure all MID Servers behind the load balancer as targets.

   c. Configure the load balancer health check to monitor the `<mid_server_install_path>/static/hello.txt` file created during MID Server setup.
   The health check ensures that there are no communication interruptions between the load balancer and the MID Server.

3. Add the `<mid_server_install_path>/static/hello.txt` file to the static folder for each MID Server.
   You can place the file in another location that doesn't require credentials, to simplify the setup.
4. Update the acc.yml configuration file with the appropriate FQDN (*`backend-url`* value) and port number, and turn off automatic MID Server selection.

```
backend-url:
 - wss://<fqdn>:<port>/ws/events
enable-auto-mid-selection: false
```

5. Restart the Agent Client Collector.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

The Agent Client Collector data flow proceeds as described in Agent Client Collector architecture.

Note:

When working with SSL, all certificates or certificate chains used by the configured MID Servers must be signed by root certificate authorities whose certificates are bundled inside a single `.pem` file. To transfer the signed certificates to the agent, specify the `.pem` file in the *`trusted-ca-file`* property of the agent's `.yml` configuration file. For information on enhancing security with SSL/TLS configuration, see the ITOM Agent Client Collector documentation material article in the Now Support Knowledge Base.

Restart an agent manually

Perform manual restart of an agent when the agent configuration file has been refreshed, or if the agent is unstable. You can perform manual restart only on agents installed in a Windows environment and for Linux-based agents that use `systemd`.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
   The **Agent Client Collectors** page appears with the list of agents.

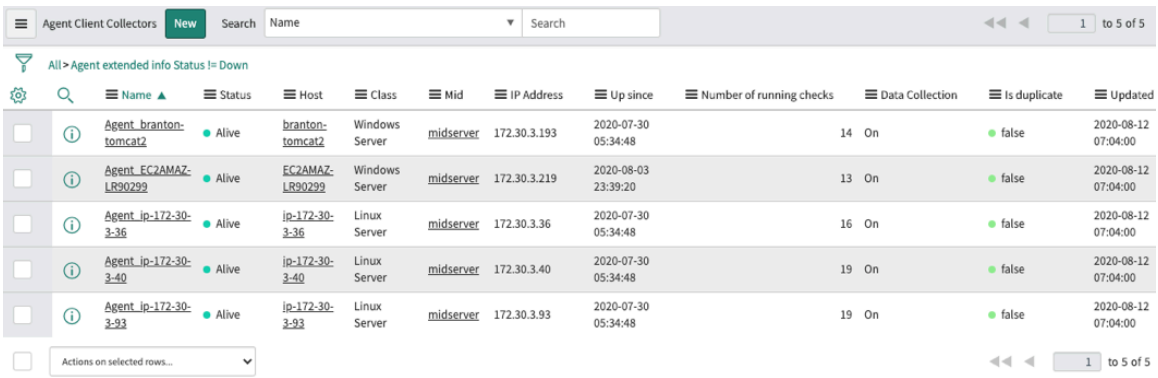

2. Select the check box of the agent you want to restart.
   You can select multiple agents.
3. In the **Actions on selected rows...** list underneath the list of agents, select **Restart agent**.
   A confirmation dialog box appears.

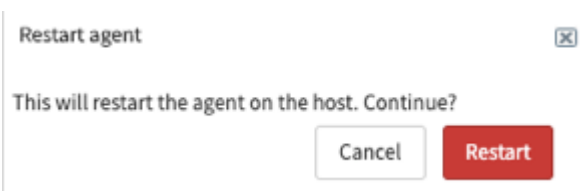

4. Select **Restart**.
   Alternatively, you can restart the agent from within the agent record by selecting the agent name and, in the **Related Links** section, select **Restart agent**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector installation on a Linux OS system

Install Agent Client Collector on a system that uses a Linux OS, either using a single-line command script, or following the installation procedure to embed in your own package distribution mechanism. The installation procedure provides consistency when using package distribution solutions.

Install Agent Client Collector on a Linux system

Install Agent Client Collector using a package distribution tool.

Before you begin

- You can manually install the Agent Client Collector on a few machines to verify that your agents contain the correct policies and checks.
- Ensure that the Agent Client Collector Listener is configured on your MID Servers, and the service is available from your target hosts, see Configure the websocket server on the MID Server.
- Verify that your server's Operation Systems and version is supported. For a list of supported Operation Systems and their versions, see ServiceNow Store Page.
- Verify whether there are restrictions or requirements to be aware of during deployment, such as specifying an account other than the default servicenow account. For more information about embedding the agent into your own automated system, see ITOM Agent Client Collector documentation material [KB1122613].
- Ensure that the MID Server and its MID Web Server and ACC Websocket Endpoint extensions are up and running.
- Enable golden image mode for cloning additional agents by setting the golden image marker located at `/tmp/acc-goldenimage`. The golden image marker takes no action during new Linux installations, as there is nothing which requires cleaning.
- Retrieve the MID Server ACC Listener information to be specified in the agent's *`backend-url`* parameter.
  1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **MID Servers**.
  2. Select a MID Server.
  3. Select the **ACC Websocket Endpoints** tab.
  4. Select a websocket endpoint.
  5. Copy the value in the **Endpoint URL** field.
- Retrieve the MID Server API key specified in the agent's *`api-key`* parameter.
  1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **MID Web Server API Key**.
  2. Select the API key to use.
  3. In the **Related Links** section, select **View API key**.
  4. Copy the API key value and close the pop-up window.

Role required: agent_client_collector_admin

About this task

During Linux installation (and upgrade), the agent executable file is enabled with Linux capabilities (CAP_SETFCAP, CAP_SETPCAP) by default. Store apps such as Agent Client Collector Log Analytics (ACC-L) can use this to grant capabilities to read the entire file system (CAP_DAC_READ_SEARCH). The system undergoes various security measures to verify that granting capabilities does not pose a security risk. These measures include double verification

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

of content origination and the plugin verification process. This procedure assumes that you are familiar with commands for Linux capabilities.

To opt out of these enhanced capabilities, run the following commands, based on your Linux OS/packaging system:

Linux enhanced capability opt-out commands

OS/Packaging System	Commands

RPM	`ACC_SKIP_CAPS=true yum / dnf localinstall`
	
	`ACC_SKIP_CAPS=true rpm -vi agent-client-collector-<version number>-x86_64.rpm`

Debian	`ACC_SKIP_CAPS=true apt-get install`
	
	`ACC_SKIP_CAPS=true dpkg -i agent-client-collector-<version number>-<distro>_amd64.deb`

SLES	`ACC_SKIP_CAPS=true zypper install`

Procedure

1. Download the relevant installation packages.
   - For manual installation:
     a. Navigate to > **Agent Client Collector** > **Deployment** > **Agent Downloads**.
     b. Download the relevant .rpm or .deb installation file.
     c. Download the relevant signature files to validate the installation files.
   - For command-line installation, run the following commands to download both the signature and installation files to your local machine:

```
curl -LO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-<version_number>-x86_64.rpm
curl -LO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-<version_number>-x86_64-rpm-rpm.zip
```

   Note: Each curl command must appear on a single line.

2. **Optional:** Verify the package signature.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Extract the installation file (if it is compressed) by running the following command: `unzip agent_client-collector-<version_number>-x86_64-rpm-rpm.zip`

b. Validate the installation file signature by running the indicated commands.

- On an RPM-based system:

```
openssl dgst -sha256 -verify {<ServiceNow
 DGST pem key>} -signature {<signature file>}
 agent-client-collector-<version number>-x86_64.rpm
```

Where `<ServiceNow DGST pem key>` is the `.pem` file extracted from the `.zip` file, and `<signature file>` is the `.bin` file extracted from the `.zip` file.

For example:

```
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-x86_64.rpm
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-x86_64-rpm-rpm.zip
$ unzip agent-client-collector-3.0.0-x86_64-rpm-rpm.zip
Archive: agent-client-collector-3.0.0-x86_64-rpm-rpm.zip
inflating: ServiceNow_Digicert_DGST.pem
extracting: agent-client-collector-3.0.0-x86_64.bin
$ openssl dgst -sha256 -verify ServiceNow_Digicert_DGST.pem
 -signature agent-client-collector-3.0.0-x86_64.bin
 agent-client-collector-3.0.0-x86_64.rpm
Verified OK
```

Note: Each command must appear on a single line.

- On a Debian-based system:

(Optional) `gpg --import ServiceNow_Digicert_Public.gpg`

`sudo gpg --verify agent-client-collector-<version number>-<distro>_amd64.deb`

For example:

```
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-debian-9_amd64.deb
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-debian-9_amd64-deb-deb.zip
$ unzip
 agent-client-collector-3.0.0-debian-9_amd64-deb-deb.zip
Archive:
 agent-client-collector-3.0.0-debian-9_amd64-deb-deb.zip
extracting: ServiceNow_Digicert_Public.gpg
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
$ gpg --import ServiceNow_Digicert_Public.gpg
gpg: /home/admin/.gnupg/trustdb.gpg: trustdb created
gpg: key 985DD52C6A0ABB45: public key "ServiceNow, Inc.
 (Signing) <seceng@servicenow.com>"
imported
gpg: Total number processed: 1
gpg: imported: 1
$ dpkg-sig --verify
 agent-client-collector-3.0.0-debian-9_amd64.deb
Processing
 agent-client-collector-3.0.0-debian-9_amd64.deb...
GOODSIG _gpgbuilder 9B928FB49771DF6C047430DD985DD52C6A0ABB45
 1665054068
```

Note: Each command must appear on a single line.

3. Install the Agent Client Collector package using the package manager associated with Linux distribution.

OS	Command
RHEL-based	`yum / dnf localinstall`
SLES	`zypper install`
Debian-based	`apt-get install`

Alternatively, if these commands aren't configured correctly, you can use the core commands that are configured to run with the package manager commands.

- RPM-based system: `# rpm -vi agent-client-collector-<version number>-x86_64.rpm`
- Debian-based system: `# dpkg -i agent-client-collector-<version number>-<distro>_amd64.deb`

Verify that the package commands are configured correctly with your system administrator.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Some file systems may have restrictions enabled; for example, `/var/` may be mounted with a *`noexec`* flag. Because the agent must execute Agent Client Collector plugins that are normally stored in the `/var/cache` directory, you must deploy the application into specific folders by customizing the installation paths using the `--relocate` option as an `.rpm` parameter.

For example: `rpm -i --relocate /var/cache=/opt/cache agent-client-collector-<version_number>-x86_64.rpm`

The following paths can be relocated:

Path	Notes
`/etc`	When updating, you must also update the *`allow-list`* parameter in the `acc.yml` file with the new path.
`<usr share location>/usr/share`	N/A
`<cache dir location>/var/cache`	Updating the `/var` directory retains all `/var` subdirectories, nested under the new directory.
`/var/log`	
`/var/run`	
`/var`	

After relocating an installation path, update the relevant `ACC_UPGRADE_CMD` string described in step 7.

Review the paths in `/usr/lib/systemd/system/acc.service` to ensure that they appear as expected.

4. When installing a .deb package, configure the agent's `acc.yml` configuration file.

   a. Copy the sample configuration file:

   ```
   # cp -p /etc/servicenow/agent-client-collector/acc.yml.example /etc/servicenow/agent-client-collector/acc.yml
   ```

   b. Rename the allow list file.

   ```
   # cp -p /etc/servicenow/agent-client-collector/check-allow-list.json.default /etc/servicenow/agent-client-collector/check-allow-list.json
   ```

   Note: This step is not relevant for .rpm packages, which come with `acc.yml` and `check-allow-list.json` files included in the base system.

5. Update the configuration file, adding *`check-allow-list.json`* to `/etc/servicenow/agent-client-collector` and copying the *`backend-url`* and *`api-key`* from the instance.
   For example:

   ```
   ---
   # Agent Client Collector configuration
   ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
backend-url:
 - "wss://YOUR_MID_ENDPOINT_HERE:YOUR_MID_PORT_HERE/ws/events"
api-key: "YOUR_API_KEY_HERE"
log-level: "info"
insecure-skip-tls-verify: false
allow-list: /
etc/servicenow/agent-client-collector/check-allow-list.json
verify-plugin-signature: true
max-running-checks: 10
disable-sockets: true
disable-api: true
statsd-disable: true
enable-auto-mid-selection: false
agent_cpu_threshold:
 cpu_percentage_limit: 25
 repeated_high_cpu_num: 3
 monitor_interval_sec: 60
 agent_cpu_threshold_disabled: false
```

The *allow-list* feature, indicating the commands permitted to be executed by the agent, is enabled.

6. Configure sudoers.

   Configuration is typically automated by your Linux sysadmin. To verify configuration correctness, configure sudoers manually before full configuration. To do so, create a new sudoers file for the agent service user by running the following command:

   `visudo -f /etc/sudoers.d/01_servicenow`

7. Add the sudoers configuration according to your Linux distribution.

   When sudo commands are executed as part of Agent Client Collector checks, the system adds the `-E` flag to preserve the environment variables. Commands are executed as `- sudo -E <command_placeholder>` in the sudoers file. Add the *SETENV* command for the Agent Client Collector service user as shown in the code snippet below. If the *Application Patterns* for Agent Client Collector is enabled, ensure sudoers is updated with privileged commands, see Discovery commands for probes and patterns and SSH credentials.

   - When installing agents on a deb/ubuntu system, run the following:

     ```
     User_Alias ACC_USERS = servicenow
     Cmnd_Alias ACC_CMD = /usr/sbin/dmidecode -s
      baseboard-serial-number,/usr/sbin/dmidecode
      -s chassis-serial-number,/usr/sbin/dmidecode
      -s system-serial-number,/usr/sbin/dmidecode -s
      system-uuid,/usr/sbin/ss -tanp,/usr/bin/systemctl start
     acc,/usr/bin/systemctl stop acc,/usr/bin/dpkg --install
      --refuse-downgrade
      --skip-same-version /
     var/cache/servicenow/agent-client-collector/upgrade/agent-cli
     ent-collector-upgrade*
     Cmnd_Alias ACC_CMD_SETENV = /usr/bin/netstat
      -ltnup,/usr/bin/ls -l /proc/*,/usr/bin/cat /proc/*
     ACC_USERS ALL = (root) NOPASSWD:ACC_CMD
     ACC_USERS ALL = (root) NOPASSWD:SETENV:ACC_CMD_SETENV
     Defaults:ACC_USERS !requiretty
     ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Cmnd_Alias ACC_UPGRADE_CMD = /usr/bin/systemd-run
 --unit=acc-upgrade
 --collect /
usr/share/servicenow/agent-client-collector/embedded/bin/ru
by /var/cache/servicenow/agent-client-collector/acc-f-
commons/bin/linux_run_upgrade.rb /usr/bin/dpkg\ --install\
 --refuse-downgrade\
 --skip-same-version /
var/cache/servicenow/agent-client-collector/upgrade/agent-cli
ent-collector-upgrade.deb /usr/bin/systemctl\ start\
 acc /usr/bin/systemctl\ stop\ acc /usr/bin/systemctl\
 daemon-reload /
var/log/servicenow/agent-client-collector/upgrade.log *
ACC_USERS ALL = (root) NOPASSWD:ACC_UPGRADE_CMD
```

- When installing agents on an RPM system, run the following:

```
User_Alias ACC_USERS = servicenow
Cmnd_Alias ACC_CMD = /usr/sbin/dmidecode -s
 baseboard-serial-number,/usr/sbin/dmidecode
 -s chassis-serial-number,/usr/sbin/dmidecode
 -s system-serial-number,/usr/sbin/dmidecode -s
 system-uuid,/usr/sbin/ss -tanp,/usr/bin/systemctl start
acc,/usr/bin/systemctl stop acc,/usr/bin/rpm
 -Uv /
var/cache/servicenow/agent-client-collector/upgrade/agent-cli
ent-collector-upgrade.rpm
Cmnd_Alias ACC_CMD_SETENV = /usr/bin/netstat
 -ltnup,/usr/bin/ls -l /proc/*,/usr/bin/cat /proc/*
ACC_USERS ALL = (root) NOPASSWD:ACC_CMD
ACC_USERS ALL = (root) NOPASSWD:SETENV:ACC_CMD_SETENV
Defaults:ACC_USERS !requiretty

Cmnd_Alias ACC_UPGRADE_CMD = /usr/bin/systemd-run
 --unit=acc-upgrade
 --collect /
usr/share/servicenow/agent-client-collector/embedded/bin/ru
by /var/cache/servicenow/agent-client-collector/acc-f-
commons/bin/linux_run_upgrade.rb /usr/bin/rpm\
 -Uv /
var/cache/servicenow/agent-client-collector/upgrade/agent-cli
ent-collector-upgrade.rpm /usr/bin/systemctl\ start\
 acc /usr/bin/systemctl\ stop\ acc /usr/bin/systemctl\
 daemon-reload /
var/log/servicenow/agent-client-collector/upgrade.log *
ACC_USERS ALL = (root) NOPASSWD:ACC_UPGRADE_CMD
```

8. Configure the agent to run as a service.

a. Safeguard resource consumption by adding the indicated values to the `/usr/lib/systemd/system/acc.service` file.

- CPUShares=128
- CPUQuota=10%
- MemoryLimit=192M

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- BlockIOWeight=10
- LimitNICE=15

For example:

```
# vi /usr/lib/systemd/system/acc.service
[Unit]
Description=Agent-Now acc
After=network-online.target
[Service]
Environment=AGENT_ROOT=/usr/share
Environment=AGENT_CACHE_ROOT=/var/cache
Environment=AGENT_CONFIG_ROOT=/etc
Environment=AGENT_LOG_ROOT=/var/log
Environment=AGENT_RUN_ROOT=/var/run
Environment=RUBYOPT=-Eutf-8
User=servicenow
Group=servicenow
ExecStart=/usr/share/servicenow/agent-client-collector/bin/a
cc-service start acc
KillMode=process
Restart=on-failure
RestartSec=1min
CPUShares=128
CPUQuota=10%
MemoryLimit=192M
BlockIOWeight=10
LimitNICE=+15
[Install]
WantedBy=network-online.target
```

b. Run the reload daemon command If you modified the service file after enabling the acc service.

```
# systemctl daemon-reload
```

c. Enable and start the service by running the following commands.

```
# systemctl enable acc
```

```
# systemctl start acc
```

Validate Linux installation

After installing the Agent Client Collector on a Linux system, validate the installation by ensuring it was completed properly.

Before you begin

Install the Agent Client Collector on a Linux system (see Install Agent Client Collector on a Linux system).

Role required: agent_client_collector_admin

Procedure

1. Run the `# systemctl status acc` command.
   The *`active (running)`* indicator appears when the service is running properly.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
# systemctl status acc
● acc.service - Agent-Now ACC
Loaded: loaded (/usr/lib/systemd/system/acc.service; enabled;
 vendor preset: disabled)
Active: active (running) since Mon 2022-04-25 14:10:24 UTC; 6
 months 23 days ago
Main PID: 2693 (acc)
Tasks: 10 (limit: 512)
Memory: 143.4M (limit: 192.0M)
CPU: 12h 57min 50.442s
CGroup: /system.slice/acc.service
 └─2693 /usr/share/servicenow/agent-client-collector/bin/acc
 start --config-file
/etc/servicenow/agent-client-collect...
```

2. View the encrypted *api-key* and *agent-key-id* in use by running the `grep -E '^api-key|agent-key-id' /etc/servicenow/agent-client-collector/acc.yml` command.

```
# grep -E
 '^api-key|agent-key-id' /
etc/servicenow/agent-client-collector/acc.yml
api-key:
 "encrypted:adskfjldfjlasd+ajflksadjmoreunreadablestringasdfdfk
djksdjaasdfas"
agent-key-id:
 b4ac2ab4ac2b4ac2b4ac2b4ac2b4ac2b4ac2b4ac2b4ac2b4ac2b4ac2b4ac2
b4a
```

Note: The *grep* command must appear on a single line.

3. View the agent's connection with its backend URL by running the `grep 'successfully connected' /var/log/servicenow/agent-client-collector/acc.log` command.

```
# grep 'successfully
 connected' /var/log/servicenow/agent-client-collector/acc.log
2022-11-17T19:03:17.94 [INFO] [agent] successfully connected to
 the url:
wss://MY_BACKEND_FQDN:MY_BACKEND_PORT/ws/events
```

Note: The *grep* command must appear on a single line.

4. Verify the unique agent ID by running the `# echo $(hexdump -e '/1 "%02x"' /var/cache/servicenow/agent-client-collector/agent_now_id)` command.

```
# echo $(hexdump -e '/1
 "%02x"' /
var/cache/servicenow/agent-client-collector/agent_now_id)
14b4443df6d2394d
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The indicated agent ID appears on the Agents page in an instance (**All** > **Agent Client Collector** > **Agents**) with **Status = Up** and **Host data collection = Collected** (both appearing in green).

Linux agents - green status

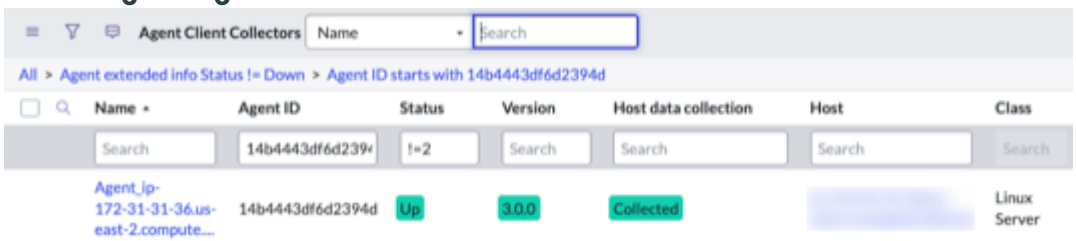

Upgrade Agent Client Collector on a Linux system

Upgrade your existing Agent Client Collector version on a system running a Linux OS if the single-line command script is not connected to the instance or you want to use enhanced customization options.

Before you begin

- The upgrade procedure is the same for both MID Server and MID-less deployments.
- Enable golden image mode for cloning additional instances by setting the golden image marker located at `/tmp/acc-goldenimage`. The golden image marker is applied only during a Linux upgrade and not during new installations.
- Role required: agent_client_collector_admin

Procedure

1. **Optional:** Back up sensitive files by running the following command.

```
tar
 -Pcvzf /root/acc_backup.tar.gz /
etc/servicenow/agent-client-collector \
/var/cache/servicenow/agent-client-collector/agent_now_id \
/usr/lib/systemd/system/acc.service
```

2. Download the .rpm/.deb agent package and signature files, as described in Install Agent Client Collector on a Linux system.
3. Upgrade the package.
   - In an RPM-based system:

     ```
     rpm -Uv agent-client-collector-<version number>-x86_64.rpm
     ```

   - In a DEB-based system:

     ```
     dpkg -i agent-client-collector-<version number>-
     <distro>_amd64.deb
     ```

4. Restore the systemd service file (if you customized the file), using either an automation tool or by reloading systemd configuration.

```
tar -C / -Pxvzf /root/acc_backup.tar.gz
 usr/lib/systemd/system/acc.service
```

```
systemctl daemon-reload
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. Change the *owner:group* setting to the correct one by running the following command.

```
chown -R
 servicenow:servicenow /
{etc,var/cache,var/log,usr/share}/servicenow/agent-client-colle
ctor
```

6. Restart the service and verify that it connects to its backend-url.

```
systemctl restart acc
tail -f /var/log/servicenow/agent-client-collector/acc.log
```

Uninstall Agent Client Collector from a Linux system manually

Uninstall the Agent Client Collector from a Linux machine manually if the command script is unavailable due to the specific agent not being connected to the instance.

Before you begin

- The uninstall procedure is the same for both MID Server and MID-less deployments.
- Role required: agent_client_collector_admin

About this task

By default, an uninstalled agent is removed from the system after 30 days. To modify this amount of time, update the *sn_agent_ci_extended_info* table, accessible from the Auto Flushes page (enter *sys_auto_flush.LIST* into the instance's search filter field).

Procedure

1. Stop the service.
   `sudo systemctl stop acc`
2. Use the relevant package management tool to remove the service.
   - yum remove agent-client-collector
   - apt-get remove agent-client-collector
   - zypper remove agent-client-collector

   Alternatively, you can run one of the following commands, based on the OS in use.

   - `rpm -e agent-client-collector`
   - `dpkg -r agent-client-collector`
3. If you ran either the `rpm` or `dpkg` command, remove any remaining files from your system.

```
# userdel -r servicenow
# groupdel servicenow
# rm
 -R /
{etc,var/cache,var/log,usr/share}/servicenow/agent-client-colle
ctor
```

Warning: Do not use the command rm -Rf, as this can cause data loss.

Perform a single-line Agent Client Collector installation on Linux

Use an efficient single-line command script to install Agent Client Collector on a machine that uses a Linux operating system. If a script is not connected to the instance or you want to use enhanced customization options, you might have to install Agent Client Collector manually.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that you have configured the Agent Client Collector web server.
- Verify that your server's OS and version matches the list of supported OS versions as described in Agent Client Collector Installation.
- Ensure that the unzip program is installed on your Linux machine.
- Determine the resource limits for the agent to be configured in the `/usr/lib/systemd/system/acc.service` file after installation.
- Ensure that you have the `openssl` and `dgst` (for RHEL/Centos/SUSE installation) packages available, to be used by the install script.
- Ensure that you have the `gpg` package available for validation (for Ubuntu/Debian installation).
- Enable golden image mode for cloning additional instances by setting the golden image marker located at `/tmp/acc-goldenimage`. The golden image marker takes no action during new Linux installations, as there is nothing which requires cleaning.
- When performing single-line installation on more than one agent, download the agent package to your local machine. Place the package on your *ftp* server, using an Ansible script. For details on mass deployment of agents, see Deploy the Agent Client Collector using a custom base image.
- In order for the Agent to retrieve the OS serial numbers and TCP connections along with associated running processes, sudo access for “dmidecode” and “ss” is required on Linux systems.

Role required: agent_client_collector_admin

Procedure

1. Retrieve MID Server websocket URL parameters.
   a. Navigate to **Agent Client Collector** > **Websocket Endpoint** and select the relevant websocket endpoint.
   b. Copy the value of the **Endpoint URL** field to be used in the installation command.
2. Copy and edit the following command line, substituting your own information for the command parameters.

```
ACC_API_KEY=<api_key> ACC_MID=wss://<mid_ip>:<websocket_port>/ws/events bash -c "$(curl -L https://<instance URL>/api/sn_agent/agents/install_agent)"
```

Command parameters

Parameter	Value
ACC_API_KEY	The API key used to access the MID Server websocket. If no ACC_API_KEY value is specified, you will be asked to enter the API key value during execution of the script.
ACC_MID	The URL of the MID Server websocket. If you are specifying multiple MID Servers, enter a string separated by commas for each MID Server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. **Optional:** If you need to create new check definitions to be included in the allow list that indicates the checks enabled to run on the agent, disable the allow list by adding *ACC_ALLOW_LIST=0* to the single-line command string.

   Note: This action compromises your system's security and is recommended only as a temporary measure until you finish updating the allow list. After completing the update, re-activate the allow list by removing *ACC_ALLOW_LIST=0* from the single-line command string.

4. **Optional:** If the certificate validation process is too time consuming or if you are using the agent for development purposes, you can bypass certificate validation by adding *ACC_VALIDATE_SIG=0* to the single-line command string.
5. Run the single-line command (specified in step 2, above) on the server where the Agent Client Collector is installed.
6. **Optional:** Configure resource limits for the agent in the `/usr/lib/systemd/system/acc.service` file, according to your business needs.

   (Optional) The default values in the file are:

   - *CPUShares*: 128
   - *CPUQuota*: 10%
   - *MemoryLimit*: 192M
   - *BlockIOWeight*: 10
   - *LimitNICE*: +15

Upgrade the Agent Client Collector manually on a Linux system

Perform a manual upgrade of your existing Agent Client Collector version on a system running a Linux OS if the single-line command script is not connected to the instance or you want to use enhanced customization options.

Before you begin

- Ensure that you have configured the Agent Client Collector web server. For more information, see Configure the websocket server on the MID Server.
- Restart the MID Web Server.
- Collect host data, to ensure that all files are up to date. For details, see Collect data from your system devices.
- Ensure that you have installed a Linux version of the Agent Client Collector. For more information, see Agent Client Collector installation on a Linux OS system.
- The upgrade procedure is the same for both MID Server and MID-less deployments.
- Role required: agent_client_collector_admin

Procedure

1. Back up the `agent-id` and configuration files.

```
mkdir -p /tmp/acc-upgrade-backup
sudo cp
 -rp /etc/servicenow/agent-client-collector/acc.yml /
tmp/acc-upgrade-backup
sudo cp
 -rp /
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
etc/servicenow/agent-client-collector/check-allow-list.json /
tmp/acc-upgrade-backup
sudo cp
 -rp /
var/cache/servicenow/agent-client-collector/agent_now_id /
tmp/acc-upgrade-backup
```

2. Uninstall the package.

   a. Stop the service.
   `sudo systemctl stop acc`

   b. Disable the service.
   `sudo systemctl disable acc`

   c. Uninstall the package:
   - For RHEL/CentOS/SUSE: `sudo rpm -e agent-client-collector`
   - For Ubuntu/Debian: `sudo dpkg -r agent-client-collector`

3. Download the installation file.
   For example, when working with an Ubuntu/Debian OS, use both of the following:

   ```
   curl -L -O https://<MID-WEBSERVER-URI:PORT>/static/
   acc_installers/agent-client-collector/glide/distribution/builds/
   package/app-signed/agent-client-collector-<version number>-
   <Ubuntu/Debian version>_amd64-deb-deb.zip -O
   ```

   ```
   https://<MID-WEBSERVER-URI:PORT>/static/acc_installers/agent-
   client-collector/glide/distribution/builds/package/app-signed/
   agent-client-collector/agent-client-collector-<version number>-
   <Ubuntu/Debian version>_amd64.deb
   ```

4. Verify and install the package.

   a. Unzip the compressed file.

   ```
   unzip agent-client-collector-2.3.0-debian-9_amd64-deb-deb.zip
   ```

   b. Validate the installer file signature.
   - RHEL/Centos/SUSE OS:

     ```
     openssl dgst -sha256 -verify {<ServiceNow DGST pem key>}
     -signature {<signature file>} agent-client-collector-<version
     number>-x86_64.rpm
     ```

   - Ubuntu/Debian OS:

     ```
     gpg --import ServiceNow_Digicert_Public.gpg
     dpkg-sig --verify agent-client_collector-<version
      number>-<Ubuntu/Debian version>_amd64.deb
     ```

   c. Install the package.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- RHEL/CentOS/SUSE: `sudo rpm -vi --force agent-client-collector-<version number>-x86_64.rpm`
- Ubuntu/Debian: `sudo dpkg -1 agent-client-collector-<version number>-<Ubuntu/Debian version>_amd64.deb`

5. Restore the `agent-id` and configuration files.

```
sudo -u servicenow cp
 -rp /tmp/acc-upgrade-backup/acc.yml /
etc/servicenow/agent-client-collector/acc.yml
sudo -u servicenow cp
 -rp /tmp/acc-upgrade-backup/check-allow-list.json /
etc/servicenow/agent-client-collector/check-allow-list.json
sudo -u servicenow cp
 -rp /tmp/acc-upgrade-backup/agent_now_id /
var/cache/servicenow/agent-client-collector/agent_now_id
```

6. Configure the agent to run as a service.
   `sudo systemctl enable acc`
7. Start the service.
   `sudo systemctl start acc`
8. View the logs to verify that the startup was successful.
   `sudo tail -f /var/log/servicenow/agent-client-collector/acc.log`

Uninstall Agent Client Collector from a Linux system using a single-line command

Uninstall the Agent Client Collector from a Linux machine by running an efficient single-line command. If the script is not connected to the instance, you might have to uninstall Agent Client Collector manually.

Before you begin

- The uninstall procedure is the same for both MID Server and MID-less deployments.
- Role required: agent_client_collector_admin

Procedure

Uninstall Agent Client Collector by running the following command in a command window where `<instance url>` is the URL of the ServiceNow instance.
`bash -c "$(curl -L https://<instance url>/api/sn_agent/agents/install_agent)" -s "--remove"`.

Note: Uninstalling the Agent Client Collector removes the `acc.yml` file and its directory from your machine.

Agent Client Collector installation on a Windows machine

When installing the Agent Client Collector on a Windows machine, either download an installation file and use a wizard to install the agent manually, or use silent installation to automate agent installation and configuration. Manual installation enables you to test the agent on a single system, while silent installation enables you to deploy the agent at scale.

Install the Agent Client Collector on a Windows machine manually

Install the Agent Client Collector on a Windows machine manually when you want to test the agent on a single system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

1. Ensure that you've configured the Agent Client Collector web server. For details, see Configure the websocket server on the MID Server.
2. Verify your server's OS and version.
3. When using Microsoft Internet Explorer, do the following to prepare the Agent Client Collector installation file:
   a. Navigate to **Server Manager** > **Local Server**.
   b. Turn off **IE Enhanced Security Configuration**.
   c. Download the installation file.
   d. Change the extension of the installation file from `.man` to `.msi`
4. Ensure that .NET Framework 4.6.2 or later is installed on the target machine.

   Note: If the .NET Framework is not running effectively, use the legacy VBSript installer by running the following command on the target machine:
   ```
   msiexec /i <msi_file_path> /quiet /qn /norestart
   USE_VBS_CA=1 ACC_API_KEY=<key_value> ACC_MID=wss://
   <mid_ip>:<websocket_port>/ws/events ACC_ALLOW_LIST=False
   ```
5. Enable golden image mode for cloning additional agents by setting the msi property *`GOLDEN_IMAGE=true`*.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **Agent Client Collector** > **Agent Downloads**.
2. To perform installation using the wizard:
   a. Download the .msi file that is accessible from the link on the UI.
   The Installation Wizard opens.
   b. Select **Next** until you reach the **MID Web Server Connection Settings** page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

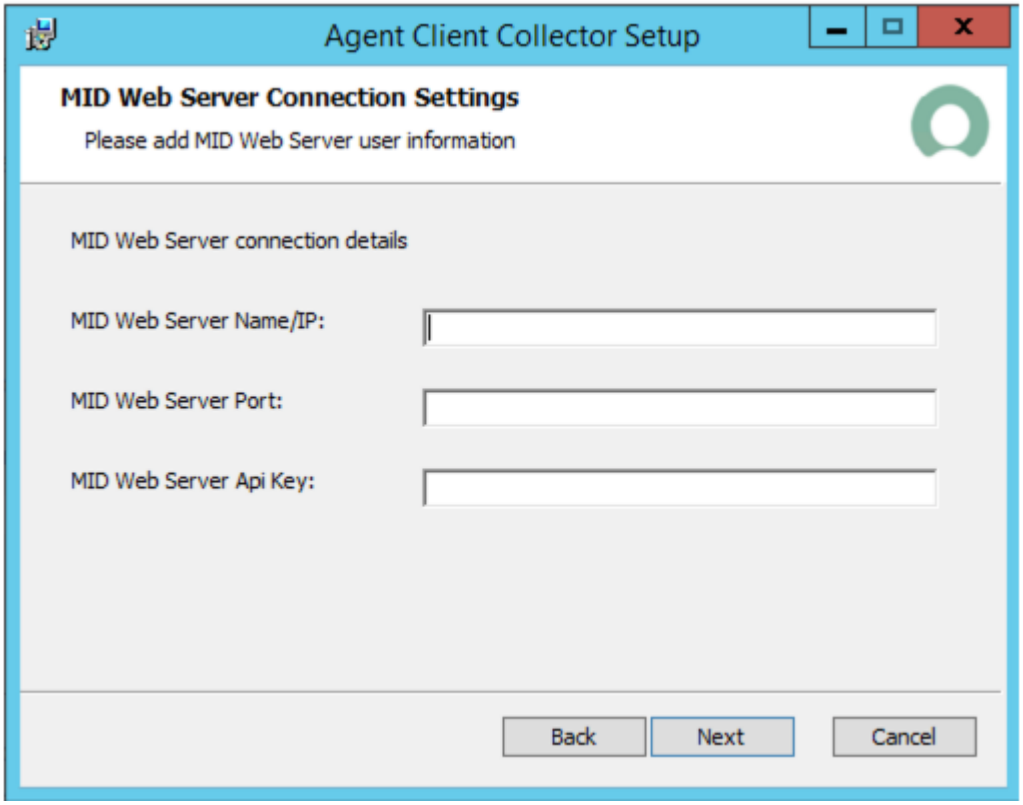

c. Enter the relevant parameters in the dialog box fields:

MID Web Server Connection Settings fields

Field name	Description
MID Web Server Name/IP	IP address of the MID Web server.
MID Web Server Port	Port number of the MID Web server. Use a number between 1024-49151.
MID Web Server API Key	API key to access the MID Web server.

d. Select **Next**.
The **Create Local User Account** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

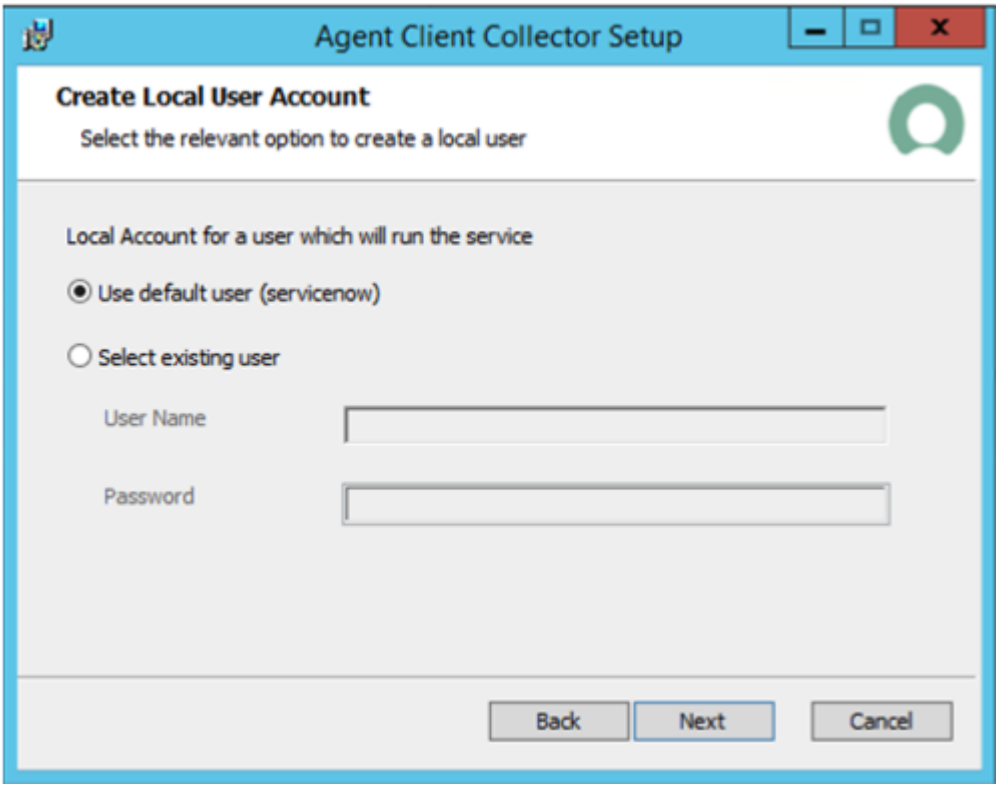

e. Select the relevant option in the dialog box.

Create Local User Account dialog box

Option	Description
Use default user (servicenow)	Default user with which to run the Agent Client Collector on a Windows machine. Selecting this option creates a local user called **servicenow** as part of the installation. The servicenow user password is randomly generated and is not stored in the system. A security event is created on the host machine which, if repeated, requires opening a support case. For details, see KB1649116 article in the HI Knowledge Base.
Select existing user	When selecting this option, configure the following subfields: ▪ **User Name**: User name to log in to the Windows environment. ▪ **Password**: Password of the selected user. The user can be one of the following: ▪ Local user (created by the virtual machine owner; provides added resiliency). ▪ Domain user (created by IT personnel).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	▪ Group Managed Service Account (gMSA): Enter the name of the group managed service account in the **User Name** field and leave the **Password** field empty. ▪ Local Service Account: Enter *LOCALSERVICE* in the **User Name** field and leave the **Password** field empty. ▪ Local System Account: Enter *SYSTEM* in the **User Name** field and leave the **Password** field empty. Due to Windows UAC restrictions, running the agent as a local SYSTEM account is necessary to perform an upgrade via *msiexec*.

Note: Using a Local System account creates the local group *ServiceNow Users*, and adds the created user to the group. The ServiceNow Users group has access to Agent Client Collector directories and has the following privileges:

- **Performance Monitor**: Comes with the base system; enables viewing all performance counters in the system.
- **Log on as a service**: Enables starting network services and services that run continuously, even when no one is logged onto the console.
- **Debug programs**: Enables debugging processes and reading process details for users outside the ServiceNow user group.

f. Select **Create User**.
The system creates a Windows user with the specified credentials, and the Agent Client Collector Configuration page appears.

g. To enable the **Allow-List** containing the checks enabled to run after setup, select the **Enable executing command allow list** check box on the resulting page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

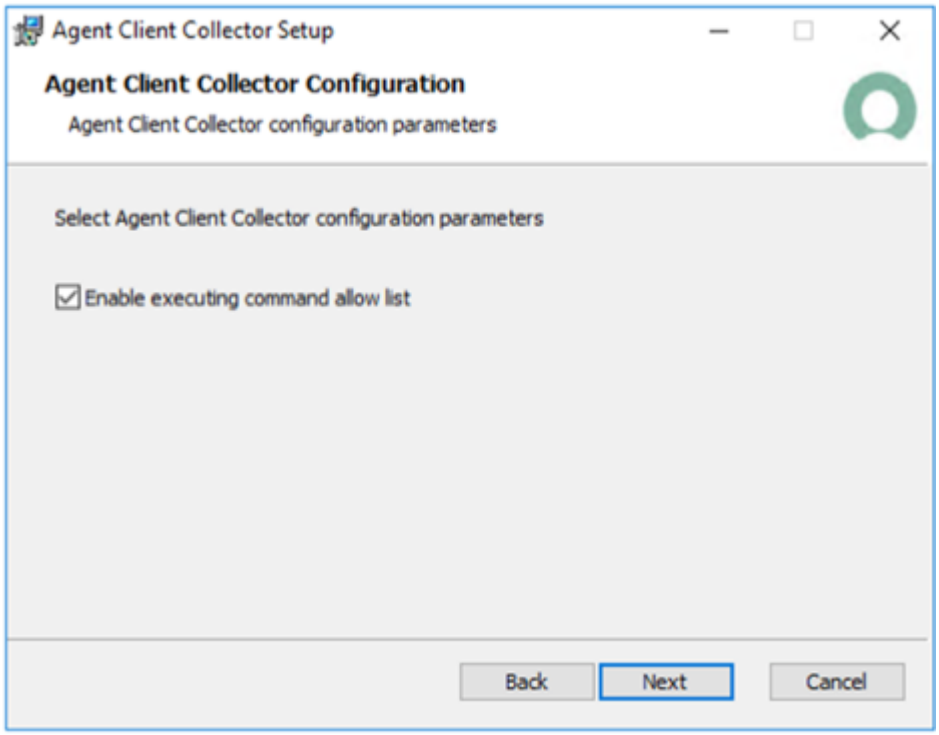

h. Select **Next** to open the installation page.

When installation completes, click **Next** to complete and exit the wizard.

After a successful installation, the service starts automatically and generates the following:

- User: Default = servicenow (an existing user can also be configured to manage the service). The servicenow user belongs to the local user group **ServiceNow Users**.
- Service: AgentClientCollector
- File location:
  - Default for installation = C:\Program Files\ServiceNow\agent-client-collector
  - Default for service files = C:\ProgramData\ServiceNow\agent-client-collector, which holds the cache, config, and log folders.

3. To run Silent installation to disable installation reports, run the command line appearing in the **Silent installation** section of the procedure, as follows:

```
msiexec /i <msi_file_path> /quiet /qn /norestart
ACC_API_KEY=<key_value> ACC_MID=wss://<mid_ip>:<websocket_port>/
ws/events ACC_ALLOW_LIST=False
```

To enable the allow list of specified checks when running Silent installation, set the *ACC_ALLOW_LIST* parameter to **True**.

4. To configure multiple MID Servers with Silent installation, run the command line appearing in the **Silent installation** section of the procedure with comma-separated values, as follows:

```
msiexec /i <msi_file_path> /quiet /qn /
norestart ACC_API_KEY=<key_value> ACC_MID="wss:/
/<mid1_ip>:<websocket1_port>/ws/events,wss://
<mid2_ip>:<websocket2_port>/ws/events"
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Insert the following parameter values in the command:

Command parameters

Parameter	Value
ACC_API_KEY	The API key used to access the MID Server websocket.
ACC_MID	The URL of the MID Server websocket.
DENY_LOCAL_LOGON	When set to *True*, prevents any user other than the default servicenow user from logging onto the service. Default value: False
HTTPS_PROXY	A fully qualified URL to a proxy server. For example: `https://proxy.url.net`
PAC_FILE	A readable filepath on the system. For example: `/path/to/servicenow/readable/file`
PASSWORD_NEVER_EXPIRES	When set to *True*, the password for the default servicenow user never expires. Default value: False
START_SERVICE	When set to *True*, the Agent Client Collector service automatically starts after installation. Default value: True
LOCALUSERNAME	Account running the Agent Client Collector as a Windows service. Possible values: ○ SYSTEM ○ LOCALSERVICE ○ DOMAIN\gMSA$ ○ string (local username) Default = `servicenow` When using the SYSTEM, LOCALSERVICE, or DOMAIN\gMSA$ accounts, you do not have to configure a password.
LOCALPASSWORD	Password for the user specified in the LOCALUSERNAME parameter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The configured MID Server values are automatically added to the `acc.yml` file. The agent connects to the first available MID Server in the list.

To configure a secondary MID server, enter the following syntax in the `backend-url` section of the `acc.yml` file, per the `.yml` standard: `- "wss://<MID IP>:<Web socket port>/ws/events"` (with the quotation marks).

Upgrade the Agent Client Collector manually on a Windows system

Perform a manual upgrade of your existing Agent Client Collector version on a system running a Windows OS.

Before you begin

Enable golden image mode for cloning additional agents by setting the msi property *GOLDEN_IMAGE=true*.

Role required: agent_client_collector_admin

Procedure

Upgrade to the new agent version, as described in Upgrade an agent in an instance.
You can use either the manual or single-line procedure. When restoring backup files, the system replaces the configuration file values.

Uninstall the Agent Client Collector manually from a Windows machine

Uninstall the Agent Client Collector from a Windows machine manually instead of using the single-line procedure. Use the manual procedure if the command script is unavailable due to the specific agent not being connected to the instance.

Before you begin

Role required: agent_client_collector_admin

Procedure

Uninstall the Agent Client Collector in the Windows Control Panel.

Install the Agent Client Collector on a Windows machine using silent installation

Install the Agent Client Collector on a Windows machine using silent installation when you want to deploy the agent at scale.

About this task

Silent installation automatically creates a local user called **servicenow** (unless another existing user is invoked with the *LOCALUSERNAME* parameter, as described in the command parameters table in the following procedure). The servicenow user password is randomly generated and is not stored in the system. A security event is created on the host machine which, if repeated, requires opening a support case. For details, see KB1649116 article in the Now Support Knowledge Base.

Before you begin

ACC installs need to access this URL: `https://install.service-now.com/`.

Enable golden image mode for cloning additional agents by setting the msi property *GOLDEN_IMAGE=true*.

Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Downloads**.
2. From the Windows Download section, download the MSI installer by selecting the Download icon ( ).

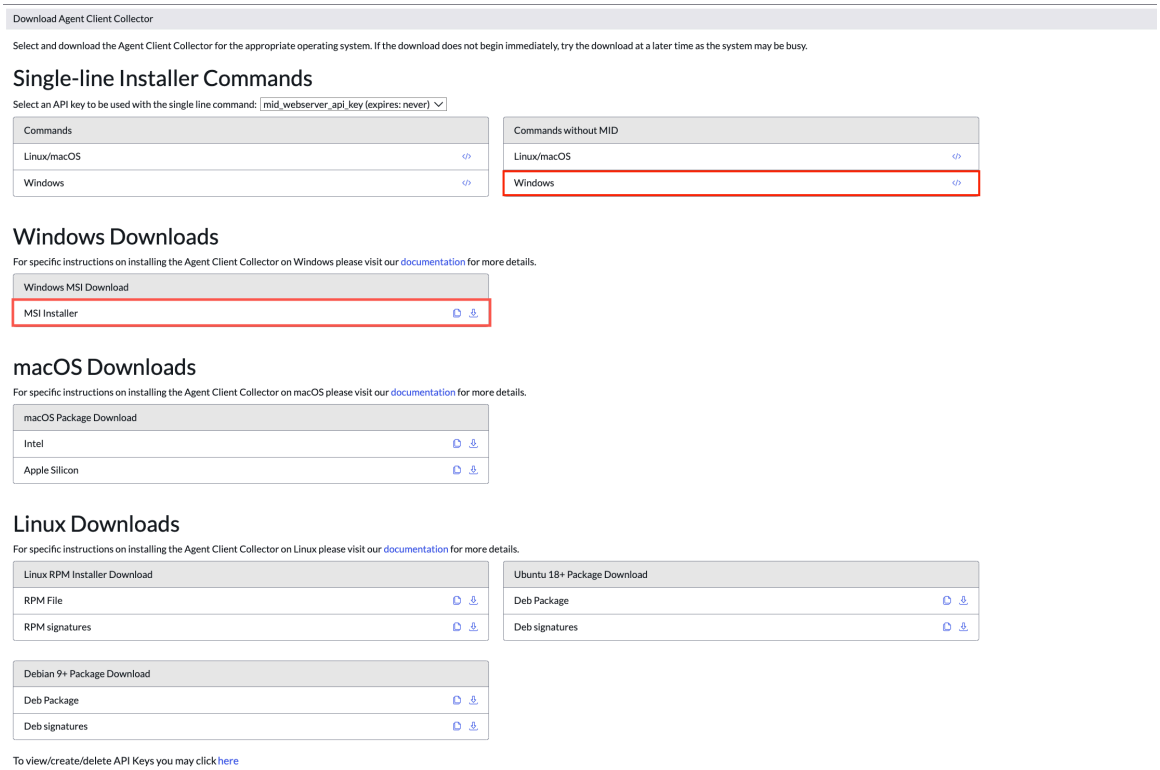

3. To run Silent installation to disable installation reports, run the command line appearing in the **Silent installation** section of the procedure, as follows:
   ```
   msiexec /i <msi_file_path> /quiet /qn /norestart
   ACC_API_KEY=<key_value> ACC_MID=wss://<mid_ip>:<websocket_port>/
   ws/events ACC_ALLOW_LIST=False
   ```

   To enable the allow list of specified checks when running silent installation, set the *ACC_ALLOW_LIST* parameter to **True**.

4. To configure multiple MID Servers with Silent installation, run the command line appearing in the **Silent installation** section of the procedure with comma separated values, as follows:
   ```
   msiexec /i <msi_file_path> /quiet /qn /
   norestart ACC_API_KEY=<key_value> ACC_MID="wss:/
   /<mid1_ip>:<websocket1_port>/ws/events,wss://
   <mid2_ip>:<websocket2_port>/ws/events"
   ```

   Insert the following parameter values in the command:

Command parameters

Parameter	Type	Value
ACC_API_KEY	String	The API key used to access the MID Server websocket.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Type	Value
ACC_MID	String	The URL of the MID Server websocket.
ACC_ALLOW_LIST	Boolean	Indicates whether the allow-list is enabled (True) or disabled (False). Default = True
START_SERVICE	Boolean	Indicates whether the agent starts automatically upon installation (true) or must be started manually (false). Default = True
LOCALUSERNAME	String	Account running the Agent Client Collector as a Windows service. Possible values: ○ SYSTEM ○ LOCALSERVICE ○ DOMAIN\gMSA$ ○ string (local username) Default = `servicenow` When using the SYSTEM, LOCALSERVICE, or DOMAIN \gMSA$ accounts: ○ You do not have to configure a password. ○ You must run the relevant .msi command, as described in Account commands for LocalService, LocalSystem, and gMSA during Windows installation.
LOCALUSERPASSWORD	String	Password used to access the agent. Default = <randomly generated password>
ISDOMAIN	Boolean	Assigns a domain user account to the Windows service. Possible values: ○ 0 - DOMAIN is ignored ○ 1 - DOMAIN is used

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Type	Value
DOMAIN	String	Indicates the user domain. Applies only to domain user accounts (not gMSA)
PASSWORD_NEVER_EXPIRES	Boolean	Indicates whether the user password expires (false) or never expires (true). Default = false
SN_PROG_ROOT	String	Indicates the installation path for the agent's application files. Enclose the path in quotes. For example: SN_PROG_ROOT="D:\ServiceNow\agent-client-collector"

The configured MID Server values are automatically added to the `acc.yml` file. The agent connects to the first available MID Server in the list.

To configure a secondary MID server, enter the following syntax in the `backend-url` section of the `acc.yml` file, per the `.yml` standard: `- "wss://<MID IP>:<Web socket port>/ws/events"` (with the quotation marks).

For the `<MID IP>` value, you can assign either the MID FQDN or CNAME. To enable high availability, use either the Auto MID Selection algorithm, or manage a virtual IP with a load balancer. For more information, see the ITOM Agent Client Collector documentation material [KB1122613] article in the HI Knowledge Base.

Uninstall Agent Client Collector from a Windows system using a single-line command

Uninstall the Agent Client Collector from a Windows machine by running an efficient single-line command. If the script is not connected to the instance, you might have to uninstall Agent Client Collector manually.

Before you begin

Role required: agent_client_collector_admin

Procedure

Uninstall Agent Client Collector by running the following command in a command window. `msiexec /quiet /x <path_to_acc_msi_file>`.

Note: Uninstalling the Agent Client Collector removes the `acc.yml` file and its directory from your machine.

Run Agent Client Collector self-diagnostic tests

Run built-in self-diagnostic tests to validate agents' configuration, connectivity, and system requirements.

Before you begin

Ensure that you have installed the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Agent Client Collector agent: Version 6.0.1 or higher
- Agent Client Collector Framework Store Application Plugin: Version 6.0.1 or higher

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select an agent record.
3. Select **Run Self Test** from the **Related Links** section.
4. Verify that tests execute automatically.
   Test results appear in the Agent Self Test Run (`sn_agent_self_test_run`) table.

   For details on the diagnostic self-tests that the agent runs, see Agent Client Collector diagnostic self tests.

Require a maintenance token for Windows uninstalls

Require using a maintenance token when uninstalling an agent from a Windows device. A maintenance token provides a layer of administrative control so that unauthorized personnel can't perform the uninstall.

Before you begin
Role required: sn_agent.token_admin

Procedure

1. Navigate to the Agent Client Collector Setup wizard, as described in Install the Agent Client Collector on a Windows machine manually.
2. Select the **Require maintenance token for uninstall** check box on the Agent Client Collector Configuration page.
3. Complete the installation and exit the setup wizard.

Note: Alternatively, you can enable requiring a maintenance token by using the following single line command:

```
msiexec /i <msi_file_path> /quiet /qn /norestart
 ACC_API_KEY=<key_value>
 ACC_MID=wss://<mid_ip>:<websocket_port>/ws/events
 ACC_ALLOW_LIST=False UNINSTALLVALIDATION=1
```

Enable uninstall validation during agent upgrade

Configure requiring a maintenance token to uninstall agents during agent upgrade. A maintenance token provides a layer of administrative control so that unauthorized personnel can't perform the uninstall.

Before you begin
Role required: sn_agent.token_admin

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Locate the *sn_agent.upgrade_enable_uninstall_validation* property and set its value to **true**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the property does not exist, select **New** to create it.

Note: If uninstall validation is enabled, setting the system property back to **false** disables the feature on the next version upgrade.

Result
The property is invoked for both selective upgrade and high volume upgrade, and the uninstall feature is enabled on Windows agents.

Create a maintenance token

Create a maintenance token to use when uninstalling an agent from a Windows device. Administrators require maintenance tokens to ensure that unauthorized employees can't perform an uninstall.

Before you begin
Role required: sn_agent.token_admin

About this task
Maintenance tokens are short-lived by design. The default expiration is after 30 minutes. To modify the expiration time:

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Select **New**.
3. Configure the following values for the token's expiration time:

Maintenance token fields

Field	Value
Name	*sn_agent.maintenance_token.expiry_minutes*
Type	integer
Value	<the amount of time, in minutes, before the token expires>

4. Select **Submit**.

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Maintenance Token**.
2. Select **New**.
3. Configure the following values for the token:

Maintenance token fields

Field	Description
Alias	Enter a descriptive name for the maintenance token.
Purpose	Select **uninstall**.
Scope	Select the token's scope:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ **global**: Token applies to all agents. ○ **targeted**: Token applies only to selected agents. When selecting this option, the **Agents** field appears.
Agents	Select the **Unlock agents** ( ) icon and select the agents for the token to apply to. This field appears only when selecting **targeted** as the token scope.

4. Select the **Activate** check box to activate the token.
5. Select **Submit**.

What to do next
Uninstall the agent using the maintenance token, as described in Uninstall an agent using a maintenance token.

Note: Ensure that you use the maintenance token before it expires.

Uninstall an agent using a maintenance token

Uninstall an agent from a Windows device using a maintenance token. Administrators require maintenance tokens to ensure that unauthorized employees can't perform an uninstall.

Before you begin
Enable requiring a maintenance token for uninstalls, as described in Require a maintenance token for Windows uninstalls.

Create a maintenance token to be used when uninstalling an agent, as described in Create a maintenance token.

Role required: sn_agent.token_admin

Procedure
1. Navigate to the Control Panel on your Windows machine and select the Agent Client Collector application to uninstall.
   After the uninstall process begins, the **Uninstall Token Required** dialog box opens.

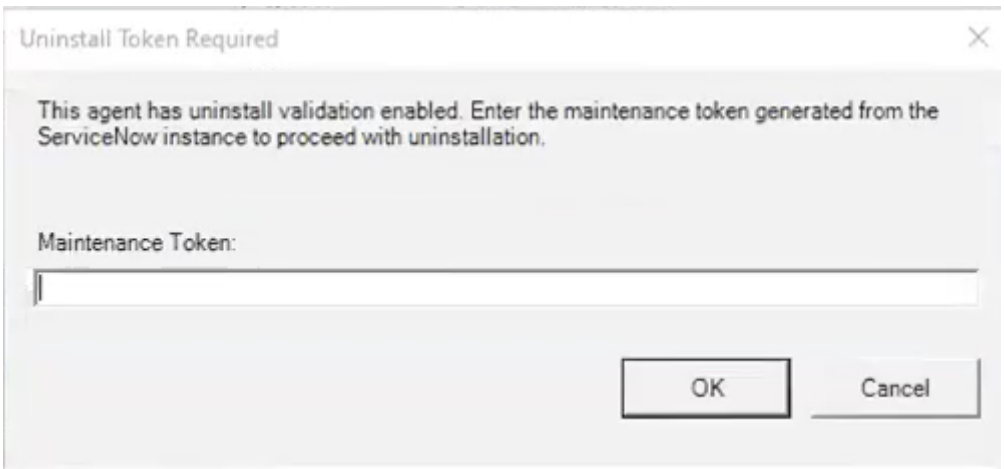

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Retrieve the maintenance token that you created.
   a. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Maintenance Token**.
   b. Select the relevant maintenance token.
   c. Select **View Maintenance Token** in the Related Links section.
   d. Select the **Copy Token** ( ) icon to copy the maintenance token value.
3. Paste the maintenance token in the **Maintenance Token** field.
4. Select **OK**.

Note: Alternatively, you can enable requiring a maintenance token by using the following single line command:

```
msiexec /quiet /x <path_to_acc_msi_file>
 UNINSTALL_TOKEN=<valid_token_value>
```

The agent is uninstalled from your Windows machine.

If you provide an invalid maintenance token, uninstall is blocked. For more information, check the uninstall logs at `<user folder>\AppData\Local\Temp\ACC_Logs`.

Prepare agent deployment on a non-persistent virtual desktop infrastructure machine

Configure the preliminary settings on an instance to enable using the agent with a non-persistent virtual desktop infrastructure machine (NPVDI) machine. NPVDI agents gather data more quickly than traditional agents not enabled for an NPVDI.

Before you begin

Role required: agent_client_collector_admin

About this task

Configurations are set on a ServiceNow instance.

Note: These configurations need to be set only once per VDI setup. You can deploy multiple VDI agents on a single setup.

Procedure

1. Create a registration key marked for light registration.
   a. Navigate to **All** > **System definition** > **Tables**.
   b. Access the *sn_agent_registration_key* table.
   c. Select **New** to create a new registration key.
   d. Set *light_registration* to **true**.
2. Add the static policies you want the VDI agent to run.
   a. Navigate to **All** > **Agent Client Collector** > **Configuration** > **VDI policies**.
   b. Select **New**.
   c. Select the policies to be added as a VDI policy in the **VDI policy** field.
   The available policies are published, non-proxy policies.
   d. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. **Optional:** Add configuration files to the agent.

   a. Navigate to **All** > **Agent Client Collector** > **Configuration** > **VDI configuration files**.

   b. Select **New**.

   c. Select the configuration files to be added to the agent in the **Configuration file** field. Configuration files that aren't mapped to a check or policy are added to the agent.

   d. Select **Submit**.

4. Grant the user the role of *agent_client_collector_user_rest*. These credentials are used when enabling an NPVDI agent.

What to do next

Enable an agent to run as an NPVDI agent, as described in Enable a non-persistent virtual desktop infrastructure agent.

Enable a non-persistent virtual desktop infrastructure agent

Configure an agent to enable it to work in a non-persistent virtual desktop infrastructure (NPVDI) environment. NPVDI agents are self-sufficient and start running checks immediately.

Before you begin

Ensure that you have set preliminary configurations in your ServiceNow instance, as described in Prepare agent deployment on a non-persistent virtual desktop infrastructure machine.

Role required: agent_client_collector_user

About this task

A non-persistent VDI (NPVDI) machine is deployed agent installed on its golden image. When working with an NPVDI agent, data collection begins within two minutes.

NPVDI machines have their file system and storage cleaned after each use. There is no record of previous activity on an NPVDI machine; the only indicator it retains throughout its existence is its FQDN.

Procedure

Install the agent MSI on the NPVDI gold image host using the following parameters:

```
msiexec /i agent.msi NP_VDI_GOLD_IMAGE=true `
  INSTANCE_URL=<instance-url> `
  INSTANCE_REST_API_USER=<basic-auth-user> `
  INSTANCE_REST_API_PASS=<basic-auth-password> `
  REGISTRATION_KEY=<light-registration-key> `
  [PAC_FILE=<pac-url-or-path>] `
  [HTTPS_PROXY=<proxy-url>]
```

For details on the MSI installation parameters, see MSI installation parameters.

The following occurs during installation.

a. The installer validates the parameters.

b. The agent downloads the NPVDI configurations (policies, checks, assets and config files) from the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. The installer registers the agent binary as a Windows group policy (GPO) shutdown script.

d. The installer marks the agent as non-persistent, making it eligible for a golden image.

Agent Client Collector upgrade overview

The Agent Client Collector Framework manages agent upgrades directly from the instance, with no manual action required on individual agent hosts.

You can upgrade Agent Client Collector agents one at a time through the UI, or upgrade all eligible agents at once using the auto-upgrade feature. After an upgrade starts, each agent downloads the new version, installs it, restarts, and reports the result back to the instance.

Before you upgrade

Configure the Agent Client Collector web server. For details, see Configure the websocket server on the MID Server.

Ensure that the MID Server, MID Web Server, and the MID Server websocket server are running.

Roles required: agent_client_collector_admin

- In a Windows environment: Local SYSTEM account
- In a Linux environment: sudo rpm/dpkg
- In a macOS environment: sudo pkg

Upgrade methods

Two upgrade methods are available:

Single agent upgrade

Upgrade one agent at a time from the agent record in the UI. Use this method to validate the upgrade process before a broader rollout. You can also select multiple agents from the list view to upgrade up to 50 agents at a time.

Mass upgrade

Upgrade all eligible agents automatically using the built-in scheduled job or a background script. The scheduled job applies rate limiting to control the rollout pace. The background script also uses rate limiting, and triggers an immediate upgrade without waiting for the next scheduled run.

Upgrade stages

Both single and mass upgrade attempts, progress through the following stages. Each stage creates a record in the Agent Upgrade Histories table.

InstanceVerification

The instance confirms the agent is eligible for upgrade. Checks include operating system support, agent version, agent status, and MID Server connectivity.

AgentVerification

The agent confirms it can perform the upgrade. Checks include permissions, available disk space, and dependencies.

Upgrade

The agent downloads the package, installs it, and restarts with the new version.

For details on the required properties for agent overview, see Agent Client Collector upgrade properties.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

High-volume upgrade does not support Agent Client Collector to MID Server communication via mTLS.

When performing high-volume upgrade, all agents that aren't using the most up-to-date version are upgraded. No upgrade is performed on agents already using the upgraded version.

No upgrade is performed on agents that are outside the application scope. For more information, see Application scope.

An agent is excluded from high-volume upgrade when you reach the failed upgrade limit for an agent. The failed upgrade limit is specified in the *sn_agent.auto_upgrade.retry_limit* system property. The default value for this property is 3.

- This property applies to both high-volume and selective upgrades that have failed.
- When an agent reaches the failed upgrade limit, an Agent Client Collector administrator can still run selective upgrade on the agent.
- You can manually reset the failed upgrade limit by selecting the **Reset failed upgrade attempts** option on the agent page (select **All** > **Agent Client Collector** > **Agents** and select an agent).
- After a successful upgrade, the upgraded agent is considered to have zero failed upgrades (for future upgrades).

Agent eligibility

An agent is eligible for auto-upgrade when it meets all of the following conditions:

- Status is **Up**.
- Version is 2.7.0 or higher and lower than the target upgrade version.
- Agent is not flagged as a duplicate.
- Has a valid MID Server or Pod reference.
- Has not exceeded the retry limit for failed upgrades.
- Runs on a supported operating system.

Package download sources

During upgrade, each agent downloads the new package from one of the following sources, depending on its connection type:

- Agents connected via Cloud Services (ICS) download from the ServiceNow CDN at `https://cdn-install.sncapps.service-now.com`.
- Agents connected via MID Server download from the MID Server directly over the local network.
- If the MID Server is unavailable, agents fall back to the Install Server at `https://install.service-now.com`. To enable agent fallback, ensure that the install server is specified on the ServiceNow instance.

If your agents use a custom download server, configure the *agent-upgrade-url-path* key in the agent's `acc.yml` file. This setting takes priority over the default sources.

Rollout strategy

Use a phased approach to reduce risk when upgrading a large agent fleet:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Upgrade one agent from the UI to confirm the process works end-to-end.
2. Upgrade 5–10 agents from the list view, selecting agents across different operating systems.
3. Perform a full rollout by enabling the scheduled job or running the background script.

Upgrade an agent in an instance

Perform selective self-upgrade instead of bulk upgrade for enhanced efficiency when working with agents that are difficult to access, such as agents deployed in the cloud. You can perform selective upgrade on up to 50 agents at a time.

Before you begin

Note: Selective upgrade is available for the agents version 2.7.0 and above.

Supported Operating Systems: Windows, Linux and macOS.

Ensure that data files are up to date by restarting the MID Server and collecting host data. For details on collecting host data, see Run host data collection on demand.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

Ensure that you have sudo permissions and that the following configurations are set:

Linux Operating Systems

OS	Commands	Agent version
CentOS,Red Hat	• `/usr/bin/systemctl start acc` • `/usr/bin/systemctl stop acc` • `<systemctl location> daemon-reload` (default location is `/usr/bin`) • `/usr/bin/rpm/ systemd-run -Uv <cache dir location>/ upgrade/agent-client-collector-upgrade.rpm` Default cache directory location: `/var/cache/ servicenow/agent-client-collector` • `unzip -o`: Unzips the signature file to validate the installation. • `openssl dgst -sha256 -verify`: Verifies use of a signature from a servicenow installation file. **Note:** If you installed the Linux agent using the `--relocate` flag (for example, `rpm -i --relocate /var/cache=/opt/cache agent-client-collector-<version_number>-x86_64.rpm`), you must have permission to execute the following commands: • `sudo /usr/`	2.7

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When connecting via a MID Server, ACC installs need to access one of the following URLs:

- `https://<MID-WEBSERVER-URI:PORT>/static/acc_installers`
- `https://install.service-now.com/`

When upgrading MID-less agents, upgrade is performed at the URL set in the *agent-upgrade-url-path* property of the acc.yml file. If no URL is set at this property, upgrade is performed by downloading the installer script from the following Content Delivery Network (CDN) URL: `https://cdn-install.sncapps.service-now.com/glide/distribution/builds/package`

Note: When using a CDN, send your download request through a proxy server, for added security. For details, see Configure a download proxy for Agent Client Collector upgrades.

Select the version to which you want to upgrade your agent.

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Select **New**.
3. Assign the following values:
   - Name = **sn_agent.agent_upgrade_version**
   - Type = **string**
   - Value = The version number to be upgraded to, in the format <major_version.minor_version.patch_version>. For example: **4.3.2**
4. Select **Submit**.

If you do not select a version, the agent automatically upgrades to the current Agent Client Collector Framework scoped app version.

Role required:

- Windows: Local SYSTEM account (Agent Client Collector service running as *Local System*)
- Linux: sudo rpm/dpkg/systemd-run
- macOS: sudo pkg

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select the agent that you want to upgrade.
   Verify that the agent has both a version number and **Status=Up**.
3. In the **Related Links** section, select **Upgrade agent**.
4. In the confirmation dialog box, select **Upgrade**.
   The **Agent Upgrade Histories** page appears, where you can monitor the upgrade progress in the **State** column, which displays either **Success** or **Failed**. The full upgrade process takes between 3-10 minutes.
5. View the log for the agent upgrade in the **Message** column.
   - For a failed upgrade on a Linux machine, navigate to the `<cache directory location>/upgrade/agent-client-collector-upgrade.rpm` file. Ensure that the configurations are set for the relevant OS, according to the Linux Operating Systems table, above.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For a failed upgrade on a Windows machine, check the relevant log file, located at `<user folder>\AppData\Local\Temp\ACC_Logs`:
  - ACC logs: `ACC_Upgrade.log`
  - MSI logs: `MSI_ACC_Upgrade.log`
- For a failed upgrade on a macOS machine, check the log file, located at `/Library/Application Support/servicenow/agent-client-collector/log/upgrade.log`.

Related topics

Upgrade multiple agents in an instance

Upgrade multiple agents in an instance

When performing selective upgrade of an agent in an instance, you can select multiple instances to upgrade at once.

Before you begin

Supported Operating Systems: Windows, Linux and macOS.

Agent Client Collector installs need to access this URL: `https://install.service-now.com/`. MID connected agents can download the installer from the MID Server (`https://<MID-WEBSERVER-URI:PORT>/static/acc_installers`).

Select the version to which you want to upgrade your agent.

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Select **New**.
3. Assign the following values:
   - Name = **sn_agent.agent_upgrade_version**
   - Type = **string**
   - Value = The version number to be upgraded to, in the format <major_version.minor_version.patch_version>. For example: **4.3.2**
4. Select **Submit**.

If you do not select a version, the agent automatically upgrades to the current Agent Client Collector Framework scoped app version.

When working with MID-less agents, ensure that the agents can communicate with the ServiceNow Content Delivery Network (CDN). For details, see Configure a download proxy for Agent Client Collector upgrades.

Role required: agent_client_collector_admin

About this task

You can select up to 50 agents at a time for upgrade. To upgrade large numbers of agents, incorporate the Agent Client Collector into your software deployment tools.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select the check box next to the agents that you want to upgrade.
   You can select up to 20 agents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the **Actions on selected rows...** box, select **Upgrade agent**.
4. In the confirmation dialog box, select **Upgrade**.

Related topics

Upgrade an agent in an instance

Run a mass upgrade via the scheduled Agent Client Collector upgrade job

Activate the built-in scheduled job to upgrade all eligible Agent Client Collector agents automatically in rate-limited batches.

Before you begin

Configure the required upgrade properties before enabling the scheduled job. For details, see Agent Client Collector upgrade properties.

Role required: agent_client_collector_admin

About this task

The **Upgrade agents incrementally** scheduled job upgrades all eligible agents on a recurring schedule. The job respects the rate limits set in the upgrade properties, upgrading agents in batches to avoid overloading MID Servers. Agents that have never failed an upgrade are prioritized over agents with prior failures.

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Search for **Upgrade agents incrementally** and select record.
3. Set **Active** to **true**.
4. **Optional:** Adjust the **Run** schedule to match your maintenance window.
   The default schedule runs every hour.
5. Select **Update**.
   The job runs on the configured schedule, upgrading agents in batches that adhere to the configured rate limits.

Run a mass Agent Client Collector upgrade using a background script

Trigger an immediate upgrade of all eligible Agent Client Collector agents without waiting for the next scheduled job run.

Before you begin

Configure the required upgrade properties before running the script. For details, see Configure Agent Client Collector upgrade properties.

Role required: agent_client_collector_admin

About this task

The background script upgrades agents connected via MID Server and agents connected via Cloud Services (ICS) in a single run. Both function calls in the script are required — omitting either skips the corresponding agent group.

The system queries all agents that are **Up**, not duplicated, have a valid MID Server or Pod connection, and have not exceeded the retry limit. Agents with no prior failures are upgraded before agents with prior failures.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Definition** > **Scripts - Background**.
2. Change the application scope to `sn_agent` (Agent Client Collector).
3. Paste the following script into the editor.

```
var duration = 60; // Duration in minutes for rate limiting
 calculations

var a = new AgentUpgradeUtil();
a.upgradeAgentsLimitedByMid(duration);  // Upgrades agents
 connected via MID Servers
a.upgradeAgentsMidless(duration);       // Upgrades agents
 connected via Cloud Services
```

   Set `duration` to the number of minutes used for rate limiting calculations. The default value of `60` applies the per-hour limits configured in the upgrade properties.

4. Select **Run script**.
   The system begins upgrading eligible agents immediately, adhering to the rate limits set in the upgrade properties.

Repeat high-volume upgrade for failed agents

If high-volume upgrade fails for specific agents, you must clear the problematic agents' history to re-enable upgrade. If the target upgrade version changes, you don't need to clear the agents' history, as the agents upgrade with the next scheduled high-volume upgrade.

Before you begin

Roles required: agent_client_collector_admin

- In a Windows environment: Local SYSTEM account
- In a Linux environment: sudo rpm/dpkg/systemd-run
- In a macOS environment: sudo pkg

Procedure

1. Navigate to the Agent Client Collectors page **All** > **Agent Client Collector** > **Agents**.
2. Select an entry for which the *`Version`* number is not the highest available version.
3. Locate the *`Agent Upgrade Histories`* tab.

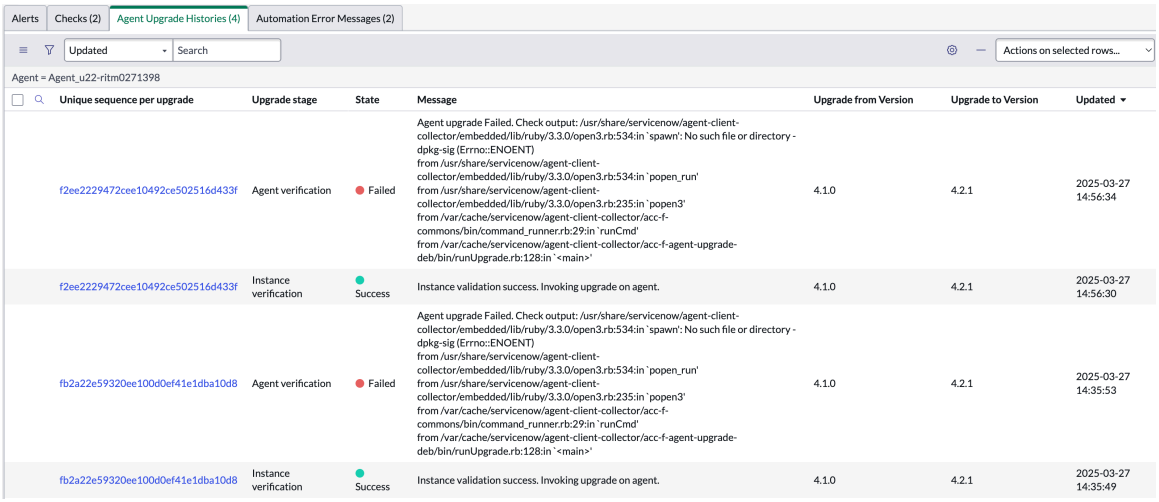

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Underneath the check box ☐, hover next to a failed upgrade entry and select the check box that appears.
5. Select the arrow next to the **Actions on selected rows...** drop-down and select **Delete**.

Result
The specified agents will be upgraded with the next scheduled high volume upgrade.

Configure proxies during MID-less upgrade using a CDN

Create proxies for added security when upgrading MID-less agents via a Content Delivery Network (CDN). You add proxy servers to the sn_agent_proxy table on the ServiceNow instance.

Before you begin
Role required: agent_client_collector_admin

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Package Download Proxies**.
2. Select **New** to create a proxy server.
3. Fill in the fields on the page.

Proxy - New record page fields

Field	Description
Name	A descriptive name to identify the proxy server.
URL	The URL of the proxy server that is accessed when the agent downloads the installer. Enter the URL in the following format: `some.example.com:8800`.
Domain	Select one of the following: ○ **leaf**: Select if domain separation is active, and select the leaf domain for the proxy to be used in. ○ **global**: Select if domain separation is not active.
Order	A number indicating which proxy is used first when the agent downloads the installer from the CDN. Lower numbers are used first.
Active	Select for the URL to be invoked. Clear for the URL to be ignored.

If no proxy server is specified, the default of *`direct`* is invoked, instructing the agent to download the installer directly from the CDN.

Configure a download proxy for Agent Client Collector upgrades

Add a proxy server entry so that agents behind a corporate proxy can download upgrade packages from the ServiceNow Content Delivery Network (CDN).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

This configuration applies only to agents connected via Cloud Services (ICS) that can't reach the CDN directly. Agents connected via MID Server download packages from the MID Server or the Install Server and don't use this proxy configuration.

Role required: agent_client_collector_admin

About this task

Proxy configuration is managed through the Package Download Proxies table. When you add, update, or remove a proxy record, the system reads all active entries sorted by the **Order** field and generates a proxy list. The updated configuration is automatically pushed to all agents via MID Server configuration sync. No manual file editing or agent restart is required.

The agent tries each proxy in order, from the lowest **Order** value to the highest. If a proxy fails, the agent moves to the next entry. To configure a direct connection as a fallback, add an entry with the URL set to `direct`.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Package Download Proxies**.
2. Select **New**.
3. Fill in the proxy record fields.

Package Download Proxies form

Field	Description
Name	Descriptive label for this proxy entry, for example US - `East Corporate Proxy`.
URL	Proxy address in `host` : `port` format, for example `proxy` . `example` . `com` : 8080. Enter `direct` to configure a direct connection with no proxy.
Order	Numeric priority. Lower numbers are tried first. Default is 100.
Active	Selected by default. Clear to disable this entry without deleting it.
Domain	The domain this proxy applies to. Use for domain-separated instances.

4. Select **Submit**.
   The system generates an updated proxy list and pushes it to all agents via MID Server configuration sync. Changes take effect without an agent restart.
5. **Optional:** Add additional proxy entries to configure failover.
   Repeat the preceding steps for each proxy, assigning a higher **Order** value to each successive entry. The agent tries each proxy in order and falls back to the next if the current one fails.

Example: Proxy with direct connection fallback

The following configuration tries a corporate proxy first. If the proxy is unavailable, the agent connects directly to the CDN.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Proxy failover example

Name	URL	Order	Active
Corporate Proxy	`proxy.example.com:8080`	100	Yes
Direct Fallback	`direct`	200	Yes

What to do next

To temporarily block all CDN downloads, clear the **Active** check box on all proxy records. With no active entries, the system blocks CDN downloads. This is useful during maintenance windows.

Configure Agent Client Collector upgrade properties

Set the system properties that control Agent Client Collector upgrade behavior, including the target version and rate limits for mass upgrades.

Before you begin

Role required: agent_client_collector_admin

About this task

Two properties are required before any upgrade can run. Additional optional properties let you control rate limits, retry behavior, and timeout thresholds.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Set the required upgrade properties.

Required upgrade properties

Property	Value	Description
sn_agent.auto_upgrade.en`true`*d*		Enables the upgrade feature. Set to **true** before running any upgrade, manual or scheduled.
sn_agent.agent_upgrade_v	Target version number, for example `5.0.1`	The version agents upgrade to. If left empty, the system uses the current Agent Client Collector Framework application version.

3. **Optional:** Set optional properties to adjust rate limits and timeout behavior.

Optional upgrade properties

Property	Default	Description
sn_agent.auto_upgrade.max_upgrades_per_hour	1000	Maximum agents upgraded per job run. Lower this value for a more gradual rollout.
sn_agent.auto_upgrades_per_mid	50	Maximum agent upgrades routed through a single MID Server per job run. Prevents overloading individual MID Servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Default	Description
sn_agent.auto_upgrade.retry_limit	3	Number of times the system retries a failed upgrade before permanently skipping the agent.
sn_agent.agent_upgrade_wait_time	30	Minutes the system waits for an agent to complete its upgrade before marking it as failed.
sn_agent.upgrade.timeout_minutes	10	Minutes allowed for the initial upgrade validation check on the agent.
sn_agent.min_upgrade_history_per_agent	2	Number of upgrade history records kept per agent. Older records are automatically removed.

Create and edit Agent Client Collector plugins

Custom plugins extend Agent Client Collector monitoring capabilities beyond the default plugins. You can create plugins for specific monitoring requirements or modify existing plugins.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Put every script or executable the plugin needs inside `bin/`.
   - For Linux and macOS assets, ensure that the agent can execute these files.
   - Make each script/executable executable before building the archive: `chmod +x bin/my_script.rb bin/helper.sh`
   - For Windows plugins that run scripts other than batch files, include a batch file in `bin/` that calls the actual script. For example, a batch file can call a PowerShell script using `powershell %~dp0\test.ps1`.
2. If the plugin requires ACC command allowlist validation, create `allow_list/allow_list.json` at the top level of the plugin folder.
   The file is a JSON array. Each entry defines an allowed executable and the arguments it may receive.

```
[
  {
    "args": [
      "-d",
      "--parent_cache_dir",
      "/var/cache/servicenow/agent-client-collector/"
    ],
    "exec": "my_script.rb",
    "skip_arguments": false
  },
  {
    "args": [
      "--log-level=(FATAL|ERROR|WARN|INFO|DEBUG)",
      "-l",
      "(FATAL|ERROR|WARN|INFO|DEBUG)"
    ],
    "exec": "helper.sh",
    "skip_arguments": false,
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    "use_regex": true
  }
```

Field meanings:

- *exec*: the file name inside `bin/` that is allowed to run.
- *args*: list of allowed arguments. Set *use_regex* to `true` when an argument is a regular expression.
- *skip_arguments* : when `true`, the agent allows the command to run but does not validate the argument list.

You can generate a starting allowlist from the instance using the **Agent Client Collector Command Allow-list Generator** page.

3. Build the `.tar.gz`.
   - On Linux or macOS:

     Open a terminal, change to the plugin folder (the folder that contains `bin/` and optionally `allow_list/`), then run:

     ```
     tar -C . -zcvf my_custom_plugin.tar.gz *
     ```

     Command options:

     - `-C .` — uses the current folder as the base, so the archive contents sit at the root level (`bin/`, `allow_list/`) instead of being wrapped in an extra parent folder.
     - `-z` — gzip compression.
     - `-c` — create the archive.
     - `-v` — verbose output.
     - `-f my_custom_plugin.tar.gz` — output file name.
     - `*` — include all files and folders in the current directory.
   - On Windows:

     Use any archiving utility that supports the `.tar.gz` format. Ensure the archive contains `bin/` at the root and optionally `allow_list/` at the root. Do not wrap the contents in an extra parent folder.
4. Before you upload, verify that the archive has the correct top-level structure:
   Run the following command:

   ```
   tar -tzf my_custom_plugin.tar.gz
   ```

   Expected output:

   ```
   bin/
   allow_list/
   bin/my_script.rb
   bin/helper.sh
   allow_list/allow_list.json
   ```

   There must be no extra parent directory such as `my_custom_plugin/bin/`.
5. Attach the `tar.gz` file to the plugin record, according to the instructions that appear on the UI page.
6. Navigate to **All** > **Agent Client Collector** > **Configuration** > **ACC Plugins**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Click **New**.
   The **Agent Client Collector Plugin - New record** page appears.
8. Enter values in the fields on the page, as described in the following table.

Agent Client Collector Plugin fields

Field	Description
Name	Plugin name. This field is populated automatically with the name of the tar.gz file attached to the plugin, and is visible only when selecting an existing record.
Description	Description of the plugin.
Operating System	Selects the operating system to be monitored by the plugin.
Platform	Selects the platform to be monitored by the plugin.
Active	Selected to activate the plugin and download it to the MID server.
Advanced	Selected to enable selecting the operating system version to be monitored by the plugin.
OS Version	Selects the operating system version to be monitored by the plugin. This field appears only when selecting the **Advanced** check box.

Secure a custom plugin with a certificate

When you customize or create an Agent Client Collector plugin, you can secure the plugin with either a third-party certificate or an internal secure certificate in the plugin's script. Official plugins are signed by an external certificate authority.

Using a signed certificate is enabled by default. To disable this feature, set the *`verify-plugin-signature`* property to **False** in the acc.yml on the host. When you disable a signed certificate, your configured plugins use no certificate.

For details on how to configure a self-signed certificate by using OpenSSL, see Enable OpenSSL secure signing for plugins. For details on how to configure a third-party certificate, consult the third-party vendor's documentation.

When the agent's current certificates can't validate a plugin, Agent Client Collector requests the updated certificates from its synchronized MID Server. You can manually refresh the agent to receive these certificates before the scheduled synchronization. For details, see Manually refresh Agent Client Collector certificates.

You can manually save the certificates in a directory on the MID Server that is connected to the installed agents. Any agent that's installed in your environment can then access the certificates from that MID Server.

For details on how to configure the certificate synchronization properties, see Synchronization properties for validating Agent Client Collector plugins.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Optimize distribution of agents to MID Servers

Optimize the distribution of agents by allowing redistribution from one MID Server to another. Agents will always be connected to the MID Server with the fastest response time.

Before you begin

Ensure that the following settings are configured:

- The *sn_agent.enable_auto_mid_selection* system property is set to true. When using the `acc.yml` file, the property is *enable_auto_mid_selection*.
- On the MID Server page (**Agent Client Collector** > **MID Server**):
  - Status: Up
  - Validated: Yes

Role required: agent_client_collector_admin

About this task

Only validated MID Servers can accept connections from agents.

Procedure

1. Navigate to **Agent Client Collector** > **Deployment** > **MID Servers**.
2. On the **MID Servers [Agent Client Collector view]** page, select the check box next to the MID Server from which you want to redistribute agents.
3. In the **Actions on selected rows** list, select **Try redistributing connected agents**.
   Note the number of agents assigned to your MID Servers.
4. Refresh the page to view the results of the redistribution attempt.
   The values in the **Number of agents** column for any affected MID Server will show a change if any agents were redistributed.

View the Agent Client Collector configuration file for an agent

View the `acc.yml` Agent Client Collector configuration file without having to access the host server by retrieving the file from an agent.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. On the Agent Client Collectors page, select an agent.
3. In the **Related Links** section, select **Grab agent config (acc.yml)**.
   A message appears on the page that indicates that the file is attaching.
4. Wait several seconds and refresh the page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The `acc.yml` configuration file appears as an attachment on the agent.

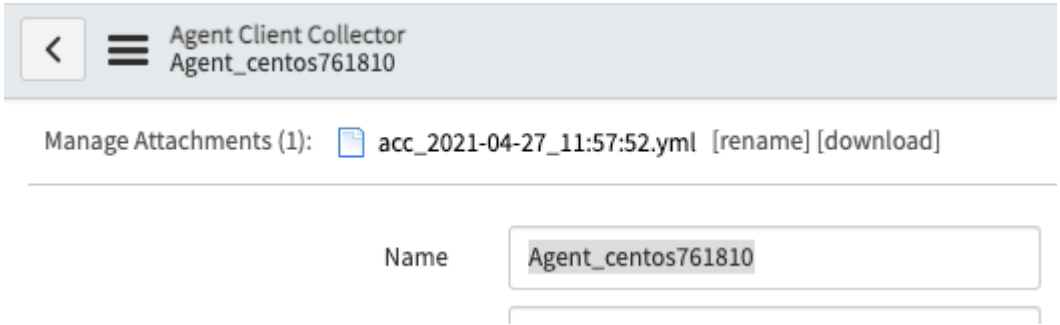

Related topics

Configuration file options

Create an Agent Client Collector configuration data file

Create an Agent Client Collector configuration data file and associate it with check definitions so that you can view the information about your instance. You can also add or delete an attachment to the configuration data files, which are used by check definitions during check executions.

Before you begin

Role required: agent_client_collector_admin

About this task

A check definition is an individual check entry in the system. For details, see Agent Client Collector configuration data files.

Configuration data file attachments contain instance data, such as configuration item (CI) information, which is used when executing a check.

Domain separation is supported in the Agent Client Collector configuration data files.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Configuration** > **Configuration Files**.
2. Select **New**.
3. Add or remove an attachment.

Action	Description
Add a file	Synchronize the record and the relevant MID Servers. **a.** Select the paper clip icon . The Attachments dialog box opens. **b.** Select **Choose file** and select a file to at tach. **c.** Close the Attachments dialog box.
Delete a file	Optionally, delete an attachment you don't want to be associated with the record. **a.** Select **Manage Attachments**. The Attachments dialog box opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action	Description
	b. Select the check box next to the attach ment that you want to remove. **c.** Select **Remove**.

The attachment name must be unique for each domain and can consist only of alphanumeric characters.

If the attachment contains invalid characters or is a duplicate of another attachment, the following situations occur:

- The configuration file is automatically set to **Active = False**.
- A message appears in the Error column indicating a problem with the attachment.

Each record can contain only one attachment.

4. **Optional:** In the **Description** field, enter a description for the configuration data file.
5. Select **Submit**.

Result

- The configuration data file is listed on the Agent Client Collector Configuration File page.
- The **Name** value is automatically set to the name of the attachment.

Checks use the configuration data file with instance data during check execution.

Validate plugins on the MID Server

Validate plugins on your MID Server to ensure that they match the plugins on your ServiceNow instance. Validating plugins ensures that the instance provides accurate data on all of your MID Server plugins.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **MID Servers**.
2. Select a MID Server entry on the MID Servers page.
3. Select the **Validate ACC Plugins** link.
   The **Queues** page appears, displaying the ECC queue.
4. Refresh the queue to view the new input record.
5. Select the record to view its results.

   Validation status is visible in the **Payload** section of the form.

   If plugins aren't validated, restart the MID Server. Plugins sync automatically when the MID Server restarts.

Validate plugins on agents

Validate plugins on your agents to ensure that they match the plugins on your ServiceNow instance. Validating plugins ensures that the instance provides accurate data on all of your agents' plugins.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: agent_client_collector_admin

About this task
Only plugins that are able to be executed are validated. For example, if you're running an agent on a macOS, only macOS plugins on your instance are validated.

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select an agent on the **Agent Client Collectors** page.
3. Select the **Validate ACC Plugins** link.
   The **Queues** page appears, displaying the ECC queue.
4. Refresh the queue to view the new input record.
5. Select the record to view its results.

   Validation status is visible in the **Payload** section of the form.

   If plugins aren't validated, navigate back to the **Agent Client Collectors** page and select the **Clear ACC Plugins** link.

   The plugins automatically re-sync the next time a command runs which requires the plugin.

Using proxy agents in Agent Client Collector
You can use a proxy agent to monitor the health and performance of your configuration items (CIs). This works even if the agent is in the cloud or external to your host server.

Proxy agents are used to monitor:

- URLs (to ensure that the URLs are available and accessible)
- Services that don't have a host server
- Closed services on which you can't configure an agent
- External databases in the cloud

One agent can monitor multiple entities and their CIs.

The following diagram shows the connection between a ServiceNow instance, the MID Server, and Agent Client Collector agents, as well as proxy agents that monitor CIs that are external to the agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Proxy agent configuration

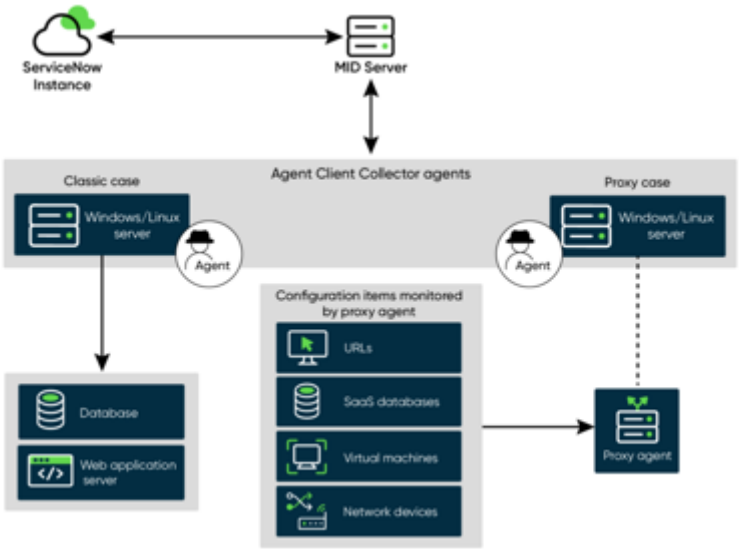

You can configure the proxy agent to do a custom check when it is monitoring external entities. To learn how to create custom checks, see Create and edit checks.

You can create a cluster of proxy agents on multiple proxy servers to monitor the services that are external to the host server. To learn how to create a proxy agent cluster, see Create a proxy agent cluster.

Use a Proxy Auto-Configuration (PAC) file to dynamically determine the appropriate proxy server to use. PAC files provide flexible and automated proxy configuration. They enable the agent to use different proxies for different destinations, connect directly to internal resources, and implement complex proxy routing logic. They also automatically handle proxy failover. For details on the parameters used for PAC file configuration, see Configuration file options.

Following is a sample PAC file:

```
function FindProxyForURL (url, host) {

   // Connect directly to internal hosts
   if (isPlainHostName(host) ||
      shExpMatch(host, "*.internal.company.com")) {
      return "DIRECT";
   }

   // use proxy for all other hosts
   return "PROXY proxy.company.com:8080:
}
```

Related topics

- Checks and policies
- Agent Client Collector Monitoring default checks and policies

Configure an agent on a proxy server

Configure an Agent Client Collector on a proxy server when monitoring services external to the host server, such as URLs or external databases in the cloud.

Before you begin

Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Proxy agents should be dedicated agents, as running other apps on a proxy agent is likely to cause performance degradation.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents** > **Policies**.
2. Select the policy containing the checks that you want to run on a proxy server.
   You must configure a proxy server for checks and policies that monitor external services. For details on these checks and policies, see Agent Client Collector Monitoring default checks and policies.
3. Select the **Proxy Settings** tab.
4. Select from the following fields.

Field	Description
Single proxy agent	Select to display the **Proxy agent** field, where you select a single proxy agent.
Multi-proxy agents by cluster	Select to enable a group of proxy agents to monitor the policy CIs. When selecting this check box, the **Agent cluster name** field appears.
Agent cluster name	Select an agent cluster to monitor the policy CIs.
Multi-proxy agents by filter	Select to filter multiple proxy agents from the agent list. When selecting this option, all proxy agents monitor the same services. For example, you may want to use this option for multiple proxy servers in different geographic locations. **Note:** ○ This field appears only when the **Hide Multi Proxy Agent by filter, if its not selected** property (*sys_ui_policy*) is set to **active=false**. ○ This field appears only for existing policies where this option was previously selected. If you remove the filter in an existing policy, the option disappears.
Multi-proxy agents by script	Select to configure a JavaScript for selecting proxy agents. When selecting this option, you can assign services to different proxy agents, which can serve as a form of load balancing.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Run checks on all proxy agents (No load balancing)	Clear this option to load balance the selected multi proxy agents that monitor the matching CIs in the policy's **Monitored CIs** tab. Select this option for each agent to monitor all matching CIs in the policy's **Monitored CIs** tab. For example, if you have five agents in a cluster or five agents using the filter/script option and 10 matching CIs in the policy, when you clear this option, each agent monitors two CIs. When selecting this option, each agent monitors all 10 of the matching CIs in the policy. For existing or active policies, this option is enabled by default. For new and inactive policies, this option is turned off by default. Select this option only in a case where you have an external endpoint for which you want to measure how quickly it is accessed from different locations.

5. Select **Activate**.

Create a proxy agent cluster

Create a cluster of agents on multiple proxy servers to monitor services external to the host server. Creating a cluster of agents enables you to assign that cluster to multiple policies instead of having to assign the agents individually to every policy.

Before you begin

Role required: agent_client_collector_admin

About this task

A proxy agent cluster monitors the CIs that match the filter specified on the policy's **Monitored CIs** tab.

A proxy agent cluster must contain at least one agent. All of the cluster's agents monitor the CIs specified in its assigned policy. No special configuration is necessary on the policy level to enable proxy cluster monitoring.

The default CPU threshold value when working with proxy agent clusters is 80%. You can modify this value in the *proxy_cpu_percentage_limit* property.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Configuration** > **Proxy Agent Cluster**.
2. Select **New**.
3. In the **Name** field, enter a descriptive name for the cluster.
4. Add an agent to the cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Agents in cluster** section, locate the **Agent** column and select **Insert a new row**.

b. Select the search icon ( ) and select an agent to add to the cluster.

c. Select the check icon ( ) to add the agent to the cluster.
Repeat this step for each agent that you want to add to the cluster.

5. Select **Submit**.

Result
The **Proxy Agent Clusters** page appears and displays the agent added to the cluster.

What to do next
Select this proxy cluster to monitor CIs specified in an agent policy. For details, see Configure an agent on a proxy server.

Assign a proxy agent cluster to a policy

Assign a proxy agent cluster to a policy when monitoring services external to the host server, such as URLs or external databases in the cloud. The agents in the cluster monitor all of the policy's CIs.

Before you begin
Ensure that you have a policy containing the checks that monitor external services. For details on these checks and policies, see Agent Client Collector Monitoring default checks and policies.

Role required: agent_client_collector_admin

About this task
Only one app should be run on a given proxy agent because running multiple apps is likely to cause performance degradation.

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Agents** > **Policies**.
2. Select the policy containing the checks that you want to run on a proxy server.
3. Select the **Proxy Settings** tab.
4. Select **Edit in Sandbox** to enable modifying fields on the form.
5. Indicate whether you are assigning a single proxy agent or a proxy agent cluster to monitor the policy CIs.
   - Assign a single proxy agent by selecting **Single proxy agent** and selecting an agent.
   - Enable a proxy agent cluster by selecting **Multi-proxy agents by cluster**. Provide the name of the cluster in the field that appears.
6. Determine whether to select multiple proxy agents from the agents list through a filter or a script.
   - To select agents through a filter, configure a filter or retain or modify an existing filter when the **Multi-proxy agents by filter** option is selected.

     Note:
     - This field appears only when the **Hide Multi Proxy Agent by filter, if its not selected** property (*sys_ui_policy*) has been set to **active=false**.
     - If you remove the filter in an existing policy, the option disappears.

   - To select agents through a script:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select **Edit in Sandbox**.

b. Clear any selected proxy options in the **Proxy Settings** tab.

c. Select **Multi-proxy agents by script**.

d. In the **Proxy script** field that appears, enter a script to assign services to different proxy agents as a form of load balancing.

7. Determine whether to load-balance the selected multi-proxy agents that monitor the matching CIs in the policy's **Monitored CIs** tab or have each agent monitor the matching CIs.
   - To load balance the selected multi-proxy agents that monitor the matching CIs, clear the **Run checks on all proxy agents (No load balancing)** option.

     For example, if you have five agents in a cluster or five agents using the filter/script option and 10 matching CIs in the policy, each agent would monitor two CIs.

     You can also enable load balancing at a later time, as described in Enable load balancing between proxy agents in a cluster.

   - To have each agent monitor all matching CIs, select the **Run checks on all proxy agents (No load balancing)** option.

     If there are 10 matching CIs in the policy, each agent would monitor all 10. For example, you might want to use this option for multiple proxy servers in different geographic locations.

     If an agent is not functioning properly, the monitored CIs are redistributed to another agent.

   Note: For existing or active policies, the **Run checks on all proxy agents (No load balancing)** option is enabled by default. For new and inactive policies, this option is turned off by default.

8. Select **Activate**.

Related topics

Enable load balancing between proxy agents in a cluster

Enable load balancing between proxy agents in a cluster

Enable load balancing between proxy agents in a cluster so that if an agent is not functioning properly, monitored CIs are redistributed to another agent. After enabling load balancing, you can view the CIs monitored by each proxy agent in a policy.

Before you begin

Ensure that you have upgraded to Tokyo patch 1 or San Diego patch 7.

Role required: agent_client_collector_user

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Configuration** > **Policies**.
2. Select the policy containing the CIs you want the proxy agents to monitor.
3. Select the **Proxy Settings** tab.
4. Enable load balancing for the proxy agents by clearing the **Run checks on all proxy agents (No load balancing)** check box.
   Load balancing is enabled for the proxy agents monitoring the CIs. For example, if you have five agents and 10 matching CIs in the policy, each agent monitors two CIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. In the **Related Links** section, select **Show distribution of proxy agent's CIs**.
   A pop-up window displays the distribution of monitored CIs among agents within the policy.

Configure a MID Server to work with Agent Client Collector Monitoring

Connect agents to a MID Server to enable Agent Client Collector Monitoring (ACC-M) functionality.

Before you begin
Role required: agent_client_collector_admin

About this task
For details on configuring Agent Client Collector for log monitoring, see Agent Client Collector Log Analytics.

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **MID Servers**.
2. Select the MID Server that you want to connect to the agent.
3. In the **Related Links** section, select **Setup ACC Monitoring**.

Result
A confirmation window appears, indicating that the MID Server is connected to the agent.

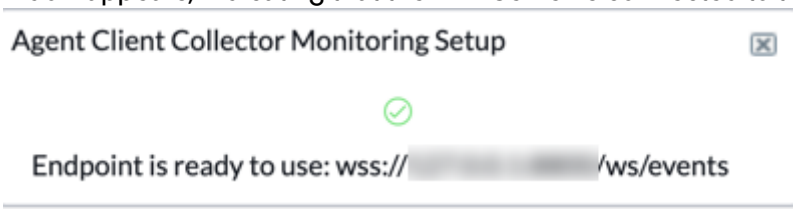

Related topics
Configure multiple MID Servers to work with Agent Client Collector Monitoring

Configure multiple MID Servers to work with Agent Client Collector Monitoring

Connect agents to multiple MID Servers to enable Agent Client Collector Monitoring functionality. Multiple MID Servers can support each other with load balancing and domain separation.

Before you begin
Role required: agent_client_collector_admin

About this task
When configuring multiple MID Servers, use the same port for all MID Servers. Using different ports requires you to add each MID Server individually, as described in Configure a MID Server to work with Agent Client Collector Monitoring.

Procedure
1. Navigate to **All** > **Agent Client Collector** > **MID Servers**.
2. Select the check boxes corresponding to the MID Servers you want to connect to your agent.
3. In the **Actions on selected rows...** window, select **Setup ACC Monitoring**.
   The option includes the number of MID Servers selected out of the number of available MID Servers. For example: **Setup ACC Monitoring (7 of 8)** indicates that 7 MID Servers were selected out of 8 available MID Servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
A confirmation window appears, indicating that the selected MID Servers are connected to the agent.

Enable log monitoring in a Linux environment

To enable monitoring logs in a Linux environment, select the relevant policy and assign specific check parameters to the policy. When log monitoring is enabled and a specified string is discovered in the log being monitored, the system creates an event.

Before you begin
Ensure that the servicenow user has access to the log file being monitored, or that the specified user can run the `sudo` command.

Role required: agent_client_collector_admin

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Select the **Linux log monitoring** policy.
3. On the **Check Instances** tab, select **util.check-logs** to enable monitoring log files not owned by the *`root`* user.
   Ensure that the Agent Client Collector *`servicenow`* user has read permissions for the log file being monitored.
4. Specify the log file to be searched by the check:
   a. In the **Check Parameter Definitions** tab at the bottom of the page, select the **file** entry in the **Name** column.
   b. In the **Default value** field, enter the path of the log file you want to search.
5. Specify the string you want to locate in the log:
   a. In the **Check Parameter Definitions** tab at the bottom of the page, select the **pattern** entry in the **Name** column.
   b. In the **Default value** field, enter the search pattern strings that you want to search for in the log file, such as `404`, `Error`, and so forth. The default values are SEVERE and `Exception`.

      Ensure that you separate multiple patterns with a pipe (|) and pass it as a parameter inside quotes. For example: `"SEVERE|404"`.
6. Click **Update** to save your check instance and return to the **Policy - Linux Log Monitoring** page.
7. Select the **util.check-logs-sudo** check to monitor logs owned by the root user.
8. Repeat steps 4 and 5 to monitor logs owned by the root user.
9. For the **util.check-logs-sudo** check, modify the `/etc/sudoers` file on the host or Linux virtual machine, as follows:
   When the agent is installed on the host's default location (`/usr/share/servicenow/agent-client-collector`) or on a Linux virtual machine, add the following string to the `/etc/sudoers` file:

```
servicenow ALL=
 SETENV: /
var/cache/servicenow/agent-client-collector/monitoring-plugin-l
inux/bin/check-log.rb *
Defaults:servicenow !requiretty
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Defaults exempt_group += servicenow
```

10. On the **Policy** page, click **Save** and then click **Publish** to activate the policy to monitor the log file.

Enable log monitoring in a Windows environment

To enable monitoring logs in a Windows environment, select the relevant policy and assign specific check parameters to the policy. When log monitoring is enabled and a specified string is discovered in the log being monitored, the system creates an event.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Select the **Windows log monitoring** policy.
3. On the **Check Instances** tab, select **os.windows.check-log** to enable monitoring Windows log files.
4. On the **Check Parameters** tab, specify the log parameters to be monitored by the check, as described in the following table:

Check parameters

Name	Value
warning	Number of times the specified pattern strings are found in the log which generates a *`warning`* event. Default = 1. For example, if the *`pattern`* value is **Exception** and one Exception event is located in the log, a *`warning`* event is generated.
critical	Number of times the specified pattern strings are found in the log which generates a *`critical`* event. Default = 2. For example, if the *`pattern`* value is **Exception** and two Exception events are located in the log, a *`critical`* event is generated.
file	Location of the log file.
pattern	Strings which are being searched for in the log. Default values are **Severe** and **Exception**. Other possible values include **404** and **Error**. Ensure that you separate multiple patterns with a pipe (|) and pass it as a parameter inside quotes. For example: `"SEVERE|404"`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select the Azure policies to activate for metric collection

Activate the Azure policies that come with the Agent Client Collector base system so that Azure virtual machines can collect metrics. These metrics help you monitor the configuration items (CIs) in the CMDB.

Before you begin

- Prepare an instance with an active MID Server and metric base.
- Install Agent Client Collector Framework and Agent Client Collector Monitoring.
- Run Discovery on the Azure cloud account to populate the CMDB with CIs to be monitored.

Role required: agent_client_collector_admin

Procedure

1. Prepare and configure the proxy agent to be used by the policy.
   For details on configuring a proxy agent, see Using proxy agents in Agent Client Collector.
2. Assign the credentials to the Azure policy.
   a. Navigate to **All** > **Agent Client Collector** > **Policies**.
   b. Select the relevant policy.
      For details on the policies that are available with the base system, see Azure cloud metrics.
   c. Select the **Credentials** tab.
   d. In the **Credential name** field, select the credentials to be invoked for the policy.
      To see the available credentials, go to the Credentials page by navigating to **All** > **Discovery** > **Credentials**.

      Ensure that you select the **Type=Azure** credentials. If you need to create credentials, select **New** and select **Azure Service Principal** credentials.
   e. In the **Credential alias** field, assign an alias to the credentials.
3. Navigate to **All** > **Agent Client Collector** > **Configuration** > **Policies** and select the relevant policy to activate.
4. Select the application scope icon and verify that you are in the Agent Client Collector Monitoring scope to edit the policy.
5. Enable the policy in Edit mode by selecting **Edit in Sandbox**.
6. Enable the proxy agent to be used to execute the check.
   a. Select the **Proxy Settings** tab.
   b. Select **Single proxy agent** and select the proxy agent to be used to execute the check from the **Proxy agent** field.
7. Select **Save**.
8. Republish the policy by selecting **Republish**.

Use the Agent Client Collector Health Dashboard

The Agent Client Collector Health Dashboard enables you to monitor the status of the agents in your system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agent Health Dashboard**.
   The **Agent Health Dashboard** page appears and displays information about the configured agents in the dashboard widgets.

   In the **Agents by Data Collection Status** widget, **Off (manual)** indicates agents manually turned off (for maintenance or troubleshooting), and **Off (auto)** indicates agents automatically turned off (after crossing the CPU usage threshold).

2. Click the refresh arrow icon on the top right corner of a widget to refresh its data, as needed.

   You can refresh all widgets on the dashboard by selecting the form header icon and then selecting **Refresh**.

Agent Client Collector Discovery

Discover CIs in your environment by using Agent Client Collector for Visibility Content (ACC-VC) Discovery. ACC-VC works with both horizontal IP-based Discovery, and you can also use push-based Discovery.

Using push-based and horizontal IP-based Discovery together

Discovery performed by Agent Client Collector for Visibility Content (ACC-VC) is compatible and can coexist with horizontal IP-based Discovery. You may have ACC installed on a given target host and still have that host as part of a horizontal IP-based Discovery schedule as well.

To avoid conflicts, ACC-VC does not perform Discovery when the target host's IP already has a CI that has been discovered by horizontal IP-based Discovery (discovery_source is ServiceNow). To override this behavior, set the system property [*sn_agent.disco_disable_ci_clobber_of_agentless_disco*] to false. This setting is ignored if horizontal IP-based Discovery was not performed over a specific time. Update the system property [*sn_agent.disco_ci_clobber_of_agentless_disco_threshold_days*] if necessary. The default is 14 days. See System properties for more details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To enable Discovery using Agent Client Collector, enable the following Discovery properties on the Discovery properties page (**All** > **Discovery Definition** > **Properties**):

Agent Client Collector Discovery properties

Field	Description
glide.discovery.domain.name.nbt	Set OS domain name by NBT or WMI. If set to **true**, Windows domain name is set by NBT. Otherwise it's set by WMI. • **Type**: true | false • **Default value**: true
glide.discovery.enable.software_simplify	Enables cleanup for software's name and version. If set to **true**, it removes any appended version or commonly found phrases from the name and removes leading 0s in the version. If set to **false**, raw data is persisted in the table. Use with the **glide.discovery.enable.software_simplify_sccm** property to avoid software data discrepancy between Discovery and SCCM. • **Type**: true | false • **Default value**: true **Note:** Once the property is set to **false**, to disable the cleanup, duplicate records are generated for a given software. Name and Version are the primary identifiers.

Using push-based Discovery and Intel EMA together

Agent Client Collector for Visibility Content (ACC-VC) can collect data for uses cases with the Intel vPro® platform when the Intel ® Endpoint Management Assistant (EMA) application is installed on Windows endpoints. You can install the Intel EMA application from the ServiceNow store. Attributes are stored in the CMDB when enabled. Currently, data for Intel EMA can only be fetched for Windows endpoints.

The Intel EMA application enables customers to configure access to the Intel Instance from their ServiceNow instance. For example, IT Support groups can open a remote console and perform some actions like power cycle from a ServiceNow incident.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Intel + ServiceNow platform

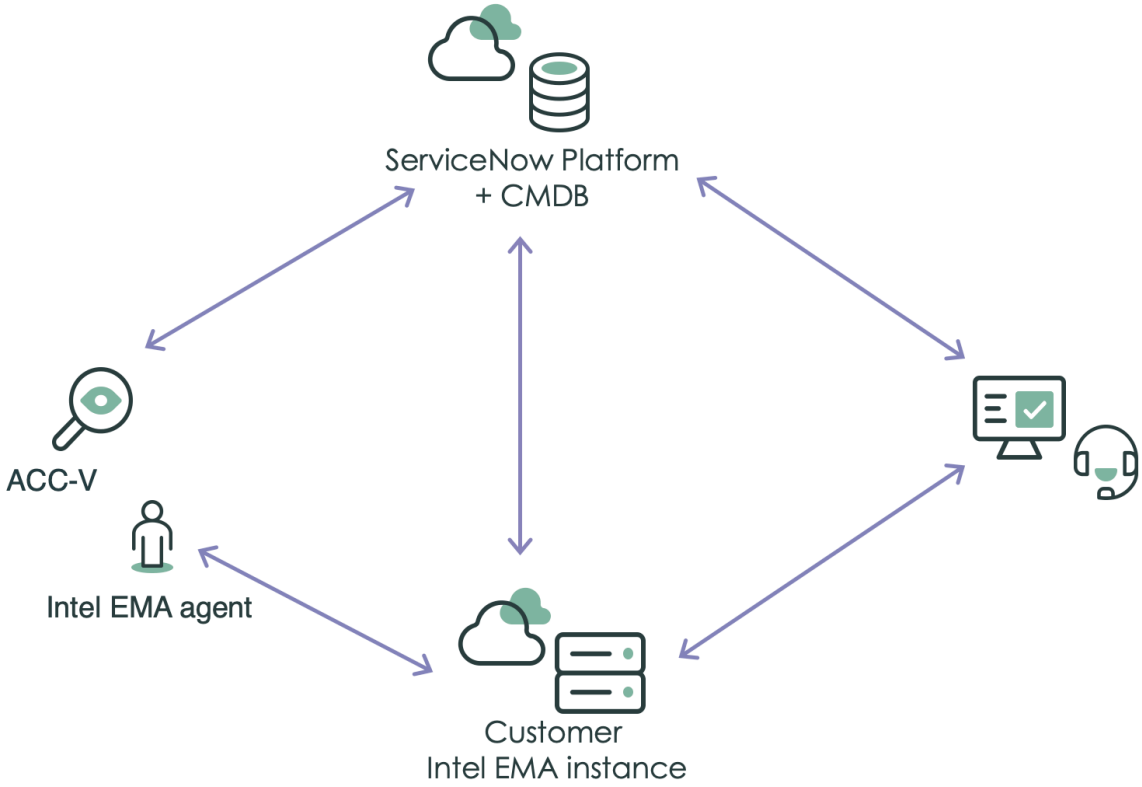

Using push-based Discovery and SAM together

Agent Client Collector for Visibility Content (ACC-VC) collects installed software data for use cases for Software Asset Management (SAM), when the SAM plugin is installed. Using push-based Discovery and SAM together can help optimize software data collection with SAM basic metering and SAM total usage metrics.

ACC-VC can capture the last accessed time for the software or applications that are installed on the target via push-based Discovery. This information along with the target CI reference, is added to the Software Update [samp_sw_usage] table.

SAM Basic metering and SAM total usage metrics are supported for both Windows and macOS.

The software usage records are domain separated. The records are populated with the domain of the MID Server that is used for the agent-based Discovery for the target.

Note: For software installations (cmdb_sam_sw_install), to avoid insertion of duplicate records, the same discovery source "ServiceNow" is being used for both push-based Discovery and horizontal IP-based Discovery.

Requirements

SAM basic metering and SAM total usage metrics

To apply SAM basic metering or SAM total usage metrics, you need the following:

- SAM plugin (com.snc.samp) enabled
- System property [*sn_acc_vis_content.persist_sam_usage_metrics*] set to true. See System properties for more details.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Write permissions enabled for the log folder in the ACC install directory.
- In a Windows environment: For maximum efficiency, run ACC with the Local System account (on the agent, set **Log On As = Local System**).

For details on SAM metering setup with the Agent Client Collector, see the Knowledge Base article KB1642676.

Software edition information

To retrieve software edition information, you need the SAM plugin (com.snc.samp) enabled.

SAM basic metering

SAM basic metering tracks when installed software was last used on an endpoint. In a Windows environment, this is determined by observing which applications are running on the machine.

How SAM basic metering works (on a Windows machine)

- Every 8 minutes, the agent checks the running process table on the endpoint and records each observed application.
- Running applications are considered used, and their last accessed time updates to the current time.
- Once every 24 hours, the agent also reads the Windows Registry Run keys (`HKLM...\Run`, `HKLM\Wow6432Node...\Run`, and `per-user HKU<SID>...\Run`) to acknowledge auto-start applications that may have started and exited before being checked. These observations are sent to the instance as part of the existing 24-hour metering payload.
- On an instance, each observed executable is matched to a product using an exact lookup in the Software Product Process (*samp_sw_product_process*) table.
- If a matching product is found, the last accessed time on the corresponding Software Installation (*cmdb_sam_sw_install*) record is updated.
- If no mapping exists in the Software Product Process (*samp_sw_product_process*) table for an executable, no last accessed time is written for the software installation record. The unmapped executable is logged for the for review.

How SAM basic metering works (on a macOS machine)

The last accessed time (*last_used* value) is collected using the native macOS metadata attribute.

If reclamation rules are enabled for the product, the last accessed time is also stored in the Software Usage (*samp_sw_usage*) table.

Use the *sn_acc_vis_content.disable_sam_reclamation_rules_for_licensable_softwares* property to define reclamation rules for licensable software, as follows:

- **True**: SAM usage is stored for all licensable software (with or without defined reclamation rules), and for any non-licensable software that has a reclamation rule defined.
- **False**: Store SAM usage according to defined reclamation rules.

Common applications supported include:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- WinZip
- Google Chrome
- Sublime Text
- Notepad++
- Autodesk
- Microsoft Office 365
- Tableau

SAM total usage metrics

SAM total usage metrics allows you to measure total usage time and total usage count on any application that has a software reclamation rule enabled.

Osquery provides a daemon executable which can run as a service, called Osqueryd. Osqueryd needs to be manually deployed for SAM total usage metrics to work properly. Each Osqueryd deployment requires the osquery.conf file, optional external packs, and initialization flags (configured in osquery.flags file) provided when starting the service. In return, the daemon service runs scheduled queries on the host and logs it into a local file system.

Note: Osquery supports filesystem-based logging by default. This configuration is provided in the osquery.conf file on any fresh Osquery installation. Osquery is not needed from an implementation or configuration viewpoint for basic metering. Total Usage metrics requires osqueryd to be installed separately.

Domain information can be collected during the data collection. This can help large organizations with multiple employee directories map software to the correct user. Currently, this is supported for Windows only. To map the software usage/assigned_to with the correct user in a domain separated environment, use the system property [sn_acc_vis_content.column_name_for_user_mapping] with a valid field name. By default, the value of this system property is empty which means it only validates the username and not the domain. You can use either of the following formats to validate username and domain: username@domain or domain\username.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

SAM total usage metrics flowchart

Using the list of processes, you can perform SAM normalization to map the processes for the relevant installed software records. This provides flexibility since installed software names and processes are not usually the same. For the list of processes in the payload, query the Software Discovery Model [cmdb_sam_sw_discovery_model] table and Software Product [samp_sw_product] table to fetch the corresponding product and publisher. Once the product is fetched, check if the reclamation rule is enabled for that product to persist the total usage time in the Software Usage [samp_sw_usage] table. See the flowchart for details.

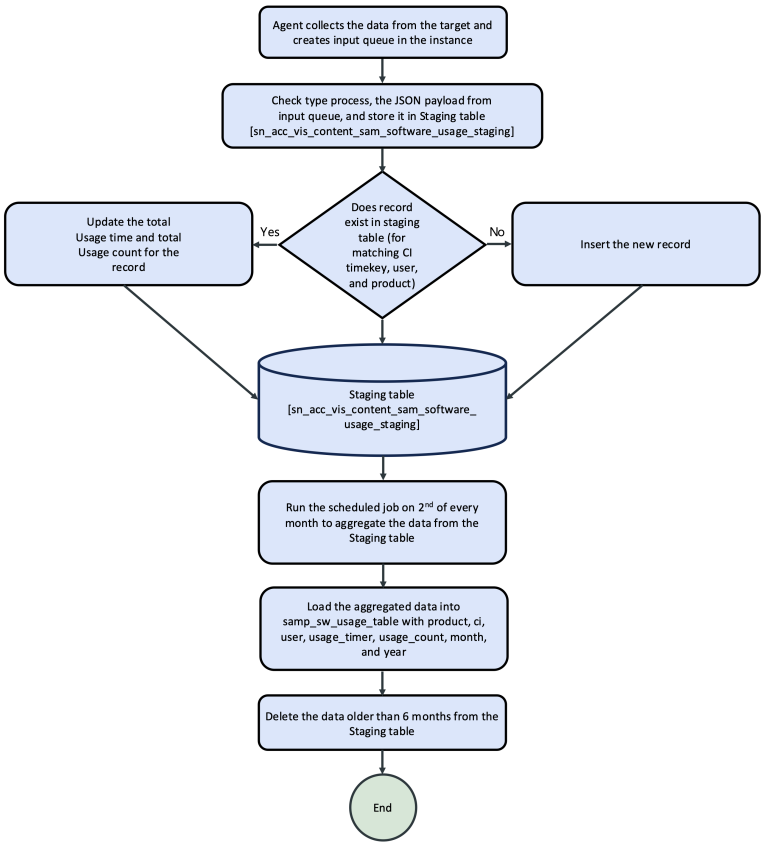

install and configure Osqueryd for Windows using the following script.

```
# Install latest osquery

$msi = "osquery-5.7.0.msi"
$url = "https://pkg.osquery.io/windows/$msi"
$dst = "$PSScriptRoot\$msi"
Invoke-WebRequest -Uri $url -OutFile $dst
# msiexec /i "$dst" /quiet /qn /norestart
Start-Process msiexec.exe -Wait "/i $dst /quiet /qn /norestart"

# Configure osqueryd service

$flags = "--logger_rotate=true
--logger_rotate_size=26214400
--logger_rotate_max_files=1
--watchdog_level=-1
--config_path=C:\Program Files\osquery\osquery-sam.conf"
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
Set-Content -Path 'C:\Program
 Files\osquery\osquery.flags.default' -Value "$flags"

$conf = @'
{
  "options": {
    "config_plugin": "filesystem",
    "logger_plugin": "filesystem",
    "utc": "true"
  },
  "schedule": {
    "sam_process_info": {
      "query": "SELECT name, pid, elapsed_time, start_time,
 user_time, system_time, username FROM processes p JOIN users
 u ON u.uid = p.uid WHERE p.elapsed_time != -1 AND u.type !=
 'special';",
      "snapshot" : true,
      "interval": 300
    },
    "system_info": {
      "query": "SELECT hostname, cpu_brand, physical_memory FROM
 system_info;",
      "interval": 3600
    }
  },
  "decorators": {
    "load": [
      "SELECT uuid AS host_uuid FROM system_info;",
      "SELECT user AS username FROM logged_in_users ORDER BY
 time DESC LIMIT 1;"
    ]
  },
  "packs": {
  }
}
'@
Set-Content -Path 'C:\Program Files\osquery\osquery-sam.conf'
 -Value "$conf"

cd 'C:\Program Files\osquery'
.\manage-osqueryd.ps1 -uninstall
.\manage-osqueryd.ps1 -install
Restart-Service osqueryd
```

For details on Windows and macOS see Configure Osqueryd schedule for SAM total usage metrics and Configure Osqueryd logs for SAM total usage metrics.

Collecting SAM metrics without osqueryd

Optionally, you can enhance efficiency by using non-osqueryd data collection when using push-based Discovery and Software Asset Management (SAM) together. When non-osqueryd data collection is invoked, data collection is automatically performed on all available agents, instead of invoking osqueryd on each agent individually.

To perform non-osqueryd data collection:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Ensure that the following permissions are configured for the relevant OS:
   - Windows: The ACC service must run as the Local System account. Set the ACC service's Log On As value to **Local System**.
   - macOS: The `servicenow` user must be able to run osqueryi without a password. For information about `servicenow` user permissions for osqueryi, see Configure ServiceNow sudoers file.
2. On the System Properties page (**All** > **System properties** > **All properties**), set the *sn_acc_vis_content.enable_sam_collection_without_osqueryd* property to **true**.

Note:

- Enable this property only when all agents are version 4.1.0 or later.
- SAM basic metering (last accessed time) works with both osqueryd and non-osqueryd configurations. The 8-minute process poll and 24-hour payload run regardless of the osqueryd deployment mode.

Software edition information

Edition information is supported for Adobe Acrobat and MS SQL server. With this feature, SAM admins can get clear visibility into the editions of their installed software. Osquery commands are used to fetch the edition information which then shows in the Software Installation [cmdb_sam_sw_install] table in the Edition Override column. For more details, see the support KB: https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB0721360

Related topics

Configure Osqueryd schedule for SAM total usage metrics

Configure Osqueryd logs for SAM total usage metrics

Configure Osqueryd schedule for SAM total usage metrics

SAM total usage metrics works by relying on the Osqueryd service running on the target host. Configure the Osqueryd service to run the required schedule Osquery on the host.

Before you begin

Role required: admin

Procedure

1. Navigate to **Osqueryd installation folder**.
2. Find and edit the osquery.conf file.
3. Add the below Osquery pack under the packs keyword.

   For Windows:

```
"packs": {
      "sam-metering":
 "C:\\ProgramData\\ServiceNow\\agent-client-collector\\cache\
\acc-visibility-modules\\bin\\sam-metering.conf"
  }
```

   For macOS:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
"packs": {
    "sam-metering": "/Library/Application\
 Support/servicenow/agent-client-collector/cache/acc-visibili
ty-modules/bin/sam-metering.conf"
  }
```

Result
You can now Configure Osquery logs for SAM total usage metrics.

Configure Osqueryd logs for SAM total usage metrics

By default, Osquery supports log rotation based on size. To enable it for SAM total usage metrics and to configure the log size and rotation, you need to add specific flags for Osqueryd service.

Before you begin
Role required: admin

Note: SAM with Osqueryd consumes memory and CPU cycles. For example, with the Agent installed on a Windows Machine having 2 CPU core, 8 GB RAM, 1000 software installations, and 500 running processes, the following was observed:

- Average CPU usage for the Agent and Osqueryd was less than 10 % CPU and maximum of 30% CPU. This will only occur when the SAM background policy is triggered. By default, the trigger happens every 480 seconds.
- Average Memory usage for Agent and Osqueryd was less than 10 MB and maximum of 26 MB was consumed.

To store the data for a single running process for two days, the log file size on an average should be 52429 bytes. The log file size must be increased if the number of running processes on the machine is higher. For example, with 500 running processes the log size would be on average 25 MB which is 26214400 bytes.

Procedure

1. Navigate to **Osqueryd installation folder**.
2. Find and edit the osquery.flags file.
3. Add the below flags with these sizing guidelines:

   Note: Modification of the log file size should be done as per the Osqueryd Log file sizing guidelines.

   For Windows:

   - --logger_rotate=true
   - --logger_rotate_size=26214400
   - --logger_rotate_max_files=1
   - --watchdog_level=1

   For macOS:

   - --logger_rotate=true
   - --logger_rotate_size=26214400
   - --logger_rotate_max_files=1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- --watchdog_level=1
- --logger_mode=0644

4. Save the file.

Result
Once the Osqueryd schedule and Osqueryd logs are configured the Osqueryd service can start.

The schedule runs the Osquery `Select name, pid, elapsed_time, start_time, user_time, system_time, username from processes p JOIN users u ON u.uid = p.uid where p.elapsed_time != -1 AND u.type !='special';` every 5 minutes (300 seconds) on the target machine. The results are logged into the log file. The log file contains snapshot entries of all the queries configured to run by the Osqueryd. This query contains all the processes attributes.

Note:

A temporary file *`marker.json`* is created in a temporary local folder on your machine in the directory:

For Windows `: <userprofile>\\AppData\\Local\\AgentClientCollector\\SAM`.

For macOS: `/Library/Application\ Support/servicenow/agent-client-collector`.

This file has read/write permissions and contains the marker data: *`Data and Last Read Unix Time stamp`*.

The Osqueryd can also be configured to write its logs to a custom directory path instead of the default directory. If you choose a custom directory, modify the check definition [samadvanced-background-log-check].

Application patterns for the Agent Client Collector

Application patterns gather details on the applications that run on the Agent Client Collector (ACC) host. Application patterns are supported only for servers, and are triggered after the Agent Client Collector host Discovery is complete.

Prerequisites for using application patterns

To enable application patterns on the Agent Client Collector, install the following:

- Agent Client Collector Framework
- Agent Client Collector for Visibility Content
- Discovery

Ensure that the agent is connected to a MID Server to enable pattern execution. Pattern execution on an agent is supported only using a MID Server. Each pattern step is executed on the agent as an ad-hoc check; ad-hoc checks are logged on the agent side as *`PatternOnTheFlyCheck`*.

Set the *`sn_agent.appl_classification_behavior`* property to **full**, and restart any MID Servers connected to the agent.

When working in a Windows environment, ensure that you're working with Powershell 3.0 or higher.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required (on the host device):

- In a Linux or macOS environment: ServiceNow user with sudoers; nopasswd enabled with regular Discovery commands.
- In a Windows environment: Local System user, SYSTEM, or a local user account where the local user account uses one of the following:
  - Logon as a service
  - Backup (SeBackupPrivilege)
  - Debug Programs privileges

Viewing supported application patterns

To view the list of patterns supported for Agent Client Collector, navigate to **All** > **Pattern Designer** > **Supported Patterns for ACC**. Application patterns are the only pattern types supported for use with Agent Client Collector.

Viewing Application Patterns executed on Agent Client Collector

You can view the application patterns that were executed on a given agent.

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select an agent.
3. Select **View** > **Discovery** from the context menu (hamburger icon).
4. Open the agent.
5. Select the **Application Pattern logs** tab on the bottom of the page.

   The list of patterns executed on the agent appears. Executed patterns are indicated with the value **true** in the Executed On Agent column.

Unsupported application pattern operations

Application patterns aren't supported in the following pattern operations:

- WMI method invocation
- Change/unchange user
- All custom operations
- SNMP query
- Put file

Agent Client Collector pattern execution properties

Configure the following properties (**All** > **System Properties** > **All Properties**) to maximize the use of application patterns with Agent Client Collector:

Application pattern properties

Property	Description
sn_agent.appl_classification_behavior	Set value to **full** to enable working with applications patterns and to access pattern commands in the allow-list. Default: simple

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application pattern properties (continued)

Property	Description
	Note: To enable application pattern execution with the allow-list, use the allow-list template for application patterns. For details, see Generate a Pattern allowlist.
sn_agent.trigger_patterns_on_endpoints	Controls whether Discovery patterns are triggered during application classification for endpoint CIs (cmdb_ci_computer). Set to true to enable pattern triggering on endpoints, which follows the same classification path as non-endpoint CI types. When set to false, endpoints skip pattern-based classification and fall back to simple classification behavior. Default: false

MID Server configuration parameter properties

mid.discovery.agent.command.mid.timeout	The maximum duration, in seconds, that the MID Server waits for a response from the agent before timing out. Default: 60 seconds
mid.discovery.agent.command.timeout	The maximum amount of time, in seconds, that the command is given to execute on the agent host. Default: 60 seconds

Using pattern designer debugging in ACC mode

You can debug application patterns that are configured in ACC mode.

1. Set the *sn_agent.appl_classification_behavior* property to **full**.
2. In the Pattern Designer, enable the debug option and select **Use ACC connection** in the Debug identification section.

The pattern log indicates whether the pattern was successfully executed on the agent.

You can also receive pattern details in the following logs:

- MID Server logs: Set *mid.log.level* to DEBUG for the mid agent log to display details on the pattern step, ACC check, and the pattern response.
- ACC logs: Search for the string *PatternOnTheFlyCheck* for details on pattern checks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle Global License Advisory Services (GLAS) data collection

Oracle Java process discovery works either with patterns (when using either a MID Server or Agent Client Collector), or with file-based discovery (using Agent Client Collector with file-based discovery). If both the pattern and file-based discovery run the same Java installation, the pattern-collected data takes precedence.

When performing Discovery with Agent Client Collector, ensure that you have the following settings:

- All users must have the Oracle GLAS store app, starting with version 1.8.4.
- Agent Client Collector users:
  - Must have the Agent Client Collector Framework store plugin.
  - Can have the Agent Client Collector for Visibility Content store plugin, for extra data.

  For information about privileged user access, see KB1705845.

> **Note:** Oracle Java discovery with ACC is applicable in all server environments. However, currently, it's partially supported for different Desktop environments. For a Windows Desktop update set, see the KB article KB1705845.

For information on java installation Discovery using Agent Client Collector for Visibility Content, see Discover java installation data using Agent Client Collector for Visibility Content file-based discovery.

For information on the collected data, see Data collected during Agent Client Collector for Visibility Content File-Based Discovery.

Application patterns support Oracle GLAS data collection in a Linux environment. For details on using patterns, see Oracle GLAS data collection using patterns.

Although Oracle GLAS supports Put file operations, the related files are deployed using Agent Client Collector plugins instead of SSH or WMI.

Related topics

Discovery patterns used by ITOM Visibility

Discover java installation data using Agent Client Collector for Visibility Content file-based discovery

Discovering java installation data using Agent Client Collector for Visibility Content file-based discovery enables you to discover file information in your system.

Before you begin

- Install the latest version of Agent Client Collector for Visibility Content.
- Install file-based discovery.
- Enable the Oracle Global License Advisory Services (GLAS) hardware data collection policy.

Role required: discovery_admin

Procedure

1. Enable file-based discovery.
   a. Navigate to **All** > **Discovery Definition** > **Configuration Console**.
   b. Enable the **File Discovery** entry.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Enable the Java Discovery property.
   a. Locate the *sn_acc_vis_content.file_discovery.enable_java_discovery* property on the System Properties page (**All** > **System Properties** > **All Properties**).
   b. Set the property value to **true**.
3. Grant elevated permissions to the `acc` user, enabling access to your java installation directories, for example:

   - `/usr/lib/jvm/*/bin/java`
   - `/usr/lib/jvm/*/*/bin/java`

   You can add additional java directories during java installation, as needed.

   - In a Linux environment, grant elevated permissions in the `/etc/sudoers` file to enable access to the java installation directories.
   - In a Windows environment, grant elevated permissions (read and execute) via the system administrator to enable access to the java installation directories.

Result

When running Discovery for java installation, data is discovered through file-based discovery and is stored in the `ora_java_audit` table. For details on the data stored in this table, see Data collected during Agent Client Collector for Visibility Content File-Based Discovery.

Discover java installation data using Agent Client Collector for Visibility Content process-based discovery

Discovering java installation data using Agent Client Collector for Visibility Content process-based discovery enables you to discover java installation information in your system.

Before you begin

- Install the latest version of Agent Client Collector for Visibility Content.
- Enable the Oracle Global License Advisory Services (GLAS) hardware data collection policy.

Role required: discovery_admin

Procedure

1. Enable the process-based discovery check definitions.
   a. Navigate to **All** > **Agent Client Collector** > **Check Definitions**.
   b. Locate the **Oracle GLAS Java Process Discovery** and **Oracle GLAS HW data collection** entries.
   c. Locate the **Active** column for each entry, select the value and choose **true**.
2. Grant elevated permissions to the `acc` user, enabling access to your java installation directories, for example:

   - `/usr/lib/jvm/*/bin/java`
   - `/usr/lib/jvm/*/*/bin/java`
   - `/usr/lib/jvm/*/bin/jcmd`
   - `/usr/lib/jvm/*/*/bin/jcmd`
   - `/var/cache/servicenow/agent-client-collector/osquery/bin/osqueryi`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can add additional java directories during java installation, as needed.

- In a Linux environment, grant elevated permissions in the `/etc/sudoers` file to access the java installation directories.
- In a Windows environment, grant elevated permissions (read and execute) via the system administrator to access the java installation directories.

Discovering DNS names using push-based discovery

CMDB owners need CIs to contain all domain system names (DNS) associated with their system. Starting in Agent Client Collector for Visibility Content (ACC-VC) version 2.3.0, ACC-VC can discover DNS name lists for Windows and Linux CIs.

When discovering DNS names associated with systems, the following occurs.

- network_adapters.rb is used to populate the DNS name for all the IP addresses using OSquery. (There is no limit.)
- Nslookupcommand is used to get the DNS name for all IP addresses.

Note: Nslookup must be configured on the target.

This creates a relationship of *`Owns::Owned by`* between the CI and the DNS names.

This information is then populated into the [cmdb_ci_dns_name] table.

Populating Assigned To attribute in Computer CI for Agent Client Collector for Visibility Content

To update the Assigned To attribute of the Computer CI, you need to collect information from the logged in user.

Before you begin

Role required: admin

You can automatically populate Assigned to for Windows endpoint devices and macOS devices, like workstations or employee laptops, as part of agent-based Discovery using ACC-VC with the following system properties. See Available system properties for more information.

- **sn_acc_vis_content.set_assigned_to**
- **sn_acc_vis_content.assigned_to_user_order**
- **sn_acc_vis_content.column_name_for_user_mapping**

Procedure

1. Create higher user privileges to fetch the logged in user details.
   Use Log on As Local System User instead of the default ServiceNow user for running the Agent Client Collector. When working in a macOS environment, use the default ServiceNow user.

   Ensure that you have Powershell to run the internal commands on your machines that fetch the logged in user details.

2. User names match is performed on the user records in the sys_user table.
   The Assigned To attribute is populated and then it creates a reference to the user. Systems logs shows a warning if no user can be found in sys_user. The user names sourced for Assigned To attribute is based on the system property: **sn_acc_vis_content.assigned_to_user_order**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Assigned_to attribute is set based on priorities.
   If assigned_to is already there, keep it as is by default. However, you can update the **sn_acc_vis_content.set_assigned_to** system property to override it. If there are any Reconciliation Rules defined, then the IRE populates the field accordingly.

Result

No new user account is created in the sys_user table. Instead, the existing user is queryied and the same is referenced as Assigned To user for a CI.

Populating users based on type for Agent Client Collector for Visibility Content

You can persist the type of users that populate the CMDB depending on your particular interests. Currently, local and system are supported for Windows, Linux, and macOS.

You can filter the types of users by using the system property: **sn_acc_visibility.persist_os_user_type**. See Available system properties for more information.

For Windows

You need to identify if a server is domain controller or not. Since domain controller machines don't have local accounts, local user information does not persist for such machines and is not populated.

For non-domain controller Windows machines, use OSquery commands to get user information through ACC-VC: select *uid*, *username*, *directory*, *type* from users where **type=local** or **type=special**.

For Linux

The userID defines if a user is a local account or system account. This is defined by *UID_MIN* and *UID_MAX*.

Use this command to get user information through ACC-VC to get *UID_MIN* and *UID_MAX* value: `grep -E '^UID_MIN|^UID_MAX' /etc/login.defs`.

Use this command to get user names: `cat /etc/passwd | awk -F':' '{print $1,$3,$6`.

For macOS

Use OSquery commands to get user information through ACC-VC: select *uid*, *username*, *directory* from users. User names that are preceded by '_' are treated as system accounts. Others are local accounts.

License key discovery

License key discovery in Agent Client Collector for Visibility Content automatically collects software license keys from the Windows registry on managed endpoints.

License key discovery enables the Agent Client Collector for Visibility Content (ACC-VC) Windows agent to read software license keys stored in the Windows registry on managed endpoints. When a device is discovered, the agent reads registry values defined by an admin-configured list of registry paths and value names, and reports the results in your instance.

The feature is off by default and must be explicitly enabled using a system property.

Why use license key discovery

License key discovery helps administrators track software license keys across managed Windows devices. By linking discovered keys to device CIs and SAM products, it supports software asset management and license tracking efforts. Administrators can detect duplicate licenses, track compliance, and manage license costs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Who can use license key discovery

Specific roles are required to configure and view license key discovery data. For details, see the Access control table on the License key discovery and access control tables page.

No role can delete license key records. If a key that was previously detected is no longer found on a device, it is marked as **Absent** rather than deleted, preserving the audit trail.

Configure license key discovery

Enable license key discovery and define the registry paths and values you want the Agent Client Collector for Visibility Content Windows agent to collect from managed endpoints.

Before you begin

The SAM Professional plugin (com.snc.samp) must be active.

The Agent Client Collector for Visibility Content Windows agent must be deployed on the endpoints.

Role required: agent_client_collector_admin or discovery_admin

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for the *sn_acc_vis_content.enable_license_key_discovery* property.
3. Set the **Value** field to `true`.
   The default value is `false`. Setting this property to `true` enables license key collection.
4. Select **Save**.
5. Navigate to the **License Registry Configuration** [sn_acc_vis_content_license_registry_config] table and create a record for each registry entry you want to collect.
   Enter the appropriate values in the fields for each record. For details about the fields, see the License Registry Configuration form table.
6. Select **Submit** to save the configuration record.
7. **Optional:** To apply configuration changes immediately, navigate to **All** > **System Definition** > **Scheduled Jobs**, search for the `Refresh License Key Config File` scheduled job, and select **Execute Now**.
   By default, this job runs daily at 8:30 a.m. UTC.

Result

The *`Refresh License Key Config File`* scheduled job reads all active configuration records and publishes them to the agent as `LicenseKeyConfig.json`. On the next agent run, the Windows agent queries the defined registry paths on each managed endpoint and sends the collected key values to your instance. Discovered keys are stored in the **License Keys** [sn_acc_vis_content_license_keys] table, linked to the device CI, the associated SAM product, and the user assigned to the CI. If a key that was previously detected is no longer found, it is marked as **Absent**.

Agent Client Collector Log Analytics setup

Configure Agent Client Collector Log Analytics on your ServiceNow instance either with or without using guided setup.

Agent Client Collector Log Analytics setup using guided setup

Agent Client Collector Log Analytics guided setup provides a sequence of tasks that help you configure Agent Client Collector Log Analytics on your instance. Using guided setup ensures that

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

you have the minimum required setup to use the application. For more information, see Set up Agent Client Collector Log Analytics using guided setup.

For more information about using the guided setup interface, see Using guided setup.

Agent Client Collector Log Analytics setup without using guided setup

If you are not using guided setup, you must perform several configuration steps manually. For more information, see Set up Agent Client Collector Log Analytics without using guided setup.

Set up Agent Client Collector Log Analytics using guided setup

Configure Agent Client Collector Log Analytics on your ServiceNow instance using guided setup.

Before you begin

- Agent Client Collector Log Analytics has a dependency on the ServiceNow® Event Management application. Set up Event Management before you configure Agent Client Collector Log Analytics. For more information, see Event Management setup.
- ServiceNow Service Operations Workspace ITOM Apps must be installed. If it is not installed already, install it manually. For more information, see Install Service Operations Workspace for ITOM application.

Role required: agent_client_collector_admin

About this task

Agent Client Collector Log Analytics guided setup provides a sequence of tasks that help you configure Agent Client Collector Log Analytics. Using guided setup ensures that you have the minimum required setup to use the application.

Note: If you are not using guided setup, you must perform several configuration steps manually. For more information, see Set up Agent Client Collector Log Analytics without using guided setup.

Procedure

1. Navigate to **All** > **Guided Setup** > **ITOM Guided Setup**.
   The IT Operations Management Guided Setup welcome screen appears.
2. Select **Get Started**.
3. In the Agent Client Collector Log Analytics pane, click **Get Started**.

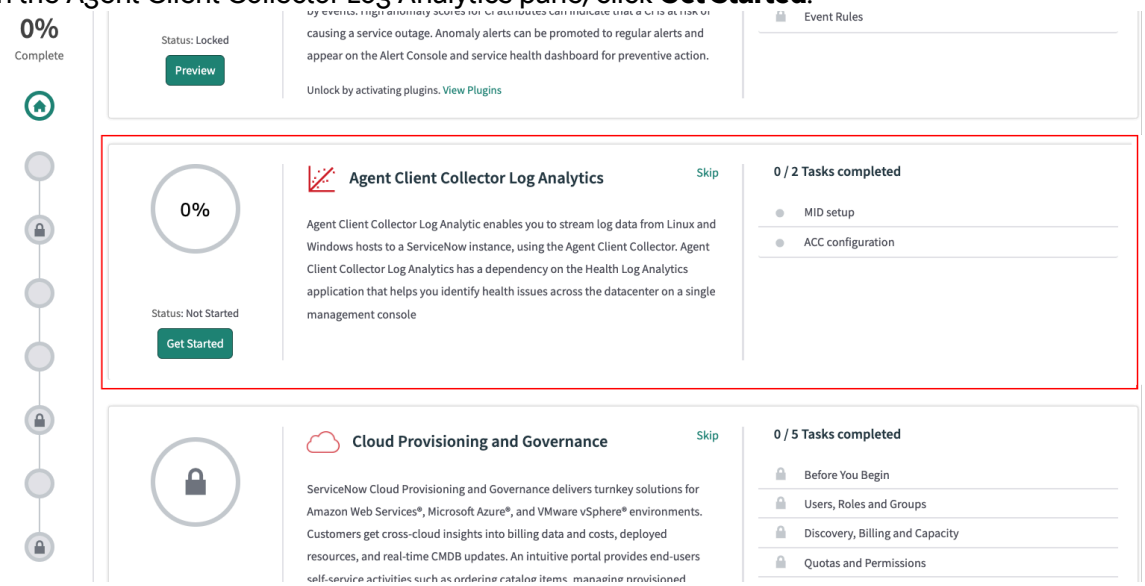

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the **MID setup** section, take the appropriate action:
   - If you have already configured a MID Server, select **Mark as Complete**.
   - If you have not configured a MID Server, select **Configure** and set up a MID Server, then select **Mark as Complete**. For more information, see Use MID Server guided setup.

   Note: To stream data into the system, you must have a MID Server with the Log Ingestion capability enabled.

   When the MID Server is configured, the status indicator for Agent Client Collector Log Analytics guided setup shows that 64% of the configuration is complete. The progress indicator shows the completion percentage for all the IT Operations Management tasks.

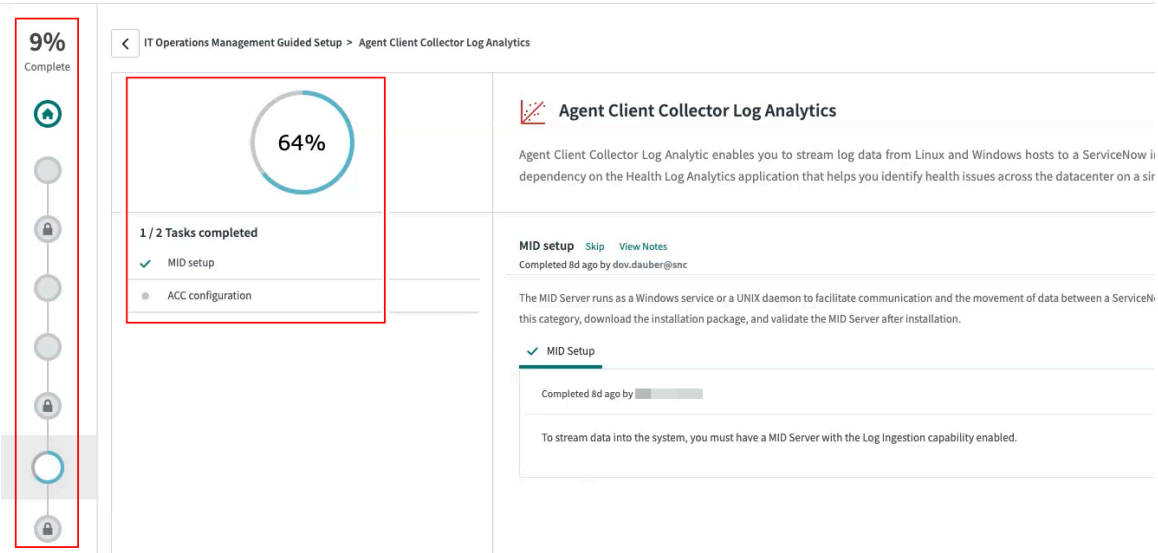

5. Configure Agent Client Collector Log Analytics.

   a. In the **ACC Configuration** section, select a setup task and read the instructions in the help pane.

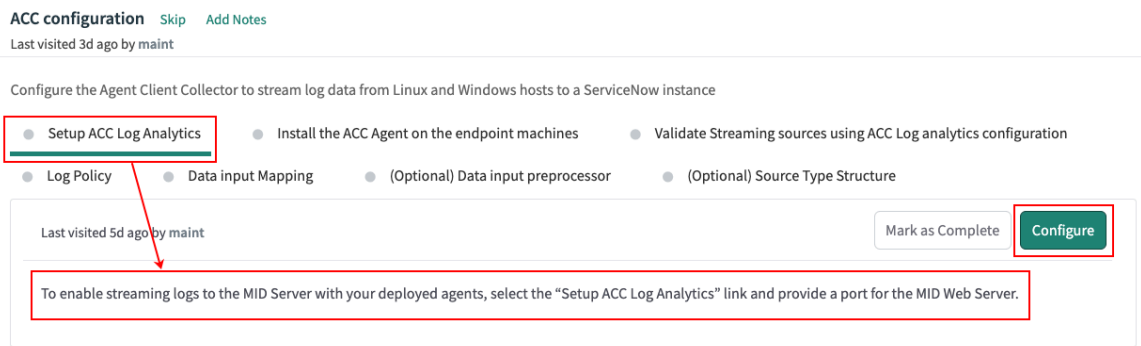

   b. Select **Configure** and complete the task.
   For detailed information about each setup task, consult the Agent Client Collector Log Analytics product documentation.

   c. Complete all required setup tasks in the task list.

   d. Select **Mark as Complete** to complete the Agent Client Collector Log Analytics configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up Agent Client Collector Log Analytics without using guided setup

Configure the Agent Client Collector to stream log data from Linux and Windows hosts to a ServiceNow instance without using guided setup.

Before you begin

Note: You can configure Agent Client Collector Log Analytics using guided setup. The guided setup ensures that you have the minimum required setup for using the application. For more information, see Set up Agent Client Collector Log Analytics using guided setup.

- The Agent Client Collector Log Analytics (ACC-L) plugin must be installed. The plugin comes with the Agent Client Collector Framework (ACC-F) and Agent Client Collector Monitoring (ACC-M) plugins. For more information, see Agent Client Collector installation.
- ServiceNow Service Operations Workspace ITOM Apps must be installed. If it is not installed already, install it manually. For more information, see Install Service Operations Workspace for ITOM application.
- Agent Client Collector Log Analytics has a dependency on the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. Streaming log data from Windows hosts is supported in Agent Client Collector Log Analytics Version 3.1.0 and later.

  Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

- The Agent Client Collector comes with the default *`servicenow`* user. Ensure that this user has read access to enable Agent Client Collector to view all the configured log paths. For example, the Agent Client Collector *`servicenow`* user that comes installed with the base system does not have permissions to view the paths to `/var/log/` in Linux and `C:\Windows\System32` in Windows. For information about configuring permissions for the *`servicenow`* user, see the ACC-L Permission Denied issues [KB1117271] article in the Now Support Knowledge Base.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **MID Servers**.
2. On the MID Server page, select a MID Server that is validated and has a Status value of **Up**.
3. On the MID Server page for the selected server, select **Setup ACC Log Analytics** in the related links section.
   Agent Client Collector Log Analytics setup starts.

   In addition, the system automatically creates an ACC data input. When it has been created, the data input record is added to the ACC data inputs table and the data input is ready to stream logs.

   Note: It is important to configure an ACC data input for every MID Server to which Agent Client Collectors can connect. For more information, see Set up additional ACC data inputs.

   The system also prompts you to set up the MID Web Server. For more information, see MID Web Server.
4. In the ACC Log Analytics Setup dialog box, configure the ports.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **MID Web Server Port** and **ACC data input Port** fields, enter the appropriate port numbers.

b. Select **OK**.
The URL with the configured endpoint appears in the dialog box.

c. Copy and save the endpoint URL for use when installing the ACC agent on the endpoint machine.

5. Navigate to **All** > **Agent Client Collector** > **MID Web Server API Key** and copy and save the MID Web Server API key for use when installing the ACC agent on the endpoint machine.
6. Install the ACC agent on the endpoint machine.

Note: If you have already installed the ACC agent, proceed to Agent Client Collector log policies.

a. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Downloads**.

b. Download the ACC agent for the appropriate operating system.
For detailed instructions, select the documentation link for the operating system. For example:

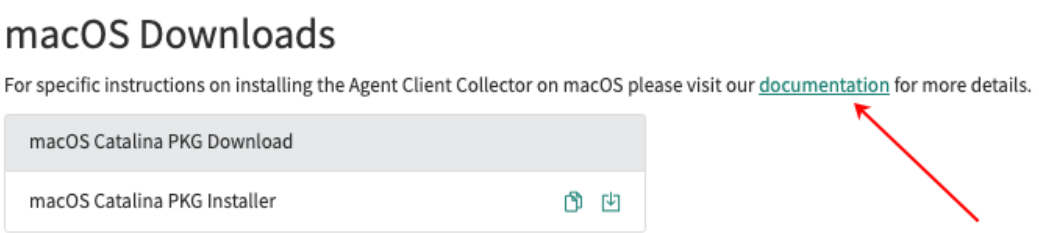

Result
Agent Client Collector Log Analytics and the ACC agent are installed. You can start using the application.

What to do next
Review the Agent Client Collector log policies.

Identify and resolve log streaming issues

Make sure that your ACC data inputs are streaming data properly by identifying and resolving log streaming issues.

Before you begin
Role required: agent_client_collector_admin

About this task
In the ACC data input record, the Streaming Sources related list shows the streaming data of that data input's sources. This data enables you to identify streaming issues and their possible cause.

For example, if the last event time for a data input's endpoint server is yesterday, the server might be down or configured incorrectly. A streaming issue might also be caused by the data input configuration file not being installed on the endpoint.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can view the streaming data of all data inputs and the MID Servers that receive logs from them in the Streaming Sources page. For more information, see Identify and resolve a log streaming issue in Health Log Analytics in the ServiceNow Health Log Analytics application documentation.

Procedure

1. Navigate to **All** > **ACC Data Inputs** and select a data input record.
2. On the data input form, select the Streaming Sources related list.
3. Review the streaming data of that data input's sources.
   For a description of the filters, see ACC data input streaming sources.
4. Resolve any data streaming issues.
   a. Navigate to **All** > **ACC Log Analytics** > **ACC Log Alerts**.
   b. Open alerts and address the specified issues.

What to do next

Proceed to the next step in the Health Log Analytics data input setup flow. For more information, see Set up additional ACC data inputs.

ACC data input streaming sources

In the ACC data input record, the Streaming Sources related list shows the streaming data of that data input's sources.

Streaming Sources related list

Filter	Description
Status	The status of the source. A red bullet indicates that this source has not streamed data in the last hour.
Last event time	The last recorded time an event arrived at the MID Server in the last one-minute interval. Health Log Analytics continuously updates the last event time. If the last event time is not up to date, data is not streaming.
Raw log lines/sec	The average number of raw log lines that streamed to the MID Server per second in the last one-minute interval. **Note:** This value represents the number of raw log lines before preprocessing.
Preprocessed log lines/sec	The average number of preprocessed log lines that streamed to the MID Server per second in the last one-minute interval.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Streaming Sources related list (continued)

Filter	Description
	Note: This value can differ from the number of raw log lines per second. For example, the difference can be a result of logs having been dropped during preprocessing.

View log shipper alerts

View all log shipper check alerts for Agent Client Collector Log Analytics (ACC-L) agents.

Before you begin

Role required: agent_client_collector_admin

About this task

In the ACC Alerts list, you can see at a glance if any issues are preventing the log shipper from collecting data. The list shows all log shipper-related alerts in one place and provides information to help you resolve the issues.

The ACC Alerts list only includes alerts about the success or failure of log shipper execution on a CI on a specific agent. These alerts are a subset of the alerts shown in the Event Management Alerts list, which includes alerts of all types. For more information about the Event Management Alerts list, see View alert information.

Log shipper alerts are generated from log shipper events with a severity of Major or Critical. These alerts are typically related to log path issues or a lack of permissions.

Note: If an **OK** event on the log shipper is received on a CI where an open log shipper alert exists, the open log shipper alert on that CI is closed.

Log shipper Alerts list

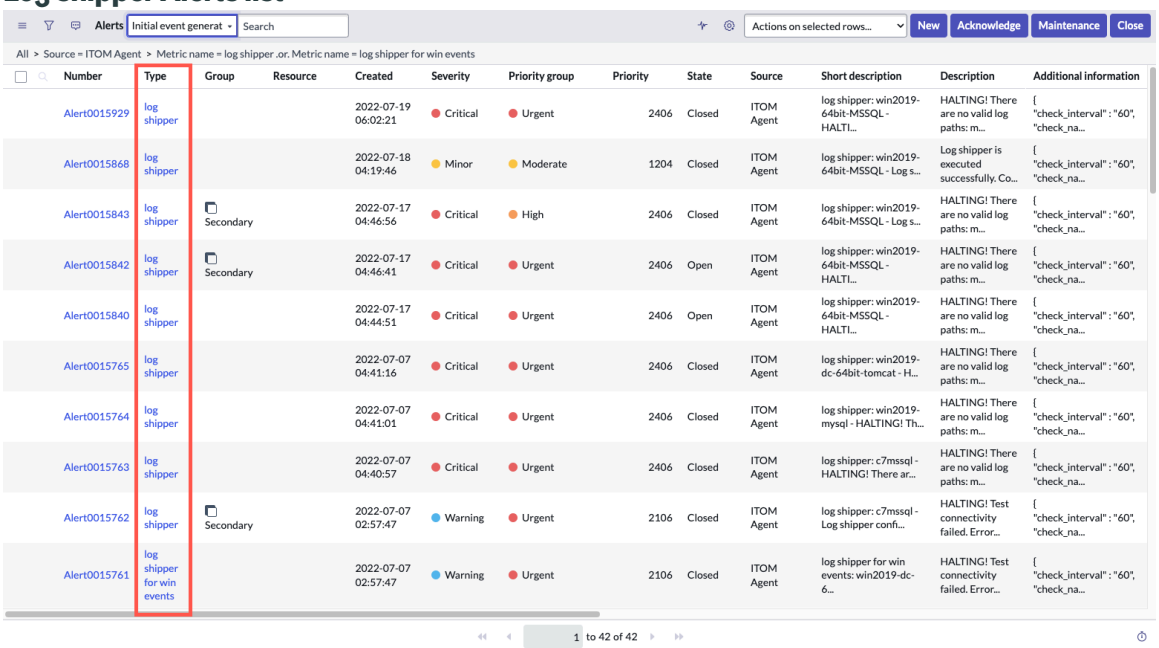

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **ACC Log Analytics** > **ACC Log Alerts**.
2. Review the information provided in the Alerts list.
   For details and a description of the columns, see Monitor incoming alerts.

Related topics

View log shipper events

View log shipper events

View all log shipper events with all severity levels for Agent Client Collector Log Analytics (ACC-L) agents.

Before you begin
Role required: agent_client_collector_admin

About this task
The ACC Events list shows all log shipper-related events in one place. The list only includes events about the success or failure of log shipper execution on a CI on a specific agent. These events are a subset of the events shown in the Event Management Events list, which includes events of all types.

Procedure

1. Navigate to **All** > **ACC Log Analytics** > **ACC Log Events**.
2. Review the information provided in the Events list.
   For details and a description of the columns, see View events.

Related topics

View log shipper alerts

Collect web usage data using Agent Client Collector for Visibility Content

Track website visits across your organization and gain visibility into web applications your users access using Agent Client Collector for Visibility Content (ACC-VC).

Before you begin
Verify that the following components are installed in your environment:

- Agent Client Collector Framework (ACC-F)
- Agent Client Collector for Visibility Content
- URL Discovery and Monitoring (sn_itom_url_disc)
- The DEX Browser Extension 4.1.0 is enabled on the host

Role required: sn_acc_vis_content.application_admin and discovery_admin

About this task
Web data collection is enabled on Windows and macOS platforms.

Web data collection is supported on Google Chrome and Microsoft Edge browsers.

ACC-VC can collect web usage data in the following modes:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Configured monitoring (tracks only pre-configured websites)
- Full monitoring (captures domain-level website access activity from supported Chromium-based browsers on managed corporate devices). Automatically captures domain-level usage data from managed Windows and macOS devices and stores it in the User Web Usage (sn_acc_vis_content_user_web_usage) table.

For details on URL Discovery, see URL Discovery insights dashboard and Discovery Admin Workspace Settings.

Procedure

1. Navigate to **System Definition** > **System Properties**.
2. Search for the property *sn_acc_vis_content.enable_full_monitoring*.
3. Set the **Value** field to `true`.
   Setting this property to true automatically activates the VISC Get URL Metrics monitoring policy and enables full monitoring of web usage data for all websites visited on Chromium-based browsers. Setting it to false deactivates the policy.
4. Select **Update**.
   The system activates the VISC Get URL Metrics monitoring policy and begins collecting web usage data from managed devices.
5. **Optional:** Configure the data retention period by searching for the property *sn_acc_vis_content.web_usage_retention_days* and setting the desired number of days.
   The default retention period is 7 days. Records older than this value are automatically deleted. Modifying this value updates the auto-flush retention period for the User Web Usage (sn_acc_vis_content_user_web_usage) and Application Domain (sn_acc_vis_content_application_domain) tables.

Result

The ACC agent on each managed device monitors URL activity from Chromium-based browsers and stores the data locally. At scheduled intervals (default: once every 24 hours), the agent runs the VISC Get URL Metrics check and sends the data to your instance. The data is processed and stored in the User Web Usage table, with daily aggregation jobs updating domain-level totals. For details on the VISC Get URL Metrics policy and check, see Agent Client Collector for Visibility Content default checks and policies.

What to do next

After enabling web usage collection, you can view the collected data in the User Web Usage (sn_acc_vis_content_user_web_usage) table and monitor aggregated domain-level statistics in the Application Domain (sn_acc_vis_content_application_domain) table. Data aggregates from the User Web Usage table to the Application Domain table, including total page views, total unique users, last accessed time, and daily page view trends. By default, this job runs daily.

For details on these tables, see Web usage data collection tables and fields.

Deploying Agent Client Collector on endpoints

When deploying the Agent Client Collector, perform deployment and management tasks on your endpoints.

Endpoint deployment connects agents directly to the ServiceNow® cloud through ITOM Cloud Services (ICS). Agents on laptops, remote offices, and distributed sites reach ICS over the internet using a secure connection, without a MID Server. ServiceNow® manages the infrastructure and issues certificates automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Choose endpoint deployment when your agents operate beyond the data center, where maintaining MID Servers isn't practical. Digital End-User Experience (DEX) requires this architecture, and Agent Client Collector for Visibility Content (ACC-VC) supports it in specific use cases.

Configuring MID-less Agent Client Collector

Configure MID-less Agent Client Collector to enable sending information through the cloud. Sending information through the cloud allows the MID Server to be used for more persistent resources.

Configure MID-less Agent Client Collector using a single-line command

Use a single-line command to set up a MID-less Agent Client Collector. You can also use the single-line command to migrate agents configured with a MID Server to have a MID-less configuration.

Before you begin

Role required:

- In a Linux or macOS environment: root
- In a Windows environment: agent_client_collector_admin

About this task

Running the single-line command includes flags, which indicate the options to be added to the `acc.yml` configuration file. These options dictate the configuration of the MID-less Agent Client Collector. If you don't specify an option in the command, the existing option in the `acc.yml` file is used. If there's no existing option, an error message indicates that the parameter is required to enable connecting to an endpoint.

Running a command validates that a connection can be made with the specified parameter (for example, the specified registration-key value). However, the command doesn't verify if the value of the parameter is valid.

Procedure

1. Run the following command in a terminal window: `acc gateway ics <flags>`.
   The available flags appear in the following table:

MID-less configuration flags

Flag	Type	Description
-y, <--acc-yml>	String	Location of the `acc.yml` being configured. Can also be used to modify and validate parameters if the file is in a different location that the default.Can also be used to modify a different file location to check if the parameters specified are valid. Optional. If no value is entered, the default location of the `acc.yml` file is used.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Flag	Type	Description
-b, <--backend-url>	String	The gRPC URL for the ITOM cloud services (ICS) gateway.
-i, <--instance-url>	String	The instance URL.
-r, <registration-key>	String	The registration key for the ITOM cloud services (ICS) gateway.
-h, help	N/A	Explains what the specified command does.

2. Restart the agent for the configuration settings to be entered into the `acc.yml` file.

Installing MID-less Agent Client Collector

Installing MID-less Agent Client Collector (ACC) enables you to send data from the agent to the instance through the cloud. Sending information through the cloud allows the MID Server to be used for more persistent resources.

Note: MID-less ACC requires ServiceNow cloud infrastructure and is not supported for customer-hosted, on-premises ServiceNow instances. For on-premises instances, install Agent Client Collector with a MID Server, as described in Configuring Agent Client Collector with a MID Server.

Before using MID-less Agent Client Collector, do the following:

1. Ensure that your ServiceNow instance has migrated from an F5 load balancer to Application Delivery Controller, version 2 (ADCv2) and has mTLS support.
   a. Enter the following command in your terminal: `curl -v https://<instance>.service-now.com/stats.do 2>&1 | grep "< Server"`.
      - If the system returns **snow_adc**, the instance is on ADCv2/ADCv3.
      - If the system doesn't return **snow_adc**, the instance isn't on ADCv2/ADCv3. To migrate the instance to ADCv2/ADCv3, contact Customer Support.
   b. Navigate to `<instance_url>/adcv2/supports_tls`

      If the URL returns a value of **true**, mTLS is supported.
2. Install the ITOM Cloud Services Core (sn_itom_cloud_svc) scoped app plugin.

   If you're using Agent Client Collector for Visibility Content, install the Agent Client Collector for Visibility Content (sn_acc_vis_content) plugin.
3. Configure an agent registration key on the instance. For details, see Configure an agent registration key.

   If registration fails between the agent and the instance, the registration key is cleared from the `acc.yml` configuration file.

Install MID-less Agent Client Collector using a single-line command in a Windows environment

Install MID-less Agent Client Collector on a Windows machine to enable sending data from the agent to the instance through the cloud. Sending information through the cloud allows the MID Server to be used for more persistent resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Install the ITOM Cloud Services Core (sn_itom_cloud_svc) plugin.
- Onboard your instance to use ITOM Cloud Services. For details, contact Customer Support.
- Configure an agent registration key. For details, see Configure an agent registration key.
- Role required: agent_client_collector_admin

About this task

For details on sending data through the cloud in a MID-less deployment, see DEX Architecture.

Procedure

1. Retrieve the publicly accessible gateway URL, based on your location.
   - AMER (Americas): `itomcnc-prod-gateway-amer.sncapps.service-now.com:443`
   - EMEA (Europe): `itomcnc-prod-gateway-emea.sncapps.service-now.com:443`
   - APAC (Asia Pacific): `itomcnc-prod-gateway-apac.sncapps.service-now.com:443`
2. On the Windows server where the agent is installed, enter the following command:

```
msiexec /i <msi_file_path> /quiet /qn /norestart
 CONNECT_WITHOUT_MID="true" ACC_CNC="<gateway_endpoint>"
 REGISTRATION_KEY="<registration_key>"
 INSTANCE_URL="https://<instance_url>"
```

   For details on the parameter values in the command, see Agent Client Collector MID-less installation command parameters.

Install MID-less Agent Client Collector manually in a Linux environment

Install MID-less Agent Client Collector manually on a Linux machine to enable sending data from the agent to the instance through the cloud. Sending information through the cloud allows the MID Server to be used for more persistent resources.

Before you begin

- Install the ITOM Cloud Services Core (sn_itom_cloud_svc) plugin.
- Onboard your instance to use ITOM Cloud Services. For details, contact Customer Support.
- Configure an agent registration key. For details, see Configure an agent registration key.
- Role required: agent_client_collector_admin

About this task

For details on sending data through the cloud in a MID-less deployment, see DEX Architecture.

Procedure

1. Download the relevant installation package.
   a. Navigate to > **Agent Client Collector** > **Deployment** > **Agent Downloads**.
   b. Download the relevant .rpm or .deb installation file.
   c. Download the relevant signature files to validate the installation files.
2. **Optional:** Verify the package signature.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Extract the installation file (if it is compressed) by running the following command: `unzip agent_client-collector-<version_number>-x86_64-rpm-rpm.zip`

b. Validate the installation file signature by running the indicated commands.

- On an RPM-based system:

```
openssl dgst -sha256 -verify {<ServiceNow
 DGST pem key>} -signature {<signature file>}
 agent-client-collector-<version number>-x86_64.rpm
```

Where `<ServiceNow DGST pem key>` is the `.pem` file extracted from the `.zip` file, and `<signature file>` is the `.bin` file extracted from the `.zip` file.

For example:

```
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-x86_64.rpm
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-x86_64-rpm-rpm.zip
$ unzip agent-client-collector-3.0.0-x86_64-rpm-rpm.zip
Archive: agent-client-collector-3.0.0-x86_64-rpm-rpm.zip
inflating: ServiceNow_Digicert_DGST.pem
extracting: agent-client-collector-3.0.0-x86_64.bin
$ openssl dgst -sha256 -verify ServiceNow_Digicert_DGST.pem
 -signature agent-client-collector-3.0.0-x86_64.bin
 agent-client-collector-3.0.0-x86_64.rpm
Verified OK
```

Note: Each command must appear on a single line.

- On a Debian-based system:

(Optional) `gpg --import ServiceNow_Digicert_Public.gpg`

`sudo gpg --verify agent-client-collector-<version number>-<distro>_amd64.deb`

For example:

```
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-debian-9_amd64.deb
$ curl -sLO
 https://install.service-now.com/glide/distribution/builds/package/app-signed/agent-client-collector-3.0.0-debian-9_amd64-deb-deb.zip
$ unzip
 agent-client-collector-3.0.0-debian-9_amd64-deb-deb.zip
Archive:
 agent-client-collector-3.0.0-debian-9_amd64-deb-deb.zip
extracting: ServiceNow_Digicert_Public.gpg
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
$ gpg --import ServiceNow_Digicert_Public.gpg
gpg: /home/admin/.gnupg/trustdb.gpg: trustdb created
gpg: key 985DD52C6A0ABB45: public key "ServiceNow, Inc.
 (Signing) <seceng@servicenow.com>"
imported
gpg: Total number processed: 1
gpg: imported: 1
$ dpkg-sig --verify
 agent-client-collector-3.0.0-debian-9_amd64.deb
Processing
 agent-client-collector-3.0.0-debian-9_amd64.deb...
GOODSIG _gpgbuilder 9B928FB49771DF6C047430DD985DD52C6A0ABB45
 1665054068
```

Note: Each command must appear on a single line.

3. Install the Agent Client Collector package using the package manager associated with Linux distribution.

OS	Command
RHEL-based	`yum / dnf localinstall`
SLES	`zypper install`
Debian-based	`apt-get install`

Alternatively, if these commands are not configured correctly, you can use the core commands that are configured to run with the package manager commands.

- RPM-based system: `# rpm -vi agent-client-collector-<version number>-x86_64.rpm`
- Debian-based system: `# dpkg -i agent-client-collector-<version number>-<distro>_amd64.deb`

Verify that the package commands are configured correctly with your system administrator.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Some file systems may have restrictions enabled; for example, `/var/` may be mounted with a *noexec* flag. Because the agent must execute Agent Client Collector plugins that are normally store in the `/var/cache` directory, you must deploy the application into specific folders by customizing the installation paths using the `--relocate` option as an `.rpm` parameter.

For example: `rpm -i --relocate /var/cache=/opt/cache agent-client-collector-<version_number>-x86_64.rpm`

The following paths can be relocated:

Path	Notes
`/etc`	When updating, you must also update the *allow-list* parameter in the `acc.yml` file with the new path.
`/usr/share`	N/A
`/var/cache`	Updating the `/var` directory retains all `/var` subdirectories, nested under the new directory.
`/var/log`	
`/var/run`	
`/var`	

Review the paths in `/usr/lib/systemd/system/acc.service` to ensure that they appear as expected.

4. Run the relevant package upgrade command, based on the cpu your machine runs with:
   - Intel cpu:

     ```
     sudo installer -pkg agent-client-collector-<version
      number>-macos_x64.pkg -target /
     ```

   - Apple silicon cpu:

     ```
     sudo installer -pkg agent-client-collector-<version
      number>-macos_arm64.pkg -target /
     ```

   This command also backs up and restores the agent-id and configuration files before upgrading.

5. In the `acc.yml` configuration file:

   a. Set the following parameters:

Parameter	Value
backend-url	"<gateway URL>"
connect-without-mid	true
instance-url	"<instance URL>"
registration-key	"<regsitration-key>"

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Value
insecure-skip-tls-verify	false

b. Comment out the *api-key* parameter.

6. Save the configuration file.

7. Configure sudoers.

Configuration is typically automated by your Linux sysadmin. To configure manually if you want to ensure correctness before full configuration, run the following.

```
# visudo -f /etc/sudoers.d/01_servicenow
User_Alias ACC_USERS = servicenow
Cmnd_Alias ACC_CMD = /usr/sbin/dmidecode -s
 baseboard-serial-number,/usr/sbin/dmidecode
 -s chassis-serial-number,/usr/sbin/dmidecode
 -s system-serial-number,/usr/sbin/dmidecode -s
 system-uuid,/usr/sbin/ss -tanp
ACC_USERS ALL = (root) NOPASSWD:ACC_CMD
Defaults:ACC_USERS !requiretty
```

8. Configure the agent to run as a service.

a. Safeguard resource consumption by adding the indicated values to the `/usr/lib/systemd/system/acc.service` file.

- CPUShares=128
- CPUQuota=10%
- MemoryLimit=192M
- BlockIOWeight=10
- LimitNICE=15

For example:

```
# vi /usr/lib/systemd/system/acc.service
[Unit]
Description=Agent-Now acc
After=network-online.target
[Service]
Environment=AGENT_ROOT=/usr/share
Environment=AGENT_CACHE_ROOT=/var/cache
Environment=AGENT_CONFIG_ROOT=/etc
Environment=AGENT_LOG_ROOT=/var/log
Environment=AGENT_RUN_ROOT=/var/run
Environment=RUBYOPT=-Eutf-8
User=servicenow
Group=servicenow
ExecStart=/usr/share/servicenow/agent-client-collector/bin/a
cc-service start acc
KillMode=process
Restart=on-failure
RestartSec=1min
CPUShares=128
CPUQuota=10%
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
MemoryLimit=192M
BlockIOWeight=10
LimitNICE=+15
[Install]
WantedBy=network-online.target
```

b. Run the reload daemon command If you modified the service file after enabling the acc service.

```
# systemctl daemon-reload
```

c. Enable and start the service by running the following commands.

```
# systemctl enable acc
```

```
# systemctl start acc
```

Install MID-less Agent Client Collector using a single-line command in a Linux environment

Install MID-less Agent Client Collector on a Linux machine to enable sending data from the agent to the instance through the cloud. Sending information through the cloud allows the MID Server to be used for more persistent resources.

Before you begin

- Install the ITOM Cloud Services Core (sn_itom_cloud_svc) plugin.
- Onboard your instance to use ITOM Cloud Services. For details, contact Customer Support.
- Configure an agent registration key. For details, see Configure an agent registration key.
- Role required: agent_client_collector_admin

About this task

For details on sending data through the cloud in a MID-less deployment, see DEX Architecture.

Procedure

1. Retrieve the agent registration key:
   a. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Registration Key**.
   b. Select the relevant agent registration key.
   c. Copy the registration key value and store it in a place where you can easily retrieve it, when needed.
2. Retrieve the publicly accessible gateway URL, based on your location.
   - AMER (Americas): `itomcnc-prod-gateway-amer.sncapps.service-now.com:443`
   - EMEA (Europe): `itomcnc-prod-gateway-emea.sncapps.service-now.com:443`
   - APAC (Asia Pacific): `itomcnc-prod-gateway-apac.sncapps.service-now.com:443`
3. On the Linux server where the agent is installed, enter the following command:

```
CONNECT_WITHOUT_MID="true" ACC_CNC="<gateway_endpoint>"
 REGISTRATION_KEY="<registration_key>"
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
INSTANCE_URL="https://<instance_url>" bash -c "$(curl -L
https://<instance_url>/api/sn_agent/agents/install_agent)"
```

For details on the parameter values in the command, see Agent Client Collector MID-less installation command parameters.

Install MID-less Agent Client Collector manually in a macOS environment

Install MID-less Agent Client Collector manually on a macOS machine to enable sending data from the agent to the instance through the cloud. Sending information through the cloud allows the MID Server to be used for more persistent resources.

Before you begin

- Install the ITOM Cloud Services Core (sn_itom_cloud_svc) plugin.
- Onboard your instance to use ITOM Cloud Services. For details, contact Customer Support.
- Configure an agent registration key. For details, see Configure an agent registration key.
- Role required: agent_client_collector_admin

About this task

For details on sending data through the cloud in a MID-less deployment, see DEX Architecture.

Procedure

1. Download the installation files to your local machine.

```
curl -L <download link from the Agent Downloads page on the
instance>
```

2. Validate the installation file signature.

```
pkgutil --check-signature agent-client-collector-<version
number>-macos-catalina_x64.pkg
```

Note: Ensure that the `pkg` certificate chain has the ServiceNow developer ID installer certificate.

3. Install the Agent Client Collector package.

```
sudo installer -pkg agent-client-collector-<version
number>-macos-catalina_x64.pkg -target /
```

4. Copy the installation files.

   a. Sample configuration file:
   ```
   sudo cp -p /Library/Application\ Support/servicenow/agent-
   client-collector/acc.yml.example /Library/Application\ Support/
   servicenow/agent-client-collector/acc.yml
   ```

   b. Sample allow list file:
   ```
   sudo cp -p /Library/Application\ Support/servicenow/agent-
   client-collector/check-allow-list.json.default /Library/
   Application\ Support/servicenow/agent-client-collector/check-
   allow-list.json
   ```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. **Optional:** To create new check definitions to be included in the allow list, first disable the allow list by commenting out the *allow-list* parameter in the `acc.yml` file.

   (Optional) `sudo -H -u _servicenow vi /Library/Application\ Support/servicenow/agent-client-collector/acc.yml`

   Note: This action compromises your system's security and is recommended only as a temporary measure until you finish updating the allow list. After completing the update, uncomment the allow list parameter in the configuration file to re-activate the allow list.

   For details on adding custom checks to an allow list, see Generate an Agent Client Collector allowlist.

6. In the `acc.yml` configuration file:

   a. Set the following parameters:

Parameter	Value
backend-url	"<gateway URL>"
connect-without-mid	true
instance-url	"<instance URL>"
registration-key	"<regsitration-key>"
insecure-skip-tls-verify	false

   b. Comment out the *api-key* parameter.

7. Save the configuration file.
8. To ensure that the configuration file points to a valid endpoint, run the following commands to restart the Agent Client Collector:

```
sudo launchctl unload
 -w /Library/LaunchDaemons/com.sn.acc.plist
sudo launchctl load -w /Library/LaunchDaemons/com.sn.acc.plist
```

   Note: By default, exiting the installation program causes the program to restart immediately. To ensure that the program remains stopped after exiting, set the KeepAlive property in the `/Library/LaunchDaemons/com.sn.acc.plist` file to **false**.

Install MID-less Agent Client Collector using a single-line command in a macOS environment

Install MID-less Agent Client Collector on a macOS machine to enable sending data from the agent to the instance through the cloud. Sending information through the cloud allows the MID Server to be used for more persistent resources.

Before you begin

- Install the ITOM Cloud Services Core (sn_itom_cloud_svc) plugin.
- Onboard your instance to use ITOM Cloud Services. For details, contact Customer Support.
- Configure an agent registration key. For details, see Configure an agent registration key.
- Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

For details on sending data through the cloud in a MID-less deployment, see DEX Architecture.

Procedure

1. Retrieve the agent registration key:
   a. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Registration Key**.
   b. Select the relevant agent registration key.
   c. Copy the registration key value and store it in a place where you can easily retrieve it, when needed.
2. Retrieve the publicly accessible gateway URL, based on your location.
   - AMER (Americas): `itomcnc-prod-gateway-amer.sncapps.service-now.com:443`
   - EMEA (Europe): `itomcnc-prod-gateway-emea.sncapps.service-now.com:443`
   - APAC (Asia Pacific): `itomcnc-prod-gateway-apac.sncapps.service-now.com:443`
3. On the macOS server where the agent is installed, enter the following command:

```
CONNECT_WITHOUT_MID="true" ACC_CNC="<gateway_endpoint>"
 REGISTRATION_KEY="<registration_key>"
 INSTANCE_URL="https://<instance_url>" bash -c "$(curl -L
 https://<instance_url>/api/sn_agent/agents/install_agent)"
```

For details on the parameter values in the command, see Agent Client Collector MID-less installation command parameters.

Configure an agent registration key

Configure an agent registration key so that you can deploy MID-less Agent Client Collector. Deploying MID-less Agent Client Collector enables you to use the MID Server for more persistent resources.

Before you begin

Install the plugins necessary for deploying MID-less Agent Client Collector.

Role required: agent_client_collector_admin

About this task

The Cleanup Agent Registration Keys scheduled job runs daily and deletes expired registration keys. The length of time for which a registration key is valid is controlled by the *sn_agent.registration_key_validity.days* property, based on when the key was created. The amount of time that the customer is notified prior to upcoming key expiration is controlled by the *sn_agent.registration_key_notification.days* property.

If a key's expiration date is approaching, create a new registration key and enter it into any existing automation jobs to prevent registration outage.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Registration Key**.
2. Select **New**.
3. In the **Name** field, enter a descriptive name for the registration key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. **Optional:** To receive notifications when a registration key is soon to expire, select a group with a valid email address in the **Ownership Group** field.
   When keys are deleted via the Cleanup Agent Registration Keys scheduled job, the indicated group receives an email notification.
5. Select **Submit**.

Result

The new agent registration key is ready to use for installation.

Related topics

Installing MID-less Agent Client Collector

Agent Client Collector Framework configuration properties

Agent Client Collector installation on a macOS system

Install Agent Client Collector on a system that uses macOS. You can either use a single-line command script or follow a manual installation procedure if the agent is not connected to the instance or you want enhanced customization options.

To enable privileged access to information on the macOS system, add the *_servcienow* user account to the admin group (or equivalent).

Incorporating the Agent Client Collector into a custom base image for mass deployment is not supported on macOS.

Perform a single-line Agent Client Collector installation on macOS when using a MID Server

Use an efficient single-line command script to install Agent Client Collector on a machine that uses macOS operating system. Use the manual installation procedure if the agent is not connected to the instance or you want enhanced customization options.

Before you begin

- Ensure that you have configured the Agent Client Collector web server.
- Verify that your server's OS and version matches the list of supported OS versions as described in Agent Client Collector Installation.
- Ensure that you have the *pkgutil* packages available, to be used by the install script.
- Enable golden image mode for cloning additional agents by setting the golden image marker located at `/Library/Preferences/acc.rc.plist`.
- Ensure that the target agents have connectivity to the following:
  - ServiceNow install server [install.service-now.com]
  - ServiceNow instance
  - The MID Server to which the agent will connect

Role required: agent_client_collector_admin

Procedure

1. Retrieve the MID Server websocket URL parameters.
   a. Navigate to **Agent Client Collector** > **Websocket Endpoint**.
   b. Select the relevant websocket endpoint.
   c. Copy the **Endpoint URL** field value to be used in the installation command.
2. Retrieve the API Key.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **Agent Client Collector** > **Deployment** > **MID Web Server API Key**.

b. Copy the **API Key** value, to be used in the installation command.

3. Run the following command on the server where the Agent Client Collector is being installed.

```
ACC_API_KEY=<api_key>
 ACC_MID=wss://<mid_ip>:<websocket_port>/ws/events
 bash -c "$(curl -L https://<instance
 URL>/api/sn_agent/agents/install_agent)"
```

Enter the following values in the command line:

Command line parameters

Parameter	Value
ACC_API_KEY	Replace the *<api_key>* string with the API key used to access the MID Server websocket. If no value is specified, you will be prompted to enter the API key value during the script's execution.
ACC_MID	Replace the *<mid_ip>* string with the URL of the MID Server websocket. When specifying multiple MID Servers, enter a string separated by commas for each MID Server.
ACC_ALLOW_LIST	(Optional) Add the string *ACC_ALLOW_LIST=0* to disable the allow list. By default, the allow list is enabled. **Note:** Disabling the allow list compromises your system's security and is recommended only as a temporary measure until you finish updating the allow list. After completing the update, re-activate the allow list by removing *ACC_ALLOW_LIST=0* from the single-line command string.
ACC_VALIDATE_SIG	(Optional) Add the string *ACC_VALIDATE_SIG=0* to disable verifying the installer certificate validation. By default, validation is enabled. Default: Validation is enabled. **Note:** Disable certification validation only when using non-standard signature validation tools.

Manually install Agent Client Collector on macOS

Install Agent Client Collector manually on a machine that uses macOS when the agent is not connected to the instance or you want enhanced customization options.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that you have configured the Agent Client Collector web server.
- Verify that your server's OS and version matches the list of supported OS versions as described in Agent Client Collector Installation.
- Enable golden image mode for cloning additional agents by setting the golden image marker located at `/Library/Preferences/acc.rc.plist`.
- Navigate to **Agent Client Collector** > **Deployment** > **Agent Downloads** and download the signed `.pkg` relevant for your OS.
- Store the `.pkg` installation package on one of your installed agents.

Role required: agent_client_collector_admin

Procedure

1. Download the installation files to your local machine.

```
curl -L <download link from the Agent Downloads page on the
 instance>
```

2. Validate the installation file signature.

```
pkgutil --check-signature agent-client-collector-<version
 number>-macos-catalina_x64.pkg
```

Note: Ensure that the `pkg` certificate chain has the ServiceNow developer ID installer certificate.

3. Install the Agent Client Collector package.

```
sudo installer -pkg agent-client-collector-<version
 number>-macos-catalina_x64.pkg -target /
```

4. Copy the installation files.

   a. Sample configuration file:
   `sudo cp -p /Library/Application\ Support/servicenow/agent-client-collector/acc.yml.example /Library/Application\ Support/servicenow/agent-client-collector/acc.yml`

   b. Sample allow list file:
   `sudo cp -p /Library/Application\ Support/servicenow/agent-client-collector/check-allow-list.json.default /Library/Application\ Support/servicenow/agent-client-collector/check-allow-list.json`

5. **Optional:** To create new check definitions to be included in the allow list, first disable the allow list by commenting out the *`allow-list`* parameter in the `acc.yml` file.

   (Optional) `sudo -H -u _servicenow vi /Library/Application\ Support/servicenow/agent-client-collector/acc.yml`

Note: This action compromises your system's security and is recommended only as a temporary measure until you finish updating the allow list. After completing the update, uncomment the allow list parameter in the configuration file to re-activate the allow list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For details on adding custom checks to an allow list, see Generate an Agent Client Collector allowlist.

6. Save the configuration file.
7. In the configuration file, connect the agent with the MID Server.

   a. Set the MID Server connection.

   ```
   "wss://<MID Server IP>:<Web Socket Port>/ws/events"
   ```

   The web socket port was configured on the MID Server when it was set as an Agent Client Collector listener.

   b. Locate the *api-key* parameter and specify the API key value inside double quotation marks, as follows:
   api-key:"111111111"
   The specified API key is encrypted the first time the agent runs.

8. To ensure that the configuration file points to a valid endpoint, run the following commands to restart the Agent Client Collector:

```
sudo launchctl unload
 -w /Library/LaunchDaemons/com.sn.acc.plist
sudo launchctl load -w /Library/LaunchDaemons/com.sn.acc.plist
```

Note: By default, exiting the installation program causes the program to restart immediately. To ensure that the program remains stopped after exiting, set the KeepAlive property in the `/Library/LaunchDaemons/com.sn.acc.plist` file to **false**.

Upgrade Agent Client Collector manually on a macOS system

Perform a manual upgrade of your existing Agent Client Collector version on a system running a macOS. Manual upgrade is useful if the single-line command script isn't connected to the instance or you want to use enhanced customization options.

Before you begin

- The upgrade procedure is the same for both MID Server and MID-less deployments.
- Enable golden image mode for cloning additional agents by setting the golden image marker located at `/Library/Preferences/acc.rc.plist`.
- Role required: agent_client_collector_admin

Procedure

1. **Optional:** Manually back up the `agent-id` and `acc.yml` configuration files by running the following commands.

```
mkdir -p /tmp/acc-upgrade-backup

cp -rp /Library/Application\
 Support/servicenow/agent-client-collector/acc.yml /
tmp/acc-upgrade-backup

cp -rp /Library/Application\
 Support/servicenow/agent-client-collector/check-allow-list.js
on /tmp/acc-upgrade-backup
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
cp -rp /Library/Application\
 Support/servicenow/agent-client-collector/agent_now_id /
tmp/acc-upgrade-backup
```

Manually backing up the configuration files provides an added layer of security by ensuring that your files are backed up before beginning the upgrade process.

2. Navigate to **All** > **Agent Client Collector** > **Deployment** > **Agent Downloads**.
3. In the **macOS Downloads** section, download the relevant agent package and signature `.pkg` files.
4. Perform package upgrade, based on the cpu your macOS runs with.
   - Intel cpu:

     ```
     sudo installer -pkg agent-client-collector-<version
      number>-macos_x64.pkg
     ```

   - Apple silicon cpu:

     ```
     sudo installer -pkg agent-client-collector-<version
      number>-macos_arm64.pkg
     ```

   This command also backs up and restores the agent-id and configuration files before upgrading.
5. Verify that the upgrade was successful.
   a. Navigate to **All** > **Agent Client Collector** > **Agents**.
   b. Verify the following values for your agent:
      - **Host data collection** = Collected
      - **Version** = <Relevant version number>

Uninstall the Agent Client Collector from macOS using a single-line command

Uninstall the Agent Client Collector from macOS machine by running an efficient single-line command. If the agent is connected to the instance, uninstall Agent Client Collector manually.

Before you begin

- The uninstall procedure is the same for both MID Server and MID-less deployments.
- Role required: agent_client_collector_admin

Procedure

Run the following command.

```
bash -c "$(curl -L https://<Instance
 URL>/api/sn_agent/agents/install_agent)" -s "--remove"
```

Note: Uninstalling the Agent Client Collector removes the `acc.yml` file and its directory from your machine. The user and group **_servicenow**, created as part of the installation process, is also removed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Uninstall the Agent Client Collector from a macOS machine manually

Uninstall the Agent Client Collector from a macOS machine manually instead of using the single-line procedure when the specific agent is not connected to the instance.

Before you begin

- The upgrade procedure is the same for both MID Server and MID-less deployments.
- Role required: agent_client_collector_admin

Procedure

1. In a command window, disable and stop the service:
   sudo launchctl unload -w /Library/LaunchDaemons/com.sn.acc.plist
2. Delete the following files/directories:

```
sudo rm /Library/LaunchDaemons/com.sn.acc.plist
sudo rm -rf /opt/servicenow/agent-client-collector
sudo rm -rf /var/run/servicenow/agent-client-collector
sudo rm -rf /Library/Application\
 Support/servicenow/agent-client-collector
```

3. Remove the **_servicenow** user and group.

```
sudo /usr/bin/dscl . -delete /Users/_servicenow
sudo /usr/bin/dscl . -delete /Groups/_servicenow
```

4. Remove the **pkg** from the list of installed packages.

```
sudo pkgutil --forget com.servicenow.agent-client-collector
```

Agent Client Collector File-Based Discovery

Agent Client Collector File-Based Discovery (FBD) scans file systems on managed endpoints to discover installed software and track file inventories.

FBD runs as a lightweight background process within the existing Agent Client Collector agent. It is configured and controlled entirely from the ServiceNow instance. Administrators define which directories to scan, which files to look for, and how often to scan. The agent discovers matching files, collects metadata such as name, path, size, and version, and sends the results back to the instance where they are stored in the appropriate CMDB and application tables.

To activate FBD, navigate to the Discovery Configuration Console (**All** > **Discovery Definition** > **Configuration Console**) and in the **File Based Discovery** section, activate the **Enable File Based Discovery** toggle switch.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery Configuration Console

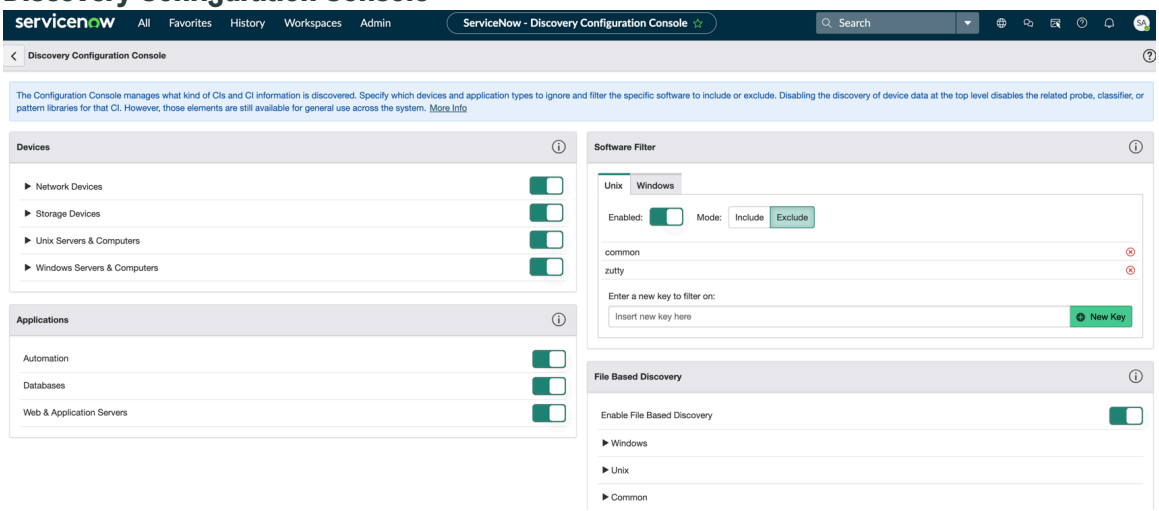

For details on the FBD configuration and system properties, see Agent Client Collector File-Based Discovery properties.

Supported platforms

FBD supports the following operating systems:

- Windows
- Linux
- macOS

Required plugins

The following plugins are required for File-Based Discovery on Agent Client Collector:

- *File-based Discovery* (com.snc.discovery.file_based_discovery)
- *Software Asset Management - File Signature Normalization* (com.snc.file_signature_normalization) — Activated automatically once the File-based Discovery plugin is active
- *Agent Client Collector - Visibility Content* (sn_acc_vis_content)
- *Software Asset Management Professional* (com.snc.samp)

Delta scanning

Delta scanning is a performance optimization that applies only to the File Management policy. It reduces the volume of data sent from the endpoint to the instance on repeat scans.

The delta scanning process works as follows:

1. **First scan** — A full scan runs. All matching files are sent to the instance. A fingerprint, which is a snapshot of all discovered file paths and sizes, is saved locally on the endpoint.
2. **Subsequent scans** — The scanner compares the current file list against the saved fingerprint and reports only the differences:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Added** — files that exist now but were not in the previous fingerprint
- **Modified** — files that exist in both but with a different size
- **Deleted** — files that were in the previous fingerprint but no longer exist on disk

3. **Periodic full scan** — After a configured number of delta scans (default: 5), a full scan is automatically forced to verify accuracy and correct any drift.

Note: Delta scanning does not apply to the SAM or SWID policies. These policies always perform a full scan on each run.

Archive file scanning

Archive file scanning extends file-based discovery to inspect ZIP and JAR archive files and report the files contained within them. This enables discovery of software artifacts such as executables and libraries that are packaged in the ZIP or JAR files on the endpoint.

Archive file scanning is disabled by default. To enable it, set the *sn_acc_vis_content.file_discovery.archive_scan_enabled* property to **true**. This property requires the **discovery_admin** role to modify.

When archive file scanning is enabled, the agent inspects each ZIP or JAR file encountered during a directory scan. The agent reads the archive's internal index to enumerate the files inside. No file content is extracted or decompressed. Only metadata is collected.

For each file found inside an archive, the agent applies the same filtering rules used for regular files on disk, including SAM allowlist and File Management extension rules. Files that pass the filters are reported alongside other discovered files.

For filtering rules and performance safeguards, see Archive file scanning filtering rules and limits.

Running process-based discovery

Running process-based discovery extends File-Based Discovery (FBD) with process-based path detection, enabling the Agent Client Collector for Visibility Content agent to detect software running outside of standard configured scan directories.

Why use running process-based discovery

Some software is never installed into the standard program directories that FBD scans. Common examples include:

- Portable applications run from a removable drive or a user folder
- User-installed tools placed in a personal directory, such as a Python interpreter under a user profile
- Command-line utilities installed into a user's local binary directory

A directory-only scan cannot detect applications installed outside standard scan directories, but they are visible as running processes. Running process-based discovery closes that gap, giving SAM teams a more complete and accurate view of the software in use across your organization. This reduces risk and improves audit readiness.

Platform coverage

Running process-based discovery supports Windows, Linux, and macOS. Coverage scope varies by operating system and depends on the privileges granted to the agent service account. For

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

platform-specific coverage details and privilege requirements, see Running process-based discovery platform coverage and properties.

Enable running process-based discovery

Enable running process-based discovery so that the Agent Client Collector for Visibility Content agent detects software running from directories outside your configured File-Based Discovery scan paths.

Before you begin

File-Based Discovery must be active on the configuration console.

The Agent Client Collector for Visibility Content Windows agent must be deployed on the endpoints.

Role required: disco_admin

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for the *sn_acc_vis_content.file_discovery.fbd_process_scan_enabled* property.
3. Set the **Value** field to `true`.
   The default value is `false`. Setting this property to `true` activates the agent policy that collects process directories.
4. Save the record.

Result

The associated agent policy activates during the next agent synchronization. From that point, the host runs the FBD SAM Process Scan Policy and begins collecting process directories every five minutes. The next daily FBD scan includes the process-discovered directories in its SAM output.

Note: A directory is removed from discovery when its software is no longer present on disk. If the software runs again, it is rediscovered within 24 hours.

What to do next

For Linux and macOS hosts, complete the privilege configuration to obtain full coverage across all users. For details, see Platform coverage.

To exclude specific directories from running process-based discovery, see Exclude directories from running process-based discovery.

Exclude directories from running process-based discovery

Exclude specific directories from running process-based discovery so that processes running from those locations aren't recorded or included in File-Based Discovery scans.

Before you begin

Running process-based discovery must be enabled. For details, see Enable running process-based discovery.

Role required: disco_admin

About this task

Some running processes have no Software Asset Management (SAM) value and only add noise to discovery output, such as security tooling or the agent's own processes. Use the FBD Process Path Exclusions [sn_acc_vis_content_fbd_proc_path_exclude] table to prevent those directories from being recorded. A set of common exclusions is included in the base system.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to the FBD Process Path Exclusions [sn_acc_vis_content_fbd_proc_path_exclude] table.
2. Select **New**.
3. Complete the fields for the exclusion record.

FBD Process Path Exclusions form

Field	Description
Path	Absolute directory prefix to exclude. Any process running from this directory or a subdirectory of it is ignored.
Platform	Operating system the exclusion applies to: **Windows** or **Unix**. Unix covers Linux and macOS. There is no combined option; specify the platform explicitly.
Active	When selected (the default), the exclusion is enforced. Clear this check box to disable the exclusion temporarily without deleting the record.
Description	Optional notes describing the purpose of the exclusion.

4. Select **Submit**.

Result

The exclusion is pushed to the relevant agents automatically. New collection cycles skip the excluded directories. Any previously collected paths that match the exclusion are removed during the next daily maintenance cycle.

Discovering a license key in a file-based discovery framework

License key discovery across Windows, Linux, and macOS enables ACC to automatically identify and create an inventory of license keys stored in files on managed computers. A license key proves that software installation is legitimate and is being used within the boundaries the vendor agreed to at the time of purchase.

What is a license key

A license key is a unique string associated with a piece of software. License keys typically encode constraints such as the number of permitted users or devices, the expiration date, or the product edition or tier purchased.

Organizations track license keys for compliance audits, renewal discussions, and licensing cost management. Without visibility into license keys, organizations can't reliably determine whether they are compliant, or whether they have purchased the right number of licenses.

Why file-based discovery is necessary

License keys are stored in different locations depending on the vendor and product architecture. The following broad storage categories exist:

Storage location	Example vendors/products	Discoverable on-machine?
Windows Registry (Windows only)	Microsoft Office, Windows operating system	Yes — via registry-based discovery (previously released)
Files on disk (cross-platform)	AutoCAD, MATLAB, Quest	Yes — via file-based discovery (current release)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Storage location	Example vendors/products	Discoverable on-machine?
Cloud or user-based authentication	Modern subscription-licensed applications	No — nothing is stored on the machine

Beginning in version 2.0.0, file-based discovery identifies license keys stored in files on any operating system, including Linux and macOS.

How file-based discovery operates

File-based license key discovery acts as a framework rather than a single reader because every vendor's license file format differs. The framework provides the underlying mechanism to find a candidate file and pass its content for processing, while administrators supply configuration that specifies what to look for and how to read it.

This architecture enables support for new vendors and file formats through configuration alone — no platform code changes are required.

Configure a license key discovery rule and write a parser script

Enable license key discovery to verify the legitimacy of your software installation. Set up support for a new vendor's license file format by creating a parser script and defining a matching rule.

Before you begin

Role required: agent_client_collector_admin or discovery_admin

About this task

To discover license keys from a vendor's license file format, you must create the following complementary configuration records:

- A parser script that knows how to extract license information from the specified file format.
- A file-matching rule that tells the framework when to invoke that parser.

The parser is created first because the file-matching rule references it.

Procedure

1. Write the parser script on the **License Parsers** page.
   a. Enter *sn_acc_vis_content_license_parser.list* into the navigation menu.
   b. Select **New**.
   c. Enter a name into the **Parser Name** field.

   The parser is a small piece of server-side JavaScript that receives the raw text content of a matched file and extracts license information from it.

   The script must extract the following information where available:

   - Product name
   - Vendor name
   - License key
   - Expiry date (if present in the file)

   The script must set a variable called `answer` to an array of result objects. Create one object per license entry found in the file — a single file can contain multiple products' licenses. If the script fails to set `answer` as an array, the parser is treated as invalid.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example structure of a result object:

```
{
  product_name: "Product Name",
  vendor: "Vendor Name",
  license_key: "ABC123DEF456",
  expiry_date: "2026-12-31"
}
```

2. Validate the parser script before using it on real data.

   Every License Parser record includes a **Test Parser** link in the **Related Links** section. To validate your script, do the following on the License Parser record page:

   a. Fill in the **sample_input** field with representative file content.

   b. Fill in the **sample_output** field with the exact result you expect the parser to produce.

   c. Select **Test Parser**.

   The Test Parser executes your script against the sample input using the exact same code path used during a real scan. It reports whether the actual result matches the expected sample output, catching mistakes before the parser processes real license files.

3. Create a License File Configuration record to define the matching rule.

   With the validated parser in place, create a License File Configuration record that describes which file to look for. Enter *sn_acc_vis_content_license_file_config.list* in the navigation menu and specify the following:

   - **display_name** — A unique name identifying the rule
   - **file_name** — The file name or partial name to match (for example, "ProductLicense")
   - **file_name_condition** — How to match the file name: `exact_match` (default), `starts_with`, `ends_with`, or `contains`
   - **file_extension** — Optional file extension to narrow the match further (do not include a leading dot before the extension)
   - **platform** — The operating system(s) the rule applies to: `all` (default), `windows`, `linux`, or `mac`
   - **parser** — Reference to the parser created in the previous steps
   - **order** — Priority when multiple rules could match the same file (default 100; lower values win)
   - **active** — Indicates whether the rule is currently in use (default = true)

4. Enable license key discovery in the system.

   Ensure that the following components are active:

   - System property `sn_acc_vis_content.enable_license_key_discovery` — Set to **true**
   - Policy `File Based Discovery Policy - License Key` — **Active**
   - Check definition `File Based Discovery - License Key` — **Active**

   After enabling these components, the agent will begin scanning for files matching your License File Configuration rule during its regular file-scanning activity.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

Your new rule is active and the agent will automatically discover license keys from files matching your configuration. The discovered licenses appear in the appropriate tables when they are successfully matched in the software catalog.

Perform Zscaler remediation

If the Zscaler application installed on your Windows or macOS agent is not running efficiently, you can stop and start the app. This process is called *remediation*. Running remediation automatically creates an incident on the agent. You can also view Zscaler statuses on the Zscaler dashboard as a graph.

By default, the **Zscaler Monitoring Policy** monitoring check runs every 30 minutes to check the status of the Zscaler app running on your agent. The Zscaler monitoring policy contains the following checks:

- *zscaler-monitoring-check*: Checks the status of the Zscaler app running on your agent.
- *zscaler-remediation-check*: Invoked only when the Zscaler app is not running efficiently. This check shuts down and restarts the app, and automatically creates an incident.

  Only one incident is created every 24 hours for a Zscaler app that does not remediate properly. For example, if the remediation check runs multiple times on an agent, new remediation records are created and associated with the same incident. After 24 hours have passed, a second incident is created. When the Zscaler app restarts successfully, the incident closes.

  Zscaler remediation produces the following output:

  - Amount of data produced per remediation: 2KB
  - Maximum number of remediations that can run in a day: 48 (one every 30 minutes)
  - Maximum amount of data produced by a single agent, per day: 96KB
  - Maximum amount of data produced by a 20K agents, per day: 1.92GB

Check Zscaler monitoring

Check the monitoring status of your Zscaler app, to verify that it is running properly.

Before you begin

Role required: agent_client_collector_admin

About this task

If the monitoring app is down, the Zscaler tunnel status is either *51 (SERVICE-DOWN-ERROR)* or *56 (SYSTEM_SOCKETS_EXHAUSTED_ERROR)*. When that happens, a remediation check runs, which shuts down and then starts the Zscaler app.

By default, the Zscaler monitoring check (*zscaler-monitoring-check*) runs automatically every 30 minutes. If you are experiencing problems with your Zscaler app, you may want to run the check more often or immediately. For details, see Customize the Zscaler monitoring check.

Procedure

1. Navigate to **All** > **Remediate Zscaler VPN** > **Monitoring Status**.
   The **Application Monitoring Status** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application Monitoring Status page

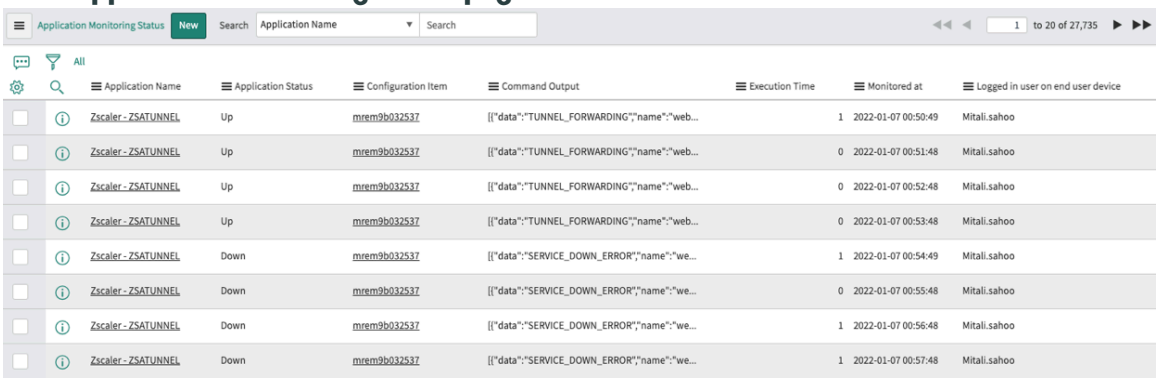

- If the Zscaler app is running properly, the **Application Status** column displays the **Up** value.
- If the Zscaler app is not running properly, the **Application Status** column displays the **Down** value.

2. If the **Application Status** column displays the **Down** value, the remediation check (*`zscaler-remediation-check`*) runs.
   - For details on customizing the amount of time after which the system verifes the Zscaler status, see Customize the Zscaler remediation wait time.
   - For details on monitoring Zscaler remediation checks, see Check Zscaler remediation.

Customize the Zscaler remediation wait time

Configure the amount of time (in seconds) by which the Zscaler remediation check verifies the Zscaler status. This amount of time indicates how often the system checks the Zscaler status after the remediation check runs.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **Agent Client Collector** > **Agents** > **Configuration** > **Check Definition**.
2. Select *`zscaler-remediation-check`*.
3. In the **Command** field, modify the value to: `zscaler_app_remediator.rb --wait_time=<time_in_sec>`, where `<time_in_sec>` is the number of seconds the system waits before checking the Zscaler status.
   By default, the system waits 30 seconds.
4. Select **Save**.
   - If the Zscaler app is running within the configured `wait_time`, the remediation result is **Success**.
   - If the Zscaler app is not running within the configured `wait_time`, the remediation result is **Failed**.

Customize the Zscaler monitoring check

Customize the frequency by which the Zscaler monitoring check runs. By default, the Zscaler monitoring check runs every 30 minutes. If your Zscaler app is not working efficiently, you can set the check to run more frequently.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: agent_client_collector_admin

About this task
For details on customizing the Zscaler monitoring check, see Perform Zscaler remediation.

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Configuration** > **Policies**.
2. Select the **Zscaler Monitoring Policy** entry.
3. Select the **Scheduling** tab.
4. In the **Interval (sec)** field, modify the value to the number of seconds for the monitoring check to wait in between executions.
   The default value is 1800.

   To run the monitoring check immediately, enter a value of **1**.

Result
Zscaler monitoring checks run according to the frequency specified in the **Interval (sec)** field.

Check Zscaler remediation

Verify that Zscaler remediation stops and starts the Zscaler app after the remediation monitoring check fails.

Before you begin
Verify that the Zscaler monitoring check returns an **Application Status** value of **Down**, indicating that the Zscaler app isn't working properly.

Role required: agent_client_collector_admin

Procedure
Navigate to **All** > **Remediate Zscaler VPN** > **Remediation Result**.

The **Application Remediation Status** page appears.

Application remediation status page

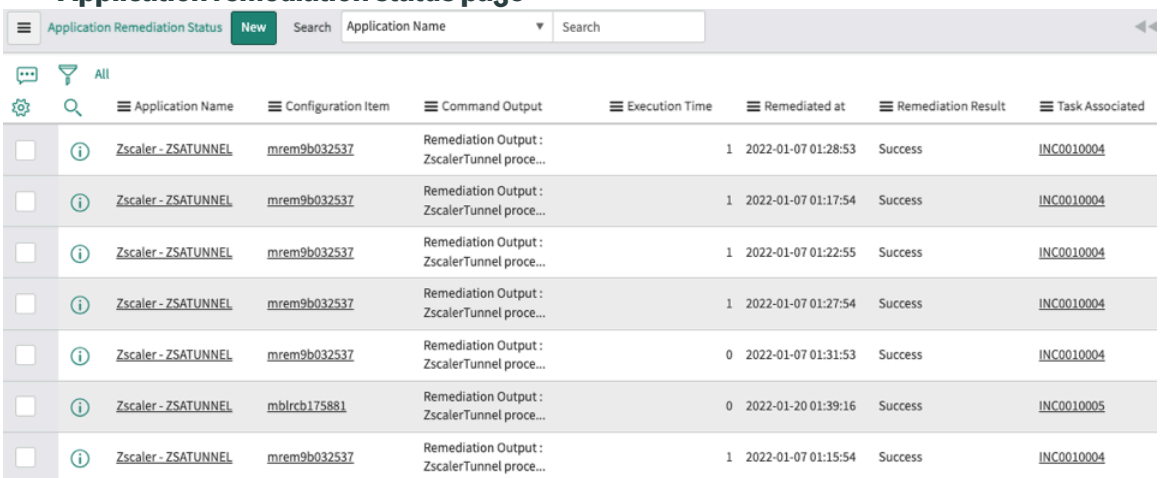

When remediation successfully shuts down and restarts Zscaler, the **Remediation Result** column has a value of **Success**. The **Task Associated** column displays a link to the incident created on remediation. When remediation is successful, the incident closes automatically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the Zscaler dashboard

Use the Zscaler dashboard to view Zscaler monitoring and remediation statuses in graph format.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Remediate Zscaler VPN** > **Dashboard**.
   The Zscaler Remediation Dasbhoard appears.

Zscaler Dashboard

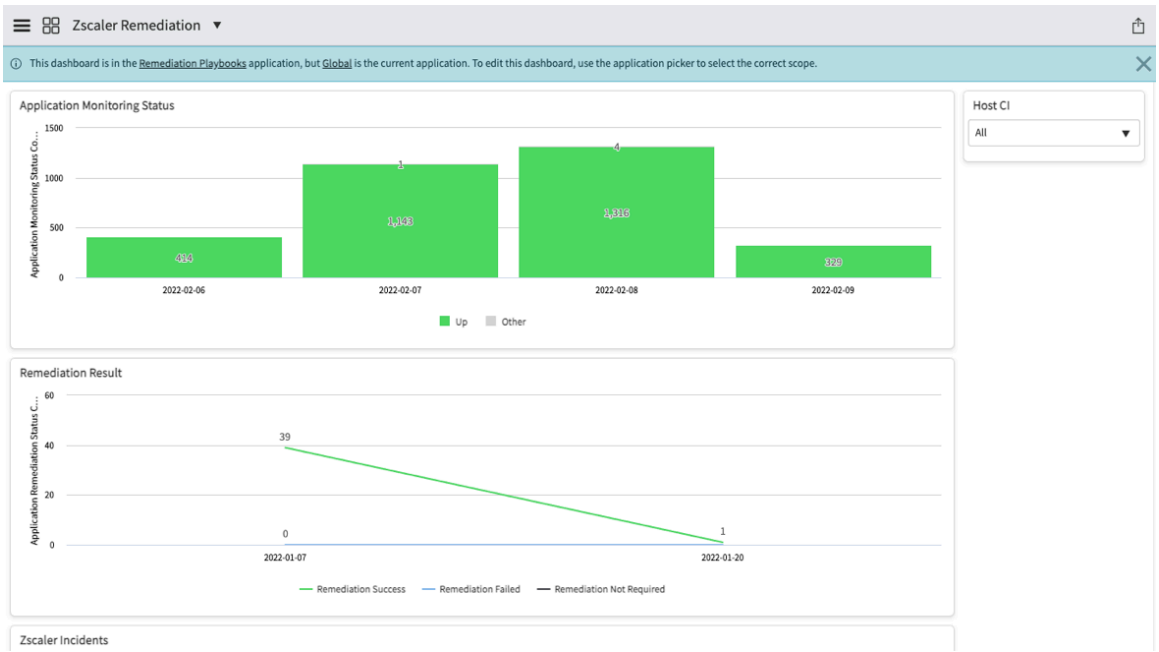

The dashboard displays graphs divided according to the following sections:

Zscaler Remediation Dashboard sections

Section	Description
Application Monitoring Status	Displays a bar graph of the number of Zscaler apps by date, and their statuses (Up, Down, Unknown, Other).
Remediation Result	Displays a line graph of the number of Zscaler remediations, and their statuses (Success, Failed, Not Required).
Zscaler Incidents	Displays a line graph of the number of Zscaler incidents created, and their statuses (Opened, Closed).

2. **Optional:** In the **Host CI** window on the upper right side of the page, select the CI whose data you want to view.
   The default value is All, indicating that the dashboard displays data from all of the CIs on the host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Choose and configure metrics to monitor

Metric Intelligence uses data sources that can monitor hundreds of metrics for all CIs. Choose which details are important for each data source type and CI, then activate or deactivate the respective monitor type to control data processing.

Before you begin

See Get started with Metric Intelligence for the high level steps of setting up Metric Intelligence.

Role required: evt_mgmt_admin

About this task

The SCOM data source, available by default, collects about 400 metric types from discovered devices such as '% Available Memory' and '% of Free Space'. Processing data from all metric types can overload the system and impact performance. It might be necessary to deactivate selected metric types for selected CIs to reduce that load.

The overall number of metrics that can be processed, depends on the number of MID Servers and their hardware configuration. For details about hardware requirements and metrics scalability, see MID Server distributed cluster for Metric Intelligence.

Note: The *Active* setting (true or false) for a metric in the Monitoring System Metric Types [sa_source_metric_type] table takes precedence over the setting for the corresponding metric in the Metric To CI Mappings [sa_metric_map] table. If a metric type in the Monitoring System Metric Types [sa_source_metric_type] table is disabled, all records related to the corresponding metrics are removed from the Metric To CI Mappings [sa_metric_map] table.

Procedure

1. Navigate to **All** > **Event Management** > **Metrics** > **Metric Types**.
   The **Monitoring System Metric Types** displays the metric types that are being collected per data source, per CI type.
2. Review the **Monitoring System Metric Types** list, and set the value in the **Active** column to **false** or **true**, as necessary.

Column	Description
Source	Data source monitoring the metric type.
Source Metric Type	Metric such as '% Free Space' that is being monitored for the CI type.
Anomaly Detection Action Level	The level of anomaly detection processing and analysis applied to CIs and metrics. The available options are: ○ **Metrics Only**: Only gather metrics from the data source without any further processing. ○ **Bounds**: Includes the Metrics Only processing level, and also creates a statistical model and shows bounds in Insights Explorer. ○ **Anomaly Scores**: Includes the Bounds processing level, and also calculates and displays anomaly scores.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
	○ **Anomaly Alerts**: Includes the Anomaly Scores processing level, and also creates anomaly alerts, as applicable. ○ **IT Alerts**: Includes the Anomaly Alerts processing level, and also creates IT alerts which are based on anomaly alerts, as applicable.
Monitored Object Type	CI type that is being monitored.
Active	Determines whether the metric data is processed by Event Management alert aggregation and RCA metrics.
Unit	Unit type that is associated with the source metric type, displayed in the Metric Intelligence reports.

Create a configuration settings rule

Configuration settings affect how metric data is processed. Configuration settings rules override the default metric processing behavior to determine the system actions when an anomaly is detected.

Before you begin

Role required: evt_mgmt_admin

About this task

A number of configuration settings determine the behavior of Metric Intelligence MID Servers. In the base system, these configuration settings are configured with default values, data types, and range of valid values. You cannot directly modify these configuration settings or add new ones. However, you can create a metric configuration rule with new configuration settings that override the default values on the MID Servers.

Then, manually apply these rules to all Metric Intelligence MID Servers in a single synchronization operation, or rely on an hourly system job to perform the synchronization.

To take effect, Metric Intelligence MID Servers must be synchronized with the updates to the configuration settings rules.

It is valid to have multiple rules for a setting that affect the same CIs, in which case:

- Rules in which a filter is defined take precedence over a global rule in which no filter has been defined.
- If multiple rules that affect the same set of CIs have the same priority, then only the latest rule to be defined is applied.
- If multiple rules with different priorities affect the same set of CIs, then rules with the highest priority are applied.

Procedure

1. Navigate to **All** > **Event Management** > **Anomaly Detection** > **Metric Config Rules**.
2. On the Metric Configuration Rules pane, click **New**, and fill out the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric Configuration Rules form

Field	Description
Name	Rule name.
Order	Rule priority within all other rules. Higher numbers represent higher priorities.
Filter by	Check box for displaying the **Rule** field, where you can specify conditions that CIs must meet for the rule to apply. For example, in the **choose field** list, select **Add Related Fields** and then add the filter **[class][is][Linux Server]**. If clear, the rule applies globally to all CIs in the Metric To CI Mapping [sa_metric_map] table.

3. Right-click the form title, and click **Save**.
4. In the Metric overridden configurations form section click **New**, fill out the form, and then click **Submit**.

Metric overridden configurations form

Field	Description
Name	Configuration setting for which to override its value. Click the icon to display the list of all configuration settings. Click the icon to display the **Metric Settings** dialog with details such as range of possible values. See the following tables (Configuration Settings and Internal Configuration Settings) for details about configuration settings.
Rule	Rule to which the created configuration setting applies.
Value	New value that overrides the default value for the specified configuration setting.

You can modify the following configuration settings in the **Name** field.

Note: The filter specified in the metric configuration rule does not apply to settings with a global scope.

Configuration Settings

Name and Description	Values	Default	Data Type	Scope
anomaly_detection_enabled Enable/disable anomaly detection.	N/A	true	boolean	CI/Metric

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name and Description	Values	Default	Data Type	Scope
Note: If anomaly_detection_action_level is set, then anomaly_detection_enabled is ignored.				
anomaly_detection_action_level Action level of anomaly analysis and processing. For more information, see Configure the action level for anomaly detection.	choices: ○ Metrics Only ○ Bounds ○ Anomaly Scores ○ Anomaly Alerts ○ IT Alerts	○ New records: Bounds ○ Upgraded records: Anomaly Alerts	choice	CI/Metric
buffer_anomaly_eviction_size Maximum number of anomalies at individual metric level that can be stored in internal buffer before sending them to instance for every CI/Metric pair.	60–1440	60	integer	Global
buffer_ci_score_eviction_size Maximum number of anomalies at CI level that can be stored in internal buffer before sending them to instance (Currently not being used)	60–1440	60	integer	Global
buffer_metric_eviction_size Maximum number of metrics that can be stored in internal buffer before sending them to instance for every CI/Metric pair.	60–1440	60	integer	Global
connection_login_timeout_secs Maximum time in seconds to log in to the local database on MID Server.	30–60	30	integer	Global
corrupt_data_count_threshold Minimum number of training points (15-minute averages)	10–100	30	integer	Global

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name and Description	Values	Default	Data Type	Scope
required to do any statistical analysis.				
deprioritize_early_batching_of_anomalous_ci Send anomalous CI information immediately or at regular interval.	N/A	false	boolean	Global
mad_model_min_days Number of days for which data must be available to consider a Median Absolute Deviation based model.	10-120	10	integer	CI/Metric
max_pool_connections_size Maximum number of connections for local database pool.	10–50	25	integer	Global
observation_time_min Expected minimum metric observation interval.	1–1440	1	integer	CI/Metric
robust_central_percentage Percentage of the residual data to compute the residual standard deviation, used for outlier detection. When set to 100 - uses the regular sample standard deviation.	50–100	90	double	Global
sparse_gap_fraction_threshold If more than this percentage of data is missing and no other class has been identified, classify as SPARSE. Do not attempt to fit a WEEKLY model.	0–100	50	double	Global
weekly_model_min_days Number of days for which data must be available in order to consider only a WEEKLY seasonality decomposition.	14-90	15	integer	CI/Metric
daily_model_min_days Number of days for which data must be available in	2-90	3	integer	CI/Metric

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name and Description	Values	Default	Data Type	Scope
order to consider only a DAILY seasonality decomposition.				
build_snpm_model Enable/disable building an SNPM data model.	N/A	true	boolean	CI/Metric
snpm_minimum_data_count Minimum number of data points required for building a stationary nonparametric model.	0 – 1e9	5000	integer	CI/Metric

The following configuration settings are for internal usage.

Internal Configuration Settings

Name and Description	Values	Default	Data Type	Scope
anomaly_memory_time_min Anomaly score calculator parameter: Memory time for abnormal situation.	1–600	45	double	CI/Metric
excess_z_score Anomaly score calculator parameter: Minimal anomalousness accumulated for outlier.	0–3	0.8	double	CI/Metric
linear_accumulator_threshold Decision Tree Threshold: ACCUMULATOR analysis	0.5–5	1	double	Global
low_freq_power_threshold Decision Tree Threshold: WEEKLY analysis	0–100	50	double	Global
low_variability_threshold Decision Tree Threshold: TRENDY analysis	0.000000000001–0.001	0.0001	double	Global
mid_freq_power_threshold Decision Tree Threshold: WEEKLY analysis	0–100	33	double	Global

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name and Description	Values	Default	Data Type	Scope
multinomial_count_threshold Decision Tree Threshold: MULTINOMIAL analysis	1–1000	40	integer	Global
non_zero_diff_threshold Decision Tree Threshold: NEAR_CONSTANT analysis	0–100	5	double	Global
normal_memory_time_min Anomaly score calculator parameter: Memory time for normal situation.	1–600	1	double	CI/Metric
normal_probability_ewma_timescale_min Anomaly score calculator parameter: Normal assessment time-scale.	1–600	15	double	CI/Metric
normal_probability_weight Anomaly score calculator parameter: Normal assessment adjustment factor.	0–1	1	double	CI/Metric
sigmoid_offset Anomaly score calculator parameter: Anomalousness to score conversion.	0–5	2.1	double	CI/Metric
sigmoid_weight Anomaly score calculator parameter: Anomalousness to score conversion.	0–5	1.2	double	CI/Metric
tiny_variability_threshold Decision Tree Threshold: NEAR_CONSTANT analysis	0–0.001	0.0000000001	double	Global
weekly_peak_hi_limit Decision Tree Threshold: WEEKLY analysis.	7–14	10	double	Global

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name and Description	Values	Default	Data Type	Scope
weekly_peak_lo_limit Decision Tree Threshold: Weekly analysis.	0.5–7	0.7	double	Global
weekly_vs_daily_log_likelihood_threshold By how much log likelihood of weekly needs to be larger than daily, to be the preferred statistical model.	100–1000	200	double	CI/Metric
daily_vs_noisy_log_likelihood_threshold By how much log likelihood of daily needs to be larger than noisy, to be the preferred statistical model.	20–1000	200	double	CI/Metric
weekly_vs_noisy_log_likelihood_treshold By how much log likelihood of weekly needs to be larger than noisy, to be the preferred statistical model.	100–1000	200	double	CI/Metric
trendy_vs_noisy_log_likelihood_threshold By how much log likelihood of trendy needs to be larger than noisy, to be the preferred statistical model.	10–1000	50	double	CI/Metric
seasonal_loess_width_in_hours Applied to the seasonal component of a weekly or daily model before making a forecast of future behavior. If set to 0, each data point in the seasonal model becomes independent of the rest of the data points.	6–24	12	double	CI/Metric
robustness Affects how outliers contribute to seasonal and trend calculations.	N/A	true	boolean	CI/Metric
snpm_min_value_threshold	-1e9 – 1e9	0	double	CI/Metric

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name and Description	Values	Default	Data Type	Scope
Minimum value of data required for building an SNPM model.				
snpm_max_observation_interval_in_sec Maximum expected observation interval required for building an SNPM model.	60 – 600000	120	integer	CI/Metric
min_std_jump_fraction Minimum ratio of locally calculated observation noise level to typical jump size that justifies recalculating a larger observation noise variance.	0.0 – 1.0	0.2	double	CI/Metric
dynamic_threshold_error_smoothing Whether to use exponentially weighted moving average to smooth the residuals in the dynamic threshold analysis.	N/A	true	boolean	CI/Metric
ewma_alpha The alpha value of the exponentially weighted moving average in dynamic threshold analysis.	1e-15 – 1.0	0.027397260274	double	CI/Metric
dynamic_threshold_beginning_smoothing_length Number of smoothed data points to set to the mean of double the smoothing length.	0 – 10000	250	integer	CI/Metric
dynamic_threshold_error_buffer_minutes Number of data points around each outlier to group together.	1 – 1000	30	integer	CI/Metric
dynamic_threshold_search_start Start value at which the optimal control factor is looked for.	0.5 – 20.0	3.0	double	CI/Metric
dynamic_threshold_search_interval Interval between search values of optimal control factor.	0.1 – 5.0	0.5	double	CI/Metric

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name and Description	Values	Default	Data Type	Scope
dynamic_threshold_search_count Number of values required for searching for optimal control factor.	1 – 50	19	integer	CI/Metric
dynamic_threshold_error_sequence_limit Maximum number of error groups for a particular control factor value when searching.	1 – 20	5	integer	CI/Metric
dynamic_threshold_minimum_data_count Minimum number of raw data points needed before attempting dynamic thresholding.	1 – 10000	5000	integer	CI/Metric
linear_seasonal_log_likelihood_threshold Threshold used in deciding whether to prefer a fitted model with linear seasonality over a model with a periodic component.	10-5000	1000	integer	CI/Metric

Related topics

Configure multiple MID Servers to work with Agent Client Collector Monitoring

Synchronize configuration settings rules

Metric Intelligence configuration settings rules contain user specified values that override default values that currently exist on Metric Intelligence MID Servers. To take effect, the Metric Intelligence MID Servers must be synchronized with the updated set of configuration settings rules.

Before you begin

Role required: evt_mgmt_admin

About this task

You can manually initiate a synchronization to immediately apply the updates on all Metric Intelligence MID Servers. Or, you can rely on a system job, that runs on an hourly recurring schedule, checking for updates to the metric configuration rules. If it is determined that there are updates, the system job sends these updates to the Metric Intelligence MID Servers.

Procedure

1. Navigate to **All** > **Event Management Rules** > **Anomaly Detection** > **Metric Config Rules**.
2. On the Metric Configuration Rules pane, click **Sync to MID**.

What to do next

Verify that all Metric Intelligence MID Servers have been synchronized with the updated set of configuration settings rules:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Navigate to **Output and Artifacts** > **ECC Queue**.
2. In the **Queues** pane, search the **Topics** column for `MetricConfigProbe`.
3. Check the **Agent** column and verify that all MID Servers are updated.

Understanding the Monitoring Technology Dashboard

The Monitoring Technology Dashboard enables you to monitor server resources for the platform you select. The dashboard enables you to identify the CIs and servers in your system with the highest resource consumption, and the most recent active alerts.

You can customize the appearance of individual widgets on the Monitoring Technology Dashboard. For example, you could change the data parameters displayed for the CPU Utilization (Linux) metric, like the date range. You can include the customized widgets when creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

The Monitoring Technology Dashboard shows the five CIs and servers with the highest metric readouts during the past six hours. It also shows the 50 most recent active alerts on your OS's servers.

To enable viewing the Monitoring Technology Dashboard, ensure that you have installed clotho version 23.1.0.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select the OS whose metrics you want to view.

Select the arrow next to the dashboard name to enabled choosing a different dashboard to view:

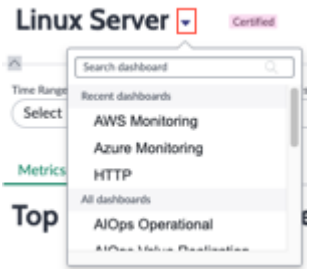

Monitoring Technology Dashboard for Linux

With the Linux Monitoring Technology Dashboard, you can monitor the health and performance of your Linux infrastructure. The dashboard enables you to identify the CIs and servers in your system with the highest resource consumption, and also the most recent active alerts.

Overview of the dashboard

The Monitoring Technology Dashboard shows the five CIs and servers with the highest metric readouts during the past six hours. The dashboard also shows the 50 most recent active alerts on your OS's servers.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **Linux Server**.

Using the Monitoring Technology Dashboard for Linux

Before using the Monitoring Technology Dashboard for Linux:

- Ensure that you're working in the Agent Client Collector Monitoring scope in your ServiceNow instance.
- Ensure that the agent_client_collector_admin and dashboard_admin roles are enabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Enable Metric Intelligence on the MID Server.
- Ensure that the **Linux OS Events** and **Linux OS Metrics** policies are enabled.

  Enabling these policies activates metric to CI mapping.
- Ensure that the clotho plugin version 1.19.0.10 or above is enabled.
- After upgrading from a previous version, follow the procedure described in the ACC-M Dashboard Filters are not working when upgraded from Tokyo/Utah to Vancouver [KB1438042] article in the NOW Support Knowledge Base.

Point to the relevant widget to view a specific point of data and all of its details in a pop-up window. Select the Alerts tab to view the 50 most recent active alerts for CIs.

You can filter the dashboard data according to the following filters:

Dashboard filters

Filter name	Description
Time range filter	The time range for which you want metric data to display on the dashboard.
Metrics by selected CI filter	The CIs for which you want metric data to display on the dashboard.
Metrics by selected OS filter	The OS for which you want metric data to display on the dashboard.

On the right side of widgets displaying a single metric, you can select the More options icon (⋮) and select **Change group by**. Select the category you want to view the metric by: Host name, Class, Operating System, or OS Version.

Linux server dashboard

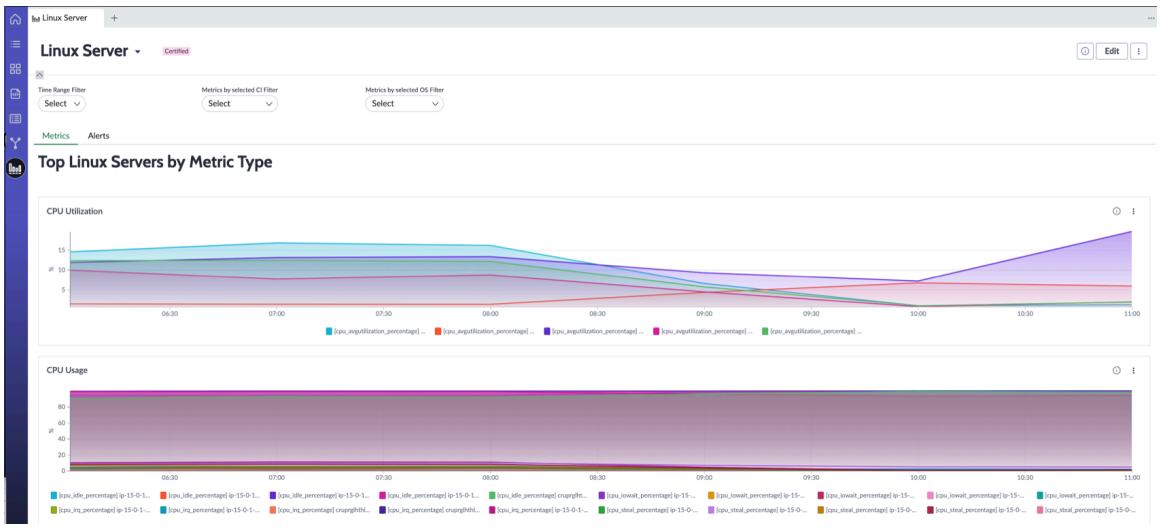

The Metrics subtab displays the top five performing CIs and servers over the past six hours, according to the following metrics.

- CPU Utilization
- System Uptime
- CPU Usage

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Load Average
- Memory
- Memory Distribution
- Processes
- Network Utilization
- Swap
- Disk Usage

You can customize the appearance of the individual widgets on the Linux Monitoring Technology Dashboard. For example, you could change the timeframe for which the alerts appear. You can also include customized widgets when you are creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

Monitoring Technology Dashboard for Windows

With the Windows Monitoring Technology Dashboard, you can monitor the health and performance of your Windows infrastructure. The dashboard enables you to identify the CIs and servers in your system with the highest resource consumption, and also the most recent active alerts.

Overview of the dashboard

The Monitoring Technology Dashboard for Windows shows the five CIs and servers with the highest metric readouts during the past six hours. The dashboard also shows the 50 most recent active alerts on your OS's servers.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **Windows Server**.

Using the Monitoring Technology Dashboard for Windows

Before using the Monitoring Technology Dashboard for Windows:

- Ensure that you're working in the Agent Client Collector Monitoring scope in your ServiceNow instance.
- Ensure that the agent_client_collector_admin and dashboard_admin roles are enabled.
- Enable Metric Intelligence on the MID Server.
- Ensure that the **Windows OS Events** and **Windows OS Metrics** policies are enabled:

  Enabling these policies activates metric to CI mapping.
- Ensure that the clotho plugin version 1.19.0.10 or above is enabled.
- After upgrading from a previous version, follow the procedure described in the ACC-M Dashboard Filters are not working when upgraded from Tokyo/Utah to Vancouver [KB1438042] article in the NOW Support Knowledge Base.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **Windows Server**.

Point to the relevant widget to view a specific point of data and all of its details in a pop-up window. Select the Alerts tab to view the 50 most recent active alerts for CIs.

You can filter the dashboard data according to the following filters:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Dashboard filters

Filter name	Description
Time range filter	The time range for which you want metric data to display on the dashboard.
Metrics by selected CI filter	The CIs for which you want metric data to display on the dashboard.
Metrics by selected OS filter	The OS for which you want metric data to display on the dashboard.

On the right side of widgets displaying a single metric, you can select the More options icon (⋮) and select **Change group by**. Select the category that you want to view the metric by: Host name, Class, Operating System, or OS Version.

Windows server dashboard

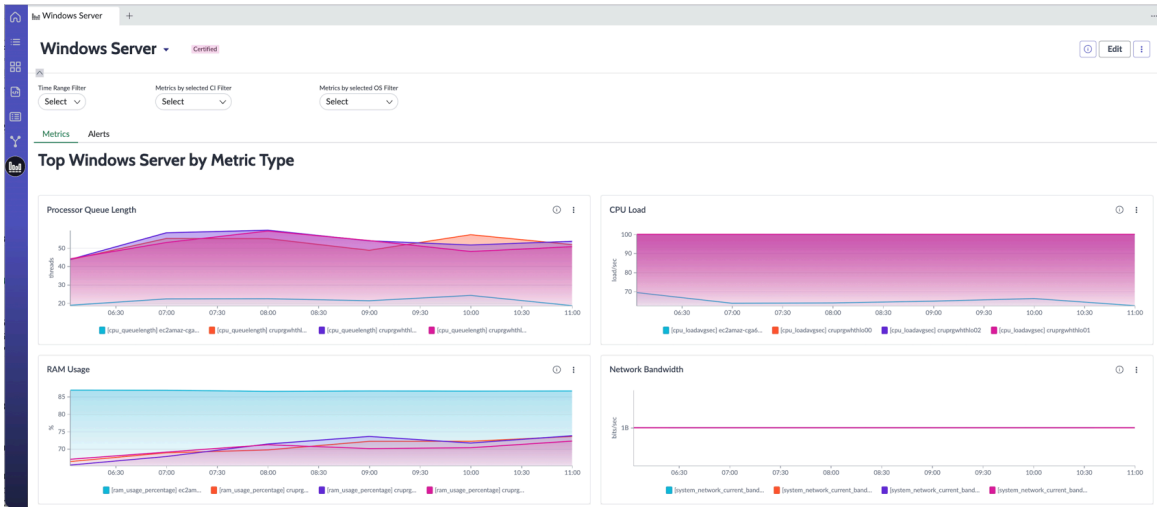

The Metrics subtab displays the top five performing CIs and servers over the past six hours, according to the following metrics.

- Processor Queue Length
- CPU Load
- System Uptime
- CPU Core
- RAM Usage
- Network Bandwidth
- Memory Distribution
- Network Flow
- Disk Usage
- Disk Read/Write
- Disk Distribution

You can customize the appearance of the individual widgets on the Windows Monitoring Technology Dashboard. For example, you could change the timeframe for which the alerts

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

appear. You can also include customized widgets when you are creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

Monitoring Technology Dashboard for HTTP

With the HTTP Monitoring Technology Dashboard you can monitor the response time of your `http` and `https` URLs. By using the dashboard, you can identify the URLs with the highest metric readouts in your system's infrastructure, and view the most recent active alerts on those URLs.

Overview of the dashboard

The HTTP Monitoring Technology Dashboard shows the five URLs with the highest metric readouts during the past six hours. The dashboard also shows the 50 most recent active alerts on those URLs.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **HTTP**.

Using the Monitoring Technology Dashboard for HTTP

Before using the Monitoring Technology Dashboard for HTTP:

- Ensure that you're working in the Agent Client Collector Monitoring scope in your ServiceNow instance.
- Ensure that the agent_client_collector_admin and dashboard_admin roles are enabled.
- Enable Metric Intelligence on the MID Server.
- Ensure that the **HTTP Entry Points Events** and **HTTP Entry Points Metrics** policies are enabled.

  Enabling these policies activates metric to CI mapping.
- Ensure that the clotho plugin version 1.19.0.10 or above is enabled.
- After upgrading from a previous version, follow the procedure described in the ACC-M Dashboard Filters are not working when upgraded from Tokyo/Utah to Vancouver [KB1438042] article in the NOW Support Knowledge Base.

The Azure Monitoring Technology Dashboard contains the tabs that are listed in the following table.

HTTP Monitoring Technology Dashboard tabs

Tab	Description
Metrics	Displays the top 5 performing metrics that are associated with the URLs in your system.
Alerts	Displays the 50 most recent alerts that are associated with the URLs in your system. Alerts help you to remediate your system by providing information on problematic behavior.

On the right side of widgets displaying a single metric, you can select the More options icon (⋮) and select **Change group by**. Select the category you want to view the metric by: Host name, Class, Operating System, or OS Version.

Point to the relevant widget to view a specific datapoint and all of its details in a pop-up window.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the **Time range filter** to filter the time range for which you want metric data to display on the dashboard.

You can customize the appearance of the individual widgets on the HTTP Monitoring Technology Dashboard. For example, you could change the data parameters that display for the Response Time metric, like the time range that the response time is displayed. You can also include the customized widgets when you're creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

Monitoring Technology Dashboard for Azure

With the Azure Monitoring Technology Dashboard, you can monitor the health and performance of your Azure infrastructure. By using the dashboard, you can identify the configuration items (CIs) and servers with the highest metric readouts, and view the most recent active alerts on those CIs.

Overview of the dashboard

The Azure Monitoring Technology Dashboard shows the five CIs and servers with the highest metric readouts during the past six hours. The dashboard also shows the 50 most recent active alerts in your Azure environment. Based on the dashboard metrics, you might want to resize your CIs and servers to improve your system performance and efficiency.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **Azure Monitoring**.

Using the Monitoring Technology Dashboard for Azure

Before using the Monitoring Technology Dashboard for Azure:

- Create an Azure credential. For details on creating Azure credentials, see Create Azure cloud credentials.
- Create a Discovery schedule, using the credential created in the previous step. For details on creating a Discovery schedule, see Discovery for Microsoft Azure.
- Ensure that you're working in the Agent Client Collector Monitoring scope in your ServiceNow instance.
- Ensure that the agent_client_collector_admin and dashboard_admin roles are enabled.
- Enable Metric Intelligence on the MID Server.
- Ensure that the **Azure VM Metrics**, **Azure Storage Account Metrics**, **Azure Load Balancer Metrics**, **Azure Gateway Metrics**, and **Azure Redis Cache Metrics** policies are enabled
- Ensure that the clotho plugin version 1.19.0.10 or above is enabled.

The Azure Monitoring Technology Dashboard contains the tabs that are listed in the following table.

Azure Monitoring Technology Dashboard tabs

Tab	Description
Virtual Machines	Displays the metric data from the virtual machines in the Azure environment.
Storage Accounts	Displays the metric data from the storage accounts in the Azure environment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Monitoring Technology Dashboard tabs (continued)

Tab	Description
Load Balancer	Displays the metric data from the load balancers in the Azure environment.
Application Gateway	Displays the metric data from the application gateway in the Azure environment.
Redis	Displays the metric data from the Redis cache in the Azure environment.

Point to the relevant widget to view a specific datapoint and all of its details in a pop-up window.

You can filter the dashboard data according to the following filters:

Dashboard filters

Filter name	Description
Time range filter	The time range for which you want metric data to display on the dashboard.
Metrics by selected CI filter	The CIs for which you want metric data to display on the dashboard.

You can customize the appearance of the individual widgets on the Azure Monitoring Technology Dashboard. For example, you could change the time that the CPU usage on virtual machines is displayed. You can also include the customized widgets when you're creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

Monitoring Technology Dashboard for AWS

With the AWS Monitoring Technology Dashboard, you can monitor the health and performance of your AWS infrastructure. Identify the configuration items (CIs) and servers with the highest metric readouts. View the most recent active alerts on those CIs and servers.

Overview of the dashboard

The AWS Monitoring Technology Dashboard shows the five CIs and servers with the highest metric readouts during the past six hours. The dashboard also shows the 50 most recent active alerts in your AWS environment.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **AWS Monitoring**.

Using the Monitoring Technology Dashboard for AWS

Before using the Monitoring Technology Dashboard for AWS:

- Create an AWS credential. For details on creating AWS credentials, see Setup AWS service accounts.
- Create a Discovery schedule, using the credential created in the previous step. For details on creating a Discovery schedule, see Access setup for AWS service accounts.
- Ensure that you're working in the Agent Client Collector Monitoring scope in your ServiceNow instance.
- Ensure that the agent_client_collector_admin and dashboard_admin roles are enabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Enable Metric Intelligence on the MID Server.
- Ensure that the **AWS Metrics**, **AWS EC2 Metrics**, and **AWS DynamoDB Metrics** policies are enabled.

  Enabling these policies activates metric to CI mapping.
- Ensure that the clotho plugin version 1.19.0.10 or above is enabled.
- After upgrading from a previous version, follow the procedure described in the ACC-M Dashboard Filters are not working when upgraded from Tokyo/Utah to Vancouver [KB1438042] article in the NOW Support Knowledge Base.

The AWS Monitoring Technology Dashboard contains the tabs that are listed in the following table.

AWS Monitoring Technology Dashboard tabs

Tab	Description
EC2	Displays the metric data from the EC2 (virtual machines) in the AWS environment.
EC2 Alerts	Displays the active EC2 (virtual machine) alerts in the AWS environment.
EBS	Displays the metric data from the EBS (storage volume) in the AWS environment.
S3	Displays the metric data from the S3 (cloud object storage) in the AWS environment.
RDS	Displays the metric data from the RDS (cloud database) in the AWS environment.
ELB	Displays the metric data from the ELB (cloud load balancing) in the AWS environment.
Dynamo DB	Displays the metric data from the Dynamo DB (cloud database) in the AWS environment.
ECS Cluster	Displays the metric data from the ECS cluster in the AWS environment.
ECS Service	Displays the metric data from the ECS service in the AWS environment.

On the right side of widgets displaying a single metric, you can select the More options icon (⋮) and select **Change group by**. Select the category you want to view the metric by: Host name, Class, Operating System, or OS Version.

Point to the relevant widget to view a specific datapoint and all of its details in a pop-up window. Select the **Alerts** tab to view the 50 most recent active alerts in your environment. Alerts help you to remediate your system by providing the information on problematic behavior.

Use the **Time range filter** to filter the time range for which you want metric data to display on the dashboard.

Use the **Metrics by selected CI filter** to select the CIs for which you want to view metrics.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can customize the appearance of the individual widgets on the AWS Monitoring Technology Dashboard. For example, you could change the time that the EC2 metrics appear. You can also include the customized widgets when you are creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

Monitoring Technology Dashboard for GCP

With the GCP Monitoring Technology Dashboard, you can monitor the health and performance of your GCP infrastructure. Identify the configuration items (CIs) and servers with the highest metric readouts. View the most recent active alerts on those CIs and servers.

Overview of the dashboard

The GCP Monitoring Technology Dashboard shows the five CIs and servers with the highest metric readouts during the past six hours. The dashboard also shows the 50 most recent active alerts in your GCP environment.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **GCP Monitoring**.

Using the Monitoring Technology Dashboard for GCP

Before using the Monitoring Technology Dashboard for GCP:

- Create a GCP credential and use it to create a Discovery schedule. For details on creating GCP credentials, see Discovery for Google Cloud Platform.
- Create a Discovery schedule, using the credential created in the previous step. For details on creating a Discovery schedule, see Discovery for Google Cloud Platform.
- Ensure that you're working in the Agent Client Collector Monitoring scope in your ServiceNow instance.
- Ensure that the agent_client_collector_admin and dashboard_admin roles are enabled.
- Enable Metric Intelligence on the MID Server.
- Ensure that the **GCP GCE Metrics** policy is enabled.

  Enabling this policy activates metric to CI mapping.
- Ensure that the clotho plugin version 1.19.0.10 or above is enabled.
- After upgrading from a previous version, follow the procedure described in the ACC-M Dashboard Filters are not working when upgraded from Tokyo/Utah to Vancouver [KB1438042] article in the NOW Support Knowledge Base.

The GCP Monitoring Technology Dashboard contains the tabs listed in the following table.

GCP Monitoring Dashboard tabs

Tab	Description
GCP	Displays graphs of the top 14 performing OS resources, such as CPU, Memory, Disk Operations and Network.
Alerts	Displays the 50 most recent alerts that are associated with the URLs in your system. Alerts help you to remediate your system by providing information on problematic behavior.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

On the right side of widgets displaying a single metric, you can select the More options icon (⋮) and select **Change group by**. Select the category you want to view the metric by: Host name, Class, Operating System, or OS Version.

Point to the relevant widget to view a specific datapoint and all of its details in a pop-up window.

Use the **Time range filter** to filter the time range for which you want metric data to display on the dashboard.

Use the **Metrics by selected CI filter** to select the CIs for which you want to view metrics.

You can customize the appearance of the individual widgets on the GCP Monitoring Technology Dashboard. For example, you could change the timeframe for which the alerts appear. You can also include customized widgets when you are creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

Monitoring Technology Dashboard for VMware vSphere

With the VMware vSphere Monitoring Technology Dashboard, you can monitor the health and performance of your VMware vSphere infrastructure. You can identify the CIs and servers with the highest metric readouts. You can also view the most recent active alerts on those CIs and servers.

Overview of the dashboard

The VMware vSphere Monitoring Technology Dashboard shows the five CIs and servers with the highest metric readouts during the past six hours. The dashboard also shows the 50 most recent active alerts in your VMware vSphere environment.

To view the dashboard, navigate to **All** > **AIOps Dashboards** > **Monitoring Technology Dashboards** and select **VMWare vSphere Monitoring**.

Using the Monitoring Technology Dashboard for VMware vSphere

Before using the Monitoring Technology Dashboard for VMware vSphere:

- Create a VMware vSphere credential and use it to create a Discovery schedule. For details on creating VMware vSphere credentials, see Discovery for VMware.
- Create a Discovery schedule, using the credential created in the previous step. For details on creating a Discovery schedule, see Discovery for VMware virtualization.
- Ensure that you're working in the Agent Client Collector Monitoring scope in your ServiceNow instance.
- Ensure that the agent_client_collector_admin and dashboard_admin roles are enabled.
- Enable Metric Intelligence on the MID Server.
- Ensure that the **VSphere Datacenter Metrics**, **VSphere Datastore Metrics**, **VSphere ESX Server Metrics** and **VSphere VirtualMachine Metrics** policies are enabled.

  Enabling these policies activates metric to CI mapping.
- Ensure that the clotho plugin version 1.1.0 or above is enabled.
- After upgrading from a previous version, follow the procedure described in the ACC-M Dashboard Filters are not working when upgraded from Tokyo/Utah to Vancouver [KB1438042] article in the NOW Support Knowledge Base.

The VMware vSphere Monitoring Technology Dashboard contains the tabs listed in the following table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

VMware vSphere Monitoring Dashboard tabs

Tab	Description
VirtualMachine	Displays graphs of the top VMware vSphere VMs by metric type.
ESX Server	Displays graphs of the top VMware vSphere ESX servers by metric type.
Datacenter	Displays graphs of the top VMware vSphere datacenters by metric type.
Datastore	Displays graphs of the top VMware vSphere datastores by metric type.
Alerts	Displays the 50 most recent alerts that are associated with the URLs in your system. Alerts help you to remediate your system by providing information on problematic behavior.

On the right side of widgets displaying a single metric, you can select the More options icon (⋮) and select **Change group by**. Select the category you want to view the metric by: Host name, Class, Operating System, or OS Version.

Point to the relevant widget to view a specific datapoint and all of its details in a pop-up window.

You can filter the dashboard data according to the following filters:

Dashboard filters

Time range filter	The time range for which you want metric data to display on the dashboard.
Metrics by selected CI filter	The CIs for which you want metric data to display on the dashboard.

You can customize the appearance of the individual widgets on the VMware vSphere Monitoring Technology Dashboard. For example, you could change the timeframe for which the alerts appear. You can also include customized widgets when you are creating a custom dashboard. For more information, see Customize the Monitoring Technology Dashboard.

Customize the Monitoring Technology Dashboard

Customize the Monitoring Technology Dashboard widgets in the Data Visualizations library so your dashboard displays information exactly as you want to see it.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to the OS dashboard whose widgets you want to customize.
   - To access the Windows Server Dashboard, navigate to **All** > **AIOps** > **Monitoring Technology Dashboard** > **Windows Server**.
   - To access the Linux Server Dashboard, navigate to **All** > **AIOps** > **Monitoring Technology Dashboard** > **Linux Server**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Select **Analytics Center** > **Data Visualizations**.
3. View all available widgets by selecting the **All** tab.
4. Select the widget that you want to change.
5. Select the ( ) icon, and select **Duplicate**.
6. In the **New name** field in the dialog box, enter a name for your duplicated widget, and select **Duplicate**.
7. In the widget's **Configuration** tab on the right, expand the relevant options and customize the information for the widget to display.

Configuration tab

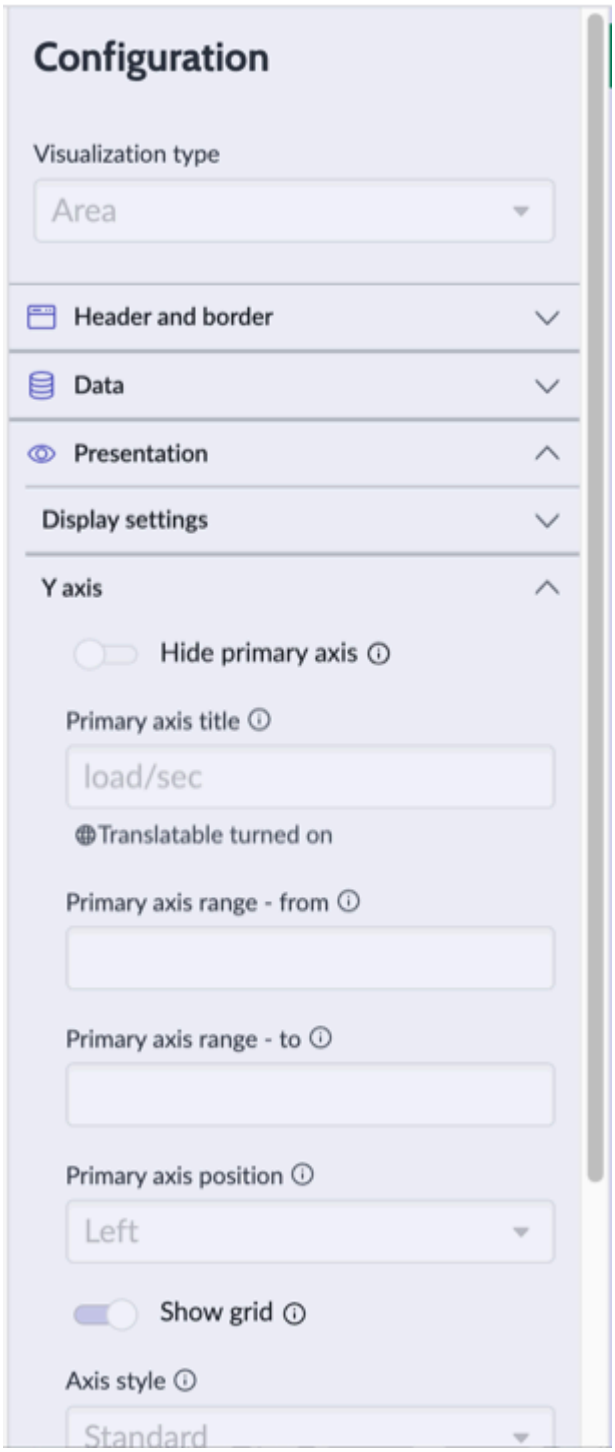

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For details on working with data visualizations, see Data visualizations in the Analytics Center Data visualizations in the Analytics Center.

8. Select **Save and close**.

Result

Your customized widgets are saved in the Data Visualizations library, and can be used to configure a customized dashboard (see Create a dashboard with the In-line editor).

Setting exclusion lists for IPs and NICs

Agent Client Collector for Visibility Content (ACC-VC) version 1.3.0 supports exclusion list for IPs and Network Interface Controllers (NICs) with a flexible mechanism for filtering out values for IPs and or NICs when creating or updating the host CI and related items.

The property [*sn_acc_vis_content.network_adapter_exclusion_list*] contains the list of regular expressions for the names and IP addresses that are excluded in the Network Adapter and IP Address tables. The value of this property is comma-separated regular expressions. Make sure that there is no comma in the regex and nic and ip_addr should be in different lines.

Sample format of the value should be:

- nic = nameRegex1, nameRegex2, nameRegex3
- ip_addr = IPRegex1, IPRegex2

Identify software editions on Windows devices

Determine which edition of software is in use on Windows devices in your environment, to maintain an accurate software inventory. Software products commonly support multiple editions, making it difficult to identify which edition is in use.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. In the navigation pane, enter `sn_acc_vis_content_software_registry_edition_config.list`.
   The **Software Edition Configurations** page appears.
2. Select **New** to set edition configuration for a Windows device.
   The **Software Edition Configuration - New record** page appears. For details on the fields displayed on the page, see Software edition configuration fields.
3. Select **Submit**.
   The configured product appears on the **Software Edition Configurations** page, and is automatically added to the `EditionRegistryConfig.json` file after the daily scheduled job runs.
4. In the navigation pane, enter `cmdb_sam_sw_install.list`.
   The **Software Installations** page opens.
5. Find the relevant software entry, and confirm the edition in the **Edition override** column.
   The edition appears only after the **Software Installed** or **SAM Discovery** policy runs. By default, these policies runs once daily.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generate a Pattern allowlist

Generate an allowlist for a selection of patterns, to configure the patterns permitted to run on an agent.

Before you begin

Ensure that the following plugins are installed on your instance:

- Discovery (com.snc.discovery)
- Pattern Designer Enhancements (com.sn_itom_pde)
- **Pattern Allowlist Generator** available as part of Agent Client Collector for Visibility Content 1.5.0 (202512)

Role required: discovery_admin or agent_client_collector_admin

Procedure

1. Navigate to **All** > **Pattern Designer** > **Patterns allowlist for ACC**.
   The **Pattern Allowlist Generator** page appears.
2. Select the type of pattern selection in the **Pattern Selection** section.
   - **All patterns**: The allowlist generates for all patterns on the agent.
   - **Specific patterns**: The allowlist generates on the agent only for the patterns you select.

     After selecting this option, choose the patterns from which you want the allowlist to be generated.
3. Select **Generate Allowlist**.
4. Select **Copy** to copy the generated allowlist.
5. Append the copied allowlist into the existing `check-allow-list.json` allowlist on the agent.

   The patterns specified in the allowlist are permitted to run on the agent.

   The allowlist's default location is:
   - Windows: `C:\ProgramData\ServiceNow\Agent Client Collector\check-allow-list.json`
   - Linux: `/etc/servicenow/agent-client-collector/check-allow-list.json`

   Note: To generate allowlist for commands nested within EVAL() in a pattern step, ensure the Pattern Designer Enhancements plugin is v3.9.1 or above.

Define temporary variables for a pattern allowlist

Define temporary variables by assigning values such as executable paths, config file paths, and so forth. Defining temporary variables ensures that runtime commands are successful.

About this task

If a Linux command executable relies on an undefined temporary variable, the command fails at runtime.

For example, for the command `$pathToCmd "/log/acc.log"`, if a value is not defined for `$pathToCmd`, the command will fail.

Before you begin

Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Pattern Designer** > **Pattern Allowlist Temporary Variables**.
2. Select **New** to create a new variable.
3. Create a record and define the temporary variable.

   For example:

Pattern Allowlist Temporary Variables

Field	Value
Value	bin/sh/testexec
Pattern	.NET Application
Operating system	Linux
Temporary variable	exePath

4. Select **Submit**.

Domain usage discovery and categorization

Agent Client Collector for Visibility (ACC-VC) domain usage discovery classifies the domains that users visit against a set of admin-defined domain signatures. It enables you to see which websites and services, such as AI tools, are being used across the endpoints in your environment.

Domain usage discovery overview

Domain usage discovery matches the domains that users visit against a set of domain-type signatures, and records a match for each combination of endpoint, domain, and user. You can create categories of any type, such as AI Tools, Productivity, or Security, and map them to the domain URLs in your signatures. Use the resulting data to see which categories of website and services are accessed from your browsers and devices.

Note: This capability reuses the existing web usage collection pipeline. No new agent-side data collection module is required.

Prerequisites

Domain usage discovery depends on the Agent Client Collector for Visibility (ACC-VC) VISC Get URL metrics policy. This policy ships turned off, so you must activate it before AI domain usage discovery can classify any domains. For more information, see Agent Client Collector for Visibility Content default checks and policies.

Domain signatures

Domain usage discovery relies on how you categorize domains. The system doesn't categorize domains automatically. To identify usage of a website or application, create a domain signature for it and assign the category that you want. Visited domains that match the signature are then reported under that category. For example, to identify usage of an AI application, create a domain signature for it and set its category to AI Tools.

How domain usage discovery works

Domain usage discovery runs inline as part of existing web usage processing:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- After the system resolves a visited domain, it checks the domain against the active domain-type signatures.
- Domain matching is case-insensitive.
- If the domain matches a signature, the system records the endpoint, domain, user, matching signature, and category.
- Existing web usage data collection and storage are unaffected. This classification happens in addition to, not instead of, standard web usage tracking.
- The system records each endpoint, domain, and user combination only one time. It doesn't update or duplicate existing records.
- The system loads the active domain signatures fresh each time it processes web usage data, so newly activated signatures are detected without requiring a restart.

Categorize domain usage

Group the domains that users visit in your environment by business relevance, such as AI tools, productivity applications, or security services.

Before you begin

Verify the Agent Client Collector for Visibility (ACC-VC) VISC Get URL metrics policy is active to enable domain usage discovery to classify any domains. For more information, see Agent Client Collector for Visibility Content default checks and policies.

Role required: discovery_admin

About this task

Domain usage discovery relies on how you categorize domains. To identify usage of a website or an application, create a domain signature and assign it to a category.

Procedure

1. Create a category.

   a. Navigate to **All** > **System Definition** > **Tables**.

   b. Open the **Entity Categories** (`sn_acc_vis_content_entity_category`) list and select **New**.

   c. On the form, fill in the fields.

Entity categories

Field	Description
Name	The display name of the category. This field is required. For example: AI Tools, Productivity, or Microsoft.
Description	Option to add a brief description of the category to help identify it.
Active	Option to activate the category.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	Note: You can also activate a saved category later from the **Entity Categories** list.

d. Select **Save**.

e. **Optional:** View the saved categories in the **Entity Categories** list.
You can activate saved categories from this list by setting Active=true for the relevant categories.

2. Create a signature.

a. Navigate to **All** > **System Definition** > **Tables**.

b. Open the **URL Domain Signature** (`sn_acc_vis_content_url_domain_signature`) table.

c. Under related links, select **Show List**.

d. Select **New**.

e. On the form, fill in the fields.

Domain Signature form

Field	Description
Domain type	The display name of the signature. The options are: External or Internal.
Category	The category to which the system assigns matching domains. For example, AI Tools. This field is required.
URL domain	The pattern to match against the URL domain. This field is required
Description	Option to add a brief description of the signature to help identify it.
Active	Option to activate the signature.

f. Select **Submit**.

Saving the signature starts a job that evaluates the visited domains against the signature and categorizes any matches.

Browser extension discovery and categorization

Agent Client Collector for Visibility (ACC-VC) browser extension discovery inventories browser extensions installed on the endpoints in your environment. When an admin-defined signature matches, ACC-VC classifies each extension into a category, such as AI Tools, Productivity, or Security.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How browser extension discovery works

Users can install browser extensions without administrator involvement. Therefore, traditional software discovery often misses them. Browser extension discovery gives you an inventory of the extensions on your endpoints, their assigned category, and the devices where each extension is installed.

Browser extension discovery supports these browsers:

- Chrome
- Edge
- Brave
- Opera
- Firefox
- Safari

Note: On Mac machines, populating the browser extension status and extended metadata columns requires FDA.

For an extension to appear in the inventory, create a signature and specify the extension ID or name to match. Assign a category, such as AI Tools, to classify the matching extensions in the inventory. ACC-VC matches each discovered extension against the active signatures, and writes the result to the Browser Extension Catalogs table.

A signature takes effect on a device the next time that the 24-hour discovery check runs on that device. The daily configuration refresh picks up signatures that are new or newly activated. If an extension matches more than one signature, ACC-VC applies only the first match, in this fixed order:

1. Extension ID signatures
2. Exact-name signatures
3. Other name-condition signatures, such as contains, starts with, and ends with.

Browser extensions are categorized or re-categorized automatically, without creating a new signature, in the following cases:

- An inactive signature is reactivated.
- New browser extensions are discovered that match an existing active signature.

Browser extension metadata fetching

Assess shadow IT and shadow AI risk in your organization by using browser extension metadata fetching in Agent Client Collector for Visibility (ACC-VC). This capability is an optional layer on top of browser extension discovery that collects each extension's metadata.

Browser extension metadata fetching overview

ACC-VC browser extension discovery identifies extensions installed across the endpoints in your environment. Browser extension metadata fetching goes a step further: it reads each extension's manifest and preference data to collect declared and runtime-granted permissions, host access, install source, and install and update timestamps. The system collects this metadata for extensions on Windows, Mac, and Linux endpoints. This feature provides security and IT teams the detail needed to assess shadow IT and shadow AI risk, rather than just the extension name and version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: On Mac machines, populating the extended metadata column requires FDA.

Data collected

The type of data browser extension metadata fetching collects depends on the browser family, as shown in the table.

Browser extension metadata by browser family

Field group	Chrome, Edge, Brave, and Opera	Firefox
Permissions	Runtime-granted permissions, falling back to permissions declared in the extension manifest	Permissions declared by the extension
Host access	Granted explicit hosts, falling back to manifest host permissions	Declared origins
Install source	Web store or side-loaded, classified from browser-reported signals	Web store or side-loaded, classified from browser-reported signals
Install and update time	Extension install and update timestamps reported by the browser	Extension install and update timestamps reported by the browser

Performance impact

Turning on browser extension metadata fetching increases the amount of data the agent reads and sends for each extension. Turn it on deliberately instead of leaving it on by default.

Note: If you turn off browser extension metadata fetching after previously turning it on, existing metadata records aren't deleted, but they stop being refreshed.

Categorize discovered browser extensions

Group discovered browser extensions in your environment by business relevance.

Before you begin

Role required: discovery_admin

About this task

Browser extension discovery finds the extensions on your endpoints and maps the category to which each extension belongs, along with the devices where it is installed. For an extension to appear in the inventory, create a signature and assign it a category to classify the matching extensions.

Important: If the system property sn_acc_vis_content.full_browser_extension_discovery is set to True, all browser extensions will be discovered, but those that don't have a matching user-defined signature will not have a category assigned to them. If this system property is set to False, only browser extensions that are matched with signatures will be discovered.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Create a category.

   a. Navigate to **All** > **System Definition** > **Tables**.

   b. Open the **Entity Categories** (sn_acc_vis_content_entity_category) table.

   c. Under related links, select **Show List**.

   d. Select **New**.

   e. On the form, fill in the fields.

   Entity Categories form

Field	Description
Name	The display name of the category. This field is required. For example: Adobe Acrobat, AI Tools, or Microsoft.
Description	Brief description of the category.
Active	When selected, activates the category. **Note:** You can also activate a saved category later from the **Entity Categories** table.

   f. Select **Submit**.

   g. **Optional:** View the saved categories in the **Entity Categories** table.

   Note: You can activate saved categories from this table by setting Active=true for the relevant categories.

2. Create a signature.
   The signature is a rule that matches extensions to a category based on the extension ID or name.

   a. Navigate to **All** > **System Definition** > **Tables**.

   b. Open the **Browser Extension Signature** (sn_acc_vis_content_browser_extension_signature) table.

   c. Under related links, select **Show List**.

   d. Select **New**.

   e. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Browser Extension Signature form

Field	Description
Name	The display name of the signature. This field is required. For example: ChatGPT.
Extension ID	The extension ID that matches the extension ID that is unique across vendors.
Extension name condition	Matching condition for the extension name. Options are: Exact match, Contains, Starts with, and Ends with.
Extension name	The pattern to match against the extension name. For example: ChatGPT.
Version condition	Matching condition for the version. Options are: Any version, Exact match, Starts with, Ends with, Contains, Less than, Greater than, Not empty, and Empty.
Version	Pattern to match against the version. **Note:** If you don't specify a version, the signature matches any version.
Category	Category assigned to matching browser extensions.
Description	Brief description of the signature. For example: Extensions related to ChatGPT.
Active	When selected, activates the signature.

f. Select **Submit**.

Submitting the signature triggers evaluation of browser extensions against it and categorization of any matches.

3. **Optional:** Fetch metadata of the discovered browser extensions.

   a. Navigate to **All** > **System Properties** > **All Properties**.

   b. Set the system property `sn_acc_vis_content.collect_extended_browser_extension_metadata` to True to collect the metadata.

   c. Optional: Review the fetched metadata.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

i. Navigate to **All** > **System Definition** > **Tables**.

ii. Open the **Browser Extension Metadata** (sn_acc_vis_content_browser_extension_metadata) table.

iii. Review the metadata.

Result

The categorized browser extensions in your environment appear in the **Browser Extension Catalogs** (sn_acc_vis_content_browser_extension_catalog) table. Each entry in the table shows the browser extension, its assigned category, extension ID, version, and the signature that matched. Filter by category to see all the browser extensions in the group.

The **Device Browser Extensions** (sn_acc_vis_content_device_browser_extension) table shows the mapping of the discovered browser extensions for each browser and device, and their profile and status.

Software packages categorization

Agent Client Collector for Visibility (ACC-VC) classifies discovered software packages in your environment into categories. This categorization removes the need to tag software records manually and provides an accurate software inventory.

Software categorization overview

Software packages the system has discovered in your environment are stored in two different tables:

- If Software Asset Management (SAM) is installed on your instance, they are stored in the **Software Discovery Models** (cmdb_sam_sw_discovery_model) table.
- If SAM is not installed, they are stored in the **CMDB CI Package** (cmdb_ci_spkg) table and the **Software package** column is populated.

ACC-VC software categorization compares each discovered software package with active, admin-defined categorization signatures. A signature is a rule that matches software to a category based on the software name or publisher. When you activate the signature, ACC-VC scans discovered software and finds matches. The system then creates entries in the Software Categorization Catalog table. Each catalog entry identifies the software package, the category assigned to it, and the signature it matched. You can filter the catalog by category to review all software in a group, such as AI tools, security-related software, or communication tools.

Software is categorized or re-categorized automatically, without creating a new signature, in the following cases:

- An inactive signature is reactivated.
- New software is discovered that matches an existing active signature.

Categorize discovered software

Group discovered installed software packages in your environment by business relevance.

Before you begin

In the System Properties, verify that the software categorization system property is set to true: `sn_acc_vis_content.sw_categorization.enabled=true`.

Role required: discovery_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Software categorization classifies discovered software packages into categories, such as AI tools, security software, communication software, or a publisher-based category. Use this procedure to define a category and a signature that matches discovered software to that category.

Procedure

1. Create a category.

   a. Navigate to **All** > **System Definition** > **Tables**.

   b. Open the **Entity Categories** (sn_acc_vis_content_entity_category) table.

   c. Under related links, select **Show List**.

   d. Select **New**.

   e. On the form, fill in the fields.

   Entity Categories form

Field	Description
Name	The display name of the category. This field is required. For example: Adobe Acrobat, AI Tools, or Microsoft.
Description	Brief description of the category.
Active	When selected, activates the category. **Note:** You can also activate a saved category later from the **Entity Categories** table.

   f. Select **Submit**.

   g. **Optional:** View the saved categories in the **Entity Categories** table.

   Note: You can activate saved categories from this table by setting Active=true for the relevant categories.

2. Create a signature.
   The signature is a rule that matches software to a category based on the software name or publisher.

   a. Navigate to **All** > **System Definition** > **Tables**.

   b. Open the **Installed Software Signature** (sn_acc_vis_content_sw_install_signature) table.

   c. Under related links, select **Show List**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Select **New**.

e. On the form, fill in the fields.

Define the match condition using the software name, publisher, or both. Either the software name or the publisher is required.

Software install signature form

Field	Description
Name	The display name of the signature. This field is required. For example: ChatGPT.
Category	Category assigned to matching software. For example: AI Tools.
Software name condition	Matching condition for the software name. Options are: Exact match, Contains, Starts with, and Ends with.
Software name	Pattern to match against the software name. Defining either Software name or Publisher is required. For example: ChatGPT.
Publisher condition	Matching condition for the publisher. Options are: Exact match, Contains, Starts with, and Ends with.
Publisher	Pattern to match against the publisher. Defining either Publisher or Software name is required.
Version condition	Matching condition for the version. Options are: Any version, Exact match, Starts with, Ends with, Contains, Less than, Greater than, Not empty, and Empty.
Version	Pattern to match against the version. **Note:** If you don't specify a version, the signature matches any version.
Description	Brief description of the signature.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Active	When selected, activates the signature.

f. Select **Submit**.

Submitting the signature triggers evaluation of the installed software against it and categorization of any matches.

Result
The categorized software appears in the **Installed Software Catalogs** (sn_acc_vis_content_sw_install_catalog) table. Each entry in the table shows the software package, its assigned category, and the signature that matched. For each record, either the **Software package** or the **SAM discovery model** column is populated, depending on whether SAM is installed on the instance. Filter by category to see all the software in the group.

Discover portable software installed by package managers

Agent Client Collector for Visibility Content (ACC-VC) can discover software on Windows, Linux, and macOS endpoints that is not discoverable by traditional ACC-VC checks and policies. ACC-VC uses third-party package managers to track essential tools and development packages for software asset management (SAM) and IT asset management (ITAM).

Portable software discovery - overview

The Agent Client Collector software discovery policies discover software that has a dedicated executable file. The Software Discovery via Running Processes check detects software installed by package managers (containing .js files or .py modules) that is not detected by other ACC-VC checks, For details on ACC-VC policies and checks, see Agent Client Collector for Visibility Content default checks and policies.

With portable software package discovery, the ACC-VC agent identifies and reports portable software installed through a package manager. This includes globally-installed npm packages, which enrich your software asset inventory with Node.js-based tools and utilities.

What gets discovered

For each package found on a host, the agent captures the following information:

- Package name: The display name configured in the Portable Software Package Discovery (*sn_acc_vis_content_process_based_sw_config*) table, or the actual npm package name if no display name is specified.
- Version: The installed version of the package. Populated only if the number is available; otherwise, this field is blank.
- Vendor/Publisher: The vendor name from the configuration. Populated only if the vendor/publisher is available; otherwise, this field is blank.
- Installation directory: The file system path where the package is installed.

Configuring npm package discovery

To begin discovering packages, you must configure package discovery. See Configure portable software discovery in Agent Client Collector for Visibility Content for step-by-step instructions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure portable software discovery in Agent Client Collector for Visibility Content

Configure portable software the Agent Client Collector for Visibility Content (ACC-VC) agent discovers by updating the config table. This ensures that only portable software that was configured is searched for and discovered.

Before you begin

Ensure that you have access to the ACC-VC scoped app in your ServiceNow instance.

Role required: discovery_admin

About this task

Software discovery is governed by the Portable Software Package Discovery (*sn_acc_vis_content_process_based_sw_config*) table. When you add or remove entries from this table, the changes are delivered to ACC agents on their next check cycle, and an agent rebuild is not needed.

Procedure

1. Enter *sn_acc_vis_content_process_based_sw_config.list* into the navigation panel.
   The **Process Based Software Discovery Configurations** page opens.
2. Configure the following fields for each software name you want the system to discover.

Field	Description
Software name	The name to match. This is either the process's own name, or the script/package name that shows up in the interpreter's command line (for example, openclaw). Not case-sensitive.
Description	Enter descriptive text indicating why you're searching for the software.
Active	Select the check box for the software to be pushed to the agent.

3. Select **Update**.
4. Activate the **Software DIscovery via Running Process** policy.
   The discovered software displays in the Software installation (*cmdb_sam_sw_install*) table with the following fields discovered:

Field	Description
Name	The name of the software you configured in the Process Based Software Config (*sn_acc_vis_content_process_based_sw_config* table.
Version	The software version. If not discoverable during initial detection, the agent tries again every 5 minutes.
Vendor/Publisher	The vendor/publisher of the software. If not discoverable during initial detection, the agent tries again every 5 minutes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Install location	The directory that the discovered software resides in.
Status	When software is successfully discovered, value is *install ok installed*.

Related topics

Discover portable software installed by package managers

ACC-VC NPVDI endpoints

Agent Client Collector for Visibility Content (ACC-VC) adjusts how its discovery and software checks run on Windows non-persistent virtual desktop infrastructure (NPVDI) endpoints, so that data collected during a short-lived session reaches your instance before the session ends.

How ACC-VC identifies an NPVDI endpoint

An NPVDI endpoint is a virtual desktop that is provisioned from a golden image for each user session and is reset or removed when the user logs off. Because the endpoint might not exist at the next scheduled discovery interval, ACC-VC changes when and how certain checks run on these endpoints so that session data isn't lost.

NPVDI support applies only to non-persistent Windows agents. Discovery of persistent Windows, macOS, and Linux devices remains unchanged. ACC-VC doesn't detect NPVDI endpoints automatically. An endpoint that isn't explicitly configured as non-persistent, including an endpoint with missing or unclear configuration, is treated as a persistent endpoint. The agent is marked as non-persistent when you install it on the golden image. For more information, see Enable a non-persistent virtual desktop infrastructure agent.

ACC-VC policies that support NPVDI

The following VDI policies list determines which policies run on NPVDI agents. ACC-VC has already added the following policies to this list, so you don't have to add them yourself. The list has no effect on persistent endpoints.

- *Enhanced Discovery*
- *SAM Discovery*
- *SAM background*
- *SAM background (Non OsqueryD)*

For more information about each policy, see Agent Client Collector for Visibility Content default checks and policies.

The *Software installed* policy isn't registered as a VDI policy because it applies only to servers. To add other published policies to an NPVDI agent, see Prepare agent deployment on a non-persistent virtual desktop infrastructure machine.

Data collection when an NPVDI session ends

The SAM Discovery policy runs on its regular schedule and also when the endpoint shuts down. Each run sends the software installation data it collects to your instance. The shutdown run makes sure that data collected since the previous run reaches your instance before the session and its data are removed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Enhanced Discovery, SAM background, and SAM background (Non OsqueryD) policies run on their regular schedule only.

Reduced Enhanced Discovery on NPVDI endpoints

Some *Enhanced Discovery* inventory modules don't have to be collected again in every short-lived session. On NPVDI endpoints, ACC-VC can skip the following modules:

- File systems
- Network adapters
- Storage devices
- Local users
- Intel EMA
- Memory modules

By default, all six modules are skipped. You control the skipped modules and whether the installed software scan results are cached within the same session with a system property. For details, see Configure ACC-VC data collection for NPVDI endpoints.

The following inventory areas run on every check and can't be skipped:

- *Basic Inventory*
- *Enhanced Inventory*
- *TCP Connections*
- *Running Processes*

Installed software caching on NPVDI endpoints

The software installations and usage metrics check discovers the software installed on an endpoint and collects usage metrics for that software, such as last used date and total usage duration.

On NPVDI endpoints, you can enable caching of the installed software inventory. When caching is enabled, the full installed software scan runs one time for each golden image. Later check runs reuse the cached results instead of rescanning. Caching is turned off by default.

Caching applies only to the installed software inventory. Software usage metrics are collected on every check run, because usage is specific to each session and can't be reused.

Behavior when NPVDI configuration is unavailable

If an NPVDI endpoint can't read its NPVDI configuration, for example, because the acc.yml file is missing or malformed, ACC-VC doesn't skip any Enhanced Discovery modules and treats caching as turned off. Discovery then runs the same way it does on a persistent endpoint.

Configure ACC-VC data collection for NPVDI endpoints

Choose which *Enhanced Discovery* modules Agent Client Collector for Visibility Content (ACC-VC) skips on Windows non-persistent virtual desktop infrastructure (NPVDI) endpoints. Optionally, choose whether installed software scan results are cached within the same session.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Verify that your instance is prepared for NPVDI agents and that the agent is installed on the golden image as a non-persistent agent. For details, see Prepare agent deployment on a non-persistent virtual desktop infrastructure machine and Enable a non-persistent virtual desktop infrastructure agent.
- ACC-VC version 2.1.1 or later must be installed.
- Role required: admin

About this task

Two system properties control how ACC-VC collects data on NPVDI endpoints. These settings apply only to Windows endpoints that are configured as non-persistent. On all other endpoints, the full *Enhanced Discovery* and installed software scan run every time, regardless of these settings. For details on each property, see ACC-VC NPVDI system properties.

Procedure

1. Navigate to **All** and enter `sys_properties.list` in the filter.
2. Open the *sn_acc_vis_content.np_vdi.enhanced_discovery_skip_modules* property.
3. Set the **Value** field to a comma-separated list of the *Enhanced Discovery* modules to skip.

   Valid values are:

   - `file_systems`
   - `network_adapters`
   - `storage_devices`
   - `local_users`
   - `intel_ema`
   - `memory_modules`

   By default, all six modules are listed. To collect a module on NPVDI endpoints, remove it from the list. Any other value is ignored.

4. Select **Update**.
5. **Optional:** Cache the installed software inventory on NPVDI endpoints.

   a. Open the *sn_acc_vis_content.np_vdi.cache_installed_software_enabled* property.

   b. Set the **Value** field to `true`.

   c. Select **Update**.

   (Optional) When caching is enabled, the full installed software scan runs one time for each golden image, and later check runs reuse the cached results. Software usage metrics are still collected on every check run.

Result

When you update either property, the instance regenerates the NPVDI configuration file (`NPVDIConfig.json`) that ACC-VC uses on NPVDI agents. A daily scheduled job also refreshes this file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Deploying Agent Client Collector on both servers and endpoints

When deploying the Agent Client Collector, there are components which are shared between servers and endpoints.

Checks and policies

A check is a combination of a command and its configuration. The check is executed on the Agent Client Collector's devices to gather data from those devices.

Checks

Checks are provided with the base system, and their commands execute scripts which provide monitoring data for your operating systems and applications. A check's default name indicates what is being monitored and measured, the entity, and the monitoring data. For example, a check named *os.linux.check-system-cpu* checks the CPU data on a Linux system. The identified command in the check runs on the monitored device, providing an output and status. Each individual check is referred to as a **check definition**. After being associated with policies, check definitions are then referred to as **check instances**.

You can customize check instances to meet your needs. For example, customize the running interval or the parameters specific to the policy, such as the login credentials to access a MySQL database. Customization of a check instance takes effect only on the check instance associated with the policy, which does not affect the original check definition or already created check instances in other policies.

The following check types are provided with the Event Management base system:

- **Event**: The check's result is transformed into an Event Management event.
- **Metric**: The values from the check result are transformed to metrics.

For details on the Agent Client Collector Framework default checks, see Agent Client Collector Framework default checks.

For details on the Agent Client Collector Monitoring default checks and policies, see Agent Client Collector Monitoring default checks and policies.

For details on the Agent Client Collector for Visibility Content default checks and policies, see Agent Client Collector for Visibility Content default checks and policies.

If checks are not running on the agent's devices, your agent may be in CPU protection mode. CPU protection mode is activated when a device's CPU consumption is too high. To view a list of all disabled checks related to the agent, navigate to the agent record on the Agent Client Collectors page (**All** > **Agent Client Collector** > **Agents**), select an agent and select the **Disabled Checks** tab at the bottom of the page.

In the `acc.yml` configuration file, you can modify the CPU protection mode thresholds which determine when a check is disabled. To re-enable a check, select the check and select **Re-enable disabled check** in the **Actions on selected rows...** drop-down. For details on the CPU protection mode thresholds, see Agent Client Collector CPU protection thresholds. For details on manually turning off data collection, see Pause Agent Client Collector data collection.

Configure the cpu_protection_behavior property to determine whether all checks enter CPU protection mode, or only the check that crosses the configured cpu_percentage_limit value with the highest CPU usage within the monitoring interval. For details, see Agent Client Collector CPU protection thresholds.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A check is marked as **Stale** if it is no longer associated with a policy. To remove a stale check, renable it by selecting the box next to the check and selecting **Re-enable disabled checks** from the **Actions on selected rows...** cell.

In the base system, the event's exit status indicates its severity, as follows:

- 0 = OK
- 1 = WARNING
- 2 = CRITICAL

You can add additional severities (such as MAJOR and MINOR) by running a customized script. The following exit statues indicate these severities:

- 13 = MAJOR
- 14 = MINOR

Privileges for running check commands

If the servicenow base system user does not have privileges to run specific check commands, do the following for the relevant operating systems:

- In a Linux system: Enable the servicenow user to run the command with `sudo` permissions. Ensure that the following sudo configuration requirements are met:
  - Disable tty and password requirements
  - Preserve all environment variables
  - Support dynamic PATH for executing commands

  Example configuration in the `/etc/sudoers` file:

```
Cmnd_Alias ACC_F = /usr/sbin/dmidecode -s
 baseboard-serial-number, /usr/sbin/dmidecode -s
 chassis-serial-number, /usr/sbin/dmidecode -s
 system-serial-number, /usr/sbin/dmidecode -s
 system-uuid, /usr/sbin/ss -tanp
servicenow ALL=(root)
 SETENV: /
var/cache/servicenow/agent-client-collector/osquery/bin/osque
ryi *, ACC_F
Defaults:servicenow !requiretty
Defaults exempt_group += servicenow
```

  Note: Command paths may vary. Consult the sudoers manual for special considerations.
  - The `SETENV:` string enables the servicenow user to preserve environment variables.
  - The `!requiretty` string disables tty.
  - Adding the servicenow user to the `exempt_group` bypasses password requirements and enables dynamic PATH for executing sudo commands.

  Ensure that you configure the *must_sudo* check parameter with a value of **true** in the check command parameters section of the check definition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- In a macOS system: Ensure that the user running the agent service belongs to a user group with privileges to query all tcp connections on the host.
- In a Windows system: Using Windows User Management, add the servicenow user to the groups with the relevant privileges, enabling the user to execute the required commands.

Executing checks that missed their execution time

A check doesn't execute if the laptop or agent is down during its scheduled execution time. When this happens, the system behavior depends on whether the check has executed previously.

- Check that never executed: The check runs at the agent reconnect time + X (where X = 0-60 minutes), and this is the permanent run time of the check.

  For example: Policy X was synced to an agent. Check A in the policy never ran previously and was scheduled to run at 2am. However, the laptop was asleep at that time. The laptop wakes up at 9:09am. After the agent is restored, it notices that check A never ran. Check A is scheduled to run between 9:09am-10:09am and ends up running at 9:54am. The check A interval is every 24 hours, so check A now is scheduled to run daily at 9:54am.

- Check that missed an existing execution time: The previously run check runs at the agent reconnect time + X (where X = 0-60 minutes), but in the future, the check reverts to its originally scheduled run time.

  For example: Check A runs once a day and is scheduled to run daily at 9:54am. The laptop was asleep or the agent was not connected from 9am until 1pm. At 1:17pm, the laptop is awake and the agent starts and connects. The agent noticed that check A did not run at 9:54am and schedules it to run between 1:17pm-2:17pm. The check ends up running at 1:37pm. However, the next time check A runs is at 9:54am, its originally scheduled time.

  This scenario takes effect only for checks running at an interval of every 60 minutes or less frequently. Checks that run more often than every 60 minutes at their scheduled interval starting from whenever the agent is up.

Policies

A **policy** is a combination of the CIs being monitored by the Agent Client Collector and the check definitions that run on those CIs.

To enable a single policy to support multiple credentials, assign a credential alias to the policy. For example, if you have MySQL servers for both Linux and Windows with different credentials, you must create separate policies for each credential type. However, if you use a credential alias, you can assign a single policy to the alias. The agent then matches the relevant credential to the application being monitored. For details on credential aliases, see Create a Connection and Credential alias.

Alerts generated by policies with an event-type check are automatically closed when the CI monitoring stops due to any of the following:

- De-activating the policy
- De-activating the check that caused the alert
- Deleting the check from the policy
- Deleting the policy
- Modifying the policy filter that determines the monitored CIs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector configuration data files

Configuration data files store dynamic instance data, such as virtual machine details, that check definitions use during execution. This ensures that checks are executed with up-to-date and accurate information about the instance being monitored.

Configuration data files structure

- Attachment based: A configuration data file contains a single attachment with instance data, such as Azure VM details.
- Synchronization: When an attachment is added or deleted, the data file is synchronized across all MID Servers configured with the Agent Client Collector. If a file or its attachment is deleted, it is also removed from the MID Server.
- Access point: The configuration data files can be accessed on the MID Server at `static/cache/config-files`.
- Size: Configuration data files cannot be larger than 10MB.

Check definitions used by configuration data files

- Check definition: Individual check entries are referred to as check definitions. For instance, *`os.linux.check-system-cpu`* checks the CPU on Linux systems.
- Accessing check definitions: Check definitions can be accessed from **All** > **Agent Client Collector** > **Configuration** > **Check Definitions**.

File association

- Download process: When a check executes, the agent downloads the configuration data file associated with that check.
- Domain-Specific Synchronization: Files are synchronized only with agents within the same domain, ensuring secure and localized data handling.

Working with configuration data files

Configuration data files serve as a repository for dynamic instance data that is used by check definitions when they are executed. These files ensure that checks are run with accurate, real-time information about the systems or instances being monitored. Configuration data files work as follows:

- Store instance-specific data: Configuration data files store dynamic instance data, such as virtual machine details, server configurations, and other system-specific information. This data is used by check definitions during execution to ensure the latest information is available.
- Create check definitions: A check definition is created with an associated configuration data file. For example, a check definition monitoring system CPU data may include configuration data detailing the system's environment, ensuring checks use accurate and relevant instance data.
- Synchronize with MID Servers: When a configuration data file is added or updated, it synchronizes across all MID Servers communicating with the Agent Client Collector. If the file or attachment is deleted, it is also removed from the MID Servers, ensuring consistency and up-to-date data availability.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Execute with updated data: During execution, the Agent Client Collector agent downloads the relevant configuration data file from the MID Server. The agent uses this data to execute the check, ensuring access to the most current and accurate instance information.
- Store and access data: Configuration data files are stored on the MID Server in the `static/cache/config-files` directory. This storage location facilitates easy access and management, ensuring the necessary data for checks is always available when needed.

Agent Client Collector API

Use the Agent Client Collector (ACC) API to create a flow that executes an `osquery` command on agents and processes the results. By leveraging the ACC API, you can automate the querying of agent data and streamline the processing of results, making it easier to monitor and manage system performance.

The system supports osquery 5.7.0, providing compatibility with a range of queries for efficient data collection.

The ACC API enables seamless integration and automation for executing queries and managing agents. Here's how it works:

1. Execute osquery commands: You can use the ACC API to send osquery commands to agents. These commands collect a wide range of system and performance data from the target machines.
2. Process results: After the osquery command is executed on the agents, the ACC API processes the results and returns the necessary data. This data can be used for monitoring, troubleshooting, or reporting purposes.
3. Support for osquery 5.7.0: The ACC API is compatible with osquery 5.7.0, which ensures it supports the latest features and improvements, enabling you to leverage advanced query capabilities for enhanced data collection.
4. Enable API features: To use the ACC API, enable the agent and policy APIs. This allows you to configure and manage various aspects of your agents, including their communication, data collection, and query execution.

For details on agent APIs, see Enable agent APIs.

For details on enabling policy APIs, see Enable policy APIs.

Agent Client Collector health instance scan suite

The Agent Client Collector (ACC) health instance scan suite consists of checks that detect anomalies and other issues that might occur on your instance. These checks ensure the overall health and performance of the ACC, proactively identifying potential problems before they impact system operations.

For details on running the ACC health instance scan, see Schedule the Agent Client Collector (ACC) health instance scan.

For details on the ACC health instance scan suite checks, see Agent Client Collector health instance scan checks.

For details on importing the `global.ACCInstanceScanUtil` script include, see the Script Include ACCInstanceScanUtil [KB1630132] article in the Now Support Knowledge Base.

Schedule the Agent Client Collector (ACC) health instance scan

Run the **Run (ACC) health instance scan** as a scheduled job to monitor the overall health of the instance receiving data from the Agent Client Collector.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: agent_client_collector_admin

About this task
This procedure describes how to schedule an ACC health instance scan. Alternatively, you can run the ACC health instance scan on demand, as follows:

1. Navigate to **All** > **Instance Scan** > **Suites**.
2. Select **ACC Health Instance Scan**.
3. Select **Execute Suite Scan**.

Procedure
1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Select the **Run ACC Health Instance Scan** entry.
3. Select the **Active** check box.
4. Set the frequency by which the job is to repeat in the **Repeat Interval** field.
   The default is 7 days.
5. Select **Execute Now** to run the **Run ACC Health Instance Scan** job.

Result
The results of the job are compiled in the Scan Results (scan_result) table. Errors are compiled in the Agents issue (sn_agent_acc_error_msg) table.

For details on the ACC health instance scan suite checks, see Agent Client Collector health instance scan checks.

Agent Client Collector certificates

Agent Client Collector certificates verify the authenticity of your agents, servers, and users. Using Agent Client Collector certificates ensures the safety of your environment.

Discovering certificates

To discover certificates on all of your host's applications with open ports, navigate to **All** > **System Properties** > **All Properties** and set the *sn_acc_vis_content.tls_ssl_scan_all_open_ports* property to **true**.

If you want to maintain certificates only for applications running on specific ports, set this property to **false**. The property's default value is **true**.

Manually refresh Agent Client Collector certificates

Refresh Agent Client Collector self-signed certificates manually to validate Agent Client Collector plugins, instead of waiting for the scheduled synchronization. For example, you can use this feature when the agent can't validate a plugin and you don't want to wait for the scheduled synchronization.

Before you begin
- Ensure that the self-signed certificate is placed in a `.zip` file on the MID Server in the following location: `agent\static\cert`
- Ensure that the certificate authority is part of a truststore, as the certificate chain validates before the customer plugin signature.
- Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **MID Servers**.
2. Select a MID Server.
3. In the Related Links section, select **Trigger ACC certificate sync**.
4. Restart the MID Server.

Result

When an agent fails to perform a certificate check, the agent synchronizes certificates available on the MID Server by pulling self-signed certificates from the install server. Certificates are downloaded by the Agent Client Collector if the agent does not validate its plugins.

Enable OpenSSL secure signing for plugins

Create a self-signed certificate for an Agent Client Collector plugin. The following procedure gives an example of how to create an x509 certificate using OpenSSL. For other certificate types, consult OpenSSL documentation.

Before you begin

- Ensure that the *verify-plugin-signature* property is set to **True** in the agent's `acc.yml` file to verify plugin signing.
- Ensure **OpenSSL** is installed on your system.

Role required: agent_client_collector_admin

About this task

Enabling an OpenSSL secure signing mechanism for plugins works with an Agent Client Collector installed on a Linux system.

Procedure

1. Create a plugin file with a `tar.gz` extension.
   For more information, see Create and edit Agent Client Collector plugins.
2. Generate your own secure self-certificate for the plugin file.

   a. Create an x509 certificate.

   ```
   openssl req -nodes -x509 -sha256 -newkey rsa:2048 -keyout
    "sign.key" -out "sign.crt" -days 365 -subj

    "/C=<CountryName>/ST=<StateOrProvinceName>/L=<Locality>/O=<O
   rganization>/OU=<OrganizationalUnit>/CN=sign"
   ```

   b. Sign the plugin file.

   ```
   openssl dgst -sha256 -sign "sign.key" -out sign.txt.sha256
    <plugin-name>.tar.gz
   ```

   Alternatively, you can sign plugins using a certificate authority. When doing so, assign the certificate .pem format and place it in the agent's `cert` directory.

   c. Verify that the signature is configured correctly.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
openssl dgst -sha256 -verify  <(openssl x509 -in
 "sign.crt" -pubkey -noout) -signature sign.txt.sha256
 <plugin-name>.tar.gz
```

If the file is valid, the output is `Verified OK`.

d. Encode the signature's certificate with base64 coding.

```
base64 sign.txt.sha256 > sign.txt.sha256_encode64.sig
```

A `sign.txt.sha256_encode64.sig` file is created.

3. Run the following commands to create a new directory and insert the `tar.gz` and `sign.txt.sha256_encode64.sig` files.
   a. `mkdir signed-plugin`
   b. `mv <plugin-name>.tar.gz signed-plugin`
   c. `mv sign.txt.sha256_encode64.sig signed-plugin`
   d. `cd signed-plugin`
4. Create another `tar.gz` file by running the same commands as you ran for the first `tar.gz` file.
   a. `tar -C . -zcvf ../<plugin-name>.tar.gz *`
   b. `cd..`

   Note: Save the new file as `../<plugin-name>.tar.gz` to avoid naming collisions with the original `<plugin-name>.tar.gz` file that exists in the current directory.
5. Upload the new `tar.gz` plugin file to the instance.
6. Set the plugin file as *`active=true`*.
7. Place the `sign.crt` file into one of the following locations:
   - The Agent Client Collector's `cert` directory, located in the `config` folder.
   - On the MID Server, compress the file into a `.zip` file and place it in the `agent/static/cert/customer` folder.
8. In the `acc.yml` file, set `verify-plugin-signature` to **true**.

What to do next

Add the self-signed `sign.crt` file to the agent host trust store, as described in Add a self-signed certificate to the OS truststore.

Add a self-signed certificate to the OS truststore

Add a self-signed certificate to the truststore of your operating system (OS). By adding a certificate to the truststore, you can verify that the certificate is authentic and that your connections are secure.

About this task

The procedure for adding a certificate to a truststore depends on the host's OS.

Adding a self-signed certificate to the truststore instructs the operating system to trust the certificate when it's establishing secure connections. This way, the applications can

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

communicate securely with services that use the self-signed certificate. If you add a certificate that isn't trustworthy, you might compromise your system's security.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. On a Windows OS:

   a. Press the Windows button + R on your keyboard.
   The **Run** field opens.

   b. Enter the string `certmgr.msc` and press **Enter**.
   The Certificate Manager opens.

   c. Expand the Trusted Root Certification Authorities directory.

   d. Select and hold (or right-click) **Certificates**.

   e. Select **All tasks** > **Import**.
   The Certificate Import Wizard opens.

   f. Follow the steps in the Certificate Import Wizard.
   For details, see Import a self-signed certificate.

   Alternatively, you can add a Windows certificate using PowerShell by running the following command as an administrator:

```
Import-Certificate -FilePath <path to your certificate file>
 -CertStoreLocation Cert:\LocalMachine\Root
```

2. On a macOS:

   a. Import the certificate into a keychain by running the following command in a terminal window:

```
bash

code:
sudo security add-trusted-cert -d -r trustRoot -k
 "/Library/Keychains/System.keychain" <path to your
 certificate file>
```

   b. Enter the administrator password when prompted.

   c. Open the Keychain Access application to verify that the certificate has been added.
   A successfully added certificate is labeled as, **This certificate is marked as trusted for...plugin verification.**

3. On a Linux OS:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Prepare the self-signed certificate file as a Privacy Enhanced Mail (PEM) file.

b. Copy the certificate to the Certificate Authority (CA) directory, which is the directory where the certificates are stored.

- On a Debian/Ubuntu-based system, run the following command:

```
sudo cp <path to certificate
 file> /usr/local/share/ca-certificates/
```

- On a Red Hat/Fedora-based system, run the following command:

```
sudo cp <path to certificate
 file> /etc/pki/ca-trust/source/anchors
```

c. Update the CA certificates:

- On a Debian/Ubuntu-based system, run the following command:

```
update-ca-certificates
```

The command reads the certificate bundle in the `/usr/local/share/ca-certificates` directory, updates the certificate bundle, and stores it in `/ext/ssl/certs/ca-certificates.crt`.

- On a Red Hat/Fedora-based system, run the following command:

```
sudo update-ca-trust extract
```

Result

Your self-signed certificate is on the OS that you are working with.

Import a self-signed certificate

Import a self-signed certificate in a Windows system by using the Certificate Import Wizard. The Certificate Import Wizard is required to complete the self-signed certificate import process on a Windows Operating System (OS).

Before you begin

Ensure that you have followed the procedure for importing a self-signed certificate in a Windows system, as described in Add a self-signed certificate to the OS truststore.

Role required: agent_client_collector_admin

Procedure

1. On the first page of the wizard, select **Next**.
   The truststore is typically located in the Trusted Root Certification Authorities directory of the certificate manager.
2. Browse to the location where you have the self-signed certificate file.
   The file is typically located in the agent's `config/cert` directory.
3. Select and hold (or right-click) the certificate file and select **Install certificate**.
4. Select **Next**.
5. Select the option to place the certificate in the Trusted Root Certification Authorities truststore.
6. Select **Next** and **Finish** to complete the wizard.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The certificate is imported into the truststore.

What to do next
If a security warning appears, select **Confirm** to confirm the import.

Revoke Agent Client Collector certificates

Stop communication between the agent and ITOM cloud services by removing an Agent Client Collector certificate. For example, there might be a security breach due to which you want to stop communication by revoking the agent's certificate.

Before you begin
Role required: agent_client_collector_admin

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select the agent for which you want to revoke the certificate.
3. Select the **Revoke Agent Certificate** button.
4. Select the reason for certificate revocation.
   For a list and explanation of the available revocation reasons, see Agent Client Collector certificate revocation reasons.
5. Select **OK**.
   The agent's certificate is revoked and the ITOM cloud services layer has its certificate revocation list (CRL) updated.

Result
Communication between the agent and ITOM cloud services is stopped.

Related topics
Add a self-signed certificate to the OS truststore

Agent certificate rotation

The Agent Client Collector certificate is valid for two years and must be rotated before it expires to avoid issues with agent connectivity. When expiration is approaching, the agent initiates a certificate rotation request.

When an Agent Client Collector's certificate expiration is approaching, system properties verify that a valid certificate is in place when the existing certificate expires.

- *certificate-rotation-days-out*: Indicates the number of days before certificate expiration that an agent attempts to rotate its certificate. Configured in the `acc.yml` configuration file. Default = 28.
- *sn_agent.certificate_rotation_days_out*: Indicates the number of days before certificate expiration that the system accepts a certificate rotation request for a specific agent. Configured on the System Properties page (**All** > **System Properties** > **All Properties**). Default = 28.

Related topics
View the Agent Client Collector configuration file for an agent

Configuration file options

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Run Certificate Discovery via Agent Client Collector for Visibility Content

Discover TLS/SSL certificates used by ports running on the agent's server. The Certificate Inventory and Management application uses this information to manage TLS/SSL certificates.

Before you begin

Verify that the latest version of Agent Client Collector for Visibility Content (ACC-VC) is installed.

Verify that the latest version of Certificate Inventory and Management is installed.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Enable the *TLS SSL Certificate Capture* policy.
3. Navigate to **All** > **System properties** > **All properties**.
4. To discover certificates on specific ports:
   a. Locate and select the *sn_acc_vis_content.tls_ssl_scanner_ports* property.
   b. Enter a comma-separated list of port numbers on which to scan for certificates in the **Choices** field.
   By default, the values are identical to those values configured for the *tls_ssl_certs* property.
5. To discover certificates on all available open ports:
   a. Locate and select the *sn_acc_vis_content.tls_ssl_scan_all_open_ports* property.
   b. Assign the value **true**.
6. **Optional:** To store the original certificate in the instance, set the *sn_acc_vis_content.tls_ssl_keep_original* property to **true**.

Result

The Agent Client Collector collects the data and discovers the TLS/SSL certificates. This data is the same as what is gathered by the IP-based *tls_ssl_probe* property.

The TLS/SSL certificates are populated in the following tables:

- Discovered Certificates [sn_disco_certmgmt_certificate_history]
- Unique Certificates [cmdb_ci_certificate]
- Installed Certificates [sn_disco_certmgmt_cmdb_installed_certificate]

A CI relationship is created between the discovered certificates and the corresponding Host CI (in this case, the agent's server).

Related topics

Certificate Inventory and Management

Run Certificate Discovery via port scans

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Collect data from your system devices

Run a check command on your server or database to gather data from those devices. Depending on the check that is invoked, collected data may be monitoring data, visibility data, log data, or user metrics.

Policies

A policy consists of the Configuration Items (CIs) monitored by the Agent Client Collector, and the checks that run on those CIs. Policies define which checks to run, as well as when and on which CIs to run them.

Checks

A policy runs a set of checks. Each check is defined with the following attributes:

- Command to run
- Check type
- Command timeout
- Plugins to be used (`tar.gz` files that are securely and automatically downloaded to the agent)

Determining when checks run

Policy checks run based on one of the following scheduling models:

- Interval-based scheduling: The minimum number of seconds between each check run. For example:
  - An interval value is at least once per hour (at least once every 3,600 seconds).
  - The agent was down during the scheduled interval when the check was to run. When the agent is back up, it runs the check within one hour.

  If a check interval is less than one hour, the check is executed based on the indicated amount of time since the check's most recent run. For example, for a check scheduled to run once every eight hours:

  - The check runs at 00:00 (midnight)
  - The check runs again at 08:00 (8 AM)

  However, if the agent is down at 8:00 and is back up at 09:03, the check runs within an hour; that is, between 09:03 and 10:03. If the check runs at 09:57, the next run is at 17:57 (5:57 PM). In these cases, the check picks the running time at random to ensure running of once every eight hours.

  A check completes its run when the results are sent to the MID Server (in a MID deployment) or ITOM Cloud Services (ICS) Pod (in a MID-less deployment).

- Cron-based scheduling: Runs at fixed times or intervals (for example, every day at 6:00 PM or every two hours).

Determining where checks run

Policies monitor the configuration items (CIs) that are mapped to the agent hosts. To determine the specific CI types to be monitored, you can create filters. For details on configuring policy filters, see Create a service filter for a policy.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Updating policy mappings

When activating a policy, Agent Client Collector Framework maps monitored CIs to a policy. This policy mapping repeats every seven days after the initial mapping. Every 15 minutes (by default), the Agent Client Collector monitors the CMDB for any changes matching the filter conditions, and updates the mapping accordingly. To modify the frequency by which mapping between the CMDB and CI is updated, modify the sn_agent_sync_filters_interval_min property on the System Properties page (**All** > **System Properties** > **All Properties**).

Verify data collection in Agent Client Collector

Collect data by gathering essential information from an agent's host system before executing any checks or policies. This process ensures that the Agent Client Collector has accurate and up-to-date data on the infrastructure, processes, and applications running on the host.

Before you begin
Role required: agent_client_collector_admin

About this task
Collected data is stored as Configuration Item (CI) data, which serves as a foundation for monitoring, alerts, and Metric Intelligence in the ServiceNow® instance.

Before checks and polices can be executed, Agent Client Collector collects data on:

- Host system details (such as hardware and the system's OS version)
- Installed applications and services
- Running processes
- Resource usage metrics (such as CPU, memory, and disk utilization)

Collected data is stored in the ServiceNow instance in either the CMDB (Configuration Management Database) or events table.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. On the Agent Client Collectors page, locate the **Host Data Collection** column for an agent. A green Collected icon Collected indicates successful data collection for the specific agent.

Create an Agent Client Collector policy

Policies consist of the CIs monitored by the Agent Client Collector and the checks that run on those CIs. When creating a policy, you configure a filter which determines the CIs on which the checks are to run. For example, a policy to run checks on all Apache web servers. You can create new policies or edit the default policies.

Before you begin
Role required: agent_client_collector_admin

About this task
When associating a check definition with a policy, a copy of the definition is created as a check instance. These values appear on the **Check Instances** tab on the bottom of the **Policy** form. Any changes you make to a check instance affect only the instance in the policy; the original check definition remains unchanged. When changing the original check definition, the changes do not apply to already created check instances. It is recommended you make a copy of a check definition if you are changing its default parameters after having initiated one or more check definitions in a policy.

Some policy checks require credentials, while others don't. For example:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The Agent Client Collector does not require credentials for checks that run on its own server. However, when accessing software requiring root or administrator access (such as MySQL), credentials are required.
- When the Agent Client Collector is used as a proxy:
  - Credentials must be invoked to allow the agent to connect to the remote device over SNMP or other protocols.
  - Increase the agent's *cpu_percentage_limit* value in the `acc.yml` file, as indicated in the blue box on the **Proxy Settings** tab.
- When the Agent Client Collector uses the HTTP endpoint:
  - If the URL does not require logging in (such as with a public website), no credentials are required.
  - If the URL does require logging in (such as with an API endpoint), credentials are required.

You assign a credential alias in the policy's **Credentials** tab, as described in the following procedure.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Click **New**.
   The **Policy** page appears.
3. Enter values in the fields on the page, as described in the following table:

Policy Definition Form

Field	Description
Name	A descriptive name for the policy.
Description	Description of the policy.
Publish status	Hard coded as **Draft**, which means that the policy has not yet been published. You cannot edit this field.
Hierarchy	Hard coded as **None**. When a child policy is added to the policy, the value changes to **Parent**. Child policies have a value of **Child**.

4. On the **Monitored CIs** tab, select one of the following options:
   - **Manual calculation**: Select to exclude the policy from scheduled policy calculation, ignoring any changes to the policy's impacted CIs.

     Select this option when the policy monitors a single CI. Otherwise, the Refresh and Publish Monitoring Policies scheduled job takes an abnormally long time to complete.
   - **Monitored CI type by filter**: Configure the following sub-fields:
     - **Monitored CI type**: Select the CI type which you want the policy's checks to monitor.
     - **Filter**: Configure a filter so that the policy's checks monitor only CI types meeting the specified criteria.
   - **Monitored CI type by script**: Enables you to specify the monitored CIs using a script. Using a script allows you to create a CI filter for tables related to each other, and not just for a single

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

table. For example, you can set a filter in both a Linux servers table and an Oracle table when searching for a CI. You return an object containing the table and filter at the end of the script.

An example of a script appearing in the **Monitored CI Script** field is:

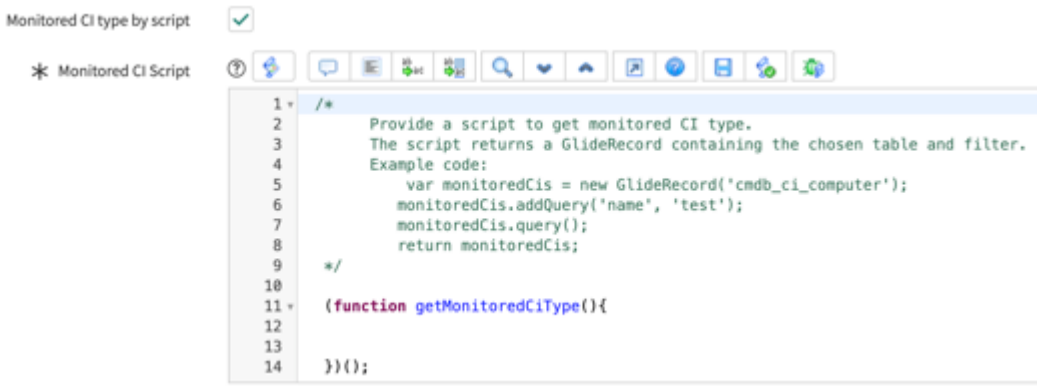

- **Monitored CI type by CMDB Group**: Enables you to specify the monitored CIs by using CMDB group queries. When selecting this option, the **Monitored CMDB group** field appears, where you select a CMDB group to associate with the policy.
  - For details on CMDB groups, see CMDB groups.
  - For details on building a CMDB query using the CMDB query builder, see Build a CMDB query using the CMDB query builder.

5. Select checks in the **Available** cell on the **Checks** tab, and move them to the **Selected** cell for them to be included in the policy.
   Checks can be selected multiple times, when you are monitoring more than one process. You can also select a group of checks in the **Filter checks by groups** field, which presents checks of the selected group in the **Available** cell.
6. **Optional:** Configure a proxy server on the **Proxy Settings** tab.

   (Optional) Configure a proxy server only when using the agent as a proxy to report data on remote machines. For details, see Configure an agent on a proxy server.
7. On the **Scheduling** tab, select one of the following:
   - **Interval-based scheduling**: Configure the time interval (in seconds) to indicate the frequency with which the policy's checks will run.
   - **Cron-based scheduling**: Enables configuring the policy to be active only during a specific time or time frames.

     a. Select the unlock icon (

) to enable selecting cron expressions.

     b. Select the search icon ( ) to select a cron expression. Select **New** to create a new cron expression, or select one of the cron expressions that come with the base system, which are:
     - Every hour nightly 5:00pm-7:00am
     - Every minute daily 8:00am-4:59pm
     - Every minute Mon-Fri 8:00am-4:59pm

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When creating a new cron expression, you must enter it in valid cron format.

The specified cron expression operates in the time zone of the machine hosting the agent.

Note: If an **Interval** or **Cron** value was specified for the check definition, that value overrides the value configured in either of these fields.

8. On the **Credentials** tab, select one of the following:
   - **Credential name**: Select credentials to be assigned to the policy. The available credentials are those displayed on the associated check definition's **Check Secure Parameter Definitions** tab.
   - **Credential alias**: Select the search icon to select a credential alias to be used by the policy's checks to connect to the monitored CI. The available options are those created on the **Connection & Credential Aliases** page. For details on credential aliases, see Credential aliases for Discovery.
9. Click **Save**.
   The following buttons appear:
   - **Publish**: Publishes the Draft policy, moving its **Publish status** to **Queued**. A queued policy is processed by the policy publishing job. The job calculates which agents run the policy and then sends the policy to the relevant MID Servers. The job then moves the policy's **Publish status** to **Published**. A published policy is active on the agent.
   - **Delete**: Deletes the policy.

Related topics

Create a policy hierarchy

Create a policy hierarchy

Policies are attached to groups of related CIs. If you have a subgroup of related CIs, you can attach a new policy to the subgroup using a policy hierarchy, without creating the new policy from scratch.

Before you begin

Role required: agent_client_collector_admin

About this task

Policy hierarchy enables you to use a policy as a template to create a new policy, which is then attached to a subgroup of the original policy's CIs. The original policy is the **parent** policy, and the new policy is the **child** policy.

Note: You cannot create a child policy when a proxy agent is in place for the parent policy.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. On the **Policies** page, select a published policy with a Hierarchy value of either **None** or **Parent**.
3. Select the **Create Child** button at the top of the page or underneath the **Monitored CIs** tab.
   A new (child) policy opens with the same parameters as the original policy. The name of the policy is *<Policy name><date and time created>*.

Note: The child policy inherits the parent policy's parameters only when it is first created; there is no continuous inheritance from the parent to the child.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following fields appear hard-coded with the indicated values:

Child policy fields

Field	Value
Publish status	Draft
Hierarchy	Child
Parent	<Name of the parent policy>

The **Order** field value indicates the priority by which the agent attaches itself to the policy. A lower number indicates higher priority, and each created child policy takes on a higher value than the policy preceding it. You can modify this value, as needed.

4. Modify the policy's parameters, as needed.
The following buttons appear on the child policy page:

Button	Description
Compare to Parent	Select to display the differences between the child policy and its parent policy.
Return to Parent	Select to leave the child policy page (without saving) and return to the parent policy.

The following buttons appear on the parent policy page:

Button	Description
Deactivate	Appears for an active policy. Enables you to select from the following: ○ Deactivate the parent policy and all of its child policies ○ Deactivate only the parent policy
Activate	Appears for an inactive policy. Enables you to select from the following: ○ Activate the parent policy and all of its child policies ○ Activate only the parent policy
Delete	Enables you to select from the following: ○ Delete the parent policy and all of its child policies ○ Delete only the parent policy

5. Select **Save**.
Child policies appear on the **Child Policies** tab of the parent policy.

Edit a published policy

A policy with a status of **published** has been passed through the MID Server and sent to the agent. To edit a published policy, you must enable editing.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: agent_client_collector_admin

About this task
You select a Published policy and then select the **Edit in Sandbox** option to enable editing of the policy. You edit the policy in the sandbox and then save, republish, or delete the policy, as needed. The following diagram illustrates the flow of editing a published policy:

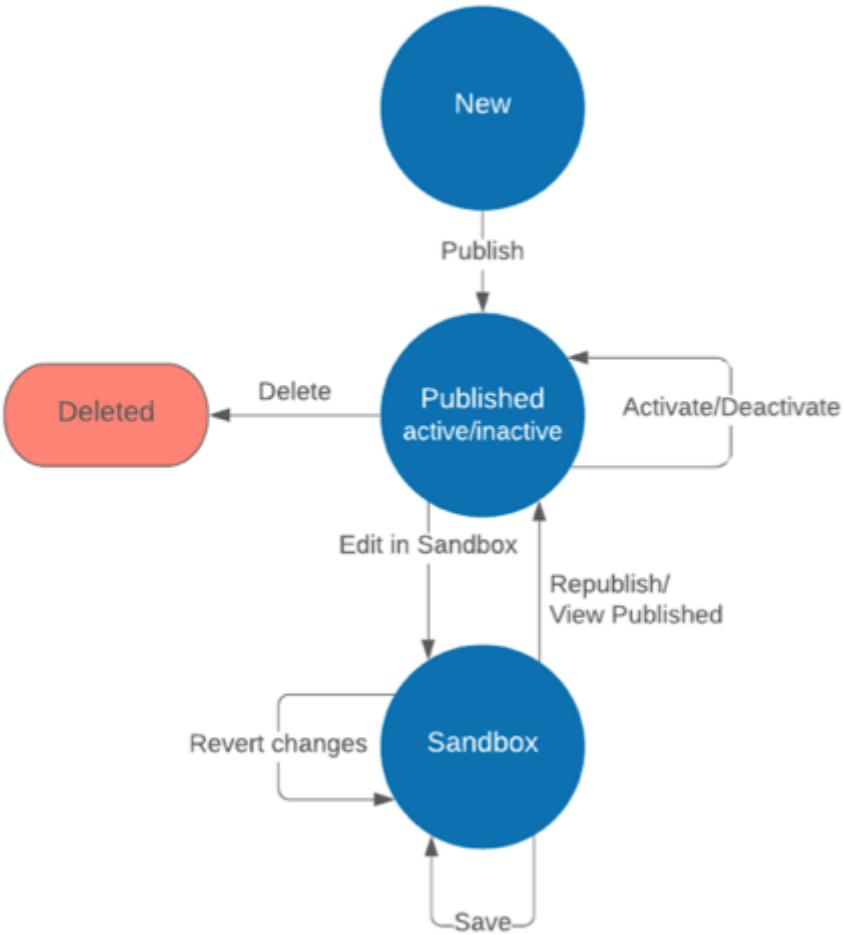

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Click a policy for which the **Publish status** is **Published**.
3. Activate or deactivate the policy, as needed, by clicking the **Activate** or **Deactivate** button.
4. Select the **Edit in Sandbox** button.
   The policy is now in **Draft view**, and its fields are editable.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:

- When you select this option, the policy's check instances are also editable.
- Although the policy is in **Draft view**, its **Publish status** remains **Published**.

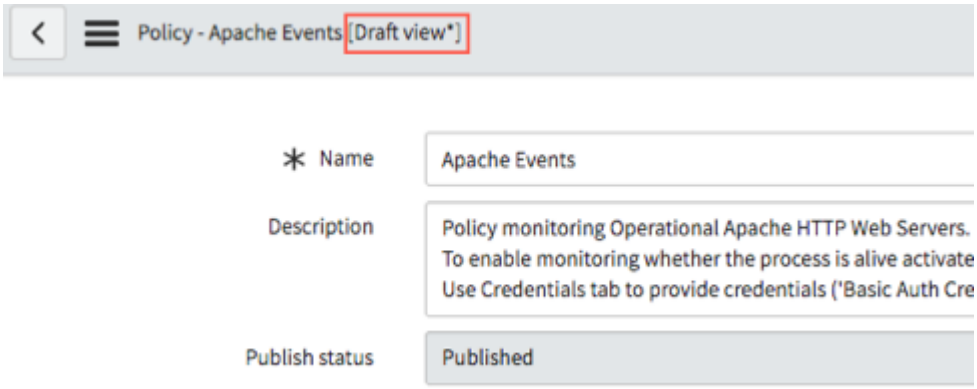

While editing the policy's fields, you can select the following options, as needed:

Edit policy buttons

Button	Description
Revert All Changes	Reverts the changes you entered into the Sandbox policy fields to the most recently published values. The policy's **Publish status** returns to **Published**, and remains in **Draft view** for editing. This button appears only after entering updates into the Sandbox policy fields.
View Published	Returns to the most recently published policy. The policy returns to **Published** view.
Save	Saves the changes you entered into the Sandbox policy fields. The policy remains in **Draft view**, and the **Publish status** is **Published***. The asterisk indicates that the current published version is not the most up to date, since there are changes that have not yet been published.
Republish	Publishes the updates you entered into the Sandbox policy fields. The policy is saved with your updates and opens as the published policy. ○ The published policy's **Publish status** is set to **Queued** if the policy changes result in either: ▪ A change in which agents are running the policy (that is, a change to the monitored CI filter). ▪ A change in the policy's check instances.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Button	Description
	The policy's publishing job recalculates Queued policies and sends them to the relevant MID Servers. ○ The published policy's **Publish status** is set to **Ready to publish** if the changes to the policy do not require any agent recalculation (that is, a change to the policy interval). The policy's publishing job sends Ready to publish policies to the relevant MID Servers, without recalculating those policies. After the policy publishing job processes a policy, the policy **Publish status** moves to **Published**. This button appears only after entering updates into the policy fields, for a policy whose **Publish status** is **Published***.
Delete	Deletes the policy. This button appears only when viewing the published policy, and deletes both the Published and Sandbox policies.

5. Select checks in the **Available** cell of the Checks tab, and move them to the **Selected** cell for them to be included in the policy.
   Checks can be selected multiple times, when you are monitoring more than one process. You can also select a group of checks in the **Filter checks by groups** field, which presents checks of the selected group in the **Available** cell.
6. In the **Check Instances** table, view the check instances that are configured for the policy, or select a check instance to edit it.
7. On the **Agents** tab, view the agents that are executing the policy.

   The list of agents is calculated based on the monitored CIs.

   This tab appears only when viewing a published policy.

Create a service filter for a policy

Configure a service filter so that the policy monitors only those services that match the filter.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
   The **Policies** page appears.
2. Select a policy.
   The details page of the policy appears.
3. Select the **Edit in Sandbox** button to enable editing the form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. On the Monitored CIs tab, select the **Filter Monitored CIs by service** check box.
   The **Application Service filter** section appears.

Application Service filter section

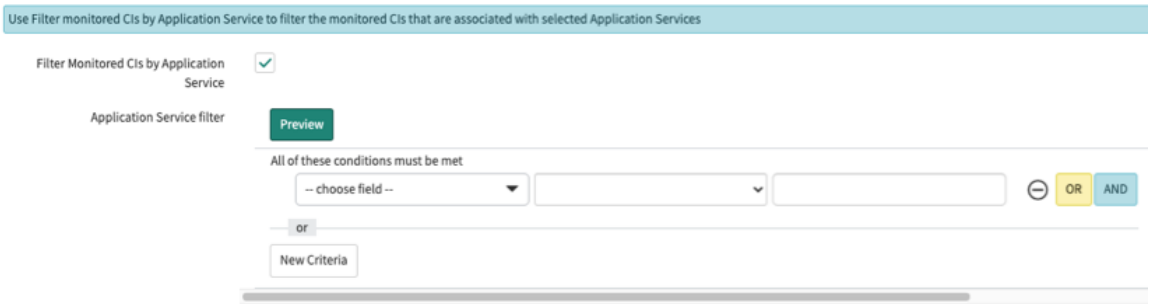

5. Configure conditions in the fields and select **Save** at the bottom of the page.
   The policy is configured with the indicated service filter. Only CIs matching the service filter are included in the policy's logs.

Export and import an Agent Client Collector policy

You can move a single or multiple Agent Client Collector policies from one ServiceNow instance to another by exporting the policy and then importing it into the target instance. Only published policies can be exported.

Before you begin

Ensure that you published the policies you want to export on the source instance.

Note: When working in a domain separated environment, imported policies are moved from the source instance domain to the identical domain on the target instance. Therefore, you must ensure that you have defined the same domains on the target instance as you have on the source instance, either by cloning the source instance or by exporting and importing domains from the source to the target instance.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Select the check box of the published policies you want to export.
3. In the **Actions on selected rows...** list underneath the list of policies, select **Export policies**.
   A confirmation dialog box appears.
4. Click **Yes**.
   The selected policies are exported as an update set file in XML format, which is downloaded to your browser. The generated update set has the following name convention: `AccPolicyExport <Current Date and Time>`. For example, `AccPolicyExport 2020-08-11 15:35:05`.
5. Import the policy to the target instance.

   a. Navigate to **System Update Sets** > **Retrieved Update Sets**.
      The **Retrieved Update Sets** page appears.

   b. Right-click on one of the table headers and select **Import XML**.
      The **Import XML** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

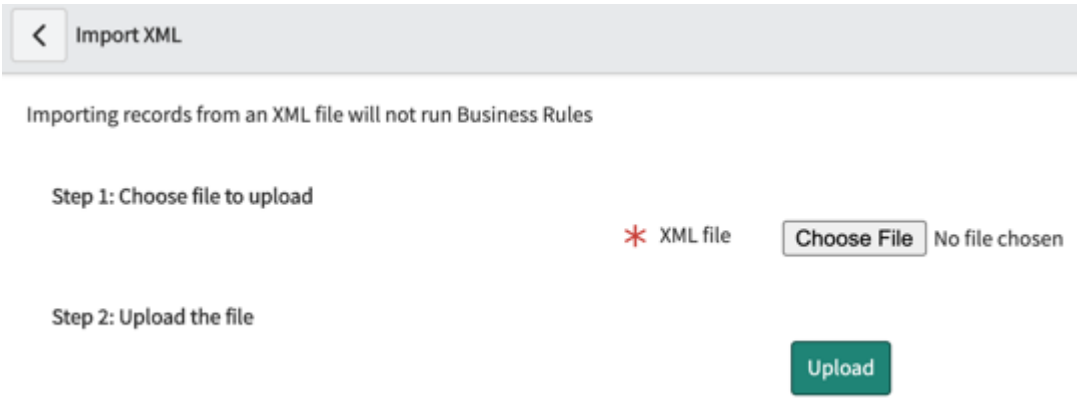

c. Click the **Choose File** button and select the previously downloaded update set to import.

d. Click **Upload**.

> **Note:** When the exported content is in a scope, the imported update set appears on the **Retrieved Update Sets** page as multiple records. The **Global** scope record is the parent, and the other records appear with the Global record as their parent update set. Only the Global update set record can initiate the import process.

e. Select the **Global** update set.

f. On the **Retrieved Update Sets** page, click **Preview Update Set Batch** to verify whether there are any conflicts with existing policies on the target instance.
Resolve any conflicts, as needed.

g. If no conflicts exist, click **Commit Update Set Batch** to validate and commit the update set.

h. Navigate to **Agent Client Collector** > **Policies** and verify that the imported policies appear on the Policies list.
The imported policies are synchronized with the relevant MID Servers on the target instance and are ready for immediate use.

> **Note:** The import process overrides any existing policies on the target instance and sets the imported policies to **Published**. However, draft copies of policies created on the target instance are not modified. Therefore, delete draft records of the imported policies after completing the import process, as these records are not synchronized with the published policies on the target instance.

Create and edit checks

You can select a check based on the type of entity (such as a server or an application) you want to monitor. You can view and edit the default checks or create new checks, as needed.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Check Definitions**.
2. Click **New**.
   The **Check Definition** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Using Ruby scripts in a Windows environment causes high CPU consumption.

3. Enter values in the fields on the page, as described in Agent Client Collector check definition page.
4. On the **Parameters** tab, enter parameter definitions, as needed.
   Parameters are used for the **Command** field when the Simple template is used, as described in Agent Client Collector check definition page.

   The information you can add for parameters is described in Check definition form parameters tab.

Test a check definition or check instance

When working with a check definition or check instance, you can test the check against its assigned agent to ensure that the check is configured properly.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Check Definitions**.
2. Select a check definition from the list.

   Note: When running a Test Check in a check instance:

   a. Navigate to **Agent Client Collector** > **Policies**.

   b. Select a policy.

   c. Select a check instance in the policy.
   Continue with the following steps in this procedure.

3. In the **Related Links** section, click **Test Check**.
   The **Test Check** dialog box appears.

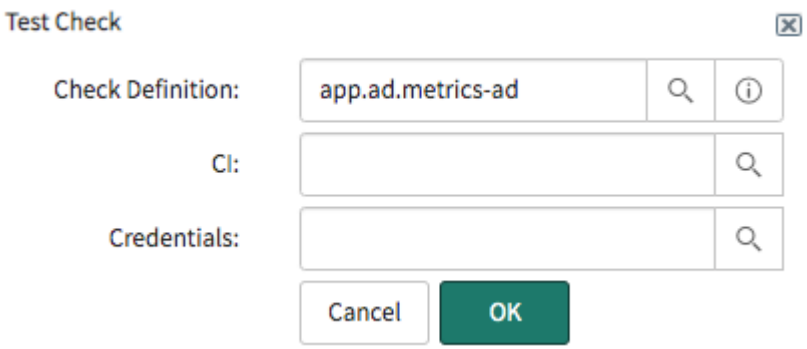

4. Configure the fields in the dialog box.

Test Check dialog box fields

Field	Description
Check Definition	Displays the name of the selected check definition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
CI	Select the CI on which to run the check. The available options are those CIs which meet both of the following criteria: ○ Belong to the group specified in the **Check group** field. ○ Belong to an agent with **Up** status. If the **Check group** field is empty, the value defaults to **cmdb_ci_computer**.
Credentials	Select the credentials of the user whom you want the agent to connect with. This field appears only if you have configured secure check parameters on the **Check Secure Parameter Definitions** tab in the Check Definition.
Proxy agent	Select a proxy agent to be associated with the check. This field appears only when working with a check instance and when one of the options has been selected on the **Proxy Settings** tab in the instance's policy.

Create check parameters for a check

Create a check parameter to add a configurable input that customizes how a check runs, without changing the underlying script or command.

Before you begin
Role required: agent_client_collector_admin

About this task
A check parameter contains:

- Check Parameter Definition: defines each check type. It declares the parameter's name, meaning, default value, and the command-line flag it maps to.
- Check Parameter Value: specify the value of the defined parameter which is specific to a check instance.

Note: The meaning and valid values of any parameter depend on the related check definition.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Check Definitions**.
2. Select a check definition from the displayed list.
   The **Check Definition** page for the selected check appears.
3. Scroll to the bottom of the page and select the **Check Parameter Definitions** tab.
4. Select **New**.
   The **Check Parameter Definition New Record** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

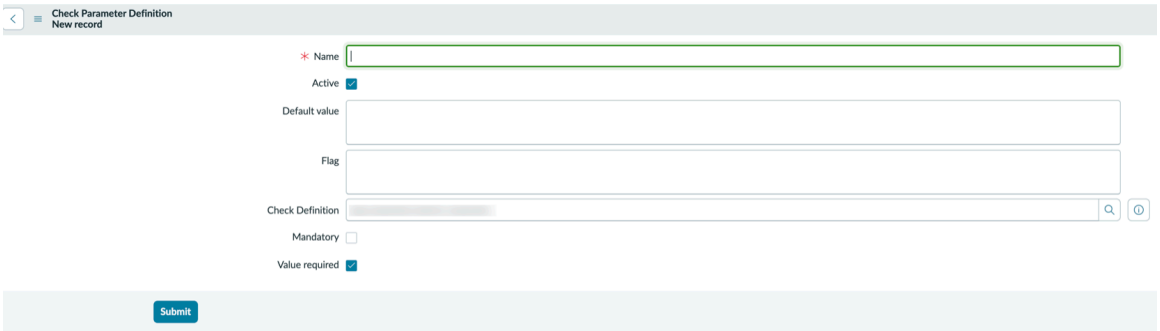

5. Configure the fields on the page.

Field	Description
Name	Name of the parameter, formatted as a refer ence prefix.
Active	Activates the check parameter on selection.
Default Value	Indicates that the parameter must contain a default value.
Flag	A prefix that must be added before the para meter value in the command. For example `-p`.
Check Definition	The name of the check definition connected to the parameter.
Mandatory	Indicates that the parameter is mandatory.
Value required	Indicates that a value must be provided for the parameter.

Note: The **Active**, **Flag**, and **Value required** options are displayed only when the **Command Auto Generation** option is selected.

6. Select **Submit**.
The configured parameter appears in the table on the **Check Parameter Definitions** tab.

   Repeat this procedure to enter multiple check parameters.

7. Enter the check parameters in the **Command** field of the check definition.
Each check definition has a **Command** field containing a template. Parameter values are substituted into this template at execution time:

```
check-kernel-parameter.rb {{if .labels.params_parameter_value}}
 -v {{.labels.params_parameter_value}} {{end}}
 {{if .labels.params_parameter_name}} -p
 {{.labels.params_parameter_name}} {{end}}
```

   - `.labels.params_<parameter_name>` refers to the value configured for that parameter.
   - `{{if ...}} ... {{end}}` adds the flag and value only when the parameter has a value. If the parameter is empty, its flag is omitted.
   - The flag prefix such as `-v`, `-p`, and so on is picked from the **Flag** field on the check parameter definition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This command is assembled automatically from the configured parameter values.

Example:

The `os.linux.check-kernel-parameter` check defines two parameters:

Example parameter values

Parameter name	Flag	Purpose
`parameter_name`	-p	The Linux kernel parameter to check. For example, `kernel.shmmni`
`parameter_value`	-v	The expected value for that kernel parameter.

When you configure `parameter_name` with the following settings:

```
Name:           parameter_name
Active:         true
Mandatory:      true
Flag:           -p
Value required: true
Value:          kernel.shmmni
```

You also specify the `parameter_value` as 4096. The check runs as follows:

```
check-kernel-parameter.rb -v 4096 -p kernel.shmmni
```

Related topics

Create a check definition

Agent Client Collector check definition page

Create secure parameters for a check

When creating a check definition or check instance, you can configure the parameters you want to be secured when the agent executes the check. During check execution, the secured parameters are obfuscated, securing their information. Only credential information is obfuscated.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Check Definitions**.
2. Select a check definition from the displayed list.
   The **Check Definition** page for the selected check appears.
3. Scroll to the bottom of the page and select the **Check Secure Parameter Definitions** tab.
4. Click **New**.
   The **Check Secure Parameter Definition New Record** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

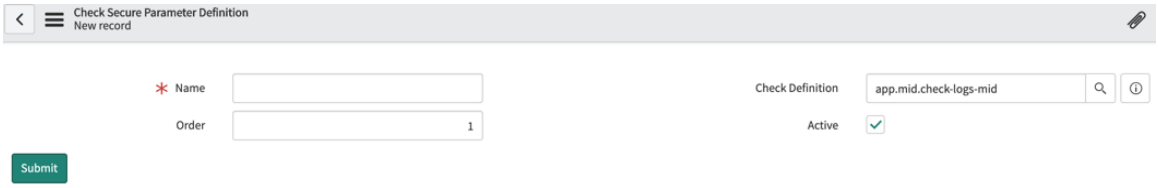

5. Configure the fields on the page.

Field Name	Description
Name	The name of the parameter, formatted as a reference prefix. For example, `cred_` is a reference prefix for the credentials table.
Check Definition	The name of the check definition connected to the parameter.
Order	A number indicating the order in which the parameter is sent to the check command/script. For example, if you configure the following parameters: ○ Name=**cred_user_name**, Order=**1** ○ Name=**cred_password**, Order=**2** The *`username`* parameter is sent first to the standard input, and the *`password`* parameter is sent second to the standard input. The *READ* command is performed first on the *`username`* parameter and then on the *`password`* parameter. You can then use *`$username`* and *`$password`* in your *Bash* script.
Active	Select the check box to activate the secure parameter.

6. Click **Submit**.
   The configured parameter appears in the table on the **Check Secure Parameter Definitions** tab.

   Repeat this procedure to enter multiple secure parameters.

7. Enter the secure parameters into the **Command prefix** field of the check definition.

Create a check type

Create a check type to execute the `osquery` command on the Agent.

Before you begin

Role required: agent_client_collector_integration or agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In an Event Management instance, enter `sn_agent_check_type.list` into the navigation bar.
   The **Check Types** page appears.
2. Click **New**.
   The `Check Type - New Record` page appears.
3. In the **Name** field, enter `osquery`.
4. In the **Instance Script** field, enter the following script:

```
for (var index = 0; index < checkResults.length; index++) {
var check = checkResults[index].check;
gs.info('result ' +
index + ': requestId: ' +
check.requestId + ' agent_name: ' +
checkResults[index].agent_id +
' ci_id: ' + check.ci_id +
' status: ' + check.status +
' output: ' + check.output);

}
```

   You can modify this script, as needed, but this initial script helps you to verify that your setup was successful.
5. Click **Submit**.

Create a check definition

Create a check definition to execute the `osquery` command on the Agent.

Before you begin

Role required: agent_client_collector_integration or agent_client_collector_admin

Procedure

1. In an Event Management instance, navigate to **Agent Client Collector** > **Check Definitions**.
2. Click **New**.
3. In the **Name** field, enter `util.osquery`.
4. In the **Check type** field, enter `osquery`.
5. In the **Command** field, enter the following script:

```
osqueryi  --logger_min_status 1 --json
 "{{.labels.params_query}} "
```

6. In the **Plugins** field, enter the `osquery` plugin.
7. In the **Parameters** section, enter the following values for a check parameter definition.

Column	Value
Name	query
Default value	select * from logged_in_users
Mandatory	true

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Click **Test check** and select one of the available agents.
   The test result appears, indicating its success or failure.

Enable checks from the community for Agent Client Collector

You can take checks from the github community and customize them for use in the Agent Client Collector (ACC), or you can compose your own scripts. Create a plugin with the customized Sensu check and install it on a ServiceNow instance.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Copy a check from the github community to your local machine.
2. Create a `tar.gz` file to be appended to your plugin:
   a. Create a folder on your local machine.
   b. Create a sub-folder called `bin`.
   c. Place the checks you want to run (scripts and executables) in the `bin` folder.
   d. In a Linux environment, give executable permissions to the checks by running *`chmod +x filename`* on each of the checks.
   e. Navigate to the parent folder (above the bin folder).
   f. In a Linux environment, run the *`tar -C <parent_folder_name> --zcvf <my_plugin_name>.tar.gz *`* command.
   In a Windows environment, use an archiving utility that supports the `tar.gz` format.
   g. In a Windows environment, ensure that you include the following scripts in the `bin` folder:
      - A batch script configured to run on the check command and initiate the run. For example, *`powershell <%~dp0>\test.ps1`*

        The *<%~dp0>* variable directs to the drive (*d*) and path (*p*) of the batch file.
      - The relevant script (for example, `test.ps1`) to be executed in the batch script.
3. Create a new ACC plugin, as described in Create and edit Agent Client Collector plugins.
4. Test the new plugin.
   a. Create a new check definition for the customized Sensu check.
   For details, see Create and edit checks.
   b. On the **Check Definition** page, navigate to the **Related Links** section and select **Generate allow list content**.
   c. Copy the entries you want to include in the allow list and add them to the agent's `check-allow-list.json` file, either manually or using an automation tool.
   The location of this file is indicated in the `allow-list` parameter of the agent's `acc.yml` file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Create a new policy and add the newly created check definition.
For details on creating a policy, see Create an Agent Client Collector policy.

e. Select the check definition and in the **Related Links** section, select **Test check**.

5. Enable credentials when testing a check:

a. Navigate to **Connections & Credentials** > **Credentials**.

b. Click **New**.

c. Select the type of credentials you want to create.
For details on creating credentials, see Create a Connection and Credential alias.

d. Enter values in the fields on the page to assign credentials to your selected credential type.

e. Click **Submit**.

6. Create secure parameters for your check, as described in Create secure parameters for a check.

Verify agent functionality

Verify that an agent is running properly by performing a self-test on the agent. If one or more of the tests fail, you can diagnose the agent problem and view a potential resolution.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select an agent that is running (**Status = Up**).
3. In the **Related Links** section, select **Run self-test**.
4. The **Agent Self Test Runs** page appears, displaying the previous self-tests run on the agent.
If the current test doesn't appear, refresh the page.
5. Hover next to the entry for the current test and select the info icon .
The **Agent Self Test Run** pop-up window appears.
6. Select **Open Record**.
The Agent Self Test Run page appears, displaying information on the self test.
7. In the **Agent Self Test Results** table at the bottom of the page, locate a test with **Status = Fail**.
8. Hover next to the entry for the failed test and select the info icon .
The **Agent Self Test Results** pop-up window appears.
9. Select **Open Record**.
The Agent Self Test Results page appears, displaying information on the failed test.

Result

In the **Agent Self Test Logs** table on the bottom of the page, you can view the level of the failed test (ERROR or WARNING), and a resolution for the failed test.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View the Agent feature matrix

The Agent Client Collector Agent feature matrix displays the availability of Agent Client Collector features. The matrix displays data in a graph and a table.

Before you begin

Role required: agent_client_collector_admin

About this task

The Agent feature matrix displays the granularity of Agent Client Collector features that are fully supported, partially supported, and not supported. A partially supported feature may have instance and framework requirements met, but runs on at least one agent version which does not meet the feature requirements.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agent Feature Matrix**.
   The **Agent feature matrix** appears.

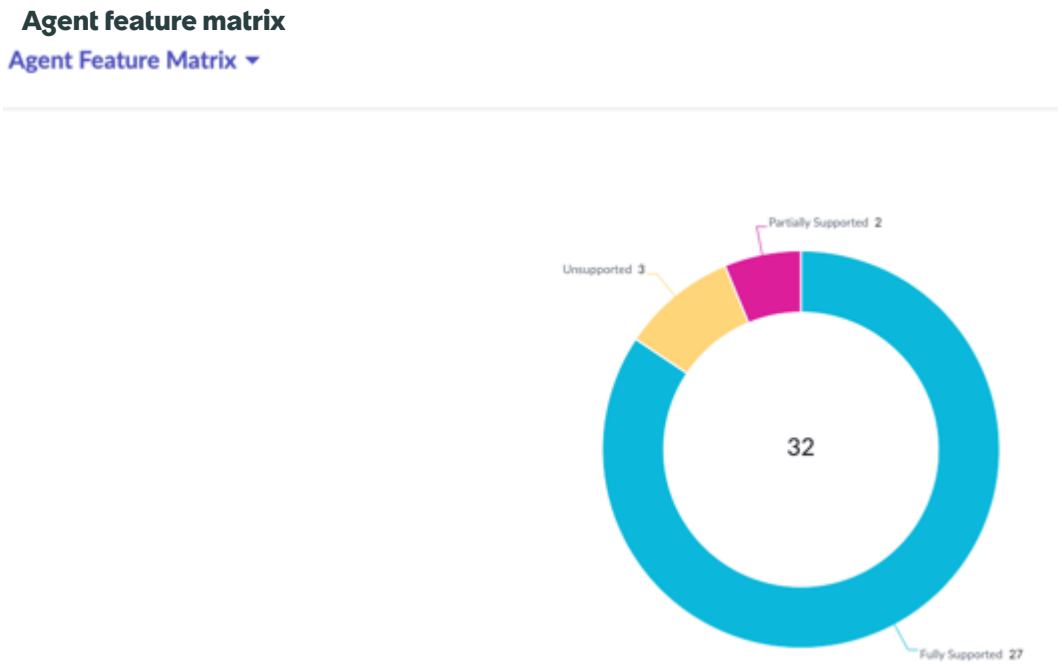

   The matrix shows the support status of Agent Client Collector features that are fully supported, partially supported, and not supported (Unsupported).

2. Select one of the categories on the graph.
   A list of all features in the category appears on the **Agent Feature Support Status** page.

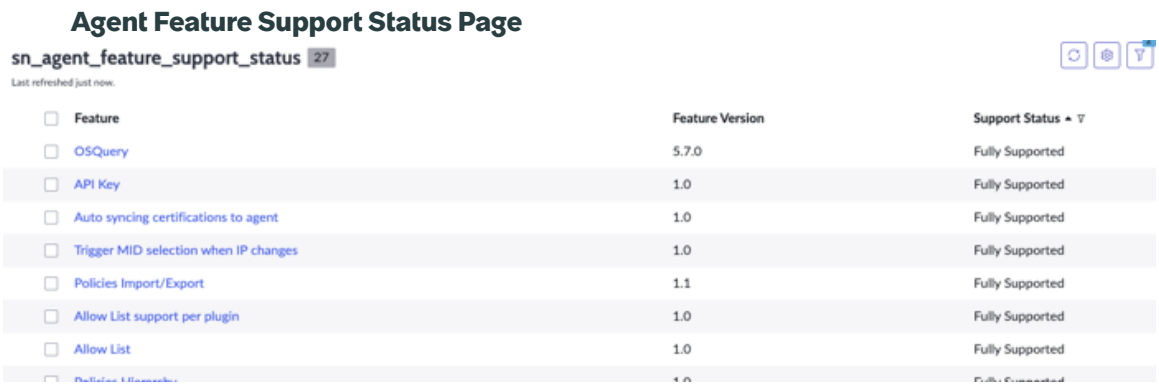

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select one of the features.
   The **Agent Feature Support Status** page for the specific feature appears:

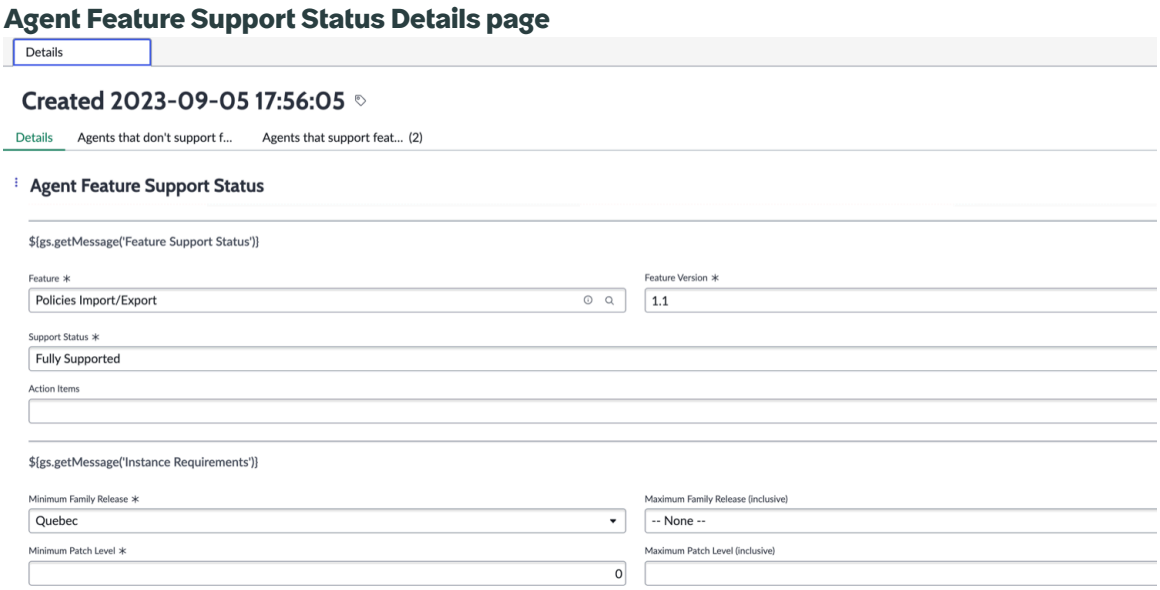

   The page displays the following:

   - Sections indicating the support status of the feature, and the various instance requirements of the feature. The **Action Items** field displays the actions you need to take to fully support a feature that is either partially supported or unsupported.
   - Instance requirements and agent framework version requirements that have been met by the agent feature.

View agent errors

Agent Client Collector (ACC) errors are visible in logs related to the agent and the ServiceNow instance. This feature provides improved visibility of agent errors, enabling faster error resolution.

Before you begin

- Ensure that the Error Framework (*com.glide.error_framework*) plugin is active.
- Ensure that the system property *sn_agent.use_glide_error_framework* is set to **true** (**All** > **System properties** > **All Properties**).
- Ensure that ACC Admin Workspace (*sn_acc_wrksp*) store application is active.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **Workspaces** > **ITOM Infra Services Workspace**.
2. Select the ACC agents icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Errors** to view the errors existing on ACC agents.

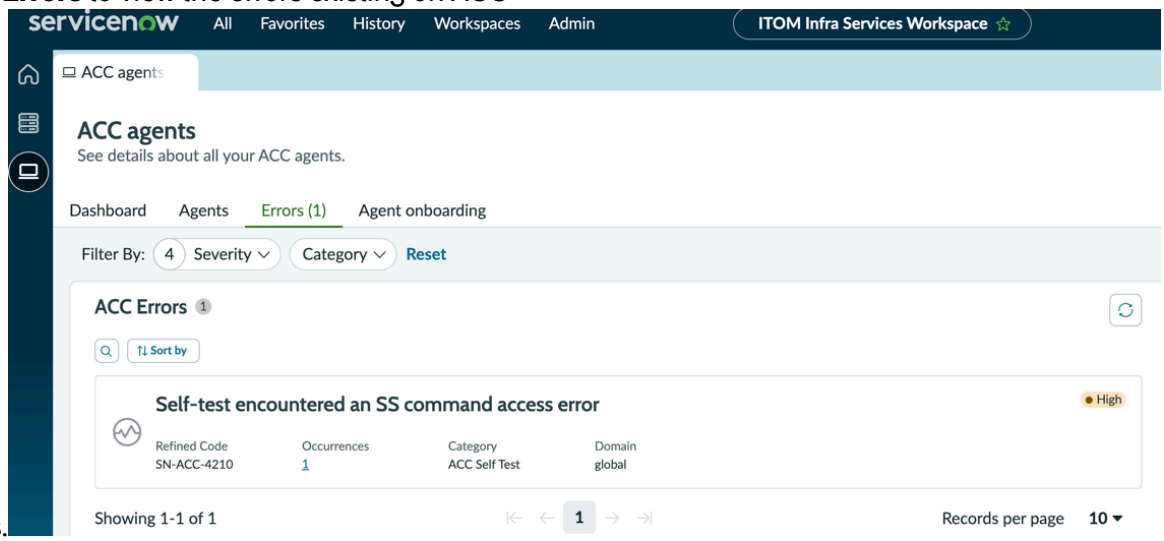

4. Filter the displayed errors by selecting severity and category values in the **Filter By** options. If you don't configure a filter, all agent errors appear.

5. Select an error to view additional information, including the error's root cause and remediation instructions.

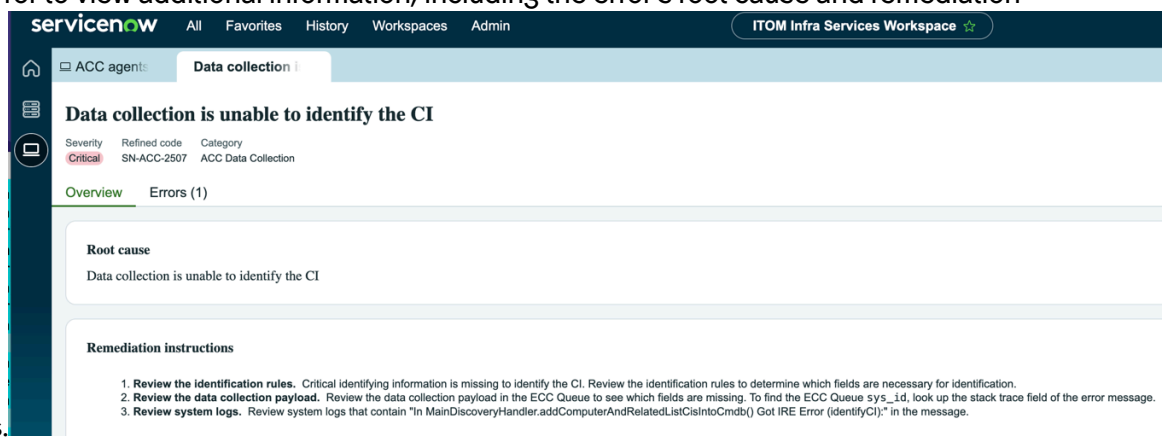

   a. Select the **Errors** subtab to view detailed information about the error.

   b. Select the information ⓘ icon to open the **Error details** panel.

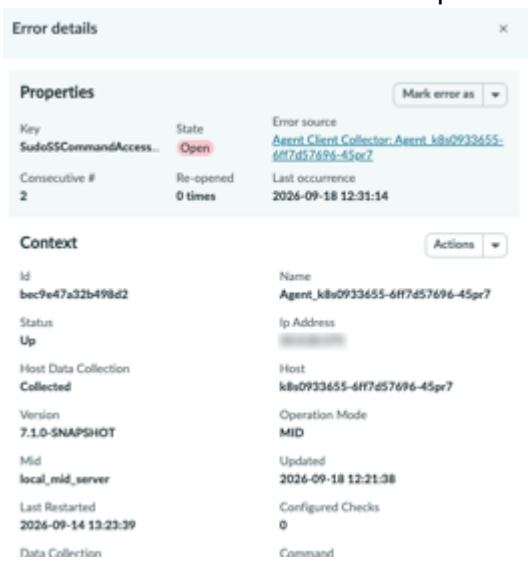

6. **Optional:** To view errors for a specific agent:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select **All** > **Agent Client Collector** > **Agents**.
The **Agent Client Collectors** table appears.

b. Select an agent.

c. Select the **Automation Error Messages** tab at the bottom of the page to view errors relating to the agent.

Using Agent Client Collector logs

Use Agent Client Collector logs to track events in your system. Log levels indicate the severity of each event.

Configure Agent Client Collector log levels

You configure the log levels that you want the Agent Client Collector logs to contain. All events that match or exceed the specified log level display in the Agent Client Collector logs.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Specify the desired log level in the `acc.yml` configuration file.
   Alternatively, you can specify the log level as an argument in the `--loglevel` command-line flag, as follows:

```
acc start --log level debug
```

2. Restart the agent using the service manager:
   - **Linux**:

```
sudo service acc restart
```

   - **Windows**:

```
sc.exe stop AgentClientCollector
sc.exe start AgentClientCollector
```

   - **macOS**:

```
sudo launchctl load
 -w /Library/LaunchDaemons/com.sn.acc.plist
```

Note: You must restart the service after changing the log level.

View the Agent Client Collector logs

You can view the activity logs for an Agent Client Collector.

About this task

Agent Client Collector logs are accessible from:

- **Linux**: `/var/log/servicenow/agent-client-collector/acc.log`
- **Windows**: The location specified by the `log-file` configuration flag. The default location is: `C:/ProgramData/ServiceNow/agent-client-collector/log/acc.log`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If the C : /ProgramData directory is hidden (such as on older Windows versions), change the Explorer filter to show hidden elements.

- **macOS**: The default location is: /Library/Application\ Support/servicenow/agent-client-collector

Before you begin
Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
   The **Agent Client Collectors** page opens.

Agent Client Collectors page

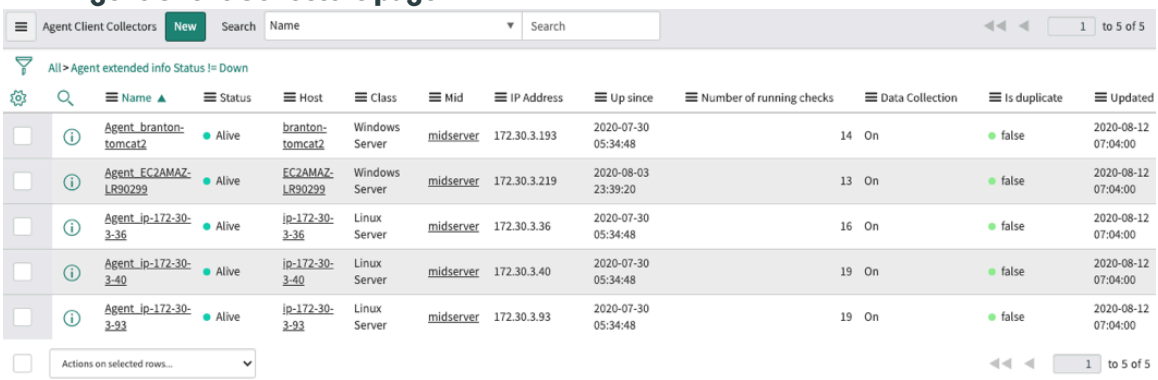

   The displayed columns are those selected when clicking the Settings icon .

   The **Host** column indicates the host name of the agent.

   The **Class** column indicates the agent's CI type.

2. Select an agent.
   The **Agent Client Collector** page opens.
3. In the **Related Links** section, click **Grab Agent Log**.
   The **Queues** page opens for the selected Agent.
4. Refresh the page and click the relevant entry with the value **Queue = input**.
   The **Payload** section on the **Queue** form displays the log for the Agent Client Collector.

Related topics

Enable log monitoring in a Linux environment

Enable log monitoring in a Windows environment

Set the agent log level

Configure the agent log level directly from your ServiceNow® instance, without needing to access the acc.yml configuration file.

About this task
Configuring the agent log level from the instance changes the log level temporarily, and doesn't update the log level in the acc.yml configuration file. When restarting the agent, the log level reverts to its status as set in the acc.yml configuration file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: agent_client_collector_admin

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select an agent.
3. In the related links, select **Set Log Level**.
4. In the Set Log Level dialog box, select the log level to invoke for the agent, and select **OK**.
5. Reset the log level to its originally configured level by doing one of the following actions:
   - In the related links, select **Reset Log Level**.
   - Restart the agent.

Enable metrics collection and evaluation

To enable a ServiceNow instance to collect and evaluate metrics, you must create a distributed MID Server cluster, associate MID Servers with the cluster, and enable Metric Intelligence for your MID Server.

Before you begin
Ensure that the following plugins are installed on your instance:

- Event Management (com.glideapp.itom.snac)
- Metric Intelligence - com.snc.sa.metric plugin
- MetricBase

To verify that the MetricBase plugin has been installed correctly:

1. Enter `stats.do` in the instance navigation pane.
2. Search for the **MetricBase Statistics** section and verify that the following appears:
   - Status: ONLINE
   - Name: clotho

Role required: agent_client_collector_admin

Procedure
1. Create a distributed MID Server cluster in your instance:

   a. Navigate to **MID Server** > **Clusters**.
   The MID Server Clusters page opens, displaying the list of configured MID Server clusters.

   b. Click **New**.
   The **MID Server Cluster - New record** page opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. In the **Name** field, assign a descriptive name for the cluster.

d. In the **Type** field, select **Distributed**.

e. Click **Submit**.
The distributed MID Server appears on the MID Server Clusters page.

2. Assign MID Servers to the distributed MID Server cluster:

a. On the MID Server Clusters page, select the distributed MID Server cluster.
The following page appears:

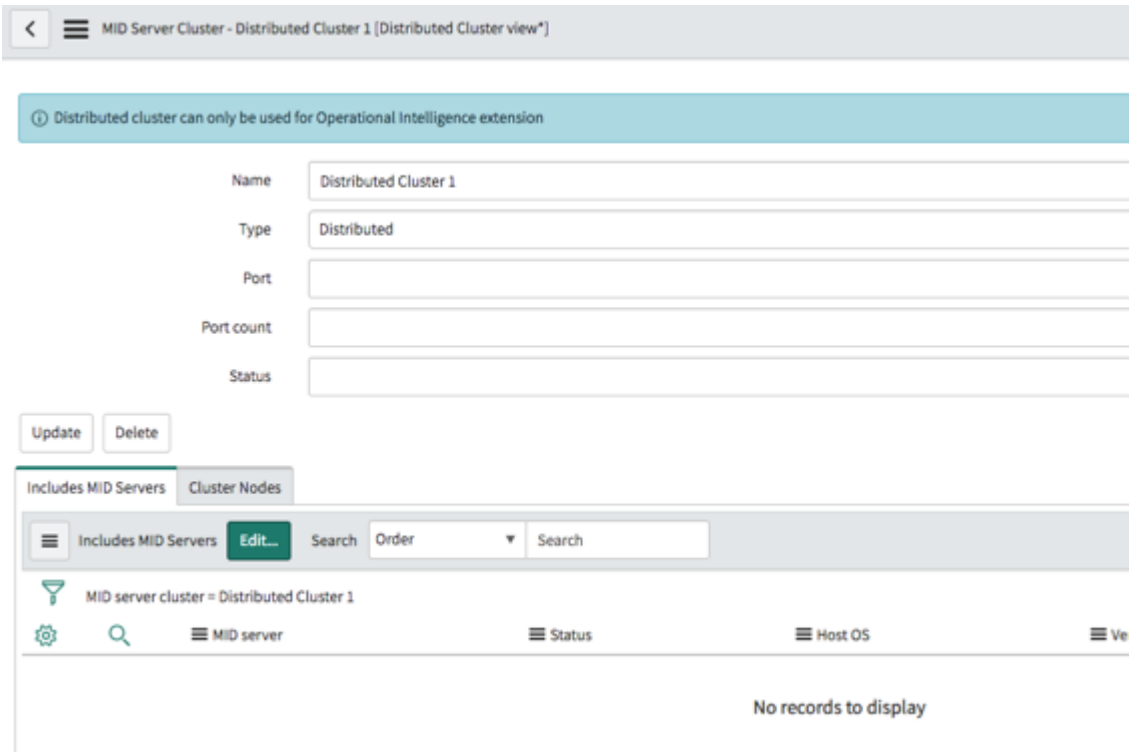

b. On the **Includes MID Servers** tab, click **Edit**.
The **Edit Members** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

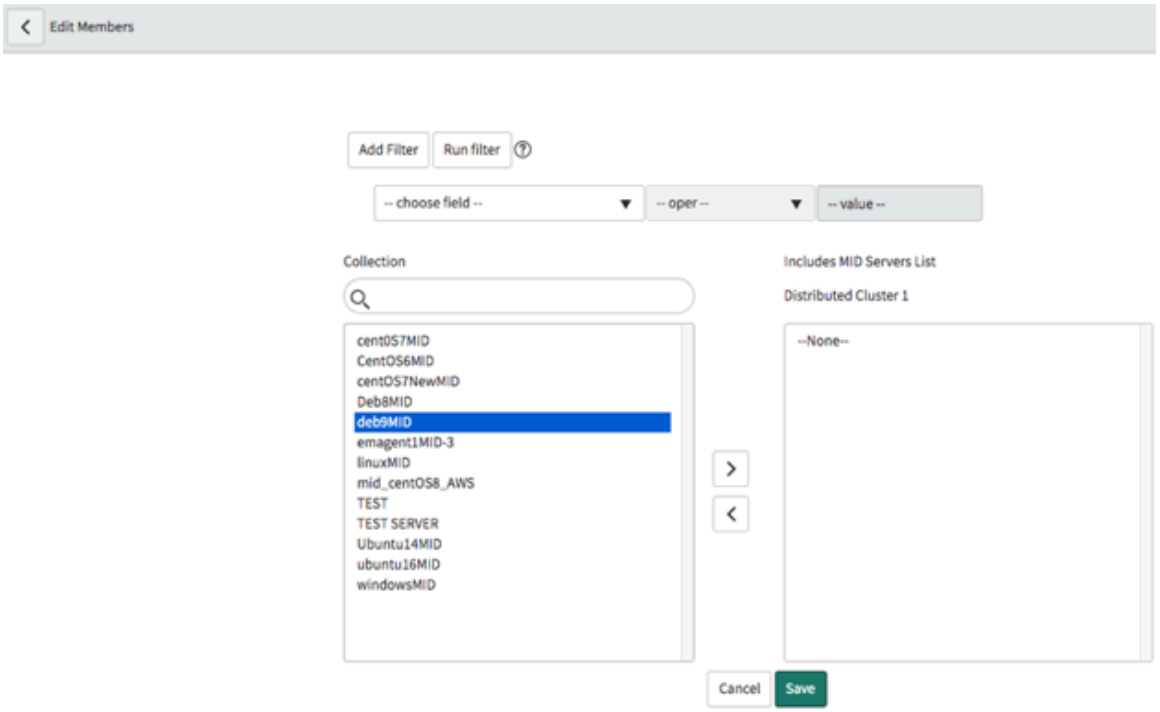

c. Select the relevant MID Servers from the left cell and click the right arrow button to move them to the right cell.
In a staging environment, it is acceptable to have only one MID Server in the cluster.

d. Click **Save**.
The **Includes MID Servers** tab displays the MID Servers that are part of the cluster.

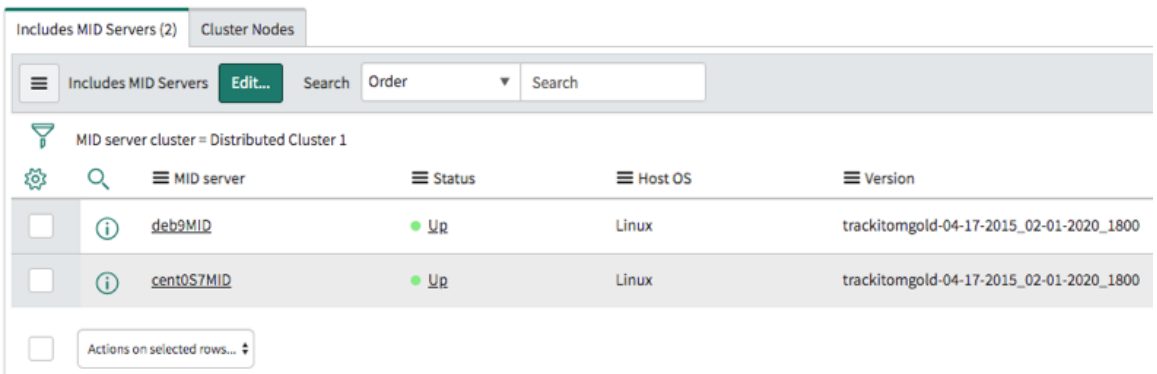

3. Enable Metric Intelligence for your MID Server:

a. Navigate to **MID Server** > **Extensions** > **Metric Intelligence**.
The **Metric Intelligence Metrics Contexts** page appears.

b. Click **New**.
The **Metric Intelligence Contexts - New record** page appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

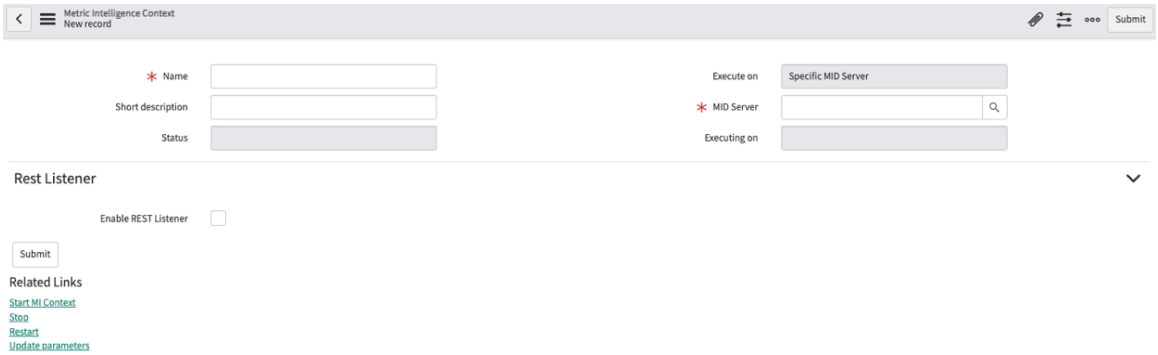

c. In the **Name** field, enter a descriptive name for the Metric Intelligence metrics extension.

d. In the **MID Server** field, select the MID Server to which you want to associate the Metric Intelligence metrics extension.

e. Select the **ACC Dedicated MID** check box to assign the specified MID Server as the only MID Server in the distributed cluster.
Selecting this option helps conserve system resources and memory.

f. In the **Related Links** section at the bottom of the page, click **Start**.
A message appears, indicating the status of the extension.

g. Click **Submit**.

h. Click **Start**.

Related topics

Limit metrics collection and evaluation

Limit metrics collection and evaluation

You can limit the metrics you send from the MID Server to the instance, either by de-activating a specific CI or an entire CI type.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. To limit metrics collection and evaluation for specific CIs:

a. Navigate to **Event Management** > **Anomaly Detection** > **Metric to CI**.
The **Metric to CI Mappings** page appears.

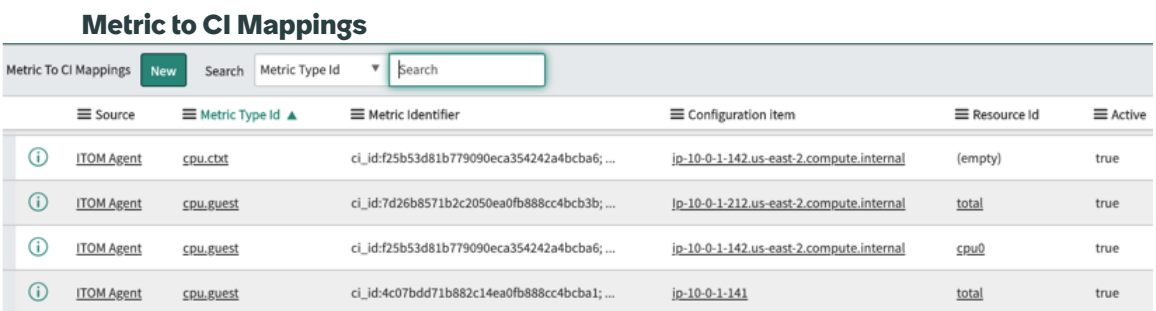

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The Metric to CI Mapping is identified by both its metric identifier and its configuration item.

b. To prevent sending metrics for a specific metric and CI, set the value of the **Active** column to **false**.

2. To limit metrics collection and evaluation for an entire CI type:

a. Navigate to **Event Management** > **Anomaly Detection** > **Metric Types**.
The **Monitoring System Metric Types** page appears.

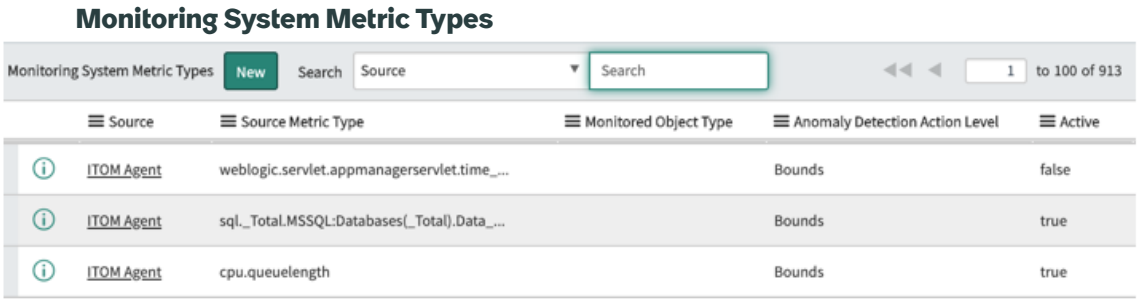

b. To prevent sending metrics for a metric type, set the value of the **Active** column to **false**.

Software filter

The Software filter lets you define rules to automatically exclude irrelevant entries from your Software Asset Management (SAM) inventory. At the same time, it keeps a complete, auditable record of everything filtered so that nothing disappears silently.

Discovery and the Agent Client Collector regularly report every application installed on a device. However, not everything found is meaningful licensing or asset-management software. For example, Operating-system shortcuts, Progressive Web App (PWA) stubs, Citrix Receiver placeholder entries, browser-extension shortcuts, and similar artifacts are frequently reported as installed software. However, they aren't applications your organization is licensing, tracking, or reconciling.

Software filter - Benefits

- Cleaner license reconciliation. Irrelevant entries no longer inflate install counts or trigger false compliance findings.
- Less manual clean-up. You no longer must search for and deactivate the same irrelevant software records after every discovery cycle.
- Full transparency. Every exclusion is logged, so you can see what was filtered, when, and by which rule.

How filtering works

Agent scans and background discovery runs produce software installs on devices in your network. Each item is checked against the set of currently active filter rules before it is written to your SAM software inventory.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If an item matches an active rule, it is excluded from the SAM inventory and a record of the exclusion is written to the review table instead.
- If an item matches no active rule, it flows through to your SAM inventory as normal. Nothing changes for software that is relevant to you.
- If the same software is reported again on a later scan, its existing review record is refreshed, not duplicated. The review table reflects the latest confirmation date rather than growing indefinitely.

A single malformed rule (for example, an invalid pattern) is skipped and logged rather than blocking the rest of the check. One bad rule doesn't accidentally exclude everything.

For details on the tables that determine which entries are labeled as irrelevant, see Software filter tables.

Modifying rules

If you deactivate or edit a custom rule, the behavior of excluded software changes on the next discovery scan:

- Software that no longer matches any active rule stops being excluded and flows into your SAM inventory as normal.
- Entries for filtered software remain in the Software Install Filter Staging table as a historical record, but are no longer refreshed with new confirmation dates.
- If you reactivate the rule or change it to match the software again, exclusion resumes on the next scan and the staging table entry is refreshed.

Create a custom filter rule

Create a custom filter rule in the Software Install Custom Filter table (samp_sw_install_custom_filter). Rules determine the criteria by which software is excluded from your Software Asset Management (SAM) workspace.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to the Software Install Custom Filter table **All** > **Software Install Custom Filter**.
2. Select **New** to create a custom filter rule.
3. Complete the fields based on the software you want to exclude.
   For details on the available fields, see Software filter fields.
4. Select **Submit**.

Result

The custom filter rule is created and takes effect on the next discovery or agent scan that reports matching software.

Example: Exclude a custom internal shortcut

Your organization has an internal tool that installs a desktop shortcut called *`Company Portal Shortcut`* on every managed Windows device. To ensure that it doesn't clutter your SAM inventory, configure the rule as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Name**: Exclude Company Portal Shortcut
- **Software Name**: Company Portal Shortcut, Condition: Exact match
- **OS Filter**: Windows
- **Discovery Source**: (leave blank)
- **Active**: On

From the next scan onward, any Windows device reporting *Company Portal Shortcut* will have that entry excluded from the SAM inventory and logged in the review table instead.

Review filtered software

Review the entries in the Software Install Filter Staging table (samp_sw_install_filter_staging) to confirm your rules are filtering only the software you expect.

Before you begin
Role required: agent_client_collector_admin

About this task
The Software Install Filter Staging table automatically records every software item that has been excluded by an active filter rule.

When the same software is reported on a later scan, its staging table entry is automatically refreshed with the new confirmation date. Entries that have not been updated indicate software that has not been reported recently by your devices.

Procedure

1. Navigate to the Software Install Filter Staging table (**All** > **Software Install Filter Staging**).
2. Review the records to verify the software that has been excluded and by which rule.
   Each entry contains the following information:
   - The software's display name, publisher, version, and discovery source.
   - The custom rule that excluded it (appears blank if it was excluded by a rule from the base-system).
   - The date it was last reconfirmed as excluded.
3. Confirm that your rules are excluding only the software you intended.
   If you find unexpected exclusions, check your active rules and adjust them as needed.

Configure the Google Cloud Platform (GCP) configuration data file

Configure the Google Cloud Platform (GCP) configuration data file to monitor metrics in your GCP environment.

Before you begin
Role required: agent_client_collector_admin

About this task
The GCP configuration data file is a .json file which contains all of the metrics collected from your GCP environment. Metrics that you add to the *acc_gcp_metrics_list.json* file are monitored by the check to which the file is attached.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Configuration** > **Configuration Files**.
2. Locate the *acc_gcp_metrics_list.json* configuration file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Download and add the GCP metrics to the configuration file that you want to be monitored. For details on the available GCP metrics, see Google Cloud Platform (GCP) metrics.
4. Select the Manage Attachments icon ( ) to attach the updated configuration file.

Configure the action level for anomaly detection

Create a configuration setting rule that refines the level of anomaly detection processing and analysis that is applied to specific CIs and metrics. Set a processing level that reflects the importance of metrics at different stages of implementation, to reduce data load if needed.

Anomaly detection consists of multiple levels of processing. At the lowest level, metric data is gathered from data sources. At the highest level, statistical models are created, anomaly scores are calculated, and anomaly alerts and IT alerts are created. You can apply the full anomaly detection process or a reduced level of processing, for specific data series.

Use the *anomaly_detection_action_level* configuration setting in a configuration setting rule to set the level of processing for specific CIs and metrics. For example, you can initially set only a small percentage of CIs and metrics to generate anomaly alerts, and later increase that percentage.

Action levels

The *anomaly_detection_action_level* configuration setting provides action levels starting at 'Metrics only' for basic anomaly processing, and ending at 'IT Alerts' for full processing and analysis.

Action levels from lowest to highest:

Metrics Only

Only gather metrics from the data source without any further processing.

Bounds

'Metrics Only' processing level and in addition creates statistical model and show bounds in Insights Explorer.

Anomaly Scores

'Bounds' processing level and in addition calculates and display anomaly scores.

Anomaly Alerts

'Anomaly Scores' processing level and in addition creates anomaly alerts as applicable.

IT Alerts

'Anomaly Alerts' processing level and in addition creates IT alerts which are based on anomaly alerts, as applicable.

Related topics

Create a configuration settings rule

Create Azure policies for metric collection

Create your own policies to enable Azure to collect the metrics for the resources that have policies that don't come with the base system. The Azure Metrics Collector check brings the metrics to the agent through the Azure Metrics Batch API.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Ensure that the configuration Items (CIs) that you want to collect the metrics for are included in the CMDB. CIs are typically populated during Discovery.
- Ensure that the *sn_agent.multi_ci_dev_mode* property is set to **true** on the **System Properties** page (**All** > **System Properties** > **All Properties**).

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Policies**.
2. Select **New**.
3. Select the application scope icon and verify that you are in the Agent Client Collector Monitoring scope.
4. In the **Name** field, enter a name for the policy.
5. Configure a filter that determines which CIs are to be monitored by the policy by selecting the **Monitored CIs** tab and the **Monitored CI type by filter** check box, as described in the following table.

Monitored CIs filter

Option	Action
Monitored CI type	Select the CI type to be monitored. The CI type must match the resource that you entered in the .json script file that contains the metrics to be monitored.
<Filter fields>	Configure the conditions that must be met for the resources in the specified Monitored CI type to be monitored. Ensure that you set the following conditions: ○ **Object ID** = <name of the resource entered in the .json script file that contains the metrics to be monitored> ○ **State** = Available The Available state indicates that resources must be available to be monitored.
Preview	Select to view the number of resources that match the specified condition. The number of matching results appears next to the **Preview** button. Select the result to view all CI records that match the condition.

6. Add the Azure Metrics Collector check to the policy.
   a. On the **Checks** tab, locate the **Azure Metrics Collector** check in the **Available** cell.
   b. Select the right arrow button to move the check to the **Selected** cell.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. Configure a proxy server to access the Azure metrics.
   a. On the **Proxy Settings** tab, select **Single proxy agent**.
   b. In the **Proxy agent** field, select the proxy agent to execute the policy and retrieve the Azure metrics.
8. Configure the credentials for the policy.
   a. Select the **Credentials** tab.
   b. In the **Credential name** field, select the **azure_metrics_collector**.
9. Select **Save**.
10. Add secure parameters to the check instance.
    a. In the policy's **Check Instances** section, select the **Azure Metrics Collector** check instance.
    b. On the **Check Secure Parameters** tab, select **New**.
    c. Add the secure check parameters that are described in the following table.

Secure check parameters

Name	Order value
cred_tenant_id	1
cred_client_id	2
cred_secret_key	3

Ensure that the **Active** check box is selected for each parameter.

    d. Select **Update**.

11. Add the `.json` configuration file to the check instance.
The file sends the list of Azure resources to the agent.
    a. Select the Azure Metrics Collector check instance.
    b. Select the **Check Parameters** subtab.
    c. Select **New**.
    d. On the form, fill in the fields.

Check Parameter form

Field	Description
Name	Descriptive name for the `.json` file to be added to the instance's configuration files (**Agent Client Collector** > **Configuration Files**).
Value	String that indicates the name of the `.json` file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	For example: `acc_azure_vm_metrics_config.json`
Flag	Value that you enter (**-c**) to indicate that the check retrieves the name of the configuration file. The configuration file contains a namespace and a list of metrics. The value must indicate the name of the file that was uploaded to the list of configuration files (**Agent Client Collector** > **Configuration Files**), in `.json` format. For example: `acc_azure_vm_metrics_config.json`.

Ensure that the **Active**, **Mandatory**, and **Value required** check boxes are selected.

e. Select **Submit**.

12. Create a `.json` file that indicates the metrics that you want your policy to monitor.

The following example displays a `.json` file for retrieving load-balancer metrics.

acc_azure_static_ms_lb_config.json

```
{
  "namespace": "Microsoft.Network/loadBalancers",
  "metrics": [
        "AllocatedSnatPorts",
        "ByteCount",
        "DipAvailability",
        "PacketCount",
        "SnatConnectionCount",
        "SYNCount",
        "UsedSnatPorts",
        "VipAvailability"
  ],
  "time_grain": "PT1M"
}
```

Add the `.json` file to the list of configuration files at **Agent Client Collector** > **Configuration Files**.

13. Create a .json file to indicate the resources that you want your policy to monitor.

a. On the **Proxy Settings** tab, select **Multi-CI Mode**.

b. Copy a script from one of the Azure cloud policies that comes with the base system.

- Cloud policies that come with the base system are described in Create Azure policies for metric collection.
- The script is located in the **Config file creation script** field of the policy's **Proxy Settings** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Paste the script into the **Config file creation script** field of your policy.

d. At the end of the script, configure the *resource_id*, *region*, and *sys_id* parameters with the values of the CI that you want the policy to monitor.

14. Select **Save** and then publish the policy by selecting **Publish**.

Result
The configured policy retrieves the resources and their metrics from the specified Azure CI and sends them to the agent.

Block event creation for non-existent entities

Invoke the *skip_events_for_nonexistent* parameter to block event creation when specific process, service, or log does not exist. Invoke this parameter if you do not want to be notified about these non-existent entities.

Before you begin
Role required: agent_client_collector_admin

About this task
A process, service, or log does not exist when the following checks do not contain valid values for the indicated parameters:

Check parameters

Check	Parameter
os.linux.check-process	*file_pid*, *daemon_name*
util.check-windows-service	*service*
os.windows-check-log	*file*
util.checks-logs	*file*
util.checks-logs-sudo	*file*

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Configuration** > **Check Definitions**.
2. Locate the check definition for which you want to block events.
3. On the **Check Parameter Definitions** tab, select *skip_events_for_nonexistent*.
4. Select the **Active** check box to enable the parameter and prevent alert creation when the process, service, or log isn't located for the specified check.
   If you want to receive notification about the missing process or service, leave the *skip_events_for_nonexistent* parameter with its default value of **false**.

Enable viewing HAProxy metrics

Configure the HAProxy metrics stats page to enable monitoring your system devices by viewing HAProxy metrics.

Before you begin
Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

HAProxy provides load balancing and reverse proxy capabilities for both HTTP and TCP requests. The HAProxy checks remove backend servers that fail several health checks in a row, and transfer incoming traffic to backend servers as part of load balancing.

Procedure

1. Access the HAProxy configuration file, located at: `/etc/haproxy/haproxy.cfg`.
2. Add the following to the bottom of the file:

```
listen frontend *:8080       #Haproxy
Monitoring run on port 8080
mode http
option forwardfor
option httpclose
stats enable
stats show-legends
stats refresh 5s
stats uri /stats        #URL for HAProxy monitoring
stats realm Haproxy\Statistics.   #Title for pop up window
stats auth admin:admin     #User and Password for login to the
 monitoring dashboard
stats admin if TRUE
default_backend backend       #This is optional for monitoring
 backend
```

3. Run the following command to restart the HAProxy:
   `sudo service haproxy restart`
4. View the HAProxy metrics on your system by viewing the HAProxy metrics stats page, accessible from the following URL:
   `http://<HAProxyServer>:8080/stats`

Related topics

HAProxy default checks and policies

Enable Apache Kafka health monitoring

Set configurations in the Apache HTTP server application to enable the Agent Client Collector to perform Apache Kafka HTTP server monitoring.

Before you begin

- Enable Java Management Extensions (JMX) for Kafka metric collection by running the following command:

```
Environment="JMX_PORT=9999"
```

- Role required: agent_client_collector_admin

Procedure

1. Enable collecting Apache Kafka performance metrics using Java Management Extensions (JMX) by running the following command:

```
JMX_PORT=9999 bin/kafka-server-start.sh config/server.properties
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Configure the admin server port by modifying the port number of the admin.serverPort property in the zookeeper.properties file.
   The default port is 8080.

Enable running of Varnish checks

Use application fingerprints to discover the Varnish CI, to enable running Varnish checks.

Before you begin
Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Home**.
2. Select the **Application Fingerprints** tile.
   The **Application Fingerprints** dashboard appears.
3. Select the **All Suggested Applications** tile.
4. On the Application Suggestions page, select the **varnishstat** entry in the **Suggested Group Name** column.
   The varnishstat details appear on the Application Suggestion page.
5. Select Discover Application.
   Discovery locates an application CI for the varnishstat entry. The CI appears on the Application Fingerprints dashboard in the **Created Applications** section.

Created applications section of Application Fingerprints dashboard

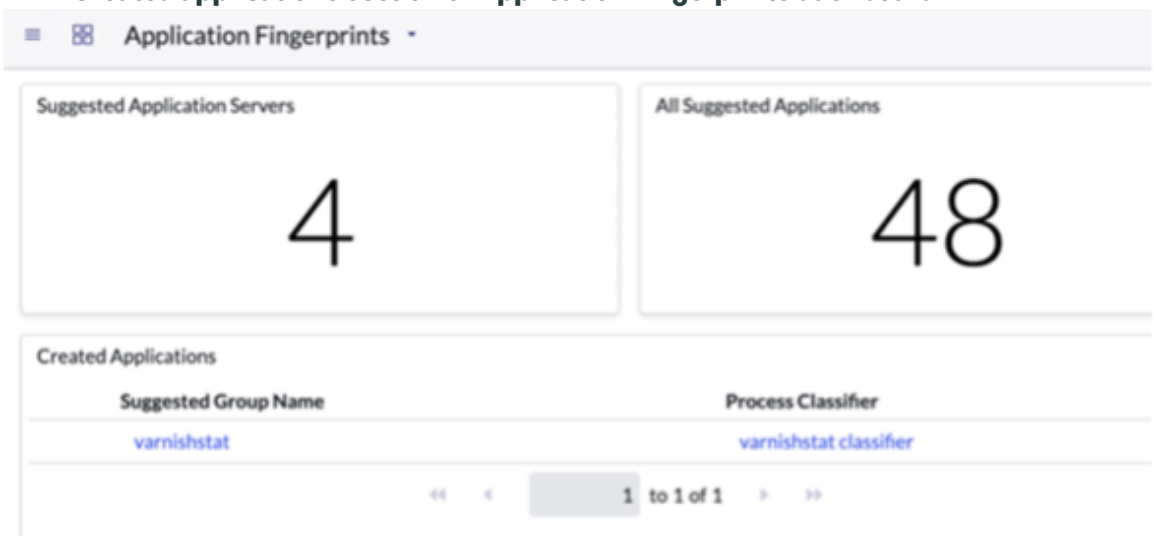

6. Run basic Discovery on the host machine.

Result
The varnishstat CI is discovered and a record is created in the u_cmdb_ci_varnishtat table.

The Varnish policies and checks are available to run.

Detecting portable applications using push-based discovery

Portable applications are those applications that don't need to be installed on the target system. Starting in ACC-VC version 2.3.0, push-based Discovery can detect portable applications, such as Firefox, VLC, Notepad++ etc, for Windows only.

Note: Only portable applications listed in https://portableapps.com/ are supported.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When the portable application is executed, it creates a running process. The application creates a path similar to the following for the configuration file 'appinfo.ini':

```
-\ <--- Directory with FirefoxPortable.exe
 +\App
 +\firefox
 +\Data
 +\profile
 +\settings
 +\plugins\ (optional)
```

Note: ServiceNow users must have Local account access to get the data in order to read appinfo.ini file.

Name, Version, and Vendor attributes are fetched from the Portable software and are found in the appinfo.ini files. appinfo.ini is the mandatory file to execute any portable application. If SAM is installed, the information is displayed in Software installation [cmdb_sam_sw_install] table. If SAM is not installed, the information shows in the Software instance [cmdb_software_instance] table.

Cloud Native Operations for Visibility

Cloud Native Operations for Visibility (CNO for Visibility) has been renamed Agent Client Collector for Kubernetes – Visibility. The term Cloud Native Operations for Visibility has been deprecated.

For the Kubernetes Visibility Agent documentation, see Kubernetes Visibility Agent.

Enabling Agent Client Collector data collection

You can perform actions which enhance data collection from the hosts on which Agent Client Collector is running.

Create a call API to send an osquery request

Create a background script to send an `osquery` request. The `osquery` request enables data collection from your host's operating system.

Before you begin

Role required: administrator

Procedure

1. Create the following background script:
   - `"MTLV9DE22016"` is the name of the server running an agent
   - `"c8bf9072539333003e76ddeeff7b12c5"` is the sys_id of the previously created check definition

```
cis = new GlideRecord("cmdb_ci_server");
    cis.addQuery("name", "MTLV9DE22016");
    cis.query();
    check = {
        "checkDefId": "c8bf9072539333003e76ddeeff7b12c5",
        "params": {
            "query": "select pid,cmdLine from processes limit
 4"
        }
    };
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
    var agentNow = new sn_agent.AgentNowHandler();
    gs.info('request_id: ' +
agentNow.runCheckForCis(cis,check,'1',60000));
```

2. Run the script to print the request ID:
request_id: `eefbff7ab75733003e766068de11a981`
3. Check the system logs for messages coming from the instance script.

```
result 0 requestId: 06c6c887b79733003e766068de11a912
 agent_name: MTLV9DE22016 ci_id:
00726c36539333003e76ddeeff7b12db status: 0 output: [
{"cmdline":"","pid":"0"},
{"cmdline":"","pid":"1"},
{"cmdline":"","pid":"45"},
{"cmdline":"","pid":"46"}
]
```

View host data collection for an agent

Run data collection on an agent's host to discover the components that are running on the host.

Before you begin

Role required: agent_client_collector_admin

About this task

The Refresh Host Data for Agents scheduled job collects host data and runs every hour on the following:

- Hosts which have not been discovered in the past 12 hours.
- Host for which Discovery failed.

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
In the **Host data collection** column, view the status of data collection on the agent's host. The possible statuses are:

Host data collection statuses

Status	Description
<empty>	Indicates that the agent has been added within the past minute and the Refresh and Publish Monitoring Policies scheduled job has not yet run to initiate data collection. The job runs every minute.
Collected (green)	Host data collection is complete.
Collecting (yellow)	Host data collection is in progress.
Collection failed (red)	Host data collection failed.

2. **Optional:** Locate the **Host data last collected** column to view when data collection was last performed on the host.

Generate an Agent Client Collector allowlist

Specify the checks to be included in the list of checks that are enabled to run on the agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: agent_client_collector_admin

About this task
The default allowlist includes all validated checks, and is included in the base system policies. If the allowlist cannot be read correctly, no checks run on the agent.

Procedure
1. Navigate to **All** > **Agent Client Collector** > **Check Definitions**.
   The **Check Definitions** page appears.
2. In the **Related Links** section at the bottom of the page, click **Generate allow-list content**.

   Note: Alternatively, you can click a check definition or check instance record and click **Generate allow-list content** to generate a complete allowlist.

   The **Agent Client Collector Command Allow-list Generator** page appears, with the list of checks scanned from the **Command** field of the check instance records.

   The values appear with the following attributes:

Command attributes

Attribute	Description
exec	The executing command.
args	Possible argument strings.
skip_arguments	Boolean. Determines whether to verify that the argument was added to the allowlist. ○ **true**: Skips argument validation and checks only for illegal characters, specifically |; \<>&'$() ○ **false**: Executes argument validation to ensure that the argument is included in the allowlist. If it is not, the command cannot execute.
use_regex	Boolean. Determines whether to evaluate the *arg* strings as a regular expression. ○ **true**: Evaluates each entry in the args array as a regular expression ○ **false**: Evaluates each entry in the args array as it is. **Note:** Ensure that the regex is useable in valid JSON format and is golang compatible.
allow_shell	Boolean. Allow for the given entry to be executed in shell mode and for check definitions to default to *execv*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Attribute	Description
	○ **true**: This allowlist entry permits execution in shell mode. ○ **false**: This allowlist entry doesn't permit execution in shell mode.

3. Copy the entries you want to include in the allowlist and add them to the agent's `check-allow-list.json` file, either manually or using an automation tool.
   The location of this file is indicated in the *`allow-list`* parameter of the agent's `acc.yml` file. The default location is:

   ○ Windows: `C:\ProgramData\ServiceNow\Agent Client Collector\check-allow-list.json`
   ○ Linux: `/etc/servicenow/agent-client-collector/check-allow-list.json`

Run host data collection for an agent

Collect data on the host being monitored by an agent. Manually collecting host data ensures that the host's files are up to date.

Before you begin
Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Underneath the check box icon ☐, hover next to an agent and select the check box that appears.
   Repeat this step for all agents for which you want to collect data.
3. Select the arrow next to the **Actions on selected rows...** drop-down and select **Collect host data**.

Result
Data is collected for the specified agents. The collection status is indicated in the page's **Host data collection** column.

Run host data collection on demand

You can perform host data collection on demand, if you need to perform data collection before the job is scheduled to run.

Before you begin
Ensure that you have sufficient permissions to create CIs.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
   The **Agent Client Collectors** page appears.
2. Select an agent with **Status = Alive**.
3. On the agent form, click the **Collect Host Data** button.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Pause Agent Client Collector data collection

If CPU consumption on your server is getting too high, you can manually turn off the Agent Client Collector data collection to pause all checks performed by the agent, except for the agent's keep alive messages.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
   The **Agent Client Collectors** page appears, displaying the list of configured agents.
2. Select the checkbox next to the relevant agents.
3. In the **Actions on selected rows...** dropdown, select **Pause data collection** to pause data collection on the agents (**Data Collection = Off (manual)**).

   Select **Resume data collection** to resume data collection on the agents.

   Alternatively, you can set the status of an agent by selecting the agent and in the **Related Links** section, select the **Pause data collection** link to pause data collection on the agent. To resume data collection on the agent, select the **Resume data collection** link.

Clear Agent Client Collector plugins

Clear the plugin and configuration file cache folder and control plugin downloads. Clearing the cache folder removes unwanted plugins that have synced to the agent.

Before you begin

Ensure that the following are installed on your system:

- Agent Client Collector Framework (ACC-F), starting with version 3.5.5
- MID Server
- CMDB
- Agent Client Collector agent, starting with version 3.5.5
- ACC Global Utilities

Ensure that the **disable_assets** system property in the `acc.yml` configuration file is set to **false**.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **Agent Client Collector** > **Agents**.
2. Select an agent.
3. In the **Related Links** section, select **Clear ACC Plugins**.
4. In the confirmation dialog box, select **Clear ACC Plugins**.
5. **Optional:** Select **Collect host data** to re-download your plugins from the instance.
   Plugins download automatically when used by their respective checks/policies.

Create a golden image for a virtual machine

Create a golden image to run on a VIrtual Desktop Infrastructure (VDI) machine. The golden image enables installing a virtual agent on VDI virtual machine, which generates data quicker than standard agents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: agent_client_collector_admin

Procedure

Navigate to **All** > > .

Agent Client Collector Monitoring

Agent Client Collector Monitoring enables you to monitor your service availability, examine the health and performance of your environment, and ensure that your infrastructure and its applications are running properly.

Get started

Note:

If checks are not running on the agent's devices, your agent may be in CPU protection mode. CPU protection mode is activated automatically when a device's CPU consumption is too high for one of two reasons: either the user is running too many checks, or the machine is under-resourced. When it happens, the agent's Data Collection status turns to Off. ACC-F generates an event and records it in a log. To manually resume data collection by the ACC agent, see Pause Agent Client Collector data collection.

The user needs to monitor the activity of the ACC to resolve the issue through manual intervention and then re-enable ACC collection. To setup the ACC monitoring, see Enable Agent Client Collector monitoring on the MID Server.

Explore	Configure	Metric Intelligence
Learn the key features and business value that the Agent Client Collector Monitoring application offers.	Configure the features and components to enable Agent Client Collector Monitoring.	Use Metric Intelligence to explore metric data and identify anomalies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Work with the Monitoring Technology Dashboard	Use proxy agents	Reference
Monitor server resources by platform.	Monitor the health and performance of your Configuration Items (CIs) when the agent is deployed outside of your host server.	Get additional details about checks and policies, and metric collection.

Troubleshoot and get help

- Ask or answer questions in the Agent Client Collector community forum.
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Agent Client Collector Monitoring

Agent Client Collector Monitoring is built on a Sensu framework which enables you to adopt and extend monitoring checks from the community. Agent Client Collector Monitoring collects data on your company's infrastructure and installed applications.

Agent Client Collector Monitoring overview

Checks and policies run on the agent's client to retrieve the relevant data, which is transformed into events or metrics, as appropriate. Events and metrics are sent from the agent to your ServiceNow instance through a Metric Intelligence MID Server and are stored in the relevant database. A single MID Server can support multiple agents (such as ACC-M and ACC-L). For details on configuring the MID Server for Metric Intelligence, see MID Server distributed cluster for Metric Intelligence.

Agent Client Collector Monitoring comes with the Event Management plugin and also requires installing the request Metricbase (Clotho DB), as described in How to request Metricbase (Clotho DB) configuration (KB0816088).

For details on the plugins installed with Agent Client Collector Monitoring, see Plugins or applications installed with ITOM AIOps.

Agent Client Collector Monitoring workflow

The following illustration describes the layout and data flow within the Agent Client Collector Monitoring application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Monitoring service health and availability with ACC-M

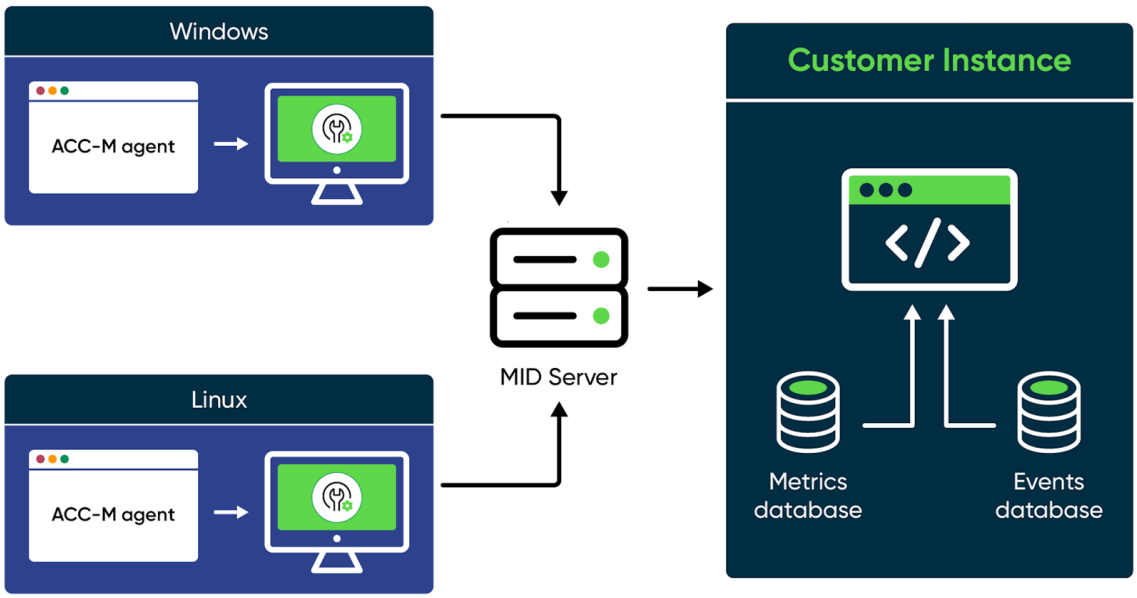

Agent Client Collector Monitoring benefits

Benefit	Feature	Users
Monitor logs which track events occurring in a Windows environment	Windows event log monitoring	Agent Client Collector administrator
Monitor operating systems and applications	Operating system and application monitoring using Agent Client Collector	Agent Client Collector administrator
Gather metrics on host data	Metric Intelligence	Agent Client Collector administrator

AI in Agent Client Collector Monitoring

AI provides generative AI features in Agent Client Collector Monitoring that help administrators diagnose agent errors and investigate agent behavior faster.

ServiceNow Otto for ITOM includes two AI features for Agent Client Collector Monitoring. Both features are accessed through the ServiceNow Otto panel and require ServiceNow Otto for IT Operations Management (ITOM) to be installed. For a summary of all AI features across ITOM AIOps, see AI in ITOM AIOps.

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

AI features in Agent Client Collector Monitoring

AI feature	Description	Use case	Resources
ACC diagnostic workflow	An agentic workflow invoked through the ServiceNow Otto panel that examines agent behavior, diagnoses issues on a specific ACC agent, and suggests remediations for	An administrator notices errors on an ACC agent and wants to investigate	Agent Client Collector (ACC)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI features in Agent Client Collector Monitoring (continued)

AI feature	Description	Use case	Resources
	specific error codes. Generates an error report based on activity logged by the Agent Client Collector Framework.	the cause and get remediation suggestions without manually reviewing logs.	diagnostic workflow
Error analysis skill	Uses generative AI to generate an analysis of errors within an ACC agent. Displays the analysis in the ServiceNow Otto panel and supports follow-up questions about a specific agent error or error code.	An administrator is reviewing an agent with recurring errors and wants an AI-generated summary to understand the error pattern before taking action.	View an error analysis by ServiceNow Otto in Agent Client Collector

Configuring Agent Client Collector Monitoring

Plan and configure your Agent Client Collector Monitoring implementation.

To configure Agent Client Collector Monitoring, follow these steps:

1. Configure a MID Server to work with Agent Client Collector Monitoring.
2. Choose and configure metrics to monitor.
3. Configure the action level for anomaly detection.
4. Create a configuration settings rule.
5. Enable metrics collection and evaluation.

Depending on the checks and metrics you want to monitor, perform the following actions to enable full check activation and metrics collection:

1. Enable viewing HAProxy metrics.
2. Enable running of Varnish checks.
3. Select the Azure policies to activate for metric collection.
4. Create Azure policies for metric collection.

Related topics

Agent Client Collector checks – credential and permission requirements

View an error analysis by ServiceNow Otto® in Agent Client Collector

View an error analysis created by generative AI. Error analyses enable asking questions on a specific agent's error or error code.

Before you begin

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select the **Agentic workflows** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select the **Agent Client Collector (ACC) Diagnostic** entry.
4. Locate the **Add AI agents that can perform these steps** widget and select the **Agent Client Collector (ACC) Diagnostic** entry.
5. Enable ServiceNow Otto® by sliding the **Allow AI specialists to access this AI agent** toggle switch to the right.

   ServiceNow Otto® retrieves information on errors within agents.

Agent Client Collector (ACC) diagnostic workflow

Enable and use AI agents to examine agent behavior through the Agent Client Collector (ACC) diagnostic workflow.

Before you begin

- Ensure that the ServiceNow Otto panel is enabled on your instance. For details, see Activate the ServiceNow Otto panel standard chat.
- Ensure that ServiceNow Otto for IT Operations Management (ITOM) is installed on your instance. For more information, see Install the ServiceNow Otto for IT Operations Management (ITOM) application.
- Enable the Agent Client Collector diagnostic workflow in AI Agent Studio.
  1. Select **All** > **AI Agent Studio** > **Create and manage**.
  2. On the **Agentic workflows** tab, select **Agent Client Collector (ACC) Diagnostic**.
  3. Select the **Select channels and status** option.
  4. In the **Engage via the** *`ServiceNow Otto`* panel, slide the **Display** toggle bar to the right (active).
- Role required: agent_client_collector_admin, agent_client_collector_user, or now_assist_panel_user

About this task

Invoke the Agent Client Collector (ACC) diagnostic workflow to:

- Question issues of an ACC agent while on an ACC form.
- Request suggestions to remediate a specific ACC error code.

Procedure

1. In a ServiceNow instance, select the ServiceNow Otto icon ( ).
2. In the ServiceNow Otto panel, enter prompts to invoke the ACC diagnostic workflow and perform remediation actions on the agent.
   Possible prompts include:
   - Diagnose issues with ACC agent <agent name>
   - Multiple errors on <agent name>, diagnose them
   - Provide suggestions to remediate ACC error code <error code number>

Result

An error report is generated, based on the activity logged by the Agent Client Collector Framework.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector Log Analytics

Agent Client Collector Log Analytics (ACC-L) enables you to stream log data from Linux and Windows hosts to a ServiceNow instance, using the Agent Client Collector.

As with Agent Client Collector Framework (ACC-F), the Agent Client Collector connects to a MID Server. A single MID Server may support multiple agents, while a single agent works with one MID Server at a time and switches to a different MID Server when necessary to provide failover protection.

Use the Agent Client Collector to stream log data to your ServiceNow® instance by means of a log shipper. The log data is mapped and structured in the ServiceNow® Health Log Analytics application. For more information, see the Health Log Analytics product documentation.

Agent Client Collector Log Analytics has a dependency on the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store. Streaming log data from Windows hosts is supported in Agent Client Collector Log Analytics Version 3.1.0 and later.

The Agent Client Collector Log Analytics (ACC-L) application, Version 3.2.0, August 2022, has a dependency on the Agent Client Collector Framework (ACC-F) application, Version 2.10.1, August 2022.

To enable Agent Client Collector Log Analytics, you enable log analytics on the Agent Client Collector, as described in Set up Agent Client Collector Log Analytics without using guided setup.

Dependencies

For a list of Agent Client Collector Log Analytics plugin dependencies, see Plugins or applications installed with ITOM AIOps.

Set up additional ACC data inputs

Data inputs for streaming log messages to your ServiceNow instance using the Agent Client Collector are created automatically as part of the Agent Client Collector setup. Set up additional ACC data inputs manually as needed using the Health Log Analytics data input setup flow.

Before you begin

- The Agent Client Collector Log Analytics application has a dependency on the following ServiceNow applications, available from the ServiceNow Store:
  - Health Log Analytics, Version 22.0.12 - December 2021 and later releases.
  - Agent Client Collector -Framework, Version 2.7.0 - December 2021 and later releases.
- Verify that a MID Server is installed and configured. For more information, see MID Server system requirements.
- The MID Server must be configured to work with the Agent Client Collector. For more information, see Enable Agent Client Collector monitoring on the MID Server.

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

Role required: evt_mgmt_admin

About this task

It is important to configure an ACC data input for every MID Server to which Agent Client Collectors can connect, even if no agent is currently connected to it. An agent works with one

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MID Server at a time, but is connected to a different MID Server when necessary to provide failover protection.

Important: As part of the standard Health Log Analytics data input setup, you can define JavaScript functions for mapping and preprocessing raw log data. The JavaScript function must be identical for all ACC data inputs, because agents will switch to a different ACC data input when the system connects them to a different MID Server to provide failover protection. The latest published JavaScript function is provided to all existing and future ACC data inputs and replaces the previous one. This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store.

Procedure

1. Set up additional ACC data inputs manually using the Health Log Analytics data input setup flow.
   For more information, see Configure ACC data inputs manually.
2. Identify and address log streaming issues to ensure that your data inputs are streaming data properly.
   For more information, see Ensure that the data input is streaming data.
3. Determine how Health Log Analytics handles raw log data that is streaming into your instance. By default, every incoming log line is auto-mapped to the correct tag. If properties aren't discovered automatically, map the data input sources manually by defining a JavaScript function. For more information, see Map the raw data.
4. **Optional:** Modify raw log data and drop or break up log messages before they are processed. Define a JavaScript function that modifies your raw log data before Health Log Analytics maps and structures it. For more information, see Edit your raw log data before processing.
5. **Optional:** Refine the source type structure to make sure that Health Log Analytics extracts and classifies all properties correctly.
   For more information, see Refine the source type structure in Health Log Analytics.

Configure ACC data inputs manually

Agent Client Collector setup automatically creates an ACC data input for streaming log messages to your ServiceNow instance. You can set up additional ACC data inputs manually as needed.

Before you begin

- The Agent Client Collector Log Analytics application has a dependency on the following ServiceNow applications, available from the ServiceNow Store:
  - Health Log Analytics, Version 22.0.12 - December 2021 and later releases.
  - Agent Client Collector -Framework, Version 2.7.0 - December 2021 and later releases.
- Verify that a MID Server is installed and configured. For more information, see MID Server system requirements.
- The MID Server must be configured to work with the Agent Client Collector. For more information, see Enable Agent Client Collector monitoring on the MID Server.

Important: Health Log Analytics does not support IPv6. To work with the application, configure the MID Server to IPv4.

Role required: evt_mgmt_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **ACC Log Analytics** > **ACC Data Inputs**.
2. Select **New**.
3. On the form, fill in the fields.
   For a description of the fields, see ACC data input configuration fields.
4. **Optional:** Select **Advanced** to set additional configuration fields.
5. Select **Save**.
   Health Log Analytics adds the data input record to the Data Inputs table.
6. Ensure that the data input is configured correctly by selecting **Test connection**.

   Health Log Analytics tries to connect the MID Server to the data repository.

   - If the connection was established, the **Test connection** button is turned off and the **Publish** button is enabled.
   - If the connection failed, the reason for the failure displays in the **Error message** field. This field displays only when a streaming error has occurred.

     Resolve the issue, select **Save** if you modified the configuration, and then select **Test connection** to test the connection again.

     Note: You can only publish the data input configuration when the connection is created successfully.

   Note: You can revert to the last published configuration by selecting **Revert Changes**. This option is available only when you're modifying a configuration that has been published previously.

7. Select **Publish** to publish the data input to the MID Server.

Result

The new data input is ready to stream logs.

Note: The data input record includes these related lists:

- Agents – Lists all active agents connected to the MID Server for this data input.
- Streaming Sources – Lists all sources for this data input and the MID Server that receives logs from it.

What to do next

Ensure that the data input is streaming data.

ACC data input configuration fields

Description of the fields on the ACC data input configuration form.

Basic configuration

Field	Description
MID	The MID Server to which the logs stream.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Basic configuration (continued)

Field	Description
	Note: • Only one ACC data input can be defined per MID Server. You can select only MID Servers with AgentClientCollector capability that support basic authentication. MID Servers that support mTLS are not listed. • The default maximum number of data inputs streaming logs to a single MID Server is 10. You can modify this number in the MID Server properties. However, when 10 non-ACC data inputs are defined for a MID Server, you can still define one ACC data input for it. This addition makes a total of 11 data inputs for that MID Server. • When you submit the form, this field becomes read-only. This field is required.
Port	The port on the MID Server. The port must be configured and active. It must not be occupied by another process. Make sure that your organization's security team opens the port before you assign it. **Note:** When you update the port, the system updates the Agent Client Collector with the new port configuration. Log streaming continues seamlessly without log loss after 1-3 minutes. This field is required.
Description	Description of the data input.

The fields in the following table show read-only information.

Field	Description
Name	The name of the data input: Agent Log Analytics. **Note:** All ACC data inputs have the same name. You can identify an ACC data input by the name of the MID Server that is defined for it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Status	The status of the data input.
Transport	The protocol used to send the log data. The ACC data input sends data using a ServiceNow Agent.
Sources count	The total number of log sources originating from all ACC data inputs together. This feature is supported in the Health Log Analytics application, Version 22.0.12 - December 2021 and later, available from the ServiceNow Store.
Disabled since	The time when the data input stopped or failed.
Last log time	The time when the last log streamed in the data input.
Error message	The streaming error. This field is populated automatically. It displays only when a streaming error has occurred.

Advanced configuration

Field	Description	Default value
Look up hostnames	Option for selecting to perform DNS lookup to resolve IPs to hostnames.	false
Use SSL	Option for selecting to use SSL.	true
Client inactivity timeout (sec)	The timeout, in seconds, to close an inactive channel.	15
Worker thread count	The number of threads that handle incoming data.	4
Default time zone	The default time zone of events. The system uses this default when the log does not specify a time zone.	GMT
Sub sample drop ratio	The ratio of events to drop.	-1
Sub sample receive ratio	The ratio of events to receive.	-1
Max length in bytes	The maximum length of log messages, in bytes.	32,766
Character encoding	The character encoding for this data input.	UTF-8

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration (continued)

Field	Description	Default value
Drop if queue is full	Option for selecting to discard logs if many processes are waiting in the queue to access the MID Server.	false

Agent Client Collector log policies

Agent Client Collector data inputs stream log messages to your ServiceNow instance using the ServiceNow Agent Client Collector. The Agent Client Collector Log Analytics (ACC-L) application provides predefined log policies for collecting log data from the CIs monitored by the ACC.

You can work with Agent Client Collector log policies in the following ways:

- Use the system's predefined log policies.

  The default ACC-L log policies consist of the monitored CIs and the check that runs on those CIs. For a complete list of the predefined log policies and monitored CI types, see Agent Client Collector Log Analytics default policies and checks.

- Edit a predefined log policy.

  For example, you can modify a default policy if your organization always uses a specific non-default log path. For more information, see Edit a log policy.

- Customize a published log policy by creating a child policy from it.

  You can use an existing log policy (the parent policy) as a template to create a new log policy (the child policy). You then customize the parameters that the child policy inherits from the parent to focus on a subgroup of its CIs or services. For example, create a child policy if your organization saves a subset of IIS Server logs to a different location. For more information, see Create a child log policy.

  Note: The child policy is inherently more specific than its parent policy. When Agent Client Collector matches a condition defined in the child policy, it will not try to match the parent policy.

- Define a new log policy when no policy exists for the CI that you want Agent Client Collector to monitor.

  For more information, see Create an Agent Client Collector log policy. For a general explanation of Agent Client Collector policies, see Checks and policies.

Agent Client Collector Log Analytics default policies and checks

Agent Client Collector Log Analytics provides various default policies and checks.

Streaming log data from Windows hosts is supported in Agent Client Collector Log Analytics Version 3.1.0 and later.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Default ACC-L policies

Name	OS	Description	Check
Apache on Linux Logs	Linux	Collects logs from Apache servers running on a Linux OS and sends them to the MID Server for further analysis.	log shipper
Apache on Windows Logs	Windows	Collects logs from Apache servers running on a Windows OS and sends them to the MID Server for further analysis.	log shipper
IIS logs	Windows	Collects logs from Microsoft IIS web servers and sends them to the MID Server for further analysis.	log shipper
JBoss/WildFly Application Server on Linux Logs	Linux	Collects logs from JBoss/WildFly application servers running on a Linux OS and sends them to the MID Server for further analysis.	log shipper
JBoss/WildFly Application Server on Windows Logs	Windows	Collects logs from JBoss/WildFly application servers running on a Windows OS and sends them to the MID Server for further analysis.	log shipper
Linux OS Logs	Linux	Collects logs from the Linux OS and sends them to the MID Server for further analysis.	log shipper
MSSQL on Linux Logs	Linux	Collects logs from MSSQL DBs running on Linux and sends them to the MID Server for further analysis.	log_shipper
MSSQL on Windows Logs	Windows	Collects logs from MSSQL DBs running on a Windows OS and sends them to the	log shipper

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	OS	Description	Check
		MID Server for further analysis.	
MySQL DB on Linux Logs	Linux	Collects logs from MySQL DBs on a Linux OS and sends them to the MID Server for further analysis.	log shipper
MySQL DB on Windows Logs	Windows	Collects logs from MySQL DBs running on a Windows OS and sends them to the MID Server for further analysis.	log shipper
Nginx on Linux Logs	Linux	Collects logs from Nginx servers running on a Linux OS and sends them to the MID Server for further analysis.	log shipper
Tomcat on Linux Logs	Linux	Collects logs from Tomcat servers running on a Linux OS and sends them to the MID Server for further analysis.	log shipper
Tomcat on Windows Logs	Windows	Collects logs from Tomcat servers running on a Windows OS and sends them to the MID Server for further analysis.	log shipper
Oracle WebLogic	Windows and Linux	Collects logs from Oracle WebLogic servers and sends them to the MID Server for further analysis.	log shipper

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	OS	Description	Check
		Note: • The log path that WebLogic uses has two environment variables, as described in the WebLogic documentation. For more information, see Environment variables. • On both Windows and Linux, Agent Client Collector Log Analytics adds a default environment variable from the WebLogic documentation to the log path. You can either use this default environment variable to get the full directory, or change it manually if you prefer to use a known direct directory. For information about customizing the environment variable, see Edit log path configurations.	
WebSphere Application Server on Linux Logs	Linux	Collects logs from WebSphere application servers running on a Linux OS and sends them to the MID Server for further analysis.	log shipper

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	OS	Description	Check
Windows OS Event Logs	Windows	Collects Windows event logs from Windows servers and sends them to the MID Server for further analysis.	log shipper for win events

Environment variables

The log path used for collecting logs from certain servers may contain environment variables. It's important to note that on Linux systems, the ACC Agent can only access environment variables that reside in the `/etc/environment` file. Add the environment variables that the agent must access to the `/etc/environment` file using the format *<NAME>='<VALUE>'*. On Windows systems, you can create environment variables in the usual manner.

Default ACC-L checks

Agent Client Collector Log Analytics (ACC-L) provides the following default **Type=Log** checks.

- The *log shipper* check launches a background process that sends logs from servers to an agent data input, according to the check's configured log path. When access permissions are lacking or the log path is configured incorrectly, the system sends an alert to the instance. For proper log configuration, ensure that the *servicenow* user has permissions to access the configured log path on the machine that is hosting the agent, and that the log path is configured correctly.
- The *log shipper for win events* check is similar to the *log shipper* check, but sends event logs from Windows OS servers to the agent data input.

> **Note:** To collect Security Windows event logs, *servicenow* users on an agent hosting machine must have admin privileges.

Create an Agent Client Collector log policy

Create an ACC log policy when no default policy exists for the CI you want Agent Client Collector to monitor.

Before you begin

- The Agent Client Collector Log Analytics (ACC-L) application, available from the ServiceNow Store, must be installed. For more information, see Agent Client Collector Installation.
- The Agent Client Collector comes with the default *servicenow* user. Ensure that this user has read access to enable Agent Client Collector to view all the configured log paths. For example, the Agent Client Collector *servicenow* user that comes installed with the base system does not have permissions to view the paths to `/var/log/` in Linux and `C:\Windows\System32` in Windows. For information about configuring permissions for the *servicenow* user, see the ACC-L Permission Denied issues [KB1117271] article in the Now Support Knowledge Base.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **ACC Log Analytics** > **ACC Log Policies**.
   The **Policies** page displays all Log Analytics policies. For a list of the policies that come with the base system, see Agent Client Collector Log Analytics default policies and checks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Select **New**.

   Note: For general information about creating an ACC policy, see Create an Agent Client Collector policy.

3. On the form, fill in the fields.

   Policy Definition Form

Field	Description
Name	A descriptive name for the policy.
Description	Description of the policy.
Publish status	Hard-coded as **Draft**, which means that the policy has not yet been published. You cannot edit this field.
Hierarchy	Hard-coded as **None**. When a child policy is added to the policy, the value changes to **Parent**. Child policies have a value of **Child**.

4. On the **Checks** tab, associate the log policy with the relevant log shipper check.
   - For Linux and Windows, except for Windows event logs, select the *`log shipper`* check definition.
   - For Windows event logs only, select the *`log shipper for win events`* check definition.
5. On the **Monitored CIs** tab, specify the CIs to which the policy applies.

   a. Choose the CI type to be monitored.
   - **Monitored CI type by filter**: Select the monitored CI type. You can narrow down the CIs that will be monitored by using filter conditions.
   - **Monitored CI type by script**: Specify the monitored CIs by using a script.
   - **Monitored CI type by CMDB Group**: Specify the monitored CIs by using CMDB group queries.

   For more information about choosing monitored CI types, see Create an Agent Client Collector policy.

   b. **Optional:** Monitor only CIs that are associated with an Application Service by selecting **Filter Monitored CIs by Application Service**.
   You can specify the Application Services to be monitored by using filter conditions. Agent Client Collector will only retrieve the logs of CIs that are associated with these Application Services.

6. Save the log policy.
   In the Check instances related list, a check instance record is created.
7. Open the relevant check instance record and then select **Edit in Sandbox**.
8. Select the Log path configurations related list.
9. Add a log path for the check instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: A check must have at least one log path configured for it to enable streaming logs. For more information about checks, see Checks and policies.

a. Select **New**.

b. On the form, fill in the fields.

Configuration form

Field	Description
Path	The full path from where the logs are streamed. You can use a wildcard. This field is required.
Component	The device type or stack layer that provides a context for the logs, used for anomaly detection and correlation. For example: Tomcat.
Source Type	Defines how Health Log Analytics handles a specific log type and parses the log data. For example: Tomcat Catalina.

c. Optional: For shipping multiline logs using Filebeat, configure the following properties.

(Optional) These parameters control how Agent Client Collector Log Analytics (ACC-L) handles messages that span multiple lines of text.

For more information, see Manage multiline messages in the Elastic documentation.

Field	Description
multiline.pattern (regex)	The regular expression to match. **Note:** You must define this property before you can configure the multiline.match and multiline.negate properties.
multiline.match	How ACC-L combines matching lines into a single log line. (Optional) The available options are None, Before, and After. Default is None.
multiline.negate	Option for determining whether the pattern identified in the log lines is negated. (Optional) The available options are None, True, and False. Default is None.

d. Optional: Define the following properties that control the Filebeat YML configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Fields	Field that enables you to include and exclude information in the output. For example, you can add a field for filtering the log data. (Optional) Add more Field rows by selecting the plus icon next to the Value field: ⊕. Remove a Field row by selecting the minus icon: ⊖. For more information, see the Log input Fields description in the Elastic documentation.
Configuration Options	Field that enables you to add configuration options to the log lines. For example, you can add the encoding to use. **Note:** Define only configuration options that are supported by Filebeat. For more information, see the Log input Configuration options description in the Elastic documentation.

e. Select **OK**.
The log path is created.

10. Select **Return to Policy**.
11. On the policy form, select **Publish**.
The policy's Publish status changes to **Published**.
12. **Optional:** Activate the policy by selecting **Activate**.

What to do next
Verify that the data input is streaming data.

Create a child log policy

Refine an existing log policy to focus on a subset of CIs or services by creating a child policy.

Before you begin
Role required: agent_client_collector_admin

About this task
A published log policy (the parent policy) can be used as a template for creating a new log policy (the child policy). You can customize the parameters that the child policy inherits from the parent to focus on a subgroup of its CIs or services.

The child policy inherits parameters from its parent policy that are modified after the child policy's creation, unless you disable this option in the child policy. You can choose to review these modified parameters in draft mode in each child policy before accepting them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The child policy is inherently more specific than its parent policy. When Agent Client Collector matches a condition defined in the child policy, it will not try to match the parent policy.

Procedure

1. Navigate to **All** > **ACC Log Analytics** > **ACC Log Policies**.
2. On the **Policies** page, select a published policy with a Hierarchy value of either **None** or **Parent**.
3. Select **Create Child**.
   A new policy opens with the same parameters as the published policy. The name of the new policy is *<Policy name><date and time created>*.

   For example, MySQL DB on Linux Logs 2022-02-13 04:44:12.

   The following fields are read-only:

   Child policy fields

Field	Value
Publish status	Draft
Hierarchy	Child
Parent	<Name of the parent policy>

   The **Order** field value indicates the priority by which the agent attaches itself to the child policy in relation to its sibling policies. A lower number indicates a relatively higher priority for this child policy. By default, each new child policy takes on a higher value (and hence, a lower priority) than the one preceding it. You can modify the order value as needed.
4. Customize the child policy's parameters.
   For a description of the parameters, see Create an Agent Client Collector log policy.
5. Select **Save** to create the child policy.
6. **Optional:** Select **Disable inherit option** to disable automatically inheriting modified parameters from the parent policy.
7. Select **Publish** to publish the child policy.

Result
The child policy is published and appears on the **Child Policies** tab of the parent policy.

Edit a log policy

You can edit published Agent Client Collector log policies. For example, modify a default log policy if your organization always uses a specific non-default log path.

Before you begin
Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **ACC Log Analytics** > **ACC Log Policies**.
2. Select a policy for which the **Publish status** is **Published**.
3. Select **Edit in Sandbox**.
   The policy is now in **Draft view**, and its fields and check instances are editable.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Although the policy is in **Draft view**, its publish status remains **Published**.

4. On the **Monitored CIs** tab, edit the CIs to which the policy applies as needed.

   a. Choose the CI type to be monitored.

   - **Monitored CI type by filter**: Select the monitored CI type. You can narrow down the CIs that will be monitored by using filter conditions.
   - **Monitored CI type by script**: Specify the monitored CIs by using a script.
   - **Monitored CI type by CMDB Group**: Specify the monitored CIs by using CMDB group queries.

   For more information about choosing monitored CI types, see Create an Agent Client Collector policy.

   b. **Optional:** Monitor only CIs that are associated with an Application Service by selecting **Filter Monitored CIs by Application Service**.
   You can specify the Application Services to be monitored by using filter conditions. Agent Client Collector will only retrieve the logs of CIs that are associated with these Application Services.

5. **Optional:** In the **Check Instances** table, select the log shipper check instance record and edit it.
   For more information, see Edit log path configurations.
6. Select **Save** to save the changes you entered into the Sandbox policy fields.
   The policy remains in **Draft view**, and the **Publish status** is **Published***. The asterisk (*) indicates that the current published version is not the most up to date one, because it contains saved changes that have not yet been published.
7. Select **Republish** to publish the updates.
   The policy returns to **Published** state and opens as the published policy.

   Note: This option appears only after entering updates into the policy fields, for a policy whose **Publish status** is **Published***.

8. Determine the active status of the policy by selecting either **Activate** or **Deactivate**.
9. **Optional:** Select **Delete** to delete both the Published and the draft policies.

Edit log path configurations

Modify a log path configuration in the log shipper check instance.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **ACC Log Analytics** > **ACC Log Policies**.
2. Select a policy record.
3. In the Check instances related list, select the log shipper check instance record.
4. Select **Edit in Sandbox** to enable editing the record.
5. On the Log path configuration related list, edit a log path configuration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select the log path configuration that you want to edit.

b. Edit the configuration.

- You can add custom fields, allowing you to include additional information in the output.
- You can add extra configuration options, such as `encoding`, `include_lines`, and `exclude_lines`, to filter your log data. For information about discovering file encoding, see the How to Detect File Encoding on Windows, Linux and Mac [KB1702770] article in the Now Support Knowledge Base.

  These configuration options impact all the files that the data input streams to your instance.

  Note: Define only configuration options that are supported by Filebeat. For more information, see the Log input Configuration options description in the Elastic documentation.

c. Select **Update** to save your edits.

6. **Optional:** Add an environment variable or edit an existing one in the **Path** field.
   Use the relevant format for its name:
   - On Linux, capitalize the name and precede it with a dollar sign: $. For example: `$HOME`. The name is case sensitive.
   - On Windows, surround the name with percentage signs: %. For example: `%HOMEPATH%`. The name is not case sensitive.

   Important: The servicenow user and ACC agent must have access to the environment variable. On Linux systems, only environment variables that reside in the `/etc/environment` file are accessible. Add the environment variable to this file using the format *<NAME>='<VALUE>'*.

7. **Optional:** Delete a log path.

   a. Check the box next to the log path configuration that you want to delete.

   Note: You can select multiple log path configurations, but a check must have at least one log path configured for it to enable streaming logs. For more information about checks, see Checks and policies.

   b. Select **Delete**.

8. Select **Return to policy** to go back to the log policy record.
9. Select **Republish** and then confirm to publish the policy again with the updated log path configurations.
   The changes take effect in a few minutes.

CI association to a dynamic CI Group in Agent Client Collector Log Analytics

When the Agent Client Collector Log Analytics (ACC-L) application discovers a CI that is not associated to an application service, it creates a dynamic CI Group. The application then associates the CI to the group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A dynamic CI Group is a group of CIs that are organized by attribute, such as the CI type. By default, Agent Client Collector log policy checks monitor only CIs that are associated to an application service. The Agent Client Collector Log Analytics application inspects the CIs that the Agent Client Collector discovers. When it discovers a CI that is not associated to an application service, the application creates a dynamic CI Group and associates the CI to it.

If the CI is later associated to an actual application service, its association to the dynamic CI Group is automatically deleted. The baseline for the CI is lost in this process. While the system is establishing the new baseline, it doesn't issue alerts for the application service to which the CI is now associated.

You can choose to disable the association of unassociated CIs to a dynamic CI group by changing the default setting of the relevant system property, *sn_accl.dynamic_service.creation.enabled*, to false. The Agent Client Collector Log Analytics application will then not associate unassociated CIs to dynamic CI Groups. In this case, logs are not streamed for CIs that are not associated to an application service.

View live CI data with Agent Client Collector

View live data for incident-related CIs through Agent Client Collector for information that can help resolve the incidents.

Before you begin

- Collecting running processes for a macOS system requires `sudo` privileges. If `sudo` privileges aren't granted, only processes run by the Agent Client Collector are collected.

  To enable `sudo` privileges for osqueryi, add the following strings to either the `/etc/sudoers` file or to an individual file in the `/etc/sudoers.d` directory:

  - `Cmnd_Alias LIVE_CI_VIEW = /Library/Application\ Support/servicenow/agent-client-collector/osquery/bin/osqueryi`
  - `_servicenow ALL=(ALL) NOPASSWD:LIVE_CI_VIEW`
- To enable the OSQuery executable to retrieve information on logged-in users in a Windows environment, the agent must run as a local SYSTEM account.

Verify that you have installed Agent Client Collector Framework and Service Operations Workspace on your instance.

Role required: itil

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the Lists icon ( ).
3. Select **Incidents** > **All**.
4. Select an incident from the list.
5. On the <name of incident> form, select the **Live CI Data** tab.
   The CI data is displayed in the cards described in the following table.

Live CI Data tab cards

Card	Description
Device details	General details about the CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Card	Description
Memory Usage	Amount of memory being used by the CI.
Top 5 running processes by CPU usage	The five CI processes using the most CPU.
Top 5 running processes by memory usage	The five CI processes using the most memory.
Disk Usage %	Percentage of disk space in use on the CI.
Logged-in users	Users logged in to the CI. Whether this data is visible depends on your OS privileges. ○ MacOS: Requires sudo privileges ○ Windows system: Requires administrator privileges ○ Linux: No special privileges are required

Related topics

View live CI data logs

Assign a CI to an incident to view live CI data

View live CI data logs

View CI logs to receive additional details on an incident when viewing live CI data. For example, viewing log information for a process with high CPU usage can help you pinpoint when the usage went up, helping you determine the root cause of the issue.

Before you begin

Ensure that you have installed Agent Client Collector Framework and Service Operations Workspace on your instance.

Role required: itil

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the Lists icon ( ).
3. Select **Incidents** > **All**.
4. Select an incident from the list.
5. Select the **Live CI Data** tab.
6. Select the **Logs** link on the right side of the pane.

Logs link

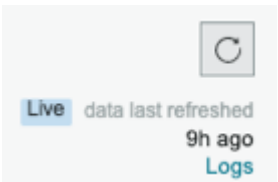

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The **data last refreshed** value follows the instance time zone and not your local system time zone. Therefore, ensure that your local system time zone matches the instance time zone.

The logs appear on the **Log** page.

Assign a CI to an incident to view live CI data

Retrieve live CI data to help troubleshoot the issues that caused an incident by assigning the problematic CI to its incident. An incident typically is associated with the CI that caused the incident, but if this was omitted during the incident creation, you can assign the CI manually.

Before you begin
Role required: itil

Procedure

1. Navigate to **Workspaces** > **Service Operations Workspace**.
2. Select the Lists icon ( ) icon.
3. Select **Incidents** > **All**.
4. Select an incident with a value of **Configuration item = (empty)**.
5. Select the **Live CI Data** tab.
6. Enter the first several characters of the CI that caused the incident in the **Primary configuration item** field.
   A list of CIs matching the string appears.
7. Select the relevant CI from the list.
   The name of the CI appears in the **Primary configuration item** field, and the CI's live data appears on the **Live CI Data** tab.

AI Control Tower

AI Control Tower (AICT) is an enterprise-grade governance control plane which provides centralized visibility, monitoring, and policy enforcement across all AI systems and platforms within an organization. AICT addresses the challenge of uncontrolled AI adoption by intercepting AI-related traffic at the network, device, or application layer and applying organization-wide security and data governance policies.

Business Challenge

As organizations adopt AI systems, multiple AI platforms are available for use. These platforms include including cloud-based large language models (LLMs), browser-based AI tools, autonomous AI agents, local offline models, and embedded AI services. Organizations currently lack the visibility and control mechanisms necessary to govern these interactions effectively.

Key challenges include:

- Limited visibility into which AI models and services are in use across the organization.
- Lack of awareness about which data is being shared with external AI systems.
- Insufficient oversight of agent-to-agent communication and data exchange.
- Risk of confidential or sensitive data being exposed to unauthorized external systems.
- Unclear permissions and access controls granted to AI tools.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Compliance gaps in AI-related interactions and data handling.
- User awareness gaps regarding the sensitivity of information shared with external AI systems.

These factors create organizational risk through potential data leakage, compliance violations, and uncontrolled shadow AI adoption.

AICT Solution

AI Control Tower functions as a comprehensive governance platform that intercepts, monitors, and controls AI-related traffic and interactions. The platform applies enterprise security and data governance policies to provide centralized oversight of all AI systems and agent activities within an organization.

AICT enables you to maintain granular visibility into AI usage patterns, classify and protect sensitive data, enforce compliance policies, and reduce the risk of unauthorized data exposure.

Key Capabilities

AI Control Tower provides the following primary functional capabilities:

Capability	Description
Visibility	Identifies LLMs and AI systems that you can interact with. Tracks agent-to-agent communication, monitors browser-based AI usage, and detects local AI model usage where applicable.
Data Awareness	Inspects outbound prompts and file uploads in accordance with policy and consent requirements. Classifies data sensitivity including personally identifiable information (PII), financial data, and confidential intellectual property. Detects when regulated or confidential data is being shared externally.
Central Policy Enforcement	Enforces organization-wide AI usage policies. Blocks access to unauthorized AI services, restricts specified data categories from external sharing, and applies conditional access controls based on risk assessment.
Risk and Compliance Monitoring	Detects potential data leakage and unauthorized autonomous agent activity. Generates audit logs and compliance reports to support regulatory requirements and organizational governance.

Business Benefits

AI Control Tower provides:

- Prevention of confidential data leakage to external AI systems.
- Improved AI governance and organizational compliance posture.
- Centralized control over AI tool adoption and shadow AI reduction.
- Increased user awareness regarding AI data sharing risks.
- Safe and monitored deployment of autonomous agents.
- Comprehensive audit trails for compliance and regulatory reporting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI Proxy

AI Proxy gives customers visibility and control over AI traffic leaving managed endpoints. It intercepts outbound calls to AI providers, attributes each call to the originating process, and streams structured events to the AI Control Tower (AICT) on your ServiceNow instance.

AI Proxy - Overview

AI Proxy operates as a lightweight HTTPS proxy embedded inside your existing Agent Client Collector (ACC). You activate AI Proxy through a configuration flag during installation. AI Proxy intercepts outbound calls to AI providers and attributes each call to the originating process. The proxy streams structured events to AICT on your ServiceNow instance. Non-AI traffic flows through without inspection.

AI provider classification

After detecting AI traffic, the AI Proxy enforces policies based on your organization's classification of AI providers.

Authorized LLM providers

These providers your organization approves for use under organizational policies. AI Proxy forwards traffic to these providers after policy inspection and enforcement.

Unauthorized LLM providers

These providers your organization does not approve. AI Proxy blocks traffic to unauthorized providers when it detects their use.

Policy enforcement and monitoring

AI Proxy tracks AI activities by intercepting and monitoring AI-related traffic. The system detects policy violations in real time. Your ServiceNow instance syncs governance policies for applicable devices and users. AI Proxy then sends this data to ServiceNow, where the system maps the data to users, devices, and departments to provide comprehensive insights into AI usage across your organization.

Set up an AI proxy

Configure an AI proxy server to monitor AI commands traveling from your Agent Client Collector (ACC) agent to the AI Control Tower (AICT).

Before you begin

- Install the ITOM Cloud Services Core (sn_itom_cloud_svc) plugin.
- Onboard your instance to use ITOM Cloud Services. For details, contact Customer Support.
- Configure an agent registration key. For details, see Configure an agent registration key.

You must configure ACC with an ICS (mid-less) architecture to enable using an AI proxy. For details on configuring ACC without a MID Server, see Configuring MID-less Agent Client Collector.

Role required: agent_client_collector_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Retrieve the publicly accessible gateway URL, based on your location.
   - AMER (Americas): `itomcnc-prod-gateway-amer.sncapps.service-now.com:443`
   - EMEA (Europe): `itomcnc-prod-gateway-emea.sncapps.service-now.com:443`
   - APAC (Asia Pacific): `itomcnc-prod-gateway-apac.sncapps.service-now.com:443`
2. Run the relevant command, depending on the OS you're using.
   - On a Windows server:

```
msiexec /i <msi_file_path> /quiet /qn /norestart
 CONNECT_WITHOUT_MID="true" ACC_CNC="<itom_gateway_url>"
 REGISTRATION_KEY="<agent_registration_key>"
 INSTANCE_URL="https://<instance_url> AI_PROXY_ENABLED="true"
 LOCALUSERNAME="LocalSystem"
```

   - On a Linux or macOS server:

```
sudo CONNECT_WITHOUT_MID="true" ACC_CNC=<itom
 gateway url> REGISTRATION_KEY=<agent
 registration key> INSTANCE_URL=<instance url>
 AI_PROXY_ENABLED="true" LOCALUSERNAME="LocalSystem"
 bash -c "$(curl -L https://<instance
 url>/api/sn_acc_shadow_ai/agents/install_agent)"
```

   For details on the parameter values in the command, see Agent Client Collector MID-less installation command parameters.

Result
The AI proxy monitors AI traffic coming from your AI tools. Traffic is either allowed to pass through to the AI Control Tower, or is blocked, depending on your configured AICT policies.

What to do next
Configure AICT policies to determine which AI resources are blocked, and for whom. For details, see .

Create AI Control Tower (AICT) policies

Create AICT policies to block personnel from using unauthorized AI tools. AICT polices are synchronized with the Agent Client Collector, which invokes the policies that run.

Before you begin
You must configure ACC with an ICS (mid-less) architecture to enable using an AICT policies. For details on configuring ACC without a MID Server, see Configuring MID-less Agent Client Collector.

Ensure that ITOM Shadow AI governance is installed in your enviornment.

Role required: agent_client_collector_admin

Procedure

1. Navigate to **All** > **AI Control Tower** > **AI Assets**.
2. In the filter on the left side of the page, expand **Govern** and select **Policies**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select **Create policy**.
4. Select a policy type in the **Choose a policy type** field.
   - A threat response: Takes an agent offline when a security threat is detected.
   - An explicit block: Blocks a user or department from using an AI agent, model, or domain.
5. Configure the relevant parameters, based on the selected policy type.
   - A threat response:
     a. Select what will happen when a threat is detected.
     b. In the **Which agents** cell, select the conditions by which the detected threat takes the agent offline.
     c. In the **sensitivity** field, select whether the agent is to be taken offline on every detection, or only above a specified threshold.
     d. Select a follow-up action in the **Do this** cell, after the agent is taken offline:
        - Create a ticket: Opens a ticket in the indicated group's queue.
        - Notify: Sends a notification to the user.
     e. Select the **Publish** button to publish the policy.
   - An explicit block:
     a. Select whether the block is to apply to a user or a department in the **Who** field.
     b. Select what is being blocked in the **What** field, either an AI agent, a model, or a domain.

        When selecting an AI agent, configure a filter for the field values that create the block, as needed.

     c. Select a follow-up action of **Notify** to send a notification to the user, in the **Add follow-up** drop down.
     d. Select the **Publish** button to publish the policy.

Result

- The action to be taken appears in the **Policy summary** cell.
- The policy is synchronized with the Agent Client Collector policies, and appears in the list of Agent Client Collector policies. Agent Client Collector invokes the policy and takes the indicated action.

Agent Client Collector reference

Reference topics provide additional information about Agent Client Collector.

Agent Client Collector Framework reference

Reference topics provide additional information about mapping and fine-tuning application services using Agent Client Collector Framework lists and forms.

Agent Client Collector user roles

Administrators can assign user roles to grant access to the Agent Client Collector. The following standard roles for the Agent Client Collector Framework (ACC-F) are included in the ServiceNow system with the Agent Client Collector installed on it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector user roles

Role	Description
agent_client_collector_admin	All permissions on Agent Client Collector, including: • Access to REST API • Add, edit, and delete policies, checks, agents, web sockets • MID Server management
agent_client_collector_integration	Access to REST API
agent_client_collector_user	• View checks and policies • Access to some REST APIs
evt_mgmt_admin	Starts the ACC listener.

Agent Client Collector Framework default checks

Agent Client Collector Framework provides default checks with the base system.

Checks

The following table lists the default checks that come with the Agent Client Collector Framework base system.

Type	Name	Description	Command
Restart	restart	Restarts the agent. Check is active by default	Restart agent
Discovery	check-discovery-basic	Performs basic discovery of the agent's host to discover: • Host CI, including IP address, serial number, and hostname. • Running processes (cmdb_running_process) of the host • TCP ports (cmdb_tcp) The check is performed by default every 12 hours (43200 seconds), and times	endpoint_discovery.rb

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Name	Description	Command
		out after 600 seconds of inactivity. Plugins: osquery, acc-f-commons, acc-f-modules	

Note: The *check-read-log* and *config-file-reader* checks are for internal use only.

Agent Client Collector diagnostic self tests

The following table describes each diagnostic self-test that runs in the Agent Self Test Run (`sn_agent_self_test_run`) table.

Agent Client Collector (ACC) diagnostic self-tests

Self-test name	Description
Allow List Validation	Checks if the allowlist is enabled and whether the allowlist file is a valid JSON file.
Cache Subdirectory Creation	Verifies that a user can create subdirectories in the cache directory. Required by ACC as it creates a directory per plugin.
MID Authentication Configuration	Checks if the MID Server authentication is configured. At least 1 authentication method is required.
MID Connection	Tests the connection to the MID Server over both HTTP(S) and WebSockets.
MID Connection Configuration	Checks if the MID connection is configured. At least 1 MID Server connection is required.
OSQuery Test	Verifies OSquery executable availability and executes a basic system query.
PAC File	Validates the configured PAC file is accessible and valid.
Powershell Version	Validates Powershell installation and version.
Ruby Script Execution	Verifies that the current user is able to execute a basic Ruby script.
Sudo SS Command Access	Validates sudo access via `ss` command.
WMI Permissions	Checks WMI permissions for process queries.

Agent Client Collector performance and footprint for URL monitoring

The following tables display the agent performance KPIs and its footprint on the host during URL monitoring data collection execution on different operation systems.

URL monitoring checks on a Linux system

Linux Cent OS 8, 2 CPUs, 4GB RAM					
Checks per minute	100	250	500	1000	1500

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

URL monitoring checks on a Linux system (continued)

Linux Cent OS 8, 2 CPUs, 4GB RAM					
Metrics (expected/ actual)	700/700	1750/1750	3500/3500	7000/7000	10500/10500
Network utilization - Tx (MID > Agent)	35 KB/s	78 KB/s	146 KB/s	286 KB/s	430 KB/s
Network utilization - Rx (MID < Agent)	15 KB/s	27 KB/s	51 KB/s	107 KB/s	116 KB/s
Network utilization - Tx (Agent > MID)	8 KB/s	23 KB/s	37 KB/s	89 KB/s	113 KB/s
Network utilization - Rx (Agent < MID)	28 KB/s	72 KB/s	142 KB/s	285 KB/s	426 KB/s
Memory consumption	22 MB	25 MB	26 MB	34 MB	39 MB
CPU of all checks	~3%	~7%	~11%	~21%	~24%
Process CPU utilization	1% - 1.5%	1.5% - 2.5%	5% - 6%	6% - 7%	12% - 13%
Host CPU utilization	~10%	~19%	~23%	~45%	~53%

Agent Client Collector certificate revocation reasons

The following table lists and describes the possible reasons for revoking an Agent Client Collector certificate to stop communication between the agent and ITOM cloud services.

Certificate revocation reasons

Reason	Description
Unspecified	Revoke a certificate without providing a specific revocation code. **Note:** This reason code doesn't provide an audit trail identifying why a certificate was revoked.
Key Compromise	You suspect that the private key associated with a certificate is compromised. For example, if a laptop belonging to a user in your organization is stolen, any private keys stored on the laptop may be compromised.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate revocation reasons (continued)

Reason	Description
CA Compromise	You suspect that a Certificate Authority's (CA) private key has been compromised and is in the hands of an unauthorized individual. If a CA's private key is revoked, the CA hierarchy considers all certificates below that CA revoked as well.
Affiliation Changed	An individual is no longer working in the organization, or the computer account to which the certificate was issued is no longer in use. Can also be used if a person changes roles within an organization and no longer requires using the certificate associated with their previous role. For example, an employee could move from the purchasing department and no longer require a certificate to authorize purchase requests.
Superseded	Indicates that a new certificate must be issued in place of a previously issued certificate. For example, if you update a certificate template and reissue certificates, you could revoke the previous certificate with this reason code.
Cessation of Operation	A server or workstation is decommissioned, and all certificates issued to the server are no longer required. You can use this revocation reason when decommissioning a CA.
Certificate Hold	A temporary revocation indicating that a CA won't validate a certificate at that specific time. **Note:** Using this reason code makes it difficult to determine whether a certificate was valid at a specific time.
Remove from CRL	Enables future unrevoking of the certificate, removing the certificate from the certificate revocation list (CRL).
Privilege Withdrawn	The certificate holder no longer has the privileges required to continue using the certificate.
AA Compromise	Indicates suspected or actual compromise of the authentication authority (AA) validated in the certificate.

Agent Client Collector check definition page

The fields to be configured on the Check Definition page, when creating a check definition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Check definition form

Field	Description
Name	A descriptive name for the check.
Check Type	Select the type of check. The available monitoring options are: • Events • Metrics
Active	Select to activate the check.
Is Proxy Valid	Select to indicate that the check's policy is set to work as a proxy. When viewing this option in the check instance, it is read-only.
Command Auto Generation	Select to automatically generate the command with the **Command Prefix** value and the **Parameters** tab values.
Command Prefix	The command that is used for automatic generation when the **Command Auto Generation** check box is selected. The prefix consists of any portion of the command which is static (does not change), such as the script name.
Command	• Enter the command that the Agent Client Collector executes. **Note:** When the **Command Auto Generation** check box is selected, this field is automatically populated with the prefix and flags of the active parameters specified in the **Parameters** tab. • The field contains parameters taken from one or more of the following: ○ Simple template: `{{.labels.params_field_name}}`: Takes values from entries in the **Parameters** tab. For example, a string of `{{.labels.params_warning}}` means that the value of the *warning* parameter is used (provided it has been configured on the **Parameters** tab). ○ Monitored CI: `{{.labels.params_ci_field_name}}`: Takes values automatically from the monitored CI, where `field_name` is replaced with the exact name of the required field in the CI table.
Configuration Data Files	Select the configuration data files to be downloaded when the check runs.
Description	Description of the check.
Exec Mode	The execution mode of the check. Available options are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Check definition form (continued)

Field	Description
	• *shell*: To be used for checks which do not support *execv*. • *execv*: Provides enhanced security during check execution. The default execution modes are: • Existing check definitions: *shell* • New check definitions: *execv* • If no value is entered in the field: *shell* When upgrading to version 2.9.0, if this field is read-only, ensure that you disable the read-only status, as described in the (KB1122498) article in the Now Support Knowledge Base.
Interval-based scheduling	When selected, enables configuring a time interval (in seconds) to wait between check executions.
Interval	The amount of time, in seconds, to wait between check executions. For example, a value of 60 means that the check runs every 60 seconds. Specified value must be an integer. Appears only when selecting **Interval-based scheduling**.
Cron-based scheduling	When selected, enables selecting the frequency with which check results are to be sent, in the **When to send the check results** field.
When to send the check results	The time or frequency to send check results. • When event severity changes (default for event checks) • Always send check results (default for non-event checks) • Do not send check results • When the check results change • (Advanced) Based on exit codes of the check: Send results only if the exit code is 240 (OK), 241 (Warning), or 242 (Critical). However, if the command output is empty, then even if one of these codes are part of the exit code, check results aren't sent.
Timeout	The amount of time, in seconds, after which the check execution stops when no output is returned. For example, a value of 60 means that when the check execution doesn't return a value for 60 seconds, the execution stops.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Check definition form (continued)

Field	Description
Check group	A record in the *sn_agent_check_group* table consisting of the group name and CMDB CI class type. Assign a check group to enable filtering only checks in the group for selection when defining a policy.
Event status change threshold	The number of consecutive times that a check's response status must happen before a new event is sent. For example, if this value is set to 5, a check whose response status changes from **OK** to **Error** generates a new event with an **Error** status only after the fifth consecutive occurrence of the status change. After an agent restart, an initial event is always sent. Default value: 5 (except for AWS event policies)
Event status repair threshold	The number of consecutive times that a check's response status must improve to close the previous event. For example, if this value is set to 3, a check whose response status changes from **Error** to **OK** closes the previous event and generates a new event with an **OK** status only after the third consecutive occurrence of the status change. Default value: 1 (except for AWS event policies)
Credential failure regex	The string from the check's output indicating that the check failed due to a credentials error. The agent then continues searching credentials that are part of the credential alias until it finds the correct one. This field is relevant (and mandatory) when the check's policy uses a credential alias. When creating a new check, enter the string manually.
Plugins	Select the plugins to be associated with the check. Once this is done, the plugins download before the check executes.

Check definition form parameters tab

The Parameters tab definitions for the **Command** field on the check definition form.

Parameters Tab

Field	Description
Name	A descriptive name for the parameter.
Default value / value	Default value / value of the parameter.
Active	Boolean. When set to **true**, the parameter is taken to the **Command** field when the **Command Auto Generation** check box is selected.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameters Tab (continued)

Field	Description
Mandatory	Boolean. When set to **true**, a parameter value must be specified for the command to run.
Flag	String. Prefix that is added before the parameter value in the **Command** field when the **Command Auto Generation** check box is selected.
Value required	Boolean. When set to **true** and when the **Command Auto Generation** check box is selected, the value of this parameter is added to the **Command** field after the **Flag** value.

Test check dialog box fields

The dialog box fields when testing a check definition.

Test Check dialog box fields

Field	Description
Check Definition	Displays the name of the selected check definition.
CI	Select the CI on which to run the check. The available options are those CIs which meet both of the following criteria: • Belong to the group specified in the **Check group** field. • Belong to an agent with **Up** status. If the **Check group** field is empty, the value defaults to **cmdb_ci_computer**.
Credentials	Select the credentials of the user whom you want the agent to connect with. This field appears only if you have configured secure check parameters on the **Check Secure Parameter Definitions** tab in the Check Definition.
Proxy Agent	Select a proxy agent to be associated with the check. This field appears only when working with a check instance and when one of the options has been selected on the **Proxy Settings** tab in the instance's policy.

Agent Client Collector data collection properties

Description of the properties that determine the behavior of Agent Client Collector data collection.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collection properties

Property	Description
sn_agent.host-data_collection_refresh_duration	The amount of time (in seconds) between agent data collections on the host. Applicable only when no host CI is assigned to the agent. Default: 43200 (12 hours)
sn_agent.disco_minimum_threshold	The minimum number of minutes since the last endpoint Discovery on a given IP before Discovery can be triggered again. This property prevents running endpoint Discovery too often, such as in cases where a network intermittently has a break or a laptop is repeatedly rebooted. Default: 60
sn_agent.disco_disable_ci_clobber	When set to *true*, indicates that data collected by the ServiceWatch or ServiceNow Discovery sources takes precedence over Agent Client Collector when the threshold in the *sn_agent.disco_ci_clobber_of_agentless_disc* property hasn't been exceeded. Default: true
sn_agent.disco_ci_clobber_of_agentless_threshold	Maximum number of days since an existing CI has been discovered using agentless Discovery where the existing CI is linked instead of being rediscovered by the agent. Takes effect when *endpoint_disco_disable_ci_clobber_of_agentl* = **true** and over 14 days (by default) have passed since a non-ACC Discovery source (such as ServiceWatch or ServiceNow) has discovered the CI. After this threshold has been exceeded, agent Discovery can execute and overwrite an existing CI regardless of whether the CI has been previously discovered by a non-ACC Discovery source. This process presents the most recently discovered data in the CMDB. Default: 14
sn_agent.ci_prefers_ip_version	Indicates whether IPv4 or IPv6 is the preferred version when saving an IP Address value in the IP Address field of the CI record. Default: 4
sn_agent.ci_serial_number.pref_order	Specify the order of the fields to locate a serial number value to be saved to the host CI's *serial_number* field during agent data

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data collection properties (continued)

Property	Description
	collection. Values must be comma-separated, without spaces, and in lowercase. Default: bios,system,uuid,baseboard,chassis
sn_agent.mac_ci_name_pref_order	Specify the order of the values by which the CI record's *Name* field is populated during macOS data collection. Values must be comma-separated, without spaces, and in lowercase. Default: name,hostname,computer_name,local_hostname
sn_agent.host_data_collection.disable_container	When set to *true*, indicates that data collection is disabled for containerized agents. Default: true

Agent Client Collector policy collection properties

Description of the properties that determine the behavior of Agent Client Collector policy collection.

Policy collection properties

Property	Description
sn_agent.full_policy_calculation_cycle_days	The number of days for a policy to reach a full calculation cycle. If the policy's last full calculation cycle was more than 7 days ago, the policy is queried to fully calculate its calculation cycle. Default: 7
sn_agent.max_number_of_monitored_cis_per_policy	The maximum number of monitored CIs that can be processed per policy. Default: 400000
sn_agent.max_num_of_monitored_cis_unassociation	The maximum number of monitored CIs unassociated with agents that can be processed per policy. Default: 1000
sn_agent.policy_force_recalc_interval	The time limit, in minutes, used to determine whether forced recalculation is required for any policies. Forced recalculation redistributes both a policy's monitored CIs and the corresponding agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Policy collection properties (continued)

Property	Description
	Default: 10
sn_agent.policy_force_recalc_required	When set to *true*, checks whether forced recalculation is required for policies. Forced recalculation redistributes a policy's monitored CIs and the corresponding agent. Default: true
sn_agent.populate_proxy_agent_monitored_cis	When set to *true*, repopulates the monitored CIs for given policies to show the distribution of CIs among proxy agents. Default: false

Agent Client Collector Framework configuration properties

Description of the properties that determine the behavior of Agent Client Collector Framework configuration.

Framework configuration properties

Property	Description
sn_agent.agent_static_import_chartreader_limit	The character limit for a given payload item. Relates only to the configuration item being processed and not the total attachment size. Default: 1000000
sn_agent.agent_upgrade_wait_time	Maximum number of minutes for an agent to successfully upgrade before entering *Warning* status. Default: 30
sn_agent.agent_version	The current ACC-F release version installed on the instance. Type: String
sn_agent.appl_classification_behavior	The characteristics and behaviors for applications during Agent Discovery. When set to *simple*, creates a shell CI. Default: simple
sn_agent.auto_upgrade.enabled	When set to true, the instance automatically upgrades existing Agent Client Collector installations to the version set in the sn_agent.agent_version property, if the agent version is lower than the property's version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Framework configuration properties (continued)

Property	Description
	Type: Boolean Default: false
sn_agent.auto_upgrade.max_upgrades_per_hour	The maximum number of Agent Client Collector upgrades possible per MID Server every hour. Type: Integer Default: 50 **Note:** Modifying this value may negatively affect system performance.
sn_agent.auto_upgrade.max_upgrades_per_hour	The maximum number of Agent Client Collector upgrades possible for all agents on all MID Servers. Type: Integer Default: 1000 **Note:** Modifying this value may negatively affect system performance.
sn_agent.config_publish.max_payload_size	The size, in bytes, that the config_publish_payload can reach on the MonitoringProbe ECC queue topic before an error is logged and the payload is skipped. Default: 15000000
sn_agent.ecc_queue.max_payload_size	The maximum allowable payload size for an ECC queue record to be processed by the agent. Default: 5000000
sn_agent.enable_auto_mid_selection	When set to **true**, the instance sends a list of MID Servers to the agents. To enable automatic selection of a MID Server from the list, set the *auto-mid-selection* property to true in the agents' acc.yml file. Automatic selection of MID Servers ensures that each agent uses the most efficient available MID Server. Default: false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Framework configuration properties (continued)

Property	Description
	Note: This property determines only whether a list of MID Servers is sent to the agent. • If this property was previously set to **true** and is then set to **false**, automatic MID selection is performed on the previously sent MID Server list if the *auto-mid-selection* property is set to **true** in the acc.yml file. • If this property is set to **true** and the *auto-mid-selection* property in the acc.yml file is set to **false**, no automatic MID Server selection is performed.
sn_agent.installer_max_endpoints	Maximum number of MID Web Server endpoint URLs to be selected during agent single line installation. Default: 3
sn_agent.keep_alive.alive_duration	The amount of time, in seconds, within which the agent's last_refreshed timestamp must be updated. After this time period, the agent status moves from *Up* to *Warning*. If the indicated time period passes again, the status moves from *Warning* to *Down*. Default: 120
sn_agent.keep_alive.disconnected_duration	The amount of time, in seconds, within which the last_keepalive timestamp must be updated. After this time period, the agent status changes from *Up* to *Disconnected*. Default: 600
sn_agent.logging.verbosity	The minimum level of messages that are logged for Agent Client Collector Framework. Messages with the indicated status and anything more severe are logged. Options: • error • warn • info • debug Default: info

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Framework configuration properties (continued)

Property	Description
sn_agent.max_sys_ids_in_query	The maximum number of sys IDs to return during GlideRecord queries. Default: 400000
sn_agent.maximum_mids_to_select	The maximum number of MID Servers that can be selected for agent setup. Default: 20
sn_agent.min_upgrade_history_per_agent	The minimum number of upgrades saved per agent. Default: 2
sn_agent.registration_key_validity_days	The number of days after creation of a registration key that the key is valid before it is deleted. Default: 90
sn_agent.registration_key_notification_days	The number of days before registration key deletion that an email notification is sent to the key's ownership group, informing of pending deletion. Default: 14
sn_agent.sync_filters_interval_min	The number of minutes after which CMDB changes are detected and then synched to the MID Server. Default: 10
sn_agent.time_ago_in_seconds_to_query_processing_sec	When processing sec_queues, only query those created within the indicated time (in seconds). Default: 43200 (12 hours)

Agent Client Collector data collection tables

Agent Client Collector performs data collection based on the scoped apps that you've installed. Agent Client Collector Framework performs basic data collection, and Agent Client Collector for Visibility Content performs enhanced data collection.

The Refresh Host Data for Agents scheduled job collects host data every hour on hosts that haven't had data collection run in the past 12 hours. The job collects the data described in the following tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector Framework data collection

Data	CMDB field
hostname	ci.name
physical memory	ci.ram
hardware vendor	ci.manufacturer
hardware model	ci.model_id
CPU name	ci.cpu_type
OS domain	ci.os_domain
OS name	ci.os
OS version	ci.os_version
IPv4 address	ci.ip_address
Default gateway	ci.default_gateway
Serial number	ci.serial_number
DNS Name	ci.name
FQDN	ci.fqdn
TCP Connections	cmdb_tcp
Running Processes	cmdb_running_process
Serial numbers **Note:** When collecting data from Linux hosts, `dmidecode` must be configured with `sudo` permissions to collect all serial numbers.	cmdb_serial_number

Agent Client Collector for Visibility Content data collection

Data	CMDB field
CPU manufacturer	ci.cpu_manufacturer
CPU speed	ci.cpu_speed
OS arch	ci.os_address_width
CPU count	ci.cpu_count
CPU core count	ci.cpu_core_count
CPU core thread	ci.cpu_core_thread
Is virtual	ci.virtual
Start date (uptime)	ci.start_date
Object ID	ci.object_id
TCP Connections	cmdb_tcp
Running Processes	cmdb_running_process

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Account commands for LocalService, LocalSystem, and gMSA during Windows installation

When installing an agent on a Windows host using a LocalService, LocalSystem, or gMSA account, you must run msi commands to enable the accounts.

Account type commands

Account type	Command
LocalSystem	`msiexec /i C:\acc-installer.msi /qn /norestart ACC_API_KEY=<API key code> ACC_MID=<MID Server URL>/ws/events ACC_ALLOW_LIST=False LOCALUSERNAME="SYSTEM"`
LocalService	`msiexec /i C:\acc-installer.msi /qn /norestart ACC_API_KEY=<API key code> ACC_MID=<MID Server URL>/ws/events ACC_ALLOW_LIST=False LOCALUSERNAME="LOCALSERVICE"`
gMSA	`msiexec /i C:\acc-installer.msi /qn /norestart ACC_API_KEY=<API key code> ACC_MID=<MID Server URL>/ws/events ACC_ALLOW_LIST=False LOCALUSERNAME="DOMAIN\gMSA$"`

Verify API is associated with a live agent

Call the *AgentNowHandler.hasAgent* API to determine if the CI is associated with a live agent. The API gets a sys_id of the CI and returns a true/false result.

The *hasAgent()* method takes a sys_id of the CI as a string parameter and returns **true** when the CI fulfills one of the following conditions:

- The CI is `sn_agent_cmdb_ci_agent` and this agent is alive.
- The CI is a host (extends `cmdb_ci_hardware`) and is associated with a live agent.
- The CI is an application (extends `cmdb_ci_appl`) and has a runs-on relation to a host associated with a live agent.

For example:

```
var agentNow = new sn_agent.AgentNowHandler();
gs.info('Agent assocation with CI? ' +
agentNow.hasAgent('158279505372b30034b8ddeeff7b1270'); // sys_id
 of agent
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration file options

Options available in the `acc.yml` configuration file.

Basic configuration options for acc.yml configuration file

Option	Type	Default	Description	Example
name	String	Agent_<value of the *hostname* command>	Agent name registered on the SN instance can be specified by the user. The result of the 'hostname' command is used as the default.	`name: <name of agent>`
backend-url	List	`wss://127.0.0.1:8800/ws/events`	List of MID Webserver endpoint URLs to communicate with. If communication cannot be configured with the first URL, the system moves to the ensuing URLs until it establishes a connection. Works when `enable-auto-mid-selection=true`	`backend-url: <mid server ip>:<websocket port>`
api-key	String	<None>	API key used by the MID Server to authenticate incoming agent connections. Value is encrypted on initial agent startup.	`api-key: <mid web server api key>`
user	String	admin	Username used for basic authentication. If this parameter is empty, the agent does not start.	`user: "agent-01"`
password	String	admin	Password used for basic authentication.	`password: <secure-password>`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Basic configuration options for acc.yml configuration file (continued)

Option	Type	Default	Description	Example
			Value is encrypted on initial agent startup. If this parameter is empty, the agent does not start.	
log-level	String	Info	Amount of logging to appear in the acc.log file. Values: • Panic • Fatal • Error • Warn • Info • Debug	`log-level: debug`
allow-list	String	/etc/servicenow/agent-client-collector/check-allow-list.json	Path to the JSON file that contains the list of check commands the agent can execute. Comment out this parameter to disable the allow-list. If this parameter is empty, the allow-list is disabled.	`allow-list: /etc/agent/check-allow-list.json`
appl_classification_behavior	list	simple	Indicates whether to enable shell CI creation on the agent. Possible values are: • simple: Indicates that shell CI creation is enabled.	`appl_classification_beha off`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Basic configuration options for acc.yml configuration file (continued)

Option	Type	Default	Description	Example
			• off: Indicates that no shell CIs are created for the application. • full: Indicates that complete Discovery of the application CIs is performed using patterns.	

Advanced configuration options for acc.yml configuration file

Option	Type	Default	Description	Example
verify-plugin-signature	Boolean	true	Verifies the plugin signature before execution. Disable when using self-signed or developmental plugins.	`verify-plugin-signature: true`
insecure-skip-tls-verify	Boolean	true	Determines whether to verify the certificate when connecting to the MID Server. Set as **false** to enforce full TLS verification.	`insecure-skip-tls-verify: true`
enable-auto-mid-selection	Boolean	true	Controls the Auto MID Selection feature to connect to the optimal MID Web Server provided by the instance.	`enable-auto-mid-selection: true`
check-command-prefer-installed	Boolean	false	Indicates the preference of executables provided within ACC plugins or executables available in the host system's PATH variable.	`check-command-prefer-installed: false`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration options for acc.yml configuration file (continued)

Option	Type	Default	Description	Example
			• false = ACC plugins • true = Executables in the host system's PATH variable	
command-log-redaction-regex	String	`'-password[=\s]([^"\s]+)|-[pP]\s([^"\s]+)|--user\s[^":\s]+:([^\s"]+)|-pw\s([^"\s]+)|-u\s[^":\s]+:([^\s"]+)|[^"\s]+/([^\s"]+)@'`	The regex to search for and to remove potentially sensitive user information from ACC logs and payloads.	`'-password[=\s]([^"\s]+)|-[pP]\s([^"\s]+)|--user\s[^":\s]+:([^\s"]+)|-pw\s([^"\s]+)|-u\s[^":\s]+:([^\s"]+)|[^"\s]+/([^\s"]+)@'`
powershell_installed	Boolean	false	Disables powershell command execution on agents.	`powershell_installed: true`
allow-list-global-only	Boolean	false	Set to **true** to enhance security by relying only on the allow list defined in the *`allow-list`* parameter you specify during configuration, ignoring allow lists bundled with ACC plugins.	`allow-list-global-only: false`
disable-assets	Boolean	false	Indicates whether a check running with an asset (plugin) dependency fetches ACC plugins from the ServiceNow® instance, or uses a copy of the plugins in its cache folder.	`disable-assets: false`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration options for acc.yml configuration file (continued)

Option	Type	Default	Description	Example
			When set to **false**, additional assets can be downloaded during check execution. Set to **true** to enhance security and ensure that no new plugins are downloaded during check execution.	
agent-upgrade-url-path	String	`https://install.servicenow.com/glide/distribution/builds/package/app-signed/`	Indicates an alternate web server URL endpoint for downloading ACC installer packages when performing selective upgrade.	`agent-upgrade-url-path: https://<ip address>:<port>/acc_installers`
certificate-rotation-days-out	Integer	28	Indicates the number of days before certificate expiration that an agent attempts to rotate its certificate.	`certificate-rotation-days-out=28`
enable-patterns-on-agent	Boolean	false	Enables gathering details on the applications which run on the Agent Client Collector. This parameter is required only when using the Agent Client Collector for pattern execution.	`enable-patterns-on-agent: true`
keepalive-filter-nics	Boolean	true	Indicates whether Network Interface Controllers	`keepalive-filter-nics: true`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced configuration options for acc.yml configuration file (continued)

Option	Type	Default	Description	Example
			(NICs) are filtered on the agent (**true**) or the MID Server (**false**) during keepalive action.	
keepalive-number_nics_per_ip_type	Integer	1	Indicates the maximum number of Network Interface Controllers (NICs) per IP type (IP4, IP6) sent with a keepalive action. The indicated number is sent for each IP type. For example, if the value is 1, a maximum of 2 NICs are sent (0-1 each for IP4 and IP6). If the value is 4, a maximum of 8 NICs are sent (0-4 each for IP4 and IP6).	`keepalive-number_nics_per_ip_type 4`

Proxy Auto-Configuration (PAC) file parameters

Option	Type	Default	Description	Example
pac-file	String (Required)	"" (empty)	Specifies the location of the PAC file to use for proxy configuration. Can be either: • A local file path (such as: `file:///etc/proxy/proxy.pac`) • A remote URL (such as: `https://proxy.company.com/proxy.pac`)	`pac-file: "https://proxy.company.com/proxy.pac"`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Proxy Auto-Configuration (PAC) file parameters (continued)

Option	Type	Default	Description	Example
pac-cache-ttl	Duration (Optional)	30m (30 minutes)	Determines how long proxy rules from the PAC file are cached in memory. Setting this to 0 disables caching.	`pac-cache-ttl: "1h" # Cache for 1 hour`
pac-refresh-interval	Duration (Optional)	30m (30 minutes)	Specifies how often the agent is to check for updates to the PAC file. Useful when the PAC file is hosted remotely and may be updated periodically.	`pac-refresh-interval: "15m" # Check for updates every 15 minutes`
pac-dial-timeout	Duration (Optional)	30s (30 seconds)	Indicates the amount of time to wait when establishing a connection through a proxy server before timing out.	`pac-dial-timeout: "10s" # 10 second timeout`
pac-reset-on-connect-failure	Boolean (Optional)	true	When set to true, the agent clears the PAC cache and attempts to refresh the PAC file if a proxy connection proxy fails. This helps recover from proxy configuration changes or temporary proxy issues.	`pac-reset-on-connect-failure: true`

Synchronization properties for validating Agent Client Collector plugins

Use the following properties when synchronizing public certificates from the MID Server to the Agent Client Collector.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Certificate synchronization - MID properties

Property	Description
mid.acc.cert_monitor_disable	Ability to disable the monitor from synchronizing the public certificates. Default value: False
mid.acc.cert_monitor_interval	Interval (in seconds) at which the monitor synchronizes public certificates from a configured endpoint URL. Default value: 86400
mid.acc.cert_monitor_target_url	Endpoint URL prefix (minus the filename) from which the monitor synchronizes public certificates. Default value: ServiceNow Install server
mid.acc.cert_monitor_target_filename	Endpoint URL target filename from which the monitor synchronizes public certificates. Default value: SNCertificates.zip

Certificate synchronization - Agent properties

Property	Description
fetch-certificates-from-mid	Status of whether the public certificates are retrieved from the MID Server for plugin asset validation. Default value: True

Agent Client Collector statuses

The following table lists and describes the Agent Client Collector statuses.

Agent Client Collector statuses

Status	Description
Up	Agent is connected and working properly.
Warning	Agent has not been refreshed in over 2 minutes (based on the value of the sn_agent.keep_alive.alive_duration property).
Down	Agent is shut down.
Restarting	Agent restart is in progress.
Upgrading	Agent upgrade is in progress.
Disconnected	Agent's MID Server is down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector statuses (continued)

Status	Description
Registered	The agent has registered successfully with the instance, a certificate was issued, and the instance has not yet received a keepalive.
Registration Failed	The agent failed to register with the instance and was not issued a certificate. See Agent Issues for more info.
Unknown	Agent's MID Server has not sent a keepalive communication for over 10 minutes (based on the value of the sn_agent.keep_alive.disconnected_duration system property).

Agent Client Collector log levels

The following table displays the available log levels which indicate the severity of Agent Client Collector events.

Log level	Description
Panic	Severe errors that cause the service to shut down unexpectedly.
Fatal	Fatal errors that cause the service to shut down (status 0).
Error	Non-fatal service error messages.
Warn	Warning messages indicating potential issues.
Info	Information messages representing service actions.
Debug	Detailed service operation messages which help troubleshoot issues.

Agent Client Collector log rotation parameters

If Agent Client Collector logs get too large, they can drain system resources. To ensure system efficiency, configure parameters in the `acc.yml` file by which to rotate logs out of the system's storage (Windows default location = `C:\ProgramData\ServiceNow\agent-client-collector\config\acc.yml`. Linux default location = `/etc/servicenow/agent-client-collector/acc.yml`).

Log rotation parameters

Parameter	Definition
log-file-and-stdout	Write logs to the `stdout` file. Default: **false**
log-retention-duration	Maximum age, in days, of the log file before it is rotated out of the system's storage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Log rotation parameters (continued)

Parameter	Definition
	Default: **72h**
log-retention-files	Maximum number of log files that can be stored before being rotated out of the system. Default: **3**
log-max-size	Maximum size, in MB, of the log file before it is rotated out of the system's storage. Default: **10**
log-level	The log level to be measured by the logs. Available options are: **Panic, Fatal, Error, Warn, Info, Debug**. Default: **Info** The specified log level represents the lowest level of events displayed in the log. For example, a user who specifies **Error** sees all Error events, as well as Fatal and Panic events.

Agent Client Collector CPU protection thresholds

When an agent meets the configured thresholds specified in the agent's `acc.yml` file, it enters CPU protection mode, either for an individual check or for all checks. Agents in CPU protection mode appear in the agent logs with the syntax `Agent Protection`.

Parameter	Description
agent_cpu_threshold_disabled	To prevent the agent from shutting itself off when CPU memory reaches a certain threshold, set value to **True**. Default=False
cpu_percentage_limit	The percentage of the agent's CPU usage that sends the agent into CPU protection mode. Default=5 **Note:** When working with Windows event and metric checks, set this value to 10.
cpu_protection_behavior	Assign **individual** for the system to disable only the individual check that passed the *cpu_percentage_limit* value. If multiple checks crossed the threshold, the check with the highest CPU is disabled.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description
	Assign **all** to disable all checks and pause data collection. Possible values: all, individual Default: all
repeated_high_cpu_num	The number of consecutive times CPU consumption exceeds the *cpu_percentage_limit* for the agent to enter CPU protection mode. Default=3
monitor_interval_sec	Indicates the frequency, in seconds, that the agent monitor runs to check if the *cpu_percentage_limit* is exceeded. Default=60
proxy_cpu_percentage_limit	The percentage of the agent's CPU usage when the agent is running proxy checks too frequently, that sends the agent into CPU protection mode. Default=80

Golden image structure and modularity

The following table outlines the structure and modularity of the golden image mode for cloning agents, based on the operating system (OS) in use.

Golden image structure by OS

	Linux	macOS	Windows
Trigger type	marker file	plist key	msi property
Fresh install handling	none	supported	supported
Service control	stop and disable	skip load	stop and disable
Identity handling	cache wipe	skip restore	skip restore
Certificate handling	cert wip	skip restore	cert wipe
Runtime state	Cache and database wipe	skip restore	Cache and database wipe

Agent Client Collector health instance scan checks

The Agent Client Collector (ACC) health instance scan suite includes checks that detect anomalies and issues on your instance. Use this reference to identify what each check does and when it applies.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC Health Instance Scan Suite checks

Check name	Description
ACC duplicate agent for host CI	Identifies configuration items (CIs) associated with multiple active agents, indicating potential issues with CI identification.
ACC Error Framework	Verifies whether the ACCErrorManager has been created, which enables associating errors to the relevant ACC scoped app. Detects whether an automation source application, automation error category, and automation error code exist for the Agent Client Collector Framework application.
ACC find modified script includes	Detects modified Agent Client Collector Framework script includes that might have been altered post-upgrade.
ACC skipped upgrades	Reports upgrades that were skipped since the last upgrade. **Note:** Skipped upgrades don't generate an error message.
ACC verify plugin folder hierarchy	Validates the correct creation of ACC Plugin folders and flags duplicate folders that may cause synchronization issues.
ACC active policies count	Counts the number of active published policies. If the number exceeds 300, consolidate the policy to improve performance.
ACC agent to MID Connection	Verifies the connection of 1000 agents per 1-GB Java Virtual Machine (JVM) to confirm the MID Server isn't overloaded. If overloaded, configure additional JVMs or redistribute agents to other MID Servers.
ACC verify sn_agent CI counts match	Confirms that the records in the sn_agent_ci_extended_info and sn_agent_cmdb_ci_agent tables match, verifying data consistency for agents.
ACC ICS Diagnostics Tool	Analyzes the ServiceNow instance and its setup with ITOM Cloud Services.
ACC automation_error_msg cleaner	Verifies that the Agent Issues table cleaner exists and is enabled, ensuring ACC-specific error records are regularly cleaned up and the table remains manageable.
ACC Verify global.ACCInstanceScanUtil exists	Detects that the `global.ACCInstanceScanUtil` file is installed for instances starting with the Xanadu release or later. Without this file,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC Health Instance Scan Suite checks (continued)

Check name	Description
	several ACC Health Instance Scan checks don't run.

Additional information

For details on importing the `global.ACCInstanceScanUtil` script include, see the Script Include ACCInstanceScanUtil [KB1630132] article in the HI Knowledge Base.

For details on running the ACC health instance scan, see Schedule the Agent Client Collector (ACC) health instance scan.

Agent onboarding configuration settings

The settings you can customize when generating an Agent Onboarding installation plan. Toggle settings to activate and enable customization.

Agent onboarding configuration fields

Field	Description	Options
Connection type	Method the agent uses to connect to the instance.	• **Use ICS (MID-less)** • **Use MID Server**
MID Server	MID Servers the agent sends data to. Up to three MID Servers can be selected. Appears only when the Connection type is set as **Use MID Server**.	Drop-down list of available MID Server instances.
Registration key	Authorization key for registering with the ServiceNow instance for ICS (MID-less) connection. Appears only when the Connection type is set as **Use ICS (MID-less)**.	Drop-down list of available registration keys.
API key	Authentication key for connecting to the MID Server. Enabled by default and can't be turned off.	Drop-down list of available API keys. Select **Add new** to create a new key.
Log level/debug mode	Controls the level of details that the agent logs. Keep at **Info** unless you're troubleshooting.	• **Error** • **Warn** • **Info** (default) • **Debug** • **Trace**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent onboarding configuration fields (continued)

Field	Description	Options
CPU thresholds	Caps agent CPU usage to protect endpoint performance. Default: 20% for Digital End-User Experience (DEX), 5% for all others.	Enter an integer
Enable PAC file configuration	URL to be used when connecting the agent through a proxy server.	Enter the Proxy Auto-Configuration (PAC) file URL.
API port	REST API port on the agent.	Enter the port number Default port number: 3031

Agent installation plan steps

After you complete agent onboarding, the system generates an installation plan. The following table describes each step in the plan and the commands or actions required.

Installation plan steps

Step	Description
Instance configured	Confirms the instance is configured and ready. Marked as Completed when the criteria monitored by the check are complete.
MID Server configured	Confirms the MID Server is configured and ready. Marked as Completed when the criteria monitored by the check are complete.
ICS configured	Confirms the ITOM Cloud Service (ICS) is configured and ready. Marked as Completed when the criteria monitored by the check are complete.
Network connectivity verified from endpoint to MID	Runs the indicated command in a terminal on the endpoint to verify network connectivity before installation.
Service Account Configuration	Runs the indicated commands in an elevated Command Prompt on the endpoint. For the "Deny log on locally" setting, use Local Security Policy (`secpol.msc`). Required for Linux and macOS, optional for Windows
Installation packages downloaded for Windows	Runs the indicated command in an elevated Command Prompt or PowerShell, from the directory where you downloaded the .msi file.
Review and update configuration file	After installation, find the agent `acc.yml` configuration file and update it with your registration (for ICS) or API key (for MID Server) and verify that the default field values match your configuration. The default field values are:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Installation plan steps (continued)

Step	Description

	`verify-plugin-signature: true`
	`max-running-checks: 10`
	`disable-secrets: true`
	`disable-api: true`
	`stated-disable: true`
	
	Note: Only the configuration options handled by the onboarding process are used.

Agent Client Collector upgrade properties

System properties that control Agent Client Collector upgrade behavior, including the target version, rate limits, retry logic, and timeouts.

Selective upgrade properties

Set these properties before running any selective upgrade. Navigate to **All** > **System Properties** > **All Properties** and search for each property name.

Selective upgrade properties

Property	Value	Description
sn_agent.auto_upgrade.enab	`true`	Enables the upgrade feature. Set to **true** before running any upgrade, manual or scheduled.
sn_agent.agent_upgrade_versio	Target version, for example `5.0.1`	The version agents upgrade to. If empty, the system uses the current Agent Client Collector Framework application version. Must be a valid version number.

Optional properties

These properties have default values suitable for most environments. Adjust them only when your environment requires different rate limits or timeout thresholds.

Optional upgrade properties

Property	Default	Description
*sn_agent.auto_upgrade.max_upgrades_per_*	1000	Maximum agents upgraded per job run across all MID Servers. Lower this value for a more gradual rollout pace.
*sn_agent.auto_upgrade.max_upgrades_per_*	50	Maximum agents upgraded through a single MID Server per job run. Prevents overloading individual MID Servers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Optional upgrade properties (continued)

Property	Default	Description
sn_agent.auto_upgrade.retry_limit	3	Number of times the system retries a failed upgrade before permanently skipping the agent. After an agent reaches this limit, reset the `failed_upgrade_attempts` field on the Agent Extended Info record to allow retries.
sn_agent.agent_upgrade_wait_time	30	Minutes the system waits for an agent to complete its upgrade before marking it as failed and transitioning the agent to **Warning** status.
sn_agent.upgrade.timeout_minutes	10	Minutes allowed for the initial upgrade validation check on the agent (AgentVerification stage).
sn_agent.min_upgrade_history_per_agent	2	Number of upgrade history records retained per agent. Older records are automatically removed.

Property usage note

The *sn_agent.agent_upgrade_version* property is specifically for overriding the upgrade target version. The *sn_agent.agent_version* property controls package download URLs, version display, and other system behaviors. Use *sn_agent.agent_upgrade_version* when you want to target a specific version for upgrades without affecting other system behavior.

Agent Client Collector upgrade error codes

Error codes generated during Agent Client Collector upgrades, with descriptions and resolution steps.

Upgrade errors appear in **All** > **Agent Client Collector** > **ACC Errors**. Filter by error codes starting with ACC - 50 to see upgrade-specific issues. Additional detail is available in the **Message** column of the Agent Upgrade Histories table. For supported platforms, see Supported platforms for Agent Client Collector auto-upgrade.

Upgrade error codes

Error code	Description	Resolution
ACC-5005	Agent operating system is not supported for auto-upgrade.	The agent can't be auto-upgraded. Manually install the new version on the agent host. For supported platforms, see Supported platforms for Agent Client Collector auto-upgrade.
ACC-5006	Agent version is earlier than the minimum required for auto-upgrade (2.7.0).	Manually upgrade the agent to version 2.7.0 or higher, then retry.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Upgrade error codes (continued)

Error code	Description	Resolution
ACC-5007	Agent is not in **Up** status.	Wait for the agent to reconnect. If the agent remains down, investigate connectivity issues on the agent host.
ACC-5013	Agent-side validation failed (AgentVerification stage).	Check the `acc.log` file on the agent host. Common causes include insufficient permissions, low disk space, and missing dependencies.
ACC-5015	MID Server is not in **Up** status.	Restore the MID Server that this agent connects through, then retry the upgrade.
ACC-5016	Upgrade timed out. The agent did not complete the upgrade within the allowed time.	Connect to the agent host and check whether the agent service is running. Review the upgrade log: • Linux and macOS: `{AGENT_LOG_DIR}/servicenow/agent-client-collector/upgrade.log` • Windows: `%TEMP%\ACC_Logs\ACC_Upgrade.log` If the service is running but the status did not update, restart the agent service.
ACC-5017	Agent has exceeded the maximum retry limit for failed upgrades.	1. Investigate the root cause in the agent logs. 2. After resolving the issue, open the agent record and navigate to the related Agent Extended Info record. 3. Set `failed_upgrade_attempts` to 0 and save. 4. Navigate to ACC Errors and resolve the ACC-5017 error for this agent. 5. The agent is picked up in the next scheduled upgrade run.
ACC-5018	Target upgrade version is invalid.	Check the *`sn_agent.agent_upgrade_version`* property. The value must be a valid version number, for example `5.0.1`.

Supported platforms for Agent Client Collector auto-upgrade

Operating systems and package types supported for Agent Client Collector auto-upgrade, and the minimum agent version required.

The minimum agent version for upgrade eligibility is 2.7.0 on all platforms. Agents on unsupported platforms are skipped with error code ACC-5005.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Supported platforms for auto-upgrade

Operating system	Package type	Requirements
Windows	MSI	Agent must run as the SYSTEM account.
Linux — Red Hat, CentOS, Oracle, Rocky, Alma, CentOS Stream	RPM	Agent must run as root or have full sudo access. For information on configuring sudo access, see Install Agent Client Collector on a Linux system.
Linux — Ubuntu, Debian, GNU/Linux	DEB	Agent must run as root or have full sudo access. For information on configuring sudo access, see Install Agent Client Collector on a Linux system.
macOS — Intel (x64)	PKG	Agent must run as root or have full sudo access. For information on configuring sudo access, see Manually install Agent Client Collector on macOS.
macOS — Apple Silicon (ARM64)	PKG	Architecture is detected automatically. Agent must run as root or have full sudo access. For information on configuring sudo access, see Manually install Agent Client Collector on macOS.

Non-supported platforms

Auto-upgrade is not supported on Linux SuSE, Solaris, or AIX. Agents on these platforms are skipped during upgrade runs and generate error code ACC-5005. To upgrade agents on unsupported platforms, manually install the new version on the agent host.

Required agent permissions

Agents must run with elevated privileges on the host to perform a self-upgrade:

- Windows: SYSTEM account
- Linux: root user or sudoers file configuration with full sudo access
- macOS: root user or sudoers file configuration with full sudo access

MSI installation parameters

The following table describes the MSI parameters used when preparing an agent for installation on a gold image and used with a non-persistent virtual desktop infrastructure (NPVDI) machine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Type	Description
NP_VDI_GOLD_IMAGE	boolean	Indicates a gold image host is enabled for the NPVDI agent. Set to **true**.
INSTANCE_URL	string	ServiceNow instance base URL for REST downloads.
INSTANCE_REST_API_USER	string	Basic-auth username for `fetch_vdi_config` REST API. This is the user created in the Prepare agent deployment on a non-persistent virtual desktop infrastructure machine procedure.
INSTANCE_REST_API_PASS	string	Basic-auth password; passed to the agent as the `ACC_REG_KEY` environment variable.
REGISTRATION_KEY	string	VDI light-registration key (must have `light_registration=true` on the instance). Passed to `acc` as the `ACC_REG_KEY` environment variable.
PAC_FILE	string	PAC file for proxy auto-configuration during downloads. Takes precedence over `HTTPS_PROXY`. Optional.
HTTPS_PROXY	string	Static HTTPS proxy URL for downloads (ignored if `PAC_FILE` is set). Applied to REST API calls only. Optional.
GOLD_IMAGE_SKIP_PROXY	boolean	Set to **true** to bypass PAC/proxy during configuration downloads and connect directly. Optional.
CUSTOM_PLUGIN_CERT_FOLDER	string	Directory containing custom plugin-signing certificates; copied into the agent config directory before download.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Type	Description
		Optional.

Agent Client Collector Monitoring reference

Reference topics provide additional information about Agent Client Collector Monitoring checks and policies.

Agent Client Collector Monitoring Footprint

The following table describes Agent Client Collector resource consumption. We supply performance testing values to specific customers upon request.

Parameters	Linux	Windows
Machine spec	Cent OS 8, 4 CPU Core, 8GB Memory	Windows Server 2019, 2 CPU, 4GB Memory
Network - (Agent > MID)	10.64 KB/s	350-400 byte/s
Network - (MID > Agent)	48 byte/s	5-15 byte/s
Average CPU Utilization	~1%	~3.89%
Max CPU per check	10 to 15%	15 to 20%
Memory overhead (avg)	~13.88 MB	~31.15 MB

Agent Client Collector Monitoring default checks and policies

Agent Client Collector Monitoring provides various default checks and policies.

Provided checks

The following tables list the default checks. You can edit these checks as needed. For more information about a specific check, see its Check Definitions page, as described in Checks and policies.

Linux

Type	Name	Description
Event	os.linux.check-process	Counts the number of running processes. You can block event creation for this check when no process is detected. For details, see Block event creation for non-existent entities.
Event	os.linux.check-system-cpu	Checks the CPU usage percentage.
Event	os.linux.check-system-cpu-load	Checks the CPU data load percentage against the configured thresholds.
Event	os.linux.check-system-disk-usage	Checks the file system's mount points and metrics.
Event	os.linux.check-system-memory-percentage	Checks the memory usage percentage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux (continued)

Type	Name	Description
Event	os.linux.check-system-memory-swap	Checks the memory swap usage.
Event	os.linux.check-threads-count	Counts the number of threads running on the system. Sends an alert if that number is greater than the Warning or Critical threshold values.
Metric	os.linux.metrics-memory-vmstat	Collects vmstat memory metrics.
Metric	os.linux.metrics-process-status	Collects process status metrics provided to the metrics parameter.
Metric	os.linux.metrics-system-cpu	Collects CPU metrics, including cores.
Metric	os.linux.metrics-system-cpu-load	Uses uptime to collect CPU load metrics.
Metric	os.linux.metrics-system-cpu-percentage	Collects CPU usage percentage metrics, including average CPU utilization percentage.
Metric	os.linux.metrics-system-disk	Collects disk metrics.
Metric	os.linux.metrics-system-disk-capacity	Collects disk capacity metrics.
Metric	os.linux.metrics-system-disk-usage	Collects disk usage metrics, providing total, used, and available disk memory, as well as the disk's percentage of memory used.
Metric	os.linux.metrics-system-memory	Collects memory metrics, providing data in KB.
Metric	os.linux.metrics-system-memory-percent	Collects memory percentage metrics.
Metric	os.linux.metrics-system-uptime	Collects system uptime in seconds.
Metric	os.linux.metrics-reboot-count-today	Collects the number of reboots performed today.

Linux Log Monitoring

Type	Name	Description
Event	util.check-logs	Enables monitoring log files owned by a regular user. You can block event creation for this check when no process is detected. For details, see Block event creation for non-existent entities.
Event	util.check-logs-sudo	Enables monitoring log files owned by a root user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux Log Monitoring (continued)

Type	Name	Description
		You can block event creation for this check when no process is detected. For details, see Block event creation for non-existent entities.

Linux Network Monitoring

Type	Name	Description
Metric	os.linux.metrics-network-interface	Retrieves all network interface related metrics for Linux servers.
Metric	os.linux.metrics-netstat.tcp	Retrieves metrics on TCP socket states from netstat. Useful on high-traffic web or proxy servers with large numbers of short-lived TCP connections coming and going.

Windows Log Monitoring

Type	Name	Description
Event	os.windows.check-log	Enables monitoring log files on a Windows server.

Windows

Type	Name	Description
Event	os.windows.check-disk-name	Takes the storage drive name as input and verifies if the drive is present. Returns a CRITICAL\WARNING\OK value based on the parameter provided.
Event	os.windows.check-event-log	Measures the Windows event log against parameter thresholds and returns a CRITICAL\WARNING\OK event.
Event	os.windows.check-event-log-count	Measures the Windows event log against parameter thresholds and returns a CRITICAL\WARNING\OK event. Provides information on the number of events that have occurred within a specified duration for a single log file and a single ID. Also indicates the filters to be applied to retrieve events for a specific single-valued windows event level and provider name. Retrieving events from multiple log files is not supported. The number of events is provided, without details of each and every event.
Event	os.windows.check-event-log-details	Collects and filters Windows event logs based on the `duration_hour`, `event_log_level` and `log_file` values. Retrieves and filters Windows event logs according to the provided parameters. It returns details about the events

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows (continued)

Type	Name	Description
		with CRITICAL, WARNING, or OK status, based on the specified severity level. Supported on Windows version 7 and higher.
Event	os.windows.check-env-variables	Checks the environmental variables using a regular expression and returns either a *WARNING* or *OK* value. If a new system variable is created after agent installation, you must restart the agent. This check can access only those user variables that are associated with the current user (used during agent installation).
Event	os.windows.check-file-hashcode-update	Takes the file's path and MD5 hashcode as an input and verifies whether file content has been modified. Returns a *CRITICAL*, *WARNING*, or *OK* value. Read permissions are required on the monitored file.
Event	os.windows.check-file-update	Takes the file's path and interval as an input and verifies whether file content has been modified. Returns a *CRITICAL*, *WARNING*, or *OK* value. Read permissions are required on the monitored file.
Event	os.windows.check-modules	Verifies whether the list of modules is present.
Event	os.windows.check-processor-queue-length	Checks Processor Queue Length.
Event	os.windows.check-system-cpu-load	Checks the CPU load by using the *typeperf* command.
Event	os.windows.check-system-disk	Checks disk usage under specific parameters.
Event	os.windows.check-system-memory	Collects RAM usage.
Event	os.windows.check-system-patch	Verifies system patch installation. Returns either a *WARNING* or *OK* value.
Event	os.windows.check-system-process	Checks the user's task list to find running processes.
Event	os.windows.check-user-account	Takes the list of user names as an input and verifies whether the user account is active. Returns a *CRITICAL*, *WARNING*, or *OK* value.
Event	os.windows.check-windows-directory	Checks if a directory exists.
Event	os.windows.check-windows-pagefile	Collects page file usage and compares it to the Warning and Critical thresholds.
Event	os.windows.check-windows-process	Checks if a user input process is running.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows (continued)

Type	Name	Description
Event	util.check-windows-service	Checks if a user-supplied service is running on Windows. Only the RUNNING status is treated as OK. Any other status generates either a WARNING or CRITICAL event, depending on your configuration. You can block event creation for this check when no process is detected. For details, see Block event creation for non-existent entities.
Metric	os.windows.metrics-process-status	Retrieves the number of running instances, the percentage of CPU utilization, and the memory usage (in kilobytes) of the specified Windows process.
Metric	os.windows.metrics-processor-queue	Collects Processor Queue Length.
Metric	os.windows.metrics-system-cpu-load	Collects the average CPU load per second.
Metric	os.windows.metrics-system-disk-usage	Collects disk usage metrics.
Metric	os.windows.metrics-system-memory-percent	Collects RAM usage per second, providing free physical, free virtual, total physical and total virtual memory percentage.
Metric	os.windows.metrics-system-network	Collects network metrics.
Metric	os.windows.metrics-system-uptime	Collects system uptime.
Metric	util.metrics-windows-perf-counters	(No description)

MID Server

Type	Name	Description
Event	app.mid.check-active	Checks if the MID Server is active, based on the occurrence of LogStatusMonitor in the MID log.
Events	app.mid.check-logs-mid	Checks the MID logs for errors.
Metric	app.mid.metrics-stats	Collects MID Server information from the log and returns metrics.

HAProxy monitoring

Type	Name	Description
Event	check-haproxy	Retrieves the HAProxy load balancer details to determine whether it is running.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HAProxy monitoring (continued)

Type	Name	Description
Metric	metrics-check-haproxy	Retrieves the HAProxy load balancer metrics to determine whether it is running.

HTTP (to be run in proxy mode)

Type	Name	Description
Event	util.check-http-response	Takes a URL or a combination of host/path/port/ssl, and checks for a 200 response that matches a pattern, if provided. Can use client certificates.
Event	util.check-http-follow-redirect	Checks that the URL doesn't exceed the maximum number of redirects.
Metric	util.metrics-http-curl	Retrieves HTTP endpoint metrics using curl.

Additional HTTP checks (no policy)

Type	Name	Description
Event	util.check-http-response-code	Raises a critical event if the URL response code does not match the regex.

Active Directory Metrics

Type	Check	Description	Command
Metric	metrics.windows.active-directory	Collects Active Directory metrics on the Windows host. Usage: winchecks metric-windows-active-dir -s hostname	`winchecks metric-windows-active-dir {{if .labels.params_scheme}} -S {{.labels.params_scheme}} {{end}}`

Apache

Type	Event/Metric	Description	Minimum version
Metric	app.apache.metrics-apache	Retrieves Apache HTTP server metrics.	2.4.x

Apache Kafka

Type	Event/Metric	Description	Minimum version
Event	kafka.check-zookeeper-status	Raises a critical event if the hosted Kafka Zookeeper is down.	ZooKeeper 3.1.2

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Kafka (continued)

Type	Event/Metric	Description	Minimum version
Event	kafka.check-topic-replicas	Raises a critical event if any topic has partitions with unknown replicas.	ZooKeeper 3.1.2
Event	kafka.check-topic-replication-factor	Raises a critical event if replication factor of at least one topic is above or below provided replication factor parameter.	ZooKeeper 3.1.2
Event	kafka.check-topic-leader	Raises a critical event if any topic has partitions with unknown leaders or unpreferred replica as leader.	ZooKeeper 3.1.2
Event	kafka.check-topic-partitions	Raises a critical event if the number of partitions for a topic is less the min_partitions param.	ZooKeeper 3.1.2
Event	kafka.check-broker-status	Raises a critical event if Kafka Broker on the host is down.	Kafka 0.10.1.0
Metric	kafka.metrics.broker	Collects Kafka Broker Metrics from the host.	Kafka 0.10.1.0
Metric	kafka.metrics.zookeeper	Collects Zookeeper Metrics from the host.	ZooKeeper 3.5.0-alpha

Azure k8s Metrics

Type	Event/Metric	Description	Command
Metric	azure.metrics-k8s-service	Provides pod, nodes, and cluster level metrics from all Azure Kubernetes services in metric form. Usage: metrics-azure-k8s-service.rb Command example: `metrics-azure-k8s-service.rb -S {{.labels.params_subscription_id}} -r westus2`	`metrics-azure-k8s-service.rb -r {{.labels.params_ci_region}} {{-s {{if.labels.params_subscrip -S {{.labels.params_subscripti {{end}}`

Azure VM Metrics

Type	Event/Metric	Description	Command
Metric	azure.metrics-vm-cpu	Provides CPU-related metrics of each VM instance in the Azure region. Command example: `./metrics-`	`metrics-azure-vm.rb -r {{.labels.params_ci_region} {{if .labels.params_subscri -S {{.labels.params_subscripti`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure VM Metrics (continued)

Type	Event/Metric	Description	Command
		`azure-vm.rb -S {{.labels.params -T cpu`	`{{end}} -T {{.labels.params_type}}`
Metric	azure.metrics-vm-disk	Provides disk-related metrics of each VM instance in the Azure region. Command example: `./metrics-azure-vm.rb -S {{.labels.params_subscription_id}} -T disk`	`metrics-azure-vm.rb -r {{.labels.params_ci_region} {{if .labels.params_subscri -S {{.labels.params_subscripti {{end}} -T {{.labels.params_type}}`
Metric	azure.metrics-vm-memory	Provides memory-related metrics of each VM instance in the Azure region. Command example: `./metrics-azure-vm.rb -S {{.labels.params_subscription_id}} -T memory`	`metrics-azure-vm.rb -r {{.labels.params_ci_region} {{if .labels.params_subscri -S {{.labels.params_subscripti {{end}} -T {{.labels.params_type}}`
Metric	azure.metrics-vm-network	Provides network-related metrics of each VM instance in the Azure region. Command example: `./metrics-azure-vm.rb -S {{.labels.params_subscription_id}} -T network`	`metrics-azure-vm.rb -r {{.labels.params_ci_region} {{if .labels.params_subscri -S {{.labels.params_subscripti {{end}} -T {{.labels.params_type}}`
Metric	azure.metrics-vm-flows	Provides flows-related metrics of each VM instance in the Azure region. Command example: `./metrics-azure-vm.rb -S {{.labels.params_subscription_id}} -T flows`	`metrics-azure-vm.rb -r {{.labels.params_ci_region} {{if .labels.params_subscri -S {{.labels.params_subscripti {{end}} -T {{.labels.params_type}}`

GCP GCE metrics

Type	Event/Metric	Description
Metric	gcp.metrics.GCE	Retrieves metrics from GCP for Google Compute Engine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP GCE metrics (continued)

Type	Event/Metric	Description
Metric	glassfish.metric-server-requests	Provides http-request metrics for all applications deployed on the Glassfish server.
Metric	glassfish.metric-server-resources	Provides metrics of the resources deployed on the Glassfish server.
Metric	glassfish.metric-transaction-service	Provides metrics for all transactions on applications deployed on the Glassfish server.
Metric	glassfish.metric-deployed-apps	Provides metric output of the applications deployed on the Glassfish server.
Metric	glassfish.metric-deployment-lifecycle	Provides lifecycle metrics applications deployed on the Glassfish server.

Glassfish

Type	Name	Description
Metric	glassfish.metric-jvm-level	Retrieves JVM metrics on the Glassfish server.
Metric	glassfish.metric-server-requests	Provides http-request metrics for all applications deployed on the Glassfish server.
Metric	glassfish.metric-server-resources	Provides metrics of the resources deployed on the Glassfish server.
Metric	glassfish.metric-transaction-service	Provides metrics for all transactions on applications deployed on the Glassfish server.
Metric	glassfish.metric-deployed-apps	Provides metric output of the applications deployed on the Glassfish server.
Metric	glassfish.metric-deployment-lifecycle	Provides lifecycle metrics applications deployed on the Glassfish server.

IBM WebSphere AS

Type	Event/Metric	Description	Minimum version
Metric	app.websphere.metrics-websphere	Returns IBM WebSphere application server metrics.	9

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IBM WebSphere AS (continued)

Type	Event/Metric	Description	Minimum version
		Ensure that the user running the agent has sufficient permissions to access the required files in the `IBM Websphere` AS installation directories.	

IIS

Type	Event/Metric	Description	Minimum version
Event	app.iis.check-iis-alive	Checks if IIS is running.	-
Event	app.iis.check-iis-current-connections	Checks the number of current connections to the IIS server and their status, depending on the input. By default, all sites are connected to the IIS server.	-
Metric	app.iis.metrics-iis	Returns IIS metrics.	-

JBoss/WildFly

Type	Event/Metric	Description	Minimum version
Metric	app.jboss.metrics-jboss	Retrieves JBoss/WildFly server metrics. Ensure that the user running the agent has sufficient permissions to access the required files in the`JBoss/WildFly` installation directories.	16, 17

Kubernetes

Type	Event/Metric	Description	Command
Event	k8s.check-kube-apiserver-available	Checks if Kubernetes API server is up.	`check-kube-apiserver-available.rb {{if .labels.params_api_server}} -s {{.labels.params_api_server}} {{end}}`
Event	k8s.check-kube-nodes-ready	Checks whether Kubernetes nodes are in ready-to-use state.	`check-kube-nodes-ready.rb {{if .labels.params_api_server}} -s {{.labels.params_api_server}} {{end}}`
Event	k8s.check-kube-pods-pending	Measures the threshold by which pods are in the Pending state. Determine the threshold by setting the	`check-kube-pods-pending.rb {{if .labels.params_api_server}} -s {{.labels.params_api_server}} {{end}}`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Kubernetes (continued)

Type	Event/Metric	Description	Command
		pending_timeout check parameter value (default = 300 seconds).	
Event	k8s.check-kube-pods-restarting	Checks the pods that are restarting. Determine the threshold by setting the check parameter value (default = 10).	`check-kube-pods-restarting.rb {{if .labels.params_api_server}} -s {{.labels.params_api_server}} {{end}} {{if .labels.params_restart}} -r {{.labels.params_restart}} {{end}}`
Event	k8s.check-kube-pods-running	Checks the pods that are in the running state. Configure the threshold in the *not_ready_time* check parameter, in seconds (default = 0).	`check-kube-pods-running.rb {{if .labels.params_api_server}} -s {{.labels.params_api_server}} {{end}} {{if .labels.params_not_ready_time}} -time {{.labels.params_not_ready_time}} {{end}}`
Event	k8s.check-kube-pods-runtime	Checks whether pods are running longer than expected. Configure threshold in the check parameter (in seconds). You can configure both WARNING and CRITICAL thresholds.	check-kube-pods-runtime.rb {{if .labels.params_critical}} -c {{.labels.params_critical}} {{end}} {{if .labels.params_warning}} -w {{.labels.params_warning}} {{end}} {{if .labels.params_api_server}} -s {{.labels.params_api_server}} {{end}}
Event	k8s.check-kube-service-available	Checks whether Kubernetes services are up and running. Add the list of services to be monitored in the services' *check_param* field as comma-separated values (without spaces). If you enter services that are not valid, the check fails.	`check-kube-pods-runtime.rb {{if .labels.params_critical}} -c {{.labels.params_critical}} {{end}} {{if .labels.params_warning}} -w {{.labels.params_warning}} {{end}} {{if .labels.params_api_server}} -s {{.labels.params_api_server}} {{end}}`
Metric	k8s.metrics-pods	Provides a pod count from all of the exposed services.	`metrics-pods.rb {{if .labels.params_api_server}} -s {{.labels.params__api_server }} {{end}}`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MsSQL

Type	Event/Metric	Description	Minimum version
Metric	app.msssql.metrics-mssql	Retrieves MSSQL database metrics.	2016
Metric	app.metrics-mssql-query	Generic MSSQL query class which accepts a query and dumps metrics.	2016

MySQL

Type	Event/Metric	Description	Minimum version
Event	app.mysql.check-mysql-alive	Checks if the MySQL database is running.	5.7, 8
Event	app.mysql.check-mysql-threads	Checks the number of running threads on the MySQL database. Depending on the input, it determines the status: OK, Warning, or Critical.	5.7, 8
Event	util.check-mysql-query	Checks the length of a MySQL query result set. Ensure that you use the native_password authentication method on the MySQL server when utilizing this check.	5.7, 8
Metric	app.mysql.metrics-mysql	Returns MySQL database metrics.	5.7, 8
Metric	app.mysql.metrics-mysql-processes	Returns MySQL database process metrics.	5.7, 8
Metric	util.metrics-mysql-query	Creates a formatted metric for the length of a MySQL query result set.	5.7, 8

MongoDB

Type	Event/Metric	Description	Minimum version
Event	app.mongodb.check-mongodb-alive	Monitors whether the MongoDB server is alive and creating alerts for the MongoDB server health status.	
Event	app.mongodb.check-mongodb-metrics	Creates alerts for any of the MongoDB metrics, based on the threshold limit.	
Metric	app.mongodb.metrics-mongodb	Returns metrics of the MongoDB server and all databases.	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Oracle

Type	Event/Metric	Description	Minimum version
Event	app.oracle.check-oracle-alive	Checks if the Oracle database is running. When using this check, add the OS user (servicenow) to the **oinstall** group to ensure that the user has execute permissions for the sqlplus tool. Ensure that the CREATE SESSION database privilege is activated.	12c, 18c
Metric	app.oracle.metrics-oracle	Returns Oracle database metrics. When using this check, add the OS user (servicenow) to the **oinstall** group to ensure that the user has execute permissions for the sqlplus tool. Ensure that the CREATE SESSION and SELECT database privileges are activated.	12c, 18c
Metric	util.metrics-oracle-query	Generic Oracle query class that accepts a query and dumps metrics. When using this check, add the OS user (servicenow) to the **oinstall** group to ensure that the user has execute permissions for the sqlplus tool. Ensure that the CREATE SESSION and SELECT database privileges are activated.	12c, 18c

Oracle WebLogic AS

Type	Event/Metric	Description	Minimum version
Metric	app.weblogic.metrics-weblogic	Returns Oracle WebLogic application server metrics. Ensure that the user running the agent has sufficient permissions to access the required files in the `WebLogic` installation directories.	12cR2

Tomcat

Type	Event/Metric	Description	Minimum version
Metric	app.tomcat.metrics-tomcat	Returns Apache Tomcat server metrics.	8.5.27, 9.x

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

App

Type	Name	Description
Event	util.check-snmp	SNMP check that compares the metric value to the default. Requires input such as server, port, SNMP community, and limits.
Metric	util.metrics-snmp	Collects metrics from an SNMP OID value.
Metric	util.metrics-snmp-bulk-sn	Collects metrics from an SNMP Table OID and returns the result as a table.

Service enhancements (can be associated with any policy)

Type	Event/Metric	Description	Command
Metric (Windows)	os.windows.metrics-service-status	Collects Windows service status, providing metric of number of services running. Only the RUNNING status is treated as OK. Any other status generates either a WARNING or CRITICAL event, depending on your configuration.	`winchecks metric-windows-service-status -S Wcmsvc --scheme hostname.proc`

Varnish

Type	Name	Description
Metric	app.varnish.metrics-varnish	Collects Varnish metrics from the host.

vSphere

Type	Name	Description
Metric	vsphere.metric-VirtualMachine	Retrieves metrics for the vSphere Virtual Machine.
Metric	vsphere.metric-Datastore	Retrieves metrics for the vSphere Datastore.
Metric	vsphere.metric-Datacenter	Retrieves metrics for the vSphere Datacenter.
Metric	vsphere.metric-ESX_Server	Retrieves metrics for the vSphere ESX Server or HostSystem

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Zscaler

Type	Name	Description
Metric	zscaler-monitoring-check	Verifies whether the Zscaler app is running properly.
Metric	zscaler-remediation-check	Shuts down and restarts the Zscaler app, and creates an incident. Runs only when the zscaler-monitoring-check fails.

CPU enhancements

Type	Event/Metric	Description	Command
Metric	os.windows.metrics-system-cpu	Collects CPU core metrics.	`winchecks metric-windows-cpu --scheme hostname.proc`

Disk enhancements - Linux events (can be associated with any policy)

Event/Metric	Description	Command
os.linux.check-disk-ioTimeWeighted	Measures the `disk ioTimeWeighted` value, returning a **CRITICAL \WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-ioTimeWeighted.rb`
os.linux.check-disk-sectorsWritten	Measures the total number of sectors written successfully, returning a **CRITICAL \WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-sectorsWritten.rb`
os.linux.check-disk-sectorsRead	Measures the disk's total number of sectors read successfully, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-sectorsRead.rb`
os.linux.check-disk-readTime	Measures the total number of milliseconds spent by all read- against thresholds, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-readTime.rb`
os.linux.check-disk-writes	Measures the total number of writes completed successfully, returning a **CRITICAL \WARNING\OK** message,	`check-disk-writes.rb`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk enhancements - Linux events (can be associated with any policy) (continued)

Event/Metric	Description	Command
	according to the thresholds specified in the parameters.	
os.linux.check-disk-reads	Measures the total number of reads completed successfully, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-reads.rb`
os.linux.check-disk-writesMerged	Measures the `disk writesMerged` value, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-writesMerged.rb`
os.linux.check-disk-readsMerged	Measures the `disk readsMerged` value, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-readsMerged.rb`
os.linux.check-disk-writeTime	Measures the total number of milliseconds spent by all write operations, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-writeTime.rb`
os.linux.check-disk-ioInProgress	Measures the `disk ioInProgress` value, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-ioInProgress.rb`
os.linux.check-disk-iotime	Measures the `disk ioTime` value, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-disk-ioTime.rb`

Disk enhancements - Windows events (can be associated with any policy)

Event/Metric	Description	Command
os.windows.check-disk-AvgSecWrite	Measures average disk writes per second, returning a **CRITICAL\WARNING\OK** message, according to the specified in the parameters.	`winchecks check-windows-AvgDisksec-write -w 1 -c 0`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disk enhancements - Windows events (can be associated with any policy) (continued)

Event/Metric	Description	Command
os.windows.check-disk-ReadBytes-sec	Measures the `disk ReadBytes` value per second, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`winchecks check-windows-DiskReadBytes-sec -w 1 -c 0`
os.windows.check-disk-WriteBytes-sec	Measures the `disk WriteBytes` value per second, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`winchecks check-windows-DiskWriteBytes-sec -w 1 -c 0`
os.windows.check-disk-AvgSecRead	Measures the average disk reads per second, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`winchecks check-windows-AvgDisksec-read -w 1 -c 0`

Disk enhancements - Windows metrics (can be associated with any policy)

Event/Metric	Description	Command
os.windows.metrics-system-disk	Collects the following disk metrics: • AvgDiskSecPerRead • AvgDiskSecPerWrite • DiskReadBytesPerSec • DiskWriteBytesPerSec	`winchecks metric-windows-disk`

Memory enhancements (can be associated with any policy)

Type	Event/Metric	Description	Command
Event (Linux)	os.linux.check-free-physical-memory	Measures the free physical memory percentage of the Linux system, returning a **CRITICAL\WARNING\OK** event message, according to the thresholds specified in the parameters.	`check-free-physical-memory -w 10 -c 5`
Event (Linux)	os.linux.check-free-virtual-memory	Measures the free virtual memory percentage of	`check-free-virtual-memory -w 10 -c 5`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Memory enhancements (can be associated with any policy) (continued)

Type	Event/Metric	Description	Command
		the Linux system, returning a **CRITICAL\WARNING\OK** event message, according to the thresholds specified in the parameters.	
Event (Windows)	os.windows.check-free-physical-memory	Measures the free physical memory percentage of the Windows system, returning a **CRITICAL\WARNING\OK** event message, according to the thresholds specified in the parameters.	`winchecks check-windows-free-physfical-memory -w 10 -c 5`
Event (Windows)	os.windows.check-free-virtual-memory	Measures the free virtual memory percentage of the Windows system, returning a **CRITICAL\WARNING\OK** event message, according to the thresholds specified in the parameters.	`winchecks check-windows-free-physical-memory -w 10 -c 5`
Metric (Windows)	os.windows.metrics-system-memory	Collects the following memory metrics: • FreePhysicalMemory • TotalPhysicalMemory • FreeVirtualMemory • TotalVirtualMemorySize • AvailableMemory • TotalVisibleMemorySize	`winchecks metric-windows-memory --scheme hostname.proc`

Process enhancements (can be associated with any policy)

Type	Event/Metric	Description	Command
Event (Linux)	os.linux.check-process-cpu	Measures the Linux process CPU usage, returning a **CRITICAL\WARNING\OK** message, according to the thresholds	`check-process-cpu.rb -p acc -c 95 -w 85`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Process enhancements (can be associated with any policy) (continued)

Type	Event/Metric	Description	Command
		specified in the parameters.	
Event (Linux)	os.linux.check-process-memory	Measures the Linux process memory usage, returning a **CRITICAL\WARNING\OK** message, according to the thresholds specified in the parameters.	`check-process-memory.rb -p acc -c 95 -w 85`
Metric (Linux)	os.linux.metrics-process-usage	Collects the Linux process status, including the CPU and memory data used by the process.	`metrics-process-usage.rb -p acc -s hostname.proc`
Event (Windows)	os.windows.check-process-cpu	Measures the Windows process CPU usage, returning a **CRITICAL\WARNING\OK** message according to the thresholds specified in the parameters.	`winchecks check-windows-process-cpu-p acc -c 95 -w 85`
Event (Windows)	os.windows.check-process-memory	Measures the Windows process memory usage, returning a **CRITICAL\WARNING\OK** message according to the thresholds specified in the parameters.	`winchecks check-windows-process-memory-p acc -c 95 -w 85`
Metric (Windows)	os.windows.metrics-process-usage	Collects the Windows process status, including the CPU and memory data used by the process.	`winchecks metric-windows-process-status -n Svchost -s hostname.proc`

Supported Agent policies

This table lists the supported Agent Client Collector for Monitoring policies that the check definitions are associated with. The policies consist of the monitored CIs and the checks that run on them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
Apache Events	Monitors operational Apache HTTP web servers. To enable monitoring if the process is alive, activate the relevant check: • For Linux: os.linux.check-process • For Windows: os.windows.check-windows-process On the **Credentials** tab, specify Basic Auth Credentials.	• os.linux.check-process • os.windows.check-windows-process
Apache Metrics	Monitors operational Apache HTTP web servers. On the **Credentials** tab, specify Basic Auth Credentials.	app.apache.metrics-apache
Apache Kafka Events	Monitors the status of Kafka Broker, Kafka topics, partictions, replicas and Zookeeper.	
Apache Kafka Metrics	Monitors and collects metrics from Kafka Broker and Kafka Zookeeper.	

Azure DB Metrics	Monitors Azure databases on an Azure datacenter.
	
	To enable the policy's checks:
	• Configure a proxy agent.
	• Add the following Azure service principal credentials:
	  ○ Tenant ID
	  ○ Client ID
	  ○ Secret Key
	• Configure the Authentication Method as **Client secret**.
	• Add Azure credentials in the policy.
	• Provide the *subscription_id* value in each check instance.
	
	The default monitored CI Type is cmdb_ci_azure_datacenter, but the binding occurs on the Cloud Database (cmdb_ci_cloud_database).
	For details on configuring a policy, see Create an Agent Client Collector policy.	

Azure k8s Service Metrics	Monitors the Azure CI, *cmdb_ci_azure_datacenter*.	azure.metrics-k8s-service

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
	Binds with the Kubernetes data cluster, *cmdb_ci_kubernetes_cluster*.	For details, see Azure Health Monitoring default checks and policies.
Azure VM Metrics	Monitors the Azure CI, *cmdb_ci_azure_datacenter*. Binds with the virtual machine instance, *cmdb_ci_vm_instance*.	• azure.metrics-vm-cpu • azure.metrics-vm-disk • azure.metrics-vm-memory • azure.metrics-vm-network • azure.metrics-vm-flows For details, see Azure Health Monitoring default checks and policies.
Basic Discovery	Runs basic ServiceNow Discovery on all agents, every 12 hours. This check extracts basic information about the host and its running processes. The host information is used to create the corresponding host in the CMDB, while the running processes create the classified application CIs. You need advanced OS privileges to use the Net stat command to discover ports. The command that you run depends on your OS, as follows: • Linux: `sudo -n netstat -antp` • Windows: `netstat -anop TCP`	check-discovery-basic
Docker Container Metrics	Collects performance metrics for Docker containers. Provide the credentials of users who are Docker user group members or have sudo rights. Provide these credentials on the **Credentials** tab of the policy form.	container.docker.metrics-docker
Docker Engine Metrics	Collects general metrics for the Docker engine. Provides the credentials of a user which is either a member of the Docker user group or has been assigned sudo rights on the **Credentials** tab. To enable running checks through this policy, add the servicenow user to the Docker user	container.docker.metrics-docker-info

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
	group on the machine where the agent is running.	
F5 SNMP Events	Example for using SNMP checks. Because the Agent doesn't run on SNMP devices, provide the following to complete this policy: • A Proxy Agent and credentials. • The host of the Proxy Agent. On the **Credentials** tab, specify the SNMP Community Credentials for connecting to the server.	util.check-snmp
F5 SNMP Metrics	Example for using SNMP metrics. Because the Agent doesn't run on SNMP devices, provide the following to complete this policy: • Proxy Agent and credentials. • The host of the Proxy Agent. On the**Credentials** tab, specify the SNMP Community Credentials for connecting to the server.	• util.metrics-snmp • util.metrics-snmp-bulk-sn
Glassfish Metrics	Runs metric checks on the Glassfish server. Glassfish metrics are supported in ITOM Agent Client Collector version 2.3.0 - August 2022, available from the ServiceNow Store.	
HAProxy Events	Retrieves the HAProxy load balancer details to determine whether it is running. HAProxy events are supported in ITOM Agent Client Collector version 2.3.0 - August 2022, available from the ServiceNow Store.	
HAProxy Metrics	Retrieves the HAProxy load balancer metrics to determine whether it is running. HAProxy metrics are supported in ITOM Agent Client Collector version 2.3.0 - August 2022, available from the ServiceNow Store.	
HTTP Entry Points Events	Runs the util.check-http-follow-redirect check on all HTTP entry points that belong to application business services. The policy filter chooses the services. The services define the entry points. • To enable this check, there is a new database view which combines the application services table with the entry points table.	util.check-http-follow-redirect

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
	• You can filter the services and receive the entry points. • When you receive an alert, it is automatically bound to the entry points. This alert affects the impact calculation, because entry points are considered. • This policy requires a Proxy Agent.	
HTTP Entry Points Metrics	For example, when the alert is bound to the entry point, the service map identifies an entry point problem. • The metric binds to the entry point CI. In previous releases, the metric was bound to the service and the entry point was the resource. • This policy requires a Proxy Agent.	util.metrics-http-curl
IIS Events	Monitors operational IIS servers.	• app.iis.check-iis-alive • app.iis.check-iis-current-connections
IIS Metrics	Monitors operational IIS servers.	app.iis.metrics-iis
JBoss/WildFly Application Server Events	Monitoring operational JBoss/WildFly application servers. To enable monitoring if the JBoss process is alive, activate os.linux.check-process. The *jboss_home* parameter in the check instance is specified as the location where JBoss is installed.	os.linux.check-process
JBoss/WildFly Application Server Metrics	Monitors operational JBoss/WildFly application servers.	app.jboss.metrics-jboss
Kubernetes	Monitors health of Kubernetes infrastructure, such as clusters, nodes, pod status, API, and service availability. When running Kubernetes checks: • Proxy agent must be enabled. • Enter Kubernetes credentials, either a username/password combination or a bearer token (but not both). • To activate and publish the policy, configure the *api_server* value in the	• k8s.check-kube-apiserver-available • k8s.check-kube-nodes-ready • k8s.check-kube-pods-pending • k8s.check-kube-pods-restarting • k8s.check-kube-pods-running

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
	check parameter definitions of each check, as one of the following: ◦ `https://127.0.0.1:<portnumber>/api` (If the API server and agent are running on the same server) ◦ `https://<api_server_url>/api` (Whether the API server and agent are running on the same or different servers). To use the localhost as your api_server, run the following command: `"kubectl proxy -port=8081 &"` • Run Discovery to discover all Kubernetes CIs. For details, see Kubernetes discovery.	• k8s.check-kube-pods-runtime • k8s.check-kube-service-available • k8s.metrics-pods
Linux log monitoring	Monitors logs on operational Linux servers.	• util.check-logs • util.check-logs-sudo
Linux OS Events	Runs Linux checks on operational Linux servers.	• os.linux.check-process • os.linux.check-system-cpu • os.linux.check-system-cpu-load • os.linux.check-system-disk-usage • os.linux.check-system-memory-percentage • os.linux.check-system-memory-swap • os.linux.check-threads-count
Linux OS Metrics	Collects Linux metrics on operational Linux servers.	• os.linux.metrics-memory-vmstat • os.linux.metrics-process-status • os.linux.metrics-system-cpu • os.linux.metrics-system-cpu-load • os.linux.metrics-system-disk • os.linux.metrics-system-disk-capacity • os.linux.metrics-system-disk-usage

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
		• os.linux.metrics-system-memory • os.linux.metrics-system-memory-percentage • os.linux.metrics.reboot_count_today
MSSQL Events	Monitors the operational MSSQL server process. Runs os.windows.check-windows-process to monitor when the server process is alive.	os.windows.check-windows-process
MSSQL Metrics	Monitors operational MSSQL server process. Runs app.mssql.metrics-mssql to gather metrics.	app.mssql.metrics-mssql
MySQL DB Events	Monitors operational MySQL instances. On the **Credentials** tab, specify Basic Auth Credentials for connecting to the database server. **Note:** Utility checks, prefixed with util, have required parameters for which you must specify a value. These parameters have no default values.	• app.mysql.check-mysql-alive • app.mysql.check-mysql-threads • util.check-mysql-query
MySQL DB Metrics	Monitors operational MySQL instances. On the **Credentials** tab, specify Basic Auth Credentials for connecting to the database server. **Note:** Utility checks, prefixed with util, have required parameters for which you must specify a value. These parameters have no default values.	• app.mysql.metrics-mysql • app.mysql.metrics-mysql-processes • util.metrics-mysql-query
Network ping metrics	Monitors the ip address discovered in the ServiceNow instance. The cmdb_ci_hardware monitored CI type is bound to the relevant server: This policy requires a proxy agent.	util.metrics-ping
Oracle DB Events	Monitors operational Oracle Database instances. On the **Credentials** tab, specify Basic Auth Credentials for connecting to the database server. **Note:** Utility checks, prefixed with util, have required parameters for which you must specify a value. These parameters have no default values.	app.oracle.check-oracle-alive

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
Oracle DB Metrics	Monitors operational Oracle Database instances. On the **Credentials** tab, specify Basic Auth Credentials for connecting to the database server. **Note:** Utility checks, prefixed with util, have required parameters for which you must specify a value. These parameters have no default values.	• app.oracle.metrics-oracle • util.metrics-oracle-query
Self-Healing Events	Runs a ping command to another host. Ensure that you run this check in proxy mode.	util.check-ping
Tomcat Events	Monitors operational Apache Tomcat servers. To enable monitoring if the Tomcat process is alive, activate the relevant check: • For Linux: os.linux.check-process • For Windows: util.check-windows-service If the Tomcat service is running under a different name, update the value of the check instance's *service* parameter to the correct value.	• os.linux.check-process • util.check-windows-service
Tomcat Metrics	Monitors operational Apache Tomcat servers. If Tomcat Metrics are running on a port other than the default port (9000), update the value of the check instance's *port* parameter to the correct port value.	app.tomcat.metrics-tomcat
Varnish Metrics	Monitors and collects metrics from Varnish. Varnish metrics are supported in ITOM Agent Client Collector version 2.3.0 - August 2022, available from the ServiceNow Store. You must perform discovery using application fingerprints before running Varnish checks. For details, see Enable running of Varnish checks.	
vSphere Metrics	Monitors and collects metrics from vSphere. vSphere metrics are supported in ITOM Agent Client Collector version 2.3.0 - August 2022, available from the ServiceNow Store. This policy requires a proxy agent.	
WebLogic Application Server Events	Monitors operational Oracle WebLogic application servers. This policy monitors if the WebLogic process running on Linux is alive.	os.linux.check-process

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
	On the **Credentials** tab, specify Basic Auth Credentials for connecting to the application server.	
WebLogic Application Server Metrics	Monitors operational Oracle WebLogic application servers. On the **Credentials** tab, specify Basic Auth Credentials for connecting to the application server. Ensure that the user running the agent has sufficient permissions to access the required files in the `WebLogic` installation directories.	app.weblogic.metrics-weblogic
WebSphere Application Server Events	Monitors operational IBM WebSphere servers. This policy monitors if the WebSphere process running on Linux is alive. On the **Credentials** tab, specify Basic Auth Credentials for connecting to the application server.	os.linux.check-process
WebSphere Application Server Metrics	Monitors operational IBM WebSphere servers. On the **Credentials** tab, specify Basic Auth Credentials for connecting to the application server.	app.websphere.metrics-websphere
Windows Log Monitoring	Monitors logs on operational Windows servers.	os.windows.check-log You can block event creation for this check when no process is detected. For details, see Block event creation for non-existent entities.
Windows OS Events	Runs Windows checks on operational Windows servers.	• os.windows.check-processor-queue-length • os.windows.check-system-cpu-load • os.windows.check-system-disk • os.windows.check-system-memory • os.windows.check-system-process • os.windows.check-windows-directory

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description	Checks
		• os.windows.check-windows-directory • os.windows.check-windows-directory • os.windows.check-windows-pagefile • os.windows.check-windows-process
Windows OS Events - Extended	Runs Windows checks to monitor directories and files on Windows servers.	• check-directory-file-count • check-directory-integrity • check-file-age • check-file-response-time • check-file-size • check-file-space • os.windows.check-directory-space
Windows OS Metrics	Collects Windows metrics on operational Windows servers.	• os.windows.metrics-processor-queue-len • os.windows.metrics-system-cpu-load • os.windows.metrics-system-disk-usage • os.windows.metrics-system-memory-percent • os.windows.metrics-system-network • os.windows.metrics-system-uptime
Zscaler Monitoring Policy	Executes Zscaler monitoring on all of a machine's CIs.	• zscaler-monitoring-check • zscaler-remediation-check

Active Directory metrics

The following table lists the metrics that are gathered as output from Active Directory checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Active Directory metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
active_dir.DRA_Inbound_Bytes_after_compressed/sec			Shows the compressed size of inbound compressed replication data (size after compression, from DSAs in other sites).
active_dir.DRA_Inbound_Bytes_before_compressed/sec			Shows the original size of inbound compressed replication data (size before compression, from DSAs in other sites).
active_dir.DRA_Inbound_Bytes_not_compressed/sec			Shows the number of incoming bytes replicated that were not compressed at the source (that is, from DSAs in the same site).
active_dir.DRA_Inbound_Bytes_Total/sec (featured metric)			Shows the total number of bytes replicated in. This counter is the sum of the number of uncompressed bytes (never compressed) and the number of compressed bytes (after compression).
active_dir.DRA_Inbound_Full_Sync_Objects_Remaining (featured metric)			Shows the number of objects remaining until the full synchronization is completed (when set).
active_dir.DRA_Inbound_Object_Updates_Remaining_in_Packet			Shows the number of object updates received in the current directory replication update packet that have not yet been applied to the local server.
active_dir.DRA_Inbound_Objects/sec			Shows the rate at which replication updates received from

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Active Directory metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			replication partners are applied by the local directory service. This counter excludes changes that are received but not applied (because, for example, the change has already been made). This indicates how much replication update activity is occurring on the server as a result of changes generated on other servers.
active_dir.DRA_Inbound_Objects_Applied/sec			Shows the number of objects received from inbound replication partners that contained no updates that needed to be applied.
active_dir.DRA_Inbound_Objects_Filtered/sec			Shows the number of objects received from neighbors through inbound replication. A neighbor is a domain controller from which the local domain controller replicates locally.
active_dir.DRA_Inbound_Properties_Applied/sec			Shows the number of properties that are updated due to the incoming property's winning the reconciliation logic that determines the final value to be replicated.
active_dir.DRA_Inbound_Properties_Filtered/sec			Shows the number of property changes that are received during the replication that we have already seen.
active_dir.DRA_Inbound_Properties_Total/sec			Shows the total number of object

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Active Directory metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			properties received from inbound replication partners.
active_dir.DRA_Inbound_Values_DNs/sec			Shows the number of object property values received from inbound replication partners that are distinguished names (DNs) that reference other objects. DN values, such as group or distribution list memberships, are generally more expensive to apply than other kinds of values.
active_dir.DRA_Inbound_Values_Total/sec			Shows the total number of object property values received from inbound replication partners. Each inbound object has one or more properties, and each property has zero or more values. Zero values indicate property removal.
active_dir.DRA_Outbound_Bytes_after_compressed/sec			Shows the compressed size of outbound compressed replication data (size after compression, from DSAs in other sites).
active_dir.DRA_Outbound_Bytes_before_compressed/sec			Shows the original size of outbound compressed replication data (size before compression, from DSAs in other sites).
active_dir.DRA_Outbound_Bytes_not_compressed/sec			Shows the number of bytes replicated

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Active Directory metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			out that were not compressed (that is, from DSAs in the same site).
active_dir.DRA_Outbound_Bytes_Total/sec (featured metric)			Shows the total number of bytes replicated out. This counter is the sum of the number of uncompressed bytes (never compressed) and the number of compressed bytes (after compression).
active_dir.DRA_Outbound_Objects/sec			Shows the number of objects that were determined by outbound replication to have no updates that the outbound partner did not already have.
active_dir.DRA_Outbound_Objects_Filtered/sec			Shows the number of objects replicated out.
active_dir.DRA_Outbound_Properties/sec			Shows the number of properties replicated out.
active_dir.DRA_Outbound_Values_DNs/sec			Shows the number of object property values containing DNs sent to outbound replication partners. DN values, such as group or distribution list memberships, are generally more expensive to read than other kinds of values.
active_dir.DRA_Outbound_Values_Total/sec			Shows the number of object property values sent to outbound replication partners.
active_dir.DRA_Pending_Replication_Synchronizations			Shows the number of directory synchronizations that are queued for this

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Active Directory metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			server but not yet processed.
active_dir.DRA_Sync_Requests_Made			Shows the number of synchronization requests made to neighbors.
active_dir.DS_Threads_in_Use (featured metric)			Shows the current number of threads in use by the directory service (which is different from the number of threads in the directory service process). This is the number of threads currently servicing client API calls; it can be used to indicate whether additional processors should be used.
active_dir.LDAP_Bind_Time (featured metric)			Shows the time, in milliseconds, taken for the last successful LDAP bind.
active_dir.LDAP_Client_Sessions (featured metric)			Shows the number of currently connected LDAP client sessions.
active_dir.LDAP_Searches/sec (featured metric)			Shows the rate at which LDAP clients perform search operations.
active_dir.LDAP_Successful_Binds/sec (featured metric)			Shows the number of LDAP binds per second.

Apache Kafka default checks and policies

Agent Client Collector provides the following policies for Apache Kafka health monitoring. Policies come with the checks specified in the indicated table. Policies and checks are available for both Windows and Linux.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Kafka Topic Events

Check	Description	Usage	Output
kafka.check-zookeeper-status	Raises a critical event if the hosted Kafka Zookeeper is down.	`commonchecks check-kafka-zk-status` [flags] Where the flags are: `-p, --port =` Zookeeper Port (default "2181"). Usage example: `commonchecks check-kafka-zk-status -p 2181`	Kafka Zookeeper Status OK: Kafka Zookeeper is Up!
kafka.check-topic-replicas	Raises critical event if any topic has partitions with unknown replicas.	`commonchecks check-kafka-replicas` [flags] Where the flags are: • `-p, --port =` Zookeeper Port (default "2181"). • `-d, --detailed` = Lists unknown replicas in each partition of a topic. • `-i, --include_list` = Comma seperated list to include topics (allows * wild character) • `-e, --exclude_list` = Comma seperated list to exclude topics (allows * wild character) Usage example: commonchecks check-kafka-replicas -H localhost -p 2181 -i "test*" -e "accTopic,*offsets" -d	<topic> has partitions with unknown replicas. Unknown replicas are: {"0":["0"],"1":["0"],"2":["0"]}. <topic> has partitions with unknown replicas. Unknown replicas are: {"0":["0"]}.
kafka.check-topic-replication-factor	Raises critical event if replication factor of at least one topic is above or below	`commonchecks check-kafka-rf` [flags]	TestTopic has replication factor 1,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Kafka Topic Events (continued)

Check	Description	Usage	Output
	provided replication factor param.	Where the flags are: • `-p, --port =` Zookeeper Port (default "2181"). • `-d, --detailed` = Lists unknown replicas in each partition of a topic. • `-i, --include_list` = Comma separated list to include topics (allows * wild character) • `-e, --exclude_list` = Comma separated list to exclude topics (allows * wild character) • `-r, --replication factor` = Expected replication factor for a topic (default 1) Examples: `commonchecks check-kafka-partitions -H localhost -p 2181 -r 2 -i "accMetrics,*Topic" -e "testTopic"`	which is less than expected: 2. accMetrics has replication factor 1, which is less than expected: 2.
kafka.check-topic-leader	Raises critical event if any topic has partitions with unknown leaders or unpreferred replica as leader.	`commonchecks check-kafka-leader` [flags] Where the flags are • `-p, --port` = Zookeeper Port (default "2181"). • `-d, --detailed` = A list of partitions in each topic with	<topic> contains, partitions with unpreferred replica as leader.(partitions with unpreferred replicas are [0]). <topic> contains, partitions with unpreferred replica as leader.(partitions with

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Kafka Topic Events (continued)

Check	Description	Usage	Output
		unknown leaders or unpreferred replicas. • `-i, --include_list` = Comma separated list to include topics (allows * wild character) • `-e, --exclude_list` =Comma separated list to exclude topics (allows * wild character) Examples: `commonchecks check-kafka-leader -H localhost -p 2181 -d -e "*offsets"`	unpreferred replicas are [0]).
kafka.check-topic-partitions	Raises critical events if number of partitions for a topic is less the min_partitions param.	`commonchecks check-kafka-partitions [flags]` Where the flags are: • `-p, --port` = Zookeeper Port (default "2181"). • `-P, --min_partitions` = Minimum partitions for a topic (default 1). • `-i, --include_list` = Comma separated list to include topics (allows * wild character) • `-e, --exclude_list` =Comma separated	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Kafka Topic Events (continued)

Check	Description	Usage	Output
		list to exclude topics (allows * wild character)	
		Usage example 1: `commonchecks check-kafka-partitions -H localhost -p 2181 -P 3`	<topic> has 1 partitions, expected at least 3. <topic> has 1 partitions, expected at least 3. <topic> has 1 partitions, expected at least 3.
		Usage example 2: commonchecks check-kafka-partitions -H localhost -p 2181 -P 3 -i "accMetrics,*Topic" -e "testTopic"	<topic> has 1 partitions, expected at least 3. <topic> has 1 partitions, expected at least 3.

Note: Values for include_list and exclude_list parameters have to be wrapped between double quotes. For example: "test1,*topic".

Apache Kafka Broker Events

Check	Description	Usage	Output
kafka.check-broker-status	Raises critical event if Kafka Broker on the host is down.	`commonchecks check-kafka-broker-status` [flags] Where the flags are: `-p, --port` = Kafka Broker port (default "9092"). Usage example: `commonchecks check-kafka-broker-status -p 9092`	Kafka Broker Status OK: Kafka Broker ubuntu20:9092 is Up!

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Apache Kafka Broker Metrics

Check	Description	Usage	Output
kafka.metrics.broker	Collects Kafka Broker Metrics from the host.	`commonchecks metric-kafka-broker` [flags] Where the flags are: • `-J, --javapath` = Java executable path (default "java"). • `-j, --jmxport` = JMX Port (default "9999") Usage example: `commonchecks metric-kafka-broker -J "/usr/bin/java" -j 9999`	hostname.Kafka.Broker.ReplicaManage 0.000 hostname.Kafka.Broker.DelayedOperati 627.000 hostname.Kafka.Broker.ControllerStats. 0.000 hostname.Kafka.Broker.RequestMetrics 0.000

Apache Kafka Zookeeper Metrics

Check	Description	Usage	Output
kafka.metrics.zookeeper	Collects Zookeeper Metrics from the host.	`commonchecks metric-kafka-zookeeper` [flags] Where the flag is: `-p, --adminserverport` = Admin Server Port (default "8085") Usage example: `commonchecks metric-kafka-zookeeper -p 8085`	hostname.Kafka.Zookeeper.outstanding 2.000 1648183249 hostname.Kafka.Zookeeper.avg_latency 1.05 1648183249 hostname.Kafka.Zookeeper.num_alive_ 1.000 1648183249 hostname.Kafka.Zookeeper.open_file_d 124.000 1648183249

Apache Kafka metrics

The following table lists the metrics that are gathered as output from Kafka checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Broker metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
Kafka.Broker.RequestMetrics.TotalTimeMs.FetchConsumer.Mean (featured metric)			Total time (in ms) to serve the Consumer request.
Kafka.Broker.RequestMetrics.TotalTimeMs.FetchFollower.Mean (featured metric)			Total time (in ms) to serve the Follower request.
Kafka.Broker.RequestMetrics.TotalTimeMs.Produce.Mean (featured metric)			Total time (in ms) to serve the Producer request.
Kafka.Broker.BrokerTopicMetrics.BytesOutPerSec.OneMinuteRate (featured metric)			Aggregate outgoing byte rate.
Kafka.Broker.BrokerTopicMetrics.BytesInPerSec.OneMinuteRate (featured metric)			Aggregate incoming byte rate.
Kafka.Broker.ReplicaManager.UnderReplicatedPartitions.Value			Number of unreplicated partitions.
Kafka.Broker.ReplicaManager.IsrExpandsPerSec.OneMinuteRate			Rate at which the pool of in-sync replicas (ISRs) shrinks.
Kafka.Broker.ReplicaManager.IsrShrinksPerSec.OneMinuteRate			Rate at which the pool of in-sync replicas (ISRs) shrinks.
Kafka.Broker.KafkaController.ActiveControllerCount.Value (featured metric)			Number of active controllers in cluster.
Kafka.Broker.KafkaController.OfflinePartitionsCount.Value			Number of offline partitions.
Kafka.Broker.ControllerStats.LeaderElectionRateAndTimeMs.OneMinuteRate (featured metric)			Leader election rate and latency.
Kafka.Broker.ControllerStats.UncleanLeaderElectionsPerSec.OneMinuteRate			Number of unclean elections per second.
Kafka.Broker.DelayedOperationPurgatory.PurgatorySize.Fetch.Value			Number of requests waiting in fetch purgatory.
Kafka.Broker.DelayedOperationPurgatory.PurgatorySize.Produce.Value			Number of requests waiting in producer purgatory.
Kafka.Broker.RequestMetrics.RequestsPerSec.Produce.OneMinuteRate (featured metric)			Number of producer requests per second.
Kafka.Broker.RequestMetrics.RequestsPerSec.FetchConsumer.OneMinuteRate (featured metric)			Number of consumer requests per second.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Broker metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
Kafka.Broker.RequestMetrics.RequestsPerSec.FetchFollower.OneMinuteRate (featured metric)			Number of follower requests per second.

Zookeeper metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
Kafka.Zookeeper.outstanding_requests (featured metric)			Number of requests queued.
Kafka.Zookeeper.avg_latency (featured metric)			Average amount of time it takes to respond to a client request (in ms).
Kafka.Zookeeper.num_alive_connections (featured metric)			Number of clients connected to ZooKeeper.
Kafka.Zookeeper.open_file_descriptor_count			Number of file descriptors in use.
Kafka.Zookeeper.pending_syncs			Number of pending syncs from followers.
Kafka.Zookeeper.followers			Number of active followers.

AWS Health Monitoring default checks and policies

Agent Client Collector provides the following default checks and policies for AWS monitoring.

AWS Metrics policy

Type	Check	Description	Usage	Usage example
Metric	aws.metrics-ec2	Returns metrics of all EC2 instances in a AWS Datacenter.	metrics-ec2.rb (options) -F, --filter FILTER String representation of the filter to filter EC2 instances by tags. -b, --batch_size BATCH_SIZE Batch size. Number of instances to	`./metrics-ec2.rb -F "{name:tag:environment,v [ENV1,ENV2]} {name:tag:Name,values: [Instance1,Instance2]}"` Provide the filter string along with double quotes ("") around the filter parameter. Set the *batch_size*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Metrics policy (continued)

Type	Check	Description	Usage	Usage example
			fetch metrics in parallel. (Default value: 10)	parameter according to the agent environment.
Metric	aws.metrics-ebs	Returns metrics of all EBS volumes in an AWS datacenter.	metrics-ebs.rb (options) -F, --filter FILTER String representation of the filter to filter EBS Volumes by Tags. -b, --batch_size BATCH_SIZE Batch size. Number of volumes to fetch metrics parallelly. (Default value: 10)	`./metrics-ebs.rb -F "{name:tag:environment,values:[ENV1,ENV2]} {name:tag:Name,values:[Volume1,Volume2]}"` Provide filter string along with double quotes ("") for the filter parameter. Set the *batch_size* parameter according to the agent environment.
Metric	aws.metrics-s3	Returns metrics of all S3 buckets in an AWS datacenter.	metrics-s3.rb (options) -b, --batch_size BATCH_SIZE Batch size. Number of S3 buckets to fetch metrics in parallel. (Default value: 10)	`./metrics-s3.rb -b 5` Set the *batch_size* parameter according to the agent environment.
Metric	aws.metrics-rds	Returns metrics of all RDS buckets in an AWS datacenter.	metrics-rds.rb (options) -F, --filter FILTER String representation of the filter to filter RDS Instances by tags. -b, --batch_size BATCH_SIZE Batch size. Number of	`./metrics-rds.rb -F "{name:tag:environment,values:[ENV1,ENV2]} {name:tag:engine,values:[MariaDB]}"` Provide the filter string with double quotes ("") around the filter parameter. Set the *batch_size*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Metrics policy (continued)

Type	Check	Description	Usage	Usage example
			databases to fetch metrics in parallel. (Default value: 10)	parameter according to the agent environment.
Metric	aws.metrics-classic-elb	Returns metrics of all classic elastic load balancers in an AWS datacenter.	metrics-elb.rb (options) -e, --exclude_lb EXCLUDE_LB Exclude metrics from load balancers whose name contains one of the comma separated strings. -b, --batch_size BATCH_SIZE Batch size. Number of databases to fetch metrics in parallel. (Default value: 10)	`./metrics-elb.rb -t classic -e "classic,itom,ip" -b 10` Provide comma separated strings in double quotes ("") for the *`exclude_lb`* parameter. Set the *`batch_size`* parameter according to the agent environment.
Metric	aws.metrics-network-elb	Returns metrics of all network elastic load balancers in an AWS datacenter.	metrics-elb.rb (options) -e, --exclude_lb EXCLUDE_LB Exclude metrics from load balancers whose name contains any of the comma separated strings. -b, --batch_size BATCH_SIZE Batch size. Number of databases to fetch metrics in parallel. (Default value: 10)	`./metrics-elb.rb -t network -e "vpc,itom,ip" -b 10` Provide comma separated strings in double quotes ("") for the *`exclude_lb`* parameter. Set the *`batch_size`* parameter according to the agent environment. Do not change the command prefix.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS Metrics policy (continued)

Type	Check	Description	Usage	Usage example
Metric	aws.metrics-application-elb	Returns metrics of all application elastic load balancers in an AWS datacenter.	metrics-elb.rb (options) -e, --exclude_lb EXCLUDE_LB Exclude metrics from load balancers whose name contains any one of the string from comma separated strings. -b, --batch_size BATCH_SIZE Batch size. Number of databases to fetch metrics in parallel. (Default value: 10)	`./metrics-elb.rb -t application -e "classic,app,ip" -b 10` Provide comma separated strings in double quotes ("") for the *`exclude_lb`* parameter. Set the *`batch_size`* parameter according to the agent environment. Do not change the command prefix.

Metric	aws.metrics-ec2-count	Returns count of EC2 instances by instance_type or status in an AWS datacenter.	metrics-ec2-count.rb (options)
			
			-s, --scheme SCHEME Metric naming scheme, text to prepend to metric. Default: aws.ec2
			
			-t, --type METRIC type Count by type: status, instance. Default: instance
				`./metrics-ec2-count.rb -t status`
				
				`./metrics-ec2-count.rb -t status -s metric.aws.ec2`

AWS DynamoDB metrics

Type	Check	Description	Usage	Command
Metric	aws.metrics-dynamodb	Monitors the health of DynamoDB tables in an AWS datacenter.	e, --exclude_tb EXCLUDE_TB Exclude metrics from tables whose name	`metrics-dynamodb.rb -r {{.labels.params_ci_reg` `{{.labels.params_exclud` `-f`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AWS DynamoDB metrics (continued)

Type	Check	Description	Usage	Command
		Binds to the DynamoDB table (cmdb_ci_dynamodb_table)	contains a string from comma separated strings. b, --batch_size BATCH_SIZE Batch size. Number of tables to fetch metrics in parallel. (Default value: 10)	{{.labels.params_exclud {{end}} {{.labels.params_batch_ -b {{.labels.params_filter {{end}}

AWS EC2 metrics

Type	Check	Description	Usage	Command
Metric	aws.metrics-ecs-cluster	Monitors ECS clusters in an AWS datacenter. Binds to the AWS ECS Cluster CI type (cmdb_ci_cloud_ecs_cluster)	-F, --filter FILTER String representation of the filter to filter ECS clusters by tags. b, --batch_size BATCH_SIZE Batch size. Number of clusters to fetch metrics in parallel. (Default value: 10)	metrics-ecs-cluster.rb -r {{.labels.params_ci_reg {{.labels.params_filter -f {{.labels.params_filter {{end}} {{.labels.params_batch_ -b {{.labels.params_filter {{end}}
Metric	aws.metrics-ecs-service	Monitors ECS services in an AWS datacenter. Binds to the AWS ECS Service CI type (cmdb_ci_cloud_ecs_service)	-F, --filter FILTER String representation of the filter to filter ECS clusters by tags. b, --batch_size BATCH_SIZE Batch size. Number of clusters to fetch metrics in parallel. (Default value: 10)	metrics-ecs-service.rb -r {{.labels.params_ci_reg {{.labels.params_filter -f {{.labels.params_filter {{end}} {{.labels.params_batch_ -b {{.labels.params_filter {{end}}

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

AWS metrics

AWS metrics

The following tables list and describe the metrics that are gathered as output from the specified AWS checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

aws.metrics-ec2 metrics

Metric Name	Description
aws.ec2.CPUCreditBalance.Average	Number of earned CPU credits that an instance has accrued since it was launched or started.
aws.ec2.CPUUtilization.Average	Average percentage of allocated EC2 compute units that are currently in use on the instance.
aws.ec2.DiskReadOps.Average (Featured metric)	Average of completed read operations from all instance store volumes available to the instance in a specified time period.
aws.ec2.DiskWriteOps.Average (Featured metric)	Completed write operations to all instance store volumes available to the instance in a specified time period.
aws.ec2.DiskReadBytes.Average (Featured metric)	Average number of bytes read from all instance store volumes available to the instance.
aws.ec2.DiskWriteBytes.Average (Featured metric)	Average number of bytes written to all instance store volumes available to the instance.
aws.ec2.NetworkIn.Average (Featured metric)	Average number of bytes received by the instance on all network interfaces. This metric identifies the volume of incoming network traffic to a single instance.
aws.ec2.NetworkOut.Average (Featured metric)	Average number of bytes sent out by the instance on all network interfaces. This metric identifies the volume of outgoing network traffic from a single instance.
aws.ec2.NetworkPacketsIn.Average	Average number of packets received by the instance on all network interfaces. This metric identifies the volume of incoming traffic in terms of the number of packets on a single instance.
aws.ec2.NetworkPacketsOut.Average	Average number of packets sent out by the instance on all network interfaces. This metric identifies the volume of outgoing traffic in terms of the number of packets on a single instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-ec2 metrics (continued)

Metric Name	Description
aws.ec2.MetadataNoToken.Average	Average number of times the instance metadata service was successfully accessed using a method that does not use a token.
aws.ec2.CPUCreditUsage.Average	Average number of CPU credits spent by the instance for CPU utilization.
aws.ec2.CPUSurplusCreditBalance.Average	Average number of surplus credits that have been spent by an unlimited instance when its *CPUCreditBalance* value is zero.
aws.ec2.CPUSurplusCreditsCharged.Average	Average number of spent surplus credits that are not paid down by earned CPU credits, and which thus incur an additional charge.
aws.ec2.StatusCheckFailed_Instance.Average	Reports whether the instance has passed the instance status check in the last minute.
aws.ec2.StatusCheckFailed_System.Average	Reports whether the instance has passed the system status check in the last minute.
aws.ec2.StatusCheckFailed.Average	Reports whether the instance has passed both the instance status check and the system status check in the last minute.

aws.metrics-ebs metrics

Metric Name	Description
aws.ebs.VolumeReadBytes.Average (Featured metric)	Average total number of bytes transferred during the period while reading.
aws.ebs.VolumeWriteBytes.Average (Featured metric)	Average total number of bytes transferred during the period while writing.
aws.ebs.VolumeReadOps.Average (Featured metric)	Average total number of read operations in a specified time period.
aws.ebs.VolumeWriteOps.Average (Featured metric)	Average total number of write operations in a specified time period.
aws.ebs.VolumeTotalReadTime.Average (Featured metric)	Average total number of seconds spent by all read operations completed in a specified time period.
aws.ebs.VolumeTotalWriteTime.Average (Featured metric)	Average of total number of seconds spent by all write operations completed in a specified time period.
aws.ebs.VolumeIdleTime.Average	Average of total number of seconds in a specified time period when no read or write operations were submitted.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-ebs metrics (continued)

Metric Name	Description
aws.ebs.VolumeQueueLength.Average	Average number of read and write operation requests waiting to be completed in a specified time period.
aws.ebs.VolumeThroughputPercentage.Average	Average percentage of I/O operations per second (IOPS) delivered of the total IOPS provisioned for an Amazon EBS volume.
aws.ebs.VolumeConsumedReadWriteOps.Average	Average total amount of read and write operations (normalized to 256K capacity units) consumed in a specified time period.
aws.ebs.BurstBalance.Average	Average percentage of I/O credits (for gp2) or throughput credits (for st1 and sc1) remaining in the burst bucket.

aws.metrics-s3 metrics

Metric Name	Description
aws.s3.NumberOfObjects.Average (Featured metric)	Average total number of objects stored in a bucket for all storage classes. Collected every 12 hours.
aws.s3.BucketSizeBytes.Average (Featured metric)	Average amount of data in bytes stored in a bucket in the STANDARD storage class. Collected every 12 hours.

aws.metrics-rds metrics

Metric Name	Description
aws.rds.BinLogDiskUsage.Average	Average amount of disk space occupied by binary logs on the primary. Applies to MySQL read replicas.
aws.rds.BurstBalance.Average	Average percent of General Purpose SSD (gp2) burst-bucket I/O credits available.
aws.rds.CPUUtilization.Average (Featured metric)	Average percentage of CPU utilization.
aws.rds.CPUCreditUsage.Average	Average number of CPU credits spent by the instance for CPU utilization.
aws.rds.CPUCreditBalance.Average	Average number of surplus credits that have been spent by an unlimited instance when its *`CPUCreditBalance`* value is zero.
aws.rds.DatabaseConnections.Average (Featured metric)	Average number of database connections in use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-rds metrics (continued)

Metric Name	Description
aws.rds.DiskQueueDepth.Average (Featured metric)	Average number of outstanding I/Os (read/write requests) waiting to access the disk.
aws.rds.EBSByteBalance%.Average	Average percentage of throughput credits remaining in the burst bucket of your RDS database.
aws.rds.EBSIOBalance%.Average	Average percentage of I/O credits remaining in the burst bucket of your RDS database.
aws.rds.FailedSQLServerAgentJobsCount.Average (Featured metric)	Average number of failed Microsoft SQL Server Agent jobs during the last minute.
aws.rds.FreeableMemory.Average (Featured metric)	Average amount of available random access memory.
aws.rds.FreeStorageSpace.Average	Average amount of available storage space.
aws.rds.MaximumUsedTransactionIDs.Average	Average amount of maximum transaction IDs that have been used. Applies to PostgreSQL.
aws.rds.NetworkReceiveThroughput.Average (Featured metric)	Average incoming (receive) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication.
aws.rds.NetworkTransmitThroughput.Average (Featured metric)	Average outgoing (transmit) network traffic on the DB instance, including both customer database traffic and Amazon RDS traffic used for monitoring and replication.
aws.rds.OldestReplicationSlotLag.Average	Average lagging size of the replica lagging the most in terms of write-ahead log (WAL) data received. Applies to PostgreSQL.
aws.rds.ReadIOPS.Average	Average number of disk read I/O operations per second.
aws.rds.ReadLatency.Average (Featured metric)	Average amount of time taken per disk I/O operation.
aws.rds.ReadThroughput.Average	Average number of bytes read from disk per second.
aws.rds.ReplicaLag.Average	Average amount of time a read replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, Oracle, PostgreSQL, and SQL Server read replicas.
aws.rds.ReplicationSlotDiskUsage.Average	Average amount of disk space used by replication slot files. Applies to PostgreSQL.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-rds metrics (continued)

Metric Name	Description
aws.rds.SwapUsage.Average	Average amount of swap space used on the DB instance. This metric is not available for SQL Server.
aws.rds.TransactionLogsDiskUsage.Average	Average disk space used by transaction logs. Applies to PostgreSQL.
aws.rds.TransactionLogsGeneration.Average	Average size of transaction logs generated per second. Applies to PostgreSQL.
aws.rds.WriteIOPS.Average	Average number of disk write I/O operations per second.
aws.rds.WriteLatency.Average (Featured metric)	Average amount of time taken per disk I/O operation.
aws.rds.WriteThroughput.Average	Average number of bytes written to disk per second.
aws.rds.CPUSurplusCreditBalance.Average	Average number of surplus credits that have been spent by an unlimited instance when its *`CPUCreditBalance`* value is zero.
aws.rds.CPUSurplusCreditsCharged.Average	Average number of spent surplus credits that are not paid down by earned CPU credits, and which therefore incur an additional charge.

aws.metrics-classic-elb metrics

Metric Name	Description
aws.classicElb.EstimatedALBActiveConnectionCount.Average (Featured metric)	Average estimated number of concurrent TCP connections active from clients to the load balancer and from the load balancer to targets.
aws.classicElb.EstimatedALBConsumedLCUs.Average	Average estimated number of load balancer capacity units (LCU) used by an Application Load Balancer. You pay for the number of LCUs that you use per hour.
aws.classicElb.EstimatedALBNewConnectionCount.Average (Featured metric)	Average estimated number of new TCP connections established from clients to the load balancer and from the load balancer to targets.
aws.classicElb.EstimatedProcessedBytes.Average	Estimated number of bytes processed by an Application Load Balancer.
aws.classicElb.DesyncMitigationMode_NonCompliant_Request_Count.Sum	Number of requests that do not comply with RFC 7230.
aws.classicElb.HTTPCode_ELB_5XX.Sum (Featured metric)	Number of HTTP 4XX client error codes generated by the load balancer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-classic-elb metrics (continued)

Metric Name	Description
aws.classicElb.HTTPCode_ELB_4XX.Sum (Featured metric)	Number of HTTP 5XX server error codes generated by the load balancer.
aws.classicElb.SurgeQueueLength.Maximum	Maximum total number of requests (HTTP listener) or connections (TCP listener) that are pending routing to a healthy instance. The maximum size of the queue is 1,024. Additional requests or connections are rejected when the queue is full.
aws.classicElb.SpilloverCount.Sum (Featured metric)	Total number of requests rejected because the surge queue is full.
aws.classicElb.RequestCount.Sum	Total number of requests completed or connections made during the specified interval.
aws.classicElb.Latency.Average (Featured metric)	Average total time elapsed, in seconds, from the time the load balancer sent the request to a registered instance until the instance started to send the response headers.
aws.classicElb.BackendConnectionErrors.Sum (Featured metric)	Total number of connections that were not successfully established between the load balancer and the registered instances.
aws.classicElb.HealthyHostCount.Average	Average number of healthy instances registered with your load balancer. A newly registered instance is considered healthy after it passes the first health check.
aws.classicElb.UnHealthyHostCount.Average	Average of number of unhealthy instances registered with your load balancer. An instance is considered unhealthy after it exceeds the unhealthy threshold configured for health checks.

aws.metrics-network-elb

Metric Name	Description
aws.networkElb.ActiveFlowCount.Average (Featured metric)	Average total number of concurrent flows (or connections) from clients to targets. This metric includes connections in the SYN_SENT and ESTABLISHED states.
aws.networkElb.ActiveFlowCount_TCP.Average (Featured metric)	Average total number of concurrent TCP flows (or connections) from clients to targets. This metric includes only connections in the ESTABLISHED state.
aws.networkElb.ActiveFlowCount_TLS.Average	Average total number of concurrent TLS flows (or connections) from clients to targets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-network-elb (continued)

Metric Name	Description
(Featured metric)	This metric includes only connections in the ESTABLISHED state.
aws.networkElb.ActiveFlowCount_UDP.Average (Featured metric)	Average total number of concurrent UDP flows (or connections) from clients to targets.
aws.networkElb.ClientTLSNegotiationErrorCount.Sum (Featured metric)	Average total number of TLS handshakes that failed during negotiation between a client and a TLS listener.
aws.networkElb.ConsumedLCUs.Average	Average number of load balancer capacity units (LCU) used by your load balancer.
aws.networkElb.ConsumedLCUs_TCP.Average	Average number of load balancer capacity units (LCU) used by your load balancer for TCP.
aws.networkElb.ConsumedLCUs_TLS.Average	Average number of load balancer capacity units (LCU) used by your load balancer for TLS.
aws.networkElb.ConsumedLCUs_UDP.Average	Average number of load balancer capacity units (LCU) used by your load balancer for UDP.
aws.networkElb.HealthyHostCount.Maximum (Featured metric)	The maximum number of targets that are considered healthy.
aws.networkElb.NewFlowCount.Sum (Featured metric)	Total number of new flows (or connections) established from clients to targets in the time period.
aws.networkElb.NewFlowCount_TCP.Sum (Featured metric)	Total number of new TCP flows (or connections) established from clients to targets in the time period.
aws.networkElb.NewFlowCount_TLS.Sum (Featured metric)	Total number of new TLS flows (or connections) established from clients to targets in the time period.
aws.networkElb.NewFlowCount_UDP.Sum (Featured metric)	Total number of new UDP flows (or connections) established from clients to targets in the time period.
aws.networkElb.PeakBytesPerSecond.Maximum (Featured metric)	Maximum average throughput (bytes per second), calculated every 10 seconds during the sampling window. This metric includes health check traffic.
aws.networkElb.PeakPacketsPerSecond.Maximum (Featured metric)	Maximum average packet rate (packets processed per second), calculated every 10 seconds during the sampling window. This metric includes health check traffic.
aws.networkElb.ProcessedBytes.Sum	Total number of bytes processed by the load balancer, including TCP/IP headers. This count

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-network-elb (continued)

Metric Name	Description
	includes traffic to and from targets, minus health check traffic.
aws.networkElb.ProcessedBytes_TCP.Sum	Total number of bytes processed by TCP listeners.
aws.networkElb.ProcessedBytes_TLS.Sum	Total number of bytes processed by TLS listeners.
aws.networkElb.ProcessedBytes_UDP.Sum	Total number of bytes processed by UDP listeners.
aws.networkElb.ProcessedPackets.Sum	Total number of packets processed by the load balancer. This count includes traffic to and from targets, including health check traffic.
aws.networkElb.TargetTLSNegotiationErrorCount.Sum (Featured metric)	Total number of TLS handshakes that failed during negotiation between a TLS listener and a target.
aws.networkElb.TCP_Client_Reset_Count.Sum	Total number of reset (RST) packets sent from a client to a target. These resets are generated by the client and forwarded by the load balancer.
aws.networkElb.TCP_ELB_Reset_Count.Sum	Total number of reset (RST) packets generated by the load balancer.
aws.networkElb.TCP_Target_Reset_Count.Sum	Total number of reset (RST) packets sent from a target to a client. These resets are generated by the target and forwarded by the load balancer.
aws.networkElb.UnHealthyHostCount.Maximum (Featured metric)	Maximum number of targets that can be considered unhealthy.

aws.metrics-application-elb metrics

Metric Name	Description
aws.applicationElb.ActiveConnectionCount.Sum (Featured metric)	Total number of concurrent TCP connections active from clients to the load balancer and from the load balancer to targets.
aws.applicationElb.ClientTLSNegotiationErrorCount.Sum (Featured metric)	Total number of TLS connections initiated by the client that did not establish a session with the load balancer due to a TLS error.
aws.applicationElb.ConsumedLCUs.Average	The average number of load balancer capacity units (LCU) used by your load balancer. You pay for the number of LCUs that you use per hour. For more information, see Elastic Load Balancing pricing.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-application-elb metrics (continued)

Metric Name	Description
aws.applicationElb.DesyncMitigationMode_NonCompliant_Request_Count.Sum (Featured metric)	Total number of requests that do not comply with RFC 7230.
aws.applicationElb.DroppedInvalidHeaderRequestCount.Average (Featured metric)	Average of requests where the load balancer removed HTTP headers with header fields that are not valid before routing the request. The load balancer removes these headers only if the routing.http.drop_invalid_header_fields.enabled attribute is set to true.
aws.applicationElb.ForwardedInvalidHeaderRequestCount.Average (Featured metric)	Average of requests routed by the load balancer that had HTTP headers with header fields that are not valid. The load balancer forwards requests with these headers only if the routing.http.drop_invalid_header_fields.enabled attribute is set to false.
aws.applicationElb.GrpcRequestCount.Sum	Total number of gRPC requests processed over IPv4 and IPv6.
aws.applicationElb.HTTP_Fixed_Response_Count.Sum	Total number of fixed-response actions that were successful.
aws.applicationElb.HTTP_Redirect_Count.Sum	Total number of redirect actions that were successful.
aws.applicationElb.HTTP_Redirect_Url_Limit_Exceeded_Count.Sum (Featured metric)	Total number of redirect actions that couldn't be completed because the URL in the response location header is larger than 8K.
aws.applicationElb.HTTPCode_ELB_3XX_Count.Sum (Featured metric)	Total number of HTTP 3XX redirection codes that originate from the load balancer. This count does not include response codes generated by targets.
aws.applicationElb.HTTPCode_ELB_4XX_Count.Sum (Featured metric)	Total number of HTTP 4XX client error codes that originate from the load balancer. This count does not include response codes generated by targets.
aws.applicationElb.HTTPCode_ELB_5XX_Count.Sum (Featured metric)	Total number of HTTP 5XX server error codes that originate from the load balancer. This count does not include any response codes generated by the targets.
aws.applicationElb.HTTPCode_ELB_500_Count.Sum (Featured metric)	Total number of HTTP 500 error codes that originate from the load balancer.
aws.applicationElb.HTTPCode_ELB_502_Count.Sum (Featured metric)	Total number of HTTP 502 error codes that originate from the load balancer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-application-elb metrics (continued)

Metric Name	Description
aws.applicationElb.HTTPCode_ELB_503_Count.Sum (Featured metric)	Total number of HTTP 503 error codes that originate from the load balancer.
aws.applicationElb.HTTPCode_ELB_504_Count.Sum (Featured metric)	Total number of HTTP 504 error codes that originate from the load balancer.
aws.applicationElb.IPv6ProcessedBytes.Sum	Total number of bytes processed by the load balancer over IPv6.
aws.applicationElb.IPv6RequestCount.Sum	Total number of IPv6 requests received by the load balancer.
aws.applicationElb.NewConnectionCount.Sum	Total number of new TCP connections established from clients to the load balancer and from the load balancer to targets.
aws.applicationElb.NonStickyRequestCount.Sum	Total number of requests where the load balancer chose a new target because it couldn't use an existing sticky session.
aws.applicationElb.ProcessedBytes.Sum (Featured metric)	Total number of bytes processed by the load balancer over IPv4 and IPv6.
aws.applicationElb.RejectedConnectionCount.Sum (Featured metric)	Total number of connections that were rejected because the load balancer had reached its maximum number of connections.
aws.applicationElb.RequestCount.Sum	Total number of requests processed over IPv4 and IPv6.
aws.applicationElb.RuleEvaluations.Sum	Total number of rules processed by the load balancer given a request rate averaged over an hour.

aws.metrics-ec2-count metrics

Metric Name	Description
aws.ec2.types.<type_name>	Count of EC2 instances for the specified type in AWS Datacenter.
aws.ec2.status.running	Count of EC2 instances in Active state in an AWS Datacenter.
aws.ec2.status.stopped	Count of EC2 instances in Stopped state in an AWS Datacenter.
aws.ec2.status.total	Count of all EC2 instances in an AWS Datacenter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-dynamodb metrics

Metric Name	Description
aws.dynamodb.ConsumedWriteCapacityUnits.Average (Featured metric)	Number of write capacity units consumed over the specified time period, for tracking how much provisioned throughput is used.
aws.dynamodb.ConsumedReadCapacityUnits.Average (Featured metric)	Number of read capacity units consumed over the specified time period, for tracking how much provisioned throughput is used.
aws.dynamodb.ProvisionedWriteCapacityUnits.Average (Featured metric)	Number of provisioned write capacity units for a table.
aws.dynamodb.ProvisionedReadCapacityUnits.Average (Featured metric)	Number of provisioned read capacity units for a table.
aws.dynamodb.SuccessfulRequestLatency.GetItem.Average (Featured metric)	Successful requests to DynamoDB or Amazon DynamoDB streams during the specified time period. Run by a GetItem operation.
aws.dynamodb.SuccessfulRequestLatency.BatchGetItem.Average	Successful requests to DynamoDB or Amazon DynamoDB streams during the specified time period. Run by a BatchGetItem operation.
aws.dynamodb.SuccessfulRequestLatency.PutItem.Average (Featured metric)	Successful requests to DynamoDB or Amazon DynamoDB streams during the specified time period. Run by a PutItem operation.
aws.dynamodb.SuccessfulRequestLatency.BatchPutItem.Average	Successful requests to DynamoDB or Amazon DynamoDB streams during the specified time period. Run by a BatchPutItem operation.
aws.dynamodb.SuccessfulRequestLatency.Query.Average	Successful requests to DynamoDB or Amazon DynamoDB streams during the specified time period. Run by a Query operation.
aws.dynamodb.SuccessfulRequestLatency.Scan.Average	Successful requests to DynamoDB or Amazon DynamoDB Streams during the specified time period. Run by a Scan operation.
aws.dynamodb.ReturnedItemCount.Scan.Average	Number of items returned by Scan operations during the specified time period.
aws.dynamodb.ReturnedItemCount.Query.Average	Number of items returned by Query operations during the specified time period.
aws.dynamodb.UserErrors.Average (Featured metric)	Requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 400 status code during the specified time period. An HTTP 400 usually indicates a client-side error, such as an invalid combination of parameters, an attempt to update a nonexistent table, or an incorrect request signature.
aws.dynamodb.SystemErrors.GetItem.Average	Requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-dynamodb metrics (continued)

Metric Name	Description
(Featured metric)	500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by a GetItem operation.
aws.dynamodb.SystemErrors.Scan.Average	The requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by a Scan operation.
aws.dynamodb.SystemErrors.Query.Average	Requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by a Query operation.
aws.dynamodb.SystemErrors.BatchGetItem.Average	Requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by a BatchGet operation.
aws.dynamodb.SystemErrors.PutItem.Average (Featured metric)	Requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by a PutItem operation.
aws.dynamodb.SystemErrors.UpdateItem.Average	The requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by an UpdateItem operation.
aws.dynamodb.SystemErrors.DeleteItem.Average	Requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by a DeleteItem operation.
aws.dynamodb.SystemErrors.BatchWriteItem.Average	Requests to DynamoDB or Amazon DynamoDB streams that generate an HTTP 500 status code during the specified time period. An HTTP 500 usually indicates an internal service error. Run by a BatchWrite operation.
aws.dynamodb.ConditionalCheckFailedRequests.Average (Featured metric)	Number of failed attempts to perform conditional writes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-dynamodb metrics (continued)

Metric Name	Description
aws.dynamodb.TimeToLiveDeletedItemCount.Average	Number of items deleted by Time to Live (TTL) during the specified time period. This metric helps you monitor the rate of TTL deletions on your table.
aws.dynamodb.TransactionConflict.Average	Rejected item-level requests due to transactional conflicts between concurrent requests on the same items.
aws.dynamodb.ReadThrottleEvents.Average (Featured metric)	Requests to DynamoDB that exceed the provisioned read capacity units for a table.
aws.dynamodb.WriteThrottleEvents.Average (Featured metric)	Requests to DynamoDB that exceed the provisioned write capacity units for a table.
aws.dynamodb.ThrottledRequests.GetItem.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for a GetItem operation.
aws.dynamodb.ThrottledRequests.PutItem.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for a PutItem operation.
aws.dynamodb.ThrottledRequests.DeleteItem.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for a DeleteItem operation.
aws.dynamodb.ThrottledRequests.UpdateItem.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for an UpdateItem operation.
aws.dynamodb.ThrottledRequests.BatchGetItem.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for a BatchGetItem operation.
aws.dynamodb.ThrottledRequests.Scan.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for a Scan operation.
aws.dynamodb.ThrottledRequests.Query.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for a Query operation.
aws.dynamodb.ThrottledRequests.BatchWriteItem.Average	Requests to DynamoDB that exceed the provisioned throughput limits on a resource for a BatchWriteItem operation.
aws.dynamodb.SuccessfulRequestLatency.GetRecords.Average	The successful requests to DynamoDB or Amazon DynamoDB Streams during the specified time period. Run by a GetRecords operation.
aws.dynamodb.ReturnedBytes.GetRecords.Average	The number of bytes returned by GetRecords operations (Amazon DynamoDB streams) during the specified time period.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-dynamodb metrics (continued)

Metric Name	Description
aws.dynamodb.ReturnedRecordsCount.GetRecords.Average	The average of stream records returned by GetRecords operations (Amazon DynamoDB streams) during the specified time period.

aws.metrics-ecs-cluster metrics

Metric	Description
aws.ecs.cluster.CPUReservation.Average (Featured metric)	Percentage of CPU units reserved by running tasks in the cluster.
aws.ecs.cluster.CPUUtilization.Average (Featured metric)	Percentage of CPU units used in the cluster.
aws.ecs.cluster.MemoryReservation.Average (Featured metric)	Percentage of memory reserved by running tasks in the cluster.
aws.ecs.cluster.MemoryUtilization.Average (Featured metric)	Percentage of memory used in the cluster.
aws.ecs.cluster.CpuReserved.Average	CPU units reserved by tasks in the cluster.
aws.ecs.cluster.CpuUtilized.Average	CPU units used by tasks in the cluster.
aws.ecs.cluster.NetworkRxBytes.Average (Featured metric)	Number of bytes received by the cluster.
aws.ecs.cluster.NetworkTxBytes.Average (Featured metric)	Number of bytes transmitted by the cluster.
aws.ecs.cluster.EphemeralStorageReserved.Average	Ephemeral storage is the volatile temporary storage attached to your instances, present only during the running lifetime of the instance.
aws.ecs.cluster.MemoryUtilized.Average (Featured metric)	Memory being used by tasks in the cluster.
aws.ecs.cluster.MemoryReserved.Average (Featured metric)	Memory reserved by tasks in the cluster.
aws.ecs.cluster.StorageReadBytes.Average (Featured metric)	Number of bytes read from storage in the cluster.
aws.ecs.cluster.StorageWriteBytes.Average	Number of bytes written to storage in the cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-ecs-cluster metrics (continued)

Metric	Description
(Featured metric)	
aws.ecs.cluster.ContainerInstanceCount.Average (Featured metric)	Number of EC2 instances running the Amazon ECS agent that are registered with a cluster.
aws.ecs.cluster.ServiceCount.Average (Featured metric)	Number of services in the cluster.
aws.ecs.cluster.TaskCount.Average	Number of tasks running in the cluster.

aws.metrics-ecs-service metrics

Metric	Description
aws.ecs.service.CPUUtilization.Average (Featured metric)	Percentage of CPU units used in the service.
aws.ecs.service.MemoryUtilization.Average (Featured metric)	Percentage of memory used in the service.
aws.ecs.service.CpuReserved.Average	CPU units reserved by tasks in the service.
aws.ecs.service.CpuUtilized.Average	CPU units used by tasks in the service.
aws.ecs.service.NetworkRxBytes.Average (Featured metric)	Number of bytes received by the service.
aws.ecs.service.NetworkTxBytes.Average (Featured metric)	Number of bytes transmitted by the service.
aws.ecs.service.EphemeralStorageReserved.Average	Ephemeral storage is the volatile temporary storage attached to your instances which is present only during the running lifetime of the instance.
aws.ecs.service.MemoryUtilized.Average	Memory being used by tasks in the service.
aws.ecs.service.MemoryReserved.Average	Memory reserved by tasks in the service.
aws.ecs.service.StorageReadBytes.Average (Featured metric)	Number of bytes read from storage in the service.
aws.ecs.service.StorageWriteBytes.Average (Featured metric)	Number of bytes written to storage in the service.
aws.ecs.service.RunningTaskCount.Average	Number of tasks currently in the RUNNING state.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

aws.metrics-ecs-service metrics (continued)

Metric	Description
(Featured metric)	
aws.ecs.service.DeploymentCount.Average	Number of deployments in an Amazon ECS service.
aws.ecs.service.DesiredTaskCount.Average (Featured metric)	Desired number of tasks for an Amazon ECS service.
aws.ecs.service.TaskSetCount.Average (Featured metric)	Number of task sets in the service.
aws.ecs.service.PendingTaskCount.Average (Featured metric)	Number of tasks currently in the PENDING state.

Azure Health Monitoring default checks and policies

Agent Client Collector provides the following default checks and policies for Azure health monitoring.

Azure k8s Metrics

Type	Check	Description	Command

Metric	azure.metrics-k8s-service	Provides pod, nodes, and cluster level metrics from all Azure Kubernetes services in metric form.
		
		Usage: metrics-azure-k8s-service.rb
		
		`-S, --subscription ID ARM Subscription ID`
		`-r, Region`
		
		Example: `metrics-azure-k8s-service.rb -S {{.labels.params_subscription_id}} -r westus2`
			`metrics-azure-k8s-service.rb -r {{.labels.params_ci_region}} {{-s {{if.labels.params_subscrip -S {{.labels.params_subscripti {{end}}`

Azure metrics

The following tables list and describe the metrics that are gathered as output from Azure checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

Azure Kubernetes Service check metrics

Metric Name	Description
apiserver_current_inflight_requests (Featured metric)	Maximum number of currently used inflight requests on the API server per request kind in the past second.
cluster_autoscaler_cluster_safe_to_autoscale	Determines whether cluster autoscaler takes action on the cluster.
cluster_autoscaler_scale_down_in_cooldown	Determines if the scale down is in cooldown; no nodes are removed during this timeframe.
cluster_autoscaler_unneeded_nodes_count	Number of nodes that are currently not needed and can be removed by the autoscaler.
cluster_autoscaler_unschedulable_pods_count	Number of pods that are currently not able to be scheduled in the cluster.
kube_node_status_allocatable_cpu_cores (Featured metric)	Total number of available CPU cores in a managed cluster.
kube_node_status_allocatable_memory_bytes (Featured metric)	Total amount of available memory in a managed cluster.
kube_node_status_condition (Featured metric)	Statuses for various node conditions.
kube_pod_status_phase (Featured metric)	Number of pods by phase.
kube_pod_status_ready (Featured metric)	Number of pods in Ready state.
node_cpu_usage_millicores (Featured metric)	Aggregated measurement of CPU utilization in millicores across the cluster.
node_cpu_usage_percentage (Featured metric)	Aggregated average CPU utilization measured in percentage across the cluster.
node_disk_usage_bytes	Number of disk space bytes in use by device.
node_disk_usage_percentage (Featured metric)	Percentage of disk space used by device.
node_memory_rss_bytes	Number of container RSS memory bytes in use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Kubernetes Service check metrics (continued)

Metric Name	Description
node_memory_rss_percentage (Featured metric)	Percentage of container RSS memory in use.
node_memory_working_set_bytes	Number of container working set memory bytes in use.
node_memory_working_set_percentage (Featured metric)	Percentage of container working set memory in use.
node_network_in_bytes (Featured metric)	Network received bytes.
node_network_out_bytes (Featured metric)	Network transmitted bytes.

Azure VM CPU metrics

Metric Name	Description
Percentage CPU (Featured metric)	The percentage of allocated compute units that are currently in use by the Virtual Machine(s).
CPU Credits Remaining	Total number of credits available to burst. Only available on B-series burstable VMs.
CPU Credits Consumed	Total number of credits consumed by the Virtual Machine. Only available on B-series burstable VM.

Azure VM Network metrics

Metric Name	Description
Network In (Featured metric)	The number of billable bytes received on all network interfaces by the Virtual Machines (Incoming Traffic) (Deprecated).
Network Out (Featured metric)	The number of billable bytes out on all network interfaces by the Virtual Machines (Outgoing Traffic) (Deprecated).
Network In Total (Featured metric)	The number of bytes received on all network interfaces by the Virtual Machines (Incoming Traffic).
Network Out Total (Featured metric)	The number of bytes out on all network interfaces by the Virtual Machines (Outgoing Traffic).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure VM Disk metrics

Metric Name	Description
Disk Read Bytes (Featured metric)	Bytes read from disk during monitoring period.
Disk Write Bytes (Featured metric)	Bytes written to disk during monitoring period.
Disk Read Operations/Sec (Featured metric)	Disk Read IOPS.
Disk Write Operations/Sec (Featured metric)	Disk Write IOPS.
Data Disk Read Bytes/sec (Featured metric)	Bytes per second read from a single disk during monitoring period.
Data Disk Write Bytes/sec (Featured metric)	Bytes per second written to a single disk during monitoring period.
Data Disk Read Operations/Sec (Featured metric)	Read IOPS from a single disk during monitoring period.
Data Disk Write Operations/Sec (Featured metric)	Write IOPS from a single disk during monitoring period.
Data Disk Bandwidth Consumed Percentage (Featured metric)	Percentage of data disk bandwidth consumed per minute.
Data Disk IOPS Consumed Percentage (Featured metric)	Percentage of data disk I/Os consumed per minute.
Data Disk Target Bandwidth (Featured metric)	Baseline bytes per second throughput Data Disk can achieve without bursting.
Data Disk Target IOPS (Featured metric)	Baseline IOPS Data Disk can achieve without bursting.
Data Disk Max Burst Bandwidth (Featured metric)	Maximum bytes per second throughput Data Disk can achieve with bursting.
Data Disk Max Burst IOPS (Featured metric)	Maximum IOPS Data Disk can achieve with bursting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure VM Disk metrics (continued)

Metric Name	Description
Data Disk Used Burst BPS Credits Percentage	Percentage of Data Disk burst bandwidth credits used so far.
Data Disk Used Burst IO Credits Percentage	Percentage of Data Disk burst I/O credits used so far.
Premium Data Disk Cache Read Hit (Featured metric)	Premium Data Disk Cache Read Hit.
Premium Data Disk Cache Read Miss (Featured metric)	Premium Data Disk Cache Read Miss.
Premium OS Disk Cache Read Hit (Featured metric)	Premium OS Disk Cache Read Hit.
Premium OS Disk Cache Read Miss (Featured metric)	Premium OS Disk Cache Read Miss.

Azure VM Flow metrics

Metric Name	Description
Inbound Flows (Featured metric)	Inbound Flows are number of current flows in the inbound direction (traffic going into the VM).
Outbound Flows (Featured metric)	Maximum creation rate of outbound flows (traffic going out of the VM).
Inbound Flows Maximum Creation Rate (Featured metric)	Maximum creation rate of inbound flows (traffic going into the VM).
Outbound Flows Maximum Creation Rate (Featured metric)	Maximum creation rate of outbound flows (traffic going out of the VM).

Azure VM Memory metrics

Metric Name	Description
VM Cached Bandwidth Consumed Percentage (Featured metric)	Percentage of cached disk bandwidth consumed by the VM.
VM Cached IOPS Consumed Percentage	Percentage of cached disk IOPS consumed by the VM.
VM Uncached Bandwidth Consumed Percentage	Percentage of uncached disk bandwidth consumed by the VM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure VM Memory metrics (continued)

Metric Name	Description
(Featured metric)	
VM Uncached Bandwidth Consumed Percentage	Percentage of uncached disk IOPS consumed by the VM.
Available Memory Byte (Featured metric)	Amount of physical memory, in bytes, immediately available for allocation to a process or for system use in the Virtual Machine.

Azure VM Metrics policy - Azure Metrics Collector check metrics

Metric name	Units	Description
Percentage CPU (featured metric)	%	The percentage of allocated compute units that are currently in use by the Virtual Machine.
CPU Credits Remaining		Number of earned CPU credits that have not been consumed by the Virtual Machine.
CPU Credits Consumed		Total number of CPU credits that have been consumed by the Virtual Machine.
Network In		Number of bytes received on all network interfaces by the instance.
Network Out		Number of bytes sent out on all network interfaces by the instance.
Network In Total (featured metric)	Bytes	The number of bytes received on all network interfaces by the Virtual Machine(s) (Incoming Traffic).
Network Out Total (featured metric)	Bytes	The number of bytes out on all network interfaces by the Virtual Machine(s) (Outgoing Traffic).
Available Memory Bytes (featured metric)	Bytes	Amount of physical memory in bytes immediately available for allocation to a process or for system use in the Virtual Machine.
VM Cached Bandwidth Consumed Percentage		Percentage of cached disk bandwidth consumed by the VM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure VM Metrics policy - Azure Metrics Collector check metrics (continued)

Metric name	Units	Description
VM Cached IOPS Consumed Percentage		Percentage of cached disk IOPS consumed by the VM.
VM Uncached Bandwidth Consumed Percentage		Percentage of uncached disk bandwidth consumed by the VM.
VM Uncached IOPS Consumed Percentage		Percentage of uncached disk IOPS consumed by the VM.
Disk Read Bytes (featured metric)	Bytes	Bytes read from disk during monitoring period.
Disk Write Bytes (featured metric)	Bytes	Bytes written to disk during monitoring period.
Disk Read Operations/Sec (featured metric)	CountPerSecond	Disk Read IOPS.
Disk Write Operations/Sec (featured metric)	CountPerSecond	Disk Write IOPS.
OS Disk Read Bytes/sec		Bytes/Sec read from disk during monitoring period.
OS Disk Write Bytes/sec		Bytes/Sec written to disk during monitoring period.
OS Disk Read Operations/Sec		Read IOPS from disk during monitoring period.
OS Disk Write Operations/Sec		Write IOPS from disk during monitoring period.
OS Disk Queue Depth		Average number of read and write requests that were queued for the selected disk during the sample period.
OS Disk Bandwidth Consumed Percentage		Percentage of data disk bandwidth consumed per minute.
OS Disk IOPS Consumed Percentage		Percentage of data disk I/Os consumed per minute.
OS Disk Target Bandwidth		Baseline bytes per second throughput data disk can achieve without bursting.
OS Disk Target IOPS		Baseline IOPS data disk can achieve without bursting.
OS Disk Max Burst Bandwidth		Maximum bytes per second throughput data disk can achieve with bursting.
OS Disk Max Burst IOPS		Maximum IOPS data disk can achieve with bursting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure VM Metrics policy - Azure Metrics Collector check metrics (continued)

Metric name	Units	Description
OS Disk Used Burst BPS Credits Percentage		Percentage of data disk burst bandwidth credits used so far.
OS Disk Used Burst IO Credits Percentage		Percentage of data disk burst I/O credits used so far.
Premium OS Disk Cache Read Hit		Premium OS disk cache read hit.
Premium OS Disk Cache Read Miss		Premium OS disk cache read miss.
Inbound Flows		Number of current flows in the inbound direction (traffic going into the VM).
Outbound Flows		Number of current flows in the inbound direction (traffic going out of the VM).
Inbound Flows Maximum Creation Rate		Maximum creation rate of inbound flows (traffic going into the VM).
Outbound Flows Maximum Creation Rate		Maximum creation rate of outbound flows (traffic going out of the VM).

Azure Load Balancer metrics policy - Azure Metrics collector check metrics

Metric name	Units	Description
AllocatedSnatPorts (featured metric)		Total number of SNAT ports allocated within the time period specified by the metric definition's *PT* (Period of Time) value.
ByteCount (featured metric)	bytes	Total number of bytes transmitted within the time period specified by the metric definition's *PT* (Period of Time) value.
DipAvailability (featured metric)	count	Average load balancer health probe status per time duration.
BlockedCount (featured metric)		Total number of packets transmitted within the time period specified by the metric definition's *PT* (Period of Time) value.
SnatConnectionCount (featured metric)	count	Total number of new SNAT connections created within the time period specified by

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Load Balancer metrics policy - Azure Metrics collector check metrics (continued)

Metric name	Units	Description
		the metric definition's *PT* (Period of Time) value.
SYNCount (featured metric)		Total number of SYN packets transmitted within the time period specified by the metric definition's *PT* (Period of Time) value.
UsedSnatPorts (featured metric)		Total number of SNAT ports used within the time period specified by the metric definition's *PT* (Period of Time) value.
VipAvailability (featured metric)	count	Average load balancer data path availability per time duration.

Azure Redis Cache policy - Azure Metrics Collector check metrics

Metric name	Units	Description
connectedclients		Number of client connections to the cache.
totalcommandsprocessed		Total number of commands processed by the cache server.
cachehits (featured metric)	count	Number of successful key lookups.
cachemisses (featured metric)	count	Number of failed key lookups.
cachemissrate		Rate of failed key lookups.
getcommands		Number of get operations from the cache.
setcommands (featured metric)	count	Number of set operations to the cache.
operationsPerSecond		Number of instantaneous operations per second executed on the cache.
evictedkeys		Number of items evicted from the cache.
totalkeys		Number of keys in the cache.
expiredkeys		Number of items expired from the cache.
usedmemory	bytes	Amount of cache memory used for key/value pairs in the cache.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Redis Cache policy - Azure Metrics Collector check metrics (continued)

Metric name	Units	Description
usedmemoryRss	bytes	Amount of cache memory used during the specified reporting interval, including fragmentation and metadata
serverLoad (featured metric)	Percent	Percentage of cycles in which the Redis server is busy processing and not waiting idle for messages
cacheWrite (featured metric)	BytesPersSecond	Amount of data written to the cache.
cacheRead (featured metric)	BytesPerSecond	Amount of data read from the cache.
allConnectionsCreatedPerSecond		Number of instantaneous connections created per second on the cache via port 6379 or 6380 (SSL).
allConnectionsClosedPerSecond		Number of instantaneous connections closed per second on the cache via port 6379 or 6380 (SSL).
allconnectedclients (featured metric)	count	Number of currently connected clients to the system.
alltotalcommandsprocessed		The cumulative count of all commands processed by the system.
alloperationsPerSecond		Total number of operations (such as read or write commands) performed per second across all instances.
allusedmemory (featured metric)	bytes	Total amount of memory being used across the system.
allusedmemoryRss (featured metric)	bytes	The total Resident Set Size (RSS), which reflects the amount of memory physically present in RAM across all instances, including swap.

Azure Gateway Metrics policy - Azure Metrics collector check metrics

Metric name	Units	Description
ApplicationGatewayTotalTime (featured metric)	MilliSeconds	Time that it takes for a request to be processed and its response to be sent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Gateway Metrics policy - Azure Metrics collector check metrics (continued)

Metric name	Units	Description
AvgRequestCountPerHealthyHost		Average request count per minute per healthy backend host in a pool.
AzwafBotProtectionMatches		Matched Bot Rules.
AzwafCustomRuleMatches		Matched Custom Rules.
AzwafSecRuleMatches		Matched Managed Rules.
AzwafTotalRequests		Total number of requests evaluated by WAF.
BackendConnectTime		Time spent establishing a connection with a backend server.
BackendFirstByteResponseTime		Time interval between start of establishing a connection to backend server and receiving the first byte of the response header, approximating processing time of backend server.
BackendLastByteResponseTime		Time interval between start of establishing a connection to backend server and receiving the last byte of the response body.
BackendResponseStatus		Number of HTTP response codes generated by the backend members. This does not include any response codes generated by the Application Gateway.
BlockedCount		Total number of blocked operations or requests
BytesReceived (featured metric)	bytes	Total number of bytes received by the Application Gateway from the clients.
BytesSent (featured metric)	bytes	Total number of bytes sent by the Application Gateway to the clients.
CapacityUnits		Capacity units consumed.
ClientRtt		Round trip time between clients and Application Gateway. Indicates how long it takes to establish connections and return acknowledgments.
ComputeUnits		Compute units consumed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Gateway Metrics policy - Azure Metrics collector check metrics (continued)

Metric name	Units	Description
CpuUtilization		Current CPU utilization of the Application Gateway.
CurrentConnections (featured metric)	count	Count of current connections established with Application Gateway.
EstimatedBilledCapacityUnits		Estimated capacity units that will be charged.
FailedRequests (featured metric)	count	Count of failed requests that Application Gateway has served.
FixedBillableCapacityUnits		Fixed Billable Capacity Units.
HealthyHostCount (featured metric)	count	Number of healthy backend hosts.
NewConnectionsPerSecond		New connections per second established with Application Gateway.
ResponseStatus		HTTP response status returned by Application Gateway.
Throughput (featured metric)	bytesPerSecond	Number of bytes per second the Application Gateway has served.
TotalRequests (featured metric)	count	Count of successful requests that Application Gateway has served.
UnhealthyHostCount (featured metric)	count	Number of unhealthy backend hosts.

Azure Storage Account policy - Azure Metrics Collector check metrics

Metric name	Units	Description
Transactions (featured metric)	count	Total number of requests made to the storage account.
Ingress (featured metric)	bytes	Amount of ingress data. This number includes ingress from an external client into Azure Storage as well as ingress within Azure.
Egress (featured metric)	bytes	Amount of egress data. This number includes egress from an external client into Azure Storage as well as egress within Azure. As a result,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Storage Account policy - Azure Metrics Collector check metrics (continued)

Metric name	Units	Description
		this number does not reflect billable egress.
SuccessServerLatency (featured metric)	milliseconds	Average latency used by Azure Storage to process a successful request.
SuccessE2ELatency (featured metric)	milliseconds	The average end-to-end latency of successful requests made to a storage service or the specified API operation.
Availability (featured metric)	percent	The percentage of availability for the storage service or the specified API operation.

Azure Blob Services metrics

Metric Name	Description
Availability	Percentage of availability for the storage service or the specified API operation. Availability is calculated by taking the TotalBillableRequests value and dividing it by the number of applicable requests, including those that produced unexpected errors. All unexpected errors result in reduced availability for the storage service or the specified API operation.
BlobCapacity (Featured metric)	Amount of storage used by the storage account's Blob service, in bytes.
BlobCount	Number of Blobs in the storage account's Blob service
Container Count	Number of containers in the storage account's Blob service.
Egress (Featured metric)	Amount of Egress data, in bytes. Includes egress from an external client into Azure Storage as well as Egress within Azure. As a result, this number does not reflect billable Egress.
Index Capacity (Featured metric)	Amount of storage used by ADLS Gen2 (Hierarchical) Index, in bytes.
Ingress (Featured metric)	The amount of Ingress data, in bytes. Includes Ingress from an external client into Azure storage, as well as ingress within Azure.
SuccessE2ELatency	End-to-end latency of successful requests made to a storage service or the specified

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure Blob Services metrics (continued)

Metric Name	Description
(Featured metric)	API operation, in milliseconds. Includes the required processing time within Azure storage to read the request, send the response, and receive acknowledgment of the response.
SuccessServerLatency (Featured metric)	Latency used by Azure storage to process a successful request, in milliseconds. This value does not include the network latency specified in **SuccessE2ELatency**.
Transactions	Number of requests made to a storage service or the specified API operation. Includes successful and failed requests, as well as requests which produced errors. Use the *ResponseType* dimension for the number of different types of responses.

Azure cloud metrics

Azure cloud metrics are gathered from Azure virtual machines (VMs) and Azure storage account policies. Collecting the cloud metrics enables you to monitor the performance of your Azure resources.

Azure cloud metrics collection - Overview

To begin collecting Azure cloud metrics, you must select the metrics that you want to add to the CMDB and set the relevant system configurations. For example, you can monitor the number of bytes sent and received from Azure resources. For details, see Agent Client Collector for Visibility Content reference.

Using Azure policies

You can use preconfigured Azure policies for metric collection, or you can create your own Azure policies.

When using a preconfigured policy, you do the following tasks:

- Prepare an instance with an active MID Server and metric base.
- Install Agent Client Collector Framework and Agent Client Collector Monitoring.
- Run Discovery on the Azure cloud account to populate the CMDB with the configuration items (CIs) to be monitored.
- Configure a proxy agent and assign credentials to the policy. Proxy agents monitor CIs for agents that are in the cloud. Credentials enable the agent to access the Azure policy.

The policies available with the base system for you to activate and configure are described in the following table.

Azure policies available with base system

Policy name	Description	Link to supported metrics
Azure Load Balancer metrics	Collects the metrics for the Azure load balancers.	Supported metrics for Microsoft Load Balancers

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Azure policies available with base system (continued)

Policy name	Description	Link to supported metrics
Azure Gateway metrics	Collects the metrics for the Azure network gateway.	Supported metrics for Microsoft Network Gateways
Azure Redis Cache metrics	Collects the metrics for the Azure Redis cache.	Supported metrics for Microsoft Redis cache
Azure Storage Account metrics	Collects the metrics for the Azure storage accounts.	Supported metrics for Microsoft storage accounts
Azure Virtual Machine (VM) metrics	Collects the metrics for the VMs, which help you monitor Azure VMs.	Supported metrics for Microsoft virtual machines

Specifying the CIs for metric collection

When creating policies, you must include the CIs that you want to collect the metrics for in the CMDB. The CIs are typically populated during Discovery, which finds applications and devices on your network and updates the CMDB accordingly.

When creating a policy, you must decide which CIs you want the policy to monitor. To enable monitoring CIs, you do the following tasks:

1. Configure the Azure Metrics Collector check in the policy.
2. Set the credentials for the Azure Metrics Collector check to enable retrieving metrics from Azure to the Agent Client Collector agent.
3. Generate a .json file that specifies the list of resources to be monitored.
4. Generate a .json file that specifies the list of metrics to be retrieved from monitored resources.

For details on the Azure policies, see Create Azure policies for metric collection.

For details on the Azure Metrics Batch API, see the Metrics Batch API page.

Cassandra default checks and policies

Agent Client Collector provides the following policies for Cassandra health monitoring. Policies come with the checks specified in the indicated table. Policies and checks are available for Linux only.

Cassandra Metrics

Check	Description	Usage	Command	Output
app.cassandra.metrics-cassandra	Returns metrics of a Cassandra database.	`metrics-cassandra-graphite.rb` Options: **-h host** Hostname where the Cassandra	`metrics-cassandra-graphite.rb --cfstats -P {{.labels.params_port}} -h {{.labels.params_host}}`	`<host>.cassandra.load 132270.08 1645107221`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cassandra Metrics (continued)

Check	Description	Usage	Command	Output
		database is running **-P port** Cassandra JMX port		

Cassandra metrics

The following table lists the metrics that are gathered as output from Cassandra checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

metrics-cassandra metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.load (featured metric)		byte	The disk space used by live data on a node.
cassandra.uptime (featured metric)		seconds	
cassandra.heap.used (featured metric)		byte	
cassandra.heap.total (featured metric)		byte	
cassandra.exceptions		count	Number of internal exceptions caught. Under normal exceptions this should be zero.
cassandra.key_cache.size		byte	Total size of occupied cache, in bytes.
cassandra.key_cache.capacity		byte	Cache capacity in bytes.
cassandra.key_cache.hits		count	Total number of cache hits.
cassandra.key_cache.requests (featured metric)		count	Total number of cache requests.
cassandra.key_cache.hit_rate (featured metric)		count	All time cache hit rate.
cassandra.compactionstats.pending_tasks		count	The number of pending compactions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.threadpool.ReadStage.active		count	Number of tasks being actively worked on by this pool. Local reads run on this thread pool on this ReadStage.
cassandra.threadpool.ReadStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.ReadStage.completed		count	Number of tasks completed.
cassandra.threadpool.ReadStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.MiscStage.active		count	Number of tasks being actively worked on by this pool. Miscellaneous tasks run on this stage.
cassandra.threadpool.MiscStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.MiscStage.completed		count	Number of tasks completed.
cassandra.threadpool.MiscStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.CompactionExecutor.active		count	Number of tasks being actively worked on by this pool. Compactions are run on these threads.
cassandra.threadpool.CompactionExecutor.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.CompactionExecutor.completed		count	Number of tasks completed.
cassandra.threadpool.CompactionExecutor.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.MutationStage.active		count	Number of tasks being actively worked on by this pool. Responsible for all other writes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.threadpool.MutationStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.MutationStage.completed		count	Number of tasks completed.
cassandra.threadpool.MutationStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.MemtableReclaimMemory.active		count	Number of tasks being actively worked on by this pool. Memtable recycling.
cassandra.threadpool.MemtableReclaimMemory.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.MemtableReclaimMemory.completed		count	Number of tasks completed.
cassandra.threadpool.MemtableReclaimMemory.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.PendingRangeCalculator.active		count	Number of tasks being actively worked on by this pool. Calculates token range.
cassandra.threadpool.PendingRangeCalculator.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.PendingRangeCalculator.completed		count	Number of tasks completed.
cassandra.threadpool.PendingRangeCalculator.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.GossipStage.active		count	Number of tasks being actively worked on by this pool. Handles gossip requests.
cassandra.threadpool.GossipStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.GossipStage.completed		count	Number of tasks completed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.threadpool.GossipStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.SecondaryIndexManagement.active		count	Number of tasks being actively worked on by this pool. Performs updates to secondary indexes.
cassandra.threadpool.SecondaryIndexManagement.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.SecondaryIndexManagement.completed		count	Number of tasks completed.
cassandra.threadpool.SecondaryIndexManagement.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.HintsDispatcher.active		count	Number of tasks being actively worked on by this pool. Performs hinted hand off.
cassandra.threadpool.HintsDispatcher.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.HintsDispatcher.completed		count	Number of tasks completed.
cassandra.threadpool.HintsDispatcher.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.RequestResponseStage.active		count	Number of tasks being actively worked on by this pool. Coordinator requests to the cluster run on this thread pool.
cassandra.threadpool.RequestResponseStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.RequestResponseStage.completed		count	Number of tasks completed.
cassandra.threadpool.RequestResponseStage.blocked		count	Number of tasks that were blocked due to queue saturation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.threadpool.ReadRepairStage.active		count	Number of tasks being actively worked on by this pool. Read repair happens on this thread pool.
cassandra.threadpool.ReadRepairStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.ReadRepairStage.completed		count	Number of tasks completed.
cassandra.threadpool.ReadRepairStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.CounterMutationStage.active		count	Number of tasks being actively worked on by this pool. Responsible for counter writes.
cassandra.threadpool.CounterMutationStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.CounterMutationStage.completed		count	Number of tasks completed.
cassandra.threadpool.CounterMutationStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.MigrationStage.active		count	Number of tasks being actively worked on by this pool. Runs schema migrations.
cassandra.threadpool.MigrationStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.MigrationStage.completed		count	Number of tasks completed.
cassandra.threadpool.MigrationStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.MemtablePostFlush.active		count	Number of tasks being actively worked on by this pool. Cleans up commit log after memtable is written to disk.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.threadpool.MemtablePostFlush.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.MemtablePostFlush.completed		count	Number of tasks completed.
cassandra.threadpool.MemtablePostFlush.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.PerDiskMemtableFlushWriter_0.active		count	Number of tasks being actively worked on by this pool. Responsible for writing a spec (there is one of these per disk 0-N).
cassandra.threadpool.PerDiskMemtableFlushWriter_0.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.PerDiskMemtableFlushWriter_0.completed		count	Number of tasks completed.
cassandra.threadpool.PerDiskMemtableFlushWriter_0.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.ValidationExecutor.active		count	Number of tasks being actively worked on by this pool. Performs validation compaction or scrubbing.
cassandra.threadpool.ValidationExecutor.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.ValidationExecutor.completed		count	Number of tasks completed.
cassandra.threadpool.ValidationExecutor.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.Sampler.active		count	Number of tasks being actively worked on by this pool. Responsible for re-sampling the index summaries of *`SStables`*.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.threadpool.Sampler.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.Sampler.completed		count	Number of tasks completed.
cassandra.threadpool.Sampler.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.MemtableFlushWriter.active		count	Number of tasks being actively worked on by this pool. Writes memtables to disk.
cassandra.threadpool.MemtableFlushWriter.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.MemtableFlushWriter.completed		count	Number of tasks completed.
cassandra.threadpool.MemtableFlushWriter.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.InternalResponseStage.active		count	Number of tasks being actively worked on by this pool. Responsible for intra-cluster callbacks.
cassandra.threadpool.InternalResponseStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.InternalResponseStage.completed		count	Number of tasks completed.
cassandra.threadpool.InternalResponseStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.ViewMutationStage.active		count	Number of tasks being actively worked on by this pool. Responsible for materialized view writes.
cassandra.threadpool.ViewMutationStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.ViewMutationStage.completed		count	Number of tasks completed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.threadpool.ViewMutationStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.AntiEntropyStage.active		count	Number of tasks being actively worked on by this pool. Builds merkle tree for repairs.
cassandra.threadpool.AntiEntropyStage.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.AntiEntropyStage.completed		count	Number of tasks completed.
cassandra.threadpool.AntiEntropyStage.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.threadpool.CacheCleanupExecutor.active		count	Number of tasks being actively worked on by this pool. Cache maintenance performed on this thread pool.
cassandra.threadpool.CacheCleanupExecutor.pending		count	Number of queued tasks queued up on this pool.
cassandra.threadpool.CacheCleanupExecutor.completed		count	Number of tasks completed.
cassandra.threadpool.CacheCleanupExecutor.blocked		count	Number of tasks that were blocked due to queue saturation.
cassandra.message_type.READ.dropped		count	Number of dropped messages. Regular reads.
cassandra.message_type.RANGE_SLICE.dropped		count	Number of dropped messages. Token range read.
cassandra.message_type._TRACE.dropped		count	Number of dropped messages. Tracing writes.
cassandra.message_type.HINT.dropped		count	Number of dropped messages. Hint replay.
cassandra.message_type.MUTATION.dropped		count	Number of dropped messages. Regular writes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.message_type.COUNTER_MUTATION.dropped		count	Number of dropped messages. Counter writes.
cassandra.message_type.BATCH_STORE.dropped		count	Number of dropped messages. Batchlog write.
cassandra.message_type.BATCH_REMOVE.dropped		count	Number of dropped messages. Batchlog cleanup (after successfully applied).
cassandra.message_type.REQUEST_RESPONSE.dropped		count	Number of dropped messages. RPC Callbacks.
cassandra.message_type.PAGED_RANGE.dropped		count	Number of dropped messages. Paged read.
cassandra.message_type.READ_REPAIR.dropped		count	Number of dropped messages. Read repair.
cassandra.system_traces.read_count		count	Number of read for this table.
cassandra.system_traces.read_latency			Local read latency for this table.
cassandra.system_traces.write_count		count	Number of writes for this table.
cassandra.system_traces.write_latency			Local write latency for this table.
cassandra.system_traces.pending_flushes		count	Estimated number of flush tasks pending for this table.
cassandra.system.read_count (featured metric)		count	
cassandra.system.read_latency (featured metric)		microseconds	
cassandra.system.write_count (featured metric)		count	
cassandra.system.write_latency (featured metric)		microseconds	
cassandra.system.pending_flushes		count	
cassandra.system_distributed.read_count		count	
cassandra.system_distributed.read_latency			

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-cassandra metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cassandra.system_distributed.write_count		count	
cassandra.system_distributed.write_latency			
cassandra.system_distributed.pending_flushes		count	
cassandra.system_schema.read_count		count	
cassandra.system_schema.read_latency			
cassandra.system_schema.write_count		count	
cassandra.system_schema.write_latency			
cassandra.system_schema.pending_flushes		count	
cassandra.system_auth.read_count		count	
cassandra.system_auth.read_latency			
cassandra.system_auth.write_count		count	
cassandra.system_auth.write_latency			
cassandra.system_auth.pending_flushes		count	

Directory Scan monitoring default checks and policies

The Agent Client Collector provides the following default checks and policies for Directory Scan monitoring.

Windows OS Events - Extended policy

Type	Check	Description	Usage and Usage Example	Output
Event	check-directory-file-count	Checks the number of files present in a given directory and compares it to the provided warning and critical thresholds. Returns a CRITICAL, WARNING, or OK event based on the critical and warning thresholds.	check-directory-file-count.rb (options) -d, --dir Absolute path to the directory to count the files (required) -w, --warning Warning threshold (required) -c, --critical Critical threshold (required)	DirectoryFileCount CRITICAL: <path to dir> has 165 files.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
		Note: • Counts the number of files inside the directory, not including files in subdirectories. • The agent requires read and execute permissions on the directory being monitored.	-H, include_hidden_files Set active to **true** to include hidden files while counting (default is **false**) **Usage Example:** `check-directory-file-count.rb -d /path/to/directory -w 50 -c 100`	
Event	check-directory-integrity	Compares the last modified time of the directory with the current time to determine if any updates have occurred within a defined time interval. Based on this comparison, the check returns a CRITICAL or OK event. **Note:** The agent requires read and execute permissions on the directory being monitored.	commonchecks check-directory-integrity (options) -d, --dirpath DIRPATH: Absolute path to the directory to check (required). -i, --interval INTERVAL: Time interval in seconds to check for recent updates (required). Default value can be set to 180 seconds. **Usage example:** `commonchecks check-directory-update -d /path/`	Common Checks CRITICAL: Directory Integrity: CHANGES DETECTED for <path to file> within the last 180 seconds (Last Modified: Wed, 04 Dec 2024 12:18:55 EST).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			`to/directory -i 180`	
Event	check-file-age	Evaluates the age of a specified file by comparing its last modification time with the current time. Raises an alert if the file exceeds the defined critical or warning age thresholds. **Note:** The agent requires read and execute permissions on the file being monitored.	commonchecks check-file-age (options) -f, --filepath FILEPATH: Absolute path to the file to check (required). -c, --critical CRITICAL: Critical age threshold in minutes for the file (required). -w, --warning WARNING: Warning age threshold in minutes for the file (required). **Usage example:** `commonchecks check-file-age -f /path/to/file.txt -c 120 -w 60`	Common Checks OK: File <path to file> age: 30 minutes.
Event	check-file-response-time	Compares the time needed to read a specified file and compares it with the critical and warning thresholds. Based on this comparison, the check returns a CRITICAL or OK event.	commonchecks check-file-response-time (options) -f, --filepath FILEPATH: Absolute path to the file to check (required). -c, --critical CRITICAL: Critical threshold in seconds for file response time (required).	Common Checks OK: File read response time: 0.0020 seconds for the file_name: /var/log/servicenow/agent-client-collector/acc.log.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
		Note: The agent requires read and execute permissions on the file being monitored.	-w, --warning WARNING: Warning threshold in seconds for file response time (required). -t, --timeout TIMEOUT: Maximum time allowed for reading the file content, specified in seconds. Time out value must be greater than the critical threshold value. **Usage example:** `commonchecks check-file-response-time -f /var/log/servicenow/agent-client-collector/acc.log -c 10 -w 5 -t 20`	
Event	check-file-size	Measures the size of a file (the actual amount of data it contains) and compares it against specified thresholds. Returns a CRITICAL, WARNING, or OK event based on the comparison of the file size and the thresholds.	commonchecks check-file-size (options) -f, --filepath FILEPATH Absolute path to the required file (required). -c, --critical CRITICAL. Critical threshold in kilobytes. Provide as a number without units (for example, 1000	Common Checks OK: For File <path to file> size: 4.72 KB is within thresholds

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
		Note: The agent requires read and execute permissions on the file being monitored.	for 1000 KB) (required) -w, --warning WARNING. Warning threshold in kilobytes. Provide as a number without units (for example, 500 for 500 KB) (required) **Usage example:** `commonchecks check-file-size -f C:\ProgramData\ServiceNow\agent-client-collector\config\acc.yml`	
Event	check-file-space	Measures the size of a file on disk against specified thresholds, returning a CRITICAL, WARNING, or OK event based on the thresholds. **Note:** The agent requires read and execute permissions on the file being monitored.	commonchecks check-file-space (options) -f, --filepath FILEPATH: Absolute path to the relevant file (required). -c, --critical CRITICAL: Critical threshold in kilobytes. Provided as a number without units (for example, 1000 for 1000 KB) (required) -w, --warning WARNING: Warning threshold in kilobytes. Provided as	Common Checks OK: File <path to file> space: 8.00 KB is within threshold

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			a number without units (for example, 500 for 500 KB) (required) -b, --blocksize BLOCKSIZE: Block size in bytes (Default = 4096) (required) **Usage example:** `commonchecks check-file-space -f C:\ProgramData\ServiceNow\agent-client-collector\config\acc.yml`	
Event	os.windows.check-directory-space	Verifies the disk space occupied by a directory's content. Returns a CRITICAL, WARNING, or OK event, based on the comparison with the given critical and warning event severity thresholds. **Note:** The agent requires read and execute permissions on the directory being monitored.	winchecks check-dir-space (options) -d, --dirpath DIRPATH: Absolute path to the directory being checked (required). -c, --critical CRITICAL: Critical disk space threshold in kilobytes (required). -w, --warning WARNING: Warning disk space threshold in kilobytes (required). -t, --timeout TIMEOUT: Maximum time allowed for	Windows Checks OK: Directory <path to file> space: 369.25 KB is within thresholds

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			directory size calculation, specified in seconds (required). **Usage example:** `winchecks check-dir-space -d /path/to/directory -c 1000 -w 500 -t 120`	

Linux OS Events - Extended policy

Type	Check	Description	Usage and Usage Example	Output
Event	check-directory-file-count	Checks the number of files present in a given directory and compares it to the provided warning and critical thresholds. Returns a CRITICAL, WARNING, or OK event based on the critical and warning thresholds.	check-directory-file-count.rb (options) -d, --dir Absolute path to the directory to count the files (required) -w, --warning Warning threshold (required) -c, --critical Critical threshold (required)	DirectoryFileCount CRITICAL: <path to dir> has 165 files.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
		Note: • Counts the number of files inside the directory, not including files in subdirectories. • The agent requires read and execute permissions on the directory being monitored.	-H, include_hidden_files Set active to **true** to include hidden files while counting (default is **false**) **Usage Example:** `check-directory-file-count.rb -d /path/to/directory -w 50 -c 100`	
Event	check-directory-integrity	Compares the last modified time of the directory with the current time to determine if any updates have occurred within a defined time interval. Based on this comparison, the check returns a CRITICAL or OK event. **Note:** The agent requires read and execute permissions on the directory being monitored.	commonchecks check-directory-integrity (options) -d, --dirpath DIRPATH: Absolute path to the directory to check (required). -i, --interval INTERVAL: Time interval in seconds to check for recent updates (required). Default value can be set to 180 seconds. **Usage example:** `commonchecks check-directory-update -d /path/`	Common Checks CRITICAL: Directory Integrity: CHANGES DETECTED for <path to file> within the last 180 seconds (Last Modified: Wed, 04 Dec 2024 12:18:55 EST).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			`to/directory -i 180`	
Event	check-file-age	Evaluates the age of a specified file by comparing its last modification time with the current time. Raises an alert if the file exceeds the defined critical or warning age thresholds. **Note:** The agent requires read and execute permissions on the file being monitored.	commonchecks check-file-age (options) -f, --filepath FILEPATH: Absolute path to the file to check (required). -c, --critical CRITICAL: Critical age threshold in minutes for the file (required). -w, --warning WARNING: Warning age threshold in minutes for the file (required). **Usage example:** `commonchecks check-file-age -f /path/to/file.txt -c 120 -w 60`	Common Checks OK: File <path to file> age: 30 minutes.
Event	check-file-response-time	Compares the time needed to read a specified file and compares it with the critical and warning thresholds. Based on this comparison, the check returns a CRITICAL or OK event.	commonchecks check-file-response-time (options) -f, --filepath FILEPATH: Absolute path to the file to check (required). -c, --critical CRITICAL: Critical threshold in seconds for file response time (required).	Common Checks OK: File read response time: 0.0020 seconds for the file_name: /var/log/servicenow/agent-client-collector/acc.log.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
		Note: The agent requires read and execute permissions on the file being monitored.	-w, --warning WARNING: Warning threshold in seconds for file response time (required). -t, --timeout TIMEOUT: Maximum time allowed for reading the file content, specified in seconds. Time out value must be greater than the critical threshold value. **Usage example:** `commonchecks check-file-response-time -f /var/log/servicenow/agent-client-collector/acc.log -c 10 -w 5 -t 20`	
Event	check-file-size	Measures the size of a file (the actual amount of data it contains) and compares it against specified thresholds. Returns a CRITICAL, WARNING, or OK event based on the comparison of the file size and the thresholds.	commonchecks check-file-size (options) -f, --filepath FILEPATH Absolute path to the required file (required). -c, --critical CRITICAL. Critical threshold in kilobytes. Provide as a number without units (for example, 1000	Common Checks OK: For File <path to file> size: 4.72 KB is within thresholds

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
		Note: The agent requires read and execute permissions on the file being monitored.	for 1000 KB) (required) -w, --warning WARNING. Warning threshold in kilobytes. Provide as a number without units (for example, 500 for 500 KB) (required) **Usage example:** `commonchecks check-file-size -f C:\ProgramData\ServiceNow\agent-client-collector\config\acc.yml`	
Event	check-file-space	Measures the size of a file on disk against specified thresholds, returning a CRITICAL, WARNING, or OK event based on the thresholds. **Note:** The agent requires read and execute permissions on the file being monitored.	commonchecks check-file-space (options) -f, --filepath FILEPATH: Absolute path to the relevant file (required). -c, --critical CRITICAL: Critical threshold in kilobytes. Provided as a number without units (for example, 1000 for 1000 KB) (required) -w, --warning WARNING: Warning threshold in kilobytes. Provided as	Common Checks OK: File <path to file> space: 8.00 KB is within threshold

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			a number without units (for example, 500 for 500 KB) (required) -b, --blocksize BLOCKSIZE: Block size in bytes (Default = 4096) (required) **Usage example:** `commonchecks check-file-space -f C:\ProgramData\ServiceNow\agent-client-collector\config\acc.yml`	
Event	os.linux.check-directory-size	Verifies the space allocated for a disk's directory and compares it against specified critical and warning thresholds. **Note:** The agent requires read and execute permissions on the directory being monitored.	linuxchecks check-directory-size (options) -d, --dirpath DIRPATH: Absolute path to the directory being checked (required). -c, --critical CRITICAL: Critical disk space threshold in kilobytes (required). -w, --warning WARNING: Warning disk space threshold in kilobytes (required). -t, --timeout TIMEOUT: Maximum time allowed for	Linux Checks CRITICAL: Directory <path to file> size: 500.00 KB exceeds critical threshold 10.00 KB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			directory size calculation, specified in seconds (required). **Usage example:**`linuxchecks check-directory-size -d path/to/directory -c 100 -w 50 -t 30`	
Event	os.linux.check-directory-space	Verifies the disk space occupied by directory content. Returns a CRITICAL, WARNING, or OK event based on the comparison with the critical and warning thresholds. **Note:** The agent requires read and execute permissions on the directory being monitored.	linuxchecks check-directory-space (options) -d, --dirpath DIRPATH: Absolute path to the directory being checked (required). -c, --critical CRITICAL: Critical disk space threshold in kilobytes (required). -w, --warning WARNING: Warning disk space threshold in kilobytes (required). -t, --timeout TIMEOUT: Maximum time allowed for directory size calculation, specified in seconds (required). **Usage example:**`linuxchecks`	Linux Checks CRITICAL: Directory <path to file> space: 374.00 KB exceeds critical threshold 10.00 KB

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux OS Events - Extended policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			`check-directory-space -d /path/to/directory -c 10 -w 5 -t 30`	

Google Cloud Platform (GCP) metrics

The following table lists and describes the metrics that are gathered by the *acc_grp_metrics_list.json* configuration data file. The file is uploaded to a check definition or check instance, and the metrics in the file are monitored by the check for its agent.

GCP metrics

Metric name	Metric type	Service	Description
CPU utilization	appengine.googleapis.com/flex/cpu/utilization	app engine	Fractional utilization of allocated CPU across an App Engine flexible environment version.
Response count	appengine.googleapis.com/http/server/response_count		Delta HTTP response count
Free disk space percent	file.googleapis.com/nfs/server/free_bytes_percent	file	Free disk space as a percentage of the total space.
Circuit Operational Status	interconnect.googleapis.com/network/interconnect/link/operational	interconnect	Whether the operational status of the circuit is 'up'. This metric is specific to individual physical circuits
Operational Status	interconnect.googleapis.com/network/interconnect/operational		Whether the operational status of the interconnect is 'up'. This metric is specific to the aggregate LACP bundle.
Instance group size	compute.googleapis.com/instance_group/size	compute	The number of VMs in the instance group.
CPU utilization	compute.googleapis.com/instance/cpu/utilization		Fractional utilization of allocated CPU on this instance. Values are typically numbers between 0.0 and 1.0 (but can

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP metrics (continued)

Metric name	Metric type	Service	Description
			exceed 1.0). Charts display the values as a percentage between 0% and 100% (or more).
CPU usage time	compute.googleapis.com/instance/cpu/usage_time		Delta CPU usage time in seconds. Cumulative CPU usage time since instance start.
CPU scheduler wait time	compute.googleapis.com/instance/cpu/scheduler_wait_time		Wait time is the time a vCPU is runnable but not scheduled to run by the hypervisor.
CPU reserved cores	compute.googleapis.com/instance/cpu/reserved_cores		Number of cores reserved for the instance.
Disk queue length	compute.googleapis.com/guest/disk/queue_length		Number of requests waiting to be handled by the disk.
Uptime	compute.googleapis.com/instance/uptime		Delta of how long the VM has been running, measured in seconds. For availability monitoring, use uptime checks instead.
Disk read bytes	compute.googleapis.com/instance/disk/read_bytes_count		Delta count of bytes read from disk.
Disk max read operations	compute.googleapis.com/instance/disk/max_read_ops_count		Maximum number of read operations that can be performed per second.
Disk write bytes	compute.googleapis.com/instance/disk/write_bytes_count		Delta count of bytes written to disk.
Disk max write operations	compute.googleapis.com/instance/disk/max_write_ops_count		Maximum number of write operations that can be performed per second.
Disk read operations	compute.googleapis.com/instance/disk/read_ops_count		Delta count of disk read IO operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP metrics (continued)

Metric name	Metric type	Service	Description
Disk write operations	compute.googleapis.com/instance/disk/write_ops_count		Delta count of disk write IO operations.
Memory balloon RAM size	compute.googleapis.com/instance/memory/balloon/ram_size		The total amount of memory allocated to the guest's balloon driver in bytes. Memory ballooning enables the hypervisor to reclaim memory from the guest VM.
Memory balloon RAM used	compute.googleapis.com/instance/memory/balloon/ram_used		Memory currently used by the guest RAM, excluding memory used by the balloon itself, measured in bytes.
Memory swap in	compute.googleapis.com/instance/memory/balloon/swap_in_bytes_count		The amount of memory (in bytes) swapped in from disk to memory.
Memory swap out	compute.googleapis.com/instance/memory/balloon/swap_out_bytes_count		The amount of memory (in bytes) swapped out from memory to disk.
Network received bytes	compute.googleapis.com/instance/network/received_bytes_count		Delta count of bytes received from the network.
Network sent bytes	compute.googleapis.com/instance/network/sent_bytes_count		Delta count of bytes sent over the network.
Network received packets	compute.googleapis.com/instance/network/received_packets_count		Delta count of packets received from the network.
Network sent packets	compute.googleapis.com/instance/network/sent_packets_count		Delta count of packets sent over the network.
Anonymous memory used	compute.googleapis.com/guest/memory/anonymous_used		Anonymous memory usage in bytes. Anonymous memory represents memory that is not backed by a file.
Memory bytes used	compute.googleapis.com/guest/memory/bytes_used		Total memory usage in bytes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP metrics (continued)

Metric name	Metric type	Service	Description
Dirty memory used	compute.googleapis.com/guest/memory/dirty_used		Dirty memory usage in bytes. Dirty pages are pages in memory that have been modified but not yet written to disk.
Page cache used	compute.googleapis.com/guest/memory/page_cache_used		Page cache memory usage in bytes. Page cache is used to cache file data.
Unevictable memory used	compute.googleapis.com/guest/memory/unevictable_used		Unevictable memory usage in bytes. Unevictable memory is memory that cannot be swapped out.
Disk operation bytes	compute.googleapis.com/guest/disk/operation_bytes_count		Delta count of bytes transferred in disk operations.
System problem state	compute.googleapis.com/guest/system/problem_state		Indicates whether the system is experiencing problems. 1 means there is a problem, 0 means no problem.
System uptime	compute.googleapis.com/guest/system/uptime		Delta of system uptime in seconds.
System problem count	compute.googleapis.com/guest/system/problem_count		Delta count of system problems detected.
CPU limit utilization	kubernetes.io/anthos/container/cpu/limit_utilization	anthos	The fraction of the CPU limit that is currently in use on the instance.
Memory limit utilization	kubernetes.io/anthos/container/memory/limit_utilization		The fraction of the memory limit that is currently in use on the instance.
Executions	cloudfunctions.googleapis.com/function/execution_count	cloudfunctions	Count of function executions broken down by status.
Tunnel established	n.googleapis.com/tunnel_established	vpn	Indicates successful tunnel establishment if > 0.
RTT latencies per internal TCP/UDP load balancer	loadbalancing.googleapis.com/l3/internal/rtt_latencies	loadbalancing	A distribution of RTT measured over TCP connections for

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP metrics (continued)

Metric name	Metric type	Service	Description
			internal TCP/UDP load balancer flows.
Unacked messages by region	pubsub.googleapis.com/subscription/num_unacked_messages_by_region	pubsub	Number of unacknowledged messages in a subscription, broken down by Cloud region.
Disabled for network	firebasedatabase.googleapis.com/network/disabled_for_overages	firebasedatabase	Indicates if the Firebase database has been disabled for network overages.
Disabled for storage	firebasedatabase.googleapis.com/storage/disabled_for_overages		Indicates if the Firebase database has been disabled for storage overages.
Database load	firebasedatabase.googleapis.com/io/database_load		Fraction of database load, grouped by type.
CPU utilization	cloudsql.googleapis.com/database/cpu/utilization	cloudsql	Current CPU utilization represented as a percentage of the reserved CPU that is currently in use.
Disk utilization	cloudsql.googleapis.com/database/disk/utilization		The fraction of the disk quota that is currently in use.
Instance state	cloudsql.googleapis.com/database/instance_state		The current serving state of the Cloud SQL instance.
Memory utilization	cloudsql.googleapis.com/database/memory/utilization		The fraction of the memory quota that is currently in use.
Server up	cloudsql.googleapis.com/database/up		Indicates if the server is up.
CPU utilization by priority	spanner.googleapis.com/instance/cpu/utilization_by_priority	spanner	Percent utilization of provisioned CPU, by priority.
Smoothed CPU utilization	spanner.googleapis.com/instance/cpu/smoothed_utilization		24-hour smoothed utilization of provisioned CPU.
Storage used	spanner.googleapis.com/instance/storage/used_bytes		Storage used in bytes. Sampled every 60 seconds.
CPU usage percent	memcache.googleapis.com/node/cpu/utilization	memcache	CPU usage percent by Memcached node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

GCP metrics (continued)

Metric name	Metric type	Service	Description
Cache memory usage	memcache.googleapis.com/node/cachememory		Bytes allotted for the Memcached cache in this node, grouped by whether that memory is used.
Used memory	redis.googleapis.com/redis.googleapis.com/stats/memory/usage_ratio	redis	Total number of bytes allocated by Redis.
Memory usage ratio	redis.googleapis.com/stats/memory/usage_ratio		Memory usage as a ratio of maximum memory.
CPU load	bigtable.googleapis.com/cluster_cpu_load	bigtable	CPU load of a cluster.
Disk load	bigtable.googleapis.com/cluster_disk_load		Utilization of HDD disks in a cluster.
Storage utilization	bigtable.googleapis.com/cluster_storage_utilization		Storage used as a fraction of total storage capacity.
HDFS storage utilization	dataproc.googleapis.com/cluster/hdfs/storage_utilization	dataproc	The percentage of HDFS storage currently used.
Unhealthy HDFS blocks by status	dataproc.googleapis.com/hdfs/unhealthy_blocks		Indicates the number of unhealthy blocks inside the cluster.
Processes	agent.googleapis.com/processes/count_by_state	processes	Count of processes in the given state. Relevant for Linux only.
Process CPU	agent.googleapis.com/processes/cpu_time		Time taken by CPU to complete a given process.
TCP connections	agent.googleapis.com/processes/network	network	Current count of TCP connections.

Glassfish default checks and policies

Agent Client Collector provides the following default checks and policies for Glassfish monitoring.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Glassfish metrics policy

Type	Check	Description	Usage and usage example	Output

Metric	glassfish.metric-jvm-level	Retrieves the Glassfish server's JVM metrics.
		
		Requires basic authentication credentials.
			`-i, --instance Glassfish Monitoring Instance Name (required)`
			`-p, --password Glassfish password`
			`-P, --port Glassfish API port`
			`--ssl Enable secure connnection to Glassfish`
			`-u –username Glassfish username`
			
			Usage example:
			`metric-glassfish-jvm-level.rb {{if .labels.params_port}} -P {{.labels.params_port}} {{end}} -i {{.labels.params_glassfish_server_ins`
				`glassfish.jvm.garbage-collectors.Scavenge.Col 200 0 1652959788`

Metric	glassfish.metric-server-requests	Provides http-request metrics for all applications deployed on the Glassfish server.	`-i, --instance Glassfish Monitoring Instance Name`	`glassfish.http-service.server.request.C 200 0 1652959788`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Glassfish metrics policy (continued)

Type	Check	Description	Usage and usage example	Output

		Requires basic authentication credentials.	`(requir ed) -p, --password Glassfish password -P, --port Glassfish API port --ssl Enable secure connnect ion to Glassfish -u –username Glassfish username`
			
			Usage example:
			`metric-glassfish-server-requests.rb -i {{.labels.params_glassfish_server_ins {{if .labels.params_port}} -P {{.labels.params_port}} {{end}}`	

Metric	glassfish.metric-server-resources	Provides metrics of the resources deployed on the Glassfish server. Requires basic authentication credentials.	`-i, --instance Glassfish Monitor ing Instance Name (requir ed) -P, --port Glassfish API port -r --`	`glassfish.resources. {{PoolName}}.AverageCon 0 1652959936`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Glassfish metrics policy (continued)

Type	Check	Description	Usage and usage example	Output

			`resource_pool Resource Pool Name --ssl Enable secure connnection to Glassfish -u –username Glassfish username -p --password Glassfish password`
			
			Usage example:
			`metric-glassfish-server-resources.rb {{if .labels.params_port}} -P {{.labels.params_port}} {{end}} {{if .labels.params_resource_pool}} -r {{.labels.params_resource_pool}} {{end}} -i {{.labels.params_glassfish_server_ins`	

Metric	glassfish.metric-transaction-service	Provides metrics for all transactions on applications deployed on the Glassfish server. Requires basic authentication credentials.	`-i, --instance Glassfish Monitoring Instance Name (required) -P, --port Glassfish API port`	`glassfish.transaction-service.ActiveCount 0 1652959977`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Glassfish metrics policy (continued)

Type	Check	Description	Usage and usage example	Output
			`--ssl Enable secure connnection to Glassfish -u –username Glassfish username -p --password Glassfish password` Usage example: `metric-glassfish-transaction-service.rb {{if .labels.params_port}} -P {{.labels.params_port}} {{end}} -i {{.labels.params_glassfish_server_ins`	
Metric	glassfish.metrics-deployed-apps	Provides metric output of applications deployed on the Glassfish server. Requires basic authentication credentials.	`-a, --apps Application deployed on Glassfish server (required) -i, --instance Glassfish Monitoring Instance Name (required) -P, --port`	`glassfish.applications.{{Deployed Application}}.Activated 0 1652959665`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Glassfish metrics policy (continued)

Type	Check	Description	Usage and usage example	Output
			`Glassfish API port --ssl Enable secure connnection to Glassfish -u –username Glassfish username -p --password Glassfish password` Usage example: `metric-glassfish-deployed-applications.rb {{if .labels.params_port}} -P {{.labels.params_port}} {{end}} -a {{.labels.params_deployed_apps}} -i {{.labels.params_glassfish_server_ins`	
Metric	glassfish.metrics-deployment-lifecycle	Provides lifecycle metrics of the Glassfish server applications. Requires basic authentication credentials.	`metric-glassfish-deployment-lifecycle.rb (options) -i, --instance Glassfish Monitoring Instance Name (required) -p --password`	`glassfish.deployment.li 2 1653022649`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Glassfish metrics policy (continued)

Type	Check	Description	Usage and usage example	Output
			Glassfish password -P, --port Glassfish API port --ssl Enable secure connnection to Glassfish -u –username Glassfish username Usage example: `metric-glassfish-deployment-lifecycle.rb {{if .labels.params_port}} -P {{.labels.params_port}} {{end}} -i {{.labels.params_glassfish_server_ins`	

GlassFish metrics

The following table lists the metrics that are gathered as output from GlassFish checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

glassfish.metrics-deployed-apps metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.applications.{Deployed Application}.TotalServletsLoaded	Deployed Application	count	Total number of Servlets loaded so far.
glassfish.applications.{Deployed Application}.ServletProcessingTimes	Deployed Application	ms	Cumulative Servlet processing times.
glassfish.applications.{Deployed Application}.ActiveServletsLoaded	Deployed Application	count	Number of Servlets loaded.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glassfish.metrics-deployed-apps metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.applications.{Deployed Application}.ExpiredSessionsTotal (featured metric)	Deployed Application	count	Total number of sessions expired so far.
glassfish.applications.{Deployed Application}.PassivatedSessionsTotal	Deployed Application	count	Total number of sessions passivated so far.
glassfish.applications.{Deployed Application}.ActivatedSessionsTotal	Deployed Application	count	Total number of sessions activatedso far.
glassfish.applications.{Deployed Application}.SessionsTotal	Deployed Application	count	Total number of sessions created so far.
glassfish.applications.{Deployed Application}.RejectedSessionsTotal (featured metric)	Deployed Application	count	Total number of sessions rejected so far.
glassfish.applications.{Deployed Application}.PersistedSessionsTotal	Deployed Application	count	Total number of sessions persisted so far.
glassfish.applications.{Deployed Application}.ActiveSessions (featured metric)	Deployed Application	count	Number of active sessions.
glassfish.applications.{Deployed Application}.RequestCount	Deployed Application	count	Cumulative number of requests processed so far
glassfish.applications.{Deployed Application}.MaxTime (featured metric)	Deployed Application	ms	Longest response time for a request; not a cumulative value, but the largest response time from among the response times.
glassfish.applications.{Deployed Application}.ErrorCount (featured metric)	Deployed Application	count	Cumulative value of the error count, with error count representing the number of cases where the response code was greater than or equal to 400.
glassfish.applications.{Deployed Application}.ProcessingTime	Deployed Application	ms	Average request processing time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glassfish.metrics-deployed-apps metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.applications.{Deployed Application}.TotalJspCount	Deployed Application	count	Total number of JSP pages loaded so far.
glassfish.applications.{Deployed Application}.JspCount	Deployed Application	count	Number of active JSP pages.
glassfish.applications.{Deployed Application}.JspErrorCount (featured metric)	Deployed Application	count	Total number of errors triggered by JSP page invocations.
glassfish.applications.{Deployed Application}.JspReloadedCount	Deployed Application	count	Total number of reloaded JSP pages.

glassfish.metrics-deployment-lifecycle metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.deployment.lifecycle.ActiveApplicationsDeployed (featured metric)		count	Total number of applications deployed so far.
glassfish.deployment.lifecycle.TotalApplicationsDeployed (featured metric)		count	Number of applications deployed.

glassfish.metric-server-requests metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.http-service.server.request.Count200		count	Number of responses with a status code equal to 200.
glassfish.http-service.server.request.Count2xx		count	Number of responses with a status code in the 2xx range.
glassfish.http-service.server.request.Count302		count	Number of responses with a status code equal to 302.
glassfish.http-service.server.request.Count304		count	Number of responses with a status code equal to 304.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glassfish.metric-server-requests metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.http-service.server.request.Count3xx		count	Number of responses with a status code equal in the 3xx range.
glassfish.http-service.server.request.Count400		count	Number of responses with a status code equal to 400.
glassfish.http-service.server.request.Count401		count	Number of responses with a status code equal to 401.
glassfish.http-service.server.request.Count403		count	Number of responses with a status code equal to 403.
glassfish.http-service.server.request.Count404		count	Number of responses with a status code equal to 404.
glassfish.http-service.server.request.Count4xx		count	Number of responses with a status code equal in the 4xx range.
glassfish.http-service.server.request.Count503		count	Number of responses with a status code equal to 503.
glassfish.http-service.server.request.Count5xx		count	Number of responses with a status code equal in the 5xx range.
glassfish.http-service.server.request.MaxTime		ms	Longest response time for a request; not a cumulative value, but the largest response time from among the response times.
glassfish.http-service.server.request.CountBytesReceived (featured metric)		count	Number of bytes received.
glassfish.http-service.server.request.CountBytesTransmitted (featured metric)		count	Number of bytes transmitted.
glassfish.http-service.server.request.RequestCount (featured metric)		count	Cumulative number of requests processed so far.
glassfish.http-service.server.request.CountOpenConnections		count	Number of open connections.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glassfish.metric-server-requests metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.http-service.server.request.ProcessingTime (featured metric)		ms	Average request processing time.

glassfish.metric-jvm-level metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.jvm.garbage-collectors.Scavenge.CollectionCount		count	Total number of collections that have occurred.
glassfish.jvm.garbage-collectors.Scavenge.CollectionTime		ms	Approximate accumulated collection elapsed time in milliseconds.
glassfish.jvm.garbage-collectors.MarkSweepCompact.CollectionCount		count	Total number of collections that have occurred.
glassfish.jvm.garbage-collectors.MarkSweepCompact.CollectionTime		ms	Approximate accumulated collection elapsed time in milliseconds.
glassfish.jvm.memory.CommittedHeapSize		bytes	Amount of memory in bytes that is committed for the Java virtual machine to use.
glassfish.jvm.memory.CommittedNonHeapSize		bytes	Amount of memory in bytes that is committed for the Java virtual machine to use.
glassfish.jvm.memory.InitHeapSize		bytes	Amount of memory in bytes that the Java virtual machine initially requests from the operating system for memory management.
glassfish.jvm.memory.InitNonHeapSize		bytes	Amount of memory in bytes that the Java virtual machine initially requests from the operating

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glassfish.metric-jvm-level metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			system for memory management.
glassfish.jvm.memory.MaxHeapSize (featured metric)		bytes	Maximum amount of memory in bytes that can be used for memory management.
glassfish.jvm.memory.MaxNonHeapSize		bytes	Maximum amount of memory in bytes that can be used for memory management.
glassfish.jvm.memory.ObjectPendingFinalizationCount		Count	Approximate number of objects for which finalization is pending.
glassfish.jvm.memory.UsedHeapSize		bytes	Amount of used memory in bytes.
glassfish.jvm.memory.UsedNonHeapSize		bytes	Amount of used memory in bytes.

glassfish.metric-transaction-service metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfish.transaction-service.ActiveCount (featured metric)		count	Provides the number of transactions that are currently active.
glassfish.transaction-service.CommittedCount (featured metric)		count	Provides the number of transactions that have been committed.
glassfish.transaction-service.RolledBackCount (featured metric)		count	Provides the number of transactions that have been rolled back.

glassfish.metric-server-resources metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfis.resources.{PoolName}.AverageConnWaitTime	PoolName	ms	Average wait-time-duration per successful connection request.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glassfish.metric-server-resources metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfis.resources.{PoolName}.ConnRequestWaitTime	PoolName	ms	The longest and shortest wait times of connection requests. The current value indicates the wait time of the last request that was serviced by the pool.
glassfis.resources.{PoolName}.NumConnAcquired	PoolName	count	Number of logical connections acquired from the pool.
glassfis.resources.{PoolName}.NumConnCreated	PoolName	count	The number of physical connections that were created since the last reset.
glassfis.resources.{PoolName}.NumConnDestroyed	PoolName	count	Number of physical connections that were destroyed since the last reset.
glassfis.resources.{PoolName}.NumConnfailedValidation	PoolName	count	The total number of connections in the connection pool that failed validation from the start time until the last sample time.
glassfis.resources.{PoolName}.NumConnFree	PoolName	count	The total number of free connections in the pool as of the last sampling.
glassfis.resources.{PoolName}.NumConnnotSuccessfullyMatched	PoolName	count	Number of connections rejected during matching.
glassfis.resources.{PoolName}.NumConnReleased	PoolName	count	Number of logical connections released to the pool.
glassfis.resources.{PoolName}.NumConnSuccessfullyMatched	PoolName	count	Number of connections successfully matched.
glassfis.resources.{PoolName}.NumConnTimedOut	PoolName	count	The total number of connections in the pool that timed out between the start time and the last sample time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

glassfish.metric-server-resources metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
glassfis.resources.{PoolName}.NumConnUsed	PoolName	count	Provides connection usage statistics. The total number of connections that are currently being used, as well as information about the maximum number of connections that were used (the high water mark).
glassfis.resources.{PoolName}.NumPotentialConnLeak (featured metric)	PoolName	count	Number of potential connection leaks.
glassfis.resources.{PoolName}.WaitQueueLength (featured metric)	PoolName	count	Number of connection requests in the queue waiting to be serviced.

HAProxy default checks and policies

Agent Client Collector provides the following default checks and policies for HAProxy monitoring.

HAProxy events policy

Type	Check	Description	Usage	Output
Event	check-haproxy	Retrieves the details of the HAProxy load balancer to determine whether it is running.	`commonchecks check-haproxy (options) -h host host -P port Port on which the HAProxy server is running -A all services all services -q Path on which HAProxy Stats URI is`	`CheckHAProxy OK: UP: 50% of 4 // services, DOWN: backend/nginx1[L4CON], backend/BACKEND`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HAProxy events policy (continued)

Type	Check	Description	Usage	Output
			configu red -use ssl ssl Usage example: check- haproxy.rb -h localhost -P 8080 -A true -q stats -use_ssl false	

HAProxy metrics policy

Type	Check	Description	Usage	Output
Metric	metrics-check-haproxy	Retrieves the HAProxy load balancer metrics to determine whether it is running.	commonche cks metrics-c heck-hapr oxy (options) -q host Path on which HAProxy Stats URI is configu red -P port Port on which the HAProxy server is running -c host host name Usage example: metrics- check- haproxy.rb -q stats	<hostname> .haproxy.f rontend.se ssion_curr ent 0 1652265 603 <hostname> .haproxy.f rontend.se ssion_to tal 0 1652265 603 <hostname> .haproxy.b ackend.ses sion_curr ent 0 1652265 603 <hostname> .haproxy.b ackend.ses sion_total 1 1652265 603

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HAProxy metrics policy (continued)

Type	Check	Description	Usage	Output
			`-P 8080 -c localhost`	

HAProxy metrics

The following table lists the metrics that are gathered as output from HAProxy checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

Front-end metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
haproxy.frontend.session_current (featured metric)		count	The current number of sessions being used to issue requests.
haproxy.frontend.session_total (featured metric)		count	Total number of sessions
haproxy.frontend.bytes_in (featured metric)		count	The total number of bytes recieved by the server in the last second from the client.
haproxy.frontend.bytes_out (featured metric)		count	The total number of bytes sent by the server in the last second.
haproxy.frontend.connection_errors (featured metric)		count	The number of requests that encountered an error trying to connect to a server.
haproxy.frontend.warning_retries		count	The number of times a connection to a server was retried.
haproxy.frontend.warning_redispatched		count	The number of times a request was redispatched to another server.
haproxy.frontend.response_1xx		count	HAProxy exposes the number of responses by HTTP status code - 1xx
haproxy.frontend.response_2xx		count	HAProxy exposes the number of responses

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Front-end metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			by HTTP status code - 2xx
haproxy.frontend.response_3xx		count	HAProxy exposes the number of responses by HTTP status code - 3xx
haproxy.frontend.response_4xx		count	HAProxy exposes the number of responses by HTTP status code - 4xx
haproxy.frontend.response_5xx		count	HAProxy exposes the number of responses by HTTP status code - 5xx
haproxy.frontend.response_other		count	HAProxy exposes the number of responses by non-HTTP application
haproxy.frontend.queue_time		count	The average queue time over the last 1024 requests.
haproxy.frontend.connect_time		count	The average connect time over the last 1024 requests.
haproxy.frontend.response_time (featured metric)		count	The response time to the client
haproxy.frontend.average_time (featured metric)		count	The average time taken for the sessions.

Back-end metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
haproxy.backend.session_current (featured metric)		count	The current number of sessions being used to issue requests.
haproxy.backend.session_total (featured metric)		count	Total number of sessions
haproxy.backend.bytes_in (featured metric)		count	The total number of bytes recieved by the server in the last second from the client.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Back-end metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
haproxy.backend.bytes_out (featured metric)		count	The total number of bytes sent by the server in the last second.
haproxy.backend.connection_errors (featured metric)		count	The number of requests that encountered an error trying to connect to a server.
haproxy.backend.warning_retries		count	The number of times a connection to a server was retried.
haproxy.backend.warning_redispatched		count	The number of times a request was redispatched to another server.
haproxy.backend.response_1xx		count	HAProxy exposes the number of responses by HTTP status code - 1xx
haproxy.backend.response_2xx		count	HAProxy exposes the number of responses by HTTP status code - 2xx
haproxy.backend.response_3xx		count	HAProxy exposes the number of responses by HTTP status code - 3xx
haproxy.backend.response_4xx		count	HAProxy exposes the number of responses by HTTP status code - 4xx
haproxy.backend.response_5xx		count	HAProxy exposes the number of responses by HTTP status code - 5xx
haproxy.backend.response_other		count	HAProxy exposes the number of responses by non-HTTP application
haproxy.backend.queue_time		count	The average queue time over the last 1024 requests.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Back-end metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
haproxy.backend.connect_time		count	The average connect time over the last 1024 requests.
haproxy.backend.response_time (featured metric)		count	The response time to the client
haproxy.backend.average_time (featured metric)		count	The average time taken for the sessions.
haproxy.backend.num_up		count	

HTTP default checks and policies

Agent Client Collector provides the following policies for HTTP health monitoring. Policies come with the checks specified in the tables below.

HTTP entry point events policy

Type	Check	Description	Usage and Usage Example	Output
Event	util.check-http-follow-redirect	Verifies that redirection links can be followed in a set number of requests.	Usage check-head-redirect.rb (options): • -A, --auth-first-only: Use basic auth on first request only. • -aws-access-key-id: AWS Access Key. Either set ENV["AWS_ACCESS_KEY_ID"] or provide it as an option on agent. • -r -aws-region: AWS Region (defaults to us-east-1). • --aws-secret-access-key: AWS Secret Access Key. Either set ENV["AWS_SECRET_ACCESS_KEY"] or provide it as an option on agent.	Check Head Redirect OK

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HTTP entry point events policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			• -R --redirect: Follow first <N> redirects • -g --get-redirects: Follow first <N> redirects with GET requests. • -s, --s3-config-bucket: S3 config bucket to get configuration. • -k, --s3-config-key: S3 config key to get configuration. • -u, --url: URL needs to be updated in `Monitoring HTTP Entrypoint/cmdb_ci_endpoint_http_list.do` for the CI. The Aws-region, aws-secret-key, aws-access-key ,s3-config, s3-config-key parameters are useful if you do not want to configure connection information in other check parameters. If a bucket and key have access to the environment in which the Sensu check executes in, provide an AWS key and token and the checks pull the specified	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HTTP entry point events policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			JSON file from S3 and merge the JSON config into the current check configuration. Usage example: `command: check-head-redirect.rb -R 10 -u 'https://servicenow.com'`	
Event	util.check-http-response	Verifies URL response time and raises a CRITICAL/WARNING event if elapsed time exceeds provided CRITICAL/WARNING thresholds. Otherwise, it raises an OK event.	Usage: check-head-redirect.rb (options) • -R, --redirect: Follow first <N> redirects. • -w --timeout_warning: Set the timeout threshold for warning in milliseconds. • -c --timeout_critical: Set the timeout threshold for critical in milliseconds. • -u, --url URL needs to be updated in `Monitoring HTTP Entrypoint/cmdb_ci_endpoint_http_list.do` for the CI. Usage example: `command: check-head-redirect.rb -R 10 -u 'https://servicenow.com'`	Check Head Redirect OK

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HTTP entry point events policy (continued)

Type	Check	Description	Usage and Usage Example	Output
			`-w 3000 -c 5000`	

HTTP entry point metrics policy

Type	Check	Description	Usage and Usage Example	Output
Metric	util.metrics-http-curl	Retrieves metrics on HTTP endpoints using curl. This check requires a proxy agent.	Usage: • -a, --curl_args "CURL ARGS": Additional arguments to pass to curl. • s, --scheme SCHEME: Metric naming scheme, text to append to metric (Default: hostname). • -u, --url: URL needs to be updated in `Monitoring HTTP Entrypoint/cmdb_ci_endpoint_http_list.do` for the CI. Usage example: `metrics-curl.rb -u myURL.com -a -Lk`	Check run successfully. Output: ws10.curl_timings.time_total 0.219622 1642749209 ws10.curl_timings.time_namelooku 0.145494 1642749209 ws10.curl_timings.time_connect 0.151103 1642749209 ws10.curl_timings.time_pretransfe 0.168569 1642749209 ws10.curl_timings.time_redirect 0.095899 1642749209 ws10.curl_timings.time_starttransfe 0.219351 1642749209 ws10.curl_timings.http_code 200 1642749209

HTTP entry point metrics

The following table lists the metrics that are gathered as output from HTTP entry point checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HTTP entry point metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
http_code (featured metric)			Status code of the request.
time_total (featured metric)		milliseconds	The total time the operation lasted.
time_namelookup (featured metric)		seconds	Time to resolve the name in the URL.
time_connect (featured metric)		seconds	Time to complete the connection to the remote host or proxy.
time_pretransfer (featured metric)		seconds	Time between start and beginning of file transfer. Includes all pre-transfer commands and negotiations.
time_starttransfer (featured metric)		seconds	Time between the start and the transfer of the first byte. Includes *time_pretransfer* and the time the server used to calculate the result.
time_redirect (featured metric)		seconds	Time for all redirection steps including: • name lookup • connect • pretransfer • transfer Shows the complete execution time for multiple redirections.

HTTP response code check

Agent Client Collector provides the following additional check for HTTP response code. This check is not associated with any policy.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

HTTP response code check

Type	Name	Description	Usage and usage example	Output

Event	util.check-http-response-code	Raises a critical event if the URL response code doesn't match the regex.
		
		Pass credentials (username and password) using secure parameters.
			`-u, --url URL needs to be updated in the Monitoring HTTP Entrypoint/cmdb_ci_endpoint_http_list.do for the CI`
			`-r, --response-code Regex string to match the response code. Otherwise, a critical event is raised. Default: 200`
			`-f, --follow_redirects Follows redirect URLs`
			
			Usage example:
			`command: check-response-code.rb -u 'https://servicenow.com' -r 2.. -f`
				`CheckResponseCode OK : URL is in OK state. HTTP Response code is 200, matches regex`

Internet Information Services (IIS) metrics

The following table lists the metrics that are gathered as output from IIS checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
iis._Total.Web_Service(_Total).Current_Connections (featured metric)		count	Total number of current connections to IIS.
iis._Total.Web_Service(_Total).Total_Get_Requests (featured metric)		count	Total of HTTP GET requests received by the IIS server.
iis._Total.Web_Service(_Total).Total_Post_Requests (featured metric)		count	Total of HTTP POST requests received by the IIS server.
iis._Total.Web_Service(_Total).Get_Requests/sec (featured metric)		count per second	Rate of all HTTP GET requests received by the IIS server.
iis._Total.Web_Service(_Total).Post_Requests/sec (featured metric)		count per second	Rate of all HTTP POST requests received by the IIS server.

Linux default checks and policies

Agent Client Collector provides the following default checks and policies for Linux Metrics monitoring.

Linux monitoring metrics checks

Linux metrics policy

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
os.linux.metrics.process-usage	proc.acc.running	process-name	Number of processes running with this name(acc)	count		
	proc.acc.cpuPercent	process-name	Percentage of cpu taken by the process.	percent		
	proc.acc.memPercent	process-name	Percentage of memory taken by the process.	percent		
os.linux.metrics.reboot-count-today	reboot.count.today	empty	Number of reboot done on today	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
os.linux.metrics-system-cpu	cpu.total.user	total	Normal processes executing in user mode; cpu.total.user is the total of the cpuN.user metrics.	count		
	cpu.total.nice	total	Niced processes executing in user mode; cpu.total.nice is the total of the cpuN.nice metrics.	count		
	cpu.total.system	total	Time the CPU spent running the kernel; cpu.total.system is the total of the cpuN.system metrics.	sec		
	cpu.total.idle	total	Total time the CPU spent in an idle state.; cpu.total.idle is the total of the cpuN.idle metrics.	sec		
	cpu.total.iowait	total	Total time the CPU spent waiting for IO operations to complete.;	sec		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			cpu.total.iowait is the total of the cpuN.iowait metrics.			
	cpu.total.irq	total	Total time that the processor is spending on handling Interrupts.; cpu.total.irq is the total of the cpuN.irq metrics.	sec		
	cpu.total.softirq	total	Time spent on servicing soft interrupt requests; cpu.total.softirq is the total of the cpuN.softirq metrics.	sec		
	cpu.total.steal	total	Total time the virtual CPU spent waiting for the hypervisor to service another virtual CPU. Only applies to virtual machines.	sec		
	cpu.total.guest	total	Total time the CPU spent running the virtual processor. Only	sec		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			applies to hypervisors.			
	cpu.total.guest_nice	total	Total time the CPU spent running as nice guest OS. cpu.total.guset_nice si the total of the cpuN.guest_nice metrics	sec		
	cpu.<cpu-core>.user	cpu-core	Time spent with normal processing in user mode.	sec		
	cpu.<cpu-core>.nice	cpu-core	Time spent with niced processes in user mode.	sec		
	cpu.<cpu-core>.system	cpu-core	Time spent running in kernel mode.	sec		
	cpu.<cpu-core>.idle	cpu-core	Time spent in vacations twiddling thumbs.	sec		
	cpu.<cpu-core>.iowait	cpu-core	Time spent waiting for I/O to completed. This is considered idle time too.	sec		
	cpu.<cpu-core>.irq	cpu-core	Time spent serving	sec		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			hardware interrupts.			
	cpu.<cpu-core>.softirq	cpu-core	Time spent serving software interrupts.	sec		
	cpu.<cpu-core>.steal	cpu-core	Time stolen by other operating systems running in a virtual environment.	sec		
	cpu.<cpu-core>.guest	cpu-core	Time spent for running a virtual CPU or guest OS under the control of the kernel.	sec		
	cpu.<cpu-core>.guest_nice	cpu-core	Total time the CPU spent running as nice guest OS.	sec		
	cpu.intr	empty	Interrupts serviced since boot time.	sec		
	cpu.ctxt	empty	Total number of context switches across all CPUs.	count		
	cpu.btime	empty	The time the system booted	sec		
	cpu.processes	empty	The number of processes	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			and threads created, which includes (but is not limited to) those created by calls to the fork() and clone() system calls.			
	cpu.procs_running	empty	The total number of processes running on all CPUs.	count		
	cpu.procs_blocked	empty	The number of processes currently blocked, waiting for I/O to complete.	count		
	cpu.cpu_count	empty	Number of cpu on the system	count		
	cpu.<cpu-core>.cores	cpu-core	The number of CPU cores	core count		
os.linux.metrics.system-cpu-load	load_avg.one	empty	The average system load over one minute.	thread count	yes	yes
	load_avg.five	empty	The average system load over five minutes.	thread count	yes	yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	load_avg.fifteen	empty	The average system load over fifteen minutes.	thread count	yes	yes
	load_avg.norm.one	empty	The average system load over one minute normalized by the number of CPUs.	thread count		
	load_avg.norm.five	empty	The average system load over five minutes normalized by the number of CPUs.	thread count		
	load_avg.norm.fifteen	empty	The average system load over fifteen minutes normalized by the number of CPUs.	thread count		
os.linux.metrics.system-cpu-percentage	cpu.avgutilization_percentage	empty	Percent of cpu was used on average	percent		
	cpu.user_percentage	empty	Percent of time total cpu was used by normal processes in user mode	percent	yes	yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	cpu.nice_percentage	empty	Percent of time all cpus used by niced processes in user mode	percent	yes	yes
	cpu.system_percentage	empty	The percent of time the CPU spent running the kernel.	percent	yes	yes
	cpu.idle_percentage	empty	Percent of time all cpus were idle	percent	yes	yes
	cpu.iowait_percentage	empty	Percent of time all cpus waiting for I/O to complete	percent	yes	yes
	cpu.irq_percentage	empty	Percent of time all cpus servicing interrupts	percent	yes	yes
	cpu.softirq_percentage	empty	Percent of time all cpus servicing software interrupts	percent	yes	yes
	cpu.steal_percentage	empty	Percent of time all cpus serviced virtual hosts operating systems	percent	yes	yes
	cpu.guest_percentage	empty	Percent of time all cpus serviced guest	percent	yes	yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			operating system			
os.linux.metrics.system-disk	disk.<disk-name>.reads	disk-name	Total number of reads completed successfully.	count	yes	yes
	disk.<disk-name>.readsMerged	disk-name	Total number of reads merged	count		
	disk.<disk-name>.sectorsRead	disk-name	Total number of sectors read successfully.	count		
	disk.<disk-name>.readTime	disk-name	Total number of milliseconds spent by all reads.	millisec		
	disk.<disk-name>.writes	disk-name	Total number of writes completed successfully.	count	yes	yes
	disk.<disk-name>.writesMerged	disk-name	Total number of writes merged	count		
	disk.<disk-name>.sectorsWritten	disk-name	Total number of sectors written successfully.	count		
	disk.<disk-name>.writeTime	disk-name	Total number of milliseconds spent by all writes.	misllisec		
	disk.<disk-name>.ioInProgress	disk-name	Total number of I/Os currently in progress	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	disk.<disk-name>.ioTime		Total time spent doing I/Os	millisec	yes	yes
	disk.<disk-name>.ioTimeWeighted	disk-name	Total time spent doing I/Os. This can provide an easy measure of both I/O completion time and the backlog that may be accumulating.	millisec		
os.linux.metrics.system-disk-capacity	disk.<file-system-name>.total	file-system-name	The total size of the file system.	byte		
	disk.<file-system-name>.used	file-system-name	The total amount of space allocated to existing files in the file system.	byte		
	disk.<file-system-name>.avail	file-system-name	The total amount of space available within the file system.	byte		
	disk.<file-system-name>.used_percentage	file-system-name	The percentage of the available space that currently allocated to all files	percent		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			on the file system.			
	disk.<file-system-name>.itotal	file-system-name	The total number of inodes on the file system.	count		
	disk.<file-system-name>.iused	file-system-name	The number of used inodes.	count		
	disk.<file-system-name>.iavail	file-system-name	The number of free (unused) inodes.	count		
	disk.<file-system-name>.iused_percentage	file-system-name	The percentage of used inodes.	percent		
os.linux.metrics-system-disk-usage	disk_usage.<disk>.total	disk-name	Total amount of space available on this disk	bytes		
	disk_usage.<disk>.used	disk-name	Total amount of space used in this disk	bytes		
	disk_usage.<disk>.avail	disk-name	Total amount of space available on this disk	bytes		
	disk_usage.<disk>.used_percentage	disk-name	The percentage of space used on this disk	percent	yes	yes
os.linux.metrics-system-memoryos.linux.metrics-system-	memory.total	empty	Total usable RAM.	KB		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
memory-percent	memory.free	empty	Total free RAM.	KB		
	memory.available	empty	An estimate of how much memory is available for starting new applications, without swapping.	KB		
	memory.buffers	empty	Temporary storage used for raw disk blocks.	KB		
	memory.cached	empty	In-memory cache for files read from disk (the page cache). Does not include mem_swapcached.	KB		
	memory.swapTotal	empty	Total amount of swap space available.	KB	yes	yes
	memory.swapFree	empty	Amount of swap space that is currently unused.		yes	yes
	memory.dirty	empty	Memory which is waiting to get written back to the disk.	KB		
	memory.swapUsed	empty	The amount	KB	yes	yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			of swap space in use.			
	memory.used	empty	The amount of RAM in use.	KB		
	memory.usedWOBuffersCaches	empty	The amount of memory in use.	KB		
	memory.freeWOBuffersCaches	empty	Value of MemAvailable from /proc/meminfo if present, but falls back to adding free + buffered + cached memory if not.	KB		
	memory.swapUsedPercentage	empty	Percent of swap space used.	percent		
	memory_percent.free	empty	Percent of free RAM	percent	yes	yes
	memory_percent.available	empty	Percent of Mem available	percent	yes	yes
	memory_percent.buffers	empty	Precent of Memory used for raw disk blocks	percent	yes	yes
	memory_percent.cached	empty	Percent of memory used for in-memory cache for files read from disk	percent	yes	yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	memory_percent.dirty	empty	Percent of memory waiting to get written back to the disk.	percent	yes	yes
	memory_percent.swapUsed	empty	Percent of swap space used.	percent	yes	yes
	memory_percent.usedWOBuffersCaches	empty	Percent of memory is being used	percent	yes	yes
	memory_percent.freeWOBuffersCaches	empty	Percent of Memory available	percent	yes	yes
os.linux.metrics-system-uptime	system.uptime(sec)	empty	The amount of time the system has been working and available.	sec		
os.linux.metrics-memory-vmstat	vmstat.nr_free_pages	empty	Pages that are currently unused by the system.	pages		
	vmstat.nr_alloc_batch	empty	pages allocated to other domains due to insufficient memory in each domain of each NUMA	pages		
	vmstat.nr_inactive_anon	empty	memory pages in each domain of each NUMA	pages		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			node that have not been accessed for a long time			
	vmstat.nr_active_anon	empty	Anonymous virtual memory pages that have been recently used	KB		
	vmstat.nr_inactive_file	empty	The memory page corresponding to the file that has not been accessed for a long time in each domain of each NUMA.	KB		
	vmstat.nr_active_file	empty	The memory page corresponding to the file that has been accesseed recently .	pages		
	vmstat.nr_unevictable	empty	The number of pages is in the unevictable (non-)LRU list	count		
	vmstat.nr_mlock	empty	Pages mapped into a VM_LOCKED VMA - are	pages		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			a class of unevictable pages.			
	vmstat.nr_anon_pages	empty	Memory mapped pages that is not part of a file.	pages		
	vmstat.nr_mapped	empty	The number of memory mapped pages.	count		
	vmstat.nr_file_pages	empty				
	vmstat.nr_dirty	empty	Pages waiting to be written to disk	pages		
	vmstat.nr_writeback	empty	Pages currently being written to disk	pages		
	vmstat.nr_slab_reclaimable	empty	Pages from the kernel slab memory usage that can be reclaimed	pages		
	vmstat.nr_slab_unreclaimable	empty	Pages from the kernel slab memory usage that cannot be reclaimed	pages		
	vmstat.nr_page_table_pages	empty	Pages allocated to page tables	pages		
	vmstat.nr_kernel_stack	empty	Amount of memory allocated	KB		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			to kernel stacks.			
	vmstat.nr_unstable	empty	The number of unstable pages in each domain of each NUMA node	count		
	vmstat.nr_bounce	empty				
	vmstat.nr_vmscan_write	empty	The number of dirty pages written back during a scan of LRU(s)	count		
	vmstat.nr_vmscan_immediate_reclaim	empty				
	vmstat.nr_writeback_temp	empty				
	vmstat.nr_isolated_anon	empty	The number of anonymous memory pages isolated in each domain of each NUMA node	count		
	vmstat.nr_isolated_file	empty	The number of pages of file storage pages isolated in each domain of each NUMA node	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	vmstat.nr_shmem	empty	The number of shared memory pages	count		
	vmstat.nr_dirtied	empty	The number of dirty pages in each domain of each NUMA node	count		
	vmstat.nr_written	empty				
	vmstat.numa_hit	empty	The number of pages that were successfully allocated to this node.	count		
	vmstat.numa_miss	empty	The number of pages that were allocated on this node because of low memory on the intended node.	count		
	vmstat.numa_foreign	empty	The number of pages initially intended for this node that were allocated to another node instead.	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	vmstat.numa_interleave	empty	The number of interleave policy pages successfully allocated to this node.	count		
	vmstat.numa_local	empty	The number of pages successfully allocated on this node, by a process on this node	count		
	vmstat.numa_other	empty	The number of pages allocated on this node, by a process on another node.	count		
	vmstat.workingset_refault	empty				
	vmstat.workingset_activate	empty				
	vmstat.workingset_nodereclaim	empty				
	vmstat.nr_anon_transparent_hugepages	empty				
	vmstat.nr_free_cma	empty	Free continuous memory allocator pages in each domain of each NUMA			

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	vmstat.nr_dirty_threshold	empty				
	vmstat.nr_dirty_background_threshold	empty				
	vmstat.pgpgin	empty	The number of pages brought in from disk	count		
	vmstat.pgpgout	empty	The number of pages written out to disk	count		
	vmstat.pswpin	empty	The number of pages brought in from swap space	count		
	vmstat.pswpout	empty	The number of pages swapped out into swap space	count		
	vmstat.pgalloc_dma	empty				
	vmstat.pgalloc_dma32	empty				
	vmstat.pgalloc_normal	empty				
	vmstat.pgalloc_movable	empty				
	vmstat.pgfree	empty	The number of pages are free since last boot	count		
	vmstat.pgactivat	empty	Number of page activations since last boot	count		
	vmstat.pgdeactivate	empty	Number of page deactivations	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			since last boot			
	vmstat.pgfault	empty	Minor faults since last boot	pages		
	vmstat.pgmajfault	empty	Major faults since last boot	pages		
	vmstat.pglazyfreed	empty				
	vmstat.pgrefill_dma	empty				
	vmstat.pgrefill_dma32	empty				
	vmstat.pgrefill_normal	empty	Number of page refills since last boot	count		
	vmstat.pgrefill_ movable	empty				
	vmstat.pgsteal_ kswapd_dma	empty				
	vmstat.pgsteal_ kswapd_dma32	empty				
	vmstat.pgsteal_ kswapd_normal	empty				
	vmstat.pgsteal_ kswapd_movable	empty				
	vmstat.pgsteal_ direct_dma	empty				
	vmstat.pgsteal_ direct_dma32	empty				
	vmstat.pgsteal_ direct_normal	empty				
	vmstat.pgsteal_ direct_movable	empty				
	vmstat.pgscan_ kswapd_dma	empty				
	vmstat.pgscan_ kswapd_dma32	empty				
	vmstat.pgscan_ kswapd_normal	empty	Number of pages	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			scanned by kswapd since boot			
	vmstat.pgscan_kswapd_movable	empty				
	vmstat.pgscan_direct_dma	empty				
	vmstat.pgscan_direct_dma32	empty				
	vmstat.pgscan_direct_normal	empty	Number of pages reclaimed since boot	count		
	vmstat.pgscan_direct_movable	empty				
	vmstat.pgscan_direct_throttle	empty				
	vmstat.zone_reclaim_failed	empty				
	vmstat.pginodesteal	empty				
	vmstat.slabs_scanned	empty				
	vmstat.kswapd_inodesteal	empty				
	vmstat.kswapd_low_wmark_hit_quickly	empty				
	vmstat.kswapd_high_wmark_hit_quickly	empty				
	vmstat.pageoutrun	empty	Number of times kswapd called page reclaim	count		
	vmstat.allocstall	empty	Number of times page reclaim was called directly (low memory)	count		
	vmstat.pgrotated	empty				

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	vmstat.drop_pagecache	empty				
	vmstat.drop_slab	empty				
	vmstat.numa_pte_updates	empty				
	vmstat.numa_huge_pte_updates	empty				
	vmstat.numa_hint_faults	empty				
	vmstat.numa_hint_faults_local	empty				
	vmstat.numa_pages_migrated	empty				
	vmstat.pgmigrate_success	empty				
	vmstat.pgmigrate_fail	empty				
	vmstat.compact_migrate_scanned	empty				
	vmstat.compact_free_scanned	empty				
	vmstat.compact_isolated	empty				
	vmstat.compact_stall	empty	The number of times a process stalls to run memory compaction so that a huge page is free for use.	count		
	vmstat.compact_fail	empty	The number of times the system tries to compact memory but failed.	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
	vmstat.compact_success	empty	The number of times the system compacted memory and freed a huge page for use.	count		
	vmstat.htlb_buddy_alloc_success	empty				
	vmstat.htlb_buddy_alloc_fail	empty				
	vmstat.unevictable_pgs_culled	empty				
	vmstat.unevictable_pgs_scanned	empty				
	vmstat.unevictable_pgs_rescued	empty				
	vmstat.unevictable_pgs_mlocked	empty				
	vmstat.unevictable_pgs_munlocked	empty				
	vmstat.unevictable_pgs_cleared	empty				
	vmstat.unevictable_pgs_stranded	empty				
	vmstat.thp_fault_alloc	empty	The number of huge pages is successfully allocated to handle a page fault.	count		
	vmstat.thp_fault_fallback	empty	The number of page fault fails to allocate a huge page and instead	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			falls back to using small pages.			
	vmstat.thp_collapse_alloc	empty	The number of collapse of a range of pages into one huge page and then successfull allocation of a new huge page to store the data.	count		
	vmstat.thp_collapse_alloc_failed	empty	The number of collapse of a range of pages into one huge page but failed the allocation.	count		
	vmstat.thp_split	empty	The number of split of a huge page into base pages	count		
	vmstat.thp_zero_page_alloc	empty	The number of successful allocation of huge zero page	count		
	vmstat.thp_zero_page_alloc_failed	empty	The number of times the kernel failed to allocate	count		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Check	Metric Name	Resource	Description	Units	Featured Metric	Anomaly Detection
			huge zero page and falls back to using small pages.			
	vmstat.balloon_inflate	empty				
	vmstat.balloon_ deflate	empty				
	vmstat.balloon_ migrate	empty				
os.linux.metrics-process-status	proc.<process>.VmSize	process-name	The total amount of virtual memory used by the process	KB		
	proc.<process>.VmRSS	process-name	The non-swapped physical memory a process has used	KB		
	proc.<process>.VmSwap	process-name	The total amount of swap space used.	KB		

Linux network monitoring checks

Note: When upgrading from an earlier version, manually add the checks in this table to the Linux metrics policy.

Linux metrics policy

Type	Check	Description	Usage and usage example	Metrics collected	Featured Metric
Metric	os.linux.metrics-network-interface	Retrieves all network interface related metrics for Linux servers.	Usage: • -x, excludeinterface: List of interfaces to exclude	• rxBytes (featured metric) • rxPackets (featured metric) • rxErrors	yes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Type	Check	Description	Usage and usage example	Metrics collected	Featured Metric
			(comma separated) • -i, includeinterface: List of interfaces to include (comma separated) • -I, includeinterfaceregex: Regex matching interfaces to include • -X, excludeinterfaceregex: Regex matching interfaces to exclude Usage example: `command: metrics-network-interface.rb`	• rxDrops (featured metric) • rxFifo • rxFrame • rxCompressed • rxMulticast • rxBytes (featured metric) • rxPackets (featured metric) • rxErrors • rxDrops (featured metric) • rxFifo • rxColls • rxCarrier • rxCompressed	
Metric	os.linux.metrics-netstat-tcp	Retrieves metrics on TCP socket states from netstat. Useful on high-traffic web or proxy servers with large numbers of short-lived TCP connections coming and going.	Usage: • -p, port: The port from which you want to receive metrics. Value range = 1-65535. • -t, type: The port type from which to receive metrics. Values=local or remote. Default = local	• tcp.UNKNOWN • tcp.ESTABLISHED • tcp.SYN_SENT • tcp.SYN_RECV • tcp.FIN_WAIT1 • tcp.FIN_WAIT2 • tcp.TIME_WAIT • tcp.CLOSE • tcp.CLOSE_WAIT • tcp.LAST_ACK • tcp.LISTEN • tcp.CLOSING	no

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Linux metrics policy (continued)

Type	Check	Description	Usage and usage example	Metrics collected	Featured Metric
			• -d, disabletcp6: Disables the tcp6 check. Enter a value to set disabletcp6 = true. Usage example: `metrics-netstat-tcp.rb`		

Linux metrics

The following table lists the metrics that are gathered as output from Linux checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

os.linux.metrics-process-usage metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
proc.acc.running			Number of processes running with this name (acc).
proc.acc.cpuPercent			Percentage of CPU taken by the process.
proc.acc.memPercent			Percentage of memory taken by the process.

os.linux.metrics-reboot-count-today metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
reboot.count.today			Number of reboots today.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-cpu metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cpu.total.user			Normal processes executing in user mode; *cpu.total.user* is the total of the *cpuN.user* metrics.
cpu.total.nice			Niced processes executing in user mode; *cpu.total.nice* is the total of the *cpuN.nice* metrics.
cpu.total.system			Time the CPU spent running the kernel; *cpu.total.system* is the total of the *cpuN.system* metrics.
cpu.total.idle			Total time the CPU spent in an idle state; *cpu.total.idle* is the total of the *cpuN.idle* metrics.
cpu.total.iowait			Total time the CPU spent waiting for IO operations to complete; *cpu.total.iowait* is the total of the *cpuN.iowait* metrics.
cpu.total.irq			Total time that the processor is spending on Interrupts; *cpu.total.irq* is the total of the *cpuN.irq* metrics.
cpu.total.softirq			Time spent servicing soft interrupt requests; *cpu.total.softirq* is the total of the *cpuN.softirq* metrics.
cpu.total.steal			Total time the virtual CPU spent waiting for the Hypervisor to

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-cpu metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			service another virtual CPU. Applies only to virtual machines.
cpu.total.guest			Total time the CPU spent running the virtual processor. Applies only to Hypervisors.
cpu.total.guest_nice			Total time the CPU spent running as nice guest OS; *cpu.total.guset_nice* is the total of the *cpuN.guest_nice* metrics
cpu.<cpu-core>.user			Time spent with normal processing in user mode.
cpu.<cpu-core>.nice			Time spent with niced processing in user mode.
cpu.<cpu-core>.system			Time spent running in kernel mode.
cpu.<cpu-core>.idle			Time spent idle.
cpu.<cpu-core>.iowait			Time spent waiting for I/O to complete. Also considered idle time.
cpu.<cpu-core>.irq			Time spent serving hardware interrupts.
cpu.<cpu-core>.softirq			Time spent serving software interrupts.
cpu.<cpu-core>.steal			Time stolen by other operating systems running in a virtual environment.
cpu.<cpu-core>.guest			Time spent running a virtual CPU or guest OS under the control of the kernel.
cpu.<cpu-core>.guest_nice			Total time the CPU spent running as nice guest OS.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-cpu metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cpu.intr			Interrupts serviced since boot time.
cpu.ctxt			Total number of context switches across all CPUs.
cpu.btime			Boot time.
cpu.processes			The number of processes and threads created, which includes (but is not limited to) those created by calls to the fork() and clone() system calls.
cpu.procs_running			The total number of processes running on all CPUs.
cpu.procs_blocked			The number of processes currently blocked and waiting for I/O to complete.
cpu.cpu_count			The number of CPUs on the system.
cpu.<cpu-core>.cores			The number of CPU cores.

os.linux.metrics-system-cpu-load metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
load_avg.one (featured metric)			The average system load over one minute.
load_avg.five (featured metric)			The average system load over five minutes.
load_avg.fifteen (featured metric)			The average system load over fifteen minutes.
load_avg.norm.one			The average system load over one minute normalized by the number of CPUs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-cpu-load metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
load_avg.norm.five			The average system load over five minutes normalized by the number of CPUs.
load_avg.norm.fifteen			The average system load over fifteen minutes normalized by the number of CPUs.

os.linux.metrics-system-cpu-percentage metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cpu.avgutilization_percentage			The average percentage of CPU used.
cpu.user_percentage (featured metric)			Percentage of time total CPU was used by normal processes in user mode.
cpu.nice_percentage (featured metric)			Percentage of time all CPUs used by niced processes in user mode.
cpu.system_percentage (featured metric)			Percentage of time the CPU spent running the kernel.
cpu.idle_percentage (featured metric)			Percentage of time all CPIUs were idle.
cpu.iowait_percentage (featured metric)			Percentage of time all CPUs waited for I/O to complete.
cpu.irq_percentage (featured metric)			Percentage of time all CPUs serviced interrupts.
cpu.softirq_percentage (featured metric)			Percentage of time all CPIs serviced software interrupts.
cpu.steal_percentage (featured metric)			Percentage of time all CPUs serviced virtual-host operating systems.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-cpu-percentage metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cpu.guest_percentage (featured metric)			Percentage of time all CPUs serviced guest operating systems.

os.linux.metrics-system-disk metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
disk.<disk-name>.reads (featured metric)			Total number of reads completed successfully.
disk.<disk-name>.readsMerged			Total number of reads merged.
disk.<disk-name>.sectorsRead			Total number of sectors read successfully.
disk.<disk-name>.readTime		milliseconds	Total number of milliseconds spent by all reads.
disk.<disk-name>.writes (featured metric)			Total number of writes completed successfully.
disk.<disk-name>.writesMerged			Total number of writes merged.
disk.<disk-name>.sectorsWritten			Total number of sectors written successfully.
disk.<disk-name>.writeTime		milliseconds	Total number of milliseconds spent by all writes.
disk.<disk-name>.ioInProgress			Total number of I/Os currently in progress.
disk.<disk-name>.ioTime (featured metric)			Total time spent on I/Os.
disk.<disk-name>.ioTimeWeighted			Total time spent on I/Os. This can provide a measure of both I/O completion time and the backlog that might be accumulating.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-disk-capacity metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
disk.<file-system-name>.total			The total size of the file system.
disk.<file-system-name>.used			The total amount of space allocated to existing files in the file system.
disk.<file-system-name>.avail			The total amount of space available within the file system.
disk.<file-system-name>.used_percentage			The percentage of the available space that is currently allocated to all files on the file system.
disk.<file-system-name>.itotal			The total number of *inodes* on the file system.
disk.<file-system-name>.iused			The number of used *inodes*.
disk.<file-system-name>.iavail			The number of free (unused) *inodes*.
disk.<file-system-name>.iused_percentage			The percentage of used *inodes*.

os.linux.metrics-system-disk-usage metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
disk_usage.<disk>.total			Total amount of space available on the disk.
disk_usage.<disk>.used			Total amount of space used in the disk.
disk_usage.<disk>.avail			Total amount of space available on the disk.
disk_usage.<disk>.used_percentage (featured metric)			The percentage of space used on the disk.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-memory metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
memory.total			Total usable RAM.
memory.free			Total free RAM.
memory.available			An estimate of how much memory is available for starting new applications without swapping.
memory.buffers			Temporary storage used for raw disk blocks.
memory.cached			In-memory cache for files read from disk (the page cache). Does not include *mem_swapcached*.
memory.swapTotal (featured metric)			Total amount of swap space available.
memory.swapFree (featured metric)			Amount of swap space that is currently unused.
memory.dirty			Memory which is waiting to be written back to the disk.
memory.swapUsed (featured metric)			The amount of swap space in use.
memory.used			The amount of RAM in use.
memory.usedWOBuffersCaches			The amount of memory in use.
memory.freeWOBuffersCaches			Value of *MemAvailable* from */proc/meminfo* if present, but falls back to free + buffered + cached memory if not.
memory.swapUsedPercentage			Percent of swap space used.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-system-memory-percent metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
memory_percent.free (featured metric)			Percentage of free RAM.
memory_percent.available (featured metric)			Percentage of memory available
memory_percent.buffers (featured metric)			Percentage of memory used for raw disk blocks.
memory_percent.cached (featured metric)			Percentage of memory used with in-memory cache for files read from disk.
memory_percent.dirty (featured metric)			Percentage of memory waiting to be written back to the disk.
memory_percent.swapUsed (featured metric)			Percentage of swap space used.
memory_percent.usedWOBuffersCaches (featured metric)			Percentage of memory used.
memory_percent.freeWOBuffersCaches (featured metric)			Percentage of memory available.

os.linux.metrics-system-uptime metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
system.uptime(sec)			The amount of time the system has been working and available.

os.linux.metrics-memory-vmstat metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmstat.nr_free_pages			Pages that are currently unused by the system.
vmstat.nr_alloc_batch			Pages allocated to other domains due to insufficient memory in each domain of each

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-memory-vmstat metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			non-uniform memory access (NUMA) node.
vmstat.nr_inactive_anon			Memory pages in each domain of each NUMA node that have not been accessed.
vmstat.nr_active_anon			Anonymous virtual memory pages that have been recently used.
vmstat.nr_inactive_file			The memory page corresponding to the file that has not been accessed in each domain of each NUMA.
vmstat.nr_active_file			The memory page corresponding to the file that has been recently accessed.
vmstat.nr_unevictable			The number of pages in the unevictable (non-)LRU list.
vmstat.nr_mlock			Pages mapped into a VM_LOCKED VMA that are a class of unevictable pages.
vmstat.nr_anon_pages			Memory mapped pages that are not part of a file.
vmstat.nr_mapped			The number of memory mapped pages.
vmstat.nr_file_pages			
vmstat.nr_dirty			Pages waiting to be written to disk.
vmstat.nr_writeback			Pages currently being written to disk.
vmstat.nr_slab_reclaimable			Pages from the kernel slab memory usage that can be reclaimed.
vmstat.nr_slab_unreclaimable			Pages from the kernel slab memory

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-memory-vmstat metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			usage that cannot be reclaimed.
vmstat.nr_page_table_pages			Pages allocated to page tables.
vmstat.nr_kernel_stack			Amount of memory allocated to kernel stacks.
vmstat.nr_unstable			The number of unstable pages in each domain of each NUMA node.
vmstat.nr_bounce			
vmstat.nr_vmscan_write			The number of dirty pages written back during a scan of LRUs.
vmstat.nr_vmscan_immediate_reclaim			
vmstat.nr_writeback_temp			
vmstat.nr_isolated_anon			The number of anonymous memory pages isolated in each domain of each NUMA node.
vmstat.nr_isolated_file			The number of file storage pages isolated in each domain of each NUMA node.
vmstat.nr_shmem			The number of shared memory pages.
vmstat.nr_dirtied			The number of dirty pages in each domain of each NUMA node.
vmstat.nr_written			
vmstat.numa_hit			The number of pages that were successfully allocated to this node.
vmstat.numa_miss			The number of pages that were allocated to this node because of low memory on the intended node.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-memory-vmstat metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmstat.numa_foreign			The number of pages initially intended for this node that were allocated to another node.
vmstat.numa_interleave			The number of interleave policy pages successfully allocated to this node.
vmstat.numa_local			The number of pages successfully allocated on this node by a process on this node.
vmstat.numa_other			The number of pages allocated on this node by a process on another node.
vmstat.workingset_refault			
vmstat.workingset_activate			
vmstat.workingset_nodereclaim			
vmstat.nr_anon_transparent_hugepages			
vmstat.nr_free_cma			Free continuous memory allocator pages in each domain of each NUMA.
vmstat.nr_dirty_threshold			
vmstat.nr_dirty_background_threshold			
vmstat.pgpgin			The number of pages brought in from disk.
vmstat.pgpgout			The number of pages written out to disk.
vmstat.pswpin			The number of pages brought in from swap space.
vmstat.pswpout			The number of pages swapped out into swap space.
vmstat.pgalloc_dma			
vmstat.pgalloc_dma32			
vmstat.pgalloc_normal			

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-memory-vmstat metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmstat.pgalloc_movable			
vmstat.pgfree			The number of pages free since the last boot.
vmstat.pgactivate			Number of page activations since the last boot.
vmstat.pgdeactivate			Number of page deactivations since the last boot.
vmstat.pgfault			Minor faults since the last boot.
vmstat.pgmajfault			Major faults since the last boot.
vmstat.pglazyfreed			
vmstat.pgrefill_dma			
vmstat.pgrefill_dma32			
vmstat.pgrefill_normal			Number of page refills since the last boot.
vmstat.pgrefill_movable			
vmstat.pgsteal_kswapd_dma			
vmstat.pgsteal_kswapd_dma32			
vmstat.pgsteal_kswapd_normal			
vmstat.pgsteal_kswapd_movable			
vmstat.pgsteal_direct_dma			
vmstat.pgsteal_direct_dma32			
vmstat.pgsteal_direct_normal			
vmstat.pgsteal_direct_movable			
vmstat.pgscan_kswapd_dma			
vmstat.pgscan_kswapd_dma32			
vmstat.pgscan_kswapd_normal			Number of pages scanned by *kswapd* since boot.
vmstat.pgscan_kswapd_movable			
vmstat.pgscan_direct_dma			
vmstat.pgscan_direct_dma32			

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-memory-vmstat metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmstat.pgscan_direct_normal			Number of pages reclaimed since boot.
vmstat.pgscan_direct_movable			
vmstat.pgscan_direct_throttle			
vmstat.zone_reclaim_failed			
vmstat.pginodesteal			
vmstat.slabs_scanned			
vmstat.kswapd_inodesteal			
vmstat.kswapd_low_wmark_hit_quickly			
vmstat.kswapd_high_wmark_hit_quickly			
vmstat.pageoutrun			Number of times *kswapd* called page reclaim.
vmstat.allocstall			Number of times page reclaim was called directly (low memory).
vmstat.pgrotated			
vmstat.drop_pagecache			
vmstat.drop_slab			
vmstat.numa_pte_updates			
vmstat.numa_huge_pte_updates			
vmstat.numa_hint_faults			
vmstat.numa_hint_faults_local			
vmstat.numa_pages_migrated			
vmstat.pgmigrate_success			
vmstat.pgmigrate_fail			
vmstat.compact_migrate_scanned			
vmstat.compact_free_scanned			
vmstat.compact_isolated			
vmstat.compact_stall			The number of times a process stalls when running memory compaction so that a huge page is free for use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-memory-vmstat metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmstat.compact_fail			The number of times the system attempted to compact memory but failed.
vmstat.compact_success			The number of times the system compacted memory and freed a huge page for use.
vmstat.htlb_buddy_alloc_success			
vmstat.htlb_buddy_alloc_fail			
vmstat.unevictable_pgs_culled			
vmstat.unevictable_pgs_scanned			
vmstat.unevictable_pgs_rescued			
vmstat.unevictable_pgs_mlocked			
vmstat.unevictable_pgs_munlocked			
vmstat.unevictable_pgs_cleared			
vmstat.unevictable_pgs_stranded			
vmstat.thp_fault_alloc			The number of huge pages successfully allocated to handle a page fault.
vmstat.thp_fault_fallback			The number of page fault fails to allocate a huge page before falling back to using small pages.
vmstat.thp_collapse_alloc			The number of pages collapsed into one huge page with the successful allocation of a new huge page to store the data.
vmstat.thp_collapse_alloc_failed			The number of pages collapsed into one huge page but failed allocation.
vmstat.thp_split			The number of base pages to split from a huge page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

os.linux.metrics-memory-vmstat metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmstat.thp_zero_page_alloc			The number of successful allocations of huge zero pages.
vmstat.thp_zero_page_alloc_failed			The number of times the kernel failed to allocate a huge zero page and falls back to using small pages.
vmstat.balloon_inflate			
vmstat.balloon_deflate			
vmstat.balloon_migrate			

os.linux.metrics-process-status metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
proc.<process>.VmSize			The total amount of virtual memory used by the process.
proc.<process>.VmRSS			The non-swapped physical memory used by a process.
proc.<process>.VmSwap			The total amount of swap space used.

MongoDB default checks and policies

Agent Client Collector provides the following policies for MongoDB health monitoring. Policies come with the checks specified in the indicated table. Policies and checks are available for both Windows and Linux.

MongoDB Events

Type	Check	Description	Usage	Command	Output
Event	app.mongodb.check-mongodb-alive	Monitors whether the MongoDB server is alive and creating alerts for the MongoDB server health status.	`commoncheck check-mongodb-alive -H db_host -p db_port` Where:	`commoncheck check-mongodb -p {{.labels.port}} -H {{.labels.host}}`	`Check mongodb Alive OK: Indicates that the server is in good health`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MongoDB Events (continued)

Type	Check	Description	Usage	Command	Output
			• `db_host` = Host where server is running. • `db_port` = Port to connect to the MongoDB server. Usage example: `command: commonchecks check-mongodb-alive -H 10.***.***.*** -p 27017`		
Event	app.mongodb.check-mongodb-metrics	Creates alerts for any of the MongoDB metrics, based on the threshold limit. To trigger the alerts for any MongoDB server metric, pass the whole metric name in the parameter.	`commonchecks check-mongodb-metrics -H Hostname -p port -w warning -c critical -d database -m MetricName` Where: • `Hostname` = Host where server is running. • `port` = Port to connect to the MongoDB server. • `warning` = Warning	`commonchecks check-mongodb-metrics -c {{.labels.p -d {{.labels.params_database}} -w {{.labels.params_warning}} -H {{.labels.params_host}} -p {{.labels.params_port}} -m {{.labels.params_metric}}`	`Check mongodb Metrics OK: Indicates that the mongodb.connection.cu value is within the acceptable threshold`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MongoDB Events (continued)

Type	Check	Description	Usage	Command	Output
			threshold value. • `critical` = Critical threshold value. • `database` = Database name. • `MetricName` = Specific metric to monitor. Usage example: `command: commonchecks check-mongodb-alive -H 10.***.***.*** -p 27017`		

MongoDB Metrics

Type	Check	Description	Usage	Command	Output

Metric	app.mongodb.metrics-mongodb	Retrieves metrics of the MongoDB server and all databases.
		
		Sample metrics:
		• `mongodb.connections` : Incoming connections
		• `mongodb.locks.Database` : Intent Exclusive lock mode
		• `mongodb.locks.Database.acquireCount_X` : Exclusive lock mode
			`commoncheck check-mongodb`
			
			Where:
			• `host` = Hostname where the MongoDB server is running.
			• `port` = Port to connect to the MongoDB server.
			• `database` = Database name.
				`commoncheck metrics-mongodb -p {{.labels.params_port}} -H {{.labels.params_host}} -d {{.labels.params_database}}`	`commoncheck<hostname>.mongodb.co`
					`20632`
					`1639498004`
					
					`mongodb.metrics.totalCreated:`
					
					`mongodb.locks.Database.acquireCount_IX`
					
					`mongodb.locks.Database.acquireCount_X`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MongoDB Metrics (continued)

Type	Check	Description	Usage	Command	Output
		• `mongodb.locks.Database.acquireCount_IS` : Intent shared lock mode			

MongoDB metrics

The following table lists the metrics that are gathered as output from MongoDB checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

metrics-mongodb metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mongodb.asserts.msg (featured metric)		count	Number of message assertions raised since the MongoDB process started. Examine the log file for more information about these messages.
mongodb.asserts.regular (featured metric)		count	Number of regular assertions raised since the MongoDB process started. Examine the log file for more information about these messages.
mongodb.asserts.rollovers		count	Number of times that the rollover counters have rolled over since the last time the MongoDB process started.
mongodb.asserts.tripwire			Number of tripwire assertions raised since the MongoDB process started.
mongodb.asserts.user		count	Number of "user asserts" that occurred since the last time the MongoDB process started.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mongodb.asserts.warning (featured metric)		count	Number of warnings raised since the MongoDB process started.
mongodb.connections.available		count	Number of unused incoming connections available.
mongodb.connections.current		count	Number of incoming connections from clients to the database server.
mongodb.connections.totalCreated		count	Count of all incoming connections created to the server.
mongodb.cursor.open.noTimeout		count	Number of open cursors with the option *DBQuery.Option.noTimeout* set to prevent timeout after a period of inactivity.
mongodb.cursor.open.pinned		count	Number of "pinned" open cursors.
mongodb.cursor.open.total (featured metric)		count	Number of cursors that MongoDB is maintaining for clients. Because MongoDB exhausts unused cursors, typically this value is small or zero. However, if there is a queue, stale tailable cursors, or a large number of operations, this value may increase.
mongodb.cursor.timedOut		count	Total number of cursors that have timed out since the server process started.
mongodb.databaseSizes.avgObjSize	admin	bytes	Average size of each document in bytes.
mongodb.databaseSizes.avgObjSize	config	bytes	Average size of each document in bytes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mongodb.databaseSizes.avgObjSize	local	bytes	Average size of each document in bytes.
mongodb.databaseSizes.collections	admin	count	Number of collections in the database.
mongodb.databaseSizes.collections	config	count	Number of collections in the database.
mongodb.databaseSizes.collections	local	count	Number of collections in the database.
mongodb.databaseSizes.dataSize	admin	bytes	Total size of the uncompressed data held in the database. The *dataSize* decreases when you remove documents. For databases using the *WiredTiger* storage engine, *dataSize* may be larger than *storageSize* if compression is enabled. The *dataSize* decreases when documents shrink.
mongodb.databaseSizes.dataSize	config	bytes	Total size of the uncompressed data held in the database.
mongodb.databaseSizes.dataSize	local	bytes	Total size of the uncompressed data held in the database.
mongodb.databaseSizes.indexes	admin	count	Total number of indexes across all collections in the database.
mongodb.databaseSizes.indexes	config	count	Total number of indexes across all collections in the database.
mongodb.databaseSizes.indexes	local	count	Total number of indexes across all collections in the database.
mongodb.databaseSizes.indexSize	admin	bytes	Sum of the space allocated to all

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			indexes in the database, including free index space.
mongodb.databaseSizes.indexSize	config	bytes	Sum of the space allocated to all indexes in the database, including free index space.
mongodb.databaseSizes.indexSize	local	bytes	Sum of the space allocated to all indexes in the database, including free index space.
mongodb.databaseSizes.objects	admin	count	Number of objects (specifically, documents) in the database across all collections.
mongodb.databaseSizes.objects	config	count	Number of objects (specifically, documents) in the database across all collections.
mongodb.databaseSizes.objects	local	count	Number of objects (specifically, documents) in the database across all collections.
mongodb.databaseSizes.storageSize	admin	bytes	This value does not decrease as you remove or shrink documents. This value may be smaller than *`dataSize`* for databases using the *`WiredTiger`* storage engine with compression enabled. *`storageSize`* does not include space allocated to indexes. See *`indexSize`* for the total index size.
mongodb.databaseSizes.storageSize	config	bytes	Sum of the space allocated to all collections in the database for

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			document storage, including free space.
mongodb.databaseSizes.total	totalStorageSize	bytes	Sum of the space allocated to all collections in the database for document storage, including free space.
mongodb.globalLock.activeClients.readers (featured metric)		count	Number of active client connections performing read operations.
mongodb.globalLock.activeClients.total		count	Total number of internal client connections to the database, including system threads as well as queued readers and writers. This metric will be higher than the total of *activeClients.readers* and *activeClients.writers* due to the inclusion of system threads.
mongodb.globalLock.activeClients.writers (featured metric)		count	Number of active client connections performing write operations.
mongodb.globalLock.currentQueue.readers (featured metric)		count	Number of operations that are currently queued and waiting for the read lock. A consistently small read queue, particularly of shorter operations, is not a cause for concern.
mongodb.globalLock.currentQueue.total		count	Total number of operations queued waiting for the lock (i.e., the sum of *globalLock.currentQueue.rea* and *globalLock.currentQueue.wri* A consistently small

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			queue, particularly of shorter operations, is not a cause for concern. The *globalLock.activeClients* readers and writers information provides context for this data.
mongodb.globalLock.currentQueue.writers (featured metric)		count	Number of operations that are currently queued and waiting for the write lock. A consistently small write queue, particularly of shorter operations, is not a cause for concern.
mongodb.globalLock.totalTime		microseconds	Time, in microseconds, since the database last started and created the *globalLock*. This is approximately equivalent to the total server uptime.
mongodb.locks.Collection.acquireCount_r		count	Number of times the collection lock was acquired in the Intent Shared (IS) lock mode.
mongodb.locks.Collection.acquireCount_w		count	Number of times the collection lock was acquired in the Intent Exclusive (IX) lock mode.
mongodb.locks.Collection.acquireCount_W		count	Number of times the collection lock was acquired in the Exclusive (X) lock mode.
mongodb.locks.Database.acquireCount_r		count	Number of times the database lock was acquired in the Intent Shared (IS) lock mode.
mongodb.locks.Database.acquireCount_w		count	Number of times the database lock was acquired in the Intent

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			Exclusive (IX) lock mode.
mongodb.locks.Database.acquireCount_W		count	Number of times the database lock was acquired in the Exclusive (X) lock mode.
mongodb.locks.Global.acquireCount_r		count	Number of times the global lock was acquired in the Intent Shared (IS) lock mode.
mongodb.locks.Global.acquireCount_w		count	Number of times the global lock was acquired in the Intent Exclusive (IX) lock mode.
mongodb.locks.Global.acquireCount_W		count	Number of times the global lock was acquired in the Exclusive (X) lock mode.
mongodb.mem.pageFaults		count	Total number of page faults. The `extra_info.page_faults` counter may increase dramatically during moments of poor performance and may correlate with limited memory environments and larger data sets. Limited and sporadic page faults do not necessarily indicate an issue.
mongodb.mem.resident (featured metric)		mebibyte	This value is roughly equivalent to the amount of RAM, in mebibyte (MiB), currently used by the database process. During normal use, this value tends to grow. In dedicated database servers, the number tends to approach the total

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			amount of system memory.
mongodb.mem.virtual		mebibyte	The quantity, in mebibyte (MiB), of virtual memory used by the mongod process.
mongodb.metrics.document.deleted		count	Total number of documents deleted.
mongodb.metrics.document.inserted		count	Total number of documents inserted.
mongodb.metrics.document.returned		count	Total number of documents returned by queries.
mongodb.metrics.document.updated		count	Total number of documents updated.
mongodb.metrics.getLastError.wtime_num (featured metric)		count	Total number of *`getLastError`* operations with a specified write concern (w) that wait for replica set members to acknowledge the write operation (w value greater than 1).
mongodb.metrics.getLastError.wtime_totalMillis (featured metric)		miliseconds	Total amount of time in milliseconds that the mongod has spent performing *`getLastError`* operations with write concern (w). These operations wait for one or more members of a replica set to acknowledge the write operation (a w value greater than 1).
mongodb.metrics.getLastError.wtimeouts (featured metric)		count	Number of times that write concern operations have timed out as a result of the *`wtimeout`* threshold to *`getLastError`*. This number

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			increments for both default and non-default write concern specifications.
mongodb.metrics.operation.scanAndOrder		count	Total number of queries that return sorted numbers that cannot perform the sort operation using an index.
mongodb.metrics.queryExecutor.scanned		count	Total number of index items scanned during queries and query-plan evaluation. This counter is the same as *`totalKeysExamined`* in the output of explain().
mongodb.metrics.queryExecutor.scannedObjects		count	Total number of documents scanned during queries and query-plan evaluation. This counter is the same as *`totalDocsExamined`* in the output of explain().
mongodb.metrics.record.moves			A document that reports on data related to record allocation in the on-disk memory files.
mongodb.metrics.repl.apply.batches_num		count	Total number of batches applied across all databases.
mongodb.metrics.repl.apply.batches_totalMillis		count	Total amount of time in milliseconds the mongod has spent applying operations from the oplog.
mongodb.metrics.repl.apply.ops		count	Total number of oplog operations applied. *`metrics.repl.apply.ops`* is incremented after each operation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mongodb.metrics.repl.buffer.count		count	Current number of operations in the oplog buffer.
mongodb.metrics.repl.buffer.maxSizeBytes		bytes	Maximum size of the buffer. This value is a constant setting in the mongod, and is not configurable.
mongodb.metrics.repl.buffer.sizeBytes		bytes	Current size of the contents of the oplog buffer.
mongodb.metrics.repl.network.bytes		count	Total amount of data read from the replication sync source.
mongodb.metrics.repl.network.getmores_num		count	Reports the total number of *`getmore`* operations, which are operations that request an additional set of operations from the replication sync source.
mongodb.metrics.repl.network.getmores_totalMillis		count	Total amount of time required to collect data from *`getmore`* operations.
mongodb.metrics.repl.network.ops		count	Total number of operations read from the replication source.
mongodb.metrics.repl.network.readersCreated		count	Total number of oplog query processes created. MongoDB creates a new oplog query any time an error occurs in the connection, including a timeout, or a network operation. Furthermore, *`metrics.repl.network.reade`* increments every time MongoDB selects a new source for replication.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mongodb.metrics.ttl.deletedDocuments		count	Total number of documents deleted from collections with a `ttl` index.
mongodb.metrics.ttl.passes		count	Total number of documents deleted from collections with a `ttl` index.
mongodb.network.bytesIn		count	Total number of bytes that the server has received over network connections initiated by clients.
mongodb.network.bytesOut		count	Total number of bytes that the server has sent over network connections initiated by clients.
mongodb.network.numRequests		count	Total number of distinct requests that the server has received. Use this value to provide context for the `network.bytesIn` and `network.bytesOut` values to ensure that MongoDB's network utilization is consistent with expectations and application use.
mongodb.opcounters.command		count	Total number of commands issued to the database since the mongod instance last started. `opcounters.command` counts all commands except the write commands: insert, update, and delete.
mongodb.opcounters.delete		count	Total number of delete operations since the mongod instance last started.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mongodb.opcounters.getmore		count	Total number of *getMore* operations since the mongod instance last started. This counter can be high even if the query count is low. Secondary nodes send *getMore* operations as part of the replication process.
mongodb.opcounters.insert		count	Total number of insert operations received since the mongod instance last started.
mongodb.opcounters.query		count	Total number of queries received since the mongod instance last started.
mongodb.opcounters.update		count	Total number of update operations received since the mongod instance last started.
mongodb.opcountersRepl.command		count	Total number of replicated commands issued to the database since the mongod instance last started.
mongodb.opcountersRepl.delete		count	Total number of replicated delete operations since the mongod instance last started.
mongodb.opcountersRepl.getmore		count	Total number of getMore operations since the mongod instance last started.
mongodb.opcountersRepl.insert (featured metric)		count	Total number of replicated insert operations since the mongod instance last started.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-mongodb metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mongodb.opcountersRepl.query (featured metric)		count	Total number of replicated queries since the mongod instance last started.
mongodb.opcountersRepl.update		count	Total number of replicated update operations since the mongod instance last started.

MS Exchange default checks and policies

Agent Client Collector provides the following default checks and policies for MS Exchange monitoring.

Type	Check	Description	Usage and usage example	Output
Metric	app.msexchange.metrics - msexchange	Retrieves Microsoft Exchange metrics.	`winchecks metric-msexchange (options) -s scheme Replaces the host name with this parameter` Usage example: `winchecks metric-msexchange {{if .labels.params_scheme}} -s {{if .labels.params_scheme}} {{end}}`	Metrics from MS Exchange server

MSSQL default checks and policies

The Agent Client Collector provides the following default checks and policies for MSSQL Metrics monitoring.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MSSQL Metrics

Type	Check	Description	Usage and Usage Example	Output
Metric	app.metrics-mssql-query	Generic MSSQL query class which accepts a query and dumps metrics.	winchecks metrics-mssql-query (options) -d, --database DATABASE Database schema to connect to (required) -q ,--query The query to execute, should return a list metric name and value, for example SELECT METRIC_NAME, VALUE FROM GV $SYSMETRIC -h, --host HOST Host name to connect to, defaults to localhost -r, --port PORT Port to connect to, defaults to 1433 -s, --scheme SCHEME Metric naming scheme, text to prepend to .$parent. $child, **Usage Example** `.winchecks metrics-mssql-query -d master -H localhost -p 1437 -q "select name,size from sys.database_files"`	Custom metrics defined in the parameters, which also contain value and timestamp.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL default checks and policies

The Agent Client Collector provides the following default checks and policies for MySQL Metrics monitoring.

MySQL Events

Type	Check	Description	Usage and Usage example	Output
Event	app.mysql.check-mysql-alive	Verifies the length of a result set from a MySQL query.	check-mysql-query-result-count.rb (options) -c, --critical COUNT COUNT critical threshold for number of items returned by the query (required) -d, --database DATABASE MySQL database (required) -h, --host HOST MySQL Host to connect to (required) -i, --ini VALUE My.cnf ini file --ini-section VALUE Section in my.cnf ini file. To enable connection to MySQL thru the .ini file, provide the values against the properties 'user' and 'password' in the **client** section in the .ini file. -P, --port PORT MySQL Port to connect to -q, --query QUERY Query to execute (required)	MysqlQueryCountCheck OK/CRITICAL/WARNING: message regarding ratio between query length and threshold values

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Events (continued)

Type	Check	Description	Usage and Usage example	Output
			-w, --warning COUNT Count warning threshold for number of items returned by the query (required) -S, --socket UNIX socket to connect to (required if host specified is 'localhost' on UNIX - like systems) **Usage Example** `check-mysql-query-result-count.rb -h 127.0.0.1 -P 3306 -d mysql -q "select * from user" -w 5 -c 8`	
Event	app.mysql.check-mysql-threads	Verifies the MySQL DB number of running threads and assigns a status of OK/WARNING/ CRITICAL depending on input values.	check-mysql-threads.rb (options) -h, --hostname HOST Hostname to login to -i, --ini VALUE My.cnf ini file --ini-section VALUE Section in my.cnf ini file (Needed if .ini path provided). To enable connection to MySQL thru the .ini file, provide the values against the properties	CheckMySQLHealth OK/Critical/ Warning and number of running threads

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Events (continued)

Type	Check	Description	Usage and Usage example	Output
			'user' and 'password' in the **client** section in the .ini file. -c, --critnum NUMBER Number of running threads upon which an alert is issued -w, --warnnum NUMBER Number of running threads upon which a warning is issued -P, --port PORT MySQL Port to connect to -S, --socket UNIX socket to connect to (required if host specified is 'localhost' on UNIX - like systems) **Usage Example** `check-mysql-threads.rb -h 127.0.0.1 -P 3306 -l 0 -m 1 -c 25 -w 20`	
Event	util.check-mysql-query	Verifies whether MySQL DB is running.	check-mysql-threads.rb (options) -h, --hostname HOST Hostname to login to -i, --ini VALUE My.cnf ini file --ini-section VALUE Section	CheckMySQLHealth OK/Critical/ Warning and number of running threads

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Events (continued)

Type	Check	Description	Usage and Usage example	Output
			in my.cnf ini file (Needed if .ini path is provided). To enable connection to MySQL thru the .ini file, provide the values against the properties 'user' and 'password' in the **client** section in the .ini file. -c, --critnum NUMBER Number of running threads upon which an alert is issued -w, --warnnum NUMBER Number of running threads upon which a warning is issued -l, --critlow NUMBER Number of running threads under which an alert is issued -m, --warnlow NUMBER Number of running threads under which a warning is issued -P, --port PORT MySQL Port to connect to -s, --socket UNIX socket to connect to (required if host specified is 'localhost'	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Events (continued)

Type	Check	Description	Usage and Usage example	Output
			on UNIX - like systems) **Usage Example** `check-mysql-threads.rb -h 127.0.0.1 -P 3306 -l 0 -m 1 -c 25 -w 20`	

MySQL Metrics

Type	Check	Description	Usage and Usage Example	Output
Metric	app.mysql.metrics-mysql	Returns metrics on the MySQL DB.	/usr/local/bin/metrics-mysql-graphite.rb (options) -h, --host HOST MySQL host to connect to (required) -i, --ini VALUE My.cnf ini file --ini-section VALUE Section in my.cnf ini file (Needed if .ini path is provided). To enable connection to MySQL thru the .ini file, provide the values against the properties 'user' and 'password' in the **client** section in the .ini file. -P, --port PORT MySQL port to connect to. -s, --scheme SCHEME Metric	MysqlQueryCountCheck OK/CRITICAL/ WARNING: message regarding ratio between query length and threshold values

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Metrics (continued)

Type	Check	Description	Usage and Usage Example	Output
			naming scheme, text to append to metric -S, --socket UNIX socket to connect to (required if host specified is 'localhost' on UNIX - like systems) -v, --verbose Show errors (if generated) regarding secondary status copies. Add -v to the Command to activate. **Usage Example** `check-mysql-threads.rb -h 127.0.0.1 -P 3306 -l 0 -m 1 -c 25 -w 20`	
Metric	app.mysql.check-mysql-threads	Verifies the MySQL DB number of running threads and assigns a status of OK/WARNING/ CRITICAL depending on input values.	/usr/local/bin/metrics-mysql-graphite.rb (options) -h, --hostname HOST Hostname to connect to (required) -i, --ini VALUE My.cnf ini file --ini-section VALUE Section in my.cnf ini file (Needed if .ini path is provided). To enable connection to MySQL thru the .ini file,	hostname.mysql.general.abortedC 7600 1649630975 hostname.mysql.general.abortedC 247 1649630975 hostname.mysql.general.txBytes 752733902 1649630975 hostname.mysql.commands.admir 1631 1649630975 hostname.mysql.commands.alter_ 0 1649630975

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Metrics (continued)

Type	Check	Description	Usage and Usage Example	Output
			provide the values against the properties 'user' and 'password' in the **client** section in the .ini file. -c, --critnum NUMBER Number of running threads upon which an alert is issued -w, --warnnum NUMBER Number of running threads upon which a warning is issued -l, --critlow NUMBER Number of running threads under which an alert is issued -m, --warnlow NUMBER Number of running threads under which a warning is issued -P, --port PORT MySQL Port to connect to -s, --socket UNIX socket to connect to (required if host specified is 'localhost' on UNIX - like systems) **Usage Example** `metrics-mysql-`	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Metrics (continued)

Type	Check	Description	Usage and Usage Example	Output
			`graphite.rb -h 127.0.0.1 -P 3306`	
Metric	app.mysql.metrics-mysql-processes	Returns various metrics regarding MySQL DB processes	/usr/local/bin/metrics-mysql-processes.rb (options) -h, --host MySQL host to connect to -i, --ini VALUE My.cnf ini file --ini-section VALUE Section in my.cnf ini file (Needed if .ini path is provided). To enable connection to MySQL thru the .ini file, provide the values against the properties 'user' and 'password' in the **client** section in the .ini file. -P, --port PORT MySQL Port to connect to -s --scheme SCHEME Metric naming scheme, text to append to metric -s, --socket UNIX socket to connect to (required if host specified is 'localhost' on UNIX - like systems)	processes, commands they're running and the databases they're running the commands on **Example:** • hostname.mysql.database.mysq 1 1649631113 • hostname.mysql.command.Daer 1 1649631113 • hostname.mysql.command.Slee 4 1649631113 • hostname.mysql.command.Que 1 1649631113

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

MySQL Metrics (continued)

Type	Check	Description	Usage and Usage Example	Output
			Usage Example `metrics-mysql-processes.rb -h 127.0.0.1 -P 3306`	

Network ping default checks and policies

Agent Client Collector provides the following default checks and policies for Network ping monitoring. Policies and checks are available for both Windows and Linux.

Network-ping metrics policy

Check	Description	Usage	Command	Output

util.metrics-ping	Runs on the ip address discovered in the ServiceNow instance. Pings the host and collects metrics.
	
	The monitored CI type is *cmdb_ci_hardware* and is bound to the relevant server.
	
	A proxy agent is required.
		metrics-ping.rb (options)
		
		`c --count COUNT Ping count`
		
		`-h Host to ping`
		
		`-t --timeout TIMEOUT Timeout`
		
		Example: ./ metrics-ping.rb -h 10.196.121.11
				`c7inc000473 83.ping.packets_transmitted 5 1675923 886`
				
				`c7inc0047 383.ping.packets_received 5 1675923 886`
				
				`c7inc0047 383.ping.packet_loss 0 1675923 886`
				
				`c7inc0047 383.ping.time 3999 1675923 886`
				
				`c7inc0047 383.ping.min 0.253`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Network-ping metrics policy (continued)

Check	Description	Usage	Command	Output
				1675923886 c7inc0047383.ping.avg 0.292 1675923886 c7inc0047383.ping.max 0.318 1675923886 c7inc0047383.ping.mdev 0.032 1675923886 **Note:** Windows does not support the *ping.time* and *ping.mdev* metrics.

Network ping metrics

The following table lists the metrics that are gathered as output from Network ping checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

network-ping metrics

Metric type	Units	Metric type description
ping.packets_transmitted	count	Number of data packets sent to the target host during the ping test.
ping.packets_received	count	Number of data packets successfully received by the target host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

network-ping metrics (continued)

Metric type	Units	Metric type description
ping.packet_loss	percentage	Percentage of data packets that failed to reach the target host.
ping.time	ms	Amount of time for the data packets to travel to the destination and back.
ping.min	ms	Shortest amount of time for a series of ping requests on the data packets.
ping.avg	ms	Average amount of time for a series of ping requests on the data packets.
ping.max	ms	Longest amount of time for a series of ping requests on the data packets.
ping.mdev	ms	Mean deviation (measure of the variability or dispersion) of the round-trip time (RTT) of the packets sent during the ping test.

Network host availability check

Agent Client Collector provides the following default check for network ping monitoring. The check is available for both Windows and Linux.

util.check-ping check

Check	Description	Usage	Command	Output
util.check-ping	Checks the network host availability via ICMP (ping) requests. Generates an OK event if the ping success ratio exceeds the Warning threshold (-W, default 50%). Generates a Warning event if the ping success ratio is between the warning threshold (-W,	`check-ping.rb -h <ip-address>` (options) `-h host`: Host to ping (Required). Default=<localhost> `-c count`: Number of ping requests sent to the host. Default=1 `-T timeout`: The timeout (in seconds) for each ping request. If the		`ICMP ping unsuccessful for host: 192.168.1.1 (successful: 0/5)`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

util.check-ping check (continued)

Check	Description	Usage	Command	Output
	default 50%) and the critical threshold (-C, default 20%). Generates a Critical event if the ping success ratio falls below the critical threshold (-C, default 20%). A proxy agent is required.	timeout expires before a reply is received, the ping is considered failed. Default=5 `-6 ipv6`: Enables IPv6 pinging. Ping uses IPv4 when this flag is not used. Default=False `-i interval`: Interval (in seconds) between each ping request. Default=1 `-W ratio`: The success ratio threshold for a Warning state. If the successful ping responses are below this ratio and above the -C ratio, the script returns a Warning. Default=0.5 `-C ratio`: The success ratio threshold for a Critical state. If the successful ping responses are below this ratio, the script returns a Critical failure. Default=0.2 `-r report`: Flag to append to the failure message of an MTR (My Trace Route) report if the ping fails.		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

util.check-ping check (continued)

Check	Description	Usage	Command	Output
		When MTR runs, a network diagnostic report generates. Default=false		

Nginx default checks and policies

Agent Client Collector provides the following policies for Nginx health monitoring. Policies come with the checks specified in the indicated table. Policies and checks are available for both Windows and Linux.

Nginx Events policy

Check	Description	Usage	Command	Output

nginx.check-nginx-alive	Verifies whether Nginx server is alive.
	Note: Use basic authentication for events.
		`check-nginx-status.rb`
		-h, --host
		Nginx hostname
		-p, --port
		Nginx port
		Note: Use basic authentication for events.
				Successful
				`CheckNginxStatus OK: Nginx is Alive and healthy`
				Failure
				`CheckNginxStatus CRITICAL: Nginx is Down`

Nginx Metrics policy

Check	Description	Usage	Command	Output

nginx.metrics-nginx-server	Retrieves the following information:
	nginx.active_connections
	Total number of active connections on server.
	nginx.accepts
		`metrics-nginx.rb`
		Options:
		`-u --url` Full URL to the server status page. For example: `https://yoursite.com/nginx_status`
				`ws19-INC0047517.nginx.active_connections 1 1640135552`
				`ws19-INC0047517.nginx.accepts 5746 1640135552`
				`ws19-INC0047517.nginx.handled`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Nginx Metrics policy (continued)

Check	Description	Usage	Command	Output
	The accepted connections. **nginx.handled** Handled connections, which normally equals the number of accepted connections. **nginx.requests** Total number of handled requests by the server. **nginx.read** Shows whether the server is reading requests. **nginx.writing** Shows whether the server is writing responses to the client. **nginx.waiting** Idle connections waiting for a request.			5746 1640135552 ws19-INC0047517.nginx.requests 9603 1640135552 ws19-INC0047517.nginx.reading 0 1640135552 ws19-INC0047517.nginx.writing 1 1640135552 ws19-INC0047517.nginx.waiting 0 1640135552

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Nginx Metrics policy (continued)

Check	Description	Usage	Command	Output
	Note: Use basic authentication for metrics.			

Nginx metrics

The following table lists the metrics that are gathered as output from Nginx checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

metrics-nginx metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
nginx.active_connections (featured metric)	admin	count	Total number of connections opened on the server, including waiting connections (*`nginx.waiting`* value). It reflects the number of current users on the website. There can be multiple connections opened for a single user.
nginx.accepts (featured metric)	admin	count	Represents the total established connections when `https://yoursite.com/nginx_status` is triggered.
nginx.handled (featured metric)	admin	count	Total number of connections handled by the Nginx server.
nginx.requests (featured metric)	admin	count	Total number of requests handled by the Nginx server.
nginx.reading (featured metric)	admin	count	Current number of connections where Nginx server is reading the request header.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-nginx metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
nginx.writing (featured metric)	admin	count	Represents the current number of connections where the Nginx server is writing the response back to the client.
nginx.waiting (featured metric)	admin	count	Number of idle client connections waiting for a request.

PostgreSQL default checks and policies

Agent Client Collector provides the following default checks and policies for PostgreSQL health monitoring.

PostgreSQL checks

Type	Check	Description	Command
Event	postgresql.check-alive	Triggers an alert on the status of the PostgreSQL database.	`commonchecks check-postgres-alive -d {{.labels.params_database}} -p {{.labels.params_port}} -H {{.labels.params_host}}`
Event	postgresql.check-connections	Triggers an alert on the total connection load on the PostgreSQL database, based on the configured threshold.	`commonchecks check-postgres-connections -d {{.labels.params_database}} -c {{.labels.params_critical}} -p {{.labels.params_port}} -H {{.labels.params_host}} -w {{.labels.params_warning}} -a {{.labels.params_use_percer`
Metric	postgresql.metric-active-connections	Provides metrics on the total active connection on the PostgreSQL database.	`commonchecks metric-postgres-connections -p {{.labels.params_port}} -d {{.labels.params_database}}`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PostgreSQL checks (continued)

Type	Check	Description	Command
			`-H {{.labels.params_host}}`
Metric	postgresql.metric-dbsize	Provides metrics on the total size for each of the server's PostgreSQL databases.	`commonchecks metric-postgres-dbsize -d {{.labels.params_database}} -H {{.labels.params_host}} -p {{.labels.params_port}}`
Metric	postgresql.metric-locks	Provides metrics on the locks in the PostgreSQL database/tables.	`commonchecks metric-postgres-locks -p {{.labels.params_port}} -d {{.labels.params_database}} -H {{.labels.params_host}}`
Metric	postgresql.metric-relation-size	Provides metrics on the database table size on the server.	`commonchecks metric-postgres-relation-size -p {{.labels.params_port}} -l {{.labels.params_limit}} -H {{.labels.params_host}} -d {{.labels.params_database}}`
Metric	postgresql.metric-statsbgwriter	Provides metrics related to buffer allocations.	`commonchecks metric-postgres-statsbgwriter -p {{.labels.params_port}} -d {{.labels.params_database}} -H {{.labels.params_host}}`
Metric	postgresql.metric-statsdb	Provides metrics related to commits/ rollbacks/tuples transactions / deadlocks/query conflicts count.	`commonchecks metric-postgres-statsdb -d {{.labels.params_database}} -H {{.labels.params_host}}`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PostgreSQL checks (continued)

Type	Check	Description	Command
			`-p {{.labels.params_port}}`
Metric	postgresql.metric-statsio	Provides table based metrics related index reads/total hits to a table index.	`commonchecks metric-postgres-statsio -s {{.labels.params_scope}} -H {{.labels.params_host}} -p {{.labels.params_port}} -d {{.labels.params_database}}`
Metric	postgresql.metric-statstable	Provides metrics related to CRUD operations.	`commonchecks metric-postgres-statstable -d {{.labels.params_database}} -H {{.labels.params_host}} -s {{.labels.params_scope}} -p {{.labels.params_port}}`

PostgreSQL metrics

The following tables list and describe the metrics that are gathered as output from the specified PostgreSQL checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

postgresql.check-connections metrics

Metric	Description
pgsql.connections.active (Featured metric)	Provides metrics on the total active connections in the PostgreSQL database.
pgsql.connections.idle (Featured metric)	Provides metrics on the total idle connections on the PostgreSQL database.

postgresql.metric-dbsize metrics

Metric	Description
pgsql.db.size (Featured metric)	Provides metrics on the total disk size utilization for each of the PostgreSQL databases on the server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

postgresql.metric-locks metrics

Metric	Description
pgsql.locks.AccessShareLock	Provides metrics on read-lock mode acquired automatically from queried tables.
pgsql.locks.ExclusiveLock	Provides metrics on read-lock mode acquired by LOCK TABLE table for IN EXCLUSIVE MODE statements.

postgresql.metric-relation-size metrics

Metric	Description
pgsql.tables_size (Featured metric)	Provides metrics on the database table size on the server.

postgresql.metric-statsbgwriter metrics

Metric	Description
pgsql.statsbgwriter.buffers_alloc	Provides metrics related to the number of allocated buffers.
pgsql.statsbgwriter.buffers_backend	Provides metrics related to the number of buffers written directly by a backend.
pgsql.statsbgwriter.buffers_backend_fsync	Provides metrics related to the number of times a backend had to execute its own fsync call. Normally the background writer handles these calls, even when the backend performs its own write.
pgsql.statsbgwriter.buffers_checkpoint	Provides metrics related to the number of buffers written during checkpoints.
pgsql.statsbgwriter.buffers_clean	Provides metrics related to the number of buffers written by the background writer.
pgsql.statsbgwriter.checkpoint_sync_time	Provides metrics related to the total amount of time spent in the portion of checkpoint processing where files are synchronized to disk, in milliseconds.
pgsql.statsbgwriter.checkpoint_write_time	Provides metrics related to the total amount of time spent in the portion of checkpoint processing where files are written to disk, in milliseconds.
pgsql.statsbgwriter.checkpoints_req (Featured metric)	Provides metrics related to the number of requested checkpoints that have been performed.
pgsql.statsbgwriter.checkpoints_timed (Featured metric)	Provides metrics related to the number of scheduled checkpoints that have been performed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

postgresql.metric-statsbgwriter metrics (continued)

Metric	Description
pgsql.statsbgwriter.maxwritten_clean	Provides metrics related to the number of times the background writer stopped a cleaning scan because it had written too many buffers.

postgresql.metric-statsdb metrics

Metric	Description
pgsql.statsdb.blk_read_time (Featured metric)	Provides metrics related to the time spent reading data file blocks by backends in this database, in milliseconds.
pgsql.statsdb.blk_write_time (Featured metric)	Provides metrics related to the time spent writing data file blocks by backends in this database, in milliseconds.
pgsql.statsdb.blks_hit (Featured metric)	Provides metrics related to the number of times disk blocks were found in the buffer cache, so that a read was not necessary. This only includes hits in the PostgreSQL buffer cache, not the operating system's file system cache.
pgsql.statsdb.blks_read	Provides metrics related to the number of disk blocks read in this database.
pgsql.statsdb.checksum_failures	Provides metrics related to the number of data page checksum failures detected in this database (or on a shared object), or 0 if data checksums are not enabled.
pgsql.statsdb.conflicts	Provides metrics related to the number of queries canceled due to conflicts with recovery in this database. Conflicts occur only on standby servers.
pgsql.statsdb.deadlocks (Featured metric)	Provides metrics related to the number of deadlocks detected in this database.
pgsql.statsdb.numbackends	Provides metrics related to the number of backends currently connected to this database. This is the only column in this view that returns a value reflecting the current state; all other columns return the accumulated values since the last reset.
pgsql.statsdb.temp_bytes	Provides metrics related to the total amount of data written to temporary files by queries in this database. All temporary files are counted, regardless of why the temporary file was created, and regardless of the *`log_temp_files`* setting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

postgresql.metric-statsdb metrics (continued)

Metric	Description
pgsql.statsdb.temp_files	Provides metrics related to the number of temporary files created by queries in this database. All temporary files are counted, regardless of why the temporary file was created (such as sorting or hashing), and regardless of the *log_temp_files* setting.
pgsql.statsdb.tup_deleted	Provides metrics related to the number of rows deleted by queries in this database.
pgsql.statsdb.tup_fetched	Provides metrics related to the number of rows fetched by queries in this database.
pgsql.statsdb.tup_inserted	Provides metrics related to the number of rows inserted by queries in this database.
pgsql.statsdb.tup_returned	Provides metrics related to the number of rows returned by queries in this database.
pgsql.statsdb.tup_updated	Provides metrics related to the number of rows updated by queries in this database.
pgsql.statsdb.xact_commit	Provides metrics related to the number of transactions in this database that have been committed.
pgsql.statsdb.xact_rollback	Provides metrics related to the number of transactions in this database that have been rolled back.

postgresql.metric-statsio metrics

Metric	Description
pgsql.statsio.heap_blks_hit (Featured metric)	Provides metrics related to the number of buffer hits in this table.
pgsql.statsio.heap_blks_read	Provides metrics related to the number of disk blocks read from this table.
pgsql.statsio.idx_blks_hit	Provides metrics related to the number of buffer hits in all indexes in this table.
pgsql.statsio.idx_blks_read	Provides metrics related to the number of disk blocks read from all indexes in this table.
pgsql.statsio.tidx_blks_hit	Provides metrics related to the number of buffer hits in this table's TOAST table index (if any).
pgsql.statsio.tidx_blks_read	Provides metrics related to the number of disk blocks read from this table's TOAST table index.
pgsql.statsio.toast_blks_hit	Provides metrics related to the number of buffer hits in this table's TOAST table (if any).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

postgresql.metric-statsio metrics (continued)

Metric	Description
pgsql.statsio.toast_blks_read	Provides metrics related to the number of disk blocks read from this table's TOAST table (if any).

postgresql.metric-statsio metrics

Metric	Description
pgsql.statstable.idx_scan	Provides metrics related to the number of index scans initiated on this table.
pgsql.statstable.idx_tup_fetch	Provides metrics related to the number of live rows fetched by index scans.
pgsql.statstable.n_dead_tup	Provides metrics related to the estimated number of dead rows.
pgsql.statstable.n_live_tup	Provides metrics related to the estimated number of live rows.
pgsql.statstable.n_tup_del	Provides metrics related to the number of deleted rows.
pgsql.statstable.n_tup_hot_upd	Provides metrics related to the number of rows HOT updated (that is, with no separate index update required).
pgsql.statstable.n_tup_ins	Provides metrics related to the number of inserted rows.
pgsql.statstable.n_tup_upd	Provides metrics related to the number of updated rows.
pgsql.statstable.seq_scan	Provides metrics related to the number of sequential scans initiated on this table.
pgsql.statstable.seq_tup_read	Provides metrics related to the number of live rows fetched by sequential scans.

RabbitMQ default checks and policies

Agent Client Collector provides the following default checks and policies for RabbitMQ health monitoring. You must perform RabbitMQ discovery before executing the checks. RabbitMQ checks are available only in a Windows environment.

RabbitMQ Events policy

Type	Check	Description	Command
Event	check-rabbitmq-alive	Verifies whether the RabbitMQ server is alive, using the REST API. If the server is down, an alert triggers.	`check-rabbitmq-alive.rb --host {{.labels.params_host}} --port {{.labels.params_port}} -v {{.labels.params_vhost}}`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

RabbitMQ Events policy (continued)

Type	Check	Description	Command
Event	check-rabbitmq-cluster-health	Verifies whether the RabbitMQ server's cluster nodes are running. If the nodes are down, an alert triggers.	`check-rabbitmq-cluster-health.rb --host {{.labels.params_host}} --port {{.labels.params_port}}`
Event	check-rabbitmq-consumers	Verifies the number of consumers on the RabbitMQ server and triggers an alert based on the configured threshold.	`check-rabbitmq-consumers.rb {{if .labels.params_warn}} --warn {{.labels.params_warn}} {{end}} {{if .labels.params_critica --critical {{.labels.params_critical}} {{end}} --host {{.labels.params_host}} --port {{.labels.params_port}}`
Event	check-rabbitmq-messages	Verifies the total number of messages queued on the RabbitMQ server and triggers an alert based on the threshold.	`check-rabbitmq-messages.rb --critical {{.labels.params_critical}} --port {{.labels.params_port}} --warn {{.labels.params_warn}} --host {{.labels.params_host}}`
Event	check-rabbitmq-network-partitions	Verifies whether the RabbitMQ network partition has occurred and triggers an alert based on the threshold.	check-rabbitmq-network-partitions.rb --host {{.labels.params_host}} --port {{.labels.params_port}}
Event	check-rabbitmq-node-health	Verifies whether the RabbitMQ server node is in a running state.	`check-rabbitmq-node-health.rb --host {{.labels.params_host}} {{if .labels.params_watchalarms}} --alarms {{.labels.params_watchalarms}} {{end}} {{if .labels.params_socketw`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

RabbitMQ Events policy (continued)

Type	Check	Description	Command
			`arn}} --swarn {{.labels.params_socketwarn}} {{end}} {{if .labels.params_memcrit}} --mcrit {{.labels.params_memcrit}} {{end}} {{if .labels.params_fdcrit}} --fcrit {{.labels.params_fdcrit}} {{end}} {{if .labels.params_socketcrit}} --scrit {{.labels.params_socketcrit}} {{end}} --port {{.labels.params_port}} {{if .labels.params_memwarn}} --mwarn {{.labels.params_memwarn}} {{end}} {{if .labels.params_fdwarn}} --fwarn {{.labels.params_fdwarn}} {{end}}`
Event	check-rabbitmq-node-usage	Checks and displays usage of the RabbitMQ server node.	`check-rabbitmq-node-usage.rb {{if .labels.params_procwarn}} --pwarn {{.labels.params_procwarn}} {{end}} --port {{.labels.params_port}} {{if .labels.params_socketw`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

RabbitMQ Events policy (continued)

Type	Check	Description	Command
			arn}} --swarn {{.labels.params_socketwarn}} {{end}} --type {{.labels.params_type}} {{if .labels.params_diskcrit}} --dcrit {{.labels.params_diskcrit}} {{end}} {{if .labels.params_fdcrit}} --fcrit {{.labels.params_fdcrit}} {{end}} {{if .labels.params_proccrit}} --pcrit {{.labels.params_proccrit}} {{end}} {{if .labels.params_diskwarn}} --dwarn {{.labels.params_diskwarn}} {{end}} {{if .labels.params_socketcrit}} --scrit {{.labels.params_socketcrit}} {{end}} --host {{.labels.params_host}} {{if .labels.params_memcrit}} --mcrit {{.labels.params_memcrit}} {{end}} {{if .labels.params_fdwarn}} --fwarn {{.labels.params_fdwarn}} {{end}} {{if .labels.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

RabbitMQ Events policy (continued)

Type	Check	Description	Command
			`params_memwarn}} mwarn {{.labels.params_memwarn}} {{end}}`
Event	check-rabbitmq-queue-drain-time	Verifies the time it will take for each queue on the RabbitMQ server to drain, based on the current message exit rate. For example, if a queue has 1,000 messages in it but only 1 message exits per second, an alert generates. The default critical level of 360 seconds has been exceeded.	`check-rabbitmq-queue-drain-time.rb --host {{.labels.params_host}} --port {{.labels.params_port}} --warn {{.labels.params_warn}} --critical {{.labels.params_critical}}`
Event	check-rabbitmq-queues-synchronised	Verifies that all mirrored queues with secondary queues are synchronised.	`check-rabbitmq-queues-synchronised.rb --host {{.labels.params_host}} --port {{.labels.params_port}}`
Event	check-rabbitmq-stomp-alive	Verifies whether the RabbitMQ server is alive and responding to STOMP.	`check-rabbitmq-stomp-alive.rb --host {{.labels.params_host}} --queue {{.labels.params_queue}} --port {{.labels.params_port}}`

RabbitMQ Metrics policy

Type	Check	Description	Command
Metric	metrics-rabbitmq-overview	Provides RabbitMQ overview statistics.	`metrics-rabbitmq-overview.rb --port {{.labels.params_port}} --host {{.labels.params_host}}`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

RabbitMQ Metrics policy (continued)

Type	Check	Description	Command
Metric	metrics-rabbitmq-queue	Provides RabbitMQ metrics per queue.	`metrics-rabbitmq-queue.rb --port {{.labels.params_port}} --host {{.labels.params_host}} {{if .labels.params_vhost}} --vhost {{.labels.params_vhost}} {{end}}`

RabbitMQ metrics

The following table lists the metrics that are gathered as output from RabbitMQ checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

metrics-rabbitmq-overview metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
rabbitmq.queue_totals.messages.count		count	Provides metrics on the total number of messages (ready plus unacknowledged).
rabbitmq.queue_totals.messages.rate (featured metric)		rate / sec	Provides metrics on the total number of messages (ready plus unacknowledged)
rabbitmq.queue_totals.messages_ready.count		count	Provides metrics on the number of messages ready for delivery.
rabbitmq.queue_totals.messages_ready.rate		rate / sec	Provides metrics on the number of messages ready for delivery.
rabbitmq.queue_totals.messages_unacknowledged.count		count	Provides metrics on the number of unacknowledged messages.
rabbitmq.queue_totals.messages_unacknowledged.rate		rate / sec	Provides metrics on the number of unacknowledged messages.
rabbitmq.message_stats.deliver_get.count		count	Provides metrics on messages

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-rabbitmq-overview metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			recently delivered to consumers.
rabbitmq.message_stats.deliver_get.rate (featured metric)		rate / sec	Provides metrics on messages recently delivered to consumers.
rabbitmq.message_stats.deliver_no_ack.count		count	Provides metrics on unacknowledged messages recently delivered to consumers.
rabbitmq.message_stats.deliver_no_ack.rate		rate / sec	Provides metrics on unacknowledged messages recently delivered to consumers.
rabbitmq.message_stats.publish.count		count	Provides metrics on recently published messages.
rabbitmq.message_stats.publish.rate		rate / sec	Provides metrics on the recently published messages.
rabbitmq.object_totals.channels.count		count	Provides metrics on the total number of channels.
rabbitmq.object_totals.connections.count (featured metric)		count	Provides metrics on the total number of connections.
rabbitmq.object_totals.consumers.count (featured metric)		count	Provides metrics on the total number of consumers.
rabbitmq.object_totals.exchanges.count		count	Provides metrics on the total number of exchanges.
rabbitmq.object_totals.queues.count (featured metric)		count	Provides metrics on the total number of queues.

metrics-rabbitmq-queue metrics

rabbitmq.queue.{{QueueName}}.avg_egress_rate (featured metric)	Queue Name	count	Provides queue specific metrics on the rate at which unacknowledged

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metrics-rabbitmq-queue metrics (continued)

			message records enter RAM for the given queue.
rabbitmq.queue.{{QueueName}}.consumers	Queue Name	count	Provides queue specific metrics on the total consumers for the given queue.
rabbitmq.queue.{{QueueName}}.messages	Queue Name	count	Provides queue specific metrics on the total messages in the given queue.
rabbitmq.queue.{{QueueName}}.drain_time	Queue Name	count	Provides queue specific metrics on the message drain time for the given queue.

URL Monitoring default checks and policies

Agent Client Collector provides policies for URL Monitoring. Polices come with embedded checks.

URL Monitoring uses the HTTP Entry Point default policies. For details, see HTTP default checks and policies.

Varnish default checks and policies

Agent Client Collector provides the following default checks and policies for Varnish monitoring.

Varnish metrics policy

Type	Check	Description	Usage	Output

Metric	app.varnish.metrics-varnish	Collects Varnish metrics from the host.	`linuxchecks metrics-varnish (options)`
			where:
			• -t = fields List of required metrics, comma separated within ""
			• -v = varnishname The Varnish instance to retrieve data from.
				`hostname.MAIN_sess_conn 175 1649420020`
				`hostname.MAIN_sess_drop 0 1649420020`
				`hostname.MAIN_client_req 173 1649420020`
				`hostname.MAIN_cache_hit 0 1649420020`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Varnish metrics policy (continued)

Type	Check	Description	Usage	Output
				hostname.MAIN_cache_hitpass 0 1649420020 hostname.MAIN_cache_miss 167 1649420020 hostname.MAIN_backend_conn 0 1649420020 hostname.MAIN_backend_unhealthy 0 1649420020 hostname.MAIN_backend_busy 0 1649420020 hostname.MAIN_backend_fail 171 1649420020 hostname.MAIN_backend_reuse 0 1649420020 hostname.MAIN_backend_recycle 0 1649420020 hostname.MAIN_threads 200 1649420020 hostname.MAIN_threads_limited 0 1649420020

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Varnish metrics policy (continued)

Type	Check	Description	Usage	Output
				hostname.MAIN_threads_created 200 1649420020 hostname.MAIN_threads_failed 0 1649420020 hostname.MAIN_thread_queue_len 0 1649420020 hostname.MAIN_sess_queued 0 1649420020 hostname.MAIN_n_expired 167 1649420020 hostname.MAIN_n_lru_nuked 0 1649420020 hostname.MAIN_backend_req 0 1649420020

Varnish metrics

The following table lists the metrics that are gathered as output from Varnish checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

Varnish metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
MAIN_sess_conn (featured metric)		count	Cumulative number of accepted client

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Varnish metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			connections by Varnish Cache.
MAIN_sess_drop (featured metric)		count	Number of connections dropped due to a full queue.
MAIN_client_req (featured metric)		count	Cumulative number of received client requests. Increments after a request is received, but before Varnish responds.
MAIN_cache_hit		count	Cumulative number of times a file was served from Varnish's cache.
MAIN_cache_hitpass		count	Cumulative number of hits for a "pass" file.
MAIN_cache_miss		count	Cumulative number of times a file was requested but was not in the cache, and was therefore requested from the backend.
MAIN_backend_conn		count	Cumulative number of successful TCP connections to the backend.
MAIN_backend_unhealthy (featured metric)		count	Cumulative number of backend connections which were not attempted because the backend has been marked as unhealthy.
MAIN_backend_busy		count	Cumulative number of times the maximum amount of connections to the backend has been reached.
MAIN_backend_fail (featured metric)		count	Cumulative number of failed connections to the backend.
MAIN_backend_reuse (featured metric)		count	Cumulative number of connections that were reused from the keep-alive pool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Varnish metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
MAIN_backend_recycle		count	Cumulative number of current backend connections which were put back to a pool of keep-alive connections and have not yet been used.
MAIN_threads (featured metric)		count	Number of threads in all pools.
MAIN_threads_limited		count	Number of times a thread needed to be created but couldn't because varnishd maxed out its configured capacity for new threads.
MAIN_threads_created		count	Number of times a thread has been created.
MAIN_threads_failed (featured metric)		count	Number of times that Varnish unsuccessfully tried to create a thread.
MAIN_thread_queue_len		count	Current queue length: number of requests waiting on worker thread to become available.
MAIN_sess_queued		count	Number of times Varnish has been out of threads and had to queue up a request.
MAIN_n_expired (featured metric)		count	Cumulative number of expired objects for example due to TTL.
MAIN_n_lru_nuked (featured metric)		count	Least Recently Used Nuked Objects: Cumulative number of cached objects that Varnish has evicted from the cache because of a lack of space.
MAIN_backend_req (featured metric)		count	Number of requests to the backend.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vSphere default checks and policies

Agent Client Collector provides the following default checks and policies for vSphere monitoring.

vSphere virtualMachine metrics policy

Type	Check	Description	Usage	Output
Metric	vsphere.metric-VirtualMachine	Retrieves metrics for a vSphere virtual machine. Must use a proxy agent.	`commonchecks metric-vsphere -H hostname -R VirtualMachine -v vm_id -r true -i 20` where: • -H = hostname Instance Hostname • -i = intervalId 20 for real time metrics; 300 for 5 minutes older aggregate metrics (default 300) • -r = resourceLevelMetric To enable resource level metrics (default=false)	`vsphere.VirtualMachine.c 8 1654755994`

vSphere Datastore Metrics policy

Type	Check	Description	Usage	Output
Metric	vSphere.metric-Datastore	Retrieves Metrics for vSphere Datastore. Must use a proxy agent.	`commonchecks metric-vsphere -H hostname -R Datastore -v vm_id -r` where: • -H = hostname Instance Hostname • -r = resourceLevelMetric To enable	`vsphere.Datastore.disk.c 21474836480 1655182836`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vSphere Datastore Metrics policy (continued)

Type	Check	Description	Usage	Output
			resource level metrics (default=false)	

vSphere Datacenter Metrics policy

Type	Check	Description	Usage	Output

Metric	vsphere.metric-Datacenter	Retrieves metrics for the vSphere datacenter.	`commonchecks metric-vsphere -H hostname -R Datacenter -v vm_id -r`
			where:
			• -H = hostname Instance Hostname
			• -r = resourceLevelMet To enable resource level metrics (default=false)
				`vsphere.Datacenter.vmop.numChangeDS.latest.total 217 1655194924`
				`vsphere.Datacenter.vmop.numChangeHost.latest.total 1373 1655194924`
				`vsphere.Datacenter.vmop.numChangeHostDS.latest.total 9 1655194924`
				`vsphere.Datacenter.vmop.numClone.latest.total 192 1655194924`

vSphere ESX Server Metrics policy

Type	Check	Description	Usage	Output
Metric	vsphere.metric-ESX_Server	Retrieves metrics for vSphere ESX server or HostSystem. Must use a proxy agent.	`commonchecks metric-vsphere -H hostname -R HostSystem -v vm_id -r`	`vsphere.HostSystem.net.broadcastRx.summation.total 63885`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vSphere ESX Server Metrics policy (continued)

Type	Check	Description	Usage	Output
			where: • -H = hostname Instance Hostname • -r = resourceLevelMet To enable resource level metrics (default=false)	1655198 166 vsphere.HostSystem.net.broadcastTx.summation.total 196 1655198 166 vsphere.HostSystem.net.bytesRx.average.total 5924 1655198 166

vSphere metrics

The following table lists the metrics that are gathered as output from vSphere checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

vsphere.metric-VirtualMachine metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
net.broadcastRx.summation	network adapter	packets	Number of broadcast packets received.
net.broadcastTx.summation	network adapter	packets	Number of broadcast packets transmitted.
net.bytesRx.average	network adapter	KiB	Average amount of data received per second.
net.bytesTx.average	network adapter	KiB	Average amount of data transmitted per second.
net.droppedRx.summation	network adapter	packets	Number of received packets dropped.
net.droppedTx.summation	network adapter	packets	Number of transmitted packets dropped.
net.multicastRx.summation	network adapter	packets	Number of multicast packets received.
net.multicastTx.summation	network adapter	packets	Number of multicast packets transmitted.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-VirtualMachine metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
net.packetsRx.summation	network adapter	packets	Number of packets received.
net.packetsTx.summation	network adapter	packets	Number of packets transmitted.
net.received.average	network adapter	KiB	Average rate at which data was received during the interval. This represents the bandwidth of the network.
net.transmitted.average	network adapter	KiB	Average rate at which data was transmitted during the interval. This represents the bandwidth of the network.
net.usage.average (featured metric)	network adapter	KiB	Network utilization (combined transmit- and receive-rates).
disk.maxTotalLatency.latest (featured metric)	drive	millisecond	Highest latency value across all disks used by the host.
disk.read.average	drive	KiB	Average number of kilobytes read from the disk each second.
disk.usage.average (featured metric)	drive	KiB	Aggregated disk I/O rate.
disk.usage.maximum	drive	KiB	Maximum disk I/O rate.
disk.usage.minimum	drive	KiB	Minimum disk I/O rate.
disk.write.average	drive	KiB	Average number of kilobytes written to the disk each second.
cpu.costop.summation	cpu-core	millisecond	Time the virtual machine is ready to run, but is unable to run due to co-scheduling constraints.
cpu.demand.average	cpu-core	megahertz	The amount of CPU resources a virtual machine would use if there were no CPU

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-VirtualMachine metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			contention or CPU limit.
cpu.demandEntitlementRatio.latest	cpu-core	percent	CPU resource entitlement to CPU demand ratio.
cpu.entitlement.latest	cpu-core	megahertz	CPU resources devoted by the ESXi scheduler.
cpu.idle.summation	cpu-core	millisecond	Total time that the CPU spent in an idle state.
cpu.latency.average	cpu-core	percent	Percent of time the virtual machine is unable to run because it is contending for access to the physical CPU(s).
cpu.maxlimited.summation	cpu-core	millisecond	Time the virtual machine is ready to run, but is not running because it has reached its maximum CPU limit setting.
cpu.overlap.summation	cpu-core	millisecond	Time the virtual machine was interrupted to perform system services on behalf of itself or other virtual machines.
cpu.readiness.average	cpu-core	percent	Percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU.
cpu.ready.summation	cpu-core	millisecond	Milliseconds of CPU time spent in ready state.
cpu.run.summation	cpu-core	millisecond	Time the virtual machine is scheduled to run.
cpu.swapwait.summation	cpu-core	millisecond	CPU time spent waiting for swap-in.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-VirtualMachine metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cpu.system.summation	cpu-core	millisecond	Amount of time spent on system processes on each virtual CPU in the virtual machine.
cpu.usage.average (featured metric)	cpu-core	percent	Percentage of CPU capacity being used.
cpu.usagemhz.average	cpu-core	megahertz	CPU usage, as measured in megahertz.
cpu.used.summation	cpu-core	millisecond	Time accounted to the virtual machine. If a system service runs on behalf of this virtual machine, the time spent by that service (represented by cpu.system) should be charged to this virtual machine. If not, the time spent (represented by cpu.overlap) should not be charged against this virtual machine.
cpu.wait.summation (featured metric)	cpu-core	millisecond	Total CPU time spent in wait state.The wait total includes time spent the CPU Idle, CPU Swap Wait, and CPU I/O Wait states.
cpu.usage.vcpus.average	cpu-core	percent	Percentage of vcpus capacity being used.
cpu.utilization.average	cpu-core	percent	Percentage of CPU capacity being used.
mem.active.average	memory	KiB	Amount of memory that is actively used, as estimated by VMkernel based on recently touched memory pages.
mem.activewrite.average	memory	KiB	Estimate for the amount of memory actively being written to by the virtual machine.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-VirtualMachine metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mem.compressed.average	memory	KiB	Kilobytes of memory that have been compressed.
mem.compressionRate.average	memory	KiB	Rate of memory compression for the virtual machine.
mem.consumed.average	memory	KiB	Amount of host physical memory consumed by a virtual machine, host, or cluster.
mem.decompressionRate.average	memory	KiB	Rate of memory decompression for the virtual machine.
mem.entitlement.average	memory	KiB	Amount of host physical memory the virtual machine is entitled to, as determined by the ESX scheduler.
mem.granted.average	memory	KiB	Amount of host physical memory or physical memory that is mapped for a virtual machine or a host.
mem.latency.average	memory	KiB	Percentage of time the virtual machine is waiting to access swapped or compressed memory.
mem.llSwapInRate.average	memory	KiB	Rate at which memory is being swapped from host cache into active memory.
mem.llSwapOutRate.average	memory	KiB	Rate at which memory is being swapped from active memory to host cache.
mem.llSwapUsed.average	memory	KiB	Space used for caching swapped pages in the host cache.
mem.overhead.average	memory	KiB	Host physical memory consumed by the virtualization

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-VirtualMachine metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			infrastructure for running the virtual machine.
mem.overheadMax.average	memory	KiB	Host physical memory reserved for use as the virtualization overhead for the virtual machine.
mem.overheadTouched.average	memory	KiB	Actively touched overhead host physical memory (KB) reserved for use as the virtualization overhead for the virtual machine.
mem.shared.average	memory	KiB	Amount of guest physical memory that is shared with other virtual machines, relative to a single virtual machine or to all powered-on virtual machines on a host.
mem.usage.average (featured metric)	memory	percent	Memory usage as percent of total configured or available memory.
mem.vmmemctl.average	memory	KiB	Amount of memory allocated by the virtual machine memory control driver.
mem.vmmemctltarget.average	memory	KiB	Target value set by VMkernal for the virtual machine's memory balloon size. In conjunction with vmmemctl metric, this metric is used by VMkernel to inflate and deflate the balloon for a virtual machine.
mem.zero.average	memory	KiB	Memory that contains 0s only. Included in shared amount. Through transparent page sharing, zero

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-VirtualMachine metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			memory pages can be shared among virtual machines that run the same operating system.
power.energy.summation	power	joule	Total energy (in joule) used since last stats reset.
power.power.average	power	watt	Current power usage.
sys.heartbeat.latest	system	number	Number of heartbeats issued per virtual machine.
sys.osUptime.latest	system	seconds	Total time elapsed, in seconds, since last operating system boot-up.
sys.uptime.latest	system	seconds	Total time elapsed since last system startup.
virtualDisk.read.average	virtual disk	KiB	Average number of kilobytes read from the virtual disk each second.
virtualDisk.write.average	virtual disk	KiB	Average number of kilobytes written to the virtual disk each second.
rescpu.actav1.latest	rescpu	percent	CPU active average over 1 minute.
rescpu.actav15.latest	rescpu	percent	CPU active average over 15 minutes.
rescpu.actav5.latest	rescpu	percent	CPU active average over 5 minutes.
rescpu.actpk1.latest	rescpu	percent	CPU active peak over 1 minute.
rescpu.actpk15.latest	rescpu	percent	CPU active peak over 15 minutes.
rescpu.actpk5.latest	rescpu	percent	CPU active peak over 5 minutes.
rescpu.runav1.latest	rescpu	percent	CPU running average over 1 minute.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-VirtualMachine metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
rescpu.runav15.latest	rescpu	percent	CPU running average over 15 minutes.
rescpu.runav5.latest	rescpu	percent	CPU running average over 5 minutes.
rescpu.runpk1.latest	rescpu	percent	CPU running peak over 1 minute.
rescpu.runpk15.latest	rescpu	percent	CPU running peak over 15 minutes.
rescpu.runpk5.latest	rescpu	percent	CPU running peak over 5 minutes.
rescpu.sampleCount.latest	rescpu	millisecond	Group CPU sample count.
rescpu.samplePeriod.latest	rescpu	millisecond	Group CPU sample period.

vsphere.metric-Datacenter metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmop.numChangeDS.latest	Virtual Machine Operations Counters	Number	Number of datastore change operations for powered-off and suspended virtual machines.
vmop.numChangeHost.latest	Virtual Machine Operations Counters	Number	Number of host change operations for powered-off and suspended virtual machines.
vmop.numChangeHostDS.latest	Virtual Machine Operations Counters	Number	Number of host and datastore change operations for powered-off and suspended virtual machines.
vmop.numClone.latest	Virtual Machine Operations Counters	Number	Number of virtual machine clone operations.
vmop.numCreate.latest (featured metric)	Virtual Machine Operations Counters	Number	Number of virtual machine create operations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-Datacenter metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
vmop.numDeploy.latest	Virtual Machine Operations Counters	Number	Number of virtual machine template deploy operations.
vmop.numDestroy.latest (featured metric)	Virtual Machine Operations Counters	Number	Number of virtual machine delete operations.
vmop.numPoweroff.latest	Virtual Machine Operations Counters	Number	Number of virtual machine power off operations.
vmop.numPoweron.latest	Virtual Machine Operations Counters	Number	Number of virtual machine power on operations.
vmop.numRebootGuest.latest (featured metric)	Virtual Machine Operations Counters	Number	Number of virtual machine guest reboot operations.
vmop.numReconfigure.latest	Virtual Machine Operations Counters	Number	Number of virtual machine reconfigure operations.
vmop.numRegister.latest	Virtual Machine Operations Counters	Number	Number of virtual machine register operations.
vmop.numReset.latest	Virtual Machine Operations Counters	Number	Number of virtual machine reset operations.
vmop.numSVMotion.latest	Virtual Machine Operations Counters	Number	Number of migrations with Storage vMotion (datastore change operations for powered-on VMs).
vmop.numShutdownGuest.latest	Virtual Machine Operations Counters	Number	Number of virtual machine guest shutdown operations.
vmop.numStandbyGuest.latest	Virtual Machine Operations Counters	Number	Number of virtual machine standby guest operations.
vmop.numSuspend.latest	Virtual Machine Operations Counters	Number	Number of virtual machine suspend operations.
vmop.numUnregister.latest	Virtual Machine Operations Counters	Number	Number of virtual machine unregister operations.
vmop.numVMotion.latest (featured metric)	Virtual Machine Operations Counters	Number	Number of migrations with vMotion (host

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-Datacenter metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			change operations for powered-on VMs).
vmop.numXVMotion.latest	Virtual Machine Operations Counters	Number	Number of host and datastore change operations for powered-on and suspended virtual machines.

vsphere.metric-Datastore metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
datastore.busResets.summation	Datastore	Number	Number of SCSI-bus reset commands issued.
datastore.commandsAborted.summation	Datastore	Number	Number of SCSI commands aborted.
datastore.datastoreIops.average	Datastore	Number	Storage I/O Control aggregated IOPS.
datastore.datastoreMaxQueueDepth.latest	Datastore	Number	Storage I/O Control datastore maximum queue depth.
datastore.datastoreNormalReadLatency.latest	Datastore	millisecond	Storage DRS datastore normalized read latency.
datastore.datastoreNormalWriteLatency.latest	Datastore	millisecond	Storage DRS datastore normalized write latency.
datastore.datastoreReadBytes.latest	Datastore	millisecond	Storage DRS datastore bytes read.
datastore.datastoreReadIops.latest	Datastore	Number	Storage DRS datastore read I/O rate.
datastore.datastoreReadLoadMetric.latest	Datastore	Number	Storage DRS datastore metric for read workload model.
datastore.datastoreReadOutstandingIO.latest	Datastore	Number	Storage DRS datastore outstanding read requests.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-Datastore metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
datastore.datastoreVMObservedLatency.latest	Datastore	microsecond	The average datastore latency as seen by virtual machines.
datastore.datastoreWriteBytes.latest	Datastore	millisecond	Storage DRS datastore bytes written.
datastore.datastoreWriteIops.latest	Datastore	Number	Storage DRS datastore write I/O rate.
datastore.datastoreWriteLoadMetric.latest	Datastore	Number	Storage DRS datastore metric for write workload model.
datastore.datastoreWriteOutstanding.latest	Datastore	Number	Storage DRS datastore outstanding write requests.
datastore.maxTotalLatency.latest	Datastore	millisecond	Highest latency value across all datastores used by the host.
datastore.numberReadAveraged.average	Datastore	Number	Average number of read commands issued per second to the datastore.
datastore.numberWriteAveraged.average	Datastore	Number	Average number of write commands issued per second to the datastore during the collection interval.
datastore.read.average (featured metric)	Datastore	KiB	Rate of reading data from the datastore.
datastore.siocActiveTimePercentage.average (featured metric)	Datastore	percent	Percentage of time Storage I/O Control actively controlled datastore latency.
datastore.sizeNormalizedDatastoreLatency.average	Datastore	microsecond	Storage I/O Control size-normalized I/O latency.
datastore.throughput.contention.average	Datastore	millisecond	Average amount of time for an I/O operation to the datastore or LUN across all ESX hosts accessing it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-Datastore metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
datastore.throughput.usage.average	Datastore	KiB	The current bandwidth usage for the datastore or LUN.
datastore.totalReadLatency.average	Datastore	millisecond	Average amount of time for a read operation from the datastore.
datastore.totalWriteLatency.average	Datastore	millisecond	Average amount of time for a write operation from the datastore.
datastore.unmapIOs.summation	Datastore	Number	Number of Unmap operations.
datastore.unmapSize.summation	Datastore	Number	Total UnMap size.
datastore.write.average (featured metric)	Datastore	KiB	Rate of writing data to the datastore.
disk.busResets.summation	disk	Number	Number of SCSI-bus reset commands issued.
disk.capacity.latest	disk	KiB	Configured size of the datastore.
disk.capacity.contention.average	disk	KiB	The amount of storage capacity overcommitment for the entity, measured in percent.
disk.capacity.provisioned.average	disk	KiB	Provisioned size of the entity.
disk.capacity.usage.average	disk	KiB	The amount of storage capacity currently being consumed by or on the entity.
disk.commands.summation	disk	Number	Number of SCSI commands issued.
disk.commandsAborted.summation	disk	Number	Number of SCSI commands aborted.
disk.commandsAveraged.average	disk	Number	Average number of SCSI commands issued per second.
disk.deltaused.latest	disk	KiB	Storage overhead of a virtual machine or a

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-Datastore metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			datastore due to delta disk backings.
disk.deviceLatency.average	disk	millisecond	Average amount of time it takes to complete an SCSI command from physical device.
disk.deviceReadLatency.average	disk	millisecond	Average amount of time it takes to complete read from physical device.
disk.deviceWriteLatency.average	disk	millisecond	Average amount of time to write from the physical device.
disk.kernelLatency.average	disk	millisecond	Average amount of time spent by VMkernel to process each SCSI command.
disk.kernelReadLatency.average	disk	millisecond	Average amount of time spent by VMkernel to process each SCSI read command.
disk.kernelWriteLatency.average	disk	millisecond	Average amount of time spent by VMkernel to process each SCSI write command.
disk.maxQueueDepth.average (featured metric)	disk	Number	Maximum queue depth.
disk.maxTotalLatency.latest	disk	millisecond	Highest latency value across all disks used by the host.
disk.numberRead.summation	disk	Number	Number of disk reads during the collection interval.
disk.numberReadAveraged.average	disk	Number	Average number of read commands issued per second to the datastore.
disk.numberWrite.summation	disk	Number	Number of disk writes during the collection.
disk.numberWriteAveraged.average	disk	Number	Average number of write commands

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-Datastore metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			issued per second to the datastore.
disk.provisioned.latest	disk	KiB	Amount of storage set aside for use by a datastore or a virtual machine. Files on the datastore and the virtual machine can expand to this size but not beyond it.
disk.queueLatency.average	disk	millisecond	Average amount of time spent in VMkernel queue (per SCSI command).
disk.queueReadLatency.average (featured metric)	disk	millisecond	Average amount of time spent in the VMkernel queue per SCSI read command.
disk.queueWriteLatency.average	disk	millisecond	Average amount of time spent in the VMkernel queue per SCSI write command.
disk.read.average	disk	KiB	Average number of kilobytes read from the disk each second.
disk.scsiReservationCnflctPct.average	disk	Number	Number of SCSI reservation conflicts for the LUN as a percent of total commands during the collection interval.
disk.scsiReservationConflicts.summation	disk	Number	Number of SCSI reservation conflicts for the LUN during the collection interval.
disk.throughput.contention.average	disk	millisecond	Average amount of time for an I/O operation to the datastore or LUN across all ESX hosts accessing it.
disk.throughput.usage.average	disk	KiB	The current bandwidth usage for the datastore or LUN.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-Datastore metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
disk.totalLatency.average (featured metric)	disk	millisecond	Average amount of time taken during the collection interval to process a SCSI command issued by the guest OS to the virtual machine.
disk.totalReadLatency.average	disk	millisecond	Average amount of time taken to process a SCSI read command issued from the guest OS to the virtual machine.
disk.totalWriteLatency.average	disk	millisecond	Average amount of time taken to process a SCSI write command issued by the guest OS to the virtual machine.
disk.unshared.latest	disk	KiB	Amount of space associated exclusively with a virtual machine.
disk.usage.average (featured metric)	disk	KiB	Aggregated disk I/O rate.
disk.used.latest	disk	KiB	Amount of space actually used by the virtual machine or the datastore. May be less than the amount provisioned at any given time. This depends on whether the virtual machine is powered off, whether snapshots have been created, and other factors.
disk.write.average	disk	KiB	Average number of kilobytes written to the disk each second.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-ESX_Server metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
cpu.coreUtilization.average (featured metric)	cpu-core	percent	CPU utilization of the corresponding core (if hyper-threading is enabled) as a percentage.
cpu.reservedCapacity.average	cpu-core	megahertz	Total CPU capacity reserved by virtual machines.
cpu.totalCapacity.average	cpu-core	megahertz	Total CPU capacity reserved by and available for virtual machines.
datastore.maxTotalLatency.latest	datastore	millisecond	Highest latency value across all datastores used by the host.
hbr.hbrDiskReadLatency.average	host-based replication	millisecond	Average Disk read latency.
hbr.hbrDiskStallLatency.average	host-based replication	millisecond	Average Disk stall latency.
hbr.hbrNetRx.average	host-based replication	KiB	Average amount of data received per second.
hbr.hbrNetTx.average	host-based replication	KiB	Average amount of data transmitted per second.
hbr.hbrNumVms.average	host-based replication	Number	Number of powered-on virtual machines running on this host that currently have host-based replication protection enabled.
mem.heap.average	memory	KiB	VMkernel virtual address space dedicated to VMkernel main heap and related data.
mem.heapfree.average	memory	KiB	Free address space in the VMkernel main heap.Varies based on number of physical devices and configuration options.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-ESX_Server metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mem.llSwapIn.average	memory	KiB	Amount of memory swapped-in from host cache.
mem.llSwapOut.average	memory	KiB	Amount of memory swapped-out to host cache.
mem.lowfreethreshold.average	memory	KiB	Threshold of free host physical memory below which ESX/ESXi will begin reclaiming memory from virtual machines through ballooning and swapping.
mem.reservedCapacity.average	memory	MiB	Total amount of memory reservation used by powered-on virtual machines and vSphere services on the host.
mem.sharedcommon.average	memory	KiB	Amount of machine memory that is shared by all powered-on virtual machines and vSphere services on the host.
mem.state.latest	memory	KiB	One of four threshold levels representing the percentage of free memory on the host.
mem.swapin.average	memory	KiB	Amount of memory swapped-in from disk.
mem.swapinRate.average	memory	KiB	Rate at which memory is swapped from disk into active memory.
mem.swapout.average	memory	KiB	Amount of memory swapped-out to disk.
mem.swapoutRate.average	memory	KiB	Rate at which memory is being swapped from active memory to disk.
mem.swapused.average	memory	KiB	Amount of memory that is used by swap. Sum of memory swapped of all

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-ESX_Server metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
			powered on VMs and vSphere services on the host.
mem.sysUsage.average	memory	KiB	Amount of host physical memory used by VMkernel for core functionality, such as device drivers and other internal uses. Does not include memory used by virtual machines or vSphere services.
mem.totalCapacity.average	memory	KiB	Total amount of memory reservation used by and available for powered-on virtual machines and vSphere services on the host.
mem.unreserved.average	memory	KiB	Amount of memory that is unreserved. Memory reservation not used by the Service Console, VMkernel, vSphere services and other powered on VMs user-specified memory reservations and overhead memory.
mem.vmfs.pbc.capMissRatio.latest (featured metric)	memory	percent	Trailing average of the ratio of capacity misses to compulsory misses for the VMFS PB Cache.
mem.vmfs.pbc.overhead.latest	memory	KiB	Amount of VMFS heap used by the VMFS PB Cache.
mem.vmfs.pbc.size.latest	memory	MiB	Space used for holding VMFS Pointer Blocks in memory.
mem.vmfs.pbc.sizeMax.latest	memory	MiB	Maximum size the VMFS Pointer Block Cache can grow to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

vsphere.metric-ESX_Server metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
mem.vmfs.pbc.workingSetLatest	memory	TiB	Amount of file blocks whose addresses are cached in the VMFS PB Cache.
mem.vmfs.pbc.workingSetMax.latest	memory	TiB	Maximum amount of file blocks whose addresses are cached in the VMFS PB Cache.
net.errorsRx.summation (featured metric)	network adapter	Number	Number of packets with errors received.
net.errorsTx.summation (featured metric)	network adapter	Number	Number of packets with errors transmitted.
net.unknownProtos.summation	network adapter	KiB	Number of frames with unknown protocol received.
power.powerCap.average	power	watt	Maximum allowed power usage.
rescpu.maxLimited1.latest	rescpu	percent	Amount of CPU resources over the limit that were refused, average over 1 minute.
rescpu.maxLimited15.latest	rescpu	percent	Amount of CPU resources over the limit that were refused, average over 15 minutes.
rescpu.maxLimited5.latest	rescpu	percent	Amount of CPU resources over the limit that were refused, average over 5 minutes.
storageAdapter.maxTotalLatency.latest (featured metric)	storage adapter	millisecond	Highest latency value across all storage adapters used by the host.
storagePath.maxTotalLatency.latest (featured metric)	storage adapter	millisecond	Highest latency value across all storage paths used by the host.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows default checks and policies

Agent Client Collector provides the following default checks and policies for Windows health monitoring.

Windows event monitoring checks

Windows OS Events policy

Check	Description	Usage and Example	Output
os.windows.check-event-log	Measures the Windows event log against parameter thresholds and returns a *CRITICAL* | *WARNING* | *OK* event.	Usage: • -w warning - Triggers a WARNING event if the event log count matching the pattern is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the event log count matching the pattern is above the CRITICAL parameter value specified in the check parameter. • -e event level - Specifies the severity level of the event. Possible values: Information, Verbose, Critical, Warning, Error. • -i - Unique event ID • -d - The duration of time, in hours, in which you want to retrieve events from the Windows event log. Usage example: `winchecks check-windows-event-log -w 5 -c 10 -e "Information" -l`	Check Event Log OK: The Event Log that matches the pattern is <matched count>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		`"Application" -d 24`	
os.windows.check-event-log-count	Measures the Windows event log against parameter thresholds and returns a CRITICAL \WARNING\OK event. Provides information on the number of events that have occurred within a specified duration for a single log file and a single ID. Also indicates the filters to be applied to retrieve events for a specific single-valued windows event level and provider name. Retrieving events from multiple log files is not supported. The number of events is provided, without details of each and every event.	Usage: • -w warning - Triggers a WARNING event if the event log count matching the pattern is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the event log count matching the pattern is above the CRITICAL parameter value specified in the check parameter. • -l log_file - The log file to be monitored. Name of the file is written in double quotation marks. • -r regex_pattern - The regex pattern which filters out the description in the event log. Written in double quotation marks. • -e event level - Specifies the severity level of the event. Possible values: Information, Verbose, Critical, Warning, Error. • -i id - Unique event ID • -d duration_hour - The duration of time, in hours, in which you want to retrieve events	Check Event Log OK: The Event Log that matches the pattern is <matched count>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		from the Windows event log. Decimal points can be used; for example, 30 minutes - 0.5. • -p provider_name - Source of the event, written in double quotation marks. Usage example: `winchecks check-windows-event-log -w 5 -c 10 -e "Information" -l "Application" -d 24`	
os.windows.check-event-log-details	Collects and filters Windows Event logs based on the `duration_hour`, `event_log_level` and `log_file` values. Retrieves and filters Windows event logs according to the provided parameters. It returns details about the events with CRITICAL, WARNING, or OK status, based on the specified severity level.	Usage: • -d duration_hour - Duration (in hours) from the current time to filter events (Default: 24). • -e event_log_level - Filter the events based on the event level. Possible values are: Information, Verbose, Critical, Warning, Error. Multiple values are comma-separated (Default: Information). For example: Information, Warning • -i id - Filters events based on the specified event IDs. For multiple IDs, values are comma-separated and enclosed in double quotation marks.	Check Event Log Details WARNING: Type: Information, Category: Application, Machine: ws19-inc0061393.LOCAL.LAB, Event_ID: 1704, Message: Security policy in the Group policy objects has been applied successfully., TimeCreated: 10/14/2024 12:09:35 AM. Type: Information, Category: Application, Machine: ws19-inc0061393.LOCAL.LAB, Event_ID: 16384, Message: Successfully scheduled Software Protection service for restart at 2124-09-20T06:25:44Z. Reason: Rules Engine, TimeCreated: 10/13/2024 11:25:44 PM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		For example: "1257, 1001" • -l log_file - Specifies the log file name to filter events. The name of the file is written in double quotation marks. Supports creating custom files and multiple values are comma-separated. (Default: Application). For example: "Application, System" • -p provider_name - The name of the event provider, enclosed in double quotation marks. • -r regex_pattern - Filters events by matching the event message with the specified pattern. Value must be enclosed in double quotation marks. • -s servicenow_event_severity - Creates a servicenow event with the value given in this parameter. Possible values are: Critical, Warning and OK. Usage example: `winchecks check-windows-event-log-details -d 24 -l Application -e Warning -r "*" -s Warning`	Type: Information, Category: Application, Machine: ws19-inc0061393.LOCAL.LAB, Event_ID: 16394, Message: Offline downlevel migration succeeded., TimeCreated: 10/13/2024 11:24:19 PM. Type: Information, Category: Application, Machine: ws19-inc0061393.LOCAL.LAB, Event_ID: 8224, Message: The VSS service is shutting down due to idle timeout., TimeCreated: 10/13/2024 11:51:36 AM.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
os.windows.check-disk-name	Takes the storage drive name as input and verifies if the drive is present. Returns a *CRITICAL | WARNING | OK* value based on the parameter provided.	winchecks check-windows-disk-name <options> -d : Disk name (Default = C) Usage example:`winchecks check-windows-disk-name -d C`	Windows Checks OK: Disk storage C is present.
os.windows.check-processor-queue-length	Measures the process queue length against thresholds and returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -w warning - Triggers a WARNING event if the processor queue length count matching the pattern is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the processor queue length count matching the pattern is above the CRITICAL parameter value specified in the check parameter. Usage example: `winchecks check-windows-processor-queue-length -w 5 -c 10`	Processor Queue Length OK: The Processor Queue length is 0.00
os.windows.check-system-cpu-load	Checks CPU Load by using typeperf. Measures the CPU load against configured thresholds and returns a CRITICAL \WARNING\OK event	Usage: • -w warning - Triggers a WARNING event if the CPU load count matching the pattern is above	CPU Load OK: The total CPU utilization is 26.92%

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
	according to the thresholds given in the accompanying parameters.	the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the CPU load count matching the pattern is above the CRITICAL parameter value specified in the check parameter. Usage example: `winchecks check-windows-cpu-load -w 85 -c 95`	
os.windows.check-system-disk	Measures the free physical memory against thresholds and returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -w warning - Triggers a WARNING event if the event log percentage matching the pattern is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the event log percentage matching the pattern is above the CRITICAL parameter value specified in the check parameter. • -t fs_type - Checks only volumes with the specified filesystem types. • -m mount_points - Checks only the specified mount points.	Disk Usage Check OK: The disk usage is %

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		• -x ignore_type - Skip volumes with the specified filesystem types. • -i ignore_mnt - Skip the specified mount points when collecting disk usage. Returns an error if a specified mount point does not exist in the system. • -r ignore_mnt_regex - Skip volumes whose names match the given regular expression. • -l ignore_label - Skip volumes whose labels match the given regular expression. Usage example: `winchecks check-windows-disk -w 85 -c 95`	
os.windows.check-system-memory-percent	Collects the RAM usage. Measures the memory usage against configured thresholds and returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -w warning - Triggers a WARNING event if the memory use percentage matching the pattern is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the memory use percentage matching the pattern is above the CRITICAL	RAM Usage OK: The total memory utilization is 84%

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		parameter value specified in the check parameter. Usage example: `winchecks check-windows-ram -w 85 -c 95`	
os.windows.check-system-process	Query running processes to find running processes that match the given arguments (pattern, name, both pattern and name. At least one must be given). Measures the running processes against configured thresholds and filters, returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -n name - Process executable name to check the process execution. • -p pattern - Pattern (sub string) to search for in the command that invoked the process. Produces valid results only if the user running the Agent owns the queried process has view permissions for the queried process. • -w warnover - Triggers a WARNING status if the query returns more processes than those specified by the argument. • -W warnunder - Triggers a WARNING status if the query returns fewer processes than those specified by the argument. • -c critover - Triggers a CRITICAL event if the query returns more processes than those specified by the argument.	Check Process OK: OK Found 1 matching running processes named explorer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		• -C critunder - Triggers a CRITICAL event if the query returns fewer processes than those specified by the argument. Usage example: `winchecks check-windows-processes -n explorer`	
os.windows.check-directory	Verifies whether a Windows directory exists.	Usage: -d --directory Path to the relevant directory; use '\' for separation. Usage example: `winchecks check-windows-directory -d dir_path`	Check Directory OK: The directory 'C:/Users/Public' exists
os.windows.check-pagefile	Collects the *`Pagefile`* usage and compares it against the WARNING and CRITICAL thresholds.	Usage: • -w warning - Triggers a WARNING event if the *`Pagefile`* usage is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the *`Pagefile`* usage is above the CRITICAL parameter value specified in the check parameter. Usage example: `winchecks check-windows-pagefile -w 75 -c 85`	Check Windows Page File OK: Page file usage at 31.63%

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
os.windows.check-free-physical-memory	Measures the free physical memory against configured thresholds and returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -w warning - Triggers a WARNING event if the free physical memory is under the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the free physical memory is under the CRITICAL parameter value specified in the check parameter. Usage example: `winchecks check-windows-free-physical-memory -w 10 -c 5`	Free Physical Memory OK: The Free Physical Memory is 20.25%
os.windows.check-free-virtual-memory	Measures the free virtual memory against configured thresholds and returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -w warning - Triggers a WARNING event if the free virtual memory is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the free virtual memory is above the CRITICAL parameter value specified in the check parameter. Usage example: `winchecks check-windows-free-virtual-`	Free Virtual Memory OK: The Free Virtual Memory is 25.66%

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		`memory -w 10 -c 5`	
os.windows.check-process-cpu	Processes CPU usage against configured thresholds and returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -p processname - Process name to collect CPU usage. • -w warning - Triggers a WARNING event if the CPU usage is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the CPU usage is above the CRITICAL parameter value specified in the check parameter. Usage example: `winchecks check-windows-process-cpu-p acc -c 95 -w 85`	Check Process CPU OK: Process CPU usage is 0.0000%
os.windows.check-process-memory	Processes memory usage against thresholds and returns a CRITICAL \WARNING\OK event according to the thresholds given in the accompanying parameters.	Usage: • -p processname - Process name to collect memory usage. • -w warning - Triggers a WARNING event if the process memory usage is above the WARNING parameter value specified in the check parameter. • -c critical - Triggers a CRITICAL event if the process	Check Process Memory OK: Process Memory usage is 0.0149%

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
		memory usage is above the CRITICAL parameter value specified in the check parameter. Usage example: `winchecks check-windows-process-memory-p acc -c 95 -w 85`	
os.windows.check-user-account	Takes the list of user names as an input and verifies whether the user account is active. Returns a CRITICAL\WARNING \OK value.	winchecks check-windows-user-disabled (options) -u : Comma separated List of User Name Usage example: `winchecks check-windows-user-disabled- u Administrator,Guest`	User Name and Status
util.windows.check-windows-service	Verifes whether a Windows Service is present and returns the actual service state and service type reported by the Windows Service Control Manager (`sc query`), for example: `Service: 'X', Status: 'RUNNING', Startup Type: 'WIN32_OWN_PROCESS'` Possible values: • RUNNING • STOPPED • START_PENDING • STOP_PENDING • CONTINUE_PENDING	Usage: • -s, service - The service to check if running (triggers an event with default severity of WARNING if the service is not running). • -e, non_running_services_event_severity - Sets the severity level for reporting events when the service is not running. The values can be either CRITICAL or WARNING (Default: WARNING). • -G, skip_events_for_nonexistent - Skips the event	`Service: '[ ]', Status: '[RUNNING/STOPPED/etc.]', Startup Type: '[WIN32_OWN_PROCESS/etc.]'`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events policy (continued)

Check	Description	Usage and Example	Output
	• PAUSE_PENDING • PAUSED	creation if the given service does not exist on the system when the parameter is active and value is true. Usage example: `winchecks check-windows-service -s W3SVC -G false`	

Windows metric monitoring checks

Windows OS Metrics policy

Check	Description	Usage and Example	Output
os.windows.check-processor-queue-length	Measures the processor queue length.	Usage: -s scheme - Replaces output's hostname + process with the given value (example: hostname.process) Usage example: `command: winchecks metric-windows-processor-queue-length --scheme hostname.proc`	win2019-dc-64bit.cpu.queuelength 0.00 1645371109
os.windows.check-system-cpu-load	Collects average CPU load per second.	Usage: -s scheme - Replaces output's hostname + process with the given value (example: hostname.process) Usage example: `command: winchecks metric-windows-cpu-load -scheme hostname.proc`	win2019-dc-64bit.cpu.loadavgsec 15.07 1645371561

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Metrics policy (continued)

Check	Description	Usage and Example	Output
os.windows.check-system-cpu	Collects the CPU core metric.	Usage: -s , scheme Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command : winchecks metric-windows-cpu -scheme hostname.proc`	win2019-dc-64bit.cpu.cpu0.cores 2 1645371681
os.windows.check-system-disk-usage	Collects the following disk usage metrics usage: • total in GB • usage in GB • avail in GB • used percentage	Usage: • -i , ignore_mnt: Comma separated list of mount points to ignore (:C) • -I, include_mnt: Comma separated list of mount points to include. • —scheme, scheme: Replaces output's hostname+process with the given value (example: hostname.process). Usage example: `command : winchecks metric-windows-disk-usage-scheme hostname.proc`	win2019-dc-64bit.disk_usage.disk_C.total(GB) 99.40 1645371774 win2019-dc-64bit.disk_usage.disk_C.used(GB) 50.72 1645371774 win2019-dc-64bit.disk_usage.disk_C.avail(GB) 48.68 1645371774 win2019-dc-64bit.disk_usage.disk_C.used_perc 51.02 1645371774
os.windows.check-system-memory-percent	Collects RAM percentage usage, Free Physical Memory percentage and Free Virtual Memory percentage.	Usage: -s, scheme - Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command : winchecks metric-`	win2019-dc-64bit.mem.free_physical_percentag 13.30 1645371856 win2019-dc-64bit.mem.free_virtual_percentage 13.93 1645371856 win2019-dc-64bit.ram.usage_percentage 86.07 1645371856

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Metrics policy (continued)

Check	Description	Usage and Example	Output
		`windows-disk-usage-scheme hostname.proc`	
os.windows.check-system-network	Collects the following active network adapter metrics: • Total bytes per sec • Packets/sec • Packets Received per sec • Packets Sent per sec • Current Bandwidth • Bytes Received per sec • Packets Received Unicast per sec • Packets Received Non-Unicast per sec • Packets Received Discarded • Packets ReceivedErrors • Packets Received Unknown • Bytes sent per sec • Packets sent unicast per sec • Packets sent non-unicast per sec • Packets outbound discarded • Packets outbound errors • Output queue length • Offloaded connections • TCP Active RSC Connections • TCP RSC Coalesced Packets per sec	Usage: -s scheme: Replaces output's hostname + process with the given value (example: hostname.process) Usage name: `command: winchecks metric-windows-network --scheme hostname.proc`	win2019-dc-64bit.system.network.Network_Inter name><metric value>Bytes_Total/sec 98742.67 1645372042 For example: win2019-dc-64bit.system.network.Network_Inter sec 98742.67 1645372042

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Metrics policy (continued)

Check	Description	Usage and Example	Output
	• TCP RSC Exceptions per sec • TCP RSC Average Packet Size		
os.windows.check-system-uptime	Collects system uptime.	Usage: -s, scheme - Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command : winchecks metric-windows-uptime --scheme hostname.proc`	win2019-dc-64bit.system.uptime(sec) 4614142.06 1645372124
os.windows.check-system-disk	Collects the following disk metrics: • AvgDiskSecPerRead • AvgDiskSecPerWrite • DiskReadBytesPerSec • DiskWriteBytesPerSec	Usage: • -i, ignore_mnt - Comma separated list of mount points to ignore (:C) • -I, include_mnt - Comma separated list of mount points to include. • —scheme - Replaces output's hostname+process with the given value (example: hostname.process). Usage example: `command : winchecks metric-windows-disk`	win2019-dc-64bit.disk._total.AvgDisksec/Read 0.000000 1645372198 win2019-dc-64bit.disk._total.AvgDisksec/Write 0.000608 1645372198 win2019-dc-64bit.disk._total.DiskReadBytes/sec 0.000000 1645372198 win2019-dc-64bit.disk._total.DiskWriteBytes/sec 34941.692255 1645372198 win2019-dc-64bit.disk.C.AvgDisksec/Read 0.000000 1645372200 win2019-dc-64bit.disk.C.AvgDisksec/Write 0.000000 1645372200

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Metrics policy (continued)

Check	Description	Usage and Example	Output
			win2019-dc-64bit.disk.C.DiskReadBytes/sec 0.000000 1645372200 win2019-dc-64bit.disk.C.DiskWriteBytes/sec 0.000000 1645372200
os.windows.check-system-memory	Collects the following disk metrics: • FreePhysicalMemory • TotalPhysicalMemory • FreeVirtualMemory • TotalVirtualMemorySize • AvailableMemory • TotalVisibleMemorySi;	Usage: -s, scheme - Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command: winchecks metric-windows-memory --scheme hostname.proc`	win2019-dc-64bit.mem.free_physical(KB) 1175440.00 1645372274 win2019-dc-64bit.mem.total_physical(KB) 8588898304.00 1645372274 win2019-dc-64bit.mem.free_virtual(KB) 1747636.00 1645372274 win2019-dc-64bit.mem.total_virtual(KB) 12263156.00 1645372274 win2019-dc-64bit.mem.available(KB) 1202032640.00 1645372274 win2019-dc-64bit.mem.total_visible(KB) 8387596.00 1645372274
os.windows.check-process-status	Collects windows process status with CPU and memory data used by the process.	Usage: • -n, process - Process name to collect status metric. • —scheme - Replaces output's hostname+process	win2019-dc-64bit.Process.Status 67 1645372421 win2019-dc-64bit.Process.CpuPercent 0 1645372421

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Metrics policy (continued)

Check	Description	Usage and Example	Output
		with the given value (example: hostname.process).	win2019-dc-64bit.Process.Memory(KB) 1226444 1645372421
os.windows.metrics-process-status	Retrieves the number of running instances, the percentage of CPU utilization, and the memory usage (in kilobytes) of the specified Windows process.	Usage: • -n, process - Process name to collect status metric. • —scheme - Replaces output's hostname+process with the given value (example: hostname.process). Usage example: `command: winchecks metric-windows-process-status -n Svchost -s hostname.proc`	WIN-R493MKFE75G.Process.Status 1 1625478491 WIN-R493MKFE75G.Process.CpuPercent 0 1625478491 WIN-R493MKFE75G.Process.MemoryKB 276 162547849

Windows OS event checks - Extended

Runs Windows extended checks on operational Windows servers. To run this policy, activate one of the checks and provide a CI filter on the policy's Monitored CIs tab to run these checks on selected CIs.

Windows OS Events - Extended policy

Check	Description	Usage and Example	Output
os.windows.check-processor-queue-length	Measures the processor queue length.	Usage: -s scheme - Replaces output's hostname + process with the given value (example: hostname.process) Usage example: `command: winchecks metric-windows-processor-queue-length`	win2019-dc-64bit.cpu.queuelength 0.00 1645371109

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Check	Description	Usage and Example	Output
		`--scheme hostname.proc`	
os.windows.check-system-cpu	Collects the CPU core metric.	Usage: -s , scheme Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command: winchecks metric-windows-cpu -scheme hostname.proc`	win2019-dc-64bit.cpu.cpu0.cores 2 1645371681
os.windows.check-system-cpu-load	Collects average CPU load per second.	Usage: -s scheme - Replaces output's hostname + process with the given value (example: hostname.process) Usage example: `command: winchecks metric-windows-cpu-load -scheme hostname.proc`	win2019-dc-64bit.cpu.loadavgsec 15.07 1645371561
os.windows.check-system-disk-usage	Collects the following disk usage metrics usage: • total in GB • usage in GB • avail in GB • used percentage	Usage: • -i , ignore_mnt: Comma separated list of mount points to ignore (:C) • -I, include_mnt: Comma separated list of mount points to include. • —scheme, scheme: Replaces output's hostname+process with the given value (example: hostname.process).	win2019-dc-64bit.disk_usage.disk_C.total(GB) 99.40 1645371774 win2019-dc-64bit.disk_usage.disk_C.used(GB) 50.72 1645371774 win2019-dc-64bit.disk_usage.disk_C.avail(GB) 48.68 1645371774 win2019-dc-64bit.disk_usage.disk_C.used_perc 51.02 1645371774

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Check	Description	Usage and Example	Output
		Usage example: `command : winchecks metric-windows-disk-usage-scheme hostname.proc`	
os.windows.check-system-memory-percent	Collects RAM percentage usage, Free Physical Memory percentage and Free Virtual Memory percentage.	Usage: -s, scheme - Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command : winchecks metric-windows-disk-usage-scheme hostname.proc`	win2019-dc-64bit.mem.free_physical_percentag 13.30 1645371856 win2019-dc-64bit.mem.free_virtual_percentage 13.93 1645371856 win2019-dc-64bit.ram.usage_percentage 86.07 1645371856
os.windows.check-system-network	Collects the following active network adapter metrics: • Total bytes per sec • Packets/sec • Packets Received per sec • Packets Sent per sec • Current Bandwidth • Bytes Received per sec • Packets Received Unicast per sec • Packets Received Non-Unicast per sec • Packets Received Discarded • Packets ReceivedErrors • Packets Received Unknown • Bytes sent per sec	Usage: -s scheme: Replaces output's hostname + process with the given value (example: hostname.process) Usage name: `command : winchecks metric-windows-network --scheme hostname.proc`	win2019-dc-64bit.system.network.Network_Inter name><metric value>Bytes_Total/sec 98742.67 1645372042 For example: win2019-dc-64bit.system.network.Network_Inter sec 98742.67 1645372042

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Check	Description	Usage and Example	Output
	• Packets sent unicast per sec • Packets sent non-unicast per sec • Packets outbound discarded • Packets outbound errors • Output queue length • Offloaded connections • TCP Active RSC Connections • TCP RSC Coalesced Packets per sec • TCP RSC Exceptions per sec • TCP RSC Average Packet Size		
os.windows.check-system-uptime	Collects system uptime.	Usage: -s, scheme - Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command: winchecks metric-windows-uptime --scheme hostname.proc`	win2019-dc-64bit.system.uptime(sec) 4614142.06 1645372124
os.windows.check-system-disk	Collects the following disk metrics: • AvgDiskSecPerRead • AvgDiskSecPerWrite • DiskReadBytesPerSec • DiskWriteBytesPerSec	Usage: • -i, ignore_mnt - Comma separated list of mount points to ignore (:C) • -I, include_mnt - Comma separated list of mount points to include.	win2019-dc-64bit.disk._total.AvgDisksec/Read 0.000000 1645372198 win2019-dc-64bit.disk._total.AvgDisksec/Write 0.000608 1645372198 win2019-dc-64bit.disk._total.DiskReadBytes/

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Check	Description	Usage and Example	Output
		• —scheme, scheme - Replaces output's hostname+process with the given value (example: hostname.process). Usage example: `command : winchecks metric-windows-disk`	sec 0.000000 1645372198 win2019-dc-64bit.disk._total.DiskWriteBytes/sec 34941.692255 1645372198 win2019-dc-64bit.disk.C.AvgDisksec/Read 0.000000 1645372200 win2019-dc-64bit.disk.C.AvgDisksec/Write 0.000000 1645372200 win2019-dc-64bit.disk.C.DiskReadBytes/sec 0.000000 1645372200 win2019-dc-64bit.disk.C.DiskWriteBytes/sec 0.000000 1645372200
os.windows.check-system-memory	Collects the following disk metrics: • FreePhysicalMemory • TotalPhysicalMemory • FreeVirtualMemory • TotalVirtualMemorySize • AvailableMemory • TotalVisibleMemorySi;	Usage: -s, scheme - Replaces output's hostname+process with the given value (example: hostname.process) Usage example: `command : winchecks metric-windows-memory --scheme hostname.proc`	win2019-dc-64bit.mem.free_physical(KB) 1175440.00 1645372274 win2019-dc-64bit.mem.total_physical(KB) 8588898304.00 1645372274 win2019-dc-64bit.mem.free_virtual(KB) 1747636.00 1645372274 win2019-dc-64bit.mem.total_virtual(KB) 12263156.00 1645372274 win2019-dc-64bit.mem.available(KB) 1202032640.00 1645372274

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows OS Events - Extended policy (continued)

Check	Description	Usage and Example	Output
			win2019-dc-64bit.mem.total_visible(KB) 8387596.00 1645372274
os.windows.check-process-status	Collects windows process status with CPU and memory data used by the process.	Usage: • -n, process - Process name to collect status metric. • —scheme, scheme - Replaces output's hostname+process with the given value (example: hostname.process).	win2019-dc-64bit.Process.Status 67 1645372421 win2019-dc-64bit.Process.CpuPercent 0 1645372421 win2019-dc-64bit.Process.Memory(KB) 1226444 1645372421

Windows metrics

The following table lists the metrics that are gathered as output from Windows checks. Entries indicated as **Featured metrics** are high-visibility metrics that are displayed in the Operator Workspace Metric tab after an alert is generated. These metrics provide the operator with additional information to help them further explore the specified issue.

Windows metrics

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
Process.Status			Windows process status including CPU and memory data.
Process.CpuPercent			Percentage of CPU used by the process.
Process.Memory(KB)		KB	Memory used by the given process.
cpu.queuelength (featured metric)			Processor queue length.
cpu.loadavgsec (featured metric)			Average CPU load per second.
disk.AvgDisksec/Read		seconds	Average time of a read from the disk.
disk.AvgDisksec/Write		seconds	Average time of a write to the disk.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
disk.DiskReadBytes/sec		byte/sec	Rate at which bytes are transferred from the disk during read operations.
disk.DiskWriteBytes/sec		byte/sec	Rate at which bytes are transferred to the disk during write operations.
disk_usage.avail(GB)		GB	Disk available.
disk_usage.used(GB)		GB	Disk used.
disk_usage.total(GB)		GB	Disk total.
disk_usage.used_percentage (featured metric)			Disk used in percentage.
mem.free_physical_percentage			Percent of free physical memory.
mem.free_virtual_percentage			Percent of free virtual memory.
ram.usage_percentage (featured metric)			Percent of RAM used.
mem.free_physical(KB)		KB	Amount of free physical memory.
mem.total_physical(KB)		KB	Total physical memory.
mem.free_virtual(KB)		KB	Amount of free virtual memory.
mem.total_virtual(KB)		KB	Total virtual memory.
mem.available(KB)		KB	Total amount of virtual memory available.
mem.total_visible(KB)		KB	Total amount of physical memory accessible to the operating system.
system.uptime(Sec)		seconds	Duration that the system has been operational.
system.network.Bytes_Received/sec (featured metric)		byte/sec	Bytes received per second for active network adapters.
system.network.Bytes_Sent/sec (featured metric)		byte/sec	Bytes sent per second for active network adapters.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
system.network.Bytes_Total/sec (featured metric)		byte/sec	Sum of the bytes received and sent per second.
system.network.Current_Bandwidth		bit/sec	Estimate of the network interface bandwidth.
system.network.Offloaded_Connections			Number of IPv4 and IPv6 TCP connections currently handled by the TCP chimney offload capable network adapter.
system.network.Output_Queue_Length			Length of the output packet queue.
system.network.Packets/sec			Rate that packets are sent and received on the network.
system.network.Packets_Outbound_Discarded			Number of outbound packets that were discarded.
system.network.Packets_Outbound_Errors			Number of outbound packets that could not be sent becuase of errors.
system.network.Packets_Received/sec			Rate that packets are received.
system.network.Packets_Received_Discarded			Number of inbound packets that were discarded.
system.network.Packets_Received_Errors			Number of inbound packets that could not be sent because of errors.
system.network.Packets_Received_Non-Unicast/sec			Rate that non-unicast (subnet broadcast or subnet multicast packets) are delivered to a higher-layer protocol.
system.network.Packets_Received_Unicast/sec			Rate that unicast (subnet) packets are delivered to a higher-layer protocol.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows metrics (continued)

Metric type	Resource (name of specific database, where relevant)	Units	Metric type description
system.network.Packets_Received_Unknown			Number of packets that were received, but discarded, because of an unknown or unsupported protocol.
system.network.Packets_Sent/sec			Rate that packets are sent.
system.network.Packets_Sent_Non-Unicast/sec			Rate that packets are requested for transmission to non-unicast (subnet broadcast or subnet multicast) addresses by higher-layer protocols. Includes packets that were discarded or not sent.
system.network.Packets_Sent_Unicast/sec			Rate that packets are requested to for transmission to subnet-unicast addresses by higher-layer protocols. Includes packets that were discarded or not sent.
system.network.TCP_Active_RSC_Connections			Number of IPv4 and IPv6 TCP connections currently receiving large packets from the RSC-capable network adapter.
system.network.TCP_RSC_Average_Packet_Size		byte	Average size of received packets across all TCP connections.
system.network.TCP_RSC_Coalesced_Packets/sec			The large-packet receive rate across all TCP connections.
system.network.TCP_RSC_Exceptions/sec			The RSC exception rate for received packets across all TCP connections.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Windows log monitoring default checks and policies

Agent Client Collector provides the following policy for Windows log monitoring.

Windows log monitoring policy

Type	Check	Description	Usage and Usage Example	Output
Event	os.windows.check-log	Enables monitoring log files on a Windows server.	Usage: • -c crit N: Critical level if pattern has a group. • -f -log-file FILE: Path to log file. • -o --warn-only: Warn instead of critical on match. • -q pattern PAT: Pattern to search for. To search for multiple patterns, separate each pattern with a pipe (|) and put inside quotes. For example: "SEVERE|404" • -w warn N: Warning level if pattern has a group.	Windows Log CRITICAL: Found 4 criticals, 0 warnings for pattern SEVERE|Exception|404|Errorin file C:\ProgramData\ServiceNow\agent-client-collector\log\acc.log

Linux log monitoring default checks and policies

Agent Client Collector provides the following policy for Linux log monitoring.

Type	Check	Description	Usage and Usage Example	Output
Event	util.check-logs	Enables monitoring log files owned by a regular user.	Usage: • -i --icase: Run a case insensitive match. • -c, --crit N: Critical level (if pattern has a group).	CheckLog CRITICAL: 0 warnings, 8 criticals for pattern SEVERE|Exception in log file /var/log/servicenow/agent-client-collector/acc.log

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Check	Description	Usage and Usage Example	Output
			• --encode-utf16u: Encode line with utf16 before matching. • -e, --encoding ENCODING-PAGE: Specific encoding page to read log file with. • -E, --exclude PAT: Pattern to exclude from matching. • -F, --filepattern FILE: Check a pattern of files, instead of one file. For REGEX, first test it on `https://rubular.com/` to get the expected outcomes and then pass it inside quotes as a parameter. For example, to get all *`.log`* extension files, pass `"(.)*\.log$"` as REGEX. • -f, --log-file FILE: Path to log file. • -l, --log-pattern PAT: Log format of each log entry: • -o, --warn-only Warn instead of critical on match. • -q, --pattern PAT Pattern to search	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Check	Description	Usage and Usage Example	Output
			for.To search for multiple patterns, separate each pattern with pipe(|) and put inside quotes (For example: "SEVERE|404"). • -r, --return: Return matched line. • -L, --return-length N: Matched line length. • -M, --return-error-limit N: Max number of returned matched lines(log entries). • -n, --name NAME Set state file dir automatically using name. • -s, --state_dir DIR Dir to keep state files under. • -w, --warn N: Warning level if pattern has a groupWarning level if pattern has a group. Usage example: `command: check-log.rb -c 2 -w 1 -q "SEVERE|Exception" -s /tmp/cache/check-log -f /var/log/servicenow/`	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Check	Description	Usage and Usage Example	Output
			`agent-client-collector/acc.log`	
Event	util.check-logs-sudo	Enables monitoring log files owned by a root user.	Usage: • -i --icase: Run a case insensitive match • -c, --crit N: Critical level (if pattern has a group) • --encode-utf16u: Encode line with utf16 before matching • -e, --encoding ENCODING-PAGE: Specific encoding page to read log file with. • -E, --exclude PAT Pattern to exclude from matching • -F, --filepattern FILE: Check a pattern of files, instead of one file. For REGEX, first test it on `https://rubular.com/` to get the expected outcomes and then pass it inside quotes as a parameter. For example, to get all *`.log`* extension files, pass `"(.)*\.log$"` as REGEX.	CheckLog CRITICAL: 0 warnings, 8 criticals for pattern SEVERE|Exception in log file /var/log/servicenow/agent-client-collector/acc.log

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Check	Description	Usage and Usage Example	Output
			• -f, --log-file FILE: Path to log file. • -l, --log-pattern PAT: Log format of each log entry: • -o, --warn-only Warn instead of critical on match • -q, --pattern PAT Pattern to search for.To search for multiple patterns, separate each pattern with pipe(|) and put inside quotes (for example: "SEVERE|404") • -r, --return: Return matched line. • -L, --return-length N: Matched line length. • -M, --return-error-limit N: Max number of returned matched lines(log entries). • -n, --name NAME: Set state file dir automatically using name. • -s, --state_dir DIR: Dir to keep state files under • -w, --warn N: Warning level if pattern has a	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Type	Check	Description	Usage and Usage Example	Output
			groupWarning level if pattern has a group. Usage example: `command: check-log.rb -c 2 -w 1 -q "SEVERE|Exception" -s /tmp/cache/check-log -f /var/log/servicenow/agent-client-collector/acc.log`	

Agent Client Collector checks – credential and permission requirements

Agent Client Collector checks require specific credentials and permissions to execute successfully. The following sections outline the application-level (DB grants, API roles) and OS-level access required for each check.

Oracle check

Learn about the credential and permission requirements for the Oracle checks `app.oracle.check-oracle-alive`, `app.oracle.metrics-oracle`, and `util.metrics-oracle-query`.

Oracle checks require two separate accounts with different permissions.

Oracle database account

The database user configured in ServiceNow® must have the following permissions:

- CREATE SESSION or CONNECT
- SELECT access to the Oracle views or tables queried by the specific check that is running

Checks and queries

Check	Queries
`check-oracle-alive`	`SELECT BANNER FROM v$version WHERE banner LIKE 'Oracle%'`
`metrics-oracle`	`SELECT TABLESPACE_NAME, sum(BYTES), sum(MAXBYTES) FROM sys.dba_data_files GROUP BY TABLESPACE_NAME`
`util.metrics-oracle-query`	Any custom query. The database account needs SELECT on the object targeted by that query.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector agent OS account

All three checks run *sqlplus* locally on the target host where the OS account runs the Agent Client Collector agent. This is a separate identity from the Oracle database account.

Agent Client Collector agent OS account needs the following access:

- Execute access to `$ORACLE_HOME/bin/sqlplus`
- Access to Oracle client libraries under `$ORACLE_HOME/lib`
- Access to Oracle configuration files if applicable, such as `tnsnames.ora` and `sqlnet.ora`

Depending on the Oracle installation, this means the Agent Client Collector agent OS account must be a member of `oinstall` or a database administrator.

Configuration Prerequisites:

- *sqlplus* installed and reachable at `$ORACLE_HOME/bin/sqlplus`
- Oracle listener reachable from the agent host through `--host`, `default localhost`, and `--port`, default 1521
- Valid Oracle database credentials

Troubleshooting

Error	Cause
Permission denied or command not found on *sqlplus*	The Agent Client Collector agent OS account lacks access to `$ORACLE_HOME` or *sqlplus*. This is not a database credential problem.
`ORA-xxxxx` errors	Oracle-side error that points to the database account or grants and not the OS account.
Custom query fails only on `util.metrics-oracle-query`	Re-check SELECT grants whenever `--query` is changed to target a new object.

MySQL check

Learn about the credential and permission requirements for the MySQL checks `app.mysql.check-mysql-alive`, `util.metrics-mysql-query`, `app.mysql.check-mysql-threads`, `app.mysql.metrics-mysql`, `app.mysql.metrics-mysql-processes`, and `util.check-mysql-query`

MySQL checks require two separate accounts with different permissions:

MySQL database account

The database user configured in the ServiceNow® credential must have the following permissions:

- Basic login privilege on the target host and port, or socket
- SELECT access to whatever schema or table is targeted by a custom query
- PROCESS privilege for checks that read server-wide thread or connection state

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Checks and additional requirements

Check	Additional requirement
`check-mysql-alive`	Login privilege only
`check-mysql-query-result-count`, `metrics-mysql-query-result-count`	SELECT access on the object targeted by the configured `--query`
`check-mysql-threads`, `metrics-mysql-processes`	PROCESS privilege. Without it, the account only sees its own connection or thread, not the server-wide count, and the check silently under-reports instead of erroring.
`metrics-mysql-graphite`	Login privilege only

Agent Client Collector Agent OS account

There is no special OS permission required. The Agent Client Collector agent OS account only needs additional access if:

- A UNIX socket is used instead of TCP through `--socket`. The agent needs read and execute access to the socket file and its parent directory.
- An `ini` file is used to source connection settings through `--ini` or `--ini_section`, `my.cnf-style`. The agent needs read access to this `ini` file.

Depending on the deployment, this typically means group membership on the socket's owning group, or read ACLs on the `my.cnf` file. It doesn't usually require root or DBA-level access.

Configuration Prerequisites:

- MySQL or MariaDB TCP endpoint that is reachable from the agent host through `--hostname` and `--port`, or through `--socket`
- Database credentials are valid and not blocked by host-based grant restrictions, such as `user@'agent-host'` compared to `user@'%'`
- For thread or process metrics, confirm PROCESS is granted by running `SHOW GRANTS FOR 'user'@'host';`

PostgreSQL check

Credential and permission requirements for the PostgreSQL checks that monitor database size, statistics, connections, and locks.

PostgreSQL database account

The database user configured in the ServiceNow credential must meet the following requirements.

- Required for every check in this group: the account must be allowed to authenticate and connect to the specific database named in `--database`. It needs CONNECT privilege on that database, and `pg_hba.conf` on the server must permit login from the agent host for that user and database combination.
- Required only for `postgresql.check-connections` and `postgresql.metric-active-connections`: the `pg_monitor` role on `Postgres 10+`, or superuser. Both checks query `pg_stat_activity` across all databases and sessions.

Configuration Prerequisites:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Postgres TCP endpoint must be reachable from the agent host.

- `pg_hba.conf` must let the agent host's source IP to authenticate for this user and database.
- The connection `sslmode=disable` is hard-coded. If the server enforces SSL/TLS only through `hostssl` entries, the check fails with an SSL-negotiation error that can look like an authentication failure.
- Confirm `pg_monitor` is granted if you're using connection or stats metrics.

HAProxy check

Credential and permission requirements for the HAProxy checks `check-haproxy` and `metrics-check-haproxy`.

HAProxy stats credential

The HAProxy stats credential are required only when the HAProxy stats page has authentication enabled.

Agent Client Collector Agent OS account

When using an HTTP or HTTPS stats page an Agent Client Collector account is not required. OS-level access is only required if using the UNIX stats socket, with `--stats` pointed at a socket path instead of an HTTP hostname.

In that case, the ACC agent OS account needs read/write access to the socket, which is often owned by `haproxy:haproxy` with restrictive permissions.

Configuration prerequisites:

- The HAProxy stats page or socket must be enabled in the HAProxy configuration through `stats socket` or `stats enable`.
- If a `stats-page` credential is configured, it must match `stats auth user:pass` in the HAProxy configuration.

Troubleshooting:

A permission denied error when connecting to the stats socket is an OS-level group or ACL issue on the ACC agent account. It is not a credential problem.

File or directory check

OS-level permission requirements for the file or directory check.

Agent Client Collector Agent OS account

The Agent Client Collector agent OS account is only required if the target file or directory is owned by another service account. The Agent Client Collector agent OS account needs the following permissions:

- Read permission, and for directories, execute or traverse permission, on the target `--directory`, `--dirpath`, or `--filepath`.
- For `check-file-response-time`, read access to the contents of the file, not just metadata, because it measures read latency.

A file or directory owned by another application, such as its own log or data directory, is frequently not world-readable. In that case, the Agent Client Collector agent account needs group membership or an ACL grant to read it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration Prerequisites:

- Confirm the target path exists and is spelled or cased correctly. Paths are case-sensitive on Linux.
- Confirm read and execute permission on the full path, including every parent directory, not just the final file or directory. A missing execute bit on a parent directory blocks traversal even if the file itself is readable.

Troubleshooting:

- A file not found result for a file that exists means a missing execute bit on a parent directory rather than a problem with the file itself.
- For files owned by another service account, add the Agent Client Collector agent OS account to the owning group, or grant an ACL, rather than loosening the file's world permissions.

Azure metrics script

When creating a new policy for Azure cloud metrics, create a .json script to determine the metrics to be monitored. The format for the script appears below, followed by a table explaining the script contents.

```
{
  "namespace": "Microsoft.Network/loadBalancers",
  "metrics": [
          "AllocatedSnatPorts",
          "ByteCount",
          "DipAvailability",
          "PacketCount",
          "SnatConnectionCount",
          "SYNCount",
          "UsedSnatPorts",
          "VipAvailability"
  ],
  "time_grain": "PT1M"
}
```

Resources and their metrics are retrievable from the Supported metrics for Azure resources page.

Header	Description
namespace	The name of the resource being monitored.
metrics	The metrics being retrieved from the monitored resource.
time_grain	The granularity of how often metrics are to be collected. Possible values are: • PT1M: Metric values are collected every minute. • PT5M: Metric values are collected every 5 minutes. • PT15M: Metric values are collected every 15 minutes. • PT30M: Metric values are collected every 30 minutes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Header	Description
	• PT1H: Metric values are collected every hour. • PT1D: Metric values are collected daily.

Windows event log filter parameters

The configurable values on the Check Parameters tab of the os.windows.check-event-log check.

Check Parameters tab values

Parameter	Type	Description
provider_name	String	Name of the provider that generated the event. **Note:** If you do not specify a *log_file* value together with the *provider_name*, the system searches all available log files, which increases the time it takes to receive results.
log_file	String	The name of the Windows event log file from which you retrieve events. Possible values are: • Application • System **Note:** If you do not specify a *provider_name* value together with the *log_file*, the system searches all events from the log file, which increases the number of retrieved events.
id	Integer	The numerical id of the event. Possible values are 0-65535.
warning	Integer	Any value above the specified parameter generates a Warning event.
event_level	String	The severity level of the event. Possible values: • Critical • Error

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Check Parameters tab values (continued)

Parameter	Type	Description
		• Warning • Information • Verbose
regex_pattern	String	The regex pattern to be used in searching the event logs. The value must be enclosed in double quotation marks. For example, "error".
duration_hour	Integer	The time period for which you want to retrieve events from the Windows event log. Value is specified in hours; fractions of hours are specified with decimals.
critical	Integer	Any value equal to or above the specified parameter generates a Critical event.

Parameters of MID Server distributed cluster

The following MID Server parameters support the MID Server Distributed cluster type.

You can access these parameters in the **Configuration Parameters** related list of a MID Server. To add and set any of these parameter, see Add a MID Server parameter.

Parameter	Description
mid.distributed.cluster.striped.threadpool.size	Size of striped threadpool in MID Server distributed cluster. • **Type**: integer • **Default value**: Number of cores on the virtual machine or server.
mid.distributed.cluster.log.level	Logging level for MID Server distributed cluster, which can help troubleshooting. However, logging impacts the performance of Metric Intelligence. An empty value designates no logging. • **Type**: integer • **Default value**: none • **Other possible values**: info, debug, trace

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameter	Description

mid.distributed.cluster.global.timeout	Global timeout in seconds for MID Server distributed cluster.
	
	• **Type**: integer
	• **Default value**: 180
	• **Other possible values**:
	  ○ **Minimum value**: 60 (not enforced)
	  ○ **Maximum value**: 300 (not enforced)

mid.distributed.cluster.node.max.off.heap.memory	The amount of off heap memory used by Metric Intelligence to process data series per minute. The default 4 GB of space is the amount of off heap memory needed for processing 800,000 data series per minute. Adjust this value proportionally to the amount of actual data series. For example, to process 400,000 data series per minute, you can change the value to 2048. • **Type**: integer • **Default value**: 4096

Scheduled jobs included with MID Server distributed cluster

The following MID Server scheduled job is included with the MID Server Distributed cluster type.

Name	Description
Validate MID distributed cluster	Validates status of the cluster node for each MID Server in the cluster. For a MID Server that is down and its status is not stopped, updates the node status to stopped. Also, validates overall status of the cluster. During this validation, the job stops the cluster if there are multiple sub-clusters in a single cluster.

Components installed with Metric Intelligence

Several types of components are installed with activation of the Metric Intelligence (com.snc.sa.metric) plugin, including tables, scheduled jobs, and properties.

Note: The Application Files table lists the components that are installed with this application. For instructions on how to access this table, see Find components installed with an application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Scheduled jobs installed

Name	Description
Metric Intelligence — Metric Learner job	Runs daily and builds the statistical models used for anomaly detection.
Metric Intelligence — Purge old metric schemas and anomaly scores	Cleans up the *sa_metric_anomaly_score* and the *sa_metric_schema_definition* tables.
Metric Intelligence — Sync metric schema to mid	Synchronizing data from the instance to the MID Server in the metric schema table *sa_metric_schema_definition*.
Metric Intelligence — Sync tables with mid	Synchronizing data from the instance to the MID Server in tables.
Metric Intelligence — Table cleanup	Cleaning up tables related to various deletion operations.
Metric Intelligence — Observation Interval Learner job	Calculates the observation interval using collected data points for all data series.
Metric Intelligence — Process Stale Anomaly Score – Daily	Daily job that cleans stale anomaly scores, for a specified window of time, in the Metric Anomaly Score [sa_metric_anomaly_score] table. Uses the properties *sa_metric.anomaly_score_stale_interval_in_days* and *sa_metric.anomaly_score_next_stale_interval_in_days* to calculate the window of time for cleaning.
Metric Intelligence — Sync CI Anomaly Scores	Keeps anomaly scores up to date with any deletion of metrics.
Metric Intelligence — Metric configuration job	Synchronizes metric configurations from the instance to all Metric Intelligence MID Servers.
Metric Intelligence — Local level kalman filter model learner job	Builds local level kalman filter models on which to run change detection.
Metric Intelligence — Anomaly Model Test	On-demand scheduled job which runs when user clicks **Run Anomaly Test**.
Metric Intelligence — Validate MID Distributed cluster	**1.** Checks the status of the cluster node. If status is **Started** and the status of the Metric Intelligence extension is not any of **Started**/**Starting**/**Warning** [warning with specific errors], the job takes action. The job changes the status of the cluster node to **Stopped**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	Description
	This happens when MID Server stops unexpectedly without notifying the instance that it is stopping. **2.** Checks the status of all cluster nodes on all MID Servers in the cluster. ○ If the status of all nodes is **Stopped** and the status of the cluster itself is **Started**, then the job changes the status of the cluster to **Stopped**. ○ If the cluster status is **Stopped** but MID Servers have a running cluster node, the job changes the cluster status to **Started**. **3.** Checks the number of Cluster Leaders: If two are detected, then the job stops the entire cluster. This happens when MID Servers are not able to connect with other members of the cluster and form multiple sub-clusters.
Event Management - process metric binding events [1000-1020]	Process Metric Intelligence metric binding events separately from Event Management events (bucket range 1000-1020).
Event Management - process metric binding events [1020-1040]	Process Metric Intelligence metric binding events separately from Event Management events (bucket range 1020-1040).
Event Management - process metric binding events [1040-1060]	Process Metric Intelligence metric binding events separately from Event Management events (bucket range 1040-1060).
Event Management - process metric binding events [1060-1080]	Process Metric Intelligence metric binding events separately from Event Management events (bucket range 1060-1080).
Event Management - process metric binding events [1080-1100]	Process Metric Intelligence metric binding events separately from Event Management events (bucket range 1080-1100).
Metric Intelligence - Stationary nonparametric model learner job	Learns the stationary nonparametric model.
Metric Intelligence - Combined Model Learner Job	Constructs information for real time anomaly detection on a subset of collected series. Calculates data frequency and builds statistical models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables installed

Table	Description
Metric Time Series Model [sa_time_series]	Statistical models built for metric data.
Alert Anomaly [em_alert_anomaly]	Anomaly alerts that were created from anomaly events that were sent by Metric Intelligence so users can review. Anomaly alerts are kept separately from regular system alerts.
Metric Anomaly Score [sa_metric_anomaly_score]	Stores anomaly scores for the metric.
Metric To CI Mapping Deleted SysIds [sa_metric_map_deleted]	SysIds of deleted metric-to-CI mapping records.
Metric To CI Mapping [sa_metric_map]	Mappings currently in effect of metric types to CIs.
SA Metric Type Registration [sa_metric_registration]	Details about metric type registration.
Metric Type Registration Deleted SysIds [sa_metric_registration_deleted]	SysIds of records that were deleted from the SA Metric Type Registration [sa_metric_registration] table.
Metric Schema Definition [sa_metric_schema_definition]	Map of metrics being received based on CI class. It is used to optimize the metric data payload being sent from the MID Server to the instance.
Metric Type [sa_metric_type]	Metric type source.
Monitoring System Metric Type Deleted SysIds [sa_source_metric_type_deleted]	SysIds of deleted monitoring system metrics.
Monitoring System Metric Type [sa_source_metric_type]	Metric types per CI class, active/inactive status, and metric source.
Metric Settings [sa_metric_config_setting]	Default configuration settings for Metric Intelligence MID Servers.
Metric Configuration Rules [sa_metric_config_rule]	Rules to override configuration settings for MID Servers.
Metric overridden configurations	Maps overridden configurations to custom defined rules.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sa_metric_config_rule_setting]	
Metric Staged Configurations [sa_metric_config_staging]	Saved encoded configurations for the MID Servers.
Metric learned parameters [sa_metric_learned_param]	Parameters learned for CI/metric combinations, that will be synced to the MID Servers.
Metric To CI Mapping Deleted Details [sa_metric_map_deleted_detail]	Details about deleted entries from metric to CI mapping table. Includes CI, metric, and SysId of the deleted entry in the Metric To CI Mappings [sa_metric_map] table.
Track unique CIs created in metric map [sa_metric_ci_tracker]	All unique CIs for which metric map entry was created.
Metric Class [sa_metric_class]	Metric classes that can be associated with specific metrics to override bounds and widths learned values.
Metric Change Score [sa_metric_change_score]	Change events information.
Anomaly Test Rules [sa_metric_anomaly_test_rule]	Anomaly test rules used for anomaly model testing.
Anomaly test results [sa_metric_anomaly_test_result]	Results from anomaly model testing, such as bounds and anomaly scores.
Metric anomaly test config customized parameters [sa_metric_anomaly_test_config_custom_param]	Customized parameters for selected series in anomaly model testing.
Metric Intelligence Properties [sa_metric_property]	Metric Intelligence properties by domain.
Track host name with metric CI [sa_metric_ci_host]	CIs that are bind with node name from the binding event.
Functions supported by metric expression language [sa_metric_expr_function]	Functions supported by metric expression language.
Argument types for functions supported by metric expression language [sa_metric_expr_func_arg]	Argument types for functions supported by metric expression language.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
Track CI Types to metric entries valid for different sources [sa_metric_ci_type_map]	Track CI Types to metric entries that are valid for different sources.
CI Type To Resource Class [sa_ci_type_to_resource_class]	Contains mappings of CI type to resource class. Used for resource binding during resource lookups given a CI and a `resource_path` in the event. The [sa_ci_type_to_resource_class] table contains some default mapping records, and the evt_mgmt_admin role is required to manage this table.
CI Resource [ci_resource]	Parent table used for resource binding.
CI Resource Hardware [ci_resource_hardware]	An extension of the CI Resource [ci_resource] table, used for resource binding.
CI Resource Application [ci_resource_appl]	An extension of the CI Resource [ci_resource] table, used for resource binding.
CI Resource Service [ci_resource_service]	An extension of the CI Resource [ci_resource] table, used for resource binding.
CI Resource VM [ci_resource_vm_object]	An extension of the CI Resource [ci_resource] table, used for resource binding.
CI Resource Database [ci_resource_database]	An extension of the CI Resource [ci_resource] table, used for resource binding.
Series Event History [sa_series_event_history]	Information about CIs start and end maintenance mode. Metric Intelligence stores historical information about the times CIs enter and exit maintenance mode if the system property sa.model_learner.maint_event_record_history_enabled is set to true (default). And if the system property sa_metric.maint_exclusion is also set to true (default), metrics from CIs that are in maintenance mode are excluded from model learning.
Metric Dashboard Metadata [sa_metric_dashboard_metadata]	Chart configurations per CI class, for the metrics that are displayed for a CI associated with an alert, in Agent Workspace.
Metric Config Anomaly Detection Staging [sa_metric_config_anomaly_detection_staging]	Internal table that is used to help calculate the action level for each entry in the Metric to CI Mapping [sa_metric_map] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Properties installed

Note: To open the System Property [sys_properties] table, enter `sys_properties.list` in the navigation filter. To open the Metric Settings [sa_metric_config_setting] table, enter `sa_metric_config_setting.list` in the navigation filter.

Property	Usage
sa.metric.map.with.ci.expiration.sec	Length of time (in seconds) that records for mapping of CIs remains in effect if the mapping to CI is found. • **Type**: numeric • **Default value**: 432000 • **Location**: System Property [sys_properties] table • **Learn more**: Metric Intelligence
sa.metric.map.without.ci.expiration.sec	Length of time (in seconds) that records for mapping of CIs remains in effect if the mapping to CI is not found. • **Type**: numeric • **Default value**: 86400 • **Location**: System Property [sys_properties] table • **Learn more**: Metric Intelligence
sa_metric.anomaly_score_stale_interval_in_days	The starting day of the window of time that the *Opera Metrics - Process Stale Anomaly Score* scheduled job uses when checking staleness status of score. The window of time spans from today, until the past nu specified in this property. • **Type**: numeric • **Default value**: 3 • **Location**: System Property [sys_properties] table
sa_metric.anomaly_score_next_stale_interval_in_days	Number of days by which to extend the window of time *Operational Metrics - Process Stale A Score - Daily* scheduled job. Applies if the job runs, but fails to return any data using the window of time specified by the *sa_metric.anomaly_score_stale_interva* property. • **Type**: numeric • **Default value**: 7 • **Location**: System Property [sys_properties] table
max_records_per_rest_request	Maximum number of records to download from the inst MID Server in each synchronization cycle.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Usage

	This property is for internal use.
	
	• **Type**: integer
	• **Default value**: 10000
	• **Other possible values**:
	  ○ **Minimum value**: 1000
	  ○ **Maximum value**: 20000
	• **Location**: Metric Settings [sa_metric_config_setting]

normalized_metric_max_queue_size	Size of metrics queue. Anomaly detector reads through
	and determines for each metric whether it is anomalou
	
	This property is for internal use.
	
	• **Type**: integer
	• **Default value**: 3000000
	• **Other possible values**:
	  ○ **Minimum value**: 1000000
	  ○ **Maximum value**: 5000000
	• **Location**: Metric Settings [sa_metric_config_setting]

normalized_metric_queue_consumer_count	Number of threads that concurrently perform anomaly
	
	This property is for internal use.
	
	• **Type**: integer
	• **Default value**: 12
	• **Other possible values**:
	  ○ **Minimum value**: 8
	  ○ **Maximum value**: 16
	• **Location**: Metric Settings [sa_metric_config_setting]

ignite_data_streamer_thread_count	Number of threads that data streamers use to write data
	cache.
	
	The data streamer creates batches from individual data
	adds these batches to the cache.
	
	This property is for internal use.
	
	• **Type**: integer
	• **Default value**: 8
	• **Other possible values**:
	  ○ **Minimum value**: 8
	  ○ **Maximum value**: 32
	• **Location**: Metric Settings [sa_metric_config_setting]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Usage

ignite_data_streamer_flush_frequency	Frequency of automatic flush of data streamers.
	
	The data streamer creates batches from individual data adds these batches to the cache.
	
	This property is for internal use.
	
	• **Type**: integer
	• **Default value**: 1000
	• **Other possible values**:
	  ○ **Minimum value**: 0
	  ○ **Maximum value**: 30000
	• **Location**: Metric Settings [sa_metric_config_setting]

ignite_data_streamer_buffer_size	Buffer size of data streamers.
	
	The data streamer creates batches from individual data adds these batches to the cache.
	
	This property is for internal use.
	
	• **Type**: integer
	• **Default value**: 102400
	• **Other possible values**:
	  ○ **Minimum value**: 1024
	  ○ **Maximum value**: 1024000
	• **Location**: Metric Settings [sa_metric_config_setting]

ignite_data_streamer_batch_size	Size of the batch that is added to a data streamer.
	
	The data streamer creates batches from individual data adds these batches to the cache.
	
	This property is for internal use.
	
	• **Type**: integer
	• **Default value**: 5000
	• **Other possible values**:
	  ○ **Minimum value**: 1
	  ○ **Maximum value**: 10000
	• **Location**: Metric Settings [sa_metric_config_setting]

ignite_data_streamer_batch_flush_frequency	Frequency of automatic flush of internal batching befor data to data streamers. The data streamer creates batches from individual data adds these batches to the cache. This property is for internal use.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Usage

	• **Type**: integer
	• **Default value**: 5000
	• **Other possible values**:
	  ○ **Minimum value**: 1000
	  ○ **Maximum value**: 30000
	• **Location**: Metric Settings [sa_metric_config_setting]

sa.metric.use.resource.binding	Enables resource binding. • **Type**: true | false • **Default value**: true • **Location**: System Property [sys_properties] table • **Learn more**: Metric binding to resources
sa.metric.display.mid.setup.action	Displays the **Setup Metric Intelligence** related link on t Server form, for a MID Server that has been validated a running. • **Type**: true | false • **Default value**: true • **Location**: System Property [sys_properties] table • **Learn more**: MID Server distributed cluster for Metric
sa.model_learner.maint_event_record_history_enabled	Stores historical information about CIs start and end ma mode, in the Series Event History [sa_series_event_his • **Type**: true | false • **Default value**: true • **Location**: System Property [sys_properties] table • **Learn more**: Understanding Metric Intelligence
sa_metric.maint_exclusion	Excludes metrics for CIs that are in maintenance mode learning. • **Type**: true | false • **Default value**: true • **Location**: System Property [sys_properties] table • **Learn more**: Understanding Metric Intelligence
sa.model_learner.maint_history_batch_size	Determines the number of data series per batch when information about CIs start and end maintenance mode Series Event History [sa_series_event_history] table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Usage
	Adjust the setting of this property if there are related pe issues. • **Type**: integer • **Default value**: 1000 • **Location**: System Property [sys_properties] table • **Learn more**: Understanding Metric Intelligence
sa_metric.scheduled_jobs_to_monitor_sys_id_list	Contains the sys_ids of the 3 scheduled jobs that self-h monitoring for Metric Intelligence monitors: • Metric Intelligence — Metric configuration job • Metric Intelligence — Sync tables with mid • Metric Intelligence — Validate MID Distributed cluste

Quick start tests for Metric Intelligence

Validate that Metric Intelligence still works after you make any configuration change such as apply an upgrade or develop an application. Copy and configure these quick start tests to pass when using your instance-specific data.

Metric Intelligence quick start tests require activating the Metric Intelligence plugin (com.snc.sa.metric).

Metric Intelligence quick start tests

Test	Description	Release version
OI: Health Metrics Collection	An individual test that verifies the following: • There is at least one Metric Intelligence Metrics extension that is running. • There are entries for the 'Health Metrics' in the Metric to CI table. • For each of the Metric Intelligence Metrics extensions that are currently running, that there is data stored in MetricBase. If the test initially fails, wait until the Metric Intelligence Metrics extension runs for at least 10-15 minutes. Then try again.	Madrid

Related topics

Quick start tests

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Event severities

The event severity numbers correspond to actual severities on the Metric Anomaly Score to Event Severity Maps page.

Event severity mapping

Event Severity Number	Event Severity
0	Clear (OK)
1	Critical
2	Major
3	Minor
4	Warning

Agent Client Collector for Visibility Content reference

Review this information for details on OS query scripts, data collected, and definition of terms.

OS module scripts for Agent Client Collector for Visibility Content (ACC-VC)

These platform-specific plugins invoke platform-specific module scripts. Each script outputs to a standard JSON payload regardless of the platform.

Operating System family/plugin	Name of module script
Linux: acc-f-modules	basic_inventory.rb (Starting with ACC-F version 2.10.1, this will not be used.)
	data_collection.rb
	running_processes.rb
	tcp_connections.rb
Windows: acc-f-modules	basic_inventory.rb (Starting with ACC-F version 2.10.1, this will not be used.)
	data_collection.rb
	running_processes.rb
	tcp_connections.rb
macOS Operating Systems: acc-f-modules	basic_inventory.rb (Starting with ACC-F version 2.10.1, this will not be used.)
	running_processes.rb
	tcp_connections.rb
Linux: acc-visibility-modules	installed_software.rb
	file_systems.rb
	storages_devices.rb (support for getting physical disk and corresponding disk partitions only)
	network_adapters.rb

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Operating System family/plugin	Name of module script
	local_users.rb
	enhanced_inventory.rb
	cloud.rb
Windows: acc-visibility-modules	installed_software.rb
	file_systems.rb
	storages_devices.rb (support for getting physical disk and corresponding disk partitions only)
	network_adapters.rb
	local_users.rb
	enhanced_inventory.rb
	cloud.rb
	sam_advanced.rb
	sam_processor.rb
	intel_ema.rb
macOS Operating Systems: acc-visibility-modules	installed_software.rb
	file_systems.rb
	storages_devices.rb (support for getting physical disk and corresponding disk partitions only)
	network_adapters.rb
	local_users.rb
	enhanced_inventory.rb
	cloud.rb

Note:

running_processes.rb and tcp_connections.rb are interdependent. For an efficient Discovery, keep both tcp_connections.rb and running_processes.rb together. Both the files are needed to get the complete data. If one is not there, it will not populate the other data.

To fetch information about all the running_processes on the macOS, you must provide sudo access to osquery. If this configuration is not made, then running_processes.rb will only fetch what is triggered by _servicenow user. For tcp_connections.rb, you must add 'sudo lsof' command in the sudoers file. If this is not done, tcp_connections.rb will only fetch what is triggered by _servicenow user.

Data collected

ACC-VC uses the same classification criteria as IP-based Discovery for classifying as computer or server. The subset of collected data includes the following categories:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Basic inventory – Starting with ACC-F version 2.10.1, this will not be used.
- Data collection - Collects the necessary data for classification and identification of a host CI. This includes host name, serial numbers, and OS information.
- Installed Software – cmdb_sam_sw_install (if SAM enabled) and cmdb_software_instance (if SAM not enabled)
- File Systems – cmdb_ci_file_system
- Storage Devices – cmdb_ci_disk and cmdb_ci_storage_device
- Serial Numbers – cmdb_serial_number
- Network Adapters – cmdb_ci_network_adapter and cmdb_ci_ip_address
- TCP Connections – cmdb_tcp
- Running Processes – cmdb_running_process (ACC-VC also classifies the running processes and creates application CIs in cmdb_ci_appl, where possible)
- Local User – cmdb_os_user (Populates the local users for all the Operating Systems that ACC-VC supports by adding a new module called local_user)
- Enhanced inventory – Collects enhanced data (such CPU info, start_date, object_id) which is not necessary for identifying a unique host.

Relationships between host computer, file systems, and storage devices

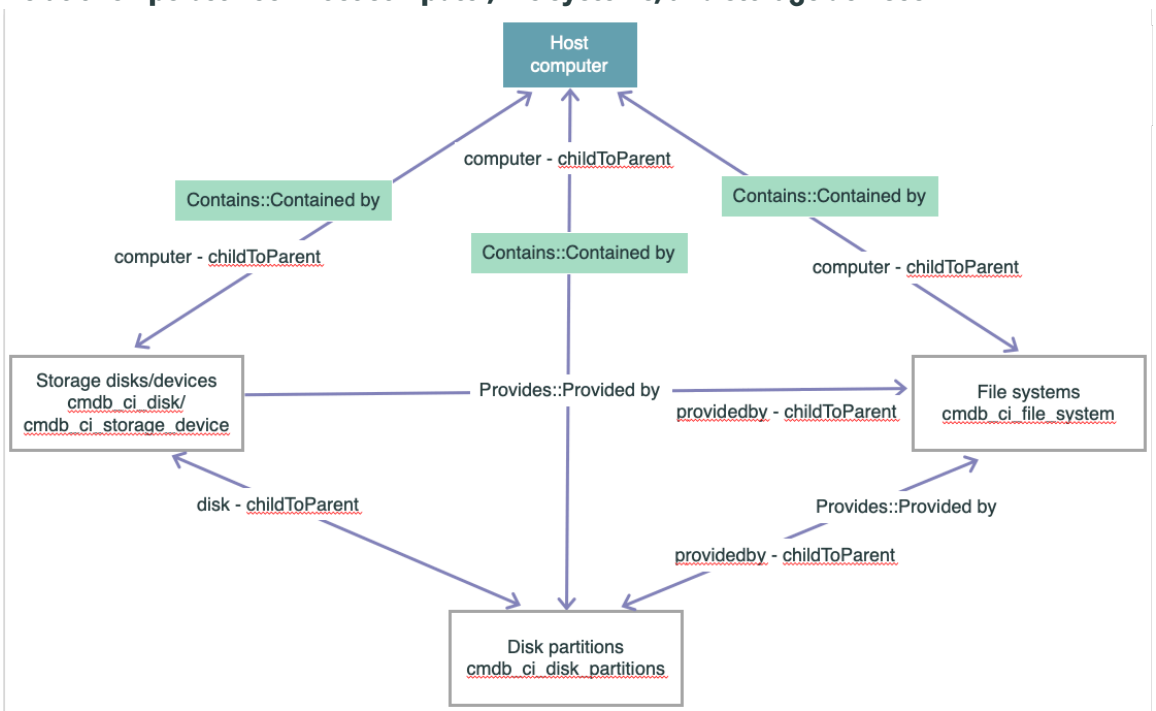

Agent Client Collector terms

Agent Client Collector (ACC)

The software component installed on target hosts that communicates with the MID Server. Sometimes referred to as Agent. This component is a ServiceNow derivative of Sensu-Go.

Agent Client Collector Framework (ACC-F)

A ServiceNow base scoped application that leverages ACC and provides core capabilities (including Check Types, Check Definitions, Policies, and so on) enabling other ACC scoped applications, including ACC-M and ACC-VC.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agent Client Collector Monitoring Monitoring (ACC-M)

A ServiceNow scoped application enabling monitoring use cases.

Agent Client Collector for Visibility Content (ACC-VC)

A ServiceNow scoped application that implements push-based Discovery leveraging ACC and ACC-F.

Horizontal IP-based Discovery

Traditional Discovery available to customers prior to ACC-VC. It discovers data via Probes and Patterns through the MID Server, requiring the Discovery Plugin.

Modules

A subset of discovered data that is populated as part of Discovery. Examples of Modules include: Data Collection, Installed Processes, Serial Numbers, File Systems, Storage Devices, Network Adapters, Running Processes, and TCP Connections.

Push-based Discovery

Discover data via ACC, ACC-F, ACC-VC, and the MID Server with a direct push of data from the target host. Discovery does not require specific IP range configuration on a Discovery Schedule or providing Discovery Credentials for target host.

Sensu-Go agent

The Free and Open Source project which ACC is derived from.

Virtual Machine Instance

A virtual target host running inside of a hypervisor, either on-premise or within a cloud service provider like AWS, GCP, or Microsoft Azure.

Virtual Machine Image

A snapshot of a live virtual machine instance persisted to a file system or cloud storage.

Agent Client Collector for Visibility Content default checks and policies

Agent Client Collector for Visibility Content (ACC-VC) provides various checks and policies as well as a business rule.

Policies

Note: ACC-VC policies execute at a frequency of once per day. The total data ingested would be approximately 572KB. This takes into consideration an average of approximately 1500 installed software applications and approximately 500 running processes other than CI data per machine.

ACC-VC policies

Name	Description	Check definitions
Enhanced Discovery	Runs on a schedule, by default every 24 hours (86400 seconds). The policy interval can be adjusted, for example to run every 4 hours (set the interval to 14400). The ACC-VC policy configuration is synced to all agents based on the policy filter defined by	Enhanced Discovery

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC policies (continued)

Name	Description	Check definitions
	ACC-VC. Update the following ACC-F system properties, if needed: • *sn_agent.disco_minimum_threshold_for_rediscovery_minu* to avoid discovering the system too frequently. • *sn_agent.disco_disable_ci_clobber_of_agentless_disco* to avoid Discovery conflicts. • *sn_agent.disco_ci_clobber_of_agentless_disco_thresho* to avoid Discovery conflicts. On Windows non-persistent VDI endpoints (NPVDI), ACC-VC can skip some inventory modules. For more information, see ACC-VC NPVDI system properties.	
SAM Discovery	Responsible for capturing the software installed on any endpoint device, such as Windows desktops or macOS servers.	Software installations and usage metrics
SAM background	Enables a background job for processing the Osqueryd logs for SAM on Windows and macOS endpoint devices.	SAM background log check
SAM background (Non OsqueryD)	Enables a background job to collect SAM information using osqueryi instead of osqueryd.	SAM Background Policy (Non OsqueryD)
Software installed	Responsible for capturing the software installed on all devices except for Windows endpoint devices. The data collected is stored in the [cmdb_sam_sw_install] table. Scheduled to run every 24 hours.	installed software
File-based Discovery background policy	Takes the config file as input from the instance to an agent. Scans the system using config file parameters and stores the output in two separate files on the agent. • FBDSAMOutput.json: Stores metadata related to the set of file names generated from the *samp_file_name* table.	File-based discovery background

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC policies (continued)

Name	Description	Check definitions
	• FBDFileOutput.json: Stores metadata related to files scanned by a wildcard extension. Runs on the agent when file-based discovery is invoked. For details, see Agent Client Collector File-Based Discovery. Default: false. To activate, navigate to **All** > **Discovery Definition** > **Configuration Console** and in the **File Based Discovery** section, activate the Enable File Based Discovery toggle switch.	
File-based Discovery - SAM	Discovers known software files on the endpoint. Uses an allowlist of recognized software filenames maintained by ServiceNow. When a file on disk matches an allowlist entry, FBD uses the FileBasedDiscovery API to identify the collected software metadata (file name, path, size, and version), identifies the software package it belongs to and records the installation on the instance. Unrecognized files are tracked in the unidentified file records table (cmdb_unidentified_file_set). Runs daily. Default: false. To activate, navigate to **All** > **Discovery Definition** > **Configuration Console** and in the **File Based Discovery** section, activate the Enable File Based Discovery toggle switch.	File-based discovery - SAM
File-based Discovery - SWID tag	Enables SWID tag scanning on a Windows, Linux or macOS platform. When enabled, the scanner looks for `.swid`, `.swidtag`, and `.cmptag` files in the	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC policies (continued)

Name	Description	Check definitions
	configured scan directories. Stores results in the following tables: • cmdb_swid_tag: Parsed SWID tag records • cmdb_file_information: Individual file records • cmdb_sam_sw_install: Individual software records Default: false. To activate, navigate to **All** > **Discovery Definition** > **Configuration Console** and in the **File Based Discovery** section, activate the Enable File Based Discovery toggle switch.	
File-based Discovery - File management	Discovers files based on customer-defined rules. Administrators configure which file extensions to look for and define filename matching rules such as exact match, starts with, ends with, or contains. This policy builds a device-level file inventory based on the organization's specific needs. Results are stored in the sn_acc_vis_content_device_file_information table. Configure rules in the File Matching rules (sn_acc_vis_contet_file_config) and File extensions (sn_acc_vis_content_file_extension) properties. For details on these properties, see Agent Client Collector File-Based Discovery properties. File Management supports delta scanning; after the initial full scan, only added, modified, and deleted files are sent on subsequent runs. For details on delta scanning, see Agent Client Collector File-Based Discovery.	File-based discovery - File management

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC policies (continued)

Name	Description	Check definitions
	Default: false. To activate, navigate to **All** > **Discovery Definition** > **Configuration Console** and in the **File Based Discovery** section, activate the Enable File Based Discovery toggle switch.	
Oracle GLAS Java Process Discovery	Discovers running Java processes on endpoints and collects installation data for the Oracle GLAS license compliance report. Runs weekly.	Oracle GLAS Java Process Discovery
Oracle GLAS HW data collection	Executes Oracle's hardware validation script on non-Oracle, non-macOS hosts to collect CPU/socket data for GLAS reporting. Runs weekly.	Oracle GLAS HW data collection
Software discovery via running processes	A background policy which discovers software without executable files in the configuration table (*sn_acc_vis_content_process_based_sw_config*) and stores them in agent's local DB. The SAM discovery policy (for Windows and macOS) and Software Installed Policy (for Linux) gathers the information collected by this policy and saves it in the *cmdb_sam_sw_install* table.	Software discovery via running processes
VISC Get application metric	Retrieves the SaaS application metrics from the agents. For details on enabling SaaS usage monitoring with ACC-VC, see the SaaS Usage Monitoring with Agent Client Collector [KB2320193] article in the Now Support Knowledge Base.	VISC Get application metric
VISC Get browser extension device init	Initializes the DEX browser extension with the host sysID.	VISC Get browser extension device init
VISC Get browser extension init	Initializes the DEX browser extension with logged-in users.	VISC Get browser extension init

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC policies (continued)

Name	Description	Check definitions
VISC Get URL metrics	Controls the collection of web usage data from Windows and macOS managed devices. Runs daily. Default: Inactive. To activate the policy, set the *sn_acc_vis_content.enable_full_monitoring* property to **true**. For details on web usage data system properties, see Web usage data collection tables and fields.	• VISC Get URL metrics • VISC Get targeted URL metric

Note: Windows endpoint devices include devices that have a Windows operating system and belong to CI class: computer.

See System properties for more details. For more details on policies, see Checks and policies.

Check type

ACC-VC has the following check types: *Enhanced Discovery*, *SAM Advanced Discovery*, and *Installed Software*.

Enhanced Discovery

This check type is responsible for invoking the EnhancedDiscoveryHandler script include that processes the payload produced by endpoint_discovery.rb as executed by ACC.

Used by File-base Discovery.

SAM Advanced Discovery

This check type is for the SAM Discovery policy that invokes the EnhancedDiscoveryHandler script include for processing the SAM data produced by the sam_advanced.rb file.

Installed Software

This check type for the *Software installed policy* that invokes the EnhancedDiscoveryHandler script include for processing the installed software data produced by the installed_software.rb file.

Check definitions

ACC-VC check definitions

Name	Description
Enhanced Discovery	Synced to all agents based on the policy filter defined by ACC-VC. The Check definition is configured to run with certain assets and determines what gets synced between the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC check definitions (continued)

Name	Description

	agent and the MID Server. For more details on policies, see Checks and policies.
	
	Note:
	
	For the agent to retrieve the OS serial numbers and TCP connections along with associated running processes, sudo access for *dmidecode* and *ss* commands is required on Linux systems. For example, this content could be added to /etc/sudoers or to an individual file in `/etc/sudoers.d/`:
	
	`Cmnd_Alias AGENT_ACC_V = /usr/sbin/dmidecode -s baseboard-serial-number,/usr/sbin/dmidecode -s chassis-serial-number,/usr/sbin/dmidecode -s system-serial-number,/usr/sbin/dmidecode -s system-uuid,/usr/sbin/ss -tanp`
	`servicenow ALL=(root) NOPASSWD:AGENT_ACC_V`
	
	For the sudoers configuration example, the *must_sudo* check parameter, and the equivalent macOS and Windows privilege requirements, see the "Privileges for running check commands" section in Checks and policies.

SAM background log check	Runs every 8 minutes and performs inline aggregation of data generated from Osqueryd logs. After collecting the data, it writes all the intermediate data results into a temporary marker file which is reused in the next run. This reuse limits the number of log files and disk space needed on target systems. **Note:** You may notice a spike in system resource consumption, as the background aggregation check runs every interval.
Software installations and usage metrics	Collects data every 24 hours.
Installed software	Fetches installed software data for all devices other than Windows and macOS endpoint devices.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC check definitions (continued)

Name	Description
File-based discovery background	Runs a file scanning background job on the agent.
File-based discovery	Fetches the file data from the agent.
Oracle GLAS Java Process Discovery	Runs `pbd_java_discovery.rb` via osquery to collect Java version, vendor, AMC, and licensing flags. Timeout: 600s. Runs on all platforms.
Oracle GLAS HW data collection	Runs `endpoint_discovery.rb --select=oracle_glas_hw_discovery` to execute Oracle's `ct_cpuq` script. Executes Oracle's hardware validation script on non-Oracle, non-macOS hosts to collect CPU/socket data for GLAS reporting. Runs weekly.
VISC Get targeted URL metric	Gathers per-day usage metrics for the monitored URL from each managed device, and sends the data back to the instance.

Business rule

The *`Enhanced Discovery - On CI Delete`* business rule triggers the Endpoint Discovery Check when the CI associated with a given CI is deleted from sn_agent_cmdb_ci_agent.

Agent Client Collector MID-less installation command parameters

Description of the command line parameters used during MID-less Agent Client Collector installation.

Installation command line parameters

Parameter	Description
CONNECT_WITHOUT_MID	Specify **true** to enable MID-less Agent Client Collector installation. If you don't specify a value, the regular MID Server websocket installation is invoked.
ACC_CNC	Specify a public gateway endpoint for Agent Client Collector.
REGISTRATION_KEY	Active and valid registration key for the agent to communicate with.
INSTANCE_URL	ServiceNow instance that the agent is registered with.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Installation command line parameters (continued)

Parameter	Description
ACC_ALLOW_LIST	Optional string; to disable the allow list, set *ACC_ALLOW_LIST=0*. By default, the allow list is enabled. **Note:** Disabling the allow list compromises your system's security and is intended as a temporary measure while you're updating the allow list. After completing the update, enable the allow list by setting **ACC_ALLOW_LIST=1**.
ACC_VALIDATE_SIG	Optional string; to disable verification of the installer certification validation, set *ACC_VALIDATE_SIG=0* that you can add to disable the verification of the installer certification validation. By default, validation is enabled. **Note:** Disable the certification validation only when using non-standard signature validation tools.

Related topics

Configure the websocket server on the MID Server

Data collected during Agent Client Collector for Visibility Content File-Based Discovery

The following tables describe the data collected during Agent Client Collector for Visibility Content file-based Discovery.

Java Audit [ora_java_audit]

Field	Description
Name	Name of the record. For example: Java@hostname
Physical Machine Name	Name of the host CI
Machine Type	Either Server(S) or Desktop(D)
Java Publisher/Runtime Vendor	Java Vendor name
Operating System	Operating system in which the java is run
Environment Type	Production or Development or QA environment
Java Release/Version/Build	Java Version
Installation Path	Absolute installation path of the java binary
Installation Date (DateTime)	The date and time of the java binary file created on the file system
Release File (Y / N)	Indicates if a *release.md* file exists in the installation

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Java Audit [ora_java_audit] (continued)

Field	Description
Processor Make/Model	Processor model name
Sockets (Integer)	Number of CPU sockets
Cores (Integer)	Number of CPU cores
V6 and earlier	Indicates if you have a 6 or earlier java version
Open Source	Indicates if you have an open-source java installation
Unlock Commercial features	Indicates if you have the commercial features (like JFR) enabled. Not collected during ACC-VC file-based discovery (FBD).
Mission Control	Indicates if the mission control plugin is enabled
AMC Agent	Indicates if the AMC agent plugin is enabled. Not collected during ACC-VC file-based discovery (FBD).
Usage Tracker	Indicates if the Usage Tracker feature is enabled
Virtual Machine Object ID	Object ID of the Virtual machine instance
PID	Running process PID
IP Address	IP address of the host/VM
First Discovered	The date of the first discovery
Most recent discovery	The date of the most recent discovery
Discovery Source	The discovery method
Application	Reference to the [cmdb_ci_appl] table
Software installation	Reference to the [cmdb_sam_sw_install] table
Java evidence	Reference to the [cmdb_ora_java_evidence] table
Host	Reference to the [cmdb_ci] table

Java Evidence [ora_java_evidence]

Field	Description
Running process	Reference to the [cmdb_running_process] table
Host CI	Reference to the [cmdb_ci_computer] table
File information	Provides a reference to the [cmdb_file_information] table. Relevant only

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Java Evidence [ora_java_evidence] (continued)

Field	Description
	for file-based discovery (FBD); in pattern-based discovery, this field remains empty.

Related topics

Discover java installation data using Agent Client Collector for Visibility Content file-based discovery

Software edition configuration fields

Description of the software edition configuration fields in use when identifying the edition of Windows software.

Software Edition Configuration - New record page

Field	Description
Registry operation	• **Read value**: Reads the values stored in the specified registry key. • **Read subkey**: Reads the key stored in the registry path.
Registry path	Windows registry path indicating where the software is registered.
Product	Software product for which the edition needs to be determined.
Active	Option to activate edition detection for the indicated product.
Priority order	Indicates the hierarchy of values that take precedence when multiple editions are discovered.

Agent Client Collector File-Based Discovery properties

Configure File-Based Discovery behavior using system properties that control scanning paths, performance throttling, and file filtering options.

Discovery settings properties

All File-Based Discovery configuration properties are configurable from the System Properties page (**All** > **System Properties** > **All Properties**).

File-Based Discovery configuration properties

Property	Description	Platf
glide.discovery.file_discovery.path.windows	Directories to scan on Windows endpoints (pipe-separated). Supports environment variables	Win

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File-Based Discovery configuration properties (continued)

Property	Description	Platf
	(%PROGRAMFILES%) and wildcards.	
glide.discovery.file_discovery.path.linux	Directories to scan on Linux/macOS endpoints (colon-separated). Supports environment variables ($HOME) and wildcards.	Linu mac
glide.discovery.file_discovery.ignore_path.windows	Directories to exclude from scanning on Windows (pipe-separated).	Win
glide.discovery.file_discovery.ignore_path.unix	Directories to exclude from scanning on Linux/macOS (colon-separated).	Linu mac
glide.discovery.file_discovery.throttle_size.windows	Number of files to process before pausing (Windows).	Win
glide.discovery.file_discovery.throttle_size.unix	Number of files to process before pausing (Linux/macOS).	Linu mac
glide.discovery.file_discovery.sleeptime.windows	Pause duration in milliseconds after each throttle batch (Windows).	Win
glide.discovery.file_discovery.sleeptime.unix	Pause duration in milliseconds after each throttle batch (Linux/macOS).	Linu mac
glide.discovery.file_discovery.skip_hidden_folders.windows	When enabled, skips hidden files and folders during scanning (Windows).	Win
glide.discovery.file_discovery.skip_hidden_folders.unix	When enabled, skips dot-prefixed (hidden) files and folders during scanning (Linux/macOS).	Linu mac

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File-Based Discovery configuration properties (continued)

Property	Description	Platf
glide.discovery.file_discovery.scan_swid.windows	Enables or disables SWID tag scanning on Windows platforms. When enabled, the scanner looks for .swid, .swidtag, and .cmptag files in the configured scan directories.	Win
glide.discovery.file_discovery.scan_swid.unix	Enables or disables SWID tag scanning on Linux and macOS platforms. When enabled, the scanner looks for .swid, .swidtag, and .cmptag files in the configured scan directories.	Linu mac

File based discovery system properties

Property	Description
File matching rules (*sn_acc_vis_content_file_config*)	Defines the filename matching rules. Each rule specifies: • File name: The pattern to match. • Condition: `exact_match`, `starts_with`, `ends_with`, or `contains`. • Extension: Reference to the *sn_acc_vis_content_file_extension* property.
File extensions (*sn_acc_vis_content_file_extension*)	Defines which file extensions to scan for File Management (such as .exe, .dll, .msi, .pkg, and so forth). Each record specifies an extension and the platform it applies to (Windows, Mac, or All).
Full Scan Mode (*sn_acc_vis_content.file_disovery.enable_file_config*)	When false (strict mode), only files that match a rule in the *File Configuration* table are discovered. When true, all files with a matching extension are discovered only if no specific rule exists for them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

File based discovery system properties (continued)

Property	Description
	Default: false
Delta Scan Frequency (*sn_acc_vis_content.file_discovery.full_scan_frequency*)	The number of delta scans to perform before forcing a full scan. This ensures data accuracy over time. For details on delta scanning, see Agent Client Collector File-Based Discovery. Default: 5
MME Types (Linux only) (*sn_acc_vis_content.file_discovery.linux_mime_types*)	Colon-separated list of MIME types to scan on Linux (for example, `application/x-pie-executable:text/x-shellscript`). On Linux, File Management uses MIME type detection instead of extension-based matching.
Running process-based discovery (*sn_acc_vis_content.file_discovery.rbd_process_scan_enabled*)	Running process-based discovery enables the ACC-VC agent to detect software running outside of standard scan directories, including applications that a directory-only scan can't detect. Default: false
Archive file scanning (*sn_acc_vis_content.file_discovery.archive_scan_enabled*)	Archive file scanning extends file-based discovery to look inside ZIP and JAR archive files and report the files contained within them. This enables discovery of software artifacts that are distributed in compressed packages without requiring the archive to be extracted on the endpoint. Default: false

Performance throttling

The throttling properties control how File-Based Discovery manages system resources during scanning operations. The throttle size determines how many files are processed in each batch, while the sleep time controls the pause duration between batches to prevent excessive system load.

Path configuration

Path properties support environment variables and wildcards for flexible directory specification. Use pipe separators (|) for Windows paths and colon separators (:) for Linux and macOS paths. The ignore path properties allow you to exclude specific directories from scanning to improve performance and avoid unnecessary data collection.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Web usage data collection tables and fields

Description of the tables and their fields used when gathering web usage data from your Windows and macOS machines.

Data collection tables

Name	Description
User Web Usage (sn_acc_vis_content_user_web_usage)	Stores per-user, per-device, per-domain web usage records, including page visit counts and data measured by each day.
Application Domain (sn_acc_vis_content_application_domain)	Stores domain-level information. Updated with aggregated totals (total page views, total users, daily trends) by the daily aggregation job. Domains not previously configured are automatically created here when collected via full monitoring.
Application (sn_acc_vis_content_application)	Stores service application records that group one or more domains. Each application is linked to a CMDB service and a source.

User web usage table fields

Name	Description
Domain (domain)	The application domain of the visited website.
Device (device_id)	The computer (cmdb_ci_computer) from which the website was accessed.
User (user)	User that accessed the website.
Last Accessed (last_accessed)	Date and time that the website was last accessed.
Page Visits (page_visits)	Number of page visits for the current user-domain-device combination.
Page Visits Daywise (page_visits_daywise)	A JSON object storing page visit counts, broken down by day for the retention period (for example, {"2026-03-15": "12", "2026-03-16": "8"})

Application Domain table fields

Name	Description
Application (application)	Application record that the domain belongs to.
Domain (domain)	Domain name that can be monitored (such as `example.com`). The display field. Must be unique.
Category (category)	The URL category (sn_acc_vis_content_url_category) for the domain.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application Domain table fields (continued)

Name	Description
Domain Type (domain_type)	Indicates whether the domain is internal or external. Default: external
CMDB endpoint http (endpoint_http)	The CMDB HTTP endpoint (cmdb_ci_endpoint_http) for SaaS applications.
Metric monitoring (monitored)	Indicates whether the domain is monitored for metrics. Default: false
Total accessed time (total_accessed_time)	Total accumulated access duration for the domain.
Total users (total_users)	Total number of unique users who accessed the domain. Updated by the daily aggregation job.
Last accessed time (last_accessed_time)	Date and time that the domain was last accessed.
Total page views (total_page_views)	Total number of page views for the domain. Updated by the daily aggregation job.
Total page views daywise (total_page_views_day_wise)	A JSON object storing total page view counts broken down by day (for example, {"2026-03-15": "120", "2026-03-16": "85"}). Updated by the daily aggregation job.

Application table fields

Name	Description
Service (ci_id)	The CMDB service (cmdb_ci_service) associated with the application. The display field.
SAMP product (samp_sw_product)	The Software Asset Management product (samp_sw_product) for the application.
Source (source)	The source (sn_acc_vis_content_source) that discovered or created the application.
Type (type)	Type of application: web or installed Default: web
Metric monitoring (monitored)	Indicates whether the application is monitored for metrics. Read-only. Default: false

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Archive file scanning filtering rules and limits

Reference information for the filtering rules and performance safeguards that apply when Agent Client Collector for Visibility Content scans ZIP and JAR archive files during File-Based Discovery.

Filtering rules

Archive entries are filtered using the same rules as files found on disk, with the following differences.

Archive file scanning filtering rules

Capability	Regular files	Archive entries
SAM allowlist matching	Yes	Yes
File Management extension rules	Yes	Yes
MIME type detection (Linux)	Yes	No. The file does not exist on disk, so the system can't inspect its content type
Version detection	Yes	No. Version can't be queried for a file that hasn't been extracted
SWID tag support	Yes	No. Archive contents aren't decompressed or read, so SWID tags inside ZIP or JAR files aren't processed

Note: On Linux, File Management matching relies on MIME type detection rather than extension rules. As a result, archive entries can only be matched by SAM rules, not by MIME type.

Performance safeguards

The following limits are enforced to prevent archive scanning from affecting agent performance or endpoint stability.

Archive file scanning performance safeguards

Safeguard	Limit	Behavior when exceeded
Maximum archive file size	500 MB	Archive is skipped entirely
Maximum entries per archive	50,000 files	Scanning stops after the limit; entries discovered so far are retained
Per-archive timeout	30 seconds	Scanning of that archive is aborted; entries discovered so far are retained
Corrupt or password-protected archives	Not applicable	Detected automatically, logged, and skipped
Nested archives (archive inside an archive)	Not applicable	Not scanned. Only top-level archives on disk are inspected
Global file limit (*maxFiles*)	Configured per policy	Archive entries count toward the global file limit; scanning stops when the limit is reached

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Software filter tables

The following table lists the software filter tables and the purpose each one serves.

Table	Who edits it	Purpose
Software Install Filter (samp_sw_install_filter)	ServiceNow	Base-system rules maintained by ServiceNow that cover common, well-known filtered software across all customers. Read-only.
Software Install Custom Filter (samp_sw_install_custom_filter)	SAM admin	Where you create, edit, and manage your own organization-specific exclusion rules for internal tools, custom shortcuts, or anything the base-system rules don't cover.
Software Install Filter Staging (samp_sw_install_filter_staging)	System-generated	An audit log of every software install that has been excluded, which rule excluded it, and the last date it was reconfirmed as irrelevant. Use this table to review and validate what is being filtered. Read-only.

Software filter fields

The following table lists the software filter fields from which you can create a filter.

Software filter fields

Field	Description
Name	A descriptive label for the rule. For example, *`Exclude company name shortcut`*.
Software name	(Optional) Enter the display name of the software. In the **Condition** dropdown, select the matching type: • Starts with • Ends with • Contains • Exact match • Regex
Publisher	(Optional) Enter the software vendor or publisher name. In the **Condition** dropdown, select the matching type: • Starts with • Ends with • Contains

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Software filter fields (continued)

Field	Description
	• Exact match • Regex
Version	(Optional) Enter the version string. In the **Condition** dropdown, select the matching type: • Starts with • Ends with • Contains • Exact match • Regex **Note:** Although this field and the **Software name** and **Publisher** fields are optional, at least one of them must have a value.
OS filter	Select the operating system the rule applies to: • All • Windows • Linux • macOS
Discovery source	(Optional) Enter the Discovery source to restrict the rule to installations reported by a specific source. Leave blank to apply the rule to installations from any source.
Active	Toggle to **Active** to enable the rule.

License key discovery and access control tables

Reference information for the tables, fields, access control, and scheduled job used by license key discovery in Agent Client Collector for Visibility Content.

Tables

License key discovery uses the following tables.

License key discovery tables

Table	Label	Purpose
sn_acc_vis_content_license_registry_config	License Registry Configuration	Admin-managed list of registry paths and value names to collect from managed endpoints.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

License key discovery tables (continued)

Table	Label	Purpose
sn_acc_vis_content_license_keys	License Keys	Stores the license key values collected from endpoint devices.

License Keys table fields

The **License Keys** [sn_acc_vis_content_license_keys] table stores the following fields for each discovered key.

License Keys form

Field	Description
Configuration Item	Device the key was collected from. Table: cmdb_ci_computer.
Product	SAM product associated with this key. Table: samp_sw_product.
User	This field is automatically set to the user assigned to the CI. Table: sys_user.
License Key Config	Configuration rule that discovered this key. Table: sn_acc_vis_content_license_registry_config.
Resolved Registry Path	Exact registry path the key was read from on the endpoint.
License Key	Discovered key value.
Last Scanned	Date the key was last detected by the agent.
Absent	Selected when the key is no longer found on the device. Records are never deleted, they are marked **Absent** instead.

License Registry Configuration form fields

The following table describes the fields in the License Registry Configuration [sn_acc_vis_content_license_registry_config] table.

License Registry Configuration form

Field	Description
Product	SAM product this license key belongs to. Links to the SAM Professional product catalog.
Registry Path	Windows registry path where the key is stored. For example, `HKLM\SOFTWARE\Microsoft\Office\16.0\Registration`.
Registry Value Name	Specific registry value to read. For example, `DigitalProductId`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

License Registry Configuration form (continued)

Field	Description
Name	This field is automatically set to a combination of the **Product** and **Registry Value Name** values.
Description	Notes about this configuration record.
Active	Indicator of whether this record is included in future collections. Clearing this check box removes the record from future collections.

Note: The combination of **Product**, **Registry Path**, and **Registry Value Name** must be unique across all configuration records.

Access control

The following table shows which roles can access each license key discovery table.

License key discovery role access

Role	License Registry Configuration table	License Keys table
ACC application admin (agent_client_collector_admin)	Read, write, create	Read only
Discovery admin (discovery_admin)	Read, write, create	Read only
SAM admin (sam_admin)	Read only	Read only
All other roles	No access	No access

Scheduled job

The following scheduled job supports license key discovery.

License key discovery scheduled job

Job	Schedule	Condition
Refresh License Key Config File	Daily at 8:30 a.m. UTC	Runs only when the *sn_acc_vis_content.enable_license_key_discovery* property is set to `true` and the SAM Professional plugin (com.snc.samp) is activated.

License key discovery tables and configuration

License key discovery uses configuration tables to define file-matching rules and extraction parsers, and stores results in dedicated tables for reporting and reconciliation.

Configuration tables

Two configuration tables work together to define and execute the license key discovery framework.

sn_acc_vis_content_license_file_config

Defines which files to treat as license files and which parser to invoke when a file is found.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column name	Type	Description
display_name	String	Unique name identifying the rule.
file_name	String	File name or partial name to match.
file_name_condition	Choice	How to match file_name: `exact_match` (default), `starts_with`, `ends_with`, `contains`.
file_extension	String	Optional file extension to further narrow the match (no leading dot before the extension).
parser	Reference	Reference to `sn_acc_vis_content_license_parser` record that processes matched files.
platform	Choice	Operating system(s) where the rule applies: `all` (default), `windows`, `linux`, `mac`.
order	Integer	Priority when multiple rules could match the same file (default 100; lower values win).
active	Boolean	Indicates whether the rule is currently in use (default = **true**).

sn_acc_vis_content_license_parser

Holds the extraction logic for one license file format. A single parser can be referenced by multiple file-matching rules.

Column name	Type	Description
parser_name	String	Unique name identifying the parser.
parsing_script	Script	Server-side JavaScript that reads file content and sets `answer` to an array of extracted results.
sample_input	String	Representative file content used with Test Parser action.
sample_output	String	Expected extraction result for sample_input, used to validate the parser.
active	Boolean	Whether the parser is currently in use (default = **true**).

Results tables

Once a file has been scanned and parsed, the results flow through three tables, each holding a distinct stage of the outcome.

sn_acc_vis_content_license_file_content

The raw, per-file result of scanning one file on one computer — captured regardless of whether a product match was later found.

Column name	Type	Description
ci	Reference	Reference to `cmdb_ci_computer` — the computer where the file was found.
license_file_config	Reference	Reference to `sn_acc_vis_content_license_file_config` — the rule that matched this file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column name	Type	Description
file_path	String	Folder path where the file was found.
file_name	String	Name of the file.
file_size	Long Integer	Size of the file in bytes.
file_version	String	File version, if available.
discovered_product_name	String	Product name extracted by the parser.
discovered_vendor	String	Vendor name extracted by the parser.
discovered_license_key	String	License key extracted by the parser.
discovered_expiry_date	Date	Expiry date extracted by the parser, if present in the file.
discovered_file_content	Long Text	Raw file content, stored only when parsing fails (for troubleshooting).
scan_id	String	Identifier of the scan that produced this result.
last_scanned	Date	When this file was last scanned.
parse_status	Choice	`parsed` or `parse_failed`.
absent	Boolean	Set to **true** when the file is no longer found in a later scan (default = **false**).

sn_acc_vis_content_license_keys

The final, resolved license key tied to a computer and a recognized software product — the table which is reported on.

Column name	Type	Description
ci	Reference	Reference to `cmdb_ci` — the computer the license key belongs to.
product	Reference	Reference to `samp_sw_product` — the matched catalog product.
user	Reference	Reference to `sys_user` — user associated with the CI, if resolved.
license_file_config	Reference	Reference to `sn_acc_vis_content_license_file_config` — the rule that led to this result (file-method only).
license_file_path	String	Path of the source license file (file-method only).
license_key	String	The resolved license key.
last_scanned	Date	When this license key was last confirmed.
absent	Boolean	Set to **true** when no longer found in a later scan (default = **false**).

Note: Columns `license_key_config` and `resolved_registry_path` also exist on this table but are populated only by the older registry-based discovery method.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

sn_acc_vis_content_license_file_unresolved

A holding area for discovered product names that could not be automatically matched to the software catalog.

Column name	Type	Description
discovered_product_name	String	Product name as extracted by the parser (unmatched).
discovered_vendor	String	Vendor name as extracted by the parser.
parser	Reference	Reference to `sn_acc_vis_content_license_parser` — the parser that produced this discovery.
normalised_product_name	Reference	Reference to `samp_sw_product` — set by the customer to manually map this discovery to a real catalog product.

Framework constraints

The following constraints are built into the framework:

- **File read permission** — The account under which the ACC agent service runs must have read permission on the folders and files being scanned. This is an operating-system-level requirement. If the agent encounters a file or folder it cannot read, it logs a permission error for that item and continues scanning.
- **File size limit** — License file content is capped at 1 MB. Files larger than this limit are skipped and not read for license key extraction.

These constraints are intentional guardrails and cannot be changed through configuration.

Product reconciliation

Once a parser successfully extracts a product name from a license file, the discovery is reconciled against the software catalog (`samp_sw_product`). If a confident match is found, a final record is created in `sn_acc_vis_content_license_keys`. If no confident match is possible, the discovery is placed in `sn_acc_vis_content_license_file_unresolved` for manual review and mapping.

Manual mappings are performed only once per unique product name. After a customer sets the **normalised_product_name** field, all future scans reporting the same discovered product name resolve automatically with no repeated manual effort.

System properties and policies

Component	Role
System property `sn_acc_vis_content.enable_license_key_discovery`	Master on/off switch (default = **false**). Must be **true** for license key data from either discovery method to be processed and stored.
Policy `File Based Discovery Policy - License Key`	Instructs the agent during its regular file-

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Component	Role
	scanning activity to watch for files matching an active License File Configuration rule. Depends on the File Based Discovery Background Policy also being active.
Check definition `File Based Discovery - License Key`	Defines the concrete agent command executed (files.rb --compact --no-delta, scoped to the license-key scan identifier) and its timeout (10 minutes)

Running process-based discovery platform coverage and properties

Platform coverage identifies which operating systems are supported and what privileges the agent needs for full coverage. The system property controls whether the feature is enabled or disabled.

Platform coverage

Coverage depends on the privileges available to the agent on each operating system.

Platform coverage for running process-based discovery

Platform	Default coverage	Requirement for full coverage
Windows	Full coverage of all running processes.	The ACC service must run as the Local System account. Set the ACC service's **Log On As** to **Local System**.
Linux	Limited to processes owned by the agent service account.	The agent service account must have permission to query process information for processes owned by other users. The `servicenow` user must be able to run `osqueryi` without a password. For more information about `servicenow` user permissions for `osqueryi`, see Configure ServiceNow sudoers file.
macOS	Limited to processes owned by the agent service account.	The agent service account must have permission to query process information for processes owned by other users. The `_servicenow` user must be able to run `osqueryi` without a password. For more information about `_servicenow` user permissions for `osqueryi`, see Configure ServiceNow sudoers file.

Note: On Linux and macOS, if the additional privilege is not granted, the feature still runs but discovers directories only for processes owned by the agent service account. The daily scan continues normally on your configured directories. No errors are raised.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

System property

The following property controls whether running process-based discovery is active. You can configure it from the System Properties page (**All** > **System Properties** > **All Properties**).

Running process-based discovery system property

Property	Default	Description
sn_acc_vis_content.file_discovery.fbd_process_scan_enabled	`false`	Primary on/off control for running process-based discovery. When set to `true`, the agent policy that collects process directories is activated and the daily FBD scan includes process-discovered directories. When set to `false`, collection stops and the daily scan uses only your configured scan directories.

ACC-VC NPVDI system properties

System properties that control how Agent Client Collector for Visibility Content (ACC-VC) runs *Enhanced Discovery* and installed software checks on Windows non-persistent virtual desktop infrastructure (NPVDI) endpoints.

Note: These properties apply only to Windows endpoints that are configured as non-persistent. On all other endpoints, they have no effect. For details on how ACC-VC behaves on NPVDI endpoints, see ACC-VC NPVDI endpoints.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ACC-VC NPVDI system properties

Property	Description
sn_acc_vis_content.np_vdi.enhanced_discovery_skip_modules	• Type: string • Default value: `file_systems,net` Comma-separated list of valid: • `file_systems` • `network_adapters` • `storage_devices` • `local_users` • `intel_ema` • `memory_modules` Any other value is ignore The following inventory a • *Basic Inventory* • *Enhanced Invento* • *TCP Connections* • *Running Process*
sn_acc_vis_content.np_vdi.cache_installed_software_enabled	• Type: true | false • Default value: false Caches the full installed check runs reuse the cac Software usage metrics,

When you update either property, the instance regenerates the NPVDI configuration file (`NPVDIConfig.json`) that ACC-VC uses on NPVDI agents. A daily scheduled job also refreshes this file.

Metric Intelligence

ServiceNow® Metric Intelligence provides the ability to capture, and then explore and analyze operational metrics data, identifying and indicating anomalies. Metric Intelligence generates anomaly alerts that can be promoted to IT alerts and appear on the Service Operations Workspace and service health dashboard. You can leverage this analysis to prevent potential service outages.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore Learn the key features and business value that the Metric Intelligence application offers.	Configure Configure the features and components to enable Metric Intelligence.
Optimize Manage metrics, events, alerts and anomalies to get the most out of your metric data.	Reference Get additional information on Metric Intelligence settings and properties.

Exploring Metric Intelligence

Learn more about using Metric Intelligence to analyze metric data and identify anomalies.

Metric Intelligence overview

Metric Intelligence helps identify and prevent potential service outages. Based on historical metric data, Metric Intelligence indicates anomalous behavior of CIs which events might not capture.

Metric Intelligence users

Users

User	Description
Event Management user [evt_mgmt_user]	Can view alerts and their underlying metrics.
Event Management administrator [evt_mgmt_admin]	Can configure all metric definitions and connector settings.
Operator [evt_mgmt_operator]	Can view all metric definitions and connector settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric Intelligence workflow

The following illustration describes the layout and data flow within the Metric Intelligence application.

Metric Intelligence benefits

Benefit	Feature	Users
Monitor your system's health, performance, and availability through automated collection of events and metrics, leveraging automated configurations.	Agent Client Collector Monitoring	NOC Operator, Event Management administrator
Reduce noise by promoting only the most meaningful anomalies.	View anomaly alerts Create metric rules	Event Management administrator
Detect anomalies with AI-based anomaly detection, either with unsupervised machine-learning abnormal pattern detection (no user intervention), or by setting deterministic alert rules (manually setting a static threshold).	How Health Log Analytics generates alerts	Event Management administrator
Improve resolution time on open alerts and incidents with raw metric data visualization.	Metric Explorer	NOC Operator, Event Management administrator

What to explore next

To learn more about configuring and using Metric Intelligence, see:

- Configuring Metric Intelligence
- Optimizing Metric Intelligence
- Metric Intelligence reference

Understanding Metric Intelligence

Use Metric Intelligence to identify and prevent potential service outages. Metric Intelligence, based on historical metric data, indicates anomalous behavior of CIs which events might not capture. Anomaly alerts can be promoted to regular IT alerts and appear on the Service Operations Workspace and service health dashboard for preventive actions.

Anomaly detection

Metric data is collected by various data sources such as SCOM, SolarWinds monitoring system, or Nagios XI server (some partially configured for metric collection by default). These monitoring systems collect metric data from the source environment regularly. Metric Intelligence captures the raw data from these monitoring systems, and uses event rules and the CMDB identification engine to map data to existing CIs and their resources. The data is then analyzed to detect anomalies and to provide other statistical scores.

Metric Intelligence uses historical metric data to build statistical models. These models facilitate projection of expected metric values along with upper and lower bounds. Metric Intelligence then uses these projections to detect statistical outliers and to calculate anomaly scores.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Anomalies are scored on a range of 0-10. High anomaly scores for CI metrics can indicate that a CI is at risk of causing a service outage.

After processing, the Insights Explorer shows metric statistics and charts, and the Anomaly Map shows correlated scores for CIs with the highest anomaly scores, across a timeline.

You may want to disable anomaly detection during system maintenance, as anomalies may be irrelevant when detected while maintenance is in progress. To do so, set the *mid.mi.anomaly_detection.disable* property to **true**.

To customize displaying anomalies for metrics classified as near-constant, contact customer support.

Metric Intelligence is available when you activate the Metric Intelligence (com.snc.sa.metric) plugin.

Terms used with Metric Intelligence

Source metric type

A metric such as '% Free Space' or 'Current Bandwidth' that can be measured by a data source for a CI. For each data source, you can choose which of all possible source metric types are processed. For example, there are about 380 source metric types that are active by default for the SCOM data source.

Anomaly

Data that is outside the control bounds is considered a statistical outlier. These outliers are used to compute an anomaly score, which is a value between 0–10 that indicates the degree to which the metric appears unlikely. When an anomaly score is above a threshold, an anomaly alert is generated. Anomaly alerts are reported separately from regular IT alerts.

Resource

A component of a CI that consists of multiple individual components of similar type, where each subcomponent can be monitored separately. For example, individual Web pages, or specific disks such as 'Disk C:' and 'Disk D:'.

Time series

A series of values (such as metric values) over a time range, associated with a CI and a metric type. Because an anomaly score is evaluated for each metric, the series of anomaly scores over a period of time are also a time series. Time series are computed by the statistical model built for a metric data series. They are used with metric data values, anomaly scores, and upper and lower control bounds.

Statistical models

Metric Intelligence jobs learn from past metric data (up to 32 days old). A model training process analyzes historical data to construct a model that projects future values. Typically, models are in effect until the next time the model learning process runs. These models are used to calculate upper and lower bounds. Incoming values that are beyond those bounds, and that deviate with statistical significance from expected values, generate anomalies. Each model is uniquely patterned and is labeled with a classifier that illustrates the general behavior of the model. This classification determines if anomaly detection can be applied. For most models, it is possible to project which future values deviate from expected values. Such models are associated with control bounds and anomaly detection can be applied (if enabled).

However, for some models, there is insufficient data to determine which values are anomalous. Anomaly detection cannot be applied without additional information (even if anomaly detection is enabled).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The learned data models are stored in the Metric Time Series Models [sa_time_series] table.

The following statistical models and classifiers are used in anomaly detection:

Time Series statistical model

After it is established, a time series model does not adjust to changes in the incoming metric data. Therefore, if the pattern of incoming data changes, those changes are likely to be identified as anomalous. Upper and lower control bounds, after they are learned, persist until the next time the learning process runs (data is learned every day).

Weekly

Data with a pattern that repeats itself over weekly intervals (seasonal model).

Requires a minimum of 15 days of data in the series, as set by the weekly_model_min_days configuration setting.

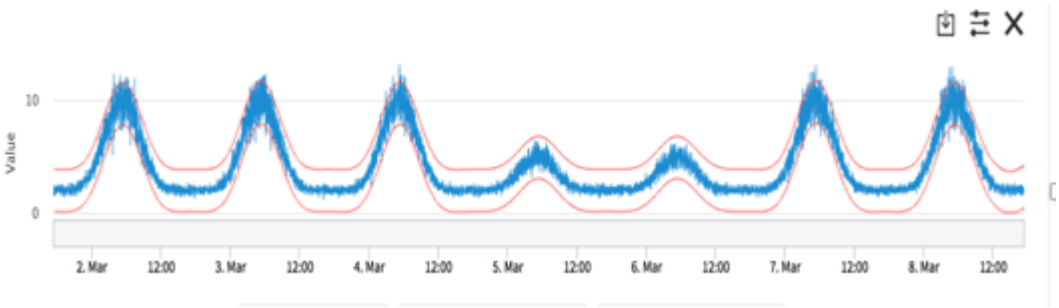

Daily

Data with a pattern that repeats itself over a daily interval (seasonal model).

Requires a minimum of 3 days of data in the series, as set by the daily_model_min_days configuration setting.

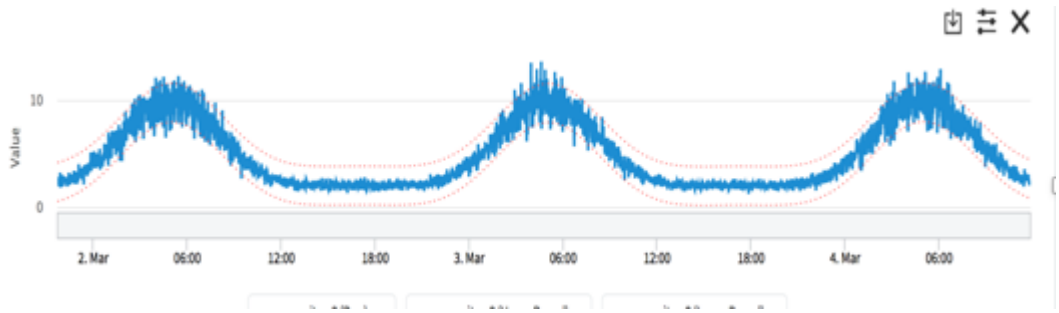

Trendy

Data that has a linear trend with some slope and with some noise.

Requires a minimum of 30 data points in the series, as set by the corrupt_data_count_threshold configuration setting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

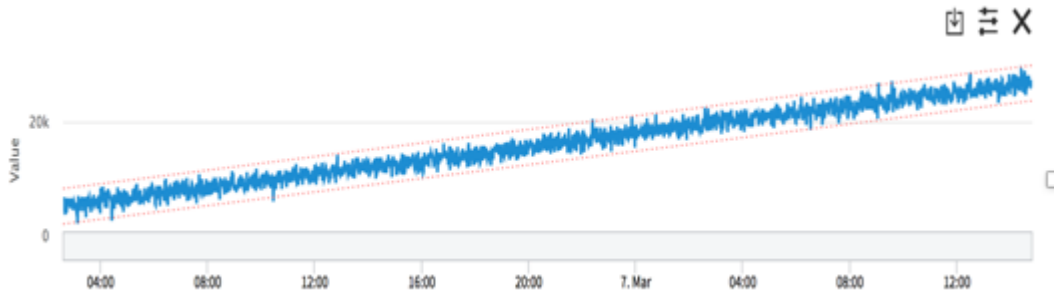

Noisy

Typical noisy data that is a basic pattern classification in a data model. The pattern cannot be identified with a specific trend or seasonality.

Requires a minimum of 30 data points in the series, as set by the corrupt_data_count_threshold configuration setting.

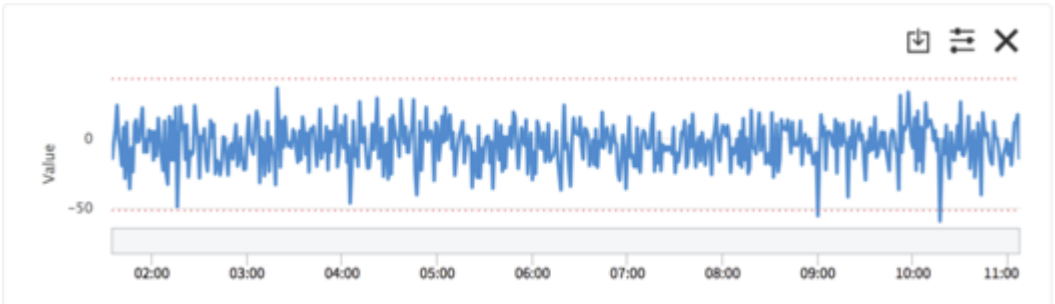

Positive clipped noisy

Similar to the noisy classifier other than the lower bound that is fixed on 0.

Requires a minimum of 30 data points in the series, as set by the corrupt_data_count_threshold configuration setting.

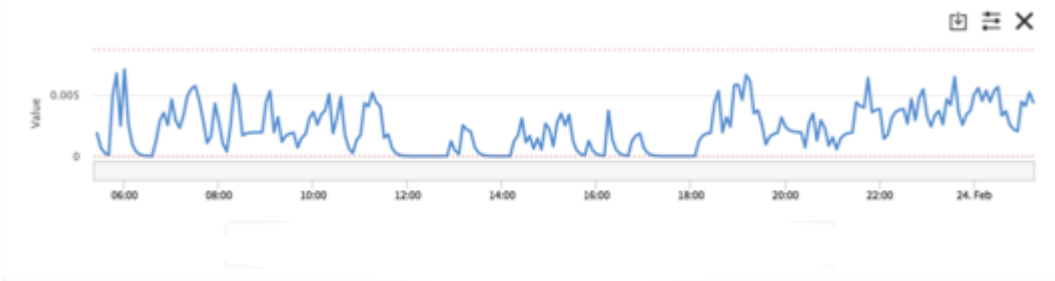

Centered noisy

Noisy data that typically spreads symmetrically between user-specified upper and lower bounds. The formula that is used to set bounds and width values, ignores the statistical data, and the lower and the upper widths have an identical value.

Requires that the number of data points in the series is zero.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To retrieve alerts for metrics with this model, create metric rules for static bounds definitions. For details, see Create metric rules.

Skewed noisy

Noisy data that is not evenly spread between user-specified upper and lower bounds, but instead tends to concentrate closer to one of the bounds. The median of the data is used to separately compute an upper width and a lower width.

Requires a minimum of one data point in the series.

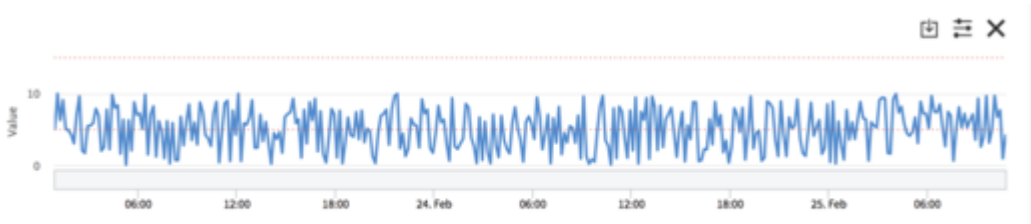

To retrieve alerts for metrics with this model, create metric rules for static bounds definitions. For details, see Create metric rules.

Skewed noisy - Generalized Extreme Value (GEV) Distribution

Noisy data that is spread unevenly between user-specified upper and lower bounds, and concentrates closer to one of the bounds. In addition, the data distribution demonstrates a long or heavy tail. The median of data derived from the tail of the distribution is used to separately compute an upper width and a lower width. There must be at least one data point in the series.

Accumulator

Data pattern similar to the trendy classifier but with a monotonous increase and without noise. For this classifier, there is no data model and no anomaly detection.

Requires a minimum of 30 data points in the series, as set by the corrupt_data_count_threshold configuration setting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

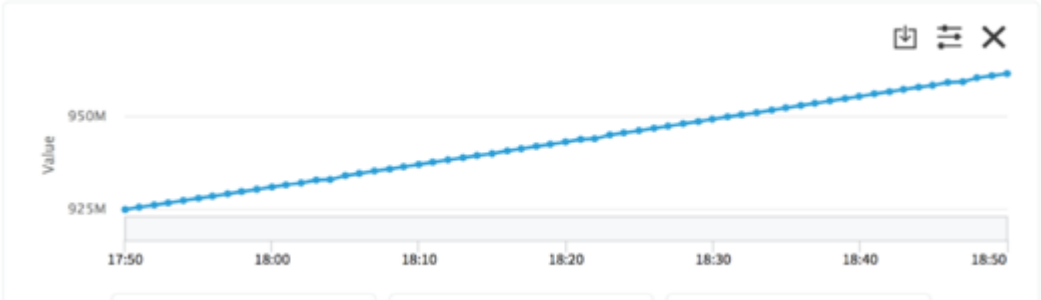

Near Constant

Nearly constant data, in which most values are a specific constant value. For this classifier, there is no data model and no anomaly detection.

Requires a minimum of 30 data points in the series, as set by the corrupt_data_count_threshold configuration setting.

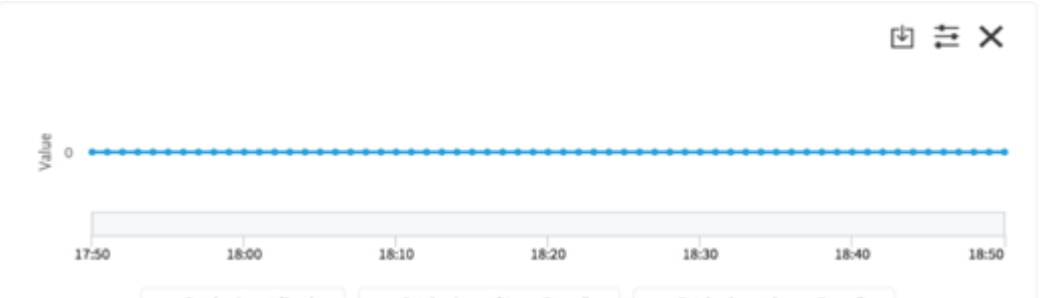

Detected Constant

Enables checking near-constant metrics for anomalies. To do so, you must move the metrics to the detected-constant model, via a system token. To add tokens to the token list and enable more metrics to be measured for anomalies, contact customer support.

Requires a minimum of 30 data points in the series, as set by the corrupt_data_count_threshold configuration setting.

Enabled when the `metric_name_analysis.add_detected_constant_candidate` system property is set to **true**. To overwrite the value of this property for a specific metric type, update the **Detected Constant Candidate** field in the Metric Types table by doing the following:

1. Enter `sa_metric_type.list` in the instance navigation field.
2. In the **Detected Constant Candidate** column, select the relevant value for the entry you want to update.

Detected Constant Candidate values

Value	Description
<empty>	An algorithm automatically decided whether the metric is to be considered for the detected constant model. This is the default value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Value	Description
Yes	The metric is considered for the detected constant model.
No	The metric is not considered for the detected constant model.

Multinomial

Data pattern in which all values are one of a relatively small number of values. For example, values are always 100 or 99.9. For this classifier, there is no data model and no anomaly detection.

Requires a minimum of 400 data points in the series, calculated as 10 times the value of the multinomial_count_threshold configuration setting.

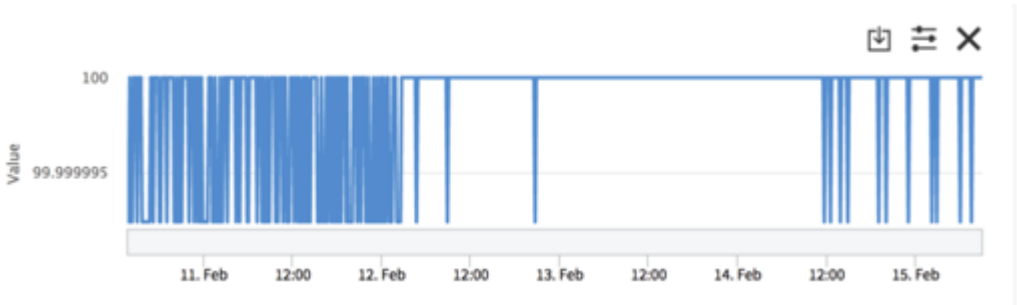

Corrupt

Data has insufficient data points to identify a pattern. For this classifier, there is no data model and no anomaly detection.

Requires that the number of data points in the series is less than the value of the corrupt_data_count_threshold configuration setting (30 by default).

Kalman Filter statistical model

Add on to the time series statistical model and applicable only to the noisy and positive noisy classifiers. This model is a general method of estimating model parameters from a stream of data where level is the only parameter in the model. The Kalman Filter model can adjust to new values in incoming metric data. When there are no clear patterns in the noise or if there is too much noise, the Kalman Filter model is not used.

Local level

When incoming data clusters around a new value according to the current control bounds, the Learner adjusts the data model to accommodate a permanent change. This clustering is detected as a new value in the data model so that most incoming data is again within the control bounds rather than anomalous. Such change detection is useful when for example, cores or memory are added to a server, which impact the baselines.

Requires a minimum of 30 data points in the series, as set by the corrupt_data_count_threshold configuration setting.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

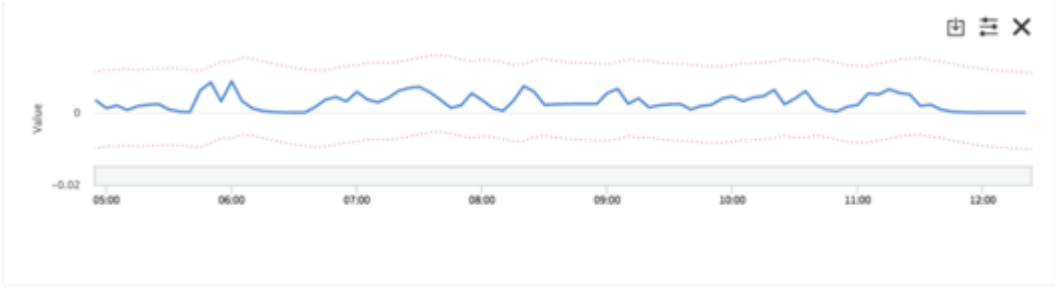

Unrecognized

When data does not fit the local level classifier, time series classifiers are used. This happens when it is not possible to adjust the variance ratio in a learned local level model to reasonable values.

Non-Parametric statistical model

Add on to the positive noisy classifier. In the nonparametric model, noise distribution is not symmetrical and does not fit any seasonal pattern. The nonparametric model creates control bounds that better fit the actual data, and once learned, the control bounds persist until the next learning cycle. This model does not adjust itself to changes in the data, and it takes longer for a deviation to be identified as an anomaly.

Stationary Non-Parametric

Data that is not time-dependent meaning that there is no significant shift in parameters such as mean and variance when shifting data in time.

Requires a minimum of 5000 data points in the series, as set by the snpm_minimum_data_count configuration setting.

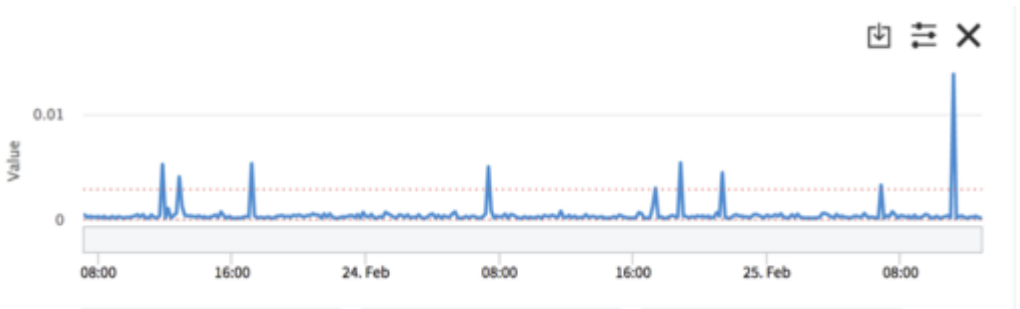

Unrecognized

When data does not fit the stationary classifier, time series classifiers are used.

Median Absolute Deviation (MAD) statistical model

An add-on to the skewed noisy classifier. In this type of data, noise distribution is not symmetrical and does not fit any seasonal pattern. In addition, data reflects a heavy or long-tailed distribution. The MAD statistical model creates control bounds that better fit the data, and once learned, the control bounds persist until the next learning cycle. Using this model enhances the deciphering of data collection with approximately 30% greater efficiency.

MID Server distributed cluster for Metric Intelligence

Using Metric Intelligence requires at least one MID Server distributed cluster which contains a single MID Server that is configured for Metric Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When using a single Metric Intelligence MID Server, that server must be added to a MID Server distributed cluster.

To support the specified throughput, create a distributed cluster with a single MID Server that meets the Metric Intelligence MID Server requirements. When using a virtual machine (VM), only a single MID Server can run that supports Metric Intelligence. When working with Agent Client Collector, a MID Server can also be connected to multiple agents (ACC-M and ACC-L).

A MID Server supporting Discovery can also be used with other applications.

For details on setting up Metric Intelligence, see Get started with Metric Intelligence.

MID Server distributed clusters

Activating Metric Intelligence (com.snc.sa.metric), automatically activates MID Server distributed cluster (com.snc.agent.distributed.cluster). A MID Server distributed cluster supports the Distributed MID Server cluster type which is used exclusively with Metric Intelligence.

If automatic creation of a MID Server distributed cluster fails, you can configure the cluster manually. For details, contact Customer Support.

Metric Intelligence MID Server

The Metric Intelligence MID Server supports the Metric Intelligence feature by processing data and detecting anomalies. The MID Server transmits batches of processed data to the instance. If an anomaly is detected, then information about the anomaly along with the raw data is sent immediately to the instance.

Each Metric Intelligence MID Server must be configured with the following settings:

- Metric Intelligence as a supported application
- The Metrics capability
- Running the Metric Intelligence Metrics extension

Note: The Metric Intelligence MID Server does not support IPv6.

The Metric Intelligence application that the Metric Intelligence MID Server is configured with enables you to add additional supported applications to the same MID Server. By default, Metric Intelligence is included in the definition of the MID Server ALL option. For information about modifying the behavior of the ALL option when selecting supported applications, see Configure applications included in ALL Applications.

If Domain Support - Domain Extensions Installer is activated, then you can configure a MID Server with the Metrics capability, per domain. In this case, metrics for a business service are processed on the MID Server that is in the same domain as the business service. Otherwise, a MID Server from the global domain is used.

Ensure that the MID Server that you want to configure for Metric Intelligence:

- Is validated. For more information, see Validate a MID Server
- Is in running state
- Meets all software, hardware, and configuration requirements

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hardware requirements for the virtual machine/host on which the MID Server runs (scaled for 5000 CIs with 100 metrics per CI)

Requirement	
One MID Server with: • Memory: 12 GB • Processor: 4 cores	Supports approximately up to a total of 1.2 million metrics (12,000 CIs with 100 metrics per CI).
One MID Server with: • Memory: 24 GB • Processor: 8 cores	Supports approximately up to a total of 1.7 million metrics (17,000 CIs with 100 metrics per CI).

Software requirements

Software	Supported versions	Additional requirements
Windows	32-bit and 64-bit versions: • Windows 2008 R2 • Windows Server 2012 R2 • Windows Server 2019 • Windows Server 2022 • Windows Server 2025	32-bit or 64-bit version of the MID Server
Linux	• Red Hat Enterprise Edition Linux 6.6 or later • CentOS Linux 6.6 or later	32-bit or 64-bit version of the MID Server

Configuration requirements

Configuration	Requirement
MID Server service logon user (required only for pulling)	Must be set to a user with read access to the SCOM database (OperationsManagerDW). **1.** On the MID Server, open the MID Server **Properties** dialog box. **2.** Click the **Log On** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuration requirements (continued)

Configuration	Requirement
	3. Select **This account**, and enter credentials for a user with the required access. 4. Click **Apply**.

MID Server throttling mechanism

The MID Server throttling mechanism protects against data bursts to the instance. Data burst protection prevents overloading the instance, while ensuring that you do not lose any data.

The MID Server throttling mechanism is enabled by default. You may want to disable this mechanism when troubleshooting, to ensure that the mechanism is not causing system malfunction.

The throttling parameters are MID Server configuration parameters. Configure these parameters on the MID Server record, on the **Configuration Parameters** tab. When you add a configuration parameter, select the matching `mid.oi` parameter from the **Parameter name** list.

- To disable this feature for metrics, set the `mid.oi.metric_token_bucket_enabled` parameter to `false`.
- To disable this feature for anomalies, set the `mid.oi.anomaly_token_bucket_enabled` parameter to `false`.

Configuration Override

Use the `../agent/conf/wrapper-override.conf` file to configure Java override parameters, such as the garbage collector and IPv4.

To configure the garbage collector, override the existing garbage collector parameter. For example, replace `-XX:+UseParallelGC` with `-XX:+UseG1GC`.

To force IPv4, add the following parameter, using the next available `wrapper.java.additional` number:

```
wrapper.java.additional.<number>=-Djava.net.preferIPv4Stack=true
```

Note: Do not renumber existing `wrapper.java.additional` lines. Parameters 101 through 200 are reserved for ServiceNow out-of-the-box parameters. Only override those parameters by uncommenting the existing lines and changing the values.

Use the `../agent/conf/wrapper.conf` file to configure memory for Metric Intelligence. Update the following parameters:

- `wrapper.java.initmemory=1024` or higher
- `wrapper.java.maxmemory=4096` or higher

Optionally, if a server has additional RAM available, for example 16 GB RAM, you can replace `4096` with `8192` in `wrapper.java.maxmemory=4096`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric Configuration Rules

Metric configuration rules enable you to determine how incoming metrics from the Agent Client Collector are to be processed.

When a new metric arrives from an Agent Client Collector, a new record is created in the Metric to CI Mappings table, with the value, **Anomaly Detection Action Level = Bounds**.

To trigger execution of a metric configuration rule, navigate to the Metric Configuration Rules table (**Event Management** > **Anomaly Detection** > **Metric Config Rules**) and select the **Sync to MID** button at the top of the page. This also synchronizes all MID Servers connected to the instance.

When executing the rule, the anomaly detection action level defined in the rule is set by the *`anomaly_detection_action_level`* parameter. This retroactively overrides the anomaly detection action level configured for the metric in the Metric to CI Mappings table.

If you do not trigger the rule execution, retroactive overriding of the *`Anomaly detection action level`* parameter is performed during execution of the hourly scheduled **Metric Intelligence - Metric configuration** job.

When configuring a metric configuration rule, rules with a filter run before rules with no filter. When multiple rules have the same priority, the most recently configured rule runs first.

For details on creating configuration settings rules, see Create a configuration settings rule.

For details on the available anomaly detection action levels, see Configure the action level for anomaly detection.

Metric bounds sensitivity

Insights Explorer automatically widens metric bounds when too many anomalies are detected, reducing false alerts so you can focus on genuinely important issues.

Reducing the number of anomalies enables you to focus on important issues while eliminating irrelevant anomalies.

The system uses measurement ranges, called 'bounds', to determine what is considered normal and abnormal metric behavior. These bounds are calculated automatically from your historical metric data.

About variation ranges

A variation range is a measurement of how much your metric values typically vary from their average. The system calculates variation ranges automatically by analyzing your historical data. For example, if your CPU usage averages 50%:

- A small variation range (like 2%) indicates that values stay consistent, typically between 48-52%.
- A large variation range (like 10%) indicates that values fluctuate more widely, typically between 40-60%.

The variation range's measurement (value) is known as standard deviation.

How bounds are set

The system sets bounds by measuring how far away from the average a value can be before it's considered abnormal. The formula is: `Average +/- (multiplier x variation range)`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The multiplier is a number that controls how wide or narrow the acceptable range is.

- Smaller multipliers create narrower bounds and detect more anomalies.
- Larger multipliers create wider bounds and detect fewer anomalies.

Example: CPU averages 50% and the standard deviation is 5%

Sample metric bounds

Multiplier	Bounds
2	40% to 60% (50%+/-10%)
3	35% to 65% (50%+/-15%)
5	25% to 75% (50%+/-25%)

Automatic adjustment process

1. High sensitivity detected: If 95% of a CI's metric values fall within a multiplier of 2, the system recognizes this as overly sensitive. As a result, too many anomalies are created, including some that aren't actually problems.
2. The system widens the bounds: When high sensitivity is detected, the system automatically adjusts from a multiplier of 2 to 5, creating wider bounds and generating fewer, more meaningful anomalies.
3. Normal sensitivity: If high sensitivity isn't detected, the system uses a multiplier of 3 for moderate bounds.

Any metric values falling outside the configured bounds are tagged as anomalies.

The multiplier values (2, 3, and 5) are set by default. For details on customizing these values, see Sensitivity bounds properties for Insights Explorer metrics.

Metric collection sources

When setting up Metric Intelligence, you must configure a source for metric collection. Metric collection can be performed either by the Agent Client Collector, or by metric connectors.

The metric collection source determines where Metric Intelligence gathers metrics from. For details on gathering metrics using Agent Client Collector, see Configuring Agent Client Collector Monitoring.

When gathering metrics with connectors, use either a SCOM or Dynatrace connector. For details on connecting metrics using a SCOM connector, see Configure the SCOM connector instance.

For details on connecting metrics using a Dynatrace connector, see Configure the Dynatrace metrics connector instance.

Metric binding

After metric data is collected, Metric Intelligence identifies the CIs and the resources to bind the data to.

The data that is collected on the MID Server is raw and does not relate to any specific CI or resource in the CMDB. To be useful, the data goes through a **normalization** process that uses CMDB identification rules and event rules to uniquely identify CIs. The process maps and binds the identified CIs to the raw data. Raw data can also be mapped and bound to resources, if binding to a CI is successful and if resource binding is configured.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Records for mapping raw data to CIs are automatically generated and remain in effect for a specified length of time determined by the properties:

- *sa.metric.map.with.ci.expiration.sec*: If the mapping to the CI was found. Set by default to be valid for five days.
- *sa.metric.map.without.ci.expiration.sec*: If mapping to the CI was not found. Set by default to be valid for 24 hours.

When similar metric data arrives within that time period, the existing mapping is used to match the data to CIs. At the end of the time period, metric-to-CI records expire. Also, a change in the event rules triggers an immediate expiration of the respective metric-to-CI records. Next time that raw metric data arrives, it will be normalized again. When Discovery adds or removes CIs, mappings are adjusted to reflect these changes at the next cycle.

All the information about the recent existing mapping is saved under metric-to-CI records in the sa_metric_map table, as shown below. You can access the Metric to CI table by navigating to **All > Event Management > Metrics > Metric to CI**.

Metric to CI table

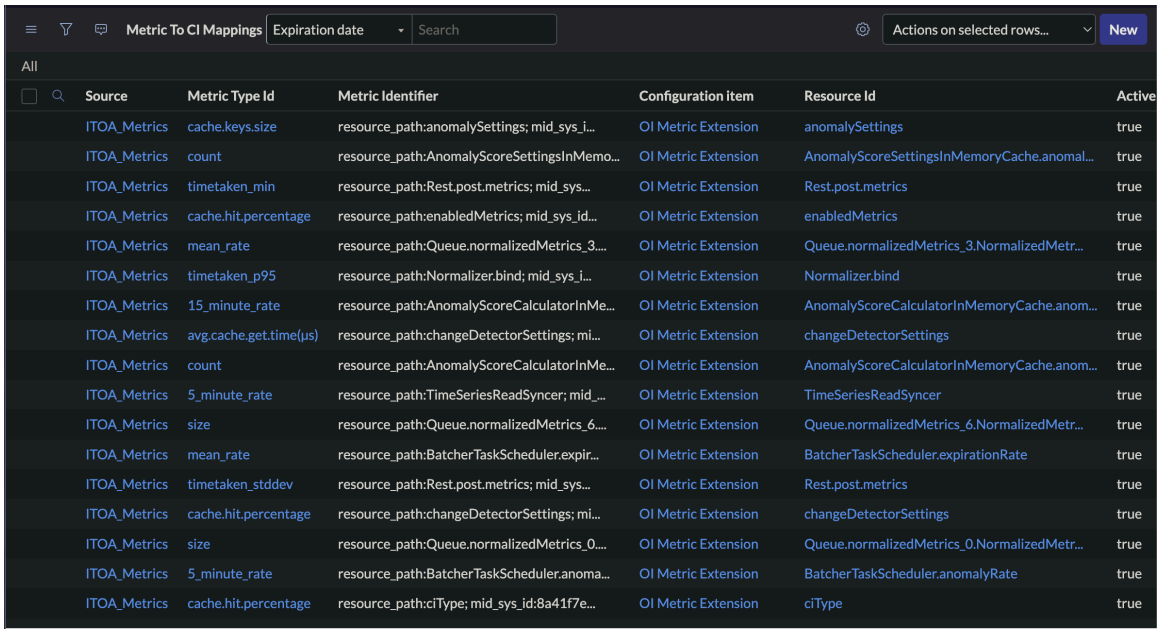

Source	Metric Type Id	Metric Identifier	Configuration item	Resource Id	Active
ITOA_Metrics	cache.keys.size	resource_path:anomalySettings; mid_sys_i...	OI Metric Extension	anomalySettings	true
ITOA_Metrics	count	resource_path:AnomalyScoreSettingsInMemo...	OI Metric Extension	AnomalyScoreSettingsInMemoryCache.anomal...	true
ITOA_Metrics	timetaken_min	resource_path:Rest.post.metrics; mid_sys...	OI Metric Extension	Rest.post.metrics	true
ITOA_Metrics	cache.hit.percentage	resource_path:enabledMetrics; mid_sys_id...	OI Metric Extension	enabledMetrics	true
ITOA_Metrics	mean_rate	resource_path:Queue.normalizedMetrics_3....	OI Metric Extension	Queue.normalizedMetrics_3.NormalizedMetr...	true
ITOA_Metrics	timetaken_p95	resource_path:Normalizer.bind; mid_sys_i...	OI Metric Extension	Normalizer.bind	true
ITOA_Metrics	15_minute_rate	resource_path:AnomalyScoreCalculatorInMe...	OI Metric Extension	AnomalyScoreCalculatorInMemoryCache.anom...	true
ITOA_Metrics	avg.cache.get.time(µs)	resource_path:changeDetectorSettings; mi...	OI Metric Extension	changeDetectorSettings	true
ITOA_Metrics	count	resource_path:AnomalyScoreCalculatorInMe...	OI Metric Extension	AnomalyScoreCalculatorInMemoryCache.anom...	true
ITOA_Metrics	5_minute_rate	resource_path:TimeSeriesReadSyncer; mid_...	OI Metric Extension	TimeSeriesReadSyncer	true
ITOA_Metrics	size	resource_path:Queue.normalizedMetrics_6....	OI Metric Extension	Queue.normalizedMetrics_6.NormalizedMetr...	true
ITOA_Metrics	mean_rate	resource_path:BatcherTaskScheduler.expir...	OI Metric Extension	BatcherTaskScheduler.expirationRate	true
ITOA_Metrics	timetaken_stddev	resource_path:Rest.post.metrics; mid_sys...	OI Metric Extension	Rest.post.metrics	true
ITOA_Metrics	cache.hit.percentage	resource_path:changeDetectorSettings; mi...	OI Metric Extension	changeDetectorSettings	true
ITOA_Metrics	size	resource_path:Queue.normalizedMetrics_0....	OI Metric Extension	Queue.normalizedMetrics_0.NormalizedMetr...	true
ITOA_Metrics	5_minute_rate	resource_path:BatcherTaskScheduler.anoma...	OI Metric Extension	BatcherTaskScheduler.anomalyRate	true
ITOA_Metrics	cache.hit.percentage	resource_path:ciType; mid_sys_id:8a41f7e...	OI Metric Extension	ciType	true

Metric binding to resources

Bind metrics to resources to simplify metric events binding by enabling binding to resources such as specific disks or web pages, in addition to binding to CIs.

Metric Intelligence models metric series in either of the following methods:

- Binding a metric series to a CI and to the metric being monitored for that CI using a 'CI/Metric' format. For example, 'win_server_001/response_time_mean'.
- Binding a metric series to a CI, a resource within that CI, and the metric being monitored for that resource, using a 'CI/Resource/Metric' format. For example, 'win_server_001/Disk — C/disk_usage'. This method is referred to as *resource binding*.

The first method enables modeling, storing, aggregating, and querying data at the CI level. However, resource binding creates resource records for specific monitored entities, such as an individual web page or a disk. Therefore, resource binding is more efficient when there are multiple resources of similar types within a CI, and metrics applicable to a category of these resources are mapped to these resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use cases:

- Common examples are disks, processors, and network interfaces. In cases in which these entities are being monitored but do not exist in the CMDB, using resources for metric binding is useful.
- Some monitoring solutions capture metric data within services such as KPIs for individual web pages in an application. In such cases where the entity being monitored is not a configuration item, metric binding to resources can be helpful.

In those situations, using resource binding results in more meaningful aggregations across similar metrics (such as: avg disk_usage for a host across all disks).

Resource binding process

The MID Server generates metric binding events that are processed by the instance. When processing a metric binding event, there is an attempt to identify the CI to which this metric binding event belongs. If this attempt is successful, then the metric binding event is bound to the identified CI and binding to a resource is attempted. The following steps describe the attempt to bind a metric binding event to a resource:

1. Identify the CI class of the CI that was bound to the metric binding event.
2. Locate the resource class which is mapped to that CI class (using the CI Type To Resource Class [sa_ci_type_to_resource_class] table).
3. Read the *resource_path* attribute value in the **additional_information** field in the metric binding event.
4. Check if a resource record exists in the resource class table. The record must have *name* equal to *resource_path* and *cmdb_ci* equal to the CI that was bound to the metric binding event.
5. If such resource record exists, then the metric binding event is bound to that resource. Otherwise, a new resource record is created with the preceding values and the metric binding event is bound to the newly created resource.

Configure resource binding

- Ensure that the sa.metric.use.resource.binding system property is set to true (default).
- For series intended to be bound to resources, ensure that the *resource_path* attribute is populated by doing either step:
  - Use an event rule to add the *resource_path* attribute to the **Additional information** field in events. For more information, see Creating an event rule to map metrics to specific CIs.
  - Populate the *resource_path* attribute as a part of the respective CI identifier when data is sent to the MID Server for processing.
- Review the default mappings in the CI Type To Resource Class [sa_ci_type_to_resource_class] table and adjust as needed. It is critical that mappings are set as desired prior to data processing.

  Managing the CI Type To Resource Class table requires the evt_mgmt_admin role.

Resource tables

Resource binding uses an underlying hierarchy structure of resources which is a subset of the [cmdb_ci] hierarchy. Metric Intelligence installs the following resource tables:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- CI Resource [ci_resource] (parent table):
- Tables that extend CI Resource [ci_resource]:
  - ci_resource_hardware
  - ci_resource_appl
  - ci_resource_service
  - ci_resource_vm_object
  - ci_resource_database

Mapping CIs to resources

Mappings of CI classes to resource classes are stored in the CI Type To Resource Class [sa_ci_type_to_resource_class] table. This table is installed with Metric Intelligence, and is used during metric binding to resources.

Mappings in the CI Type To Resource Class table (default)

CI class	Resource class
cmdb_ci_hardware	ci_resource_hardware
cmdb_ci_appl	ci_resource_appl
cmdb_ci_service	ci_resource_service
cmdb_ci_database	ci_resource_database
cmdb_ci_vm_object	ci_resource_vm_object

Creating an event rule to map metrics to specific CIs

Create event rules to map incoming raw metric data to specific CIs, to optionally modify metric names, and to populate the *`resource_path`* attribute for resource binding.

The system generates events for every new raw metric data for which there are no CI binding specifications. You can create event rules that filter those events to map the incoming metric data to specific CIs. In such event rules, you can for example define a filter that is based on the monitoring object (such as SCOM).

To use resource binding, the *`resource_path`* attribute of events must be populated. You can configure an event rule to ensure that the *`resource_path`* attribute is included in the **Additional information** field in events. For more information about configuring and using resource binding, see resource binding.

Use the *`resource_path`* attribute to identify the resource within the bound CI. For example, if a metric is collected for the C drive of the `win_server_001` CI, set *`resource_path`* to a value such as `Disk - C`. The metric identifier can then use the `CI/Resource/Metric` format, such as `win_server_001/Disk - C/disk_usage`.

For example, in the Event Rule designer you can update the event rule **SCOM IIS Web Server** as follows:

- On the **Event Filter** tab, define filter to match only IIS Web Servers.
- On the **Binding** tab, set **Class** to `Microsoft IIS Web Server`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Event rules
- View metric to CI and resource binding
- Get started with Metric Intelligence

CIs in maintenance mode

Configure anomaly detection to exclude metrics for CIs that are in maintenance mode from model learning.

You can change the settings of system properties to include or exclude metric data for CIs while they are in maintenance mode.

When the system property sa.model_learner.maint_event_record_history_enabled is set to true (default), Metric Intelligence stores historical information about the times CIs enter and exit maintenance mode. When the system property sa_metric.maint_exclusion is set to true (default), metrics from CIs that are in maintenance mode are excluded from model learning.

Anomaly model testing

Use anomaly model testing to apply and evaluate anomaly detection for a small set of CIs and metrics, using actual metric data. Compare test results to expected results, then fine-tune the anomaly detection model before enabling anomaly detection for the tested CIs and metrics in the production environment.

For example, before enabling anomaly detection for Linux Servers, you can use anomaly model testing with actual metric data collected for those Linux Servers. Existing alerts might indicate an anomaly for a specific metric and CI. You can create an anomaly test rule to test anomaly detection for those CIs and metrics, and then check if the expected anomaly is detected. You can also examine the bounds calculated by the anomaly model to see if they match expected bounds. After you fine-tune the test model, you can enable anomaly detection for Linux Servers.

Anomaly model testing supports up to 20 time series, which can include 20 metric types for the same CI, the same metric type for 20 different CIs, or any combination of those options.

View test results in the Insights Explorer

Upon completion of a test run, a URL to an Insights Explorer becomes available. The anomaly model testing results appear in this trimmed-down version of the Insights Explorer. This Insights Explorer is pre-loaded with the charts and scores for the metric series specified in the anomaly test rule. You cannot add any CMDB groups or application services to this Insights Explorer.

Each chart in the Insights Explorer, contains the Anomaly Test Statistics section in its **Chart Settings**. This section is specific to anomaly model testing and provides the following options:

- **Show Bounds**: Displays a computed upper and lower bound for the metric in an anomaly model testing, based on learning of past metric values. The upper and lower bounds always appear together.
- **Show Anomaly Scores**: Displays anomaly scores on a virtual axis of 0-10, for an anomaly model test. Color code is based on the score thresholds that are defined in the Anomaly Score to Event Severity Map [sa_metric_anomaly_score_to_event_severity_map] table.

Advanced Promotion Engine

You can create a definition for the Advanced Promotion Engine so that you can define the conditions for promoting the anomaly alerts to the All Alerts table. By promoting the alerts that meet the conditions, only the most relevant anomaly alerts move to the table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Advanced Promotion Engine moves the most relevant alerts from the list of anomaly alerts to the All Alerts table. When a group of anomaly alerts on a configuration item (CI) meets the conditions that you set in the definition that you created for the Advanced Promotion Engine, the grouping mechanism groups the promoted alerts in the All Alerts table. For example, you can create a condition indicating that only the anomaly alerts that belong to a specific CI class with a minimum severity of Major are to be promoted to the All Alerts table.

The All Alerts table is the table that contains the issues that need your immediate attention. The anomaly alerts on the different metrics within a CI in a specific time period indicate an issue that needs to be promoted to the All Alerts table.

Anomaly alerts indicate a deviation from projected metric values for monitored CIs. For details, see View anomaly alerts.

The Advanced Promotion Engine is activated automatically as part of the Metric Intelligence application.

You can use the Advanced Promotion Engine definitions that come with the base system (for the Linux server and for the Windows server), or you can create definitions on the different CI classes. The base system definitions indicate that the CIs on the Linux and Windows servers are operational, and those CIs that have a minimum severity of Major are to be promoted to the All Alerts table.

All Advanced Promotion Engine tasks require the evt_mgmt_admin role. The evt_mgmt_user role enables you to view definitions and properties only.

Metric Explorer

Metric Explorer provides an easy-to-navigate interface where service agents can view the health of a CI associated with an alert. Health details for a CI include various metric charts with control bounds, and aggregations in a time period.

The set of metric charts that appear for a CI class in Metric Explorer is determined by metric chart configurations in the Metric Dashboard Metadata [sa_metric_dashboard_metadata] table. The Metric Dashboard Metadata table is initially populated with configurations for common classes and important metrics for those classes. A user with the admin or evt_mgmt_admin role can update the default metric chart configurations in the Metric Dashboard Metadata table.

There can be up to 10 active metric charts configured for a class at any point.

Chart types

A metric chart that appears in the metric charts form in Metric Explorer can be a **time-series** chart or a **single-score** chart.

Time-series chart

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

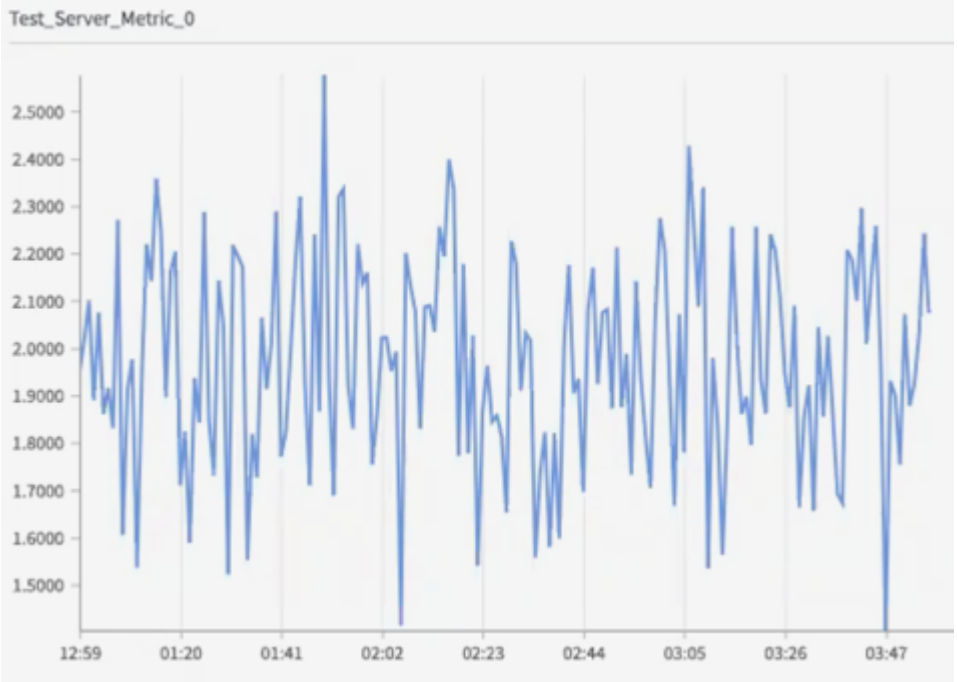

The time-series chart type is a graph of all data points in the specified time range, for the metrics of the CI associated with the chart. Time-series charts are similar to charts in Insights Explorer, and depending on settings, show lower and upper bounds.

Single-score chart

The single-score chart type shows a single number which is an aggregation of the data points in the time series for the metric and CI of the chart. The aggregation is an average of the data points, or the maximum, minimum, or last data point in the time series.

Templates

The ITOM Service Operations Workspace Metric Explorer template enables metric chart drill-down from Digital Analytics (DA) dashboards to Metric Explorer. For details, see Express list overview on the ServiceNow® Developer portal.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Self-health monitoring for Metric Intelligence

Use Event Management self-health monitors to monitor the health of Metric Intelligence infrastructure components and processes, and to alert about potential issues. Self-health monitoring allows you to proactively remediate issues and minimize data loss.

Metric Intelligence is pre-configured to use Event Management self-health monitors to ensure that essential components and processes are functioning properly. If a monitor detects a failure, then the Event Management self-health monitoring system creates an event for the issue. These events are then converted into an Event Management alert or get appended to a previous alert. Metric Intelligence self-health monitors run every 60 seconds.

For example, if an Metric Intelligence MID Server distributed cluster experiences an issue, an event is generated with the status of **Major**. Then, this event is converted into an alert. Later, when the issue is resolved, a new event is generated with the status of **Info** and the original alert is closed.

Monitored components and processes

Component/Process Check	Error Condition
Metric Intelligence extension is running properly	Monitor checks the *status* field of the *ecc_agent_ext_context_metric* table to detect a state of **Error** or **Warning**.
Metric Intelligence MID Server distributed cluster is running properly	Monitor checks the *status* field of the *ecc_agent_cluster* table to detect a state of **Error**, **Segmented**, or**Stopped**. **Note:** For a state of **Stopped**, an event is created only if the *error_message* field is not empty.
Metrics are collected by active connectors that are configured for metrics collection	For each connector with the *Metrics collection* flag set to **true**: Monitor checks the status of the *last_kpi_status* field in the *em_connector_instance* table to detect a status of **Error**.
Essential Metric Intelligence scheduled jobs are active	Monitor checks the following essential scheduled jobs, to detect if any is inactive: • Metric Intelligence — Metric configuration job • Metric Intelligence — Sync tables with mid • Metric Intelligence — Validate MID Distributed cluster

For more information about enabling or disabling an Metric Intelligence self-health monitor, see Configure a self-health monitor.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

Self-health monitors for Event Management

Scheduled jobs installed with Metric Intelligence

Domain separation and Metric Intelligence

Domain separation is supported in Metric Intelligence. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

How domain separation works in Metric Intelligence

When domains are separated in Metric Intelligence, users can view and manage alerts and events only in their own (tenant) domain. After the binding process has finished, the domain of the configuration item (CI) is used so that users within that domain can explore metric data. A domain column is present for Metric Intelligence tables that are provided in the base system. The domain column shows the name of the domain to which the CI belongs.

Note: If the domain column does not appear in the list, click (Personalize List) and add the required column. In addition, you can add a column that displays the domain path.

The following Metric Intelligence features have limited domain separation support.

Feature	Support
Metric types	Supported at global level. The instance owner can control the metric types for all domains.
Metric type registration	Supported at global level. The instance owner can control the metric type registration for all domains.
Metric class	Supported at global level.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Feature	Support
	The instance owner can control the metric classes that are used for all domains.
Configuration settings rules	Supported at global level. The instance owner can control the configuration settings rules for all domains

Related topics

Domain separation for service providers

Configuring Metric Intelligence

Plan and configure your Metric Intelligence implementation.

Configuration overview

Metric Intelligence configuration typically follows these steps:

1. Install the AIOps Experience plugin [sn_sow_aiops].
2. Manually configure a MID Server for Metric Intelligence.
3. Activate Metric Intelligence.
4. Stream metrics to the instance, via the relevant data source or connector.
5. Choose and configure metrics to monitor.
6. Ensure that metrics are properly bound to CIs.
7. Configure the action level for anomaly detection.
8. Define an advanced promotion rule, to determine the conditions for promoting an anomaly alert to an IT alert.

Get started with Metric Intelligence

Complete the following setup and initial configuration steps to start using Metric Intelligence.

Before you begin

Activate the Metric Intelligence plugin (com.snc.sa.metric).

Role required: evt_mgmt_admin

About this task

Metrics have an expiration date. You can continue receiving data without interruption when a metric's expiration date arrives by pushing ahead the expiration date. For details, contact customer support.

Metric Intelligence guided setup provides a sequence of tasks that help you configure Metric Intelligence on your ServiceNow instance. To open Metric Intelligence guided setup, navigate to **Guided Setup** > **ITOM Guided Setup**, and click **Continue**. Then scroll to **Metric Intelligence**, and click **Continue**. For more information about using the guided setup interface, see Using guided setup.

Alternatively, you can complete the following procedure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Configure a MID Server for Metric Intelligence.
   For details, see MID Server and MID Server distributed cluster for Metric Intelligence.
2. Use the Automated Metric Intelligence setup.
3. Choose a source for collecting metrics.
4. Choose and configure metrics to monitor.
5. Create event rules to map raw metric data to specific CIs and Create an event rule to bind metric events to host CIs.
6. About 10 minutes after the Metric Intelligence Metrics extension has started, run the quick start tests for Metric Intelligence to verify functionality.

Activating Metric Intelligence

Metric Intelligence enables you to capture, explore and analyze operational metrics data, identifying and indicating anomalies. This analysis can prevent potential service outages.

View and work with Metric Intelligence data in the Insight Explorer, Agent Workspace, and Service Operations Workspace interfaces.

For details on the plugins activated with Metric Intelligence, see Plugins or applications installed with ITOM AIOps.

Automated Metric Intelligence setup

When you initially set up Metric Intelligence, you can automate many of the setup steps. For example, the automated setup configures a MID Server for Metric Intelligence and creates a MID Server distributed cluster which contains that MID Server.

Before you begin

Role required: evt_mgmt_admin

About this task

When the following conditions are met, the related link **Setup Metric Intelligence** appears on the MID Server form, enabling automated setup:

- The MID Server has been validated.
- The MID Server is in a running state.
- The property *sa.metric.display.mid.setup.action* is set to true (default value).

The Metric Intelligence automated setup performs the following setup tasks:

1. Configures the MID Server with the *Operation Intelligence* application
2. Adds the *Metrics* capability to the MID Server
3. Creates a new MID Server cluster of type 'distributed' set with a default port of 47100 for Apache Ignite
4. Adds the MID Server to the new MID Server distributed cluster
5. Creates a Metric Intelligence Metrics Extension and starts it
6. Opens the Metric Intelligence Metrics Context form. On this form, you can click the **Setup REST Listener** UI action to create the MID Web Server extension configured with the 'Enable REST Listener' option and default port 8097

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If automatic setup fails, you can create a MID Server distributed cluster manually, as described in the Use the **Setup Metric Intelligence** link only to create the first MID Server distributed cluster that contains a single MID Server, with default values. After the automated setup completes, you have to manually create any additional MID Server distributed clusters or add Metric Intelligence MID Servers.

Procedure

1. Navigate to **All** > **MID Server** > **Servers**.
2. Select the MID Server that you want to configure as a Metric Intelligence MID Server.
3. On the MID Server form, select **Setup Metric Intelligence** in the Related Links section.
4. In the newly created Metric Intelligence Metrics Extension form:
   a. Wait for the Status value to be **Started**.
   b. If configuring the MID Server for pushing, select the UI action **Setup REST Listener**.
   A notification appears about the MID Server extension being created. The **MID Web Server Extension** field appears and **Enable REST Listener** is checked. You can preview the MID Server Extension record to verify that its Status value is **Started**.

Related topics

MID Server distributed cluster for Metric Intelligence

Manually configure a MID Server for Metric Intelligence

To use Metric Intelligence, configure at least one MID Server with Metric Intelligence as a supported application, with the Metrics capability, and which runs the Metric Intelligence Metrics extension. Then, add that Metric Intelligence MID Server as a member to a MID Server distributed cluster.

Before you begin

Ensure that the MID Server meets all software, hardware, and configuration requirements to be configured for Metric Intelligence.

Role required: To access the MID Server - mid_server. To configure a MID Server in an instance (for example, to add a supported application), refer to the MID Server documentation.

Procedure

1. Navigate to **All** > **MID Server** > **Servers**.
2. Click the MID Server that you want to configure for Metric Intelligence.
3. Add the Metric Intelligence application:
   a. At the center of the MID Server form, click **Supported Applications**.
   b. In the **Supported Applications** section, click **Edit**.
   c. In the slushbucket, select **OperationalIntelligence** and click the **>** add button.
   d. Click **Save**.
4. Add the Metrics capability:
   a. At the center of the MID Server form, click **Capabilities**.
   b. In the **Capabilities** section, click **Edit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. In the slushbucket, select **Metrics** and click the '>' add button.

d. Click **Save**.

5. Click **Update**.

What to do next

Create a MID Server distributed cluster, and add the Metric Intelligence MID Server as a member to that cluster.

Manually configure the Metric Intelligence extension

Configure the MID Server Metric Intelligence extension to enable the MID Server to pull raw metrics from external systems. The extension detects anomalies and reports them to the instance along with raw data. This MID Server Metric Intelligence extension is required and must be running in order for your system to collect Metric Intelligence data.

Before you begin

See Get started with Metric Intelligence for the high level steps of setting up Metric Intelligence.

Role required: evt_mgmt_admin

About this task

The MID Server Metric Intelligence extension normalizes the raw data and then transmits the data to the instance. If the extension detected any anomalies above a specified threshold, they are sent to the instance. The extension runs for as long as it is enabled. This provides a persistent connection to the MID Server to constantly listen for raw Metric Intelligence data from external systems.

The MID Server Metric Intelligence extension:

- Receives raw metric data, batches them and sends them to the instance at specified intervals.
- Detects anomalies and sends a report to the instance.
- The information that is sent to the ServiceNow instance is relevant for Metric Intelligence, so data processing is efficient.

The MID Server Metric Intelligence extension does not provide any API calls. However, when the **Enable REST Listener** option is selected, the extension adds a handler for the supported REST APIs.

Note:

- Only upon configuring the Metric Intelligence extension, the Metric Intelligence MID Servers actually join the distributed cluster for Metric Intelligence.
- After the initial configuration, the first metric is not included in the metrics data.
- There is a delay of one minute in receiving metric information from the synchronization of the instance with the MID Server.

Procedure

1. Navigate to **All** > **MID Server** > **Profiles and Deployments** > **Metric Intelligence**.
2. In the Metric Intelligence Metrics Contexts list, click **New**.
3. Fill in the fields, as appropriate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Name	A unique name for this extension for easy identification.
Short description	A description of this extension.
Status	This field is auto-populated with the status of the extension. The field is blank until the extension is started. After issuing a command to the extension, one of the following values is displayed: ○ **Started**: The extension is running. ○ **Stopped**: The extension is not running. ○ **Offline**: The MID Server is down. ○ **Error**: The extension failed with an error (the error message is displayed in **Error Message**). ○ **Warning**: A run-time exception has occurred. The extension continues to work.
Error Message	Message describing any error that causes a command, such as **Start** or **Stop**, to fail. This field appears when the value in the **Status** field is **Error** and also when the value in the **Status** field is **Warning**.
Execute on	Location for running this extension. When configuring this option for use with Metric Intelligence, set to **Specific MID Server**. Then specify the MID Server with *`Operationalintelligence`* as a supported application and the *`Metrics`* capability.
MID Server	The name of the designated MID Server, which is configured with *`Operationalintelligence`* as a supported application and the *`Metrics`* capability.
Executing on	[Read-Only] The name of the MID Server on which the extension is running. This field shows the name of the MID Server even if the MID Server is down. If the user stops the extension, this field is empty.

4. When using the Push method for collecting Metric Intelligence data, the MID Server Metric Intelligence extension must be configured with the **Enable REST Listener** enabled. This option enables a listener so that a REST endpoint can receive raw metric data. The raw metric data is then placed in the regular data flow where the data is sent to the instance and the anomaly detector looks for anomalies. When selected, it adds a handler to the web server to listen for any metrics that are pushed to the MID Server. When this option is selected, the Web Server

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

extension, which starts a Web Server on the MID Server, must also be configured. For more information, see Configure the MID Web Server extension.

5. Right-click the form heading and select **Save**.
6. Under **Related Links** click **Start** to save the Metric Intelligence data in this extension and start the extension.

Commands available in the MID Server Metric Intelligence extension

Related Link	Description
Start	Starts the extension on the configured MID Server if it is currently not running.
Stop	Stops the running extension on the configured MID Server. No action is taken if the extension is not running.
Restart	Stops, then starts the extension on the configured MID Server.
Test	The test is not relevant to Metric Intelligence. Parameters are not tested or validated when **Test** is run.
Update parameters	Sends the latest saved parameters to the extension.

Activate the Advanced Promotion Engine

Before you can use the Advanced Promotion Engine feature, you must activate it from the list of plugins.

Before you begin

Role required: evt_mgmt_admin

About this task

To activate the *var.adv-promo-eng*, you also need to install a plugin for ACC-M, which includes Metric Intelligence plugin, as described in Exploring Agent Client Collector Monitoring.

If you don't plan to use ACC-M, install Metrics Intelligence plugin only, as described in Get started with Metric Intelligence.

Procedure

1. Navigate to **System Definition** > **Plugins**.
2. Locate **Advanced Promotion Engine** (id: sn_metric_promote).
3. Select the entry and review the app details on the **Store Application Advanced Promotion Engine** form.
4. Select **Activate/Upgrade** in the Related Links section.
5. Select **Activate**.

Configure domain separation for Metric Intelligence

Configure Metric Intelligence for domain separation to create logically defined domains that limit unauthorized access to data.

Before you begin

Role required: evt_mgmt_admin and evt_mgmt_integration

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Activate the Domain Support – Domain Extension Installer plug-in.
2. Configure the MID Server for Operational Intelligence as a user that belongs to the desired domain with the role 'mid_server' assigned to it.
3. Configure a MID Server distributed cluster in the domain you desire and add Operational Intelligence MID Servers that belong to the same domain.
   For more information about setting up MID Servers for Metric Intelligence , see MID Server distributed cluster for Metric Intelligence.

Note: The Event Management's binding events are configured using the domain of the logged-in MID Server user. Once the process has finished, the record's domain is updated to use the domain from the CI so that any user within that domain can see the metric data.

Optimizing Metric Intelligence

You can perform a variety of tasks to optimize your use of Metric Intelligence.

Create an anomaly test rule

To use the anomaly model testing, create an anomaly test rule in which you specify up to 20 metric series to test anomaly detection for. Run the anomaly test, and after it completes, use the provided URL to open the Insights Explorer which is pre-loaded with the model testing results.

Before you begin

Role required: evt_mgmt_user

Procedure

1. Navigate to **All** > **Event Management** > **Anomaly Detection** > **Anomaly Model Testing**.
2. Select an existing anomaly test rule to edit and run again, or click **New** and fill out the Anomaly Test Rules form.

Field	Description
Name	Name for the model test rule.
Applies to	The table from which to select CI/metric pairs for the model test rule. The table used is Metric To CI Mapping [sa_metric_map].
Rule	Filter conditions to specify pairs of CIs/metrics for the anomaly model testing. If the rule yields more than 20 pairs of CIs/metrics, then the list is truncated to 20.
Start date/End data	The time period for which to apply the anomaly detection testing. The time period must be within the past 7 days.
Status	Status of the test, automatically updated by the system during a test run. Status is initially Scheduled and eventually changes to Complete.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
URL	URL to the Insights Explorer in which the anomaly model test results are pre-loaded.
Notes	Logging summary, automatically updated by the system with details of each test run of the anomaly test rule.

3. Click **Run Anomaly Test**.
   During the test run, the **Status**, **URL**, and **Notes** fields are automatically updated to reflect the progression of the job.
4. After **Status** changes to Complete, click the **URL** link to open the Insights Explorer to see the results of the test.

Result

The bounds and the chart for the anomaly model test, appear in faint colors in Insights Explorer. Also, the upper and lower bounds that are calculated by the test, use a line pattern that is different than line pattern used for actual anomaly detection results. These differences let you distinguish between actual anomaly detection and anomaly model test results, if both are selected in the chart settings.

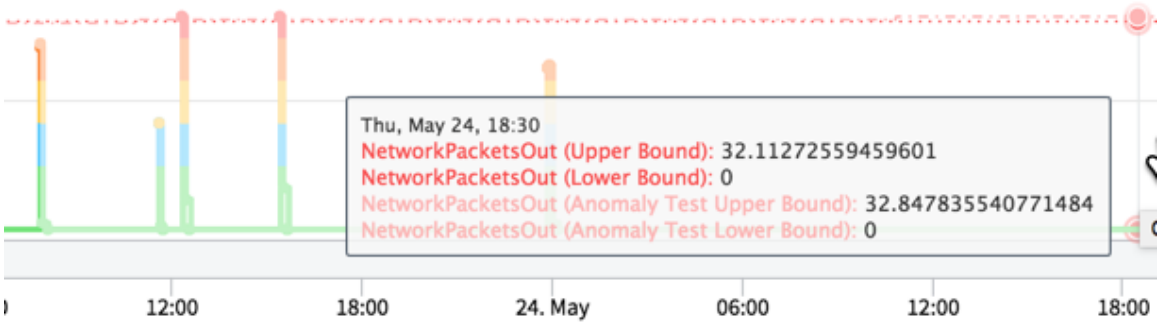

What to do next

- Modify the date range for the test or choose different time series, and click **Run Anomaly Test** again to continue and fine-tune the model.
- Click the **Chart Settings** icon for a chart to toggle the display of statistics and aggregations on the chart. Select **Show Bounds** and **Show Anomaly Scores** in the Anomaly Test Statistics and in the Anomaly Statistics sections, to display test results next to actual anomaly detection results.

  Note: Anomaly Model Testing results might be very similar to actual anomaly detection results. If you configure the chart to display both at the same time, as you move the mouse over the chart, it might be difficult to identify points of differences.

- When you are satisfied with the anomaly test results, enable anomaly detection in the production environment, for the metrics and CIs specified in the test. See Choose and configure metrics to monitor for more details.

Related topics

View metric values in the Insights Explorer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View metric to CI and resource binding

View the metric to CI and resource binding results, including details for failed bindings which you can use to mitigate the failure. Unmapped metrics are excluded from anomaly detection until properly mapped.

Before you begin

Role required: evt_mgmt_admin

If resource binding is not enabled, then no resource binding results appear.

About this task

Metric Intelligence uses event rules to map metric data with the specific CI or resource that is associated with the metric event. The results of these mappings are stored in the Metric to CI Mappings [sa_metric_map] table. If mapping fails, then the binding status for the metric indicates the general reason for the failure and further details about the failure are provided. Examine the details for a mapping failure and mitigate the problem to improve accurate and effective metric data processing.

- If mapping to a CI is successful, then the mapped CI appears in the **Configuration Item** column. If resource binding was attempted but failed, then the mapped CI does not appear, even if the mapping to a CI was successful. Missing or duplicate CIs can cause CI mapping to fail.
- If mapping to a resource is successful, then the mapped resource appears in the **Resource** column. Resource binding can fail if a resource table mapped to the CI type is missing or if the *resource_path* attribute is missing in the event.

Records in the CI Mappings [sa_metric_map] table remain in effect for 24 hours (by default, if mapping to the CI is not successful), or for 5 days (by default, if mapping to the CI is successful). Later, if within that time period raw data arrives for a metric/CI pair that already has a record, the existing mapping is used to match the data to an existing CI. After Metric to CI Mappings records expire, incoming new raw data requires remapping. These records expire when:

- The 24-hour or 5-day cycle ends.
- An event rule has changed — Triggering an immediate expiration of the respective mapping record.

Adjusting mappings to reflect the addition or removal of CIs by Discovery, takes effect only upon the beginning of the next cycle.

Procedure

1. Navigate to **All** > **Event Management** > **Anomaly Detection** > **Metric to CI**.
   The **Metrics to CI Mappings** page appears.

Column	Description
Source	Data source for the metric.
Metric Type Id	Metric that is measured for the CI, such as disk% and cpu%.
Metric Identifier	A string in the form of CI/resource/metric that uniquely identifies a CI and a metric.
Configuration Item	The CI in the CMDB that is mapped to a metric. A missing value indicates that:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Column	Description
	○ The mapping of the metric to a CI was unsuccessful. ○ Resource binding is enabled, but has failed. **Configuration Item** is missing even though binding to a CI was successful.
Resource	The resource that is mapped to a metric. If resource binding is enabled, then a missing value indicates that the mapping of the metric to a resource was unsuccessful.
Expiration Date	Record expiration date.
Anomaly Detection Action Level	The level of anomaly detection processing and analysis applied to CIs and metrics. The available options are: ○ **Metrics Only**: Only gather metrics from the data source without any further processing. ○ **Bounds**: Includes the Metrics Only processing level, and also creates a statistical model and shows bounds in Insights Explorer. ○ **Anomaly Scores**: Includes the Bounds processing level, and also calculates and displays anomaly scores. ○ **Anomaly Alerts**: Includes the Anomaly Scores processing level, and also creates anomaly alerts, as applicable. ○ **IT Alerts**: Includes the Anomaly Alerts processing level, and also creates IT alerts which are based on anomaly alerts, as applicable.
Binding Status	Status of the mapping of the metric to a CI or resource. A **CI Found** status indicates that the mapping was successful, in which case the mapped CI appears in the **Configuration Item** column. Any other status indicates that the mapping failed, and a general reason for the failure is provided. Further details about the failure are provided in the **Binding Failure Reason** column.
Binding Failure Reason	Specific details about the mapping that failed, such as class and host names, and the queries that were used. Click the respective **Event** link to see further details in the **Processing Notes** field.

2. Examine the **Binding Status** and **Binding Failure Reason** as appropriate.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the details provided to remediate the problem:

Binding Status	Remediation
Host CI Missing	Run Discovery or use any other method to populate the CMDB with the missing host CI.
Non-Host CI Missing	Run Discovery or use any other method to populate the CMDB with the missing non-host CI.
Non-Host CI Duplicate	Use the Duplicate CI Remediator to remediate the duplication.
Dependent CI Duplicate	Use the Duplicate CI Remediator to remediate the duplication.
Dependent CI Missing	Run Discovery or use any other method to populate the CMDB with the missing dependent CI.
Duplicate Metric Map Entry Exists	Check **Binding Failure Reason** and compare the metric identifiers associated with the duplicate metric map entries. If metric identifiers have changed (for example, an identifier was added), set the entry for this series in the CI Mappings [sa_metric_map] table to expire. You can also delete the entry. Also check the associated event rule and ensure that it extracts the attributes correctly and that it does not extract the same CI for two different metric identifiers. This error might appear when 2 different metric identifiers map to the same CI/metric.
Resource ID Not Found	Verify that the CI Type To Resource Class [sa_ci_type_to_resource_class] table contains a mapping record for the respective CI type, or for one of its parents. Then ensure that a table with the name of resource class for this record, exists.

What to do next

Group **Binding Failure Reason** by **Binding Status** to create a summary report with the number of records per failure category. Use this report to efficiently remediate the problems.

Related topics

Creating an event rule to map metrics to specific CIs

Disable a metric for a CI

Metric Intelligence can be configured to collect significant amounts of data, some of which might not be necessary. To improve performance, you can disable a specific metric for a specific CI to stop processing data related to the specified metric and CI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: evt_mgmt_admin

About this task

Note: The *Active* setting (true or false) for a source metric in the Monitoring System Metric Types [sa_source_metric_type] table takes precedence over the setting for the corresponding metric in the Metric To CI Mappings [sa_metric_map] table. If a source metric type in the Monitoring System Metric Types [sa_source_metric_type] table is disabled, all records related to the corresponding metrics are removed from the Metric To CI Mappings [sa_metric_map] table.

Procedure

1. Navigate to **All** > **Event Management** > **Anomaly Detection** > **Metric to CI**.
2. Locate the record for the metric/CI that you want to disable, and set its **Active** column to **false**.

Disable event collection

For performance reasons, you might want to disable the collection of events from a data source from which metric data is also collected. You can disable event collection for data sources such as Nagios XI server, SolarWinds monitoring system, and Zabbix server.

Before you begin
Role required: evt_mgmt_admin

Note: To restart event collection after it has been disabled, you will need to restore the connector's original script.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Instances**.
2. Select a connector instance for the data source for which to disable event collection.
   For example, select a connector instance for Nagios, SolarWinds, or Zabbix.
3. Set **Event collection schedule (seconds)** to a large value such as 20,000,000.
   Such setting avoids unnecessary communication with the MID Server, optimizing performance.
4. Click **Update**.
5. Navigate to **MID Server** > **Script Includes** and extend the connector's script (such as *NagiosMetrics_JS*, *SolarWindsJS*, or *ZabbixJS*) by overriding the *execute* method.
   Add a code snippet such as the following, and then click **Update**:

Example: Snippet for Nagios

```
var NagiosJSCustom = Class.create();

// Extending Nagios default script
NagiosJSCustom.prototype = Object.extendsObject(NagiosJS, {

// Override execute method
                    execute: function() {
        var retVal = {};
        retVal['status'] = SUCCESS.toString();
       retVal['error_message'] = "Disabled by custom Class,
 'Last event collection status' will be 'error' and this is
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
 normal. To achieve success status, enable event pull in custom
 class(NagiosJSCustom)";

        return retVal;
    });
```

For the SolarWinds connector use:

```
var SolarWindsJSCustom = Class.create();

// Extending solarwinds default script
SolarWindsJSCustom.prototype =
 Object.extendsObject(SolarWindsJS, {
```

For the Zabbix connector use:

```
var ZabbixJSCustom = Class.create();

// Extending zabbix default script
ZabbixJSCustom.prototype = Object.extendsObject(ZabbixJS, {
```

Configure anomaly score thresholds

Metric Intelligence scores anomalies on a range 0-10. This range is broken down to the five levels of event severities, each represented by a different color in the Insights Explorer and in the Anomaly Map. You can configure the anomaly score threshold for each level of severity.

Before you begin

Role required: evt_mgmt_admin

About this task

Anomaly score thresholds are defined in the table Anomaly Score to Event Severity Map [sa_metric_anomaly_score_to_event_severity_map]. For example, by default, anomaly scores from 8 through 9 are severity 2 (Major). These anomalies are displayed in dark orange in the Insights Explorer and in the Anomaly Map.

For information on the event severities that each number corresponds to, see Event severities.

Procedure

1. Navigate to **All** > **Event Management** > **Anomaly Detection** > **Anomaly Score to Severity Mapping**.
2. In the **Metric Anomaly Score to Event Severity Maps** list view, double-click the event severity that you want to modify.
3. Modify the range of anomaly scores that map to the respective event severity.
4. Click **Update**.

View anomaly alerts

Anomaly alerts indicate deviation from projected metric values for monitored CIs. Anomaly alerts are separate from regular IT alerts. They appear in the Service Operations Workspace but not in Express List. You can define an anomaly alert promotion rule to generate an IT alert that is based on anomaly alerts.

Before you begin

Role required: evt_mgmt_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The statistical model is used to calculate standard deviations, upper and lower bounds, and statistical outliers which are then used to detect anomalies. An anomaly occurs when metric values are out of the projected values according to the statistical model. The system monitors the frequency and persistence of statistical outliers across time to compute a score between 0-10 that indicates how abnormal a deviation is.

Metric Intelligence constantly generates anomaly alerts whenever the anomaly score is above zero. If there is a score that is above four and which has changed from the previous score, then it is sent to the instance. On the instance, the entire sequence of anomaly scores over time can then be displayed in the Insights Explorer.

Metric Intelligence monitors historical data to ensure that metric bounds are accurate and that the alerts close when the anomalous behavior stops. Metric Intelligence also ignores irrelevant anomalies based on the metric's labels. For example, if the available CPU increases from its upper bound of 50% to 60%, the *available* label is identified as an indicator. The behavior does not register as an anomaly. This is because the increase in available CPU represents a favorable outcome, even though the metric of 60% would otherwise represent an anomaly.

Procedure

1. Navigate to **All** > **Event Management** > **Metrics** > **Anomaly Alerts**.
2. In the **Alerts Anomalies** list view, double-click an alert that you want to view.
   Information about an anomaly alert includes the metric value, promotion parameter, and upper and lower bounds (lower bounds cannot be less than 0). Events relevant to the alert appear in the alert's work notes. Description includes details about why the anomaly alert was generated, the metric name, and the anomaly score.

What to do next

Create an Advanced Promotion Engine Definition, to define the conditions which must be met to promote anomaly alerts to IT alerts. For details, see Create a definition for the Advanced Promotion Engine.

View metric values in the Insights Explorer

Metric Intelligence calculates statistics for CI metric data. Insights Explorer displays these metric values as metric charts for the CIs in the CMDB. Insights Explorer lets you overlap any metrics for any CIs in a single chart to create a multi-layered view of metric values across a time range.

Before you begin

Role required: evt_mgmt_user

About this task

The Insights Explorer provides a canvas to which you can drag various metrics for CIs and create charts. You can place a single metric on the canvas, for an individual metric for a CI. For a more comprehensive chart, you can layer several metrics for a single CI or for different CIs to create a multi-metric chart. In such layered chart, each metric is distinguished by a unique color, which lets you compare metric values across CIs within a specific time range.

Insights Explorer displays metric boundaries according to the most recently configured model. Therefore, legacy anomalies from a previous build may still appear as anomalies, even though according to the current boundaries, the alert does not qualify as an anomaly.

You can also access the Insights Explorer from the business service map. For details, see View metrics from a business service map.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Insights Explorer

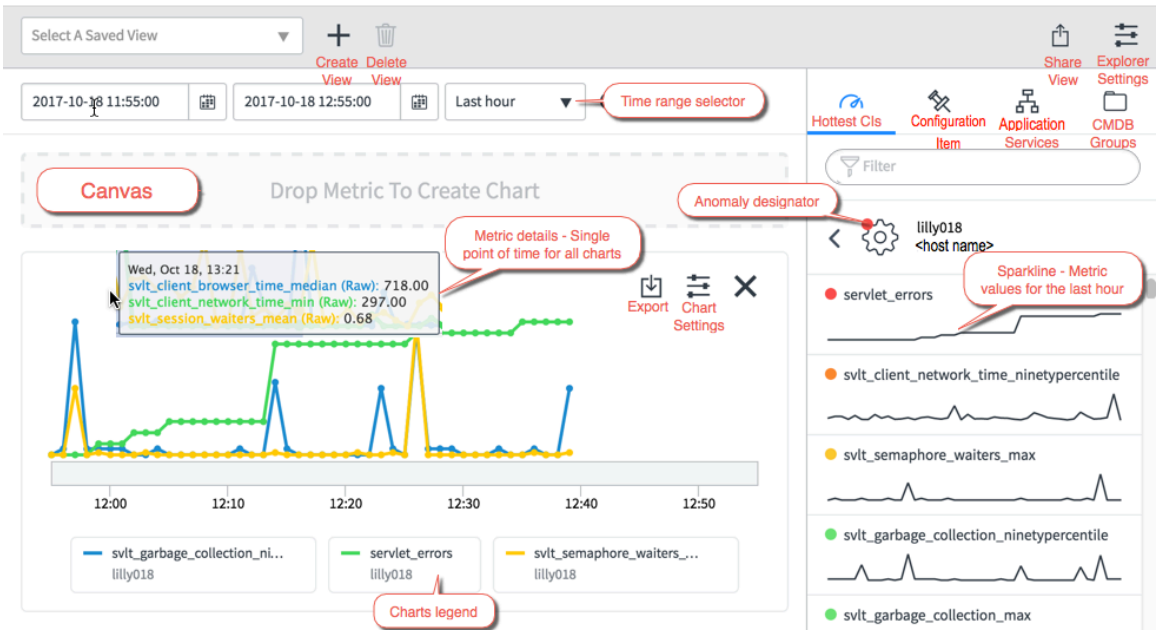

If the host name of the CI is available, it appears underneath the CI name. Host name is populated by retrieving the node information from the generated binding event as part of metric binding.

To navigate to a CI for which to display metrics:

- Use Hottest Configuration Items ( ) to quickly access the 10 most anomalous CIs.
- Use Configuration Items ( ) to create a separate custom list of any CIs from the CMDB and then add the metrics for these CIs to the canvas.
- Use Application services ( ) to create a custom list of application services that let you drill into the CIs of these services. You can then add metrics for these CIs to the canvas.
- Use CMDB Groups ( ) to create a custom list of CMDB Groups that let you drill into the CIs of these groups. You can then add metrics for these CIs to the canvas.

The following colors, which are based on color settings for alert severities, indicate anomaly severities:

Red

Critical severity (highest severity).

Orange

Major severity.

Yellow

Minor severity.

Blue

Warning severity (lowest severity).

Green

Informational. No severity.

The dots at the upper left on the CI icons indicate general anomaly tracking status:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Grey: Metrics are being tracked for the CI, but none of those metrics have ever had an anomaly.
- Green: Metric for CI had at least one anomaly in the past, but anomaly score is currently 0.
- Colored: Corresponds to the last recorded severity level according to the configuration of the color band ranges.
- No dot: There are no metrics for the CI and there is no data to drill into.

> **Note:** **Configuration Items** and **Application Services** are likely to include CIs for which there is no metric data.

Metric Intelligence does not track changes to an anomaly score if it is within the green band range.

As you create charts, a legend for each series chart automatically appears underneath the chart. For a single-metric chart, the legend has entries for the aggregations in the chart. For a multi-metric chart, the legend has entries for the metrics in the chart.

Modifications in the Insights Explorer remain during the session, and once the Insights Explorer is refreshed it returns to its initial state.

Procedure

1. Navigate to **All** > **Event Management** > **Metrics** > **Insights Explorer**.
2. Click a CI to drill down to its associated list of metrics, ordered by their most recent anomaly scores.

   For each metric, a sparkline displays metric values for the CI for the last hour. Point to the sparkline to display the metric value at each point of time.

   If resource binding is enabled, then the list of series being monitored for a CI also contains series that belong to a resource and metric. The naming convention for such series is '<resource name>/<metric name>'. For a series that is only bound to a CI, the naming convention is '<metric name>'.

3. Drag a metric to the **Drop Metric To Create Chart** area on the canvas and point to the chart to display the exact metric value at each point of time.
   The metric chart displays data for the specified time period, according to the chart settings. By default, raw data is displayed. If the time period is longer than a week or earlier than the last week, then average data is displayed. Drag subsequent metric to the canvas as follows:
   - In the **Drop Metric To Create Chart** area, so it is displayed individually in a new chart.
   - Onto an existing chart, in the **Replace Chart** highlighted block to overwrite the existing chart.
   - Onto an existing chart, in the **Add Metric** highlighted block to overlap it with the existing metrics in the chart.

     When you display multiple metrics in a single chart, the same aggregation selection applies to all the metrics, and statistics settings are not available. For example, if **Average** is selected in the **Chart Settings**, then averaged values are displayed for all the metrics in the chart.

4. Use the **Search** box to search and filter for specific configuration items, application services, or CMDB groups that are not listed by default.
   In the SEARCH RESULTS section, click the + icon next to an item that you want to add to the ADDED ITEMS section. Or, use advanced search:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Select and then click **Add Configuration Item** to add any specific CIs. In the Add Configuration Items dialog box, select a **CI Class Type** such as Linux Server.
- Select and then click **Add Application Service** to add all CIs included in a specified application services. In the Add Application Services dialog box, select `Discovered Service` or `Manual Service` as the **CI Class Type** and then specify filter conditions to narrow down the search.

  Note: You can add an application services only if its *`Operational Status`* is 'operational'.

- Select and then click **Add CMDB Group** to add all CIs included in a specified CMDB group. In the CMDB Groups dialog box specify filter conditions such as **[GroupName] [starts with] [d]**, to narrow down the search.

After selecting the items that you want to add, click **Add**.

5. Add related CIs:

   a. Click the **Add related CIs** ( ) icon next to a CI.
   The Add Related Configuration Items dialog box appears, displaying all relationships in which the selected CI is a parent or a child.

   b. To further filter the list, you can define specific conditions that must be met.

   c. Select the relationships from which you want to add the parent or the child CI, whichever is not the selected CI.

   d. Click **Add**.

   The CIs from the selected relationships, which are not the selected CI, are added to the list of CIs in the Insights Explorer.

6. Point to a legend item and click its '**X**' icon to remove the corresponding metric from the chart, or to disable the corresponding series from the **Chart Settings**.
7. Click the cog **Settings** icon to open application settings:
   - Enable, disable, or set the frequency for refreshing anomaly data on the Insights Explorer.
   - Activate or deactivate the display of related lists such as All Alerts, Incidents, and Problems. For more information, see Create or modify map indicators.
8. Click the **Chart Settings** icon to toggle the display of statistics and aggregations on the chart. Enabling or disabling an item to add or to remove metrics from a chart, also updates the legend of the chart to reflect the change.

Chart Setting	Description
Show Bounds	Displays a computed upper and lower bound for the metric, based on learning of past values. The upper and lower bounds always display together.
Show Anomaly Scores	Displays anomaly scores on a virtual axis of 0-10. Color code is based on the score thresholds that are defined in the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart Setting	Description
	Anomaly Score to Event Severity Map [sa_metric_anomaly_score_to_event_severity_map] table.
Average	Displays aggregated average metric values calculated in one hour intervals. Aggregated averages are displayed as follows: ○ For data that is up to 8 days old: Displays 10-minute aggregation windows. ○ For data that is up to 94 days old: Displays 10-minutes aggregation windows. ○ For older data: Displays 1-hour aggregation windows.
Minimum	Displays aggregated minimum metric values calculated per one hour time periods. For data that is up to 8 days old: Displays 10-minute aggregation windows.
Maximum	Displays aggregated maximum metric values calculated for one hour time periods. For data that is up to 8 days old: Displays 10-minute aggregation windows.
Raw	Displays the raw, unaggregated metric values. These values are kept only for the last 8 days.

What to do next

- Point to the icon of a CI in the right-hand side bar to display the CI type.
- Zoom in or out by changing the default **Last Hour** time range selection. Select one of the preset time periods to display anomaly scores for the last **6 Hours** for example, or specify a custom time period for up to 90 days back.

  To zoom in, highlight a section of a chart. Point to the upper left corner of the section and drag the cursor to the lower right corner. The time range of the chart changes accordingly. If there are other charts on the canvas, they are all automatically synchronized to display data for the same time range. The time range in the title bar automatically updates to reflect the new time range.
- Right-click on a CI in the right-hand side bar and select **View Form** to open its CI form.
- Click the map icon next to an application service listed in the **Application Services** tab, to open its application service map.
- Click Remove ( ) next to a CI to remove it from the Insights Explorer.
- Click the Export icon in a chart to download it as a .png or .svg image, or as a .pdf document.

Related topics

Create an Insights Explorer view

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create an Insights Explorer view

You can create a custom view in the Insights Explorer that saves any metric charts that were added to the canvas, and any added configuration items. After you create a view, you can share a link to the view.

Before you begin

Role required: TSOM Assurance admin

Procedure

1. Navigate to **All** > **Event Management** > **Metrics** > **Insights Explorer**.
2. Click **Create View** ( ), and enter a **Name** in the Create New View dialog box.
   If you are already creating a view, then you can choose **Create With Existing** to load an existing view to work with rather than starting with a blank canvas.
3. Click **Create**.
4. Add items to be included in the view:
   a. Drag metrics and create charts on the canvas.
   b. Add configuration items.
   While modifying a view, it is automatically saved every time you update it. The time settings are not saved in views, and when a saved view is loaded, its time period is always set to **Last hour** displaying current information.

What to do next

- Share a view: Click **Share View** ( ), and then in the Share dialog box, click **Copy Link**. Send or paste the link to make it accessible to whoever you want to share the view with. A shared link captures the time range of the view, the charts that were added to the canvas, and any added configuration items.
- Display a saved view: Open the **Select a Saved View** drop down list, and select a previously saved view that you want to display.

View metric charts in Agent Workspace

Agent Workspace for Metric Intelligence provides easy access to metric charts for a CI. On an alert form, you can drill down to important metric charts for the CI that is associated with the alert.

Before you begin

- The CI that you want to view metric charts for, must be associated with an alert.
- The Metric Dashboard Metadata [sa_metric_dashboard_metadata] table must contain active chart configurations for the CI class for which you want to view metric charts.

Role required: workspace_user and itil

About this task

The set of metric charts, chart types, and other chart characteristics, is determined by the chart configurations in the Metric Dashboard Metadata [sa_metric_dashboard_metadata] table, for the CI class.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **Agent Workspace** and click **Agent Workspace Home**.
2. In Agent Workspace, click **List** and in the **Lists** panel, select one of the **Alerts** categorized lists such as **All Alerts** or **Top Priority**.
3. In the alerts list view, open an alert.
   The alert must have a **Configuration item** associated with it.
4. In the **Details** tab of the alert form, click **Metrics**.

Result

All metric charts that are configured for the CI associated with the alert, and that are active, appear on the metric charts form in the **Metrics** tab. The calculation of the default time range for the charts is based on the times that the first and last events that generated the alert, occurred.

What to do next

- Modify the time range for the charts by changing **Start Time** and **End Time**. When you modify the time range, all charts are updated to reflect the new time range.
- Click the cog wheel **Open settings** icon and change the settings for all charts:
  - **Show bounds**: Show the bounds in the chart. Even if enabled, there are no bounds in a chart for aggregated resources.
  - **Show legend**: Show chart legend.
- Click **Insights Explorer** to open the Insights Explorer in a new browser tab. That Insights Explorer opens with the current CI pre-selected and with the current time range settings in Agent Workspace.

Configure metric charts for Agent Workspace

Configure the set of Metric Intelligence metric charts, chart types, and other chart characteristics, that appear in Agent Workspace.

Before you begin

Role required: admin or evt_mgmt_admin

About this task

Up to 10 metric charts for a class can be configured to appear in the Metric Intelligence metric charts form in Agent Workspace. Metric chart configurations are stored in the Metric Dashboard Metadata [sa_metric_dashboard_metadata] table, determining which metric charts appear for a CI that is associated with an alert. By default, the Metric Dashboard Metadata contains pre-configured metric charts for several classes.

If resource binding is enabled and you specify a **Resource Type** in the metric chart configuration, then all the resources for the CI are included in the chart. The class table for the resource class (such as *ci_resource_metric_extension*) that you specify is used for retrieving all the resources for the CI. Then, for each point in time in the data series, the values of resources are aggregated into a single data point. This data point is the average of the metric values of all resources for that point of time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Chart aggregation

Aggregation is applied differently, depending on the setting of **Chart Type***, and on whether resources are included in the chart as indicated by the setting of* **Resource Type***:*

Chart type	Resources	Aggregation
Time_Series	No	No aggregation
Time_Series	Yes	Resources are averaged per data point.
Single_Score	No	Data points are aggregated according to the **Chart Aggregate** setting.
Single_Score	Yes	Resources are averaged per data point and then data points are aggregated according to the **Chart Aggregate** setting.

Procedure

1. Navigate to **All** > **Event Management** > **Administration** > **Metrics View Configuration**.
2. In the Metric Dashboard Metadatas (Dashboard Metadata) list view, select an existing configuration or click **New**.
3. Fill out the Metric Dashboard Metadata form.

Field	Description
Active	Include the chart on the Metric Intelligence charts form in Agent Workspace.
Resource Type	Resource class associated with the CI and for which to retrieve and aggregate all resources.
Configuration Item Type	CI class that the chart applies to such as *cmdb_ci_metric_extension*.
Chart Aggregate	Type of aggregation that is applied for single-score charts. Aggregation can be the average of the data points in the time series. Or, the minimum, maximum, or the last data point in the time series.
Source Metric Type	Metric such as '% Free Space' that is being monitored for the CI type.
Order	Order in which the metric chart appears within all metric charts specified for the class. For proper appearance and alignment, all single-score charts always appear after any time-series charts, regardless of order settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Chart Type	○ Time_Series: Charts all the data points in the time series for the metric and CI of the chart. ○ Single_Score: A single number which is the aggregation of all data points in the time series for the metric and CI of the chart. Aggregation is determined by the setting of **Chart Aggregate**.

View Metric Intelligence self-health statistics

Monitor the status of components and processes of Metric Intelligence. Use the XMLStats page to view statistics and diagnostic details that can help with troubleshooting issues with Metric Intelligence.

Before you begin

Role required: evt_mgmt_admin

Procedure

1. Open a browser and enter the URL to the instance.
   Type in: `http://<instance name>.service-now.com`
2. Append the following text to the instance URL: `xmlstats.do?include=XMLStatsOI`.
   The complete URL format is: `http://<instance name>.service-now.com/xmlstats.do?include=XMLStatsOI`
3. Hit Enter.

Result

Details are separated between two sections, instance level statistics and node level statistics. Details in the node section are specific to the node in the URL.

Node section

Entry	Description
batches_recieved	Number of batches of metrics that the node received from the MID Server.
batches_processed	Number of batches of metrics processed on the node.

Details in the instance section are for the entire instance, showing the status of the self-health monitors.

Instance section

Entry	Description
scheduled_jobs_monitor	Status of the self-health monitor that checks if the essential Metric Intelligence scheduled jobs are active.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Instance section (continued)

Entry	Description
distributed_cluster_monitor	Status of the self-health monitor that checks if the Metric Intelligence MID Server distributed cluster is running properly.
oi_extension_monitor	Status of the self-health monitor that checks if the Metric Intelligence extension is running properly.
metrics_collection_monitor	Status of the self-health monitor that checks if metrics are being collected by active connectors that are configured for metrics collection.

Verify that the definitions for the Advanced Promotion Engine are active

Verify that the Advanced Promotion Engine definitions that come with the base system are active.

Before you begin

Role required: evt_mgmt_admin

About this task

The active definitions promote a group of anomaly alerts to the All Alerts table when all of the definition's conditions are met.

You can't modify the conditions of the definitions that come with the base system. For information on creating customized definitions, see Create a definition for the Advanced Promotion Engine.

Procedure

1. Navigate to **All** > **Advanced Promotion Engine** > **Advanced Promotion Definitions**.
2. Select the Advanced Promotion Engine definition that you want to be active:
   - Windows Server
   - Linux Server
3. Verify that the **Active** check box is selected.
4. Select **Save**.

Create a definition for the Advanced Promotion Engine

Create a definition for the Advanced Promotion Engine so that you can define the conditions for promoting the anomaly alerts to IT alerts. By promoting the alerts that meet the conditions, only the most relevant alerts move to the table.

About this task

IT Alerts appear in the Express List and are addressed by an operator.

Only metrics that are associated with a metric configuration rule that has **Anomaly score level = Anomaly alerts** are considered for promotion to IT alerts. For details on metric configuration rules, see Metric Configuration Rules.

Before you begin

Ensure that the Metric Intelligence plugin (com.snc.sa.metric) is installed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: evt_mgmt_admin

Procedure

1. Navigate to **All** > **Advanced Promotion Engine** > **Advanced Promotion Definitions**.
2. Select **New**.
3. In the **Name** field, enter a descriptive name for the definition.
4. In the **Order** field, set the execution order for the definition.
   The definitions with a lower order are executed first.
5. Activate the definition by selecting the **Active** check box.
6. On the **CIs** tab, fill in the fields for the CI alerts that you want to consider for promotion.

CIs tab fields

Field	Description
CI class	CI class alerts to be promoted to the All Alerts table.
Filter	Filter to restrict which CIs that the definition applies to. Select **New Criteria** to configure the conditions that the definition applies to. You can add more conditions by selecting either **OR** or **AND**.
Skip CIs that were created less than this number of days ago	Number of days in the past that any CI created during that time is ignored. For example, a value of 7 indicates that any CIs created less than 7 days ago are not considered for promotion to the All Alerts table.

7. On the **Severity** tab, enter the minimum severity of alerts to be considered for promotion.
   Alerts with the indicated severity level and higher levels are considered for promotion.
8. On the **Number of Alerts** tab, in the **Number of alerts on CI** field, enter the minimum number of a CI's alerts with an identical severity level (or higher) to be considered for promotion.
   For example, a value of 3 means at least three alerts with the indicated severity (or higher) must exist on a CI.
9. On the **Time Window** tab:
   - In the **Consider alerts updated at least this number of minutes ago** field, enter the minimum number of minutes old for an alert to be considered for promotion.

     For example, if the value of the field is 5, only alerts that are updated at least five minutes ago are considered for promotion.
   - In the blue box, select the number of minutes link to modify the amount of minutes during which alerts are to be considered for promotion.

     The **Properties** page opens and you modify the value in the **The width of the time...** field.

     Default = 30.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example:

- The current time is 10:05.
- In the **The width of the time...** field, you entered a value of 30.
- In the **Consider alerts updated at least this number of minutes ago** field, you entered 5.

Alerts that occur between 9:30-10:00 are considered for promotion.

What to do next

Ensure that a metric configuration rule exists which includes metrics you want to be promoted to an IT alert. Verify that the anomaly score level of the rule is **Anomaly alerts**. For details on configuring a metric configuration rule, see Create a configuration settings rule.

View the promoted events

View the anomaly alerts that the Advanced Promotion Engine promoted to IT events. Viewing IT events enables you to track the source of the event.

Before you begin

Role required: evt_mgmt_user

About this task

The system generates an IT event when an anomaly alert is promoted. Those events then become IT alerts. For details, see View the promoted anomaly alerts.

Procedure

1. Navigate to **All** > **Advanced Promotion Events**.
2. On the Events page, view the IT events that are created from the anomaly alerts by the Advanced Promotion Engine.

View the promoted anomaly alerts

View the anomaly alerts that the Advanced Promotion Engine promoted to events, which are then processed as alerts in the All Alerts table. Viewing alerts in the All Alerts table enables you to track the source of the alert.

Before you begin

Role required: evt_mgmt_user

About this task

The Event Management alert grouping feature takes the anomaly alerts that are promoted to the IT events and groups them into new IT alerts.

The Alerts table displays all of the relevant issues in the system.

Procedure

1. Navigate to **All** > **Advanced Promotion Engine** > **Anomaly Alerts**.
2. In the Is Advanced Promoted column header, view the alerts with a value of **true**.
   The IT alert that is created from the anomaly alert appears in the IT Alert column.
3. **Optional:** Select an alert to receive more information about the alert.
   For details on the displayed additional information, see Alerts in Service Operations Workspace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create metric rules

Create a metric rule to determine the metric, the threshold values and ranges, and the severity level of generated events. You can add a filter that determines the Configuration Items (CI) types or resources for which you receive alerts. For example, you can configure a filter to receive alerts only for Windows servers.

Before you begin

Role required: evt_mgmt_admin

About this task

When a metric crosses or falls within the range of a configured threshold, the system creates a single event, which is eventually processed into an IT alert. The severity of the event (and alert) is determined by a defined threshold, and the event is bound to a CI.

Procedure

1. Navigate to **All** > **Event Management** > **Metrics** > **Metric Rules**.
2. Select **New Metric Rule**.
   The **CI selection** page opens.
3. Select the CI class on which to run the rule in the **CI Class** field.
   Only CIs containing metrics are available for selection.
4. **Optional:** Apply a filter within the selected CI class.
   For example, if you selected a CI class of **Linux Server**, you can apply a filter. The filter indicates that the rule applies only to servers with a host name that starts with a specified string.
5. Select **Next**.
   The **Alert criteria** page opens.
6. Select the metric and criteria to create an alert.

   a. Select the alert severity that you want to configure in the **Alert severity** field.
      Choose from the following:
      - Critical
      - Major
      - Minor
      - Warning

   b. Select the metric for the alert severity to apply to in the **Metric** field.

   c. Select the comparison operator (equals, is greater than, is greater than or equals, and so forth) in the **Comparison operator** field.

   d. Enter a value in the **Value** field.

   For example, you can configure criteria to indicate that when the metric **node_rx_packets** is greater than 10, a Critical alert is generated.

   To define additional alert severities, select **Add another alert severity** and repeat this step.

7. Select when to generate the alert in the **Alert when** section.
   Choose from the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **All sampled values must breach the criteria in the following time period**.

  Select the time period, in minutes.

- **Average of sampled values in the following time period must breach the criteria**.

  Select the time period, in minutes.

- **Any time the sampled values breach the criteria**.

8. Select **Next**.
   The **Name and activation** page opens.
9. Verify the name and details of the metric rule.

   To activate the rule, ensure that the **Active** toggle switch

is activated.
10. Select **Finish**.

Create an event rule to bind metric events to host CIs

Create event rules to explicitly allow binding of metric events to host CIs.

Before you begin

Role required: evt_mgmt_admin

About this task

Starting with the London release, Metric Intelligence no longer relies on the default Event Management binding workflow which binds nodes to host CIs. Event Management is configured to apply the default host binding algorithm when an event rule cannot find a CI to bind to, based on event data. In Metric Intelligence, with a given event rule and two different events, one event might bind to the correct CI while another cannot. In this case, the same metric types would be created for the selected CI and for the host level CI, causing confusion.

For example: With an event rule that binds the disk percentage metric to a storage disk, one event can find a disk CI and another event cannot. This results in the disk percentage metric being created for a disk CI and for a host CI.

To avoid this confusion, Event Management default host binding is disabled for metric events in the base system. To allow host binding for metric events, explicitly create the following event rule in the Event Management event rule designer. The underlying binding functionality of the event rule designer is used to map metrics to CIs.

Procedure

1. Navigate to **All** > **Event Management** > **Rules** > **Event Rules** to open the event rule designer.
2. Click **New** and ensure that the **Active** toggle switch is set to the right (Active).
3. On the **Event Rule Info** tab:
   a. Set **Source** to the data source that this rule applies to.
   b. **Optional:** Enter a **Description** such as `Bind metrics at host level`.
4. On the **Event Filter** tab, set a filter condition to indicate that **Classification** is **Metric**.
5. Click **Save**.

Related topics

Create or edit an event rule

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Metric Intelligence reference

Reference topics provide additional information about Metric Intelligence settings and properties.

Insights Explorer settings

You can configure Insights Explorer settings to control for example, the refresh frequency or the display of related records.

Click the cog **Settings** icon to open the Application settings dialog box.

Setting	Description
Refresh interval	Enable and set the frequency for refreshing metrics data on the Insights Explorer.
Indicators	Enable any indicator to display related records such as incidents or problems for CIs. The respective icon displays on the timeline in the metric chart, indicating the number of records for the CI. A toggle is available only if the **Active in Metrics** field is selected on the Map Indicator form for the respective table. For example, the **Incident** toggle is available only if **Active in Metrics** is selected on the Map Indicator form for the incident [Incident] table.

Related topics

Dependency Views map indicators

Sensitivity bounds properties for Insights Explorer metrics

The following properties can be configured to customize system behavior when detecting sensitive bounds for Insights Explorer metrics. To invoke changes to the default values, navigate to **System Properties** > **All Properties** and add the indicated properties (they are not visible by default).

Sensitive Bounds Properties

Property Name	Description
sa_metric.sensitive_bounds_calculation_enabled	Enable/disable sensitive bounds calculation. Default: True
sa_metric.sensitive_bounds_percentile	Percentile value of sensitivity bounds calculation. If the configured percentage of metric values falls within the number of standard deviations (STDs) indicated in the *sa_metric.sensitive_bounds_sensitivity_fact* property, sensitive bounds are detected and sensitivity calculation changes to the value specified in the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sensitive Bounds Properties (continued)

Property Name	Description
	sa_metric.sensitive_bounds_control_factor property. Default: 95%
sa_metric.sensitive_bounds_control_factor	The number of standard deviations (STDs) from the mean metric value by which metric bounds are measured after the system detects sensitive bounds. Default: 5.0
sa_metric.sensitive_bounds_sensitivity_factor	The number of standard deviations (STDs) from the mean metric value by which the system tests for sensitive bounds. Default: 2.0

Related topics

Metric bounds sensitivity

Access to cloud environments for ITOM products

If your organization deploys IT assets on the cloud, ITOM products must access your cloud environment to collect information about the IT infrastructure.

Related topics

Discovery for cloud environment

Cloud resource discovery references

Install and configure MID Servers to access cloud environments

Install and configure the MID Servers correctly to enable ITOM products to access to the cloud resources.

Before you begin

Role required: discovery_admin or sn_cmp.cloud_admin (for Cloud Provisioning and Governance)

About this task

Procedure

1. For IP-based discovery, Installing the MID Server on a server inside the cloud environment, inside a Virtual Private Cloud (VPC).

   Example

   For example, if you are configuring access to AWS service accounts using IAM roles, install the MID Server at the Amazon EC2 instance.

   Important: If you configure AWS credentials at ServiceNow AI Platform, the location of the MID Server doesn't matter.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If your organization has multiple VPC segments, deploy the MID Servers in a manage-mentor transit VPC.

2. Activate IP-based discovery jobs to access the compute layer using agentless credentials.
3. Configure the MID Server with the capabilities necessary for the cloud environment you are planning to discover.

   You can select the `ALL` capability or configure specific capabilities for tighter security.

   Add Cloud Management in **supported application** in MID Server for Cloud Provisioning and Governance requests to be processed.

Cloud environment	MID Server capability required by Cloud Discovery	MID Server capability required by Cloud Provisioning and Governance
Amazon AWS Cloud	`AWS` or `Cloud Management` or `ALL`	`AWS` or `Cloud Management`
Microsoft Azure Cloud	`Azure` or `Cloud Management` or `ALL`	`Azure` or `Cloud Management`
Google Cloud	`Google` or `Cloud Management` or `ALL`	`Google` or `Cloud Management`
IBM Cloud Platform	`IBM` or `Cloud Management` or `ALL`	N/A
Oracle Cloud Infrastructure (OCI)	`OCI` or `Cloud Management` or `ALL`	N/A
OpenStack	`OpenStack` or `Cloud Management` or `ALL`	N/A
Red Hat Virtualization (RHV)	`RHV` or `Cloud Management` or `ALL`	N/A

4. If the MID Servers must go through an unauthenticated proxy server to access cloud resources, modify the `agent\conf\wrapper-override.conf` file, as follows:

```
wrapper.java.additional.1=-Dhttp.proxySet=true
        wrapper.java.additional.2=-Dhttp.proxyHost=<proxyHost>
        wrapper.java.additional.3=-Dhttp.proxyPort=<proxyPort>
        wrapper.java.additional.4=-DuseProxy=true
```

5. If you're using an authenticated proxy server, configure the MID Server parameters in the `config.xml` file.

   For a description of the parameter values, see the **Proxy server parameters** section in MID Server parameters.

Related topics

Proxy server configuration for MID Servers used for Cloud Discovery and Cloud Provisioning and Governance

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Proxy server configuration for MID Servers used for Cloud Discovery and Cloud Provisioning and Governance

You can configure the MID Server for Cloud Provisioning and Governance to use a proxy server. Using a proxy server supports all cloud-based activities such as running Discovery, billing downloads, provisioning virtual machines, and running life-cycle operations on virtual machines.

Proxy server limitations

- VMware Cloud is not supported.
- Windows or Linux platforms only are supported.
- The Google Cloud Console integration does not support downloading billing data.
- Remote PowerShell scripts cannot be executed.
- Custom APIs may not work.

Note:

SSH calls should not be routed through a proxy server, when using a private cloud environment that is on the same side of the proxy as the MID Server.

Set *mid.capi.skip_proxy_for_ssh* to true to skip it. See MID Server properties used by Service Mapping for more information.

mid.capi.skip_proxy_for_ssh:

- Determines whether the MID Server skips the proxy server when routing SSH calls.
- Type: true | false
- Default value: false

Supported proxy server authentication for Cloud Provisioning and Governance

Proxy Server type	Authentication type
HTTP/HTTPS	No authentication
SOCKS5	No authentication
HTTP/HTTPS	Basic authentication
SOCKS5	Basic authentication
HTTP/HTTPS	NTLM

Supported proxy server configurations for Cloud Provisioning and Governance

Supported configuration settings

Configuration	Operating system	Proxy server	Authentication mode
Configuration 1	Linux	None	Not Applicable
Configuration 2	Windows	Squid (HTTPS)	None
Configuration 3	Linux	Squid (HTTPS)	Local
Configuration 4	Windows	Squid (HTTPS)	Active Directory

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Supported configurations

Cloud provider	Functionality	Configuration 1	Configuration 2	Configuration 3	Configuration 4
Amazon AWS Cloud, Microsoft Azure Cloud, and IBM Cloud Platform	Cloud Discovery	Yes	Yes	Yes	Yes
	VM Provisioning	Yes	Yes	Yes	Yes
	VM life cycle	Yes	Yes	Yes	Yes
	CFT provisioning	Yes	Yes	Yes	Yes
	Billing download	Yes	Yes	Yes	Yes
	Scripts using SSH	Yes	Yes	Yes	Yes
	Scripts using PowerShell	Not Applicable	Not Applicable	Not Applicable	Not Applicable
	Cloud events	Not Applicable	Not Applicable	Not Applicable	Not Applicable
Google Cloud Platform	Cloud Discovery	Yes	Yes	Yes	Yes
	VM Provisioning	Yes	Yes	Yes	Yes
	VM life cycle	Yes	Yes	Yes	Yes
	CFT provisioning	Yes	Yes	Yes	Yes
	Billing download	No	No	No	No
	Scripts using SSH	Yes	Yes	Yes	Yes
	Scripts using PowerShell	Not Applicable	Not Applicable	Not Applicable	Not Applicable
	Cloud events	Not Applicable	Not Applicable	Not Applicable	Not Applicable
VMware Cloud	Not Applicable				

Related topics

- Install a MID Server on Windows
- Install a MID Server on Linux
- MID Server parameters

Cloud Request Retry Configuration

If a request is throttled by a cloud provider during Discovery, Cloud Request Retry Configuration provides a customizable method to retry requests. Discovery and Service Mapping Patterns

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

includes a retry configuration for AWS and Azure. You can customize the included configuration or create your own.

Discovery admins and Cloud admins can access the request retry configuration at **All** > **Discovery** > **Cloud Request Retry Configuration**. You can create one configuration for each provider.

When a request is throttled, the retry framework uses retry configuration defined for the provider to handle retries before returning the final response to the ApiCommand classes:

- AwsApiCommand
- AzureApiCommand

The retry configurations are synced with the MID Servers through the MID Server property, *`mid.cloud.discovery.retry.configuration`*.

There are the following retry strategies:

- Exponential backoff
- Response header backoff
- Custom backoff

Exponential backoff

For the following example configuration:

Setting	Value
Max retries	3
Response codes	429
Base delay in ms	1000
Max delay in ms	10000
Additional delay window in ms	1500

The Exponential backoff retry strategy works as follows:

- 1st retry—the backoff multiplier is randomly selected between 0 and 1. The max delay value is 400 ms (400 * 1).
- 2nd retry—the backoff multiplier is randomly selected between 0 and 3. The max delay value is 1200 ms (400 * 3).
- 3rd retry—the backoff multiplier is randomly selected between 0 and 7. The max delay value is 2800 ms (400 * 7).

On subsequent retries, if the delay exceeds 10000 (the max delay), 10000 will be used as initial delay.

Once the initial delay is generated, jitter is added to the delay. The jitter window is defined by the **Additional delay window in ms** field. The system selects a random value between 0 and 1500 and adds it to the initial delay.

If the initial delay is 500, the final delay (with jitter) can be a value between 500 and 2000 ms.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Response header backoff

For the following example configuration:

Setting	Value
Max retries	3
Response codes	429
Response header	Retry-After
Response header delay unit	Seconds
Additional delay window in ms	1500

The Response header backoff strategy works as follows:

- Fetch the value of header *`Retry-After`* from the server response.
- Convert the *`Retry-After`* to milliseconds by multiplying by 1000.

Once the initial delay is generated, jitter is added to the delay. The jitter window is defined by the **Additional delay window in ms** field. The system selects a random value between 0 and 1500 and adds it to the initial delay.

If the initial delay is 2000, the final delay (with jitter) can be a value between 2000 and 3500 ms.

Custom backoff

With a custom backoff retry strategy, you define the **Max retries** and **Response codes** and create your own **Mid script include** that defines how requests are retried using the *`getDelay()`* function. For more information, see Script includes.

Set up a cloud service account

Create and configure ServiceNow cloud accounts at ServiceNow AI Platform for the corresponding Cloud vendor service accounts.

Related topics

Cloud Discovery service account form reference

Set up Oracle Cloud infrastructure (OCI) service accounts

Create and configure cloud service accounts at ServiceNow AI Platform for the corresponding Oracle Cloud Infrastructure (OCI) service accounts.

OCI service account is a type of identity used to authenticate and authorize access to OCI resources and services. It provides a way for applications, services, and users to securely interact with OCI resources. Service accounts are often used for automated processes, applications, or services running within OCI, allowing them to access resources according to the defined permissions and policies.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

OCI access and permission using policies

Oracle Cloud Infrastructure (OCI) access and permission using policies provide privileges for secure cloud resource utilization and management in OCI.

For information on OCI permissions, go to the Oracle Help Center and search for the "Managing Access to Resources" article.

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Create Oracle API credentials

Create Oracle API credentials on the ServiceNow AI Platform to enable access to your Oracle resources during Oracle discovery.

Before you begin

Verify that your OCI user has at least read-only access to all Oracle services. For information on access permissions, consult your system administrator.

Retrieve the following properties from the Oracle Cloud Console:

- Tenancy Oracle Cloud Identifier (Tenancy OCID)
- User OCID
- Key fingerprint
- Private key

For instructions on retrieving these properties, go to the Oracle Help Center and search for the "Required Keys and OCIDs" article.

Role required: discovery_admin

Procedure

1. Navigate to **All** > **Discovery** > **Credentials**.
2. Select **New**.
3. Select **Oracle API Credentials**.
4. On the form, fill in the fields with the property values you retrieved.

Oracle API Credentials form

Field	Value
Name	Unique and descriptive name for the Oracle credentials.
Active	Option to select to use the credential.
Tenancy Id	Oracle Cloud Identifier (OCID) for your tenancy.
User Id	OCID of the OCI user.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Key Fingerprint	Fingerprint of the public key.
Private Key	RSA private key that is used for API signing authentication.

5. Select **Submit**.

What to do next
Create OCI service accounts. For more information, see Create OCI service accounts.

Create OCI service accounts

Create Oracle Cloud Infrastructure (OCI) service accounts on the ServiceNow AI Platform to access your Oracle account during Oracle discovery.

Before you begin
Verify that Oracle API credentials have been created. For more information, see Create Oracle API credentials.

Retrieve the Compartment ID from **Identity & Security** > **Compartments** in the Oracle Cloud Console.

For GovCloud accounts, confirm that Discovery and Service Mapping Patterns is using at least version 1.29.0.

For UK Sovereign Cloud accounts, confirm that Discovery and Service Mapping Patterns is using at least version 1.35.0.

Role required: discovery_admin

Procedure
1. In the navigation filter, enter `cmdb_ci_cloud_service_account.list`.
2. Select **New**.
3. In the **Discovery Credentials** field, begin to enter the credential name and select it from the list.
4. On the form, fill in the remaining fields.

Cloud Service Account form

Field	Description
Name	A unique name for the OCI service account.
Account ID	The compartment ID, which is the Oracle Cloud Identifier (OCID) of the compartment associated with this service account.
Datacenter URL	The URL of the datacenter. Example URLs:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ For commercial cloud: `https://{service}.ap-mumbai-1.oraclecloud.com` ○ For GovCloud: `https://{service}.us-gov-ashburn-1.oraclegovcloud.com` ○ For UK Sovereign Cloud: `https://{service}.uk-gov-london-1.oci.oraclegovcloud.uk`
Datacenter Type	The type of datacenter where the account is hosted, which should be `OCI Datacenter [cmdb_ci_oci_datacenter]`.

5. Select **Submit**.

What to do next

Schedule an OCI cloud discovery. For more information, see Create an OCI Discovery schedule in Discovery Admin Workspace.

Setup AWS service accounts

Create and configure cloud service accounts at ServiceNow AI Platform for the corresponding Amazon Web Services (AWS) service accounts.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Ensure that you are familiar with the hierarchy of AWS service accounts in your environment. For example, if there are AWS Organizations, set up the management and member accounts at ServiceNow AI Platform to reflect that hierarchy. You can set up the AWS service accounts of the following types:

- **Discrete account**: Standalone account, with no management account.
- **Management account**: Management account that may or may not contain member accounts (subaccounts).

  Note: Some ServiceNow UI screens refer to management accounts as **Master** accounts and to the management account field as **Parent account**.

- **Member account**: Subaccount that belongs to the management account.

Access setup for AWS service accounts

Cloud Discovery and Cloud Provisioning and Governance need access to resources in the Amazon Web Services (AWS) service accounts. Learn about different methods of configuring such access.

Cloud Discovery and Cloud Provisioning and Governance access resources in AWS service accounts through MID Servers. You must authorize inbound traffic to Amazon EC2 instances from the MID Server for setting up initial communication. For more information, see Configure security group inbound rules using the AWS Management Console.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Types of AWS credentials

There are permanent and temporary AWS credentials that you can use for configuring access to AWS service accounts.

Permanent

The permanent credentials are the actual AWS credentials for the service account that you add to the Connections and Credentials module of ServiceNow AI Platform. While it might be time consuming to manage credentials on ServiceNow AI Platform, you avoid the complex configurations involved in using temporary credentials.

Temporary

The temporary credentials are generated by the AWS Security Token Service (AWS STS) for IAM roles. After you configure IAM roles for AWS accounts, the MID Server accesses AWS resources with these temporary credentials. You can use the default IAM role, *`OrganizationAccountAccessRole`*, or create custom IAM roles.

Assuming IAM roles in a large AWS organization is more convenient and offers better security than using large numbers of permanent credentials for all AWS accounts. Temporary credentials are only acquired on behalf of a service account when there's no permanent credential specified for that service account in the Service Accounts [cmdb_ci_cloud_service_account] table.

The MID Server uses the AssumeRole action in the AWS Security Token Service API to assume a member account role. Parameters passed to this API determine which additional security restrictions are applied to the role when it accesses the AWS resources.

By default, the MID Server is configured to assume the *`OrganizationAccountAccessRole`*, which grants temporary credentials to all the members of a primary account. This action occurs automatically if no permanent credentials exist for the member accounts. This configuration doesn't apply any additional security or restrict access to any resources in member accounts.

By default, the ServiceNow instance caches temporary credentials for member accounts for 60 minutes. This interval enables the horizontal discovery process to run multiple times without generating new credentials during each discovery. You can avoid credential caching or modify the caching period using MID Server properties.

IAM roles and permissions

To enhance security provided by the default AWS `OrganizationAccountAccessRole` role, you can customize the AWS roles that MID Servers can assume to receive temporary credentials for member accounts. You can configure additional permissions to improve security and customize the way that the member account's role is assumed when discovering cloud resources.

Methods of granting access

For the purposes of configuring access for AWS accounts, the following terms are used:

Trusting accounts

The trusting accounts don't have permanent AWS credentials. You configure the trust relationship for IAM roles in these accounts to rely on other accounts for access.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Trusted accounts

The trusted accounts are used by the trusting accounts for access. The ServiceNow UI refers to the trusted accounts as accessor accounts.

Typically, you set up access to the AWS accounts in your organization using the following methods:

Configuring access for a single account

- Configure AWS credentials on the ServiceNow AI Platform to avoid the complexity of creating and setting up IAM roles. For more information, see Configure access to the AWS accounts using permanent AWS credentials.
- Create an AWS service account on the ServiceNow AI Platform. For more information, see Create AWS service accounts.

Configuring access for an account that trusts an accessor account with AWS credentials

Setting up any AWS account to rely on a trusted account with AWS credentials

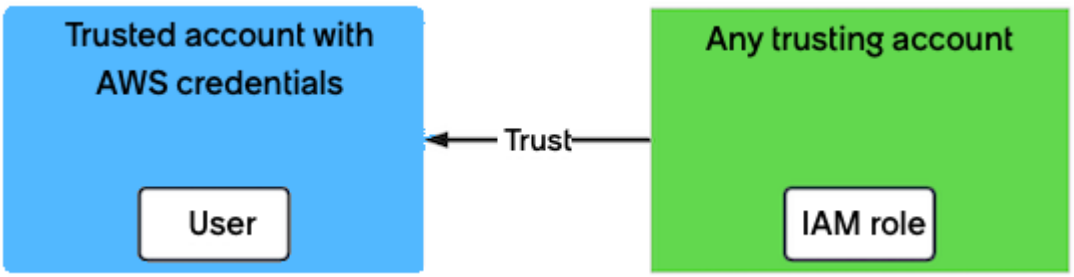

- Configure any type of account—discrete (independent), management, or member—to rely on a trusted account with AWS credentials for access.
- Configuring an IAM role belonging to the trusting accounts to trust the user of the trusted account enables using only one set of AWS credentials for providing access to multiple AWS accounts.

For more information, see Configure temporary credential access for trusted AWS accounts.

Configuring access for an account that trusts an accessor account without AWS credentials

Setting up any AWS account to rely on a trusted account without AWS credentials

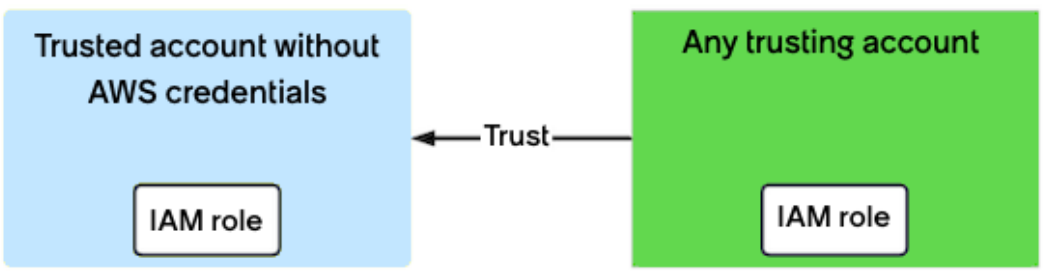

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Configure any type of account—discrete (independent), management, or member—to rely on a trusted account without AWS credentials for access.
- Configure an account without AWS credentials using an IAM role and permissions to access the trusting service account.
- Set up the IAM role of the trusting account to grant access to the IAM role of the trusted account.

For more information, see Configure credential-less access using trusted AWS accounts.

Configuring access for AWS member accounts by using a trust chain from the accessor through the management account.

Configuring member accounts to use their management account for access

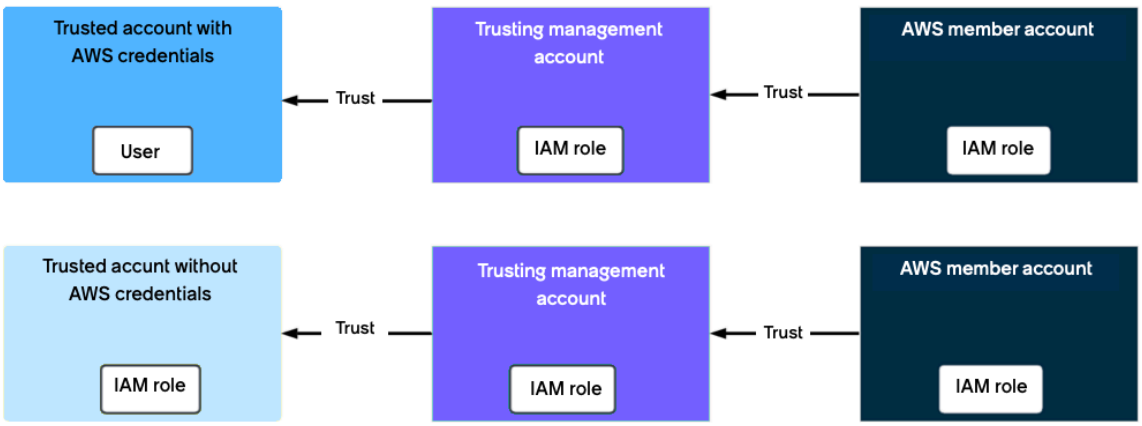

- Configure the IAM roles of the member accounts to trust the IAM role of their management account. For more information, see Configure access for trusting AWS member accounts in trust chain.
- Configure the management accounts to trust the accessor accounts with or without AWS credentials. For more information, see:
  - Configure temporary credential access for trusted AWS accounts.
  - Configure credential-less access using trusted AWS accounts.

How Cloud Discovery determines which credentials to use

Member account trusts management account

Cloud Discovery uses the following logic to determine which credentials to use to discover AWS cloud resources in member accounts:

1. If permanent credentials are defined for the member account in the Cloud Service Account [cmdb_ci_cloud_service_account] table, Discovery uses those credentials. The Cloud Service Accounts [cmdb_ci_cloud_service_account] table contains the information on the service account types, like management or member, and their credentials.
2. If no permanent credentials are defined for the member account, Discovery checks the Cloud Service Account AWS Org Assume Role Params [cloud_service_account_aws_org_assume_role_params] table for any special parameters associated with the member account. If parameters exist in that table,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Discovery uses the temporary credentials acquired from specifying a role and its parameters in the AWS Security Token Service API AssumeRole action.

3. If no special parameters are associated with the member account in the [cloud_service_account_aws_org_assume_role_params] table, Discovery checks that table for parameters associated with the management account. If parameters exist that define a role for the management account, Discovery uses the temporary credentials provided by that role.
4. If no special parameters are present in the [cloud_service_account_aws_org_assume_role_params] table for either management or member accounts, Discovery uses the defaults defined for the `OrganizationAccountAccessRole` role.

Member or management account trusts accessor account

1. If permanent credentials are defined for the member or management account in the Cloud Service Account [cmdb_ci_cloud_service_account] table, Discovery uses those credentials. The Cloud Service Accounts [cmdb_ci_cloud_service_account] table contains the information on the service account types, like management or member, and their credentials.
2. If no permanent credentials are defined for the account, Discovery checks the Cloud Service Account AWS Cross Assume Role Params [cloud_service_account_aws_cross_assume_role_params] table for any special parameters associated with the account. If parameters exist in that table, Discovery uses the temporary credentials acquired from specifying a role and its parameters in the AWS Security Token Service API AssumeRole action.

Related topics

Step-by-Step Guide to Configure AWS Cloud Discovery in ServiceNow Using Patterns [KB2184482]

AWS Cloud Discovery Setup – Use Case Guide with CloudFormation Templates [KB2221157]

Configure access to the AWS accounts using permanent AWS credentials

To securely access data on your provider account, the Discovery process must present appropriate credentials. To make the credentials available to Discovery and Cloud Provisioning and Governance, you first create a user with programmatic access in the AWS Management Console. You then securely store the credentials in a service account at ServiceNow AI Platform.

Before you begin

Roles required:

- AWS Management Console administrator
- For Cloud Discovery: discovery_admin
- For Cloud Provisioning and Governance: admin or sn_cmp.cloud_admin

About this task

Cloud providers often use different names for accounts, regions, and credential settings. Because the ServiceNow application supports several cloud providers, the app uses general-purpose names for the settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. **Optional:** To create a user for Cloud Provisioning and Governance (for programmatic access to your AWS resource and billing data), perform the following steps.

   a. On the AWS Management Console, navigate to **Identity and Access Managements (IAM)** > **Access Management** > **Users**.

   b. On the IAM Resources portal, select **Users**.

   c. Select **Add user**.

   d. On the Details page, configure the user settings, and then select **Next**.

   Details page

Field	Description
User name	Name for the programmatic user. For example, `servicenowcloud`.
Access type	Select **Programmatic access**.

   e. On the Permissions page, attach the user to a policy.
   Configure the following settings and then select **Next**.

Field	Description
Set permissions for <user name>	Select **Attach existing policies directly**.
Attach one or more policies ...	Select the appropriate policy. **Note:** The **AdministratorAccess** policy has the most powerful permission level, including permission to provision cloud resources. The policy enables the same access that would be granted to the instance if you were not using IAM and used your AWS account Access Key ID and Secret Access Key. You might instead prefer to create a policy or combine multiple policies to grant the appropriate permission level. See Control AWS access and permissions using policies for details.

   f. On the Review page, verify your selections and then select **Create user**.

   g. On the Security Credentials page, copy the secret access key and the access key ID.

   Important: Do not leave the page until you have completed both steps. The **Secret access key** value does not appear again. You need to paste the values that you generate in these steps into a Cloud Provisioning and Governance form.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

i. Select **Show** to display the Secret access key. Copy the value.

ii. Select **Download .csv** to save the CSV-format file that contains the user name, Access key ID, and the Secret access key value. You create the file as a backup in the case that you lose the values. Verify that the file was created and then store the file securely.

2. On the ServiceNow AI Platform, configure AWS credentials.

   a. In the ServiceNow instance, navigate to **Discovery** > **Credentials**.

   b. Select **New**, select **AWS Credentials**, enter a unique and meaningful **Name** (for example, `Cloud Provisioning Account`), and then fill in the form.

AWS Credentials form fields

Field	Input value
Name	Unique and descriptive name for the AWS credentials.
Active	Option to use the credential.
Access Key ID	The **Access key ID** that you generated on the AWS Management Console, such as: APIAIOSFODNN7EXAMPLE.
Secret access key	The **Secret access key** that you generated on the AWS Management Console, such as: wPalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY.

   c. Select **Update** or **Submit**.

What to do next

Create an AWS service account on the ServiceNow AI Platform. For more information, see Create AWS service accounts.

Create AWS service accounts

Create AWS service accounts on the ServiceNow AI Platform to access your AWS account during AWS discovery.

Before you begin

- Retrieve the user account number, which can be found in the AWS Management Console by expanding the list under the AWS account name.
- Confirm that Discovery Admin Workspace is using at least version 1.10.0. The **Discovery** > **Cloud Service Accounts** navigation module isn't available with earlier versions. To access **Cloud Service Accounts** with an earlier version, enter in the navigation filter: `cmdb_ci_cloud_service_account.list`.

Role required: discovery_admin

About this task

A service account is a secure record on your instance that stores the credential and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter. A cloud account is the logical

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

representation in cloud management of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts—even service accounts from different providers. For each service account, you specify which datacenter to include in the cloud account.

Procedure

1. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.
2. Select **New**.
3. On the form, fill in the fields.

Cloud Service Account form

Field	Description
Name	The unique and meaningful name for this service account.
Account ID	12-digit user account number. Expand the list under the AWS account name on the AWS Management Console to view the number. **Note:** In the **Account ID** field, remove the hyphen characters (-) from the number. **IAM user account number** 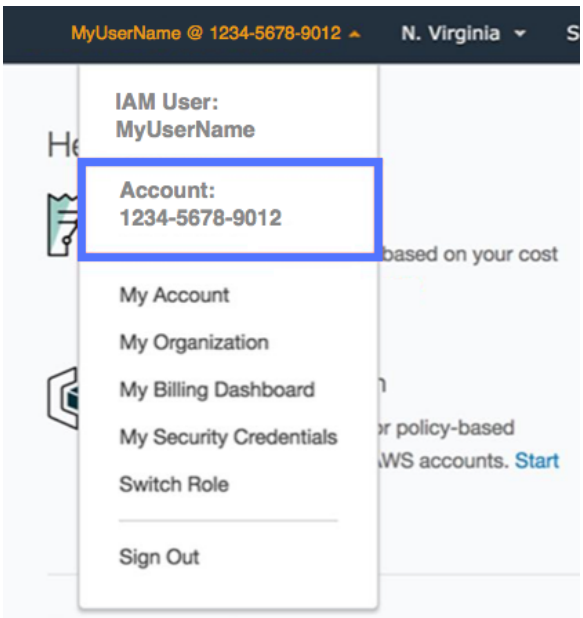
Discovery credentials	The credentials needed for ServiceNow applications to access this account. You may configure this field at a later stage, while configuring access to AWS accounts. ○ If you configured AWS credentials at ServiceNow AI Platform , select the magnifying glass icon, and then select the name of the relevant AWS credential. ○ To use other AWS accounts to access this account, leave the field blank. For example, you don't need to specify the AWS credentials for accounts assuming IAM roles or member accounts using their management account for access.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Datacenter URL	URL of the datacenter. This field is required for AWS China region and AWS GovCloud (US) accounts. For example: ○ AWS China region: `https://organizations.cn-northwest-1.amazonaws.com.cn` ○ AWS GovCloud (US): `https://ec2.us-gov-west-1.amazonaws.com`
Datacenter type	Type of the datacenter where the account is hosted. Select **AWS datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of Discovery on the datacenter.
Parent account	Name of the management account (also referred to as the master account in some ServiceNow UI screens) that represents the organization in AWS that this member account belongs to. This field appears when you select **AWS datacenter**. Leave the field empty if this account is not part of an AWS organization.
Is master account	Management account flag. This field is also referred to as the master account flag in some ServiceNow UI screens. This field appears when you select **AWS datacenter** from the Datacenter Type list. Select the check box to associate the AWS service account with the management account. Select this check box only for accounts that you previously configured in the AWS Management Console as management accounts with some member accounts belonging to them. See the AWS documentation for information on AWS Organizations. **Note:** You must set up the correct permission in AWS or the Organization role for a standard credential. For more information, see Control AWS access and permissions using policies.
Accessor account	Name of the trusted account. Configure this field only for accounts that don't use permanent AWS credentials and rely on IAM roles for access.

4. Select **Submit**.

Configure temporary credential access for trusted AWS accounts

Configure the trusting account whose resources need to be accessed, to rely on the trusted account using the Identity and Access Management (IAM) role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Familiarize yourself with the Amazon documentation on Creating a role to delegate permissions to an IAM user.
- Decide which AWS account is going to be the trusted account. You use the trusted account to configure temporary credentials for Cloud Discovery using IAM roles. The trusted account that you use to access other accounts using IAM roles is referred to as an accessor account.
- For trust chain setups, confirm that the member account trusts the management account before the management account trusts an accessor account. For more information, see Configure access for trusting AWS member accounts in trust chain.
- Confirm that Discovery Admin Workspace is using at least version 1.10.0. The **Discovery** > **Cloud Service Accounts** navigation module isn't available with earlier versions. To access **Cloud Service Accounts** with an earlier version, enter in the navigation filter: `cmdb_ci_cloud_service_account.list`.

Role required:

- For Cloud Discovery: discovery_admin
- For Cloud Provisioning and Governance: admin or sn_cmp.cloud_admin

About this task

During this configuration, you create an IAM role for the trusting account, and then configure the trusted service account for the trusting account at ServiceNow AI Platform. Finally, you associate the IAM role you created for the trusting account with the trusting account itself.

Setting up any AWS account to rely on a trusted account with AWS credentials

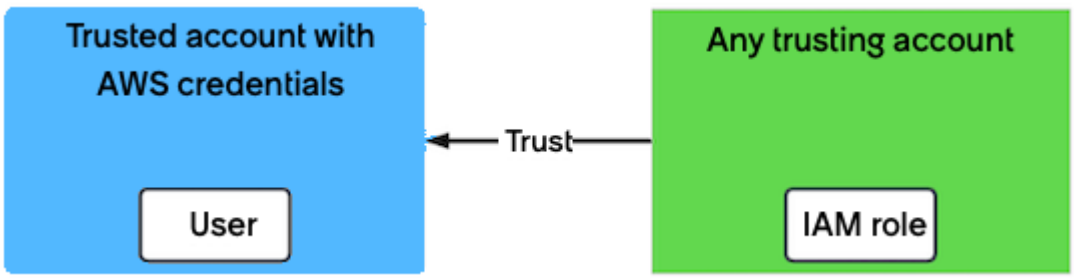

Procedure

1. Create an IAM role for the trusting account and configure the trust relationship between the user assuming this role and the trusted (accessor) account.

   a. Log in to the trusting account on the AWS Management Console.

   b. Create and configure the IAM role specifying the trusted (accessor) account ID in the **Account ID** field.
   For information on creating AWS roles, see the Amazon documentation.

   c. On the Summary page for the IAM role, select the **Trust Relationships** tab.

   d. Select **Edit trust relationship**.
   The Edit Trust Relationship page opens showing the policy document.

   e. Set the *AWS* parameter to the full user ARN of the trusted (accessor) account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

f. Verify that the *`Action`* value is set to *`sts:AssumeRole`*.

g. Select **Update Trust Policy**.

2. On the ServiceNow AI Platform, configure the trusted service account.

a. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.

b. Selelct **New**.

c. On the form, fill in the fields.
For a description of the field values, see Create AWS service accounts.

d. Select **Submit**.

3. On the ServiceNow AI Platform, configure the trusting service account.

a. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.

b. Select **New**.

c. In the **Accessor account** field, enter the name of the trusted account.

d. On the form, fill in the remaining fields.
For a description of the field values, see Create AWS service accounts.

e. Select **Submit**.

4. On the ServiceNow AI Platform, assign the AWS IAM role to the trusting account, using the relevant form, based on the relationship to the trusted account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

<table>
<tr><th>Trusted account type</th><th>Steps</th></tr>
<tr><td>Management account</td><td>
a. Navigate to All > Cloud Provisioning and Governance > Organization Access Parameters > AWS Org Assume Role Parameters.
b. Select New.
c. On the form, configure only the following fields for the trusting member account:
Cloud Service Account AWS Org Assume Role Params form
<table>
<tr><th>Field</th><th>Definition</th></tr>
<tr><td>Access role name</td><td>Name of the IAM role created for the trusting account.
▪ If IAM roles are the same across all member accounts: Enter the full ARN using an asterisk (*) as a wildcard for the account ID in the format: <code>arn:aws:iam::*:role/MemberRoleName</code>.
For example: <code>arn:aws:iam::*:role/SN_MEMBER_ACCOUNT_ROLE</code>.
▪ If IAM roles are different across member accounts: Enter the full ARN of the specific IAM role for each member account in a separate entry.</td></tr>
<tr><td>Cloud service account</td><td>Name of the trusting account for which you are providing access using the IAM role.
▪ If IAM roles are the same across all member</td></tr>
</table>
</td></tr>
</table>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

<table>
<tr><th>Trusted account type</th><th>Steps</th></tr>
<tr><td></td><td>
<table>
<tr><th>Field</th><th>Definition</th></tr>
<tr><td></td><td>accounts: Enter the management account name.
▪ If IAM roles are different across member accounts: Enter each member account in a separate entry.</td></tr>
</table>
d. Select Submit.
</td></tr>
<tr><td>Member or discrete account</td><td>
a. Navigate to All > Cloud Provisioning and Governance > Organization Access Parameters > AWS Cross Assume Role Parameters.
b. Select New.
c. On the form, configure only the following fields for the trusting account:
Cloud Service Account AWS Cross Assume Role Params form
<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td>Access role name</td><td>Name of the IAM role created for the trusting account.</td></tr>
<tr><td>Cloud service account</td><td>Name of the trusting account for which you are providing access using the IAM role.</td></tr>
</table>
d. Select Submit.
</td></tr>
</table>

What to do next

Verify that ServiceNow applications can access the trusting service account using the IAM role:

1. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.
2. Select the trusting AWS service account.
3. Under **Related Links**, select **Create Discovery Schedule**.
4. In the Discovery Manager Cloud Discovery page, select **Test Account**.
   - If the connection is successful, a message displays indicating the account validation is successful.
   - If the connection isn't successful, an error message displays indicating the cause of failure.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure credential-less access using trusted AWS accounts

Set up a trusted credential-less account that other AWS accounts can rely on for access.

Before you begin

- Familiarize yourself with the Amazon documentation on Creating a role to delegate permissions to an IAM user.
- Decide which AWS account is going to be the trusted account. You use the trusted account to configure temporary credentials for Cloud Discovery using IAM roles. The trusted account that you use to access other accounts using IAM roles is referred to as an accessor account.
- If you're setting up a trust chain, confirm that the member account trusts the management account. The management account must also trust the accessor account. For more information, see Configure access for trusting AWS member accounts in trust chain.
- Confirm that Discovery Admin Workspace is using at least version 1.10.0. The **Discovery** > **Cloud Service Accounts** navigation module isn't available with earlier versions. To access **Cloud Service Accounts** with an earlier version, enter in the navigation filter: `cmdb_ci_cloud_service_account.list`.

Role required:

- For Cloud Discovery: discovery_admin
- For Cloud Provisioning and Governance: admin or sn_cmp.cloud_admin

About this task

To use an account without AWS credentials, you must first configure that account with an IAM role and permissions to access the trusting service account. Then, you set up the IAM role of the trusting account to grant access to the IAM role of the trusted account.

Setting up any AWS account to rely on a trusted account without AWS credentials

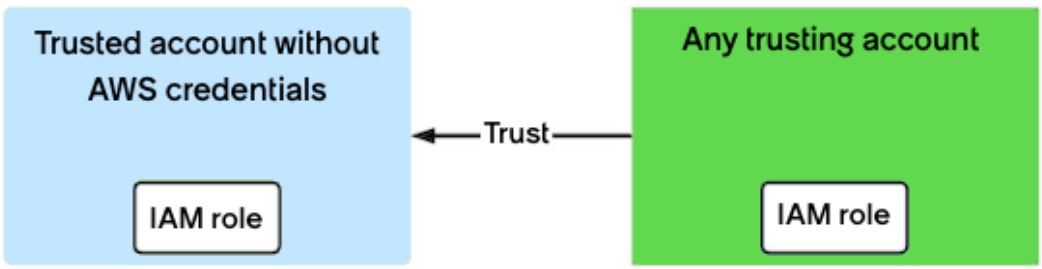

Procedure

1. Configure an IAM role for the trusting account.
   a. Log into the trusting account on the AWS Management Console.
   b. Create an IAM role for this account.
      Use the account ID of the trusted account while creating this IAM role. For operational information about working with AWS roles, refer to the Amazon documentation.
   c. Create a ReadOnlyAccess policy and attach it to the newly created IAM role.
2. Configure the IAM role for the trusted account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Log into the AWS Management Console using the credentials of the account that you want to set up as a trusted account.

b. Create an IAM role by choosing the **AWS service** option.

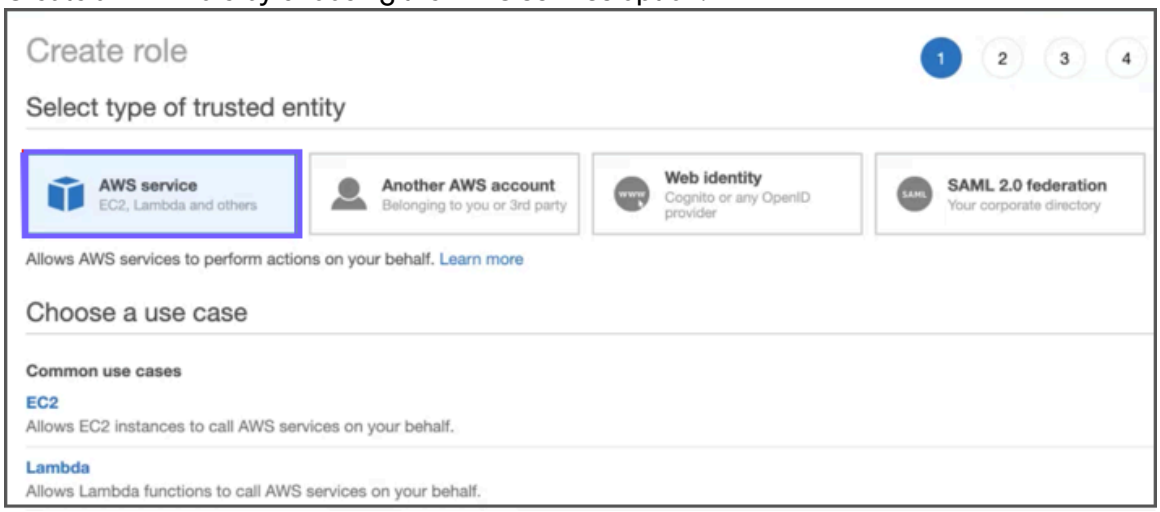

c. Create a read-only access policy for the trusted account IAM role.
For more information, see Control AWS access and permissions using policies.

d. Create an additional policy to grant this IAM role access to resources in trusting accounts:

- Set the *`Action`* parameter to `sts:AssumeRole`
- Set the *`Resource`* parameter to the ARN of the trusting account role that you created in 1.b.

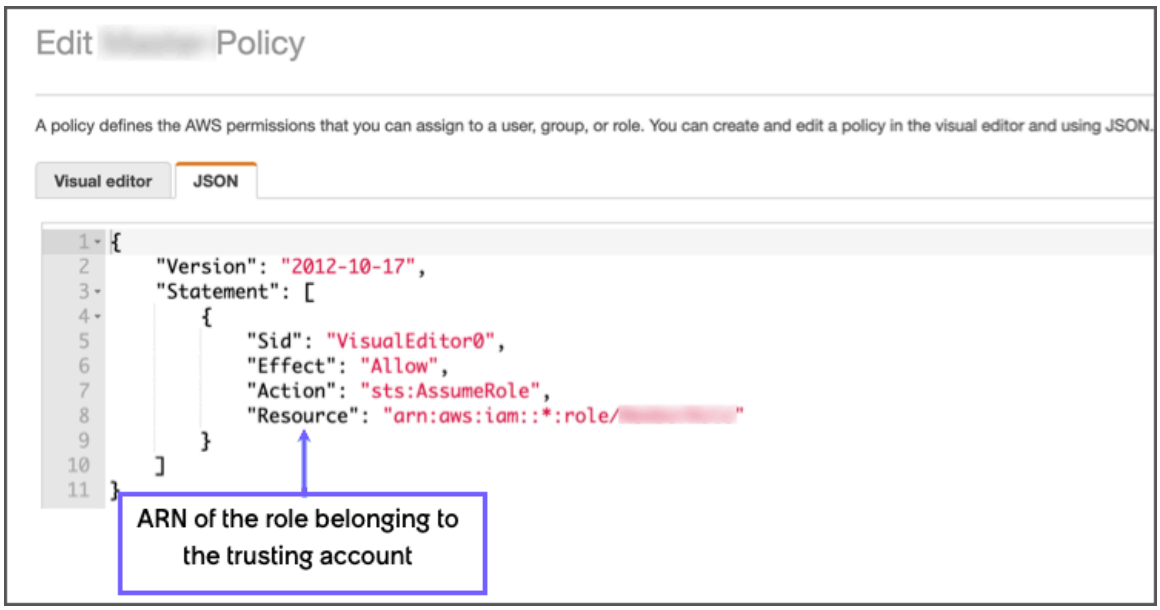

e. Attach the newly created role to the relevant Amazon EC2 instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

By default, when you attach an IAM role to an EC2 instance, it creates a trust relationship between this role and the EC2 instance.

3. Configure the trusting service account to grant access to the IAM role belonging to the trusted account.

   a. Log into the trusting account on the AWS Management Console.

   b. Navigate to the IAM role you created for this account as described in 1.b.

   c. Edit the Trust Relationship for this IAM role as follows:

   - Set the *Action* parameter to `sts:AssumeRole`.
   - Set the *AWS* parameter to the ARN of the trusted account role that you created in 2.b.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Configure the MID Server for AWS IAM roles.
   For more information, see Configure the MID Server for AWS IAM roles.
5. On the ServiceNow AI Platform, configure the trusted service account.
   a. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.
   b. Selelct **New**.
   c. On the form, fill in the fields.
      For a description of the field values, see Create AWS service accounts.
   d. Select **Submit**.
6. On the ServiceNow AI Platform, configure the trusting service account.
   a. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.
   b. Select **New**.
   c. In the **Accessor account** field, enter the name of the trusted account.
   d. On the form, fill in the remaining fields.
      For a description of the field values, see Create AWS service accounts.
   e. Select **Submit**.
7. On the ServiceNow AI Platform, assign the AWS IAM role to the trusting account, using the relevant form, based on the relationship to the trusted account.

Trusted account type	Steps
Management account	**a.** Navigate to **All** > **Cloud Provisioning and Governance** > **Organization Access Parameters** > **AWS Org Assume Role Parameters**. **b.** Select **New**. **c.** On the form, configure only the following fields for the trusting member account: **Cloud Service Account AWS Org Assume Role Params form** **Field** / **Definition** Access role name — Name of the IAM role created for the trusting account. ▪ If IAM roles are the same across all member accounts:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

<table>
<tr><th>Trusted account type</th><th>Steps</th></tr>
<tr><td></td><td>
<table>
<tr><th>Field</th><th>Definition</th></tr>
<tr><td></td><td>Enter the full ARN using an asterisk (*) as a wildcard for the account ID in the format: <code>arn:aws:iam::*:role/MemberRoleName</code>.

For example: <code>arn:aws:iam::*:role/SN_MEMBER_ACCOUNT_ROLE</code>.

▪ If IAM roles are different across member accounts: Enter the full ARN of the specific IAM role for each member account in a separate entry.</td></tr>
<tr><td>Cloud service account</td><td>Name of the trusting account for which you are providing access using the IAM role.

▪ If IAM roles are the same across all member accounts: Enter the management account name.

▪ If IAM roles are different across member accounts: Enter each member account in a separate entry.</td></tr>
</table>
d. Select Submit.</td></tr>
<tr><td>Member or discrete account</td><td>a. Navigate to All > Cloud Provisioning and Governance > Organization Access Parameters > AWS Cross Assume Role Parameters.

b. Select New.</td></tr>
</table>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Trusted account type	Steps
	c. On the form, configure only the following fields for the trusting account: **Cloud Service Account AWS Cross As sume Role Params form** **Field** / **Description** Access role name — Name of the IAM role created for the trusting account. Cloud service account — Name of the trusting account for which you are providing access using the IAM role. **d.** Select **Submit**.

What to do next

Verify that ServiceNow applications can access the trusting service account using the IAM role:

1. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.
2. Select the trusting AWS service account.
3. Under **Related Links**, select **Create Discovery Schedule**.
4. In the Discovery Manager Cloud Discovery page, select **Test Account**.
   - If the connection is successful, a message displays indicating the account validation is successful.
   - If the connection isn't successful, an error message displays indicating the cause of failure.

Configure access for trusting AWS member accounts in trust chain

Configure access for AWS member accounts by using a trust chain from the accessor through the management account.

Before you begin

- Familiarize yourself with the Amazon documentation on Creating a role to delegate permissions to an IAM user.
- Ensure that you know which AWS member accounts are assigned to the same management account. You use the management account for configuring temporary credentials for cloud discovery using IAM roles.
- Confirm that Discovery Admin Workspace is using at least version 1.10.0. The **Discovery** > **Cloud Service Accounts** navigation module isn't available with earlier versions. To access **Cloud Service Accounts** with an earlier version, enter in the navigation filter: `cmdb_ci_cloud_service_account.list`.

Role required:

- For Cloud Discovery: discovery_admin
- For Cloud Provisioning and Governance: admin or sn_cmp.cloud_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

You can configure access for AWS member accounts by using a trust chain from the accessor through the management account. The accessor account either has AWS credentials or uses a credential-free method.

Configuring member accounts to use their management account for access

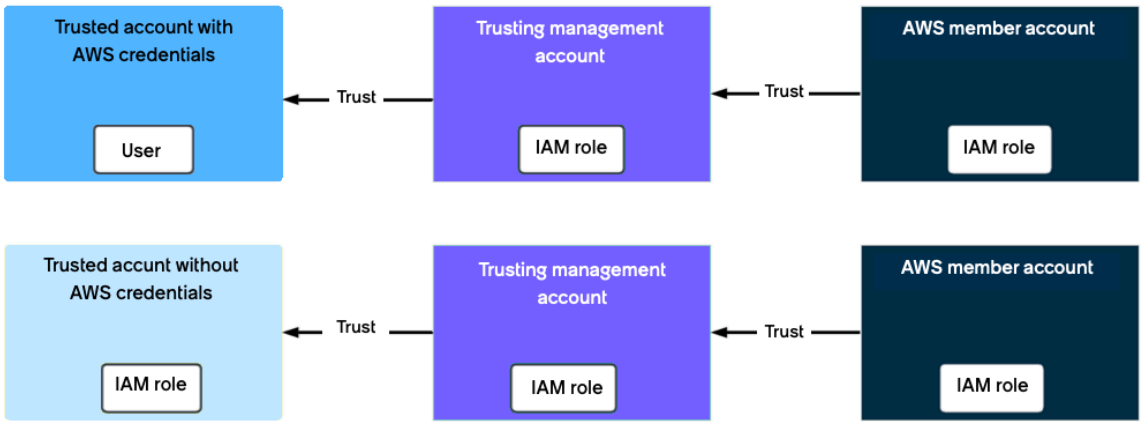

Procedure

1. Create an IAM role for the member account and configure the trust relationship between the user assuming this role and the trusted management account.

   a. Log into the AWS Management Console using the credentials of the member account for which you are configuring access.

   b. Create and configure the IAM role specifying the management account ID in the **Account ID** field.
   For operational information about creating AWS roles, refer to the Amazon documentation.

   c. On the **Summary** page for the IAM role, click the **Trust Relationships** tab.

   d. Click **Edit trust relationship**.
   The Edit Trust Relationship page opens showing the policy document.

   e. Edit the trust relationship as follows:
   - Set the *`Action`* parameter to `sts:AssumeRole`
   - Set the `AWS` parameter to the full role ARN of the management account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

f. Click **Update Trust Policy**.

2. On the ServiceNow AI Platform, configure the trusting member's service account.

   a. Navigate to **All** > **Discovery** > **Cloud Service Accounts**.

   b. Select **New**.

   c. In the **Parent account** field, enter the name of the management account.

   d. On the form, fill in the remaining fields.
   For a description of the field values, see Create AWS service accounts.

   e. Select **Submit**.

3. On the ServiceNow AI Platform, assign the AWS IAM role to the member account.

   Important: Perform this step only if you created custom IAM roles. By default, the OrganizationAccountAccessRole role is assigned to the member's trusting management account and you don't need to assign the OrganizationAccountAccessRole role to a service account.

   a. Navigate to **All** > **Cloud Provisioning and Governance** > **Organization Access Parameters** > **AWS Org Assume Role Parameters**.

   b. Select **New**.

   c. On the form, configure only the following fields for the trusting member account:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cloud Service Account AWS Org Assume Role Params form

Field	Definition
Access role name	Name of the IAM role created for the trusting account. ▪ If IAM roles are the same across all member accounts: Enter the full ARN using an asterisk (*) as a wildcard for the account ID in the format: `arn:aws:iam::*:role/MemberRoleName`. For example: `arn:aws:iam::*:role/SN_MEMBER_ACCOUNT_ROLE`. ▪ If IAM roles are different across member accounts: Enter the full ARN of the specific IAM role for each member account in a separate entry.
Cloud service account	Name of the trusting account for which you are providing access using the IAM role. ▪ If IAM roles are the same across all member accounts: Enter the management account name. ▪ If IAM roles are different across member accounts: Enter each member account in a separate entry.

d. Select **Submit**.

What to do next

Configure the trusting management account and the trusted accessor account.

- For information on configuring accessor accounts with AWS credentials, see Configure temporary credential access for trusted AWS accounts.
- For information on configuring accessor accounts without AWS credentials, see Configure credential-less access using trusted AWS accounts.

Configure the MID Server for AWS IAM roles

Configure the MID Server to retrieve the temporary security credentials associated with an IAM role.

Before you begin

Role required: discovery_admin or sn_cmp.cloud_admin (for Cloud Provisioning and Governance)

Procedure

1. Navigate to **All** > **Discovery** > **MID Servers**.
2. Select a MID Server installed on an Amazon EC2 instance within the relevant AWS service account.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Open the **Configuration Parameters** related list.
4. Click **New**.
5. In the configuration parameter form, select *mid.aws.instance_profile_name* in the **Parameter name** field.
6. Enter the name of the IAM role attached to the EC2 Instance in the **Value** field.

   Note: AWS Cloud Discovery for "mid.aws.instance_profile_name" accepts both "role name" and "full ARN" as valid parameter values. For example, both of the following values are acceptable:

   - arn:partition:service:region:account-id:resource-type/resource-id
   - resource-id

   The MID Server uses the same IAM role that is set up for the corresponding AWS service account. For more information, see Configure temporary credential access for trusted AWS accounts.

   For operational information about creating AWS roles, refer to the Amazon documentation on Creating a role to delegate permissions to an IAM user.

7. Click **Submit**.
   The new parameter is listed in the MID Server record.

Control AWS access and permissions using policies

Configure policies with the necessary level of permissions to provide access to the AWS resources for Cloud Discovery and Cloud Provisioning and Governance.

Before you begin

Familiarize yourself with creating IAM users and user policies. Refer to the AWS documentation.

Identify the level of permissions required for discovering the resources using Cloud Discovery and managing the resources with Cloud Provisioning and Governance. For example:

- Additional S3 permissions are needed as AWS uses an S3 bucket to store and run the cloud formation templates to deploy stacks using CloudFormation. These permissions are needed even if you run the same CloudFormation template from the AWS console.

  Note: `s3:ListAllMyBuckets` must be scoped account-wide, not to a specific bucket ARN.

- Read-only access is required to do pattern discovery with the custom policy.

AWS CloudFormation provisions and configures resources by making calls to the AWS services that are described in your template. For more information, see How does CloudFormation work?

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Roles required:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- AWS Management Console administrator
- For Cloud Discovery: discovery_admin
- For Cloud Provisioning and Governance: admin or sn_cmp.cloud_admin

About this task

The policy you configure defines the AWS permissions that you can assign to a user, group, or a role. Depending on the AWS services you use, specify actions for the service to perform discovery and provisioning and life-cycle actions using Cloud Provisioning and Governance.

Procedure

1. Log in to the AWS Management Console using the credentials of the relevant service account. If you used an IAM user or an IAM role to provide access using temporary credentials, log in using the same IAM user or IAM role.
2. Open the user record in the instance for the relevant user.
3. Define a user policy in AWS Management Console using either of the following methods:
   - Grant administrator access to the instance. This method gives the same level of access as using the Access Key ID and Secret Access Key gives. Attach the `AdministratorAccess` policy to the user profile.

     Note: To create a user policy that supports only Cloud Discovery rather than the provisioning of cloud resources, attach the `ReadOnlyAccess` policy instead.

   - Create a custom policy with a descriptive name. In the Policy Document field, write the code that includes the APIs that this policy allows running.

     Note: The following JSON sample is a partial example showing permissions for Cloud Discovery. It isn't a complete list of all required permissions.

```
{
   "Version":"2012-10-17",
   "Statement":[
      {
         "Sid":"PermissionsNeededForDiscovery",
         "Effect":"Allow",
         "Action":[
            "account:ListRegions",
            "apigateway:GET",
            "application-autoscaling:Describe*",
            "autoscaling-plans:Describe*",
            "autoscaling:Describe*",
            "autoscaling:DescribeAutoScalingGroups",
            "autoscaling:DescribeLaunchConfigurations",
            "cloudformation:Describe*",
            "cloudformation:DescribeStack*",
            "cloudformation:List*",
            "cloudwatch:Get*",
            "cloudwatch:List*",
            "config:ListDiscoveredResources",
            "dynamodb:Describe*",
            "dynamodb:DescribeTable",
            "dynamodb:ListGlobalTables",
            "dynamodb:ListTables",
            "ec2:Describe*",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
            "ec2:ReportInstanceStatus",
            "ecs:Describe*",
            "ecs:List*",
            "eks:Describe*",
            "eks:List*",
            "elasticache:Describe*",
            "elasticache:List*",
            "elasticfilesystem:Describe*",
            "elasticloadbalancing:Describe*",
            "lambda:Get*",
            "lambda:List*",
            "organizations:Describe*",
            "organizations:DescribeOrganization",
            "organizations:List*",
            "rds:Describe*",
            "redshift:Describe*",
            "route53:GetHostedZone",
            "route53:List*",
            "s3:ListAllMyBuckets",
            "s3:GetBucketLocation",
            "s3:GetBucketAcl",
            "s3:GetEncryptionConfiguration",
            "s3:GetBucketPolicy",
            "s3:GetReplicationConfiguration",
            "s3:GetBucketNotification",
            "tag:GetResources",
            "sdb:GetAttributes",
            "sns:GetEndpointAttributes",
            "sns:GetSubscriptionAttributes",
            "sns:List*",
            "sns:ListPlatformApplications",
            "sns:ListSubscriptions",
            "sns:ListSubscriptionsByTopic",
            "sns:ListTopics"
         ],
         "Resource":"*"
      }
   ]
}
```

- The following JSON sample is a partial example showing permissions for Cloud Discovery and Cloud Provisioning and Governance. It isn't a complete list of all required permissions.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "PermissionsNeededForCPG",
            "Effect": "Allow",
            "Action": [
                "account:ListRegions",
                "apigateway:GET",
                "application-autoscaling:Describe*",
                "autoscaling-plans:Describe*",
                "autoscaling:Describe*",
                "autoscaling:DescribeAutoScalingGroups",
                "autoscaling:DescribeLaunchConfigurations",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                "cloudformation:Describe*",
                "cloudformation:DescribeStack*",
                "cloudwatch:Get*",
                "cloudwatch:List*",
                "config:ListDiscoveredResources"
                "dynamodb:Describe*",
                "dynamodb:DescribeTable",
                "dynamodb:ListGlobalTables",
                "dynamodb:ListTables",
                "ec2:Describe*",
                "ec2:ReportInstanceStatus",
                "ecs:Describe*",
                "ecs:List*",
                "eks:Describe*",
                "eks:List*",
                "elasticache:Describe*",
                "elasticache:List*",
                "elasticfilesystem:Describe*",
                "elasticloadbalancing:Describe*",
                "lambda:Get*",
                "lambda:List*",
                "organizations:Describe*",
                "organizations:DescribeOrganization",
                "organizations:List*",
                "rds:Describe*",
                "redshift:Describe*",
                "route53:GetHostedZone",
                "route53:List*",
                "s3:ListAllMyBuckets",
                "s3:GetBucketLocation",
                "s3:GetBucketAcl",
                "s3:GetEncryptionConfiguration",
                "s3:GetBucketPolicy",
                "s3:GetReplicationConfiguration",
                "s3:GetBucketNotification",
                "tag:GetResources",
                "sdb:GetAttributes",
                "sns:GetEndpointAttributes",
                "sns:GetSubscriptionAttributes",
                "sns:List*",
                "sns:ListPlatformApplications",
                "sns:ListSubscriptions",
                "sns:ListSubscriptionsByTopic",
                "sns:ListTopics"
            ],
            "Resource": "*"
        },
        {
            "Sid":
 "MinimalPermissionsNeededForEc2ProvisioningThroughCloudForma
tion",
            "Effect": "Allow",
            "Action": [
                "cloudformation:CreateStack",
                "cloudformation:CreateUploadBucket",
                "cloudformation:DeleteStack",
                "cloudformation:DescribeStackEvents",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                "cloudformation:DescribeStackResources",
                "cloudformation:DescribeStacks",
                "cloudformation:GetTemplate",
                "cloudformation:GetTemplateSummary",
                "cloudformation:List*",
                "cloudformation:UpdateStack",
                "cloudformation:ValidateTemplate"
                "ec2:*",
                "s3:ListAllMyBuckets",
                "s3:GetBucketLocation",
                "s3:GetBucketAcl",
                "s3:GetEncryptionConfiguration",
                "s3:GetBucketPolicy",
                "s3:GetReplicationConfiguration",
                "s3:GetBucketNotification",
                "tag:GetResources"
            ],
            "Resource": "*"
        }
    ]
}
```

Related topics

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_create-console.html#access_policies_create-start

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_understand-service-summary.html

Set up Azure service accounts

Create and configure cloud service accounts at ServiceNow AI Platform for the corresponding Microsoft Azure accounts.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Before you begin

Ensure the configuration of Azure credentials

For more information, see Create Azure cloud credentials and Configuration steps for Azure Service Principal User and Credentials (KB2202765)

Ensure that you are familiar with the hierarchy of the Azure service accounts in your environment.

Set up the service accounts at ServiceNow AI Platform to reflect the hierarchy of your environment. You can set up the Azure service accounts for the following Azure account types:

- Management group.
- Subscription.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: discovery_admin

About this task

A service account is a secure record on your instance that stores the credentials and access information for your provider account. Discovery uses the information to access your provider account to get data on each resource in each specified datacenter. A cloud account is the logical representation in cloud management of all or part of your managed cloud infrastructure. A cloud account can include multiple service accounts—even service accounts from different providers. For each service account, you specify which datacenter to include in the cloud account.

Procedure

1. In the navigation filter, enter `cmdb_ci_cloud_service_account.list`.
2. Select **New**.
3. Fill in the fields by the following order:
   a. Select the datacenter.
   b. Select the check box of the **Is management account** if your are creating a service account for a management group.
   c. Select **Should pull events** to preform event-based discovery. For more information, see Microsoft Azure Alert driven discovery

Cloud Service Account form

Field	Description
Name	The unique and meaningful name for this service account.
Account ID	If you are creating a service account for a management group, fill in the Azure management group ID that belongs to the management group. If you are creating a service account for a subscription, fill in the Azure subscription ID.
Discovery credentials	The credentials needed for ServiceNow applications to access this account.
Datacenter URL	URL of the datacenter. Configuring this field is required for Azure Gov Cloud (US) accounts. For example, https://management.usgovcloudapi.net/.
Datacenter type	Type of the datacenter where the account is hosted. Select **Azure datacenter**.
Datacenter discovery status	Auto-generated value: Status and timestamp of the last execution of discovery on the datacenter.
Parent account	When setting up a service account for an Azure subscription, fill in the management group ID. When the **Is management account** check box is selected, the **Parent account** field is not applicable

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Accessor account	Not applicable for Azure service accounts.

4. Select **Submit**.

Related topics

Discovery for Microsoft Azure

Set up Google Cloud Platform service accounts

Create and configure cloud service accounts at ServiceNow AI Platform for the corresponding Google Cloud Platform (GCP) service accounts.

GCP service account is a unique account typically used by an application or compute workload, such as a Compute Engine instance, rather than a user. A service account is identified by its email address, which is unique to the account. It's used to authenticate and authorize interactions between your application and various GCP services.

Verify the REST API permissions

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

Configure the Google API credentials for Cloud Discovery

Create Google API credentials to allow Cloud Discovery to access the data on your Google account securely.

Before you begin

Ensure that the latest versions of the following applications are installed:

- Discovery and Service Mapping Patterns
- Visibility Content
- CMDB CI Class Models

Role required: discovery_admin

Procedure

Configure the GCP credentials either by importing a JSON file or by filing in the record manually.

Option	Actions
Import the JSON file	**a.** When receiving the email from GCP, down load the JSON file. **b.** Navigate to **Discovery** > **Credentials**. **c.** Select **New** and then select **Google API Cre dentials**. **d.** Select **Import JSON File**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Actions
	e. Select **Choose File** and choose the JSON file you downloaded. **f.** Select **Upload Credential Data**.
Configure the credentials manually	**a.** Navigate to **Discovery** > **Credentials**. **b.** Select **New** and then select **Google API Credentials**. **c.** Select **New**. **d.** In the form, fill in the fields according to the form description. For description of the form fields, see Google API Credentials Form. **e.** Select **Submit**.

Google API Credentials Form

Fill in the Google API Credentials form to configure the credentials.

Field	Description
Name	Set a unique name for the Google Cloud API credentials.
Active	Option for enabling this credential for discovery. Select this check box to enable discovery.
Applies to	MID Server to which you apply the credentials. Select **All MID servers**.
Order	Order in which the platform runs this credential as it attempts to log in to devices. A smaller number indicates that the credential appears higher in the list. Establish the credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number, or none, the instance tries the credentials in a random order. The default value is 100.
Email	Your GCP email address.
Secret Key	Secret key that you generated on the Google Cloud Console.
Credential alias	A credential alias that is used by the credentials to authorize actions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create GCP service accounts

Create Google Cloud Platform (GCP) service accounts on the ServiceNow AI Platform to access your GCP account during GCP discovery.

Before you begin

Verify that Google API credentials have been created. For more information, see Configure the Google API credentials for Cloud Discovery.

Retrieve the **project_id** value from the JSON key file that is associated with the service account.

Role required: discovery_admin

Procedure

1. In the navigation filter, enter `cmdb_ci_cloud_service_account.list`.
2. Select **New**.
3. In the **Discovery Credentials field**, begin to enter the credential name and select it from the list.
4. On the form, fill in the remaining fields.

Cloud Service Account form

Field	Description
Name	A unique name for the GCP service account.
Account Id	GCP project ID.
Datacenter URL	URL of the datacenter, which should remain empty.
Datacenter Type	Type of the datacenter where the account is hosted, which should be Google Datacenter [cmdb_ci_google_datacenter].

5. Select **Submit**.

What to do next

Schedule a GCP cloud discovery. For more information, see Create a discovery schedule in Cloud Discovery Workspace.

GCP access and permission using policies

Google Cloud Platform (GCP) access and permission using policies provide privileges for secure cloud resource utilization and management in GCP.

Download the Cloud Discovery patterns spreadsheet so you can grant user permissions required for running the Discovery patterns. In addition to permissions, the spreadsheet also includes useful information such as pattern names, types, CI Classes, and links to vendor documentation. New patterns are available quarterly, so check periodically to be sure you have the latest version of the spreadsheet.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.